

AI-Enhanced Adaptive Virtual Screening Platform Enabling Exploration of 69 Billion Molecules Discovers Structurally Validated FSP1 Inhibitors

Christoph Gorgulla^{1,2,3,29,+,*}, Domiziana Cecchini^{30,+}, AkshatKumar Nigam^{4,24,+}, Ming Tang^{26,27,29,+}, Joana Reis^{3,+}, Matt Koop⁹, Andrea Gottinger³⁰, Callum Robert Nicoll³⁰, Abhilash Jayaraj^{1,3}, Süleyman Selim Çınaroğlu^{21,32}, Ricarda Törner^{3,36}, Yehor Malets²², Minko Gehev¹², Krishna M. Padmanabha Das^{1,29}, Kelly Churion²⁹, Jongwan Kim²⁹, Hyuk-Soo Seo³, Sirano Dhe-Paganon³, Christopher Secker^{8,18}, Mohammad Haddadnia¹⁹, Alexander Hasson^{20,28}, Minkai Li²⁵, Abhishek Kumar^{23,29}, Roni Levin-Konigsberg⁴, Eun-Bee Choi³⁴, Geoffrey I. Shapiro³⁴, Huel Cox 3rd³, Luke Sebastian³, Chelsea Braithwaite³, Puspala Bashyal³, Dmytro S. Radchenko⁵, Aditya Kumar¹⁰, Pierre-Yves Aquilanti^{9,35}, Henry Gabb¹¹, Amr Alhossary^{13,33}, Gerhard Wagner¹, Alán Aspuru-Guzik¹⁴⁻¹⁷, Yurii S. Moroz^{5,6,7}, Charalampos G. Kalodimos²⁹, Konstantin Fackeldey^{8,10}, Andrea Mattevi^{30*}, Haribabu Arthanari^{1,3,*}

Identifying potent lead molecules for specific targets remains a major bottleneck in drug discovery. As structural information about proteins becomes increasingly available, ultra-large virtual screenings (ULVSs) which computationally evaluate billions of molecules offer a powerful way to accelerate early-stage drug discovery. Here, we introduce AdaptiveFlow, an open-source platform designed to make ULVSs more accessible, scalable, and efficient. AdaptiveFlow provides free access to a screening-ready version of the Enamine REAL Space (1), the largest library of ready-to-dock, drug-like molecules, containing 69 billion compounds that we prepared using the ligand preparation module of the platform. A key innovation of the platform is its use of a multi-dimensional grid of molecular properties, which helps researchers explore and prioritize chemical space more effectively and reduce the computational costs by a factor of approximately 1000. This grid forms the basis of a new method for identifying promising regions of chemical space, enabling systematic exploration and prioritization of compound libraries. An optional active learning component can further accelerate this process by adaptively steering the search toward molecules most likely to bind a given target. To support a broad range of applications, AdaptiveFlow is compatible with over 1,500 docking methods. The platform achieves near-linear scaling on up to 5.6 million CPUs in the AWS Cloud, setting a new benchmark for large-scale cloud computing in drug discovery. Using this approach, we identified nanomolar inhibitors of two disease-relevant targets: ferroptosis suppressor protein 1 (FSP1) and poly(ADP-ribose) polymerase 1 (PARP-1) (2, 3). By leveraging newly solved crystal structures of FSP1 in complex with NAD⁺, FAD, and coenzyme Q₁, we validated these hits experimentally and determined the first co-crystal structures of FSP1 bound to small-molecule inhibitors, enabling insights into inhibitor binding mechanisms previously unknown. With its high scalability, flexibility, and open accessibility, AdaptiveFlow offers a powerful new resource for discovering and optimizing drug candidates at an unprecedented scale and speed.

Introduction

A central challenge in early-stage drug discovery is the identification and optimization of initial hit and lead compounds that bind specifically and potently to biological macromolecules. Historically, experimental high-throughput screening (HTS) has served as the primary method for discovering such hits (4). While these approaches have yielded many successful leads, they are inherently constrained by high costs, long timelines, and the need for robust, scalable assays. These limitations are further compounded by the vastness of chemical space: the number of synthetically accessible, drug-like small molecules is estimated to exceed 10⁶⁰ (5), yet typical HTS campaigns screen only hundreds of thousands of compounds. This minuscule sampling poses a major bottleneck - particularly for challenging targets such as protein–protein interaction interfaces or allosteric sites, where binding pockets are often shallow, dynamic, or poorly defined. Virtual screenings have emerged as a compelling alternative, offering the ability to computationally evaluate extremely large numbers of molecules. Virtual screening surpassed millions of compounds over a decade ago (6) and, in recent years, has expanded to hundreds of millions to billions of molecules in silico,

¹Department of Biological Chemistry and Molecular Pharmacology, Harvard Medical School, Harvard University, Boston, USA. ²Department of Physics, Faculty of Arts and Sciences, Harvard University, Cambridge, USA. ³Department of Cancer Biology, Dana-Farber Cancer Institute, Boston, USA. ⁴Klyne, San Francisco, USA. ⁵Enamine Ltd, Kyiv, Ukraine. ⁶Chospace LLC, Kyiv, Ukraine. ⁷Taras Shevchenko National University of Kyiv, Kyiv, Ukraine. ⁸Zuse Institute Berlin (ZIB), Berlin, Germany. ⁹Amazon Web Services, Seattle, USA. ¹⁰Institute for Mathematics, Technical University Berlin, Berlin, Germany. ¹¹Intel, Santa Clara, USA. ¹²Google, Mountain View, USA. ¹³Wesleyan University, Middletown, USA. ¹⁴Department of Computer Science, University of Toronto, Toronto, Ontario, Canada. ¹⁵Vector Institute for Artificial Intelligence, Toronto, Ontario, Canada. ¹⁶Department of Chemistry, University of Toronto, Toronto, Ontario, Canada. ¹⁷Canadian Institute for Advanced Research, Toronto, Ontario, Canada. ¹⁸Max Delbrück Center for Molecular Medicine, Berlin, Germany. ¹⁹Department of Molecular Genetics, University of Toronto, Toronto, Canada. ²⁰Department of Oncology, University of Oxford, Oxford, UK. ²¹Department of Biochemistry, University of Oxford, UK. ²²V. P. Kukhar Institute of Bioorganic Chemistry and Petrochemistry, National Academy of Science of Ukraine, Kyiv, Ukraine. ²³Department of Biosciences and Bioengineering, Indian Institute of Technology Guwahati, India. ²⁴Department of Computer Science, Stanford University, Stanford, USA. ²⁵Harvard College, Cambridge, USA. ²⁶The University of Queensland, Frazer Institute, Faculty of Medicine at the Translational Research Institute Australia, Brisbane, Australia. ²⁷Queensland University of Technology, School of Biomedical Sciences, Centre for Genomics and Personalised Health at the Translational Research Institute Australia, Brisbane, Australia. ²⁸Mathematical Institute, University of Oxford, Oxford, UK. ²⁹Department of Structural Biology, St. Jude Children's Research Hospital, Memphis, Tennessee, USA. ³⁰Department of Biology and Biotechnology, University University of Pavia, Italy. ³¹Klyne AI, CA, USA. ³²Department of Bioengineering, Ege University, Bornova, 35040, Türkiye. ³³Pulse for Integrated Solutions GmbH, Unterschleißheim, Bavaria, Germany. ³⁴Department of Medical Oncology, Dana-Farber Cancer Institute, Boston, MA, USA. ³⁵NVIDIA, Santa Clara, CA, USA. ³⁶Department of Chemistry, University of Zurich, Switzerland.

⁺These authors contributed equally.

*email: c.gorgulla@gmail.com, andrea.mattevi@unipv.it, hari@hms.harvard.edu

41 giving rise to ultra-large virtual screens (ULVSs). Compared to traditional HTS, ULVSs dramatically reduce cost and time,
42 while expanding the scope of chemical space that can be explored. In addition to identifying potent candidate compounds, virtual
43 screening can provide insights into binding mechanisms and structure–activity relationships, making it especially valuable
44 for targets that are difficult to modulate using conventional experimental approaches.

45 In recent years, multiple studies have demonstrated the remarkable effectiveness of ULVSs in identifying potent hits
46 across a wide range of protein targets, including enzymatic active sites, orthosteric sites of GPCRs, and protein–protein inter-
47 action interfaces. In 2019, Lyu and colleagues (7) targeted the D4 dopamine receptor and AmpC in a pioneering work using
48 ULVSs and obtained picomolar binders directly from the screen. In 2020, the melatonin receptor MT1 was successfully targeted
49 with ULVS, resulting in compounds with sub-nanomolar potency (8). In the same year, ULVSs were used for the first time to
50 target protein–protein interactions. That study performed the first screen of one billion compounds to target the KEAP1-NRF2
51 interaction, leading to the identification of inhibitors in the sub-micromolar range (9). In 2021, the $\sigma 2$ receptor was targeted
52 by ULVSs, leading to low-nanomolar ligands (10). Notably, in 2020, the first open-source, fully automated platform dedicated
53 to ULVSs called VirtualFlow was published (9, 11). This platform allows users to routinely perform ULVS with a library of
54 approximately 1.4 billion compounds in a "ready-to-dock format". In 2020, there were approximately 2 billion on-demand
55 molecules available. Driven by the success of billion-compound screens and advancements in synthesis and computational
56 technologies, the number of available on-demand molecules, with an established synthetic route, has expanded rapidly, with
57 current estimates reaching up to 3 trillion compounds across multiple vendors (1).

58 While ultra-large virtual screenings (ULVSs) enable the exploration of vast chemical libraries, significant challenges
59 remain in making these searches efficient, affordable, and intelligent. Navigating such immense chemical space to rapidly
60 identify promising hits requires new strategies that go beyond brute-force docking. The financial cost alone is a major bottleneck
61 - docking one billion compounds using standard cloud-based CPU providers typically costs \$20,000–30,000, (9) rendering it
62 impractical as chemical databases continue to grow exponentially, with newer libraries exceeding one trillion compounds, and
63 growing (7). Moreover, effective virtual screening requires the ability to flexibly integrate a wide range of docking algorithms,
64 each built on different theoretical foundations and requiring distinct input formats, to better match the characteristics of diverse
65 target classes. Compounding these challenges is the need to fully leverage modern computational infrastructure: while many
66 existing platforms are limited to CPUs, there is increasing demand to exploit both CPU clusters and GPU architectures to
67 scale performance and reduce turnaround time. These limitations - high cost, limited flexibility, and insufficient scalability
68 - highlight the urgent need for a new solution. To end this, we developed AdaptiveFlow, an open-source platform designed
69 to intelligently navigate ultra-large chemical spaces, seamlessly incorporate diverse docking methods, and operate efficiently
70 across heterogeneous computing environments, including cloud-scale CPU and GPU resources.

71 AdaptiveFlow is a next-generation, open-source platform purpose-built for navigating ultra-large chemical libraries
72 comprising 69 billion drug-like molecules with scale, speed, and intelligence. A key innovation of the platform is the imple-
73 mentation of Adaptive Target-Guided Virtual Screens (ATG-VSs), which use an 18-dimensional grid of molecular properties
74 to identify and hone in on the most promising regions of chemical space, favorable for a chosen target site, substantially re-
75 ducing computational cost relative to exhaustive docking. An optional active learning layer can be added to further refine and
76 accelerate this process by iteratively selecting the most promising compounds of each tranche. To demonstrate the utility of
77 AdaptiveFlow, we applied it to two therapeutically relevant targets: ferroptosis suppressor protein 1 (FSP1), a key regulator
78 of lipid peroxidation and ferroptotic cell death with emerging roles in cancer resistance (2), and poly(ADP-ribose) polymerase
79 1 (PARP-1), a well-established DNA repair enzyme and clinical target in oncology (3). In both cases, the platform identi-
80 fied nanomolar inhibitors, which were experimentally validated through high-resolution co-crystal structures confirming target
81 engagement.

82 AdaptiveFlow

83 Deploying ultra-large virtual screenings (ULVSs) at the scale of billions of compounds requires not only massive computational
84 resources but also software systems that can efficiently coordinate the preparation, management, and docking of libraries across
85 heterogeneous high-performance computing (HPC) environments. To meet these requirements, we developed AdaptiveFlow,
86 an open-source platform built to streamline, scale, and intelligently guide ULVS workflows from end to end. AdaptiveFlow
87 combines deep flexibility with performance, offering support for diverse ligand and receptor types, GPU acceleration, and
88 modular integration of classical and machine learning-based docking methods.

89 AdaptiveFlow introduces major enhancements across three core components: the AdaptiveFlow Ligand Preparation
90 (AFLP) module for efficient preprocessing of massive libraries; the AdaptiveFlow for Virtual Screening (AFVS) engine, which
91 supports over 1,500 docking methods; and the AdaptiveFlow Unity (AF Unity) module, which integrates these components
92 into a unified, modular workflow adaptable to both CPU and GPU-based infrastructures (overview in Fig. 1 and in Extended
93 Data Table 1). These modules have been engineered to support both traditional and emerging use cases, while optimizing for
94 large-scale parallel execution on cloud and on-premise infrastructures. AdaptiveFlow is natively compatible with cloud services
95 such as AWS (Extended Data Fig. 1), where it demonstrates near-linear scaling up to 5.6 million virtual CPUs (Extended Fig.
96 2), enabling efficient exploration of ultra-large libraries. The use of spot (preemptible) instances further reduces the costs.

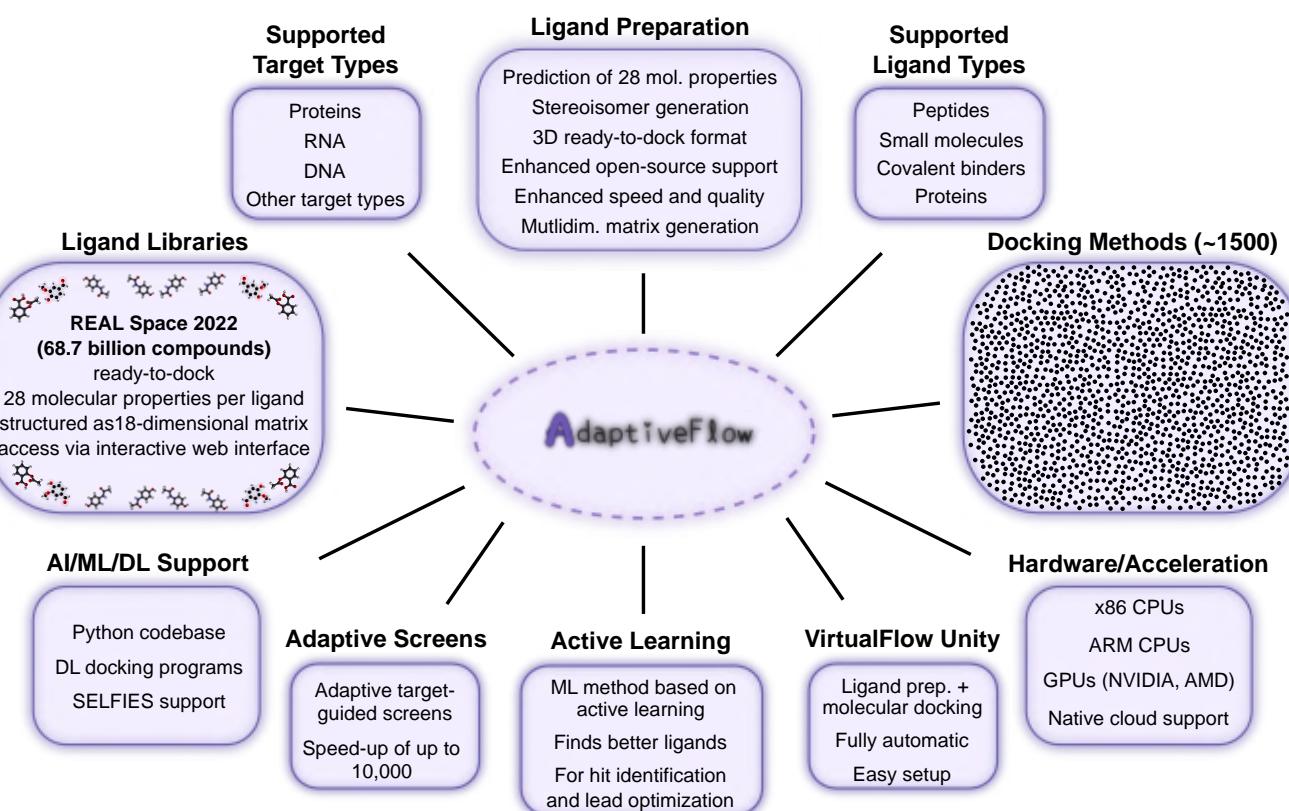


Fig. 1. Overview of features of AdaptiveFlow. AdaptiveFlow introduces several key advances that significantly expand the capabilities of ULVS, including, i) curation of the Enamine REAL Space library, comprising 69 billion molecules in a ready-to-dock format compatible in multiple chemical file formats, ii) integration of an extensive suite of approximately 1,500 docking methods, enabling the targeting of a broad range of receptor types, including RNA and DNA, iii) support for emerging hardware architectures such as ARM CPUs, along with GPU acceleration to enhance computational performance, iv) advanced ligand preparation features, including stereoisomer enumeration and the calculation of 28 molecular properties, to enable the creation of more comprehensive screening libraries, and v) built-in support for machine learning and deep learning approaches, including compatibility with SELFIES and neural network-based docking methods. The integrated module, AdaptiveFlow Unity, further streamlines workflows by combining ligand preparation and docking into a single environment. Importantly, AdaptiveFlow also introduces a novel screening paradigm known as Adaptive Target-Guided Virtual Screens (ATG-VSs), which enables efficient, target-focused exploration of ultra-large compound libraries at a fraction of the computational cost typically associated with conventional ULVSs.

97 The AFLP module provides a powerful pipeline for preparing massive compound libraries for docking. In addition
98 to generating 3D conformers, it enumerates stereoisomers and tautomers, validates geometries, and calculates 28 molecular
99 properties (see Extended Data Table 2, 3; Supplementary Fig. 1, 2, 3). These properties are used to disperse the molecules
100 within a multidimensional grid—up to 18 dimensions—based on physicochemical descriptors. This grid-based representation
101 forms the foundation for Adaptive Target-Guided Virtual Screens (ATG-VSs), which allow users to intelligently select and
102 prioritize chemically diverse subspaces relevant to a given target, thereby optimizing computational resources focusing on the
103 most promising regions of chemical space. Furthermore, an optional active learning component can be layered on top of ATG-
104 VSs to iteratively refine the search using data-driven feedback, dynamically selecting molecules predicted to yield the highest
105 information gain or binding likelihood (Extended Data Fig. 3, 4, 5, 6, 7).

106 The AFVS module enables large-scale docking campaigns using over 1,500 distinct docking methods, assembled from
107 more than 40 individual docking programs, pose prediction algorithms, and scoring functions. AFVS supports a wide range of
108 ligand and target types, including peptides, DNA, and RNA, and includes recent advances in machine learning-based docking
109 such as DiffDock and TANKBind (12, 13), which have shown promise in speed and accuracy, especially for blind or flexible
110 docking scenarios (Extended Data Table 4, 5, 6; Supplementary Table 1, 2; Supplementary Fig. 4).

111 To unify and simplify large-scale workflows, we developed AdaptiveFlow Unity (AFU), a new module that integrates
112 the functionalities of AFLP and AFVS into a cohesive interface. AFU enables users to perform ligand preparation and docking
113 within a single automated pipeline, dramatically reducing setup complexity and increasing accessibility for non-specialists.
114 This streamlined design supports rapid iteration and facilitates integration with other tools in computational drug discovery and
115 AI-driven design workflows (Supplementary Fig. 5, 6, 7, 8, 9, 10, and Methods: AdaptiveFlow Unity).

116 **Enamine REAL Space**

117 Virtual screening performance scales significantly with library size, because larger screening sizes tend to yield both higher
118 potencies and improved hit rates (i.e., the number of confirmed hits divided by the number of tested compounds) (7, 9–11).

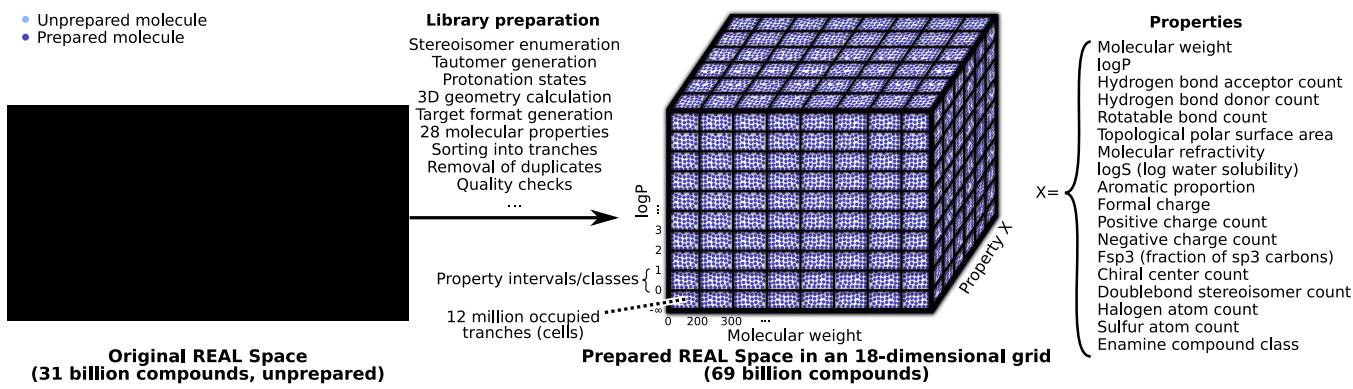


Fig. 2. AdaptiveFlow version of the Enamine REAL Space. The initial version of the REAL Space (enumerated release, 2022q1–2) comprised approximately 31 billion molecules represented as SMILES strings in an unprocessed format—e.g., without full stereoisomer enumeration or structural standardization. A selected subset of 18 of these properties was used to partition the library into an 18-dimensional property matrix, where each dimension was discretized into multiple intervals. In this figure, the high-dimensional matrix is illustrated as a 3D schematic for clarity, with three representative properties each split into 6–10 intervals, resulting in $8 \cdot 6 \cdot 6 = 288$ tranches. Each box represents a tranche that contains ligands sharing the corresponding property ranges. This structure enables flexible selection of compound subsets and serves as the foundation for adaptive, target-guided virtual screens (ATG-VSSs) that are introduced in this work.

Until recently, the largest publicly available ready-to-dock libraries, including the 2018 version of the REAL Database (9) and the ZINC20 library (14), each contained approximately 1.5 billion molecules. To enable virtual screens of an order of magnitude larger scale, we prepared the 2022 version of the Enamine REAL Space into a ready-to-dock format, comprising 68.7 billion commercially available on-demand small molecules.

The current version of the Enamine REAL Space (2022q1–2, accessed November 15, 2022 (15)) contains 31,507,987,117 enumerated molecules (becoming 69 billion after ligand preparation, e.g. due to stereoisomer and tautomer enumeration). These compounds are synthetically accessible derivatives of 137,000 established building blocks, assembled using 167 internally validated one-pot synthetic protocols developed at Enamine (16–25). This combinatorial framework enables high-throughput compound generation with a reported synthetic success rate of 82%, based on 386,000 experimental reactions conducted in 2021. The library is highly drug-like, with 99% of molecules satisfying Lipinski's Rule of Five (26). Structurally, it offers exceptional chemical diversity, comprising 1,098,811,629 unique Murcko scaffolds.

Following ligand preparation with AFLP, the number of ready-to-dock molecules expanded to 68.7 billion in the REAL Space. This transformation involved a comprehensive set of preprocessing steps, including stereoisomer and tautomer enumeration, desalting, neutralization, protonation state prediction, 3D conformer generation, and conversion into formats compatible with a broad range of docking programs. In addition, 28 physicochemical properties and cheminformatic descriptors were calculated for each compound (Extended Data Table 2). To organize this vast chemical space and enable efficient navigation, a subset of 18 key molecular properties was selected to construct an 18-dimensional matrix, in which each axis corresponds to a discrete descriptor such as molecular weight, logP, hydrogen bond donor and acceptor counts, rotatable bond count, topological polar surface area (TPSA), aqueous solubility (logS), aromatic ring count, molecular refractivity (MR), formal charge, positive and negative charge counts, fraction of sp^3 -hybridized carbons (fsp^3), number of chiral centers, halogen and sulfur atom counts, and stereoisomer count. Each dimension was discretized into multiple property intervals, resulting in a combinatorial matrix of 24.4 billion potential tranches (i.e., matrix elements). Molecules from the REAL Space occupy approximately 12 million of these tranches, allowing for highly granular, property-based selection of compound subsets. The 18 properties were chosen to balance chemical relevance for drug discovery with computational efficiency, as excessive dimensionality would lead to sparsity and reduce practical tractability. The binning schemes and occupancy statistics for each dimension are provided in Extended Data Fig. 3.

The REAL Space is not only vast, but also chemically rich and structurally diverse, particularly in its representation of three-dimensional and stereochemically complex molecules (14, 27). Notably, over **50.5 billion molecules (74%)** contain at least one chiral center, contributing to what is often referred to as “nonflat” chemical space—a feature increasingly associated with favorable pharmacological properties (28, 29). Ligand–receptor binding is inherently three-dimensional, and the presence of multiple stereoisomers allows subtle shape complementarity to be exploited. In this context, differences in docking scores across stereoisomers can serve as a proxy for specific target engagement, aiding in the prioritization of high-affinity binders. The REAL Space also includes a substantial number of halogenated compounds—**20.2 billion in total**, including **17.4 billion containing fluorine atoms**—which are well suited for ^{19}F -NMR-based binding assays that do not require isotopically labeled proteins. Despite its breadth, the 18-dimensional matrix remains sparsely populated, revealing large underexplored regions of chemical space. These unoccupied tranches present compelling opportunities for prospective synthesis campaigns to populate novel chemical spaces, enabling the development of molecules to selectively target previously undruggable regions of the proteome.

To facilitate exploration of this high-dimensional chemical library, we developed two interactive web interfaces (Ex-

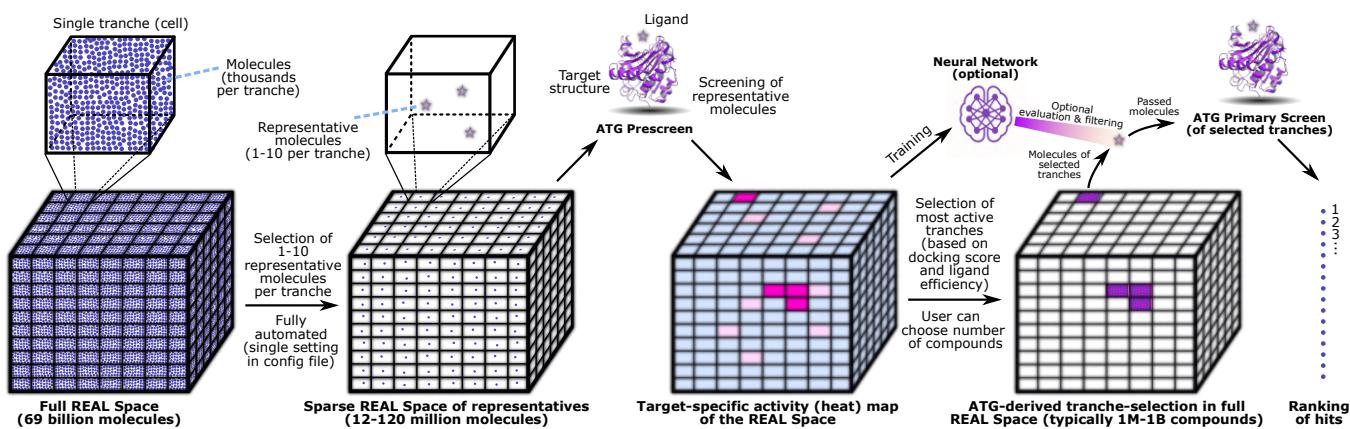


Fig. 3. Conceptual overview of Adaptive Target-Guided Virtual Screens (ATG-VSs). The REAL Space, divided into an 18-dimensional grid, populates over 12 million individual tranches (cells) represented in the figure schematically as a 3-dimensional matrix and each of these tranches has on average 5600 molecules. During the first step of ATG-VSs, the AFVS module of AdaptiveFlow selects a few (1-10) representative molecules in each tranche and docks these molecules to the target receptor. This screening mode (and stage) is referred to as a prescreen and allows AdaptiveFlow to detect the tranches of the library that contain molecules of high potential binding affinity to the given target receptor/binding site. AdaptiveFlow then selects these tranches and screens them fully against the target receptor. In addition, an optional machine learning (ML) model can be trained on the ATG Prescreen results, which can evaluate the molecules of the selected tranches during the ATG Primary Screen.

158 tended Data Fig. 5 and 6), each built with different rendering frameworks to optimize performance across devices and browsers.
 159 These interfaces present the 18-dimensional matrix as a dynamic two-dimensional projection, allowing users visualize the ma-
 160 trix. Each of these two axes of the two dimensional projection is based on one of the 18 properties that define the full matrix,
 161 and the user can select the two properties. Although the visualization is limited to 2D projections, the range of the 18 properties
 162 (of the corresponding 18 dimensions) can be selected interactively independently of the visualization. Selected subsets can
 163 then be used directly for virtual screens via AFVS. The REAL Space is available in PDB, PDBQT, MOL2, and SDF formats,
 164 ensuring compatibility with most structure-based docking programs. Additionally, the library is distributed in SMILES, SELF-
 165 IES, and Parquet formats to support cheminformatics, generative modeling, and large-scale analytics. Both the full library
 166 and user-defined subsets can be accessed via the [AdaptiveFlow project homepage](#) or through the [AWS Open Data repository](#),
 167 enabling direct use in the cloud without the need for local storage.

168 Adaptive Ultra-Large Virtual Screenings

169 Screening ultra-large ligand libraries containing billions of molecules is computationally intensive, especially when using
 170 classical physics-based docking programs. As the number of compounds in the library increases into the tens of billions, the
 171 cost of screening the library becomes prohibitively expensive.

172 Virtual screening is a process of searching a vast chemical space for the best candidate that matches the structural and
 173 biophysical properties of the target site in a given receptor. However, the fraction of compounds in an ultra-large screen that will
 174 optimally engage a specific site in the receptor is often very small and share specific molecular characteristics. To be cost- and
 175 time-effective, it is desirable to focus the screen on the optimal chemical space that is likely to contain hits for a given target.
 176 To rapidly narrow the search space for a given target, we introduce the Adaptive Target-Guided Virtual Screening (ATG-VS)
 177 technique. This approach consists of the following steps:

- 178 **1. Library Decomposition.** On average, a tranche in the current library has approximately 5600 molecules, and for each
 179 tranche up to ten representative ligands are selected for a pilot screen. The number of representatives per tranche (1-10)
 180 can be selected by the user. Since there are approximately 12 million tranches (that contain molecules from the REAL
 181 Space), if one representative per tranche is selected, the total number of molecules screened in the prescreen would be
 182 approximately 12 million for the entire library. If more than one representative molecule, per tranche, is chosen by the
 183 user then these additional molecules are automatically selected to be diverse from each other within the given tranche.
- 184 **2. Screening Representative Molecules (ATG Prescreen).** Instead of screening the entire library, a prescreen of only
 185 the selected representative ligands is performed. In addition, instead of screening representative ligands from all the
 186 12 million tranches, users can apply an optional dynamic filtering procedure that corresponds to the 18-dimensional
 187 properties of the ligands. Here, users can manually apply criteria that the molecules/tranches need to satisfy, e.g. based on
 188 the properties of the targeted site, which can significantly reduce the number of dockings required in the prescreen. After
 189 the prescreen, the sub-space (tranches) of the ligand library that have the best docking scores among the representative
 190 ligands are selected for the next screening stage.
- 191 **3. Selection of Active Tranches.** AdaptiveFlow automatically selects the most promising tranches based on docking scores
 192 of the ATG Prescreen. The total number of ligands to be screened in the next stage can be specified by the user.

193 **4. Active Cell Screening (ATG Primary Screen).** In the next stage, the tranches that were selected after the prescreen
194 are fully screened (e.g., in Fig. 3, the three tranches that contain the best ligands are marked purple). To further re-
195 fine this selection, a machine learning classification model (see **Methods: Machine Learning Classifier for Tranche**
196 **Prioritization**) is employed to identify the most promising molecules within each tranche. The model is trained on
197 Morgan fingerprints (size 1024) of the prescreened compounds, using the docking scores obtained during the prescreen
198 as labels. It applies a threshold-based classification approach, where the 75th percentile of docking scores serves as the
199 decision boundary. Only molecules classified as high-confidence binders are prioritized for full docking, ensuring that
200 computational resources are allocated efficiently to the most promising candidates.

201 ATG-VSs significantly reduce the dimension of the chemical space to be screened, thereby reducing the required com-
202 putation time. AFVS can also prepare new and/or proprietary ligand libraries with ATG-VS support, which requires that the
203 prepared ligand libraries are structured into a multidimensional tranche matrix with pre-assigned representative molecules for
204 each tranche. An optional third-stage screen can be performed to rescore the best hits with higher accuracy, such as by including
205 receptor flexibility. AdaptiveFlow can carry out all three stages in a fully automated workflow.

206 To investigate the effectiveness of the various parameters including the number of representative molecules from each
207 tranche, and the docking exhaustiveness (which is a measure of how extensively the ligands' conformational space is sampled),
208 we carried out multiple ATG prescreens on an array of targets. We examined the following scenarios, i) one representative
209 molecule per tranche and exhaustiveness of 1, ii) 10 representative molecules and exhaustiveness of 1, and iii) 1 representative
210 molecule with exhaustiveness of 10. The higher the exhaustiveness parameter the more thoroughly the conformational space
211 is searched, but also the more computation time is needed. The results of these trials can be seen in Extended Data Fig. 8.
212 As shown in the figure, the results (based on the distribution of the top docking scores, and the ranking of the tranches within
213 each property category) are relatively similar for all three scenarios. Therefore, ATG prescreens with exhaustiveness 1 and one
214 representative molecule per tranche are expected to be sufficient for many applications. In this case, an ATG prescreen involves
215 docking 12 million molecules to explore the 69 billion compound library, which is over 5000 times less computational effort
216 than screening all molecules of the library. If a second stage screen with higher accuracy is added in the context of multistaging,
217 it can be beneficial to increase the exhaustiveness (9).

218 To systematically evaluate the performance of Adaptive Target-Guided Virtual Screens (ATG-VSs) relative to standard
219 ultra-large virtual screens (ULVSs), we conducted benchmarking experiments across ten diverse protein targets, including
220 kinases, phosphatases, GPCRs, and protein–protein interaction interfaces (Extended Data Table 7). These evaluations were
221 performed in two settings: (1) test-based benchmarks using a 5-million-compound subset of chemical space (Fig. 4), and
222 (2) large-scale production runs using the full 69-billion-compound Enamine REAL Space (Extended Data Figs. 9, 10). The
223 production-scale benchmarks compared standard ULVSs to ATG-VS without the optional active learning component, while the
224 subset-based studies also assessed ATG-VS with active learning enabled. Across these targets, we evaluated performance under
225 three key scenarios:

226 (i) Comparison to standard ULVSs (production-scale benchmarks): We evaluated performance at full production scale
227 (Extended Data Figs. 9, 10) by comparing the top 50 docking scores obtained from two approaches: (1) a standard ULVS
228 involving 100 million randomly selected compounds from the full Enamine REAL Space, and (2) the ATG-VS method without
229 active learning. For ATG-VS, we conducted a prescreen of 12 million compounds—comprising 10 representative ligands per
230 tranche across the entire library—followed by a primary screen of up to 1 million compounds. To assess the impact of screening
231 depth, we systematically varied the number of compounds in the primary screen (10K, 100K, and 1M) across ten targets.
232 Active learning was not applied in this setting to isolate and evaluate the impact of tranche prioritization alone; however, based
233 on the test-based benchmarks we expect that incorporating active learning could yield further improvements in efficiency and
234 enrichment. We found that ATG-VS achieves comparable or superior top docking scores to standard ULVSs while requiring
235 dramatically fewer docking evaluations—highlighting its effectiveness and scalability in production-scale virtual screening
236 campaigns.

237 (ii) Impact of machine learning (test-based benchmarks): To further evaluate the benefit of incorporating an ML classi-
238 fier into the ATG workflow, we conducted experiments using a sub-sampled version of the REAL Space comprising 5 million
239 compounds. We compared three approaches: (1) a random screen of 1 million compounds ("1M Random" in Fig. 4); (2)
240 the ATG-VS method, which involved decomposing the library based on 18 molecular properties, performing a prescreen, and
241 selecting compounds for a primary screen—with a total of 100,000 docking evaluations across both stages ("100K ATG"); and
242 (3) an enhanced ATG-VS approach in which a machine learning model, trained on docking scores and Morgan fingerprints
243 from the prescreen, was used to prioritize compounds for the primary screen ("100K ATG+AL"). It should be noted that in
244 the case of ATG and ATG+AL scenarios the total number of docking evaluations including the prescreen was 100K, compared
245 to a million in the random screen. As shown in Fig. 4, the ML-augmented ATG-VS consistently improved the enrichment of
246 high-affinity binders relative to the standard ATG-VS. Though both version of the ATG screen sampled only 100K molecules
247 they outperformed the random screen sampling 1 million molecules.

248 While ATG-based screens consistently outperformed random screening, the extent of improvement varied depending on
249 the biophysical properties of the target. Proteins with complex, feature-rich binding pockets, particularly those combining ex-

250 tensive hydrophobic surfaces with well-defined polar subpockets, derived the greatest benefit from ATG-based prioritization and
251 ML enhancement. Taken together, these findings underscore the scalability, efficiency, and adaptability of ATG-VS. By intel-
252 ligently narrowing the search space—and optionally integrating machine learning for compound prioritization—AdaptiveFlow
253 enables high-performance virtual screening at the scale of tens of billions of molecules.

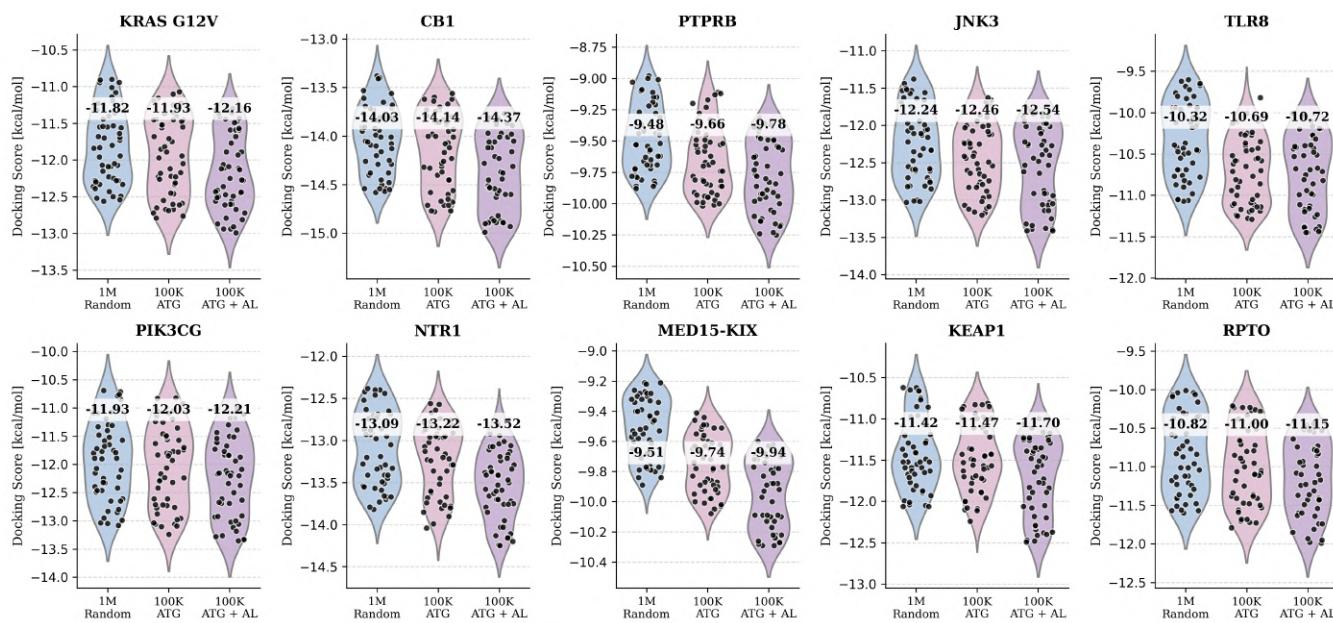


Fig. 4. Benchmarks Involving Adaptive Target-Guided (ATG) Primary Virtual Screens and Standard ULVSs. Violin plots of the docking scores (the more negative, the better) of the top 50 virtual screening hits obtained using three approaches: (i) standard ultra-large virtual screening (ULVS) of 1 million random compounds (blue), (ii) adaptive target-guided (ATG) screening with 100,000 molecules (rose), and (iii) ATG screening augmented with an additional active learning step (purple). Black dots represent individual data points, and the numbers above the violin plots indicate the mean docking scores of the top 50 hits. For most target proteins, ATG primary screens (100,000 molecules) with or without active learning performed comparably to standard ULVS (1 million random compounds), achieving a 10-fold reduction in the number of compounds screened while maintaining similar docking performance.

Machine Learning and Cloud Support

254 AdaptiveFlow introduces expanded capabilities designed to support artificial intelligence (AI) and machine learning (ML) 255 workflows across molecular biology and drug discovery. In addition to integrating state-of-the-art ML-based docking 256 programs—including EquiBind (30), DiffDock (13), and TANKBind (12)—the platform offers several features aimed at 257 facilitating the development and application of ML models. The AFLP module computes a wide range of physicochemical 258 properties commonly used as features or optimization targets in ML pipelines, including the Quantitative Estimate of Druglikeness (QED) 259 (31) and octanol–water partition coefficient (logP) (32). AFLP also supports the generation of SELFIES (Self-Referencing 260 Embedded Strings) molecular representations (33, 34), which are widely adopted in generative modeling due to their syntactic 261 robustness and ability to encode valid molecules across arbitrary token sequences (35) (Extended Data Table 2, 3, 3, 4, 5, 6, 262 7; Supplementary Table 1, 2). Additionally, AFVS introduces numerous deep learning-based docking approaches by 263 modularly assembling scoring functions and pose prediction tools into comprehensive docking workflows. To further 264 streamline ML-driven workflows the AFU module provides a fully integrated pipeline for ligand preparation and docking. This 265 unified interface simplifies data generation for model training and enables seamless incorporation of docking scores as 266 objective functions in generative or optimization-based ML models (36). With support for approximately 1,500 docking methods 267 accessible through a single command-line interface, AFU eliminates the need for users to individually configure and learn 268 multiple docking programs, thus significantly lowering the barrier to entry for ML and drug discovery communities alike. Together, 269 these features position AdaptiveFlow as a versatile platform not only for applying ML methods to virtual screening but also for 270 developing new ML architectures and pipelines for structure-based drug discovery (Extended Data Fig. 11, 12).

271 In parallel with its support for machine learning, AdaptiveFlow is optimized for efficient deployment in cloud computing 272 environments. The platform runs natively on the popular workload manager Slurm (37), widely supported across providers like 273 Google Cloud and Microsoft Azure, and AWS Batch (38). In this study, we demonstrate near-linear scalability on AWS using 274 over 5.6 million virtual CPUs, setting a new global benchmark for parallelization in cloud-based drug discovery. This level of 275 scalability enables the completion of ultra-large virtual screens involving billions of compounds within hours—timelines that 276 would otherwise require weeks or months on traditional on-premise compute clusters. By combining cloud-native performance 277 with ML-native features, AdaptiveFlow serves as a high-performance and accessible foundation for next-generation AI-driven 278 drug discovery.

280 Experimental validation of AdaptiveFlow

281 To experimentally validate the capabilities and performance of AdaptiveFlow, we selected two therapeutically significant target
282 proteins, ferroptosis suppressor protein 1 (FSP1) and poly(ADP-ribose) polymerase 1 (PARP1). FSP1 is a recently charac-
283 terized oxidoreductase that plays a central role in protecting cells from ferroptosis—an iron-dependent form of regulated cell
284 death triggered by membrane lipid peroxidation (39–42). FSP1 maintains pools of reduced coenzyme Q₁₀ and vitamin K hydro-
285 quinone using NAD(P)H, operating independently of GPX4 to suppress oxidative membrane damage. Inhibiting FSP1 disrupts
286 this defense, sensitizing cells—especially redox-imbalanced cancer cells (43)—to ferroptosis. Despite its importance, only a
287 few FSP1 inhibitors have been described and none of their binding to FSP1 has been structurally characterized (2, 44, 45),
288 making it an ideal target for exploring novel inhibitor space using AdaptiveFlow. We also targeted PARP1, a well-established
289 target involved in DNA single-strand break repair, for benchmarking AdaptiveFlow on a mechanistically distinct, clinically
290 validated enzyme class.

291 **Structural Studies of FSP1.** To streamline hit discovery and structural studies, we employed ancestral sequence reconstruc-
292 tion—a method known to yield highly stable proteins while preserving functional and structural fidelity to modern homologs
293 (46, 47). Using this approach, we reconstructed the tetrapod ancestor of human FSP1, which shares 76% sequence identity
294 and retains all active-site residues (Supplementary Fig. 11, 12). The ancestral FSP1 variant expressed efficiently in *E. coli*
295 (~100 mg/L culture), bound FAD, exhibited high thermostability ($T_m = 62$ °C; Supplementary Fig. 12B), and retained full
296 enzymatic activity (Supplementary Table 3; Supplementary Fig. 13 and 14). Crucially, it yielded high-quality, well-diffracting
297 crystals (Supplementary Table 4), enabling detailed structural analyses.

298 We determined high-resolution crystal structures of FSP1 bound to FAD, to FAD and NAD⁺, and in complex with FAD,
299 NAD⁺ and coenzyme Q₁ (Fig. 5; Supplementary Fig. 15a; Supplementary Table 4). The overall structure of the ancestral
300 protein is very similar to that of the human enzyme (Supplementary Fig. 15b). NAD⁺ triggers a 2 Å shift of the flavin ring
301 and its binding mode is fully consistent with a hydride transfer mechanism for flavin reduction as the nicotinamide C4 atom
302 is located right in front of the isoalloxazine N5 (Fig. 5B). Coenzyme Q₁ binds between two aromatic side chains of Phe354
303 and Y290, with its quinone moiety at the edge of the pyrimidine ring of the flavin in a canyon leading from the protein surface
304 (Fig. 5C-D). Specifically, quinone oxygen engages the N(3)H group of the flavin through a hydrogen bond. This geometry and
305 the proximity between the two molecules can enable the facile electron transfer from the flavin to coenzyme Q. Comparison
306 between the FSP1-FAD-NAD⁺ and FSP1-FAD-NAD⁺-coenzyme Q₁ complexes shows that NAD⁺ binding is unaffected by the
307 presence of coenzyme Q. Thus, it is demonstrated that coenzyme Q and NAD(P)⁺ can simultaneously bind to the protein and
308 this structural attribute perfectly aligns with the ternary-complex mechanism evidenced by the kinetics data (Supplementary
309 Fig. 14). The structure also suggests that there is ample space for the oxygen to access the flavin N5 locus where its reaction
310 with the flavin is known to take place. This feature explains why FSP1 can also operate as a NAD(P)H oxidase (48). This
311 detailed experimental visualization of the NAD(P)H and quinone binding modes and their interactions provided the foundation
312 for inhibitor discovery through in silico screenings targeting the substrate binding sites.

313 **Targeting FSP1 with AdaptiveFlow.** To assess the effectiveness of AdaptiveFlow, we screened FSP1 inhibitors targeting its
314 coenzyme Q site using AdaptiveFlow. Our high-resolution crystal structure of FSP1 in complex with FAD and NAD⁺ (PDB
315 code: 9IFT) was used as receptor. Maestro from Schrödinger (51) was applied to prepare the receptor's structure for screening,
316 by removing solvent molecules, assigning missing hydrogen bonds, correcting protonation states at pH 7.4, and assigning
317 correct bond orders. AutoDockTools (v.1.5.7) (52) was utilized to convert the prepared structure from the PDB format to the
318 PDBQT format. The coenzyme Q site of FSP1 (with FAD and NAD⁺ bound to the corresponding sites) was targeted. In
319 the ATG Prescreen, 1 representative compound was used per tranche (resulting in 12 million dockings), and based on these
320 results 10 million compounds were chosen for the primary screen using the ATG-VS method (without ML feature). Basic
321 filters were applied to the top hit candidates, such as logP less than 5, LogS greater than -5, Topological Polar Surface Area
322 (TPSA) smaller than 140 Å², number of stereo centers not greater than 3, and compounds with problematic functional groups
323 and predicted toxicity (mutagenic compounds, compounds with reproductive effects, tumorigenic compounds) were excluded.
324 In total, 42 compounds from 37 different clusters with zero alerts suggested by SwissADME (53) were ordered from Enamine.
325 33 compounds were synthesized and experimentally tested.

326 **Experimental Validation: Inhibitors extend from the coenzyme Q site to the NAD site.** We assessed the potential bind-
327 ing and inhibitory activities of the synthesized 33 virtual hits from AdaptiveFlow, using thermal shift and NADH-depletion
328 assays. Two of the hits elicited stabilizing shifts of at least 1.5 °C in the protein melting temperature and inhibited FSP1 oxi-
329 doreductase activity by at least 50% at a 10 μM concentration. These two compounds were found to potently inhibit coenzyme
330 Q-reduction activity (Table 1; Supplementary Fig. 16). Their K_i values are comparable to, or better than, those reported for
331 previously described FSP1 inhibitors (0.283 μM for afi-FSP1-1 and 0.777 μM for afi-FSP1-2; Supplementary Table 5; Supple-
332 mentary Fig. 17, 18) (2, 44, 45). Critically, a cellular thermal shift assay (CETSA; (54)) performed on FSP1-expressing HEK293
333 cells shows that afi-FSP1-1 effectively engages human FSP1 in cells, as evidenced by a thermal shift of 2.8 °C (Supplementary
334 Fig. 19). We determined the crystal structure of the most potent compound (afi-FSP1-1) bound to FSP1-FAD-NAD⁺ which

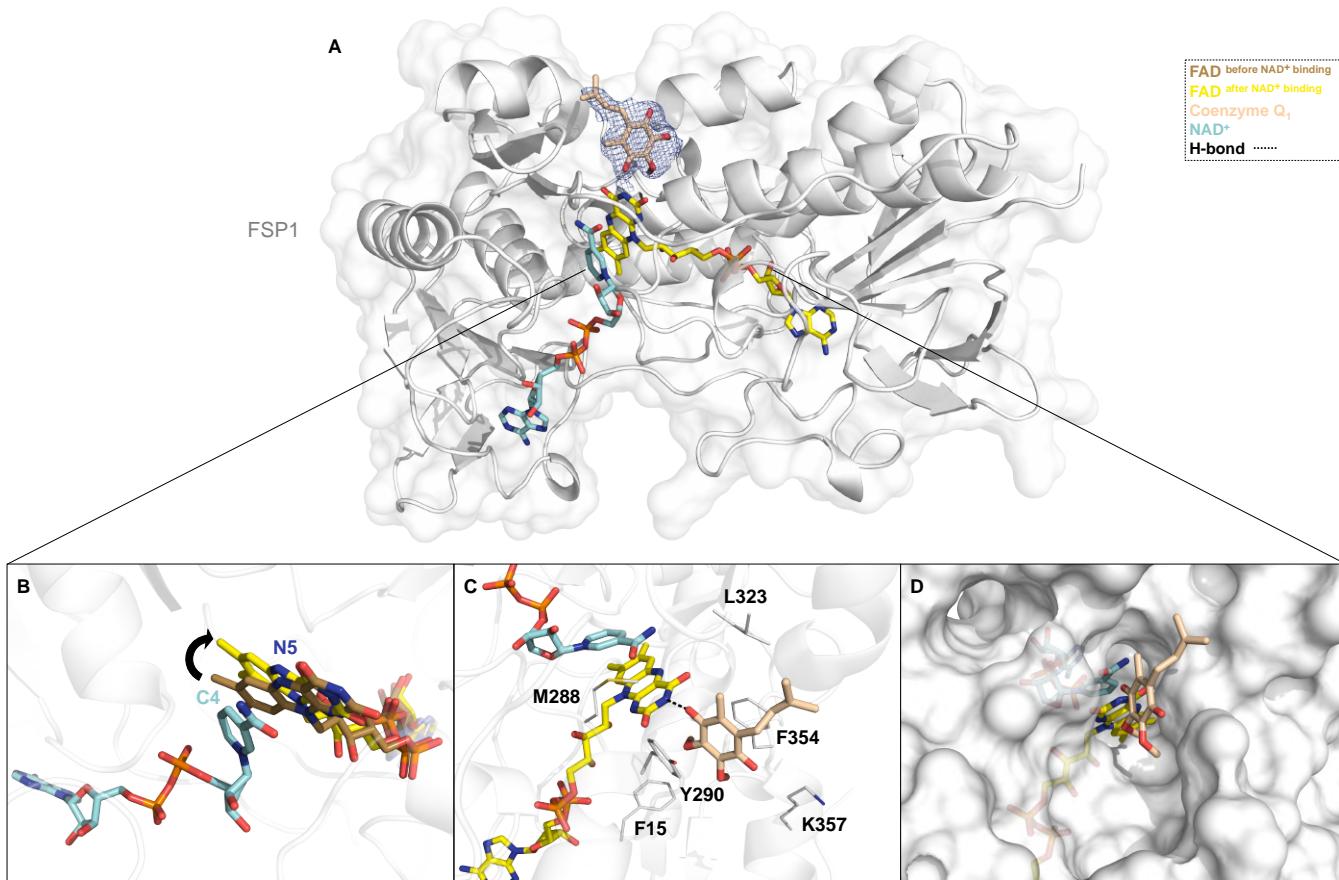


Fig. 5. Co-Crystal Structures of FSP1 in Complex with its Cofactors and Substrate. We determined three crystal structures of FSP1: (i) bound to FAD alone, (ii) bound to FAD and NAD⁺, and (iii) bound to FAD, NAD⁺, and coenzyme Q₁. (A) Overall structure of the fully assembled FSP1 complex with FAD (yellow carbons), NAD⁺ (cyan), and coenzyme Q₁ (beige). The polder omit map for coenzyme Q₁ is contoured at the 3.0 σ level. (B) Superposition of structures (i) and (ii) reveals that NAD⁺ binding induces a rotation of the FAD isoalloxazine ring. The FAD conformation prior to NAD⁺ binding is shown in brown, and the conformation after NAD⁺ binding is shown in yellow. No evidence for 6-hydroxyflavin formation was observed, in contrast to prior reports for human FSP1 (49, 50) (see also Supplementary Fig. 15A). (C) In the ternary complex, coenzyme Q₁ interacts with Tyr290 and Phe354 via edge-to-edge aromatic stacking and forms a hydrogen bond (dashed black line) with the flavin's pyrimidine ring. (D) NAD⁺ and coenzyme Q₁ occupy distinct but converging channels that meet at the FAD active site, illustrating the architecture of the substrate-binding cavity.

reveals that the inhibitor occupies the coenzyme Q binding site in a mode consistent with the predicted docking pose (Fig. 6A-B; Supplementary Fig. 15a). With its elongated conformation, afi-FSP1-1 extensively interacts with the protein by contacting several aromatic side chains, including Tyr290 and Phe354—residues that sandwich the quinone ring of coenzyme Q. Tyr290 is involved in a π -stacking interaction with the benzene ring of the isochroman moiety, whereas Phe354 and Phe15 contribute to the binding via hydrophobic interactions with the piperidine and a π -stacking interaction with the triazole group. The central carbonyl oxygen of afi-FSP1-1 forms an H-bond interaction with the N(3)H group atom of the flavin akin to the interaction established by the carboxyl oxygen of the coenzyme Q's ubiquinone head (Fig. 6B and 5C). Furthermore, the inhibitor's central carbonyl oxygen consistently interacts with the backbone amine group (NH) of Leu323 via a water bridge (Supplementary Fig. 20). Despite its proximity with the inhibitor, the binding conformation of NAD⁺ is identical to that observed both in the absence of a ligand and in the complex with coenzyme Q₁. Similar features characterize the binding compound afi-FSP1-2 whose bulkier structure enables even more extended interaction with the protein (Supplementary Fig. 20, 21). Expanding on these results, we conducted a structure activity relationship study. The three-dimensional structure of the FSP1/afi-FSP1-1 complex was used as input for an additional round of in silico screening. This procedure returned 20 analogues of compound afi-FSP1-1 that were experimentally analyzed (Supplementary Fig. 22). Rewardingly, four additional molecules (afi-FSP1-3-6) displayed inhibitory activities with K_i values in two-digit nanomolar for achiral molecules to low micromolar range for molecules tested as either a racemic mixture or a diastereomeric mixture. Notably, these compounds share an aromatic-amide-saturated heterocycle core as displayed also by afi-FSP1-1 and 2 (Table 1). Consistent with this common feature, the crystal structures of afi-FSP1-3 and afi-FSP1-4 bound to FSP1-FAD-NAD⁺ show that the ligands occupy the coenzyme Q site with their amide carbonyl oxygen engaged by the flavin N(3)H group via a hydrogen bonds and further interacting through a water molecule with the backbone amine group of L323 in FSP1 as observed for afi-FSP1-1 (Fig. 6B-D; Supplementary Fig. 20, 21). Although compound afi-FSP1-4 replaces the dimethyl-tetrahydropyran portion of the isochroman group of afi-FSP1-1 with a methylthiomethyl-triazole, the two inhibitors bind in a very similar way with their common phenyl-triazole-piperidine moiety interacting with Phe15 and Phe354. Conversely, the binding of afi-FSP1-3 is more distinct as it adopts a less extended,

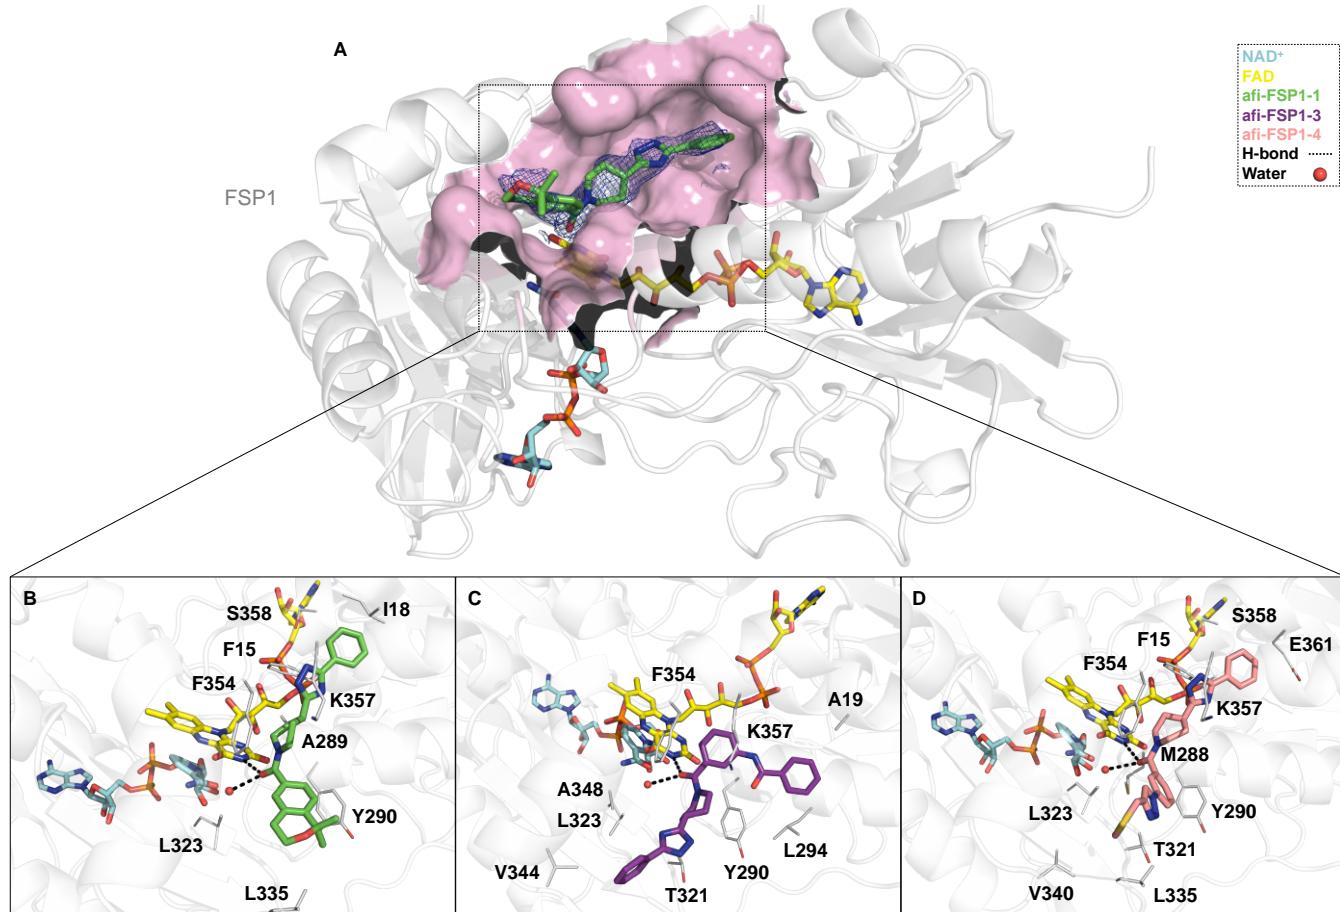


Fig. 6. Co-Crystal Structures of Inhibitors binding to FSP1. (A) Overall structure of FSP1 in complex with inhibitor afi-FSP1-1. The polder omit map for the inhibitor is contoured at the 3.0 σ level. (B-D) Close-up views of the binding site with afi-FSP1-1 (B, green), afi-FSP1-3 (C, violet), and afi-FSP1-4 (D, pink) bound in the region normally occupied by coenzyme Q₁, adjacent to FAD (yellow) and NAD⁺ (cyan). All inhibitors displace coenzyme Q₁ and interact directly with the catalytic cavity.

58 U-shaped conformation laying on Tyr290. These binding modes are validated by mutagenesis experiments (Supplementary
 59 Table 6, 7; Supplementary Fig. 20, 21) and correlate with the measured K_i values. Compounds afi-FSP1-1 and afi-FSP1-4
 60 exhibit practically identical K_i s, whereas compound afi-FSP1-3 displays a twofold decrease in potency. Besides being the best
 61 inhibitors, compounds afi-FSP1-1, afi-FSP1-2, and afi-FSP1-3 exhibit a 10-12-fold reduced efficacy against the NADH oxidase
 62 activity of FSP1, while compound afi-FSP1-4 exhibited no efficacy at all (Supplementary Fig. 24; Supplementary Table 8).
 63 The weaker inhibition reflects the retention of NAD(P)H binding in the inhibited protein as shown by crystal structures and
 64 indicates that reactivity with oxygen is only mildly or not affected by inhibitor binding in the coenzyme Q site. We posit that
 65 preferentially inhibiting the quinone-reducing activity while sparing the reactivity with oxygen, might be pharmacologically
 66 beneficial. Inhibitors of this type may convert the enzyme from being anti-ferroptotic (coenzyme Q reducing activity) to being
 67 pro-ferroptotic (uninhibited lipid-damaging ROS production). Future work is needed to explore and validate this concept in
 68 FSP1 pharmacology.

69 **Applying AdaptiveFlow to PARP1.** To further validate the effectiveness of the AdaptiveFlow platform, we selected hu-
 70 man Poly(ADP-ribose) polymerase 1 (PARP1) as a benchmark target due to its well-characterized structure, critical role in
 71 DNA repair, and relevance in cancer therapy. Following the ATG-prescreen and a primary screen of 100 million molecules,
 72 160 promising candidates were synthesized for further experimental verification, as detailed in the Methods section. A two-
 73 concentration highthroughput enzymatic screen identified seven PARP1 inhibitors, four of which (iParp1, iParp2, iParp3, and
 74 iParp4) showed sub-250 nM IC₅₀ values and were selected for further characterization. Protein NMR confirmed direct binding
 75 to the active site of PARP1, with chemical shift perturbations supporting the predicted docking poses. X-ray crystallography
 76 revealed the 2.05 Å structure of the PARP1 catalytic domain bound to a hydrolyzed form of iParp1, validating its binding mode
 77 and key interactions. Clonogenic assays in BRCA1-deficient TNBC cells demonstrated selective cytotoxicity for iParp1 and
 78 iParp2, though reduced cellular activity, especially for iParp1, was attributed to poor membrane permeability due to hydroly-
 79 sis. Remarkably, AdaptiveFlow identified PARP1 inhibitors with submicromolar to low nanomolar potency, including iParp1,
 80 which exhibited an IC₅₀ of 8.8 nM—comparable to the FDA-approved drug olaparib. This result is particularly notable given
 81 that our hits were derived from a purely virtual screen without any prior knowledge of known inhibitors or medicinal chemistry

Compound	Structure	K_i (μM)	K_d (μM)
afi-FSP1-1		$0.283^b \pm 0.05$	0.098
afi-FSP1-2		0.777 ± 0.22	0.266
afi-FSP1-3		0.520 ± 0.11	0.184
afi-FSP1-4		0.284 ± 0.07	0.198
afi-FSP1-5		1.00 ± 0.32	n.d.
afi-FSP1-6		3.94 ± 0.71	n.d.

Table 1. Competitive K_i and Binding K_d Values of the Most Potent FSP1 Inhibitors. K_i values were determined by monitoring NADH oxidation at 340 nm (pH 7.2) using variable NADH concentrations (5–200 μM), 100 μM coenzyme Q₁, and 0.1–0.2 μM FSP1 (Supplementary Fig 16). K_d values were derived from independent binding experiments based on surface plasmon resonance (SPR, Supplementary Fig. 18). The experiments were performed using the racemic mixture of compound afi-FSP1-5 and the diastereomeric mixture of compound afi-FSP1-6. Inhibitors' concentrations ranged from micromolar to low nanomolar. For comparison, Supplementary Table 9 lists IC₅₀ values obtained using the indirect resazurin assay, commonly used in previous literature (2, 44, 45).

optimization of the hit. Structural similarity analyses confirmed that several hits shared key scaffolds with known PARP1 inhibitors, while others represented new chemical motifs, demonstrating AdaptiveFlow's capacity to identify both validated and novel chemotypes. Details can be found in the Supplementary Information. Overall, this validation exemplifies AdaptiveFlow's ability to efficiently and accurately discover drug-like molecules with clinically relevant potency, streamlining early-stage drug discovery. See Supplementary Information; Supplementary Fig. 26, 27, 28, 29, 31, 30; Supplementary Table 10.

Discussion

The accessible chemical space for drug discovery is expanding at an unprecedented pace, with the number of purchasable, on-demand molecules growing from approximately 2 billion in 2020 to trillions of molecules. Yet, even this remarkable growth represents only a minuscule fraction of the estimated 10^{60} drug-like molecules that populate theoretical chemical space (5). Identifying high-quality hits within this vast and ever-expanding landscape remains a central challenge in early-stage drug discovery. To address this, we curated the largest ready-to-dock library to date, the Enamine REAL Space—comprising 69 billion commercially accessible small molecules, fully preprocessed and made available in multiple formats for structure-based virtual screening.

While ultra-large libraries increase the likelihood of identifying potent binders and enhance true-hit rates (7, 9–11, 55, 56), exhaustively screening billions of compound impose substantial computational costs, even with today's cloud-based computing capabilities. AdaptiveFlow addresses this challenge through Adaptive Target-Guided Virtual Screening (ATG-VS), a highly efficient prescreening strategy that narrows docking efforts to target-specific, chemically diverse subspaces enriched for high-affinity candidates. When applied to the 69-billion-compound REAL Space, ATG-VS reduces screening costs by up to 1,000-fold compared to exhaustive searches, while preserving strong enrichment of top hits. Additionally, AdaptiveFlow incorporates deep learning-based docking and GPU acceleration, achieving a further 10- to 100-fold increase in throughput (57, 58). To experimentally validate the practical utility of AdaptiveFlow, we focused on two biologically and clinically relevant targets. First, we targeted ferroptosis suppressor protein 1 (FSP1), a redox enzyme involved in lipid peroxide detoxification and resistance to ferroptosis, and identified a series of nanomolar inhibitors. Structural characterization of FSP1 in complex with cofactors and ligands revealed that both NAD(P)H and coenzyme Q can simultaneously bind in catalytically competent orientations, providing mechanistic insights into its function. The identified hits were validated through crystallography and biochemical assays. We also applied AdaptiveFlow to poly(ADP-ribose) polymerase 1 (PARP1), a clinically validated oncology

408 target central to DNA damage repair. Our ultra-large screen yielded seven inhibitors, including iParp1, which displayed an IC₅₀
409 comparable to that of the FDA-approved drug olaparib (59). These results underscore the power of AdaptiveFlow in rapidly
410 identifying potent hits across diverse protein classes and demonstrate the advantage of adaptive screening over brute-force enum-
411 eration. Notably, while our previous identification of nanomolar KEAP1–NRF2 inhibitors required an unbiased screen of 1.3
412 billion molecules, the ATG-guided screen for FSP1 achieved comparable success with just 22 million docking instances. Al-
413 though the effectiveness of chemical space focusing is target-dependent, our benchmarking and experimental studies highlight
414 the potential of ATG-VS to tailor the search space to specific binding pockets, significantly enhancing screening efficiency. As
415 chemical space continues to expand, focusing on target-specific subspaces will no longer be optional, but essential.

416 With the infrastructure now in place, adding molecules to the library and directing them to their respective tranches
417 is a computational process that scales linearly. As new tranches, beyond the 12 million currently populated are filled, only
418 representatives from these newly populated tranches need to be incorporated into the prescreen set. The open-source nature of
419 AdaptiveFlow enables seamless adaptation, extension, and development of the platform by both academic and industry users.
420 With the increasing availability of high-resolution structural data, enabled by techniques such as cryo-EM and de novo struc-
421 ture prediction, combined with site-specific functional insights from gene editing technologies, *in silico* screening is poised to
422 deliver chemical tools capable of modulating precise biomolecular functions. As modalities like molecular glues and targeted
423 protein degraders mature, the ability to identify small molecules that bind specific protein surfaces with functional consequences
424 is becoming increasingly valuable. AdaptiveFlow’s modular architecture supports over 1,500 docking protocols, encompassing
425 combinations of scoring functions, sampling algorithms, and ML-based pose predictors such as DiffDock and TANKBind. The
426 platform also natively supports SELFIES representations, physicochemical property calculations, and active learning–driven
427 workflows, making it fully compatible with generative models and reinforcement learning pipelines (60). As generative models
428 continue to advance, platforms like AdaptiveFlow will be essential for rapidly evaluating designed molecules against biolog-
429 ically relevant structural targets. In conclusion, AdaptiveFlow offers a unified, open-source, and massively scalable solution
430 for ultra-large virtual screening. It provides seamless integration of conventional and AI-driven docking methods, access to the
431 largest ready-to-dock chemical library, and adaptive strategies for efficient navigation of trillion-scale chemical space. With
432 demonstrated success across diverse therapeutic targets, AdaptiveFlow stands as a next-generation infrastructure poised to
433 accelerate the design and discovery of future therapeutics.

434 References

1. Oleksandr O. Grygorenko, Dmytro S. Radchenko, Igor Dziuba, Alexander Chuprina, Kateryna E. Gubina, and Yurii S. Moroz. Generating Multibillion Chemical Space of Readily Accessible Screening Compounds. *iScience*, 23(11):101681, 2020. ISSN 25890042. doi: 10.1016/j.isci.2020.101681.
2. Sebastian Doll, Florencio Porto Freitas, Ron Shah, Maceler Aldrovandi, Milena Costa da Silva, Irina Ingold, Andrea Goya Grocin, Thamara Nishida Xavier da Silva, Elena Panzilius, Christina H Scheel, et al. Fsp1 is a glutathione-independent ferroptosis suppressor. *Nature*, 575(7748):693–698, 2019.
3. Christopher J Lord and Alan Ashworth. Parp inhibitors: Synthetic lethality in the clinic. *Science*, 355(6330):1152–1158, 2017.
4. Steven M Mennen, Carolina Alhambra, C Liana Allen, Mario Barberis, Simon Berritt, Thomas A Brandt, Andrew D Campbell, Jesús Castaño, Alan H Cherney, Melodie Christensen, et al. The evolution of high-throughput experimentation in pharmaceutical development and perspectives on the future. *Organic Process Research & Development*, 23(6):1213–1242, 2019.
5. Regine S. Bohacek, Colin McMartin, and Wayne C. Guida. The art and practice of structure-based drug design: A molecular modeling perspective. *Medicinal Research Reviews*, 16(1):3–50, 1996. ISSN 01986325. doi: 10.1002/(SICI)1098-1128(199601)16:1<3::AID-MED1>3.0.CO;2-6.
6. Sally R Ellingson, Sivanesan Dakshanamurthy, Milton Brown, Jeremy C Smith, and Jerome Baudry. Accelerating virtual high-throughput ligand docking: current technology and case study on a petascale supercomputer. *Concurrency and Computation: Practice and Experience*, 26(6):1268–1277, 2014.
7. Jiankun Lyu, Sheng Wang, Trent E Balias, et al. Ultra-large library docking for discovering new chemotypes. *Nature*, 2019. ISSN 1476-4687. doi: 10.1038/s41586-019-0917-9.
8. Reed M Stein, Hye Jin Kang, John D McCorry, Grant C Glafelter, Anthony J Jones, Tao Che, Samuel Slocum, Xi-Ping Huang, Olena Savych, Yurii S Moroz, Benjamin Stauch, Linda C Johansson, Vadim Cherezov, Terry Kenakin, John J Irwin, Brian K Shoichet, Bryan L Roth, and Margarita L Dubocovich. Virtual discovery of melatonin receptor ligands to modulate circadian rhythms. *Nature*, feb 2020. ISSN 0028-0836. doi: 10.1038/s41586-020-2027-0.
9. Christoph Gorgulla, Andras Boeszoermenyi, Zi-fu Wang, Patrick D Fischer, Paul W. Coote, Krishna M. Padmanabha Das, Yehor S Malets, Dmytro S Radchenko, Yurii S Moroz, David A Scott, Konstantin Fackeldey, Moritz Hoffmann, Iryna Iavniuk, Gerhard Wagner, and Haribabu Arthanari. An open-source drug discovery platform enables ultra-large virtual screens. *Nature*, 580(7805):663–668, apr 2020. ISSN 0028-0836. doi: 10.1038/s41586-020-2117-z.
10. Assaf Alon, Jiankun Lyu, Joao M. Braz, Tia A. Tummino, Veronica Craik, Matthew J. O’Meara, Chase M. Webb, Dmytro S. Radchenko, Yurii S. Moroz, Xi-Ping Huang, Yongfeng Liu, Bryan L. Roth, John J. Irwin, Allan I. Basbaum, Brian K. Shoichet, and Andrew C. Kruse. Structures of the $\sigma 2$ receptor enable docking for bioactive ligand discovery. *Nature*, 600(7890):759–764, dec 2021. ISSN 0028-0836. doi: 10.1038/s41586-021-04175-x.
11. Christoph Gorgulla. *Free Energy Methods Involving Quantum Physics , Path Integrals , and Virtual Screenings: Development, Implementation and Application in Drug Discovery*. PhD thesis, Freie Universität Berlin, 2018.
12. Wei Lu, Qifeng Wu, Jixian Zhang, Jiahua Rao, Chengtao Li, and Shuangjia Zheng. TANKBind: Trigonometry-Aware Neural Networks for Drug-Protein Binding Structure Prediction. *bioRxiv*, 2022. doi: 10.1101/2022.06.06.495043v1.
13. Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Diffdock: Diffusion steps, twists, and turns for molecular docking. *arXiv preprint arXiv:2210.01776*, 2022.
14. John J Irwin, Khanh G Tang, Jennifer Young, Chinziray Dandarchuluun, Benjamin R Wong, Munkhzul Khurelbaatar, Yurii S Moroz, John Mayfield, and Roger A Sayle. ZINC20—A Free Ultralarge-Scale Chemical Database for Ligand Discovery. *Journal of Chemical Information and Modeling*, 60(12):6065–6073, dec 2020. ISSN 1549-9596. doi: 10.1021/acs.jcim.0c00675.
15. Enamine Ltd. Real compounds collection, 2022. Accessed November 15, 2022.
16. Andrey V Bogolubsky, Yurii S Moroz, Pavel K Mykhailiuk, Dmitry S Granat, Sergey E Pipko, Anzhelika I Konovets, Roman Doroschuk, and Andrey Tolmachev. Bis (2, 2, 2-trifluoroethyl) carbonates as a condensing agent in one-pot parallel synthesis of unsymmetrical aliphatic ureas. *ACS Combinatorial Science*, 16(6):303–308, 2014.
17. Andrey V Bogolubsky, Yurii S Moroz, Pavel K Mykhailiuk, Sergey E Pipko, Anzhelika I Konovets, Irina V Sadkova, and Andrey Tolmachev. Sulfonyl fluorides as alternative to sulfonyl chlorides in parallel synthesis of aliphatic sulfonamides. *ACS Combinatorial Science*, 16(4):192–197, 2014.
18. Andrey V Bogolubsky, Yurii S Moroz, Pavel K Mykhailiuk, Dmitry M Panov, Sergey E Pipko, Anzhelika I Konovets, and Andrey Tolmachev. A one-pot parallel reductive amination of aldehydes with heteroaromatic amines. *ACS Combinatorial Science*, 16(8):375–380, 2014.
19. Andrey V Bogolubsky, Yurii S Moroz, Pavel K Mykhailiuk, Eugeniy N Ostapchuk, Alexander V Rudnichenko, Yurii V Dmytriv, Anna N Bondar, Olga A Zaporozhets, Sergey E Pipko, Roman A Doroschuk, et al. One-pot parallel synthesis of alkyl sulfides, sulfoxides, and sulfones. *ACS Combinatorial Science*, 17(6):348–354, 2015.
20. Andrey V Bogolubsky, Yurii S Moroz, Pavel K Mykhailiuk, Sergey E Pipko, Anton V Zhemera, Anzhelika I Konovets, Olena O Stepaniuk, Inna S Myronchuk, Yurii V Dmytriv, Roman A Doroschuk, et al. 2, 2, 2-trifluoroethyl chlorooxacetate universal reagent for one-pot parallel synthesis of n-1-aryl-n-2-alkyl-substituted oxamides. *ACS Combinatorial Science*, 17(10):615–622, 2015.
21. Andrey Tolmachev, Andrey V Bogolubsky, Sergey E Pipko, Alexander V Grishchenko, Dmytro V Ushakov, Anton V Zhemera, Oleksandr O Vinylchuk, Anzhelika I Konovets, Olga A Zaporozhets, Pavel K Mykhailiuk, et al. Expanding synthesizable space of disubstituted 1, 2, 4-oxadiazoles. *ACS combinatorial science*, 18(10):616–624, 2016.
22. Andrey V Bogolubsky, Yurii S Moroz, Olena Savych, Sergey Pipko, Angelika Konovets, Maxim O Platonov, Oleksandr V Vasylchenko, Vasyl V Hurnach, and Oleksandr O Grygorenko. An old story in the parallel synthesis world: an approach to hydantoin libraries. *ACS Combinatorial Science*, 20(1):35–43, 2018.

478 23. Olena Savych, Yuliya O Kuchkovska, Andrey V Bogolyubsky, Anzhelika I Konovets, Kateryna E Gubina, Sergey E Pipko, Anton V Zhemera, Alexander V Grishchenko, Dmytro N Khomenko, Volodymyr S Brovarets, et al. One-pot parallel synthesis of 5-(dialkylamino) tetrazoles. *ACS combinatorial science*, 21(9):635–642, 2019.

479 24. OO Grygorenko, DS Radchenko, I Dziuba, A Chuprina, KE Gubina, and YS Moroz. Generating multibillion chemical space of readily accessible screening compounds. *iscience* 2020, 23 (11), 101681, 2020.

480 25. CA Lipinski, F Lombardo, BW Dominy, and PJ Feeney. In vitro models for selection of development candidatesexperimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Adv Drug Deliv Rev*, 23(1):3–25, 1997.

481 26. Christopher A Lipinski, Franco Lombardo, Beryl W Dominy, and Paul J Feeney. Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Advanced drug delivery reviews*, 23(1-3):3–25, 1997.

482 27. Anna Tomberg and Jonas Bostrom. Can easy chemistry produce complex, diverse, and novel molecules? *Drug Discovery Today*, 00(00):1–8, oct 2020. ISSN 13596446. doi: 10.1016/j.drudis.2020.09.027.

483 28. Frank Lovering, Jack Bikker, and Christine Humbel. Escape from flatland: increasing saturation as an approach to improving clinical success. *Journal of medicinal chemistry*, 52(21):6752–6756, 2009.

484 29. Frank Lovering. Escape from flatland 2: complexity and promiscuity. *MedChemComm*, 4(3):515–519, 2013.

485 30. Hannes Stärk, Octavian Ganea, Lagnajit Pattanaik, Regina Barzilay, and Tommi Jaakkola. Equibind: Geometric deep learning for drug binding structure prediction. In *International Conference on Machine Learning*, pages 20503–20521. PMLR, 2022.

486 31. G. Richard J. Bickerton, Gaia Valeria Paolini, Jérémie Besnard, Sorel Muresan, and Andrew L. Hopkins. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4:290–8, 2012.

487 32. Scott A Wildman and Gordon M Crippen. Prediction of physicochemical parameters by atomic contributions. *Journal of chemical information and computer sciences*, 39(5):868–873, 1999.

488 33. Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alan Aspuru-Guzik. Self-referencing embedded strings (selfies): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1(4):045024, oct 2020. doi: 10.1088/2632-2153/aba947.

489 34. Mario Krenn, Qianxiang Ai, Senja Barthel, Nessa Carson, Angelo Frei, Nathan C Frey, Pascal Friederich, Théophile Gaudin, Alberto Alexander Gayle, Kevin Maik Jablonka, et al. Selfies and the future of molecular string representations. *Patterns*, 3(10):100588, 2022.

490 35. Daniel Flam-Shepherd, Kevin Zhu, and Alán Aspuru-Guzik. Language models can learn complex molecular distributions. *Nature Communications*, 13(1):1–10, 2022.

491 36. AkshatKumar Nigam, Robert Pollice, Gary Tom, Kjell Jorner, John Wiles, Luca Thiede, Anshul Kundaje, and Alan Aspuru-Guzik. Tartarus: A benchmarking platform for realistic and practical inverse molecular design. *Advances in Neural Information Processing Systems*, 36, 2024.

492 37. SLURM. <https://slurm.schedmd.com>, Accessed: 2018-12-22.

493 38. AWS Batch. <https://aws.amazon.com/batch/>, Accessed: 2022-12-31.

494 39. Ashley R Brown, Tal Hirschhorn, and Brent R Stockwell. Ferroptosis—disease perils and therapeutic promise. *Science*, 386(6724):848–849, 2024.

495 40. LE Crooks, J Watts, J Hoenninger, M Arakawa, L Kaufman, H Guenther, and D Feinberg. Thin-section definition in magnetic resonance imaging. technical concepts and their implementation. *Radiology*, 154(2):463–467, 1985.

496 41. Mike Lange, Michele Wölk, Cody E Doubravsky, Joseph M Hendricks, Shunji Kato, Yurika Otoki, Benjamin Styler, Kiyotaka Nakagawa, Maria Fedorova, and James A Olzmann. Fsp1-mediated lipid droplet quality control prevents neutral lipid peroxidation and ferroptosis. *bioRxiv*, pages 2025–01, 2025.

497 42. Eikan Mishima, Junya Ito, Zijun Wu, Toshitaka Nakamura, Adam Wahida, Sebastian Doll, Wulf Tonnus, Palina Nepachalovich, Elke Eggenhofer, Maceler Aldrovandi, et al. A non-canonical vitamin k cycle is a potent ferroptosis suppressor. *Nature*, 608(7924):778–783, 2022.

498 43. Helmut Sies, Ryan J Mailloux, and Ursula Jakob. Fundamentals of redox regulation in biology. *Nature Reviews Molecular Cell Biology*, 25(9):701–719, 2024.

499 44. Joseph M Hendricks, Cody E Doubravsky, Eddie Wehri, Zhipeng Li, Melissa A Roberts, Kirandeep K Deol, Mike Lange, Irene Lasheras-Otero, Jeremiah D Mopper, Scott J Dixon, et al. Identification of structurally diverse fsp1 inhibitors that sensitize cancer cells to ferroptosis. *Cell chemical biology*, 30(9):1090–1103, 2023.

500 45. Toshitaka Nakamura, Eikan Mishima, Naoya Yamada, André Santos Dias Mourão, Dietrich Trümback, Sebastian Doll, Jonas Wanninger, Elena Lytton, Peter Sennhenn, Thamara Nishida Xavier da Silva, et al. Integrated chemical and genetic screens unveil fsp1 mechanisms of ferroptosis regulation. *Nature structural & molecular biology*, 30(11):1806–1815, 2023.

501 46. Raine ES Thomson, Saskya E Carrera-Pacheco, and Elizabeth MJ Gillam. Engineering functional thermostable proteins using ancestral sequence reconstruction. *Journal of Biological Chemistry*, 298(10), 2022.

502 47. Callum R Nicoll, Marta Massari, Marco W Fraaije, Maria Laura Mascotti, and Andrea Mattevi. Impact of ancestral sequence reconstruction on mechanistic and structural enzymology. *Current Opinion in Structural Biology*, 82:102669, 2023.

503 48. Teresa Benedetta Guerriere, Alessandro Vancheri, Ilaria Ricotti, Stefano A Serapian, Daniel Eggerichs, Dirk Tischler, Giorgio Colombo, Maria L Mascotti, Marco W Fraaije, and Andrea Mattevi. Dehydrogenase versus oxidase function: the interplay between substrate binding and flavin microenvironment. *ACS catalysis*, 15:1046–1060, 2025.

504 49. Sitoa Zhang, Shengsong Gou, Qian Zhang, Xin Yong, Boyi Gan, and Da Jia. Fsp1 oxidizes nadph to suppress ferroptosis. *Cell research*, 33(12):967–970, 2023.

505 50. Shijian Feng, Xiaofang Huang, Dan Tang, Xiaoyu Liu, Liang Ouyang, Dehua Yang, Kunjie Wang, Banghua Liao, and Shiqian Qi. The crystal structure of human ferroptosis suppressive protein 1 in complex with flavin adenine dinucleotide and nicotinamide adenine nucleotide. *MedComm*, 5(3):e479, 2024.

506 51. Schrödinger, LLC. Maestro, 2024. Schrödinger Release 2024-1.

507 52. GM Morris, R Huey, and AJ Olson. Using autodock for ligand-receptor docking. *current protocols in bioinformatics*. *Unitas*, 8(14):5, 2008.

508 53. Antoine Daina, Olivier Michelin, and Vincent Zoete. Swissadme: a free web tool to evaluate pharmacokinetics, drug-likeness and medicinal chemistry friendliness of small molecules. *Scientific reports*, 7(1):42717, 2017.

509 54. Daniel Martinez Molina, Rozbeh Jafari, Marina Ignatushchenko, Takahiro Seki, E Andreas Larsson, Chen Dan, Lekshmy Sreekumar, Yihai Cao, and Paer Nordlund. Monitoring drug target engagement in cells and tissues using the cellular thermal shift assay. *Science*, 341(6141):84–87, 2013.

510 55. David E. Gloriam. Bigger is better in virtual drug screens. *Nature*, 566(7743):193–194, feb 2019. ISSN 0028-0836. doi: 10.1038/d41586-019-00145-6.

511 56. AkshatKumar Nigam, Matthew FD Hurley, Fengling Li, Eva Konkočová, Martin Klíma, Jana Trýlová, Robert Pollice, Süleyman Selim Çınaroglu, Roni Levin-Konigsberg, Jasemine Handjaya, et al. Drug discovery in low data regimes: Leveraging a computational pipeline for the discovery of novel sars-cov-2 nsp14-mtase inhibitors. *bioRxiv*, pages 2023–10, 2023.

512 57. Diogo Santos-Martins, Leonardo Solis-Vasquez, Andreas F Tillack, Michel F Sanner, Andreas Koch, and Stefano Forli. Accelerating autodock4 with gpus and gradient-based local search. *Journal of chemical theory and computation*, 17(2):1060–1073, 2021.

513 58. Shidi Tang, Ruiqi Chen, Mengru Lin, Qingde Lin, Yanxiang Zhu, Ji Ding, Haifeng Hu, Ming Ling, and Jiansheng Wu. Accelerating autodock vina with gpus. *Molecules*, 27(9):3041, 2022.

514 59. Giorgio Valabrega, Giulia Scotti, Valentina Tuninetti, Arianna Pani, and Francesco Scaglione. Differences in parp inhibitors for the treatment of ovarian cancer: mechanisms of action, pharmacology, safety, and efficacy. *International Journal of Molecular Sciences*, 22(8):4203, 2021.

515 60. Benjamin Sanchez-Lengeling and Alán Aspuru-Guzik. Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400):360–365, 2018.

516 61. Noel M O’Boyle, Michael Banck, Craig A James, Chris Morley, Tim Vandermeersch, and Geoffrey R Hutchison. Open Babel: An open chemical toolbox. *Journal of Cheminformatics*, 3(1):33, December 2011. ISSN 1758-2946. doi: 10.1186/1758-2946-3-33.

517 62. David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36, 1988. doi: 10.1021/ci00057a005.

518 63. Greg Landrum et al. Rdkit: Open-source cheminformatics, 2006.

519 64. Martin Buttenschoen, Garrett M. Morris, and Charlotte M. Deane. PoseBusters: AI-based docking methods fail to generate physically valid poses or generalise to novel sequences. *arXiv*, 2023. doi: 10.1039/D3SC04185A.

520 65. Amr Alhossary, Stephanus Daniel Handoko, Yuguang Mu, and Chee-Keong Kwoh. Fast, accurate, and reliable molecular docking with QuickVina 2. *Bioinformatics*, 31(13):2214–2216, jul 2015. ISSN 1460-2059. doi: 10.1093/bioinformatics/btv082.

521 66. Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Diffdock: Diffusion steps, twists, and turns for molecular docking. *arXiv preprint arXiv:2210.01776*, 2022.

522 67. Hannes Stärk, Octavian Ganea, Lagnajit Pattanaik, Regina Barzilay, and Tommi Jaakkola. Equibind: Geometric deep learning for drug binding structure prediction. In *International Conference on Machine Learning*, pages 20503–20521. PMLR, 2022.

523 68. Andrew P McNutt, Paul Francoeur, Rishal Aggarwal, Tomohide Masuda, Rocco Meli, Matthew Ragoza, Jocelyn Sunseri, and David Ryan Koes. Gnina 1.0: molecular docking with deep learning. *Journal of cheminformatics*, 13(1):1–20, 2021.

524 69. Jacob D Durrant and J Andrew McCammon. Nnscore 2.0: a neural-network receptor-ligand scoring function. *Journal of chemical information and modeling*, 51(11):2897–2903, 2011.

525 70. Haiping Zhang, Linbu Liao, Konda Mani Saravanan, Peng Yin, and Yanjie Wei. Deepbindrg: a deep learning based method for estimating effective protein-ligand affinity. *PeerJ*, 7:e7362, 2019.

526 71. Mostafa Karimi, Di Wu, Zhangyang Wang, and Yang Shen. Deepaffinity: interpretable deep learning of compound–protein affinity through unified recurrent and convolutional neural networks. *Bioinformatics*, 35(18):3329–3338, 2019.

527 72. Anthony J Clark, Pratyush Tiwary, Ken Borrelli, Shulu Feng, Edward B Miller, Robert Abel, Richard A Friesner, and Bruce J Berne. Prediction of protein–ligand binding poses via a combination of induced fit docking and metadynamics simulations. *Journal of chemical theory and computation*, 12(6):2990–2998, 2016.

528 73. Edward B Miller, Robert B Murphy, Daniel Sindhikara, Kenneth W Borrelli, Matthew J Grisewood, Fabio Ranalli, Steven L Dixon, Steven Jerome, Nicholas A Boyles, Tyler Day, et al. Reliable and accurate solution to the induced fit docking problem for protein–ligand binding. *Journal of Chemical Theory and Computation*, 17(4):2630–2639, 2021.

529 74. Ercheng Wang, Huiyong Sun, Junmei Wang, Zhe Wang, Hui Liu, John ZH Zhang, and Tingjun Hou. End-point binding free energy calculation with mm/pbsa and mm/gbsa: strategies and applications in drug design. *Chemical reviews*, 119(16):9478–9508, 2019.

564 75. Dominykas Lukauskis, Marley L Samways, Simone Aureli, Benjamin P Cossins, Richard D Taylor, and Francesco Luigi Gervasio. Open binding pose metadynamics: An effective approach for the
565 ranking of protein–ligand binding poses. *Journal of Chemical Information and Modeling*, 2022.

566 76. Callum R Nicoll, Laura Alvigini, Andrea Gottinger, Domiziana Cecchini, Barbara Mannucci, Federica Corana, María Laura Mascotti, and Andrea Mattevi. In vitro construction of the coq metabolism
567 unveils the molecular determinants of coenzyme q biosynthesis. *Nature catalysis*, 7(2):148–160, 2024.

568 77. Sudhir Kumar, Glen Stecher, Michael Suleski, and S Blair Hedges. Timetree: a resource for timelines, timetrees, and divergence times. *Molecular biology and evolution*, 34(7):1812–1819, 2017.

569 78. Kazutaka Kato, John Rozewicki, and Kazunori D Yamada. Mafft online service: multiple sequence alignment, interactive sequence choice and visualization. *Briefings in bioinformatics*, 20(4):
570 1160–1166, 2019.

571 79. Alexandros Stamatakis. Raxml version 8: a tool for phylogenetic analysis and post-analysis of large phylogenies. *Bioinformatics*, 30(9):1312–1313, 2014.

572 80. Frédéric Lemoine, J-B Domeleau, Entelis, Eduan Wilkinson, Damien Correia, M Dávila Felipe, Túlio De Oliveira, and Olivier Gascuel. Renewing felsenstein's phylogenetic bootstrap in the era of
573 big data. *Nature*, 556(7702):452–456, 2018.

574 81. María Laura Mascotti. Resurrecting enzymes by ancestral sequence reconstruction. In *Enzyme Engineering: Methods and Protocols*, pages 111–136. Springer, 2021.

575 82. Ziheng Yang. Paml 4: phylogenetic analysis by maximum likelihood. *Molecular biology and evolution*, 24(8):1586–1591, 2007.

576 83. Geeta N Eick, Jamie T Bridgman, Douglas P Anderson, Michael J Harms, and Joseph W Thornton. Robustness of reconstructed ancestral protein functions to statistical uncertainty. *Molecular
577 biology and evolution*, 34(2):247–261, 2016.

578 84. Xavier Robert and Patrice Gouet. Deciphering key features in protein structures with the new endscript server. *Nucleic acids research*, 42(W1):W320–W324, 2014.

579 85. Federico Forneris, Roberto Orru, Daniele Bonivento, Laurent R Chiarelli, and Andrea Mattevi. Thermofad, a thermofluor®-adapted flavin ad hoc detection system for protein folding and ligand
580 binding. *The FEBS journal*, 276(10):2833–2840, 2009.

581 86. Wolfgang Kabsch. xds. *Acta Crystallographica Section D: Biological Crystallography*, 66(2):125–132, 2010.

582 87. Jon Aguirre, Mihaela Atanasova, Haroldas Bagdonas, Charles B Ballard, Arnaud Baslé, James Beilstein-Edmands, Rafael J Borges, David G Brown, J Javier Burgos-Mármol, John M Berrisford,
583 et al. The ccp4 suite: integrative software for macromolecular crystallography. *Biological Crystallography*, 79(6):449–461, 2023.

584 88. Ana Casaral, Bernhard Lohkamp, and Paul Emsley. Current developments in coot for macromolecular model building of electron cryo-microscopy and crystallographic data. *Protein Science*, 29
585 (4):1055–1064, 2020.

586 89. Eric F Pettersen, Thomas D Goddard, Conrad C Huang, Elaine C Meng, Gregory S Couch, Tristan I Croll, John H Morris, and Thomas E Ferrin. Ucsf chimera: Structure visualization for
587 researchers, educators, and developers. *Protein science*, 30(1):70–82, 2021.

588 90. Amr Alhosary, Stephanus Daniel Handoko, Yuguang Mu, and Chee-Keong Kwoh. Fast, accurate, and reliable molecular docking with quickvina 2. *Bioinformatics*, 31(13):2214–2216, 2015.

589 91. Thomas Sander, Joel Freyss, Modest Von Korff, and Christian Rufener. Datawarrior: an open-source program for chemistry aware data visualization and analysis. *Journal of chemical information
590 and modeling*, 55(2):460–473, 2015.

591 92. BioSolveIT GmbH. infinisee version 6.2.0, 2024.

592 93. BioSolveIT GmbH. Seesar, 2024.

593 94. D.A. Case, H.M. Aktulga, K. Belfon, I.Y. Ben-Shalom, J.T. Berryman, S.R. Brozell, D.S. Cerutti, T.E. Cheatham, III, G.A. Cisneros, V.W.D. Cruzeiro, T.A. Darden, N. Forouzesh, M. Ghazimirsaeed,
594 G. Giambasu, T. Giese, M.K. Gilson, H. Gohlik, A.W. Goetz, J. Harris, Z. Huang, S. Izadi, S.A. Izmailov, K. Kasavajhala, M.C. Kaymak, I. Kolossváry, A. Kovalenko, T. Kurtzman, T.S. Lee, P. Li,
595 Z. Li, C. Lin, J. Liu, T. Luchko, R. Luo, M. Machado, M. Manathunga, K.M. Merz, Y. Miao, O. Mikhailovskii, G. Monard, H. Nguyen, K.A. O'Hearn, A. Onufriev, F. Pan, S. Pantano, A. Rahnamoun,
596 D.R. Roe, A. Roitberg, C. Sagui, S. Schott-Verdugo, A. Shajan, J. Shen, C.L. Simmerling, N.R. Skrynnikov, J. Smith, J. Swails, R.C. Walker, J. Wang, J. Wang, X. Wu, Y. Wu, Y. Xiong, Y. Xue, D.M.
597 York, C. Zhao, Q. Zhu, and P.A. Kollman. Amber 2024, 2024.

598 95. D.A. Case, H.M. Aktulga, K. Belfon, D.S. Cerutti, G.A. Cisneros, V.W.D. Cruzeiro, N. Forouzesh, T.J. Giese, A.W. Götz, H. Gohlik, S. Izadi, K. Kasavajhala, M.C. Kaymak, E. King, T. Kurtzman,
599 T.-S. Lee, P. Li, J. Liu, T. Luchko, R. Luo, M. Manathunga, M.R. Machado, H.M. Nguyen, K.A. O'Hearn, A.V. Onufriev, F. Pan, S. Pantano, R. Qi, A. Rahnamoun, A. Risheh, S. Schott-Verdugo,
600 A. Shajan, J. Swails, J. Wang, H. Wei, X. Wu, Y. Wu, S. Zhang, S. Zhao, Q. Zhu, T.E. Cheatham III, D.R. Roe, A. Roitberg, C. Simmerling, D.M. York, M.C. Nagan, and K.M. Merz Jr. Ambertools.
601 *J. Chem. Inf. Model.*, 63:6183–6191, 2023. doi: 10.1021/acs.jcim.3c01153.

602 96. Ryne C. Johnston, Kun Yao, Zachary Kaplan, Monica Chelliha, Karl Lesswing, Sean Seekins, Shawn Watts, David Calkins, Jackson Chief Elk, Steven V. Jerome, Matthew P. Repasky, and John C.
603 Shelley. Epik: pka and protonation state prediction through machine learning. *Journal of Chemical Theory and Computation*, 19(8):2380–2388, 2023. doi: 10.1021/acs.jctc.3c00044. PMID:
604 37023332.

605 97. Omri D. Abarbanel and Geoffrey R. Hutchison. Qupkake: Integrating machine learning and quantum chemistry for micro-pka predictions. *Journal of Chemical Theory and Computation*, 20(15):
606 6946–6956, 2024. doi: 10.1021/acs.jctc.4c00328. PMID: 38832803.

607 98. Garrett M Morris, Ruth Huey, William Lindstrom, Michel F Sanner, Richard K Belew, David S Goodsell, and Arthur J Olson. Autodock4 and autodocktools4: Automated docking with selective
608 receptor flexibility. *Journal of computational chemistry*, 30(16):2785–2791, 2009.

609 99. Diogo Santos-Martins, Stefano Forli, Maria João Ramos, and Arthur J Olson. Autodock4zn: an improved autodock force field for small-molecule docking to zinc metalloproteins. *Journal of chemical
610 information and modeling*, 54(8):2371–2379, 2014.

611 100. Jerome Eberhardt, Diogo Santos-Martins, Andreas F Tillack, and Stefano Forli. Autodock vina 1.2. 0: New docking methods, expanded force field, and python bindings. *Journal of Chemical
612 Information and Modeling*, 61(8):3891–3898, 2021.

613 101. Oleg Trott and Arthur J Olson. Autodock vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of computational
614 chemistry*, 31(2):455–461, 2010.

615 102. Pradeep Anand Ravindranath, Stefano Forli, David S Goodsell, Arthur J Olson, and Michel F Sanner. Autodockfr: advances in protein-ligand docking with explicitly specified binding site flexibility.
616 *PLoS computational biology*, 11(12):e1004586, 2015.

617 103. Junkai Ji, Jin Zhou, Zhanfang Yang, Qizheng Lin, Jianqiang Li, and Carlos A Coello Coello. Autodock koto: A gradient boosting differential evolution for molecular docking. *IEEE Transactions on
618 Evolutionary Computation*, 2022.

619 104. Liang Rao, Ning-Xin Jia, Jun Hu, Dong-Jun Yu, and Gui-Jun Zhang. Atpdock: a template-based method for atp-specific protein–ligand docking. *Bioinformatics*, 38(2):556–558, 2022.

620 105. Kai Zhu, Kenneth W Borrelli, Jeremy R Greenwood, Tyler Day, Robert Abel, Ramy S Farid, and Edward Harder. Docking covalent inhibitors: a parameter free approach to pose prediction and
621 scoring. *Journal of chemical information and modeling*, 54(7):1932–1940, 2014.

622 106. William J Allen, Trent E Balias, Sudipto Mukherjee, Scott R Brozell, Demetri T Moustakas, P Therese Lang, David A Case, Irwin D Kuntz, and Robert C Rizzo. Dock 6: Impact of new features and
623 current docking performance. *Journal of computational chemistry*, 36(15):1132–1156, 2015.

624 107. Bernd Kramer, Matthias Rarey, and Thomas Lengauer. Evaluation of the flexx incremental construction algorithm for protein–ligand docking. *Proteins: Structure, Function, and Bioinformatics*, 37
625 (2):228–241, 1999.

626 108. Xiaocong Yang, Yang Liu, Jianhong Gan, Zhi-Xiong Xiao, and Yang Cao. Fitdock: protein–ligand docking by template fitting. *Briefings in Bioinformatics*, 23(3), 2022.

627 109. Woong-Hee Shin and Chaek Seok. Galaxydock: protein–ligand docking with flexible protein side-chains. *Journal of chemical information and modeling*, 52(12):3225–3232, 2012.

628 110. Thomas A Halgren, Robert B Murphy, Richard A Friesner, Hege S Beard, Leah L Frye, W Thomas Pollard, and Jay L Banks. Glide: a new approach for rapid, accurate docking and scoring. 2.
629 enrichment factors in database screening. *Journal of medicinal chemistry*, 47(7):1750–1759, 2004.

630 111. Marcel L Verdonk, Jason C Cole, Michael J Hartshorn, Christopher W Murray, and Richard D Taylor. Improved protein–ligand docking using gold. *Proteins: Structure, Function, and Bioinformatics*,
631 52(4):609–623, 2003.

632 112. Kin Meng Wong, Hio Kuan Tai, and Shirley WI Siu. Gwovina: A grey wolf optimization approach to rigid and flexible receptor docking. *Chemical biology & drug design*, 97(1):97–110, 2021.

633 113. Yumeneng Yan, Huanyu Tao, Jiahua He, and Sheng-You Huang. The hdock server for integrated protein–protein docking. *Nature protocols*, 15(5):1829–1852, 2020.

634 114. Yumeneng Yan, Huanyu Tao, and Sheng-You Huang. Hsymdock: a docking web server for predicting the structure of protein homo-oligomers with cn or dn symmetry. *Nucleic acids research*, 46
635 (W1):W423–W431, 2018.

636 115. Hongjian Li, Kwong-Sak Leung, and Man-Hon Wong. idock: A multithreaded virtual screening tool for flexible ligand docking. In *2012 IEEE Symposium on Computational Intelligence in
637 Bioinformatics and Computational Biology (CIBCB)*, pages 77–84. IEEE, 2012.

638 116. Kai-Cheng Hsu, Yen-Fu Chen, Shen-Rong Lin, and Jinn-Moon Yang. igemdock: a graphical environment of enhancing gemdock using pharmacological interactions and post-screening analysis.
639 *BMC bioinformatics*, 12(1):1–11, 2011.

640 117. Ni Liu and Zhibin Xu. Using ledock as a docking tool for computational drug design. In *IOP Conference Series: Earth and Environmental Science*, volume 218, page 012143. IOP Publishing, 2019.

641 118. Cherayathumadom M Venkatachalam, Xiaohui Jiang, Tom Oldfield, and Marvin Waldman. Ligandfit: a novel method for the shape-directed rapid docking of ligands to protein active sites. *Journal
642 of Molecular Graphics and Modelling*, 21(4):289–307, 2003.

643 119. Brian Jiménez-García, Jorge Roel-Touris, Miguel Romero-Durana, Miquel Vidal, Daniel Jiménez-González, and Juan Fernández-Recio. Lightdock: a new multi-scale approach to protein–protein
644 docking. *Bioinformatics*, 34(1):49–55, 2018.

645 120. Zhiwei Ma and Xiaoqin Zou. Mdock: A suite for molecular inverse docking and target prediction. In *Protein-Ligand Interactions and Drug Design*, pages 313–322. Springer, 2021.

646 121. Ming Liu and Shaoming Wang. Mcdock: a monte carlo simulation approach to the molecular docking problem. *Journal of computer-aided molecular design*, 13(5):435–451, 1999.

647 122. Gabriela Bitencourt-Ferreira and Walter Filgueira de Azevedo. Molegro virtual docker for docking. In *Docking screens for drug discovery*, pages 149–167. Springer, 2019.

648 123. Fang Bai, Sha Liao, Junfeng Gu, Hualiang Jiang, Xicheng Wang, and Honglin Li. An accurate metalloprotein-specific scoring function and molecular docking program devised by a dynamic
649 sampling and iteration optimization strategy. *Journal of Chemical Information and Modeling*, 55(4):833–847, 2015.

650 124. Marcus CK Ng, Simon Fong, and Shirley WI Siu. Psovina: The hybrid particle swarm optimization algorithm for protein–ligand docking. *Journal of bioinformatics and computational biology*, 13 (03):1541007, 2015.

651 125. Oliver Korb, Thomas Stützle, and Thomas E Exner. Plants: Application of ant colony optimization to structure-based drug design. In *International workshop on ant colony optimization and swarm intelligence*, pages 247–258. Springer, 2006.

652 126. Bin Li and Daisuke Kihara. Protein docking prediction using predicted protein-protein interface. *BMC bioinformatics*, 13(1):1–17, 2012.

653 127. Dima Kozakov, Ryan Brenke, Stephen R Comeau, and Sandor Vajda. Piper: an fft-based protein docking program with pairwise potentials. *Proteins: Structure, Function, and Bioinformatics*, 65 (2):392–406, 2006.

654 128. Nafisa M Hassan, Amr A Alhossary, Yuguang Mu, and Chee-Keong Kwoh. Protein-ligand blind docking using quickvina-w with inter-process spatio-temporal integration. *Scientific reports*, 7(1): 1–13, 2017.

655 129. Sergio Ruiz-Carmona, Daniel Alvarez-Garcia, Nicolas Foloppe, A Beatriz Garmendia-Doval, Szilveszter Juhasz, Peter Schmidke, Xavier Barril, Roderick E Hubbard, and S David Morley. rdock: a fast, versatile and open source program for docking ligands to proteins and nucleic acids. *PLoS computational biology*, 10(4):e1003571, 2014.

656 130. Gordon Lemmon and Jens Meiler. Rosetta ligand docking with flexible xml protocols. In *Computational Drug Discovery and Design*, pages 143–155. Springer, 2012.

657 131. Li-Zhen Sun, Yangwei Jiang, Yuanzhe Zhou, and Shi-Jie Chen. Ridock: a new method for predicting rna–ligand interactions. *Journal of chemical theory and computation*, 16(11):7173–7183, 2020.

658 132. Nicolas Majeux, Marco Scarsi, Joannis Apostolakis, Claus Ehrhardt, and Amedeo Caflisch. Exhaustive docking of molecular fragments with electrostatic solvation. *Proteins: Structure, Function, and Bioinformatics*, 37(1):88–105, 1999.

659 133. David Ryan Koes, Matthew P Baumgartner, and Carlos J Camacho. Lessons learned in empirical scoring with smina from the csar 2011 benchmarking exercise. *Journal of chemical information and modeling*, 53(8):1893–1904, 2013.

660 134. Anita K Nivedha, David F Thieker, Spandana Makeneni, Huimin Hu, and Robert J Woods. Vina-carb: improving glycosidic angles during carbohydrate docking. *Journal of chemical theory and computation*, 12(2):892–901, 2016.

661 135. Mathew R Koebel, Grant Schmadeke, Richard G Posner, and Suman Srimulla. Autodock vinaxb: implementation of xsbf, new empirical halogen bond scoring function, into autodock vina. *Journal of cheminformatics*, 8(1):1–8, 2016.

662 136. Ruth Huey, Garrett M Morris, Arthur J Olson, and David S Goodsell. A semiempirical free energy force field with charge-based desolvation. *Journal of computational chemistry*, 28(6):1145–1152, 2007.

663 137. Jayashree Srinivasan, Thomas E Cheatham, Piotr Cieplak, Peter A Kollman, and David A Case. Continuum solvent studies of the stability of dna, rna, and phosphoramidate- dna helices. *Journal of the American Chemical Society*, 120(37):9401–9409, 1998.

664 138. Maciej Wójcikowski, Pedro J Ballester, and Paweł Siedlecki. Performance of machine-learning scoring functions in structure-based virtual screening. *Scientific Reports*, 7(1):1–10, 2017.

665 139. Hao-Yang Liu, Irwin D Kuntz, and Xiaqin Zou. Pairwise gb/sa scoring function for structure-based drug design. *The Journal of Physical Chemistry B*, 108(17):5453–5462, 2004.

666 140. Rodrigo Quiroga and Marcos A Villarreal. Vinardo: A scoring function based on autodock vina improves scoring, docking, and virtual screening. *PLoS one*, 11(5):e0155183, 2016.

667 141. Susanta Das, Mor Shimsi, Keren Raz, Neta Nitoker Eliaz, Anil Ranu Mhashal, Tamar Ansbacker, and Dan T Major. Enzydock: Protein–ligand docking of multiple reactive states along a reaction coordinate in enzymes. *Journal of Chemical Theory and Computation*, 15(9):5116–5134, 2019.

668 142. Mark McGann. Fred pose prediction and virtual screening accuracy. *Journal of chemical information and modeling*, 51(3):578–596, 2011.

669 143. Noel M. O’Boyle, Michael Banck, Craig A. James, et al. Open Babel: An open chemical toolbox. *Journal of Cheminformatics*, 3(1):1–14, 2011. ISSN 1758-2946. doi: 10.1186/1758-2946-3-33.

670 144. Hannah Farmer, Nuala McCabe, Christopher J Lord, Andrew N Tutt, Damian A Johnson, Tobias B Richardson, Manuela Santarosa, Krystyna J Dillon, Ian Hickson, Charlotte Knights, et al. Targeting the dna repair defect in brca mutant cells as a therapeutic strategy. *Nature*, 434(7035):917–921, 2005.

671 145. Marie-France Langelier, Jamie L Planck, Swati Roy, and John M Pascal. Structural basis for dna damage–dependent poly (adp-ribosyl) ation by human parp-1. *Science*, 336(6082):728–732, 2012.

672 146. Tom EH Ogden, Ji-Chun Yang, Marianne Schimpl, Laura E Easton, Elizabeth Underwood, Philip B Rawlins, Michael M McCauley, Marie-France Langelier, John M Pascal, Kevin J Embrey, et al. Dynamics of the hd regulatory subdomain of parp-1: substrate access and allostery in parp activation and inhibition. *Nucleic acids research*, 49(4):2266–2288, 2021.

673 147. Thomas Helleday. The underlying mechanism for the parp and brca synthetic lethality: clearing up the misunderstandings. *Molecular oncology*, 5(4):387–393, 2011.

674 148. Junko Murai, Shar-yn Huang, Benu Brata Das, Amelie Renaud, Yiping Zhang, James H Doroshow, Jiuping Ji, Shunichi Takeda, and Yves Pommier. Trapping of parp1 and parp2 by clinical parp inhibitors. *Cancer research*, 72(21):5588–5599, 2012.

675 149. Uday Kiran Velagapudi, Marie-France Langelier, Cristina Delgado-Martin, Morgan E Diolaiti, Sietske Bakker, Alan Ashworth, Bhargav A Patel, Xuwei Shao, John M Pascal, and Tanaji T Talele. Design and synthesis of poly (adp-ribose) polymerase inhibitors: impact of adenosine pocket-binding motif appendage to the 3-oxo-2, 3-dihydrobenzofuran-7-carboxamide on potency and selectivity. *Journal of medicinal chemistry*, 62(11):5330–5337, 2019.

676 150. Fons Elstrott, Antoinette Hollestelle, Jord HA Nagel, Michael Gorin, Marijke Wasielewski, Ans Van Den Ouwendal, Sofia D Merajver, Stephen P Ethier, and Mieke Schutte. Brca1 mutation analysis of 41 human breast cancer cell lines reveals three new deleterious mutants. *Cancer research*, 66(1):41–45, 2006.

677 151. Neil Johnson, Shawn F Johnson, Wei Yao, Yu-Chen Li, Young-Eun Choi, Andrea J Bernhardy, Yifan Wang, Marzia Capelletti, Kristopher A Sarosiek, Lisa A Moreau, et al. Stabilization of mutant brca1 protein confers parp inhibitor and platinum resistance. *Proceedings of the National Academy of Sciences*, 110(42):17041–17046, 2013.

678 152. Harry L Morgan. The generation of a unique machine description for chemical structures-a technique developed at chemical abstracts service. *Journal of chemical documentation*, 5(2):107–113, 1965.

679 153. Schrödinger LLC New York NY. Schrödinger release 2024-1: Maestro. <https://newsite.schrodinger.com/platform/products/maestro/>, 2024.

680 154. Garrett M. Morris, Ruth Huey, William Lindstrom, Michel F. Sanner, et al. AutoDock4 and AutoDockTools4: Automated docking with selective receptor flexibility. *Journal of Computational Chemistry*, 30(16):2785–2791, dec 2009. ISSN 01928651. doi: 10.1002/jcc.21256.

681 155. Paul D Leeson and Brian Springthorpe. The influence of drug-like concepts on decision-making in medicinal chemistry. *Nature reviews Drug discovery*, 6(11):881–890, 2007.

682 156. John S Delaney. Esol: estimating aqueous solubility directly from molecular structure. *Journal of chemical information and computer sciences*, 44(3):1000–1005, 2004.

683 157. Claudia Lamanna, Marta Bellini, Alessandro Padova, Goran Westerberg, and Laura Maccari. Straightforward recursive partitioning model for discarding insoluble compounds in the drug discovery process. *Journal of medicinal chemistry*, 51(10):2891–2897, 2008.

684 158. Wenyu Wei, Srinivasulu Cherukupalli, Lanlan Jing, Xinyong Liu, and Peng Zhan. Fsp3: A new parameter for drug-likeness. *Drug Discovery Today*, 25(10):1839–1845, 2020.

685 159. Aixia Yan and Johann Gasteiger. Prediction of aqueous solubility of organic compounds based on a 3d structure representation. *Journal of chemical information and computer sciences*, 43(2): 429–434, 2003.

686 160. David C Kombo, Kartik Tallapragada, Rachit Jain, Joseph Chewning, Anatoly A Mazurov, Jason D Speake, Terry A Hauser, and Steve Toler. 3d molecular descriptors important for clinical success. *Journal of chemical information and modeling*, 53(2):327–342, 2013.

687 161. Soon-Keat Ooi, Shigeo Sato, Chieri Tomomori-Sato, Ying Zhang, Zhihui Wen, Charles AS Banks, Michael P Washburn, Jay R Unruh, Laurence Florens, Ronald C Conaway, et al. Multiple roles for parp1 in alc1-dependent nucleosome remodeling. *Proceedings of the National Academy of Sciences*, 118(36):e2107277118, 2021.

688 162. Marcin J Suskiewicz, Florian Zobel, Tom EH Ogden, Pietro Fontana, Antonio Ariza, Ji-Chun Yang, Kang Zhu, Lily Bracken, William J Hawthorne, Dragana Ahel, et al. Hpf1 completes the parp active site for dna damage-induced adp-ribosylation. *Nature*, 579(7800):598–602, 2020.

689 163. Frank Delaglio, Stephan Grzesiek, Geerten W Vuister, Guang Zhu, John Pfeifer, and AD Bax. Nmrpipe: a multidimensional spectral processing system based on unix pipes. *Journal of biomolecular NMR*, 6:277–293, 1995.

690 164. Wim F Vranken, Wayne Boucher, Tim J Stevens, Rasmus H Fogh, Anne Pajon, Miguel Llinas, Eldon L Ulrich, John L Markley, John Ionides, and Ernest D Lue. The ccpn data model for nmr spectroscopy: development of a software pipeline. *Proteins: structure, function, and bioinformatics*, 59(4):687–696, 2005.

691 165. Airlie J McCoy, Ralf W Grosse-Kunstleve, Paul D Adams, Martyn D Winn, Laurent C Storoni, and Randy J Read. Phaser crystallographic software. *Journal of applied crystallography*, 40(4): 658–674, 2007.

692 166. Paul D Adams, Pavel V Afonine, Gábor Bunkóczki, Vincent B Chen, Ian W Davis, Nathaniel Echols, Jeffrey J Headd, L-W Hung, Gary J Kapral, Ralf W Grosse-Kunstleve, et al. Phenix: a comprehensive python-based system for macromolecular structure solution. *Acta Crystallographica Section D: Biological Crystallography*, 66(2):213–221, 2010.

693 167. Paul Emsley and Kevin Cowtan. Coot: model-building tools for molecular graphics. *Acta crystallographica section D: biological crystallography*, 60(12):2126–2132, 2004.

724 Methods

725 AdaptiveFlow

726 AdaptiveFlow is an in silico platform consisting of three modules (AFLP, AFVS, and AFU) and ready-made ligand libraries
727 for virtual screens. In this section, general platform-wide features have been described that span multiple modules.

728 **Python Code-base.** AdaptiveFlow has been entirely rewritten in Python, a programming language that enables the imple-
729 mentation of more sophisticated and advanced code while reducing maintenance work. This rewrite has also made it easier for
730 scientists and programmers to join the open-source project and/or extend AdaptiveFlow independently, as Python is the most
731 widely used programming language in scientific communities. In contrast, the previous version of AdaptiveFlow was written
732 in Bash.

733 **Parallelism.** In addition to being rewritten in Python, a key goal for AdaptiveFlow was to improve parallelism in order to
734 decrease the overall time required for screening. AdaptiveFlow has increased the levels of parallelism compared to the previous
735 version. To help manage this increased parallelism, it maintains the use of *collections*, that group around 1,000 ligands together,
736 reducing the overhead of scheduling individual ligands for processing.

737 AdaptiveFlow also introduces the concept of *subjobs*, which are atomic units of computational work that include one or
738 more collections for processing. These subjobs are automatically generated by AdaptiveFlow to match the runtime requirements
739 specified in the job configuration. Subjobs are then combined into larger groups called *workunits*, which are designed to allow
740 schedulers to efficiently process the high levels of parallelism generated by AdaptiveFlow. These work units are sent as job
741 arrays to the end schedulers for processing. By optimizing for parallelism, the team was able to achieve 5.6 million concurrent
742 vCPUs running AdaptiveFlow.

743 **Native Cloud Support.** The AFLP and AFVS modules of AdaptiveFlow can run on computer clusters as well as on the cloud
744 by using a batch system. Currently, two batch systems are supported, Slurm (37) and AWS Batch (38). Slurm is able to run on
745 many cloud infrastructures, including Google Cloud, Microsoft Azure, as well as AWS. Within AWS, AdaptiveFlow uses AWS
746 Batch and S3 object storage. CloudFormation is used to set up AdaptiveFlow and all necessary components. The AWS-native
747 architecture allows AdaptiveFlow to utilize millions of virtual CPUs (vCPUs) in parallel, while the previous version was only
748 demonstrated to scale up to 160,000 vCPUs.

749 AdaptiveFlow is able to utilize various features from AWS in addition to supporting traditional clusters, making it easier
750 to use and increasing the scale of computations available to researchers. Some of these features include:

- 751 • **Scheduling and Orchestration:** AWS Batch is a cloud-native scheduler and orchestrator for containerized jobs. As part
752 of running a job, AdaptiveFlow can make use of AWS Batch to automatically scale cloud infrastructure up and down as
753 needed and run all of the individual subjobs. AdaptiveFlow can also automatically generate the Docker container used
754 for processing, eliminating the need to understand containerization or other infrastructure.
- 755 • **Object Storage (Amazon S3):** In addition to the usage of standard cluster filesystems, AdaptiveFlow also supports the
756 usage of object storage for both ingesting of input datasets and the output of the resulting data from the processing.
- 757 • **Spot Capacity Awareness and Spot Resilience:** Spot capacity is unused computing capacity that cloud providers make
758 available at significantly discounted prices compared to normal on-demand access. In exchange for a lower price, a
759 workload running on a spot instance may be reclaimed (stopped) with a short warning. AdaptiveFlow is designed to
760 take advantage of this capacity model by using shorter running subjobs (approximately 30 minutes) and automatically
761 re-running jobs that have prematurely exited through integration with AWS Batch. Additionally, through integration with
762 AWS Batch, AdaptiveFlow is able to select instances that are less likely to be interrupted through the "spot capacity
763 aware" scheduling method.

764 The architecture of AdaptiveFlow when deployed on AWS is shown in Supplementary Fig. 1. The architecture makes
765 use of multiple *Compute Environments*, each of which may contain thousands of instances (compute nodes). The instances
766 are automatically allocated and deallocated when the job starts and ends, respectively. Individual subjobs are allocated to the
767 instances by AWS Batch until all subjobs have been completed. These subjobs run as Docker containers on the instances.

768 **ARM.** AdaptiveFlow is also able to run on the compute ARM architecture when using ARM-compatible choices in the settings.
769 AFLP is fully ARM compatible, and AFVS is ARM compatible when docking programs are used that either make the source
770 code available or provide binaries for ARM. ARM-based computing infrastructure can be more cost-efficient than traditional
771 x86-based architectures.

772 Enamine REAL Space

773 The REAL Space (enumerated, version 2022q1-2) was prepared using AFVS. JChem Suite was used for neutralizing the
774 molecules, stereoisomer enumeration, tautomer generation, protonation state prediction, 3D coordinate calculation, and
775 the calculation of multiple molecular properties, JChem Suite 18.20.0, ChemAxon (<https://www.chemaxon.com>).
776 Open Babel (version 2.3.2) was used as backup program for protonation state prediction, 3D coordinate calculation (61),
777 and target format conversion (all formats except SELFIES). The pip-installable SELFIES package (<https://github.com/aspuru-guzik-group/selfies>) was used for the conversion of molecules between SMILES and SELFIES. An
778 overview of the processing steps and preparation methods used is shown in Supplementary Table 2.
779

780 **Multiple Formats.** AdaptiveFlow provides the REAL Space library in ready-to-dock 3D formats including PDB, PDBQT,
781 MOL2, and SDF. These formats are compatible with a wide range of docking programs for small molecules. The total size
782 of the ready-to-dock library is approximately 50 terabytes in compressed form for each of these formats, totaling around 200
783 terabytes in compressed form (uncompressed, the library is around 2 petabytes). In addition to these ready-to-dock formats,
784 the REAL Space library is also provided in the SELFIES (33), SMILES (62), and Parquet formats. These formats are useful
785 for postprocessing the results of virtual screening (e.g., for easy representation of each molecule), and they can also be used
786 as standalone libraries for other purposes such as in the AIML community, where SMILES and SELFIES are the standard
787 molecular input file formats. The Parquet format of the REAL Space only stores the properties of the molecules, which can be
788 useful for fast library-wide queries.

789 AdaptiveFlow for Ligand Preparation (AFLP)

790 AFLP is the module of AdaptiveFlow dedicated to the preparation of ligand libraries into a ready-to-dock format and can handle
791 libraries of any size, including ultra-large libraries. AFLP requires a batch system and supports Slurm as well as AWS Batch
792 as batch systems. AFLP is configured via a configuration file (an example configuration file is shown in Supplementary Listing
793 I). Details on the workflow are available in Supplementary Information Section D. The new version of AFLP has the following
794 features.

795 **Preparation into Ready-To-Dock Format.** The primary task of AFLP is to prepare ligands into a ready-to-dock format.
796 To accomplish this, AFLP desalts the original ligands (in case they are salts), neutralizes them, generates stereoisomers and
797 tautomers, predicts the protonation states, computes the minimum-potential energy conformation of the ligand (3D coordinates),
798 optionally validates the generated 3D conformers, and converts them into the desired target formats. In total over 100 output
799 formats are supported. A detailed overview of the processing steps can be seen in Supplementary Fig. 7. The support input
800 formats are SMILES, SELFIES, as well as amino acid sequences. AFLP uses external tools such as ChemAxon's JChem, Open
801 Babel, and RDKit (63), to carry out the preparation steps. For many of the preparation steps, the user can choose which external
802 tool should be used. Details on which tools are available for each processing step can be found in Supplementary Table 3.

803 **Molecular Properties Calculation.** AFLP enables researchers to calculate a variety of physicochemical properties and molec-
804 ular descriptors for the molecules in the input ligand library, in addition to its primary task of preparing them for molecular
805 docking. By calculating properties such as QED (Quantitative Estimate of Druglikeness), logP, molecular weight, and others (a
806 full list of supported properties can be found in Supplementary Table 2), researchers can more effectively select and prioritize
807 ligands, which can help reduce costs and the time required for screens using AFVS. AFLP uses mostly external tools to calculate
808 the molecular properties, such as ChemAxon's JChem (<https://chemaxon.com/calculators-and-predictors>),
809 Open Babel (61), and RDKit (for details see Supplementary Table 3). By combining the results of AFLP with those of AFVS,
810 researchers can gain valuable insights into the structural and chemical characteristics of their ligands relative to specific target
811 proteins, which can help them design and optimize new molecules with improved affinity and specificity.

812 **Automatic Tranche Assignments.** In AdaptiveFlow, ligand libraries are organized into tranches. Each tranche corresponds
813 to a property of the ligands and is divided into multiple intervals or categories. For example, the molecular weight of the ligands
814 can be used as a tranche type, and can then be partitioned into several categories or intervals such as 0-200, 200-300, 300-400,
815 400-500, and 500-∞ daltons. Multiple tranches allow for the partitioning of a ligand library into a multidimensional grid or
816 table, with one dimension corresponding to each tranche type (ligand property).

817 While the previous version of AdaptiveFlow required that the input ligand collection was already organized and struc-
818 tured in the tranche format, AdaptiveFlow can automatically reorganize the input ligand library into a new tranche output format
819 based on the computation of specific molecular properties of the ligands (see the corresponding paragraph above). This feature
820 is useful because most ligand libraries (e.g., from chemical compound vendors) are initially in the SMILES format and not in
821 the tranche format. AFLP can then be used to prepare the entire library in the multidimensional grid/tranche format for AFVS.

822 **Support of SELFIES.** SELF-referencing Embedded Strings (SELFIES) (33, 34) has become one of the most popular line
823 notation formats for small molecules within machine-learning communities. The primary reason is that SELFIES are fully
824 robust, meaning that every SELFIES string corresponds to a valid molecule, which is not the case for SMILES. Thus generative

825 ML models cannot create invalid molecules when using SELFIES. AFLP now supports SELFIES as an input format for the
826 molecules in addition to SMILES. Additionally, the entire REAL Space is free to be downloaded in the SELFIES format,
827 enabling the creation of powerful datasets.

828 **Stereoisomer Enumeration.** One of the main tasks of AFLP is to prepare ligand libraries from unprepared formats (e.g.,
829 SMILES) into ready-to-dock formats. The new version of AFLP enables the enumeration of stereoisomers for the ligands,
830 which was not possible in the previous version. Accurate stereochemistry is essential for virtual screenings using structure-
831 based molecular docking, but many ligand libraries in unprepared formats (e.g., from chemical compound vendors) do not
832 contain stereochemical information. The ability to enumerate stereoisomers allows AFLP to include this information in the
833 prepared ligand library, ensuring that the virtual screening results are accurate.

834 **Extended Open Source Support.** The previous version of AFLP required external tools (from ChemAxon) that are not open
835 source for some of the steps when preparing ligands in a ready-to-dock format. The new version provides the option to choose
836 open-source alternatives for each of the steps to prepare molecules into a ready-to-dock format. In particular, Open Babel and
837 RDKit are supported by AdaptiveFlow as open-source alternatives.

838 **Enhanced Speed.** AFLP has been improved in terms of speed compared to the original version. Firstly, this has been achieved
839 by AFLP using its own Java code to call the necessary functions of ChemAxon's JChem package via Nailgun, which reduces
840 the number of calls to Nailgun significantly. Additionally, timeout parameters that can be adjusted by users have been added
841 to the control file, allowing AFLP to skip ligands that require unusually long processing times during preparation. These
842 improvements make AFLP more efficient, robust and able to handle larger and more chemically complex ligand libraries.

843 **Enhanced Quality Checks.** AFLP checks the validity of each ligand that it prepares in the ready-to-dock format by calculating
844 the potential energy of the ligand in 3D format using Open Babel. This allows AFLP to automatically remove ligands that have
845 an excessively high potential energy, which can be an indication that the compound was damaged or corrupted during the
846 preparation process. Additionally, AFLP supports optional plausibility checks on generated 3D conformers using PoseBusters
847 (64). This helps to ensure the integrity and accuracy of the prepared ligand library.

848 **Application.** AFLP was used to prepare the REAL Space (version 2022q1-2). It was run on AWS using over 5.6 million Intel
849 vCPUs in parallel. Spot instances were used, and the used capacity was sustained without significant preemption. In total less
850 than 0.1 percent of the used CPU hours were interrupted due to preemption. The scaling behavior of AFLP was found to be
851 perfectly linear (as shown in Supplementary Fig. 2). This demonstrates the efficiency and effectiveness of AFLP in preparing
852 large ligand libraries in a ready-to-dock format.

853 **AdaptiveFlow for Virtual Screening (AFVS)**

854 AFVS is the module of AdaptiveFlow dedicated to ultra-large virtual screens but can also screen libraries of smaller sizes. It
855 requires a batch system and supports Slurm as well as AWS Batch. AFVS is configured via a configuration file (an example
856 configuration file is shown in Supplementary Listing J). AFVS and AFU share many virtual screening features, which are
857 described in a later section. AFVS and AFLP share the same parallelization mechanisms when using Slurm or AWS Batch, and
858 therefore the scaling behavior of AFVS is expected to be the same as for AFLP, which was shown to be perfectly linear up to
859 5.6 million CPUs (see also Supplementary Fig. 2).

860 **Parquet Output Format.** AFVS has the ability to store output data in the Apache Parquet format, an open-source column-
861 oriented data storage format. When using this format, AFVS can take advantage of any database and query system that can
862 read this format, including Amazon Athena, a serverless service that allows users to efficiently query output score files stored in
863 Amazon S3 (object storage) on AWS. Additionally, AFVS has a feature for automatically postprocessing ATG Prescreens using
864 Amazon Athena, enabling researchers to efficiently prepare the input data for the ATG primary screens. Tools are provided
865 with AFVS to facilitate common queries, but researchers can also create their own queries using the open-source format of the
866 dataset.

867 **ATG-VSs and Benchmarks.** Following completion of the ATG Prescreen, tranches can be selected for the ATG Primary
868 Screen using two main strategies. In the first, individual tranches from the 18-dimensional REAL Space matrix can be man-
869 ually selected, offering fine-grained control over specific chemical subspaces. In the second, selection can be restricted to a
870 hyperrectangular region of the grid. This is achieved by identifying the top-performing intervals (i.e., bins) within each of the
871 18 molecular property dimensions and computing their Cartesian product to form a structured, intersecting subset of tranches.
872 This hyperrectangle approach enables chemically diverse yet targeted selection based on multiple favorable properties. In
873 both selection modes, users define the total number of molecules to include in the ATG Primary Screen, and AdaptiveFlow
874 automatically prioritizes the tranches containing the best-performing representative ligands.

875 To evaluate the effectiveness of ATG-VS, we performed two types of benchmark studies: (1) full-scale production
876 benchmarks using the entire Enamine REAL Space, and (2) smaller-scale "test set-based" benchmarks designed for rapid

877 prototyping and detailed comparison across different ATG configurations. In the test set-based benchmarks, we randomly
878 selected synthetic reactions from the Enamine space to create a compound subset totaling approximately 5 million molecules.
879 For each of the 10 protein targets tested (listed in Supplementary Table 7), a different random subset was used to ensure unbiased
880 evaluation across targets and screening scenarios.

881 The test set-based benchmarks allowed us to systematically assess the benefits of ATG-VS in three configurations:
882 standard ULVS (random selection), ATG without active learning, and ATG augmented with the optional ML-based active
883 learning component. In contrast, the full-scale production benchmarks were conducted using the entire 69 billion-compound
884 library, but without the active learning step, to minimize computational overhead while validating the performance of ATG-
885 VS at scale. All docking in these benchmarks was performed using QuickVina 2 with default parameters, unless otherwise
886 specified. For the “100K ATG” benchmark configuration presented in Fig. 8, tranches for the primary screen were selected
887 using the hyperrectangle method, enabling a balanced, property-informed sampling of chemical space.

888 **Machine Learning Classifier for Tranche Prioritization.** To improve the efficiency of Adaptive Target-Guided Virtual
889 Screens (ATG-VSs), the AdaptiveFlow platform incorporates an optional machine learning (ML) classification step into the
890 primary screening stage. This step prioritizes compounds that are most likely to be high-affinity binders, thereby reducing
891 the number of molecules subjected to full docking. The ML classifier is applied after the ATG prescreen and is used to filter
892 undocked molecules within the selected tranches prior to docking.

893 The classifier is implemented as a fully connected feedforward neural network (FNN) designed to distinguish between
894 promising and less likely binders. Compounds are represented using Morgan fingerprints of length 1,024 and radius 2, gener-
895 ated via RDKit (v2023.03.1). The model is trained on docking scores obtained from all representative molecules used in the
896 prescreen. These scores are converted into binary classification labels using a percentile-based thresholding strategy: the top
897 25% of compounds with the most negative docking scores are labeled as positive examples (“high-confidence binders”), while
898 the remaining compounds are labeled as negatives. To mitigate class imbalance, the majority class is undersampled during
899 training.

900 The neural network architecture comprises four fully connected layers: an input layer of 1,024 nodes, followed by
901 hidden layers of 512, 256, and 128 nodes respectively, each using ReLU activation. The output layer consists of a single
902 sigmoid-activated node that outputs the predicted binding probability. Hyperparameters and layer configurations were selected
903 through limited tuning using cross-validation on the prescreen data. The model is trained with the Adam optimizer (learning
904 rate 0.001), binary cross-entropy loss, and early stopping based on validation loss with a patience of 5 epochs. A validation
905 split of 10% is used, with training capped at 50 epochs and a batch size of 256, although convergence typically occurs earlier.

906 Once trained, the model is applied to all undocked compounds within the selected tranches. Compounds with predicted
907 binding probabilities greater than 0.5 are retained for full docking, while the rest are filtered out. This filtering step typically
908 reduces the number of compounds subjected to docking by 30–70%, depending on the target, while maintaining or improving
909 hit enrichment.

910 The ML classification module is fully integrated into the AdaptiveFlow Unity (AFU) environment and can be activated
911 through a configuration parameter. The training and inference pipeline is implemented in PyTorch (v2.0.1) and supports GPU
912 acceleration for scalable and efficient deployment.

913 **AdaptiveFlow Unity**

914 AdaptiveFlow Unity is the all-in-one version of AdaptiveFlow, combining both ligand preparation and virtual screening into a
915 single streamlined workflow. It is the third module besides AFLP and AFVS. This new module adds many conveniences for
916 workflows where both ligand preparation and virtual screenings have to be carried out in tandem. Notably, AFU is designed
917 to be run as a package, supporting approximately 1500 docking methods that can be run independently of job schedulers (see
918 Supplementary Data Table 1). A conceptual overview of the workflow of AFU can be found in Supplementary Fig. 11, and a
919 detailed UML workflow diagram is shown in Supplementary Fig. 5. Specifically, as input, a user specifies the molecule to be
920 docked (as a SMILES, SELFIES, or amino-acid sequence), the protein file, the docking site, and the docking program to be run.
921 Based on the user’s choices, error checks in the form of missing docking-choice-specific files are performed, and calculations
922 are run. As a result of a successful calculation, the docking score and pose are returned. AFU can be run in standalone mode,
923 where it is run via the command line and specifies the options via a configuration file (an example configuration file is shown
924 in Supplementary Data Fig. K). This mode is for example useful for scientists in the drug discovery communities who want to
925 carry out molecular dockings in the most convenient and simple way, without requiring computational expertise, while having a
926 large number of docking methods available. In addition, AFU can be run in API mode to directly interface with other programs
927 and code, which can for instance be useful for the machine learning communities that develop new methods for drug discovery.

928 AdaptiveFlow Unity comes in two versions. The first is dedicated to single workstations or compute nodes without
929 multinode parallelization. VirtualFlow Unity Parallelized (AFUparr) is the parallelized version of VirtualFlow Unity. Designed
930 to operate seamlessly on SLURM systems, AFUparr allows users to easily parallelize the workflow on a larger number of
931 CPUs and GPUs. Execution within SLURM environments is highly optimized, with computations distributed in parallel across
932 multiple CPUs and nodes. This design ensures efficient linear scaling relative to the number of molecules provided.

933 **Joint Features of AFVS and AFU**

934 The AFU and AFVS modules support the same docking methods and therefore share all docking-specific features.

935 **Supported Docking Methods.** AdaptiveFlow supports approximately 40 docking methods, listed in Tables 4, 5, 6. These
936 include software that performs both pose prediction and scoring (Table 4), solely pose prediction (Table 6) and only scoring
937 (Table 5). Notably, within AdaptiveFlow, the scoring functions can be combined with pose prediction methods, giving rise to
938 approximately 1500 methods (listed in Table 1). We note that many of the newly supported docking software possess special
939 features such as protein-protein/RNA docking, explicit parameterization of solvents/metals, and covalent docking. In Tables
940 4,5,6, we describe these as "Special features". Additionally, a key consideration for ULVS is the timing required for carrying
941 out a single calculation. As such, in Fig. 12, we provide timing estimates for running a single calculation, averaging over 25
942 independent runs for all docking programs.

943 **Special Types of Ligands.** In addition to traditional small molecule ligands, AdaptiveFlow supports the use of peptides,
944 protein ligands, and covalently-bonded ligands in AFVS. Peptides are of increasing interest as therapeutics, and several docking
945 programs already support their use. For protein-protein/RNA docking, specialized programs require two protein/RNA targets
946 in order to generate co-complexes. Covalent docking involves reduced flexibility for certain parts of the input molecule during
947 pose sampling, and a few specialized docking programs support this ligand type (see Supplementary Table 4).

948 **Enhanced Performance.** The QuickVina 2 molecular docking tool is used to accelerate AutoDock Vina with no loss of
949 accuracy (65). QuickVina 2 execution was profiled in an effort to improve performance even further. To establish baseline
950 performance, the open-source GitHub repository was cloned and the qvina_1buffer branch was compiled using the GNU C++
951 compiler (version 9.4), the Boost C++ libraries (version 1.79, available from <https://www.boost.org>), and the default
952 parameters (see <https://qvina.github.io/compilingQvina2.html> for the build instructions). A single docking
953 experiment of Avanafil binding to unsolvated 5FNQ was run through the Intel VTune Profiler to look for performance bottle-
954 necks. Unfortunately, the execution profile was flat. The most time-consuming function accounted for less than 20% of the
955 total run-time. Tuning the source code would give rapidly diminishing returns at the risk of affecting accuracy, so compiler-
956 level optimization was the best option to improve performance. QuickVina 2 was rebuilt using the Intel oneAPI DPC++/C++
957 Compiler (version 2022.2.0.20220730) and aggressive optimization: automatic vectorization, interprocedural optimization, and
958 relaxed floating-point constraints. Profile-guided optimization was also tried, but it did not improve performance. The final
959 executable was 18% faster than the default builds with no loss of accuracy. The new binaries are provided on the QuickVina
960 and AF homepages.

961 **GPU Support.** Several of the docking programs support graphics processing units (GPUs), which can strongly reduce computa-
962 tion time and cost. For pose prediction, these include softwares: AutoDock-GPU (57), QVina2-GPU (58), QVina2-W-GPU
963 (58), Vina-GPU (58). Additionally, deep-learning methods such as TANKBind (12), DiffDock (66), EquiBind (67) use GPU-
964 compatible neural networks for direct pose prediction of several ligands in parallel. Additionally, Gnina (68), NNScore 2.0
965 (69), DeepAffinity (70) and DeepBindRG (71) use GPU-compatible deep learning models for scoring protein-ligand com-
966 plexes. AdaptiveFlow supports GPU compatibility of all these software. The relative speedups of these methods for single
967 calculations can be compared within Fig. 12.

968 **Molecular Dynamics-Based Methods.** Molecular docking is a computationally fast way to predict binding poses and
969 binding-free energies, but it has sometimes trouble distinguishing between compounds with similar binding affinities. Molecular
970 docking algorithms often prioritize speed over accuracy by using fast pose generation and approximate scoring functions
971 to save time, especially when screening extensive collections of compounds against a protein target. This trade-off reduces the
972 effectiveness of the algorithms. More precise algorithms based on molecular dynamics simulations are sometimes used down-
973 stream of docking in drug discovery processes to improve the number of accurate predictions. The combination of molecular
974 docking and molecular dynamics can yield results near experimentally determined crystal structures (72, 73). AFVS includes
975 methods for selecting and scoring poses based on molecular dynamics simulations, such as MM/PBSA and MM/GBSA, for
976 evaluating docking poses and predicting binding affinities (74). MM/PBSA and MM/GBSA are popular for binding free energy
977 prediction because they are more accurate than molecular docking and less computationally intensive than alchemical methods.
978 AdaptiveFlow also supports Binding Pose Metadynamics (BPMD), which uses a stability score based on the resistance of the
979 ligand to perturbation away from the initial pose to rerank predicted binding poses from docking (72, 75). The OpenBPMD
980 algorithm can accurately predict binding poses within 2 Å RMSD of the crystallographic complex structure for many systems
981 and efficiently rerank docked poses. The MM/PBSA, MM/GBSA, and BPMD methods can be used within AdaptiveFlow as
982 rescoring and pose refinement methods.

983 **Enhanced Quality Checks.** The docking pose of each ligand screened by AFU is checked for corruption by evaluating
984 the coordinates and computing the potential energy of the docked compound with Open Babel. Any corrupted compounds
985 are automatically removed. AFVS supports optional plausibility checks on generated docking poses of Vina-based docking

986 methods using PoseBusters, an open-source tool, which checks for chemical as well as intramolecular ligand and intermolecular
987 validity of ligand-protein complexes (64).

988 Experimental Validation Studies Involving FSP1

989 Materials and methods used in the experimental validation of the molecules for FSP1 are described below.

990 **Chemicals.** All chemicals and proteins employed for binding and enzymatic assays were purchased from Sigma Aldrich,
991 except for 2,3,5-trimethylcyclohexa-2,5-diene-1,4-dione (α -tocopherolquinone hydrophilic head) purchased from Synthonix
992 and menaquinone (Vitamin K2) from Supelco Analytical. FSP1 inhibitor candidates were purchased from Enamine; iFSP1
993 from Medchemexpress; viFSP1 and FSEN1 from Cayman Chemical.

994 **Phylogenetic Inference and Ancestral Sequence Reconstruction.** Phylogenetic analysis and ancestral sequence recon-
995 struction were carried out following the methods described in (76). Briefly, human FSP1 sequence (UniprotID: Q9BRQ8) was
996 employed as query for homology searches with BLASTP. Datasets were constructed vetting all chordate classes according to
997 TimeTree (77) by mining at least two species with fully sequenced genomes. Multiple sequence alignments were constructed in
998 MAFFT v7 (78) and manually trimmed for single sequence insertions/terminal extensions. Neighbour-joining guide trees were
999 constructed in MEGA v10.2 to assess the quality of datasets under construction. Once the working multiple-sequence align-
1000 ments were obtained, best-fit substitution models and gamma distribution values (α) were calculated using ProtTest (data for
1001 each protein phylogeny are shown in Supplementary Fig. 11B). Maximum likelihood phylogenies were inferred using RaxML
1002 v8.2.10 (HPC-PTHREADS module (79)), using rapid bootstrap analysis and searching for the best-scoring ML tree, with 500
1003 bootstraps replicates and the given best-fit model under gamma distribution. When required, a species tree was used to constrain
1004 the phylogeny, this was constructed using TimeTree (timetree.org). Once the phylogenies were inferred, the bootstrap values
1005 were subjected to transfer bootstrap expectation values using BOOSTER online (80). Figtree v1.4.2 was employed for analysing
1006 and visualizing the trees. Ancestral sequence reconstruction was performed employing PAML v4.9a (CODEML module) as
1007 marginal reconstruction, using the phylogenies obtained previously, empirical amino acid substitution model (model=2) and
1008 JTT substitution matrix, four gamma categories and re-estimation of gamma shape parameter (81, 82). The distribution of the
1009 posterior probabilities for each of the ancestral states was analysed at the node corresponding to the tetrapod ancestor in the
1010 phylogeny. Sites that displayed posterior probabilities <0.8 were considered ambiguously reconstructed when alternative states
1011 displayed posterior probabilities >0.2 (83). The length of the targeted nodes was treated by Fitch's parsimony. To visualize the
1012 degree of sequence conservation across tetrapod organisms, multiple sequence alignments were generated using ESPript 3 (84).

1013 **Construct design and expression of FSP1.** Plasmid (Genscript) coding for His6-SUMO-ancestral FSP1, His6-SUMO-
1014 ancestral L323A FSP1 and His6-SUMO-ancestral F15A FSP1 were transformed by heat shock into *E. coli* BL21(DE3) cells (25
1015 s, 42°C). Cells from a single colony were pre-inoculated into 100 mL LB broth containing 100 μ L mL⁻¹ kanamycin and grown
1016 overnight at 37°C. Pre-cultures were inoculated 100x into 1 L Terrific Broth and grown at 37 °C, 200 r.p.m. for ca. 3 h until
1017 the optical density (OD600) reached 0.5-0.7. Protein expression was then induced with isopropyl β -D-1-thiogalactopyranoside
1018 (0.1 mM final), and incubated for 16 h at 30 °C and 200 r.p.m. Cells were harvested by centrifugation (5,000 g, 15 min, 10 °C),
1019 flash frozen in liquid nitrogen, and stored at -20 °C.

1020 **Protein purification.** Cells (ca. 29 g) were re-suspended 1:5 (grams:milliliters) in Buffer A (50 mM HEPES pH 7.2, 250 mM
1021 NaCl, 10% (v/v) glycerol) supplemented with 1 mM FAD and protease inhibitors (1 mM phenylmethylsulphonyl fluoride, 10
1022 μ M leupeptin, 10 μ M pepstatin) and 5 μ g DNase I/g cells. Cell resuspensions were stirred at 4 °C for 25 min before cell
1023 lysis, which was performed with a high-pressure homogeniser (Emulsiflex c-3, ATA Scientific) in three cycles. Cell lysates
1024 were centrifuged (4 °C, 56,000 g, 1 h) using an Avanti J-26 XP centrifuge equipped with a JA-25.15 rotor (Beckman Coulter).
1025 The resulting supernatant was filtered through a 0.45 μ m filter and loaded on a gravity column containing Nickel Sepharose
1026 HP IMAC resin (Cytiva) pre-equilibrated with Buffer A. The crude extract was passed onto the column twice. The resin
1027 was then washed with at least five column volumes of Buffer A. The wash up of the resin was performed adding 5 column
1028 volumes of Buffer A supplemented with increasing concentrations of imidazole, 5 to 30 mM. The protein was then eluted with
1029 at least ten column volumes of Buffer B (Buffer A supplemented with 300 mM imidazole). The purified protein sample was
1030 concentrated using an Amicon Ultra (Merck) with a 30 kDa cut-off up to 2.5 mL. The imidazole was then removed with a
1031 PD-10 batch desalting column (Cytiva) pre-equilibrated with Buffer A. To cleave the His6-SUMO tag, the desalted sample was
1032 concentrated down to 800 μ L and then added of a His6-tagged SUMO protease (1.2 mg mL⁻¹) to a volume ratio of 100:1 and
1033 incubated overnight at 4 °C on a rotating wheel. The sample was then loaded onto an Äkta Pure system (Cytiva) equipped
1034 with a Ni HiTrap-HP 5 mL column (Cytiva) pre-equilibrated in Buffer A and proteins were eluted out using an imidazole
1035 concentration gradient of 2% Buffer B corresponding to 6 mM imidazole. The elution of the tag-less protein was monitored
1036 with a multiwavelength detector following the absorbance at 280 nm and 450 nm. The sample was loaded onto an Äkta Pure
1037 system (Cytiva) equipped with a Superdex 200 10/300 GL (Cytiva) pre-equilibrated in 100 mM NaCl, 50 mM HEPES pH
1038 7.2. The purified protein was then concentrated down to about 1 mM, measuring the concentration using a NanoDrop ND-100

1039 UV/Vis spectrophotometer (Thermo Scientific) based on the absorption at 280 nm and at 458 nm (ϵ_{280} nm predicted by
1040 ProtParam tool - Expasy). Purity of the sample was evaluated by SDS-PAGE. The purified protein was finally flash frozen in
1041 liquid nitrogen and stored at -80 °C.

1042 **Enzyme kinetics.** The activity of FSP1 and mutants was monitored using 0.2-1 μ M protein (estimated by 458 nm absorbance)
1043 in a 150 μ L final volume of Buffer A. Reactions were followed in 10.00 mm quartz cuvettes (Hellma) and a Cary Eclipse
1044 Fluorescence Spectrophotometer (Agilent) or a Cary 100 UV-vis spectrophotometer (Varian) equipped with a thermo-stated
1045 cell holder (T = 25 °C). Reactions were started by adding NADH and rates were determined by following NADH consumption,
1046 caused by the oxidation to NAD⁺ (excitation 340 nm, emission 460 nm or monitoring the absorbance at 340 nm). A calibration
1047 line was built by measuring the fluorescence at known NADH concentrations. GraphPad Prism 9 was used to perform both
1048 linear and non-linear regression.

1049 **Superoxide assay.** Superoxide radicals generated by oxidation of NADH were detected by the reduction of nitroblue tetra-
1050 zolum to nitroblue diformazane using Tris-HCl buffer (16 mM, pH 8.0) 50 μ M nitroblue tetrazolium, and 0.2 μ M final FSP1
1051 (final volume 100 μ L). The reaction was started by adding various concentrations of NADH. The absorbance was measured at
1052 560 nm using a ClarioStar plate reader.

1053 **Enzyme inhibition.** K_i were measured following NADH oxidation by monitoring the absorbance 340 nm with Cary 100 UV-
1054 vis spectrophotometer (Varian). 0.2-1 μ M purified FSP1 was added to a final volume of 150 μ L of 50 mM HEPES buffer pH
1055 7.2. Inhibitors were tested with concentrations varying from micromolar to low nanomolar.

1056 **Resazurin assay.** The activity of FSP1 and mutants was monitored using 0.2-0.1 μ M (final, absorbance at 458 nm) protein
1057 in a 100 μ L 50 mM HEPES buffer pH 7.2 and 50 μ M resazurin. The activity was monitored by the formation of resorufin
1058 (excitation 572 nm, emission 583 nm) using varying concentrations of iFSP1 and viFSP1 (2, 44, 45), aF-FSP1-1, aF-FSP1-2,
1059 aF-FSP1-3, aF-FSP1-4, aF-FSP1-5 and aF-FSP1-6. The reaction was started by adding 200-100 μ M NADH. The fluorescence
1060 was read using a ClarioStar plate reader, with a gain of 800.

1061 **Nano Differential Scanning Fluorimetry.** Thermo-stability analyses were carried out using a TychoTMNT.6 system (Nan-
1062 oTemper GmbH). The protein was diluted to 1 mg mL⁻¹ in Buffer A and pre-incubated with ligands at room temperature for
1063 15 minutes before experiments. Melting curves were obtained following the intrinsic fluorescence of tryptophan and tyrosine
1064 residues (emission 350 nm and 330 nm, respectively) applying a temperature gradient from 35 °C to 95 °C for a time range of
1065 three minutes. Data were analysed as F350/F330 ratio and derived to determine the inflection temperature (T_m). All measure-
1066 ments were performed in duplicate. The following condition were tested: FSP1 in presence of 1 mM FAD, NAD⁺ or 1 mM
1067 CoQ₁ or 1 mM α -tocopherolquinone hydrophilic head or 1 mM vitamin K3 or 10 μ M inhibitors or 5% DMSO as a control.

1068 **ThermoFAD.** Experiments were performed with 20 μ M FSP1 after 10 min incubation with inhibitors (10 μ M) or DMSO in
1069 HEPES buffer (50 mM, pH 7.2) (85). The temperature gradient was set to 30–90 °C with fluorescence detection at a constant
1070 rate 0.1 °C/s at 450 ± 30 nm excitation and 535 ± 30 nm emission (BioRad MiniOpticon Real-Time PCR System). Data were
1071 normalized and fitted with sigmoidal non-linear regression using GraphPad Prism 9.

1072 **Cellular thermal shift assay.** 30x10⁶/mL HEK293 cells were incubated with DMSO or 10 μ M inhibitor for 1 h at 37 °C in 2
1073 mL of culture medium. After centrifugation at 1,200 g for 5 min at room temperature, cell pellets were washed twice in PBS,
1074 re-suspended at 30x10⁶/mL, and aliquoted in PCR-strips at 50 μ L per tube. Samples were then heated (ranging from 38.7 °C
1075 to 86.2 °C) for 3 min in a T100 Thermal Cycler (Bio-Rad). Cells were lysed by adding 1% (v/v) triton X-100, and proteases
1076 inhibitors (see Protein purification) to each tube for 60 min at 4 °C. After centrifugation at 20,000 g for 20 min at 4 °C, each
1077 supernatant was analysed by loading 15 μ L of each sample onto a SDS gel (Bio-Rad, Mini-PROTEAN TGX). Gel bands were
1078 transferred to a Trans-Blot Turbo Mini 0.2 μ m PVDF (polyvinylidene difluoride) membrane (Bio-Rad) using a Trans-Blot Turbo
1079 Transfer System (Bio-Rad). Western blot analysis was performed by immunoblotting the PVDF membrane with Anti-AIFM2
1080 Antibody, clone 1I15 ZooMAb® Rabbit Monoclonal (Sigma-Aldrich), following the addition of goat horseradish peroxidase-
1081 conjugated antirabbit secondary antibody (Millipore). Chemiluminescent detection was performed by using the ECL detection
1082 kit (Bio-Rad, ChemiDoc MP image). Melting curves were generated by plotting the band intensities of each heated sample
1083 against the respective temperatures (°C). Data were first normalized by setting the highest and lowest intensity value in each set
1084 as 100 and 0%, respectively. Non-linear regression was performed by applying a Boltzmann sigmoidal dose-response model
1085 within GraphPad Prism 9. The T_{agg} values obtained were the results of two independent replicates and the mean ± standard
1086 deviation was calculated.

1087 **Crystallisation and structural determination of FSP1.** Crystallisation was performed using 13 g/mL in 50 mM HEPES
1088 pH 7.2, 1 mM FAD, 1 mM NAD⁺ with or without 1 mM coenzyme Q₁. Small crystals or crystalline precipitate appeared in a
1089 few days under several conditions of the Mopheus screening kit (Molecular Dimensions) using the vapour diffusion sitting drop
1090 methods. The best looking crystals grew in HEPES/MOPS 0.1 M pH 7.5, 12.5% PEG 1000, 12.5% PEG 3350, 12.5% MPD,

1091 0.06 M MgCl₂, 0.06 M CaCl₂ at 20 °C with a ratio of 1:1 in a sitting drop. After harvesting, they were smashed to generate
1092 small crystals, and used for microseeding. Specifically, the drops were composed of 0.15 µL protein, 0.05 µL seeds, and 0.2 µL
1093 reservoir from the following commercial screens.

1094 • FSP1-FAD structure; 0.06 M MgCl₂, 0.06 M CaCl₂, 0.1 M HEPES/MOPS (acid) pH 7.5; 20% v/v PEG500*MME, 10%
1095 PEG20000.

1096 • FSP1-FAD-NAD⁺ structure; 0.06 M MgCl₂, 0.06 CaCl₂, 0.1 M HEPES/MOPS (acid) pH 7.5, 12.5% PEG1000, 12.5%
1097 PEG3350, 12.5% MPD.

1098 • FSP1-FAD-NAD⁺-coenzyme Q₁ structure; 0.06 M MgCl₂, 0.06 CaCl₂, 0.1 M imidazole/MES monohydrate (acid) pH
1099 6.5, 20.0% PEG500* MME, 10% PEG20000. Crystals were then soaked with 20 mM coenzyme Q₁ added to the crys-
1100 tallisation drop.

1101 • FSP1-FAD-NAD⁺-afi-FSP1-1; 0.06 M MgCl₂, CaCl₂, 0.1M imidazole/MES monohydrate (acid) pH 6.5, 12.5% MPD,
1102 12.5% PEG1000, 12.5% PEG3350. Crystals were then soaked with 1 mM afi-FSP1-1 added to the crystallisation drop.

1103 • FSP1-FAD-NAD⁺-afi-FSP1-3; 0.12 M diethylene glycol, 0.12 M triethylene glycol, 0.12 tetraethylene glycol, 0.12 pen-
1104 taethylene glycol, 0.1 M tris (base)/bicine pH 8.5, 12.5% MPD, 12.5% PEG1000, 12.5% PEG3350. Crystals were then
1105 soaked with 10 mM afi-FSP1-3 added to the crystallisation drop.

1106 • FSP1-FAD-NAD⁺-afi-FSP1-4; 0.06 M MgCl₂, 0.06 M CaCl₂, 0.1 M imidazole/MES monohydrate (acid) pH 6.5, 12.5%
1107 MPD, 12.5% PEG1000, 12.5% PEG3350. Crystals were then soaked with 10 mM afi-FSP1-4 added to the crystallisation
1108 drop.

1109 • FSP1(L323A)-FAD-NAD⁺-afi-FSP1-2; 0.06 M MgCl₂, 0.06 M CaCl₂, 0.1 M HEPES/MOPS (acid) pH 7.5, 20% ethylene
1110 glycol, 20% PEG8000. Crystals were then soaked with 1 mM afi-FSP1-2 added to the crystallisation drop.

1111 Electron densities from co-crystallisations with previously reported inhibitors (Supp. Table 5) did not show any ligand
1112 binding. Data were collected at the European Synchrotron Radiation Facility (Grenoble, France) and processed with the XDS
1113 and CCP4 packages (86, 87). The phase problem was solved by molecular replacement using an Alpha Fold model and CCP4
1114 programs. Model building and refinement were conducted using COOT and Refmac5 (88). Figures were then generated using
1115 ChimeraX (89), PyMOL (DeLano Scientific; www.pymol.org).

1116 **SPR Binding Assays.** SPR binding assays were carried out using the Biacore 8K SPR system (Cytiva) with the running
1117 buffer (50 mM HEPES, pH 7.2, 150 mM NaCl, 0.05% Tween-20, and 2% DMSO) under controlled conditions at 25 °C.
1118 The His-Sumo–FSP1 protein was immobilized on a preconditioned and activated Series S NTA sensor chip (Cytiva). Pre-
1119 conditioning involved treatment with 350 mM EDTA, followed by activation with 0.5 mM NiCl₂. To enhance stability and
1120 minimize signal drift, the chip was further activated using EDC (1-ethyl-3-(3-dimethylaminopropyl) carbodiimide) and NHS
1121 (N-hydroxysuccinimide) before coupling the protein to the surface. The protein was immobilized to a response level of 7500
1122 RU.

1123 Compounds were applied in a three-fold serial dilution series (6, 18.9, 60, 189, 600, 1897, 6000, 18973, and 60000 nM)
1124 using a single-cycle kinetics (SCK) setup. Compound plating and DMSO normalization were performed using the Tecan D-
1125 300e digital dispenser (Tecan Life Sciences). Binding was monitored for 60 s during the association phase and 120 s during the
1126 dissociation phase. Reference channel signals and buffer blanks were used to correct the raw binding signals (RU).

1127 Binding data were collected from three independent experiments; however, only a single representative dataset is shown
1128 for each compound in the manuscript. Data processing and binding analysis were performed using Biacore Insight software
1129 (Cytiva). Steady-state binding levels (R_{eq}) were plotted against compound concentrations and fitted to a one-site binding model
1130 using the Levenberg–Marquardt algorithm, enabling accurate determination of binding affinities for each compound.

1131 **Virtual Screens.** The AdaptiveFlow open-source drug-discovery platform was used to carry out the virtual screen. The size
1132 of the docking box was set to 14 Å × 18 Å × 16 Å. The FSP1-FAD-NAD⁺ complex structure was held rigid during the virtual
1133 screens (ATG prescreen and primary screen). QuickVina 2 (90) using a hybrid empirical and knowledge-based scoring function
1134 was used as the docking program. The exhaustiveness value was set to 1. The number of replicas was set to 1. The virtual
1135 screen was run in Amazon Web Services (AWS). Postprocessing of the results including clustering with the similarity limit set
1136 to 0.8 was done with DataWarrior (v6.0.0) (91).

1137 **Discovery of Inhibitor Analogs.** Analog Hunter and Scaffold Hopper modules in the infiniSee (92) module from BioSolveIT
1138 were used to screen analogs of compound afi-FSP1-1. Top-ranking 5000 compounds from each module were exported to
1139 SeeSAR (93) from BioSolveIT and docked with the FlexX algorithm in the general docking mode to the coenzyme Q binding
1140 site of the FSP1/FAD/NAD/afi-FSP1-1 complex structure (PDB code: 9IFU) with compound afi-FSP1-1 removed. The docked

1141 poses were scored with the HYDE scoring function. Top-ranking 168 compounds with similar binding pose as compound afi-
1142 FSP1-1 were tested with 100-ns-long molecular dynamics (MD) simulations with the PMEMD module (CUDA version) of the
1143 AMBER 22 package and protein.ff19SB force field for protein (94). The Molecular Mechanics/Generalized Born Surface Area
1144 (MM/GBSA) binding free energies between protein and the ligands that remained in the binding pocket were calculated by
1145 the MMPBSA.py script in AmberTools 23 (95). The top-ranking 20 analogs with the lowest MM/GBSA binding free energies
1146 were ordered from Enamine for testing.

1147 **Data Availability**

1148 The ready-to-dock Enamine REAL Space can be explored online via an interactive web interface available on the AdaptiveFlow
1149 homepage [http://adaptive-flow.ai/](https://adaptive-flow.ai/). The data can be accessed directly via an AWS S3 bucket provided by AWS
1150 Open Data (<https://aws.amazon.com/opendata>). During the review, the reviewers can access the data via the following
1151 bucket: `s3://vf-library-access-pu-9xyrt6ot89qneyzyycp9ippn8qhkuse2b-s3alias`. Additional information is available on the AWS Open Data website of this dataset https://aws.amazon.com/marketplace/pp/prodview-m32p7engdxk7k?sr=0-1&ref_=beagle&applicationId=AWSMPContessa, as well as on the
1152 AdaptiveFlow homepage <https://adaptive-flow.ai/real-space-2022q12> (free registration required).

1153 **Code Availability**

1154 The source code of AdaptiveFlow is released under the GNU GPL v2.0, a free and open-source license. The code can be
1155 accessed via GitHub. The AdaptiveFlow GitHub repos can be found at <https://github.com/LigandUniverse>,
1156 which contains multiple repositories. The AFLP and AFVS modules of AdaptiveFlow can be found in AFLP (<https://github.com/LigandUniverse/AFLP>) and the AFVS (<https://github.com/LigandUniverse/VFVS>)
1157 GitHub repositories, respectively. The modules AdaptiveFlow Unity and AdaptiveFlow Unity Parallelized can be found in
1158 the GitHub repo named AFU (<https://github.com/LigandUniverse/AFU>) and AFUparr (<https://github.com/LigandUniverse/AFUparr>), respectively.

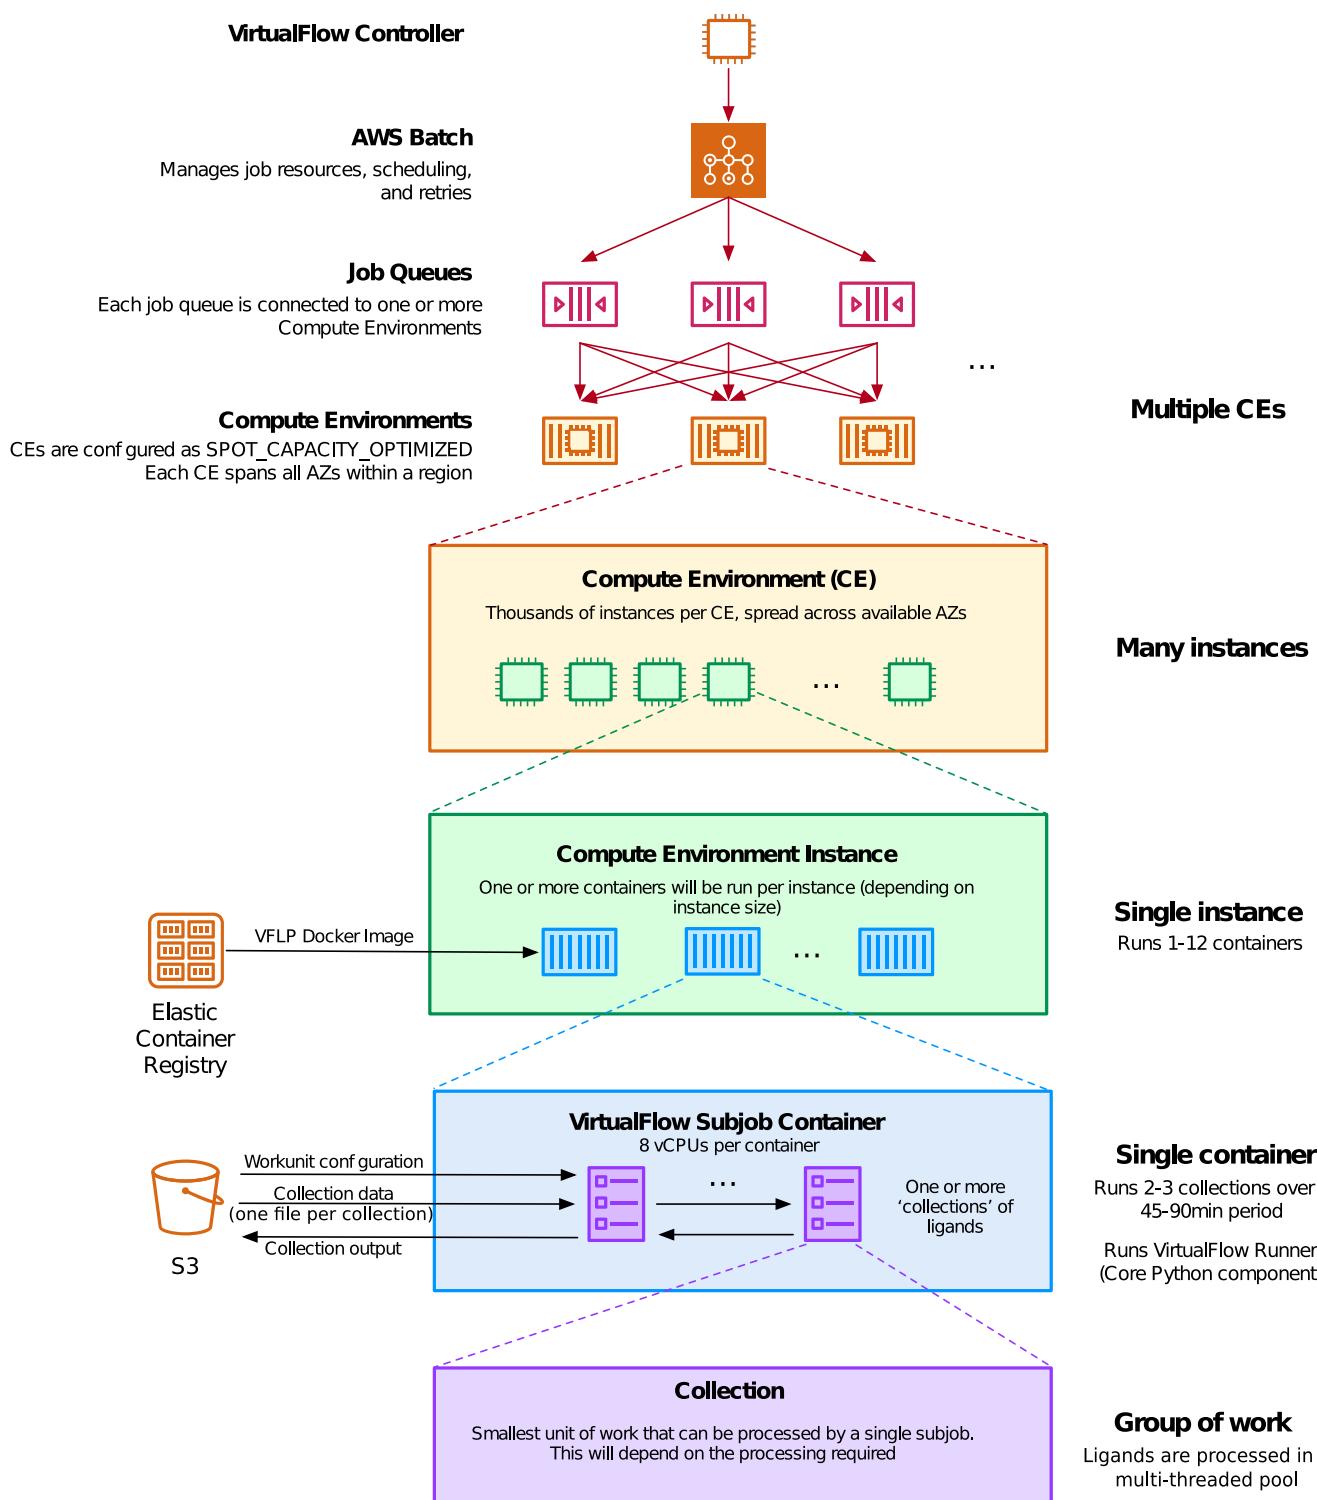
1159 **Acknowledgments**

1160 The authors would like to thank ChemAxon for a free academic license for JChem. AK.N. thanks Anshul Kundaje and Michael
1161 Bassik for their support. The authors thank AWS for research credits for the AWS Cloud and thank Nelson Gonzales, Stephen
1162 Litster, Mayank Thakkar, Robert Wang, Arun Subramaniyan, and Christine Tsien Silvers for technical and administrative
1163 support within AWS. The authors would like to thank Akanksha Bilani and Caroline Rhoades from Intel for their support,
1164 and Brigitte Klein for proofreading the manuscript. AK.N. acknowledges funding from the Bio-X Stanford Interdisciplinary
1165 Graduate Fellowship (SGIF). C.S. acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research
1166 Foundation) under Germany's Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project
1167 ID: 390685689). H.A. acknowledges funding from NIH (GM136859) and a gift from J. Goldberg, Dana-Farber Donor. A.M
1168 acknowledges funding from the European Research Council (ERC Advanced Grant, MetaQ no. 101094471), and by the
1169 Associazione Italiana per la Ricerca sul Cancro (AIRC) Investigator Grant (no. 28754). This work was supported by ALSAC,
1170 the fundraising and awareness organization of St. Jude Children's Research Hospital, and the Blue Sky Kinase Project at St.
1171 Jude Children's Research Hospital.

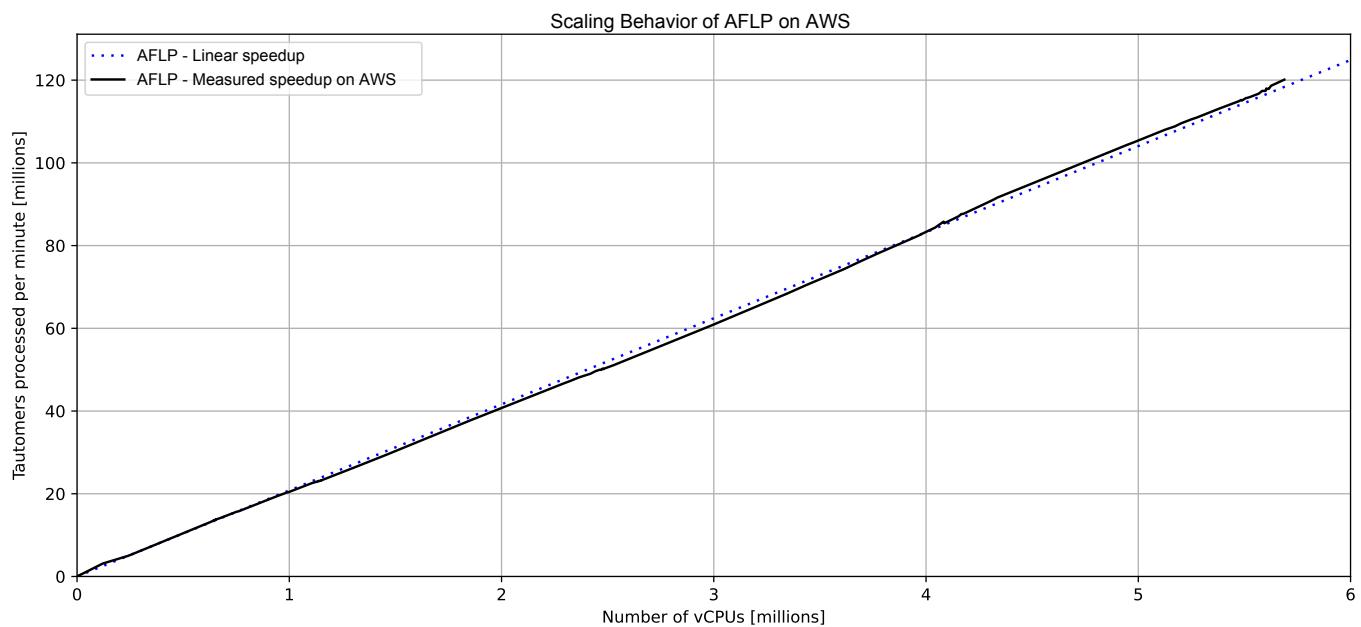
1172 **Competing Interests**

1173 The authors declare the following competing interests: D.S.R. works for Enamine, a company that is involved in the synthesis
1174 and distribution of drug-like compounds. Y.S.M is a scientific advisor to Chemspace LLC and VP of Sales and Marketing at
1175 Enamine Ltd. C.G. and G.W. are cofounders of Virtual Discovery Inc., which is a fee-for-service company for computational
1176 drug discovery. M.G. works for Google, a company providing cloud computing services. M.K. and P.-Y.A. work for Amazon
1177 Web Services, a company providing cloud computing services.

1182 **Extended Data**



Extended Data Fig. 1. Compute Stack deployed by AdaptiveFlow in AWS. The login node is used by the user to control the workflow in AWS via the workload manager AWS Batch. AWS Batch has multiple Job Queues to which jobs are submitted by AdaptiveFlow, and each job queue deploys multiple Compute Environments that can spin up and use a large number of virtual machines (instances) of specific types. On each instance, one or more subjobs are run, and each subjobs deploys a docker container that carries out the computations of AdaptiveFlow by executing AdaptiveFlow Runner (a core python component of AdaptiveFlow). Each subjob typically processes a few ligand collections, where each ligand collection contains multiple ligands (e.g. 1000).



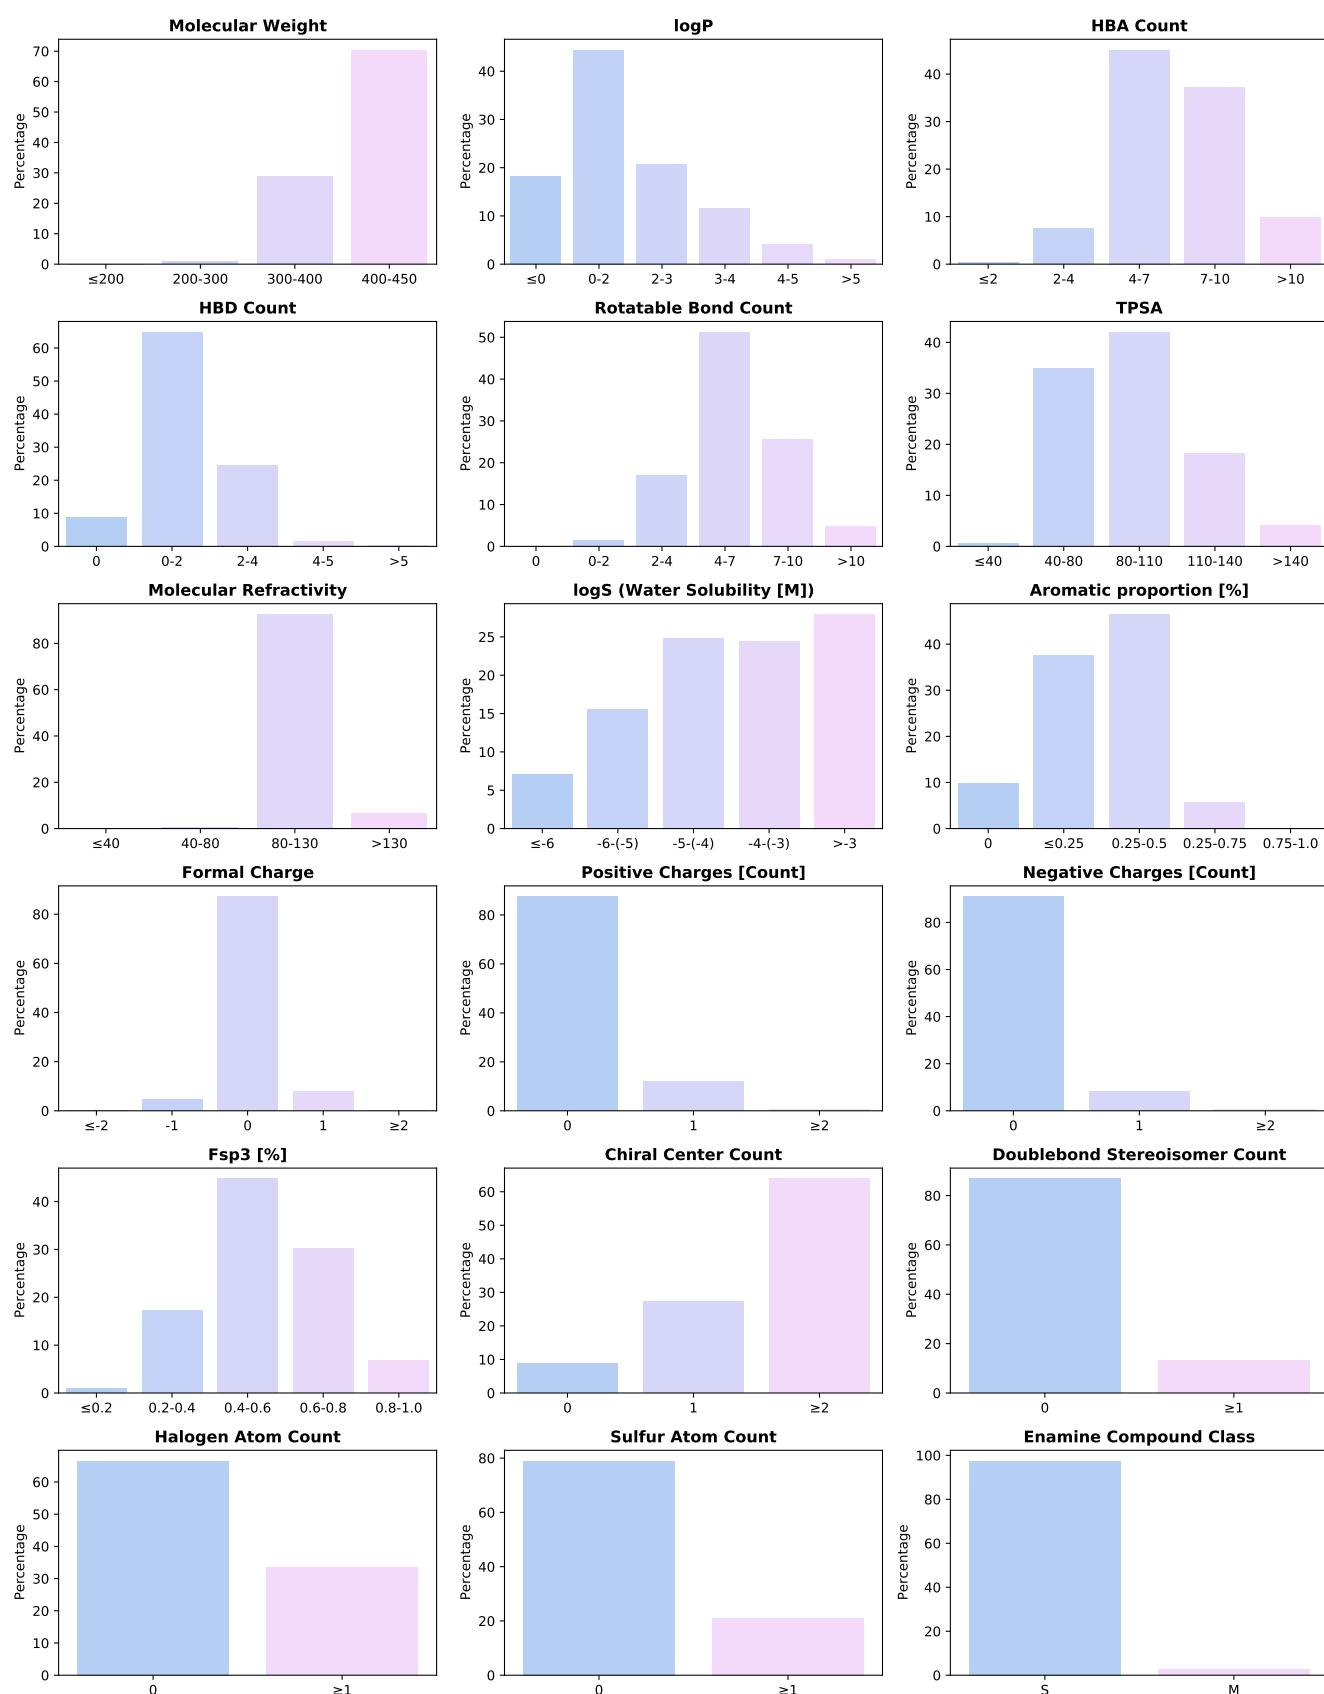
Extended Data Fig. 2. Scaling behavior of AFLP on AWS. The scaling behavior is perfectly linear even when using extremely large numbers of CPUs. Up to 5.6 million vCPUs were used in parallel across multiple AWS regions. As can be seen in the plot, the scaling behavior is perfectly linear, implying there is no slowdown even when using extremely large numbers of vCPUs.

Molecular Property (X)	Tranche ID	A	B	C	D	E	F
MW	1	X ≤ 200 3.8E+6	200 < X ≤ 300 564.9E+6	300 < X ≤ 400 19.8E+9	400 < X ≤ 450 48.4E+9	X > 450 0.0E+0	
logP	2	X ≤ 0 12.5E+9	0 < X ≤ 2 30.5E+9	2 < X ≤ 3 14.2E+9	3 < X ≤ 4 8.0E+9	4 < X ≤ 5 2.8E+9	X > 5 706.3E+6
HBA count	3	X ≤ 2 270.9E+6	2 < X ≤ 4 5.1E+9	4 < X ≤ 7 31.0E+9	7 < X ≤ 10 25.5E+9	X > 10 6.8E+9	
HBD count	4	X = 0 6.0E+9	0 < X ≤ 2 44.5E+9	16.8E+9	2 < X ≤ 4 1.2E+9	4 < X ≤ 5 277.2E+6	X > 5
Rotatable bond count	5	X = 0 2.0E+6	0 < X ≤ 2 966.3E+6	2 < X ≤ 4 11.6E+9	4 < X ≤ 7 35.2E+9	7 < X ≤ 10 17.6E+9	X > 10 3.3E+9
TPSA	6	X ≤ 40 377.6E+6	40 < X ≤ 80 24.0E+9	80 < X ≤ 110 28.9E+9	110 < X ≤ 140 12.6E+9	X > 140 2.9E+9	
Molecular refractivity	7	X ≤ 40 79.2E+3	40 < X ≤ 80 540.6E+6	80 < X ≤ 130 63.6E+9	X > 130 4.6E+9		
logS	8	X ≤ -6 4.9E+9	-6 < X ≤ -5 10.7E+9	-5 < X ≤ -4 17.1E+9	-4 < X ≤ -3 16.8E+9	X > -3 19.2E+9	
Aromatic proportion	9	X = 0 6.8E+9	0 < X ≤ 0.25 25.9E+9	0.25 < X ≤ 0.5 31.9E+9	0.25 < X ≤ 0.75 4.0E+9	0.75 < X ≤ 1.0 49.7E+6	
Formal charge	10	X ≤ -2 158.3E+6	X = -1 3.1E+9	X = 0 59.9E+9	X = 1 5.4E+9	X ≥ 2 140.9E+6	
Positive charge count	11	X = 0 60.1E+9	X = 1 8.2E+9	X ≥ 2 341.7E+6			
Negative charge count	12	X = 0 62.5E+9	X = 1 5.7E+9	X ≥ 2 484.1E+6			
Fsp3	13	0 ≤ X ≤ 0.2 647.0E+6	0.2 < X ≤ 0.4 11.9E+9	0.4 < X ≤ 0.6 30.8E+9	0.6 < X ≤ 0.8 20.7E+9	0.8 < X ≤ 1.0 4.7E+9	
Chiral center count	14	X = 0 6.1E+9	X = 1 18.7E+9	X ≥ 2 43.9E+9			
Doublebond stereoisomer count	15	X = 0 59.6E+9	X ≥ 1 9.1E+9				
Halogen atom count	16	X = 0 45.7E+9	X ≥ 1 23.0E+9				
Sulfur atom count	17	X = 0 54.2E+9	X ≥ 1 14.5E+9				
Enamine compound class	18	X = S 66.9E+9	X = M 1.8E+9				

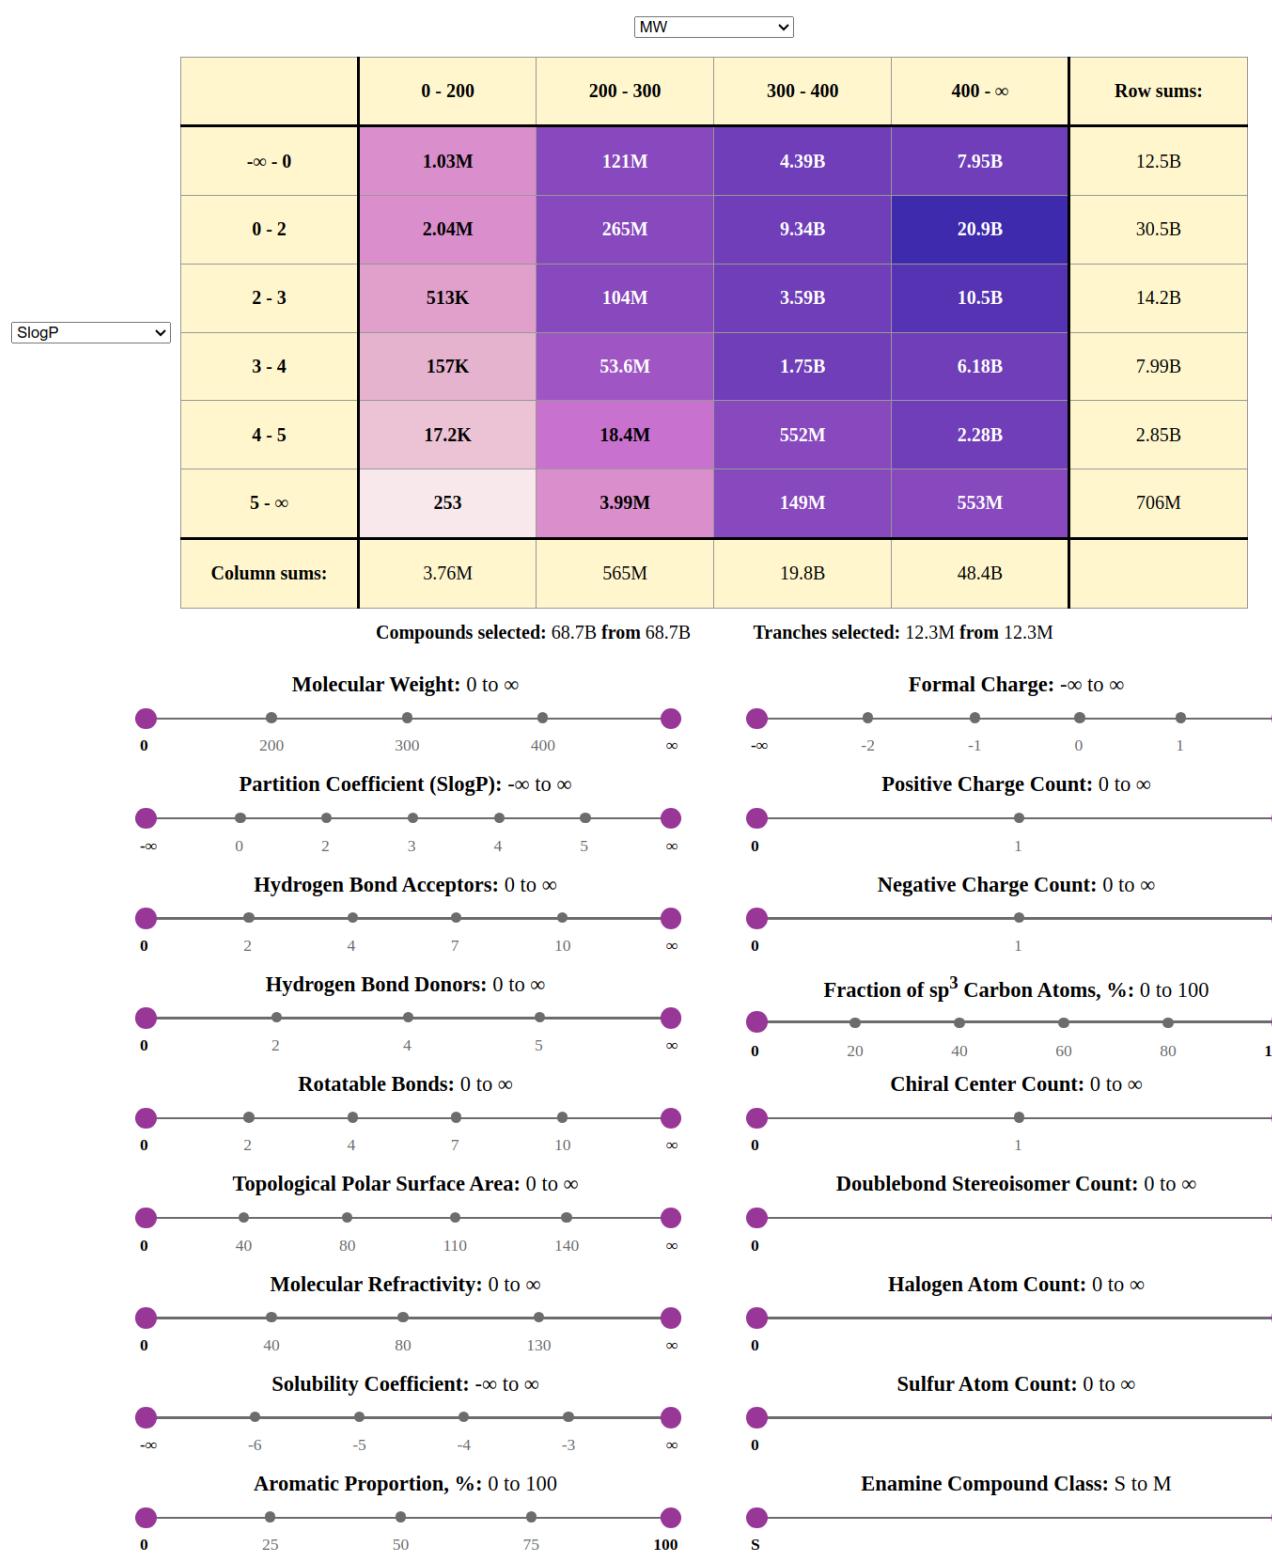
Color Scale



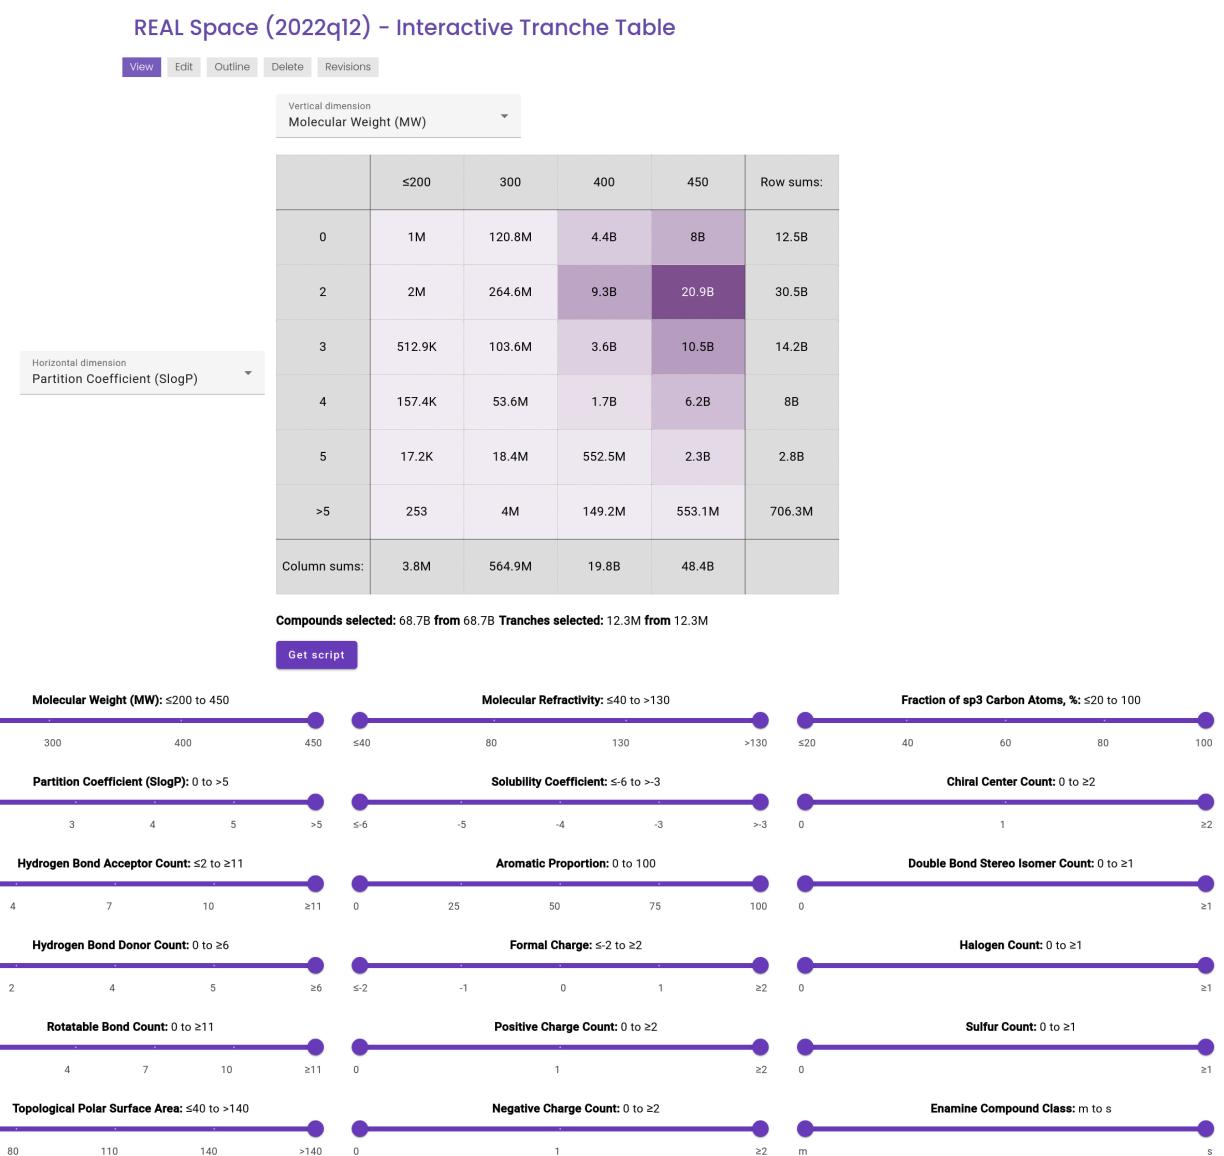
Extended Data Fig. 3. Overview of 18-dimensional tranche table. Molecular properties associated with each tranche, the partitions that are used for each tranche property, and the number of compounds per partition interval for each property (colored as a heat map).



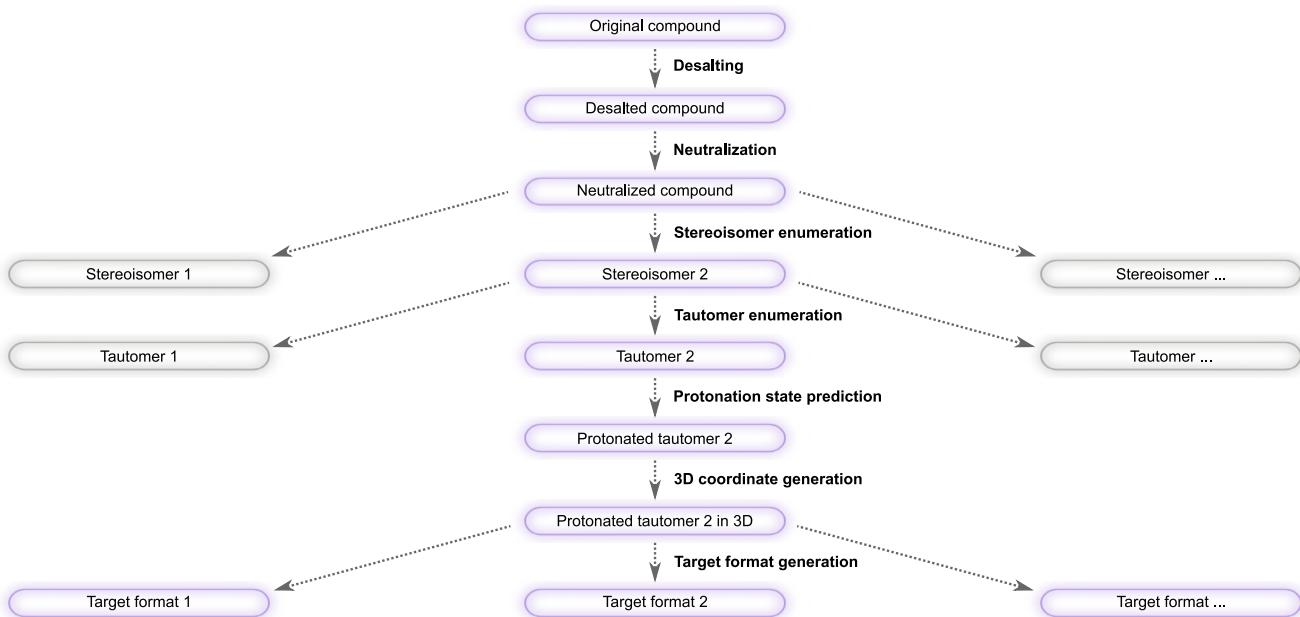
Extended Data Fig. 4. Overview of the distribution of the molecular properties of the Enamine REAL Space (version 2022q1-2). Each subplot shows the distribution of the molecular property associated with each dimension of the 18-dimensional grid into which the REAL Space was partitioned.



Extended Data Fig. 5. Webinterface (version 1) of the REAL Space. Shown is an interactive web interface that is made available on the AdaptiveFlow homepage (<https://adaptive-flow.ai/real-space-2022q12-interface-2>) that allows selecting subsets of the REAL Space partitioned into an 18-dimensional matrix. Each dimension corresponds to a molecular property and has a slider that can be adjusted to select the desired property ranges. The number of selected molecules is updated automatically in the web interface. Effectively, the sliders allow the selection of multidimensional rectangles within the 18-dimensional matrix.



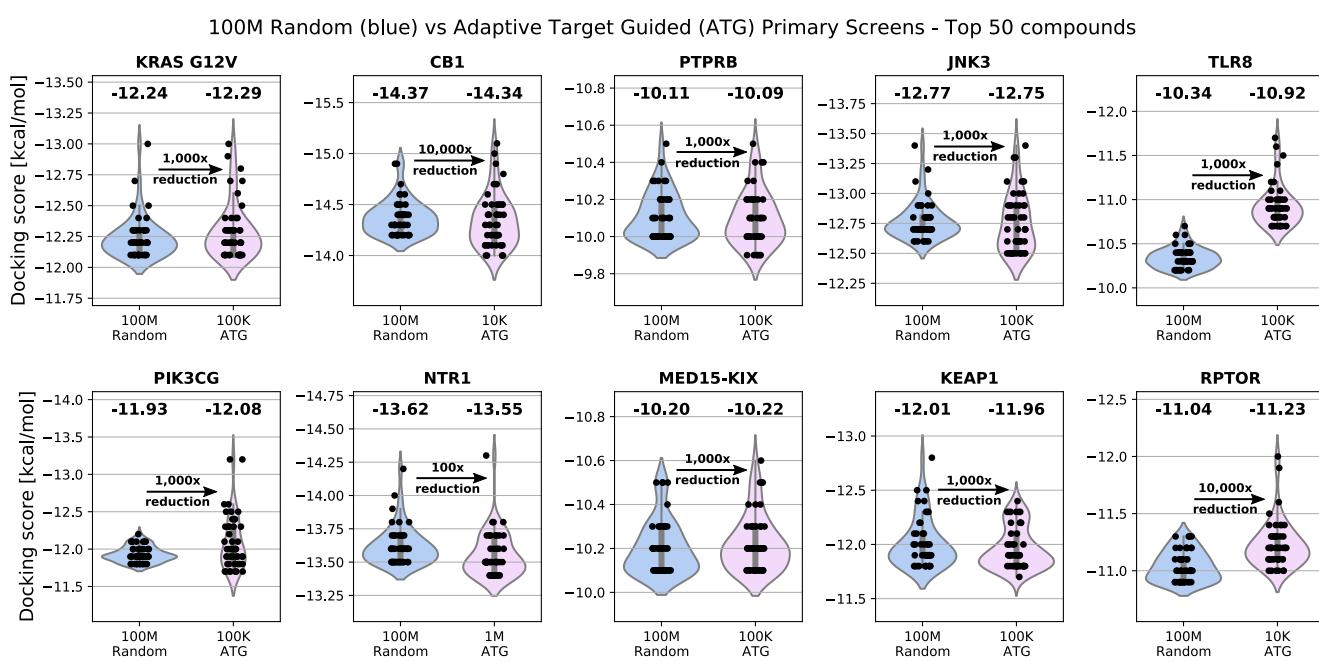
Extended Data Fig. 6. Webinterface (version 2) of the REAL Space. Shown is an interactive web interface that is made available on the AdaptiveFlow homepage (<https://adaptive-flow.ai/real-space-2022q12-interactive-tranche-table>) that allows selecting subsets of the REAL Space partitioned into an 18-dimensional matrix. Each dimension corresponds to a molecular property and has a slider that can be adjusted to select the desired property ranges. The number of selected molecules is updated automatically in the web interface. Effectively, the sliders allow the selection of multidimensional rectangles within the 18-dimensional matrix.



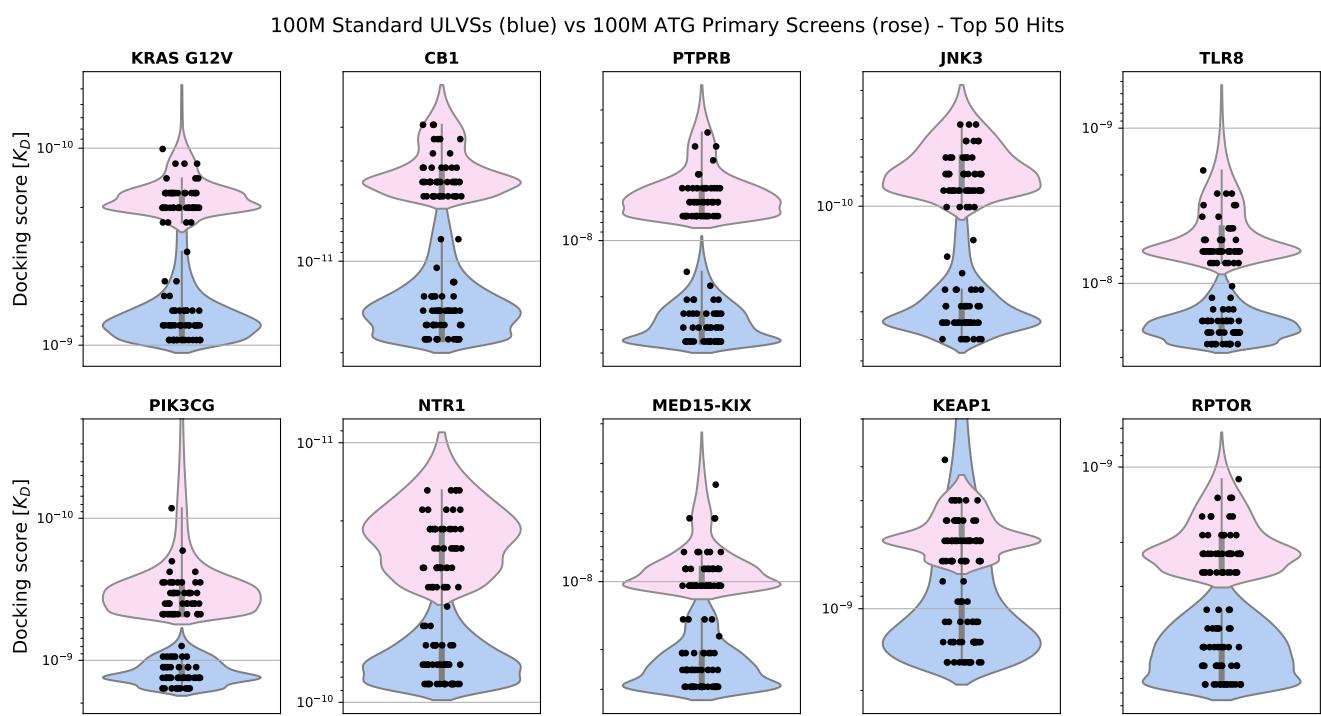
Extended Data Fig. 7. AFLP - Overview of the Ligand Preparation Procedure. Ligands are initially in a format in line notation (e.g. SMILES, SELFIES, or amino acid sequences). The preparation steps include desalting, neutralization, stereoisomer enumeration, tautomer enumeration, protonation state prediction, 3D coordinate generation, and target format generation. Stereoisomer and tautomer generation can result in multiple chemical species. Highlighted in purple are the steps of a molecule from the beginning to the end of the preparation procedure. Colored in white are the intermediate molecules/tautomers whose subsequent preparation steps are not shown in the diagram.

Property (X)	Tranche		Reps/Tranche=1, Exh=1		Reps/Tranche=10, Exh=1		Reps/Tranche=1, Exh=10	
	Letter	Range	Average Score	Rank	Average Score	Rank	Average Score	Rank
Molecular weight	A	X ≤ 200	-5.69868	4	-5.66977	4	-5.75664	4
	B	200 < X ≤ 300	-7.2417	3	-7.20452	3	-7.31959	3
	C	300 < X ≤ 400	-7.79181	1	-7.76165	1	-7.95967	2
	D	400 < X ≤ 450	-7.76098	2	-7.74938	2	-8.0356	1
	A	X ≤ 0	-7.67313	5	-7.66211	6	-7.86196	6
	B	0 < X ≤ 2	-7.68707	4	-7.6674	5	-7.87787	5
	C	2 < X ≤ 3	-7.7057	3	-7.68094	3	-7.90026	4
	D	3 < X ≤ 4	-7.71324	2	-7.69173	2	-7.90955	2
	E	4 < X ≤ 5	-7.7178	1	-7.70489	1	-7.92332	1
	F	X > 5	-7.67281	6	-7.66768	4	-7.90419	3
logP	A	X ≤ 2	-7.22028	5	-7.20911	5	-7.32189	5
	B	2 < X ≤ 4	-7.38033	4	-7.36258	4	-7.53528	4
	C	4 < X ≤ 7	-7.61231	3	-7.59707	3	-7.8139	3
	D	7 < X ≤ 10	-7.79858	2	-7.78421	2	-8.01877	2
	E	X > 10	-8.00413	1	-7.9967	1	-8.22418	1
	A	X = 0	-7.6844	3	-7.66601	4	-7.86789	4
	B	0 < X ≤ 2	-7.70107	2	-7.68555	1	-7.89879	2
	C	2 < X ≤ 4	-7.70723	1	-7.68371	2	-7.90943	1
	D	4 < X ≤ 5	-7.68301	4	-7.6664	3	-7.88586	3
	E	X > 5	-7.63874	5	-7.62658	5	-7.83489	5
Rotatable bond count	A	X = 0	-7.71737	4	-7.7189	4	-7.85219	4
	B	0 < X ≤ 2	-7.94585	2	-7.90628	2	-8.06441	2
	C	2 < X ≤ 4	-7.99004	1	-7.9602	1	-8.11996	1
	D	4 < X ≤ 7	-7.81481	3	-7.79696	3	-7.97132	3
	E	7 < X ≤ 10	-7.35753	5	-7.36396	5	-7.6019	5
	F	X > 10	-6.72586	6	-6.74264	6	-7.2096	6
	A	X ≤ 40	-7.18856	5	-7.18394	5	-7.2971	5
	B	40 < X ≤ 80	-7.489	4	-7.47825	4	-7.66254	4
	C	80 < X ≤ 110	-7.72677	3	-7.71709	3	-7.94197	3
	D	110 < X ≤ 140	-7.90033	2	-7.88201	2	-8.13229	2
TPSA	E	X > 140	-8.10076	1	-8.08286	1	-8.32545	1
	A	X ≤ 40	-5.44523	4	-5.38075	4	-5.52602	4
	B	40 < X ≤ 80	-7.31169	3	-7.25362	3	-7.40034	3
	C	80 < X ≤ 130	-7.73766	2	-7.71934	2	-7.9346	2
	D	X > 130	-7.79962	1	-7.79242	1	-8.07744	1
	A	X ≤ -6	-7.89126	1	-7.87113	1	-8.125	1
	B	-6 < X ≤ -5	-7.83936	2	-7.82017	2	-8.04633	2
	C	-5 < X ≤ -4	-7.75784	3	-7.73812	3	-7.95445	3
	D	-4 < X ≤ -3	-7.67255	4	-7.65238	4	-7.86527	4
	E	X > -3	-7.56383	5	-7.55104	5	-7.75079	5
Molecular refractivity	A	X = 0	-7.00124	5	-6.97872	5	-7.2258	5
	B	0 < X ≤ 0.25	-7.47794	4	-7.4629	4	-7.71839	4
	C	0.25 < X ≤ 0.5	-7.69386	3	-7.687	3	-7.88349	3
	D	0.25 < X ≤ 0.75	-8.07776	2	-8.07503	2	-8.23953	2
	E	0.75 < X ≤ 1.0	-8.38983	1	-8.39619	1	-8.51087	1
	A	X ≤ -2	-8.26463	1	-8.28754	1	-8.43463	1
	B	X = -1	-7.95327	2	-7.94973	2	-8.1436	2
	C	X = 0	-7.66391	3	-7.64593	3	-7.85809	3
	D	X = 1	-7.53418	4	-7.50755	4	-7.7401	4
	E	X ≥ 2	-7.26159	5	-7.21597	5	-7.4746	5
Positive charge count	A	X = 0	-7.64767	3	-7.62676	3	-7.83845	3
	B	X = 1	-7.73326	2	-7.72519	2	-7.93725	2
	C	X ≥ 2	-7.7854	1	-7.78236	1	-7.9847	1
	A	X = 0	-7.47956	3	-7.46405	3	-7.677	3
	B	X = 1	-7.85469	2	-7.86841	2	-8.05519	2
	C	X ≥ 2	-8.23447	1	-8.26297	1	-8.41197	1
	A	0 ≤ X ≤ 0.2	-8.25291	1	-8.24892	1	-8.39246	1
	B	0.2 < X ≤ 0.4	-7.9144	2	-7.91246	2	-8.09817	2
	C	0.4 < X ≤ 0.6	-7.57134	3	-7.5644	3	-7.78127	3
	D	0.6 < X ≤ 0.8	-7.27781	4	-7.26214	4	-7.50539	4
Negative charge count	E	0.8 < X ≤ 1.0	-7.02762	5	-6.99253	5	-7.26721	5
	A	X = 0	-7.68648	2	-7.66472	3	-7.88864	2
	B	X = 1	-7.68156	3	-7.66601	2	-7.87407	3
	C	X ≥ 2	-7.72642	1	-7.70654	1	-7.91929	1
	A	X = 0	-7.7409	1	-7.72163	1	-7.92645	1
	B	X ≥ 1	-7.58906	2	-7.56916	2	-7.81019	2
	A	X = 0	-7.6665	2	-7.65526	2	-7.88305	2
	B	X ≥ 1	-7.7349	1	-7.70825	1	-7.90392	1
	A	X = 0	-7.79078	1	-7.77764	1	-7.98746	1
	B	X ≥ 1	-7.5662	2	-7.53497	2	-7.7627	2
Sulfur atom count	A	X = S	-7.67793	2	-7.65586	2	-7.89126	2
	B	X = M	-7.72188	1	-7.71391	1	-7.89299	1
Enamine compound class	A	X = S	-7.67793	2	-7.65586	2	-7.89126	2
	B	X = M	-7.72188	1	-7.71391	1	-7.89299	1

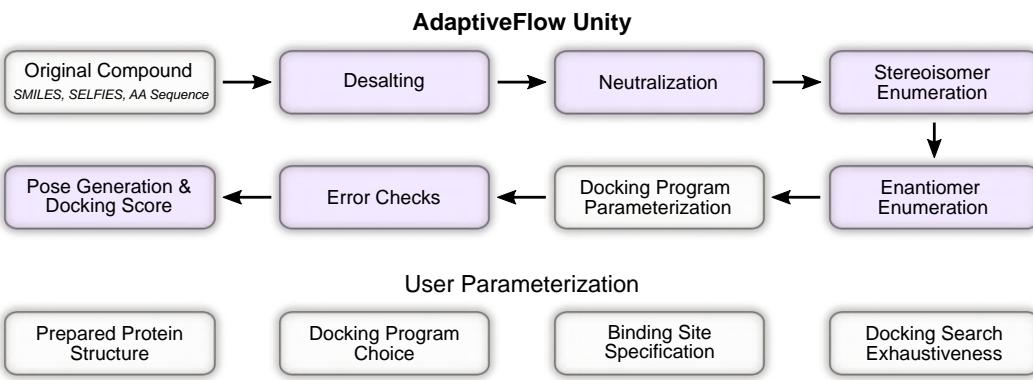
Extended Data Fig. 8. Adaptive Target-Guided Prescreening Benchmarks. Results of several prescreens for a test system within the Adaptive Target-Guided Virtual Screening (ATG-VS) approach, where different settings were used. In the first run (left), one representative per tranche was used with docking exhaustiveness 1. In the second run (middle), 10 representatives per tranche are used with exhaustiveness 1. In the third run (right), one representative per tranche was used with exhaustiveness 10. The average docking scores were computed for each tranche of each property (averaging over all other properties/dimensions). For each property, the tranches have been ranked within their property realm based on the average docking scores for the tranches. As can be seen, the patterns and rankings are very similar to each other, indicating that one representative with exhaustiveness 1 can be sufficient for the prescreenings.



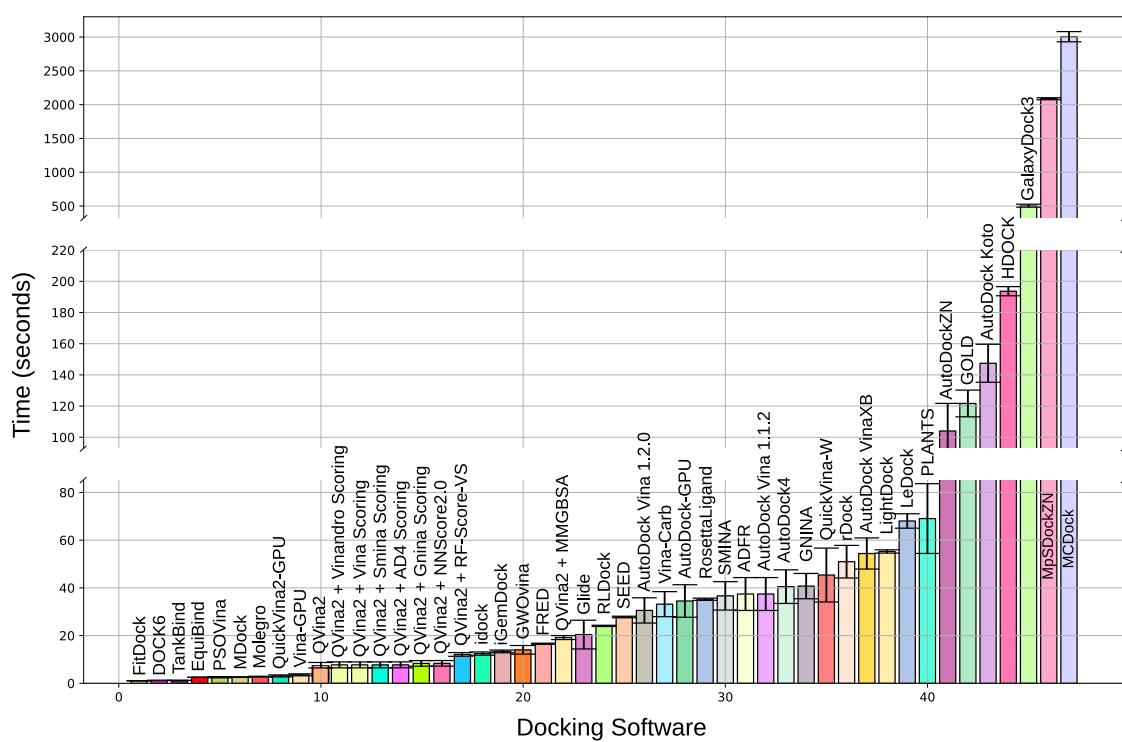
Extended Data Fig. 9. Benchmarking large-scale virtual screening strategies. Violin plots compare the distribution of docking scores for the top 50 molecules identified by two approaches: (1) a standard ultra-large virtual screen of 100 million randomly selected compounds, and (2) the Adaptive Target-Guided (ATG) screening strategy developed in this work. For each ATG run, a prescreen of 11 million representative compounds was used to identify the most promising regions of chemical space, followed by a focused primary screen whose size is indicated for each target. ATG consistently achieves comparable or superior docking scores using a fraction of the computational budget.



Extended Data Fig. 10. Comparison of Adaptive Target-Guided (ATG) Primary Virtual Screens and standard ULVS. 100 million compounds were screened with standard ULVS of randomly selected compounds, as well as with ATGs where the 100 million compounds to be screened were selected based on the ATG prescreen. Violin plots of the docking scores of the top 50 virtual screening hits obtained by standard ULVSs (blue) and ATG-VS (rose). Stripplots of the data points are overlaid. The docking scores of the top 50 compounds of the adaptive screens are considerably better for all target proteins than for the standard ULVSs.



Extended Data Fig. 11. Conceptual Overview of AdaptiveFlow Unity. AF Unity takes as input a user-specified string representation of a molecule and processes the provided molecule by desalting, neutralizing, and generating stereoisomer (with RDKit (63)). Subsequently, based on the docking program choice, all ligands are converted into 3D conformers (with OpenBabel (61)) using a default protonation state of 7.4. Finally, built-in error checks are performed depending on a user's docking program choice, returning the docking pose and scoring function value. Steps in which the parameters of the user are involved are colored in white.



Extended Data Fig. 12. Docking Time Estimates. Timing estimates for running a single calculation averaged over 25 independent runs for all docking programs. The calculations were run using AFU.

Feature		AdaptiveFlow
General Features		
Codebase		Python
Native cloud support		available (AWS)
Linear scaling behavior		demonstrated up to 5.6 million CPUs
AdaptiveFlow Ligand Libraries		
Total size of ligand library		68.7 billion
Additional predicted ligand properties		28
PDB format		available
PDBQT format		available
SDF format		available
MOL2 format		available
SMILES format		available
SELFIES format		available
Parquet format		available
Sparse version (for ATG-VSSs)		available
Interactive web interface		yes (18 dimensions)
Accessibility	AdaptiveFlow website and AWS Open Data	
AdaptiveFlow for Ligand Preparation (AFLP) Module		
Stereoisomer enumeration		supported
External open-source tools for ligand preparation		extended support
SMILES as input format		supported
SELFIES as input format		supported
Amino acid sequences as input format		supported
SELFIES as output format		supported
Prediction of physical molecular properties		supported (27 properties)
Automatic multidimensional matrix generation		extended support
Quality control		extended support
Enhanced speed		extended support
AdaptiveFlow for Virtual Screening (AFVS) Module		
Adaptive target-guided virtual screens		supported
- Machine learning enhancement		supported
Docking methods		1500
Deep learning-based scoring functions		supported
Deep learning-based pose prediction methods		supported
Molecular dynamics-based methods		supported
Protein as target		supported
RNA as target		supported
DNA as target		supported
Covalent docking		extended support
Protein ligands (protein-receptor docking)		supported
Peptide ligands		supported
Parquet output format		supported
Docking quality control		supported
ARM supported		supported
Enhanced x86 performance		available (for several docking methods)
GPU support		supported
AdaptiveFlow Unity (AFU) Module		
Standalone mode		available
API mode		available

Extended Data Table 1. Features Supported by AdaptiveFlow. Features grouped by AdaptiveFlow module: General, AFLP, AFVS, and AF Unity. AdaptiveFlow has significantly more features compared to the original version.

Index	Property	AdaptiveFlow Parameter	Preparation Method	Used for 18-Dimensional Matrix
1	molecular weight	mw_file	Input library	yes
2	logP	logp_jchem	JChem	yes
3	hydrogen bond acceptor count	hba_jchem	JChem	yes
4	hydrogen bond donor count	hbd_jchem	JChem	yes
5	rotatable bond count	rotb_jchem	JChem	yes
6	topological polar surface area	tpsa_jchem	JChem	yes
7	molecular refractivity	mr_jchem	JChem	yes
8	solubility [molar]	logs	JChem	yes
9	number of bonds	bondcount_jchem	JChem	no
10	aromatic proportion	ap	JChem	yes
11	ring count	ringcount	JChem	no
12	aromatic ring count	aromaticringcount	JChem	no
13	logD	logd	JChem	no
14	formal charge	formalcharge	AdaptiveFlow	yes
15	positive charge count	positivechargecount	AdaptiveFlow	yes
16	negative charge count	negativechargecount	AdaptiveFlow	yes
17	fraction of sp3 carbon atoms	fsp3_file	Input library	yes
18	total atom count	atomcount_jchem	JChem	no
19	heavy atom count	heavyatomcount_file	Input library	no
20	chiral center count	chiralcentercount	JChem	yes
21	doublebond stereoisomer count	doublebondstereoisomercount	JChem	yes
22	halogen atom count	halogencount	JChem	yes
23	sulfur atom count	sulfurcount	JChem	yes
24	nitrogen and oxygen atom count	NOcount	JChem	no
25	electronegative atom count	electronegativeatomcount	JChem	no
26	Enamine compound class	enamine_compound_class	Input library	yes
27	quantitative estimate of druglikeness (QED)	qed_rdkit	RDKit	no
28	Murcko scaffold	scaffold_rdkit	RDKit	no

Extended Data Table 2. Calculated Properties of Ligands in the Enamine REAL Space. All properties of the ligands that were determined/calculated during the preparation of the REAL Space (version 2022q1-2) with AFVS. A few properties (molecular weight, fsp3, heavy atom count, Enamine compound class) were available in the input library and thus could be directly used without any calculations.

Property	Available Calculation Methods
Preparation of Ligands into Ready-To-Dock Format	
Desalting	AFLP (internal)
Neutralization	JChem, Open Babel
Stereoisomer Generation	JChem, RDKit
Tautomer Generation	JChem, Open Babel, RDKit
Protonation State Calculation	JChem, Open Babel, Epik, QupKake
3D Coordinate Conversion	JChem, Open Babel, RDKit
Target Format Conversion	Open Babel (for all formats but SELFIES) and SELFIES Python Package (only for SELFIES)
Calculation of Additional Molecular Properties	
Molecular weight	JChem, Open Babel, Input Library
logP	JChem, Open Babel, Input Library
HBA count	JChem, Open Babel, Input Library
HBD count	JChem, Open Babel, Input Library
Rotatable bond count	JChem, Open Babel, Input Library
TPSA	JChem, Open Babel, Input Library
Molecular refractivity	JChem, Open Babel, Input Library
Solubility [molar]	JChem, Open Babel, Input Library
Bond count	JChem, Input Library
Aromatic proportion	JChem, Input Library
Ring count	JChem, Input Library
Aromatic ring count	JChem, Input Library
logD	JChem, Input Library
Formal charge	AFLP (internal)
Positive charge count	AFLP (internal)
Negative charge count	AFLP (internal)
Fraction of sp ³ carbon atoms (fsp ³)	JChem, Input Library
Total atom count	JChem, Open Babel
Heavy atom count	Input Library
Chiral center count	JChem, Input Library
Doublebond stereoisomer count	JChem
Halogen atom count	JChem, Input Library
Sulfur atom count	JChem, Input Library
NO atom count	JChem, Input Library
Electronegative atom count	JChem, Input Library
Enamine compound class	Input Library
Quantitative Estimate of Druglikeness (QED)	RDKit
Murcko scaffold	RDKit

Extended Data Table 3. Supported Processing Steps of AFLP. AFLP can prepare molecules into a ready-to-dock format, and calculate additional properties of each molecule. For many of these calculations, external tools (JChem from ChemAxon, RDKit, Epik from Schrödinger (96), QupKake (97) and Open Babel) are used, while for a few AFLP calculates them internally. Many properties can also be used directly from the input library if available, avoiding the need to calculate them during the runtime.

Docking program	License model	Special features	DL-based	Hardware support	Reference
AutoDock-GPU	Open Source	GPU-accelerated version of AutoDock4	✗	GPU	(57)
AutoDock4	Open Source	Improved modelling of waters and macrocycles	✗	CPU	(98)
AutoDockZN	Open Source	Special support for zinc-coordinating ligands	✗	CPU	(99)
AutoDock Vina 1.2.0	Open Source	Python support; macromolecule support	✗	CPU	(100)
AutoDock Vina 1.1.2	Open Source	Improved optimization, scoring and multithreading	✗	CPU	(101)
AutoDockFR (ADFR)	Open Source	Improved receptor side-chain flexibility modelling	✗	CPU	(102)
AutoDock Koto	Open Source	Improved sampling	✗	CPU	(103)
ATPdock	Open Source	ATP-specific protein-ligand docking	✗	CPU	(104)
CovDock	Licence Required	Covalent docking	✗	CPU	(105)
DOCK 6.0	Free for academics	RNA-ligand, DNA-ligand, protein-protein docking	✗	CPU	(106)
FlexX	Licence Required	Covalent docking, Ensemble docking	✗	CPU	(107)
FitDock	Academic Licence	Incorporation of template fitting for docking	✗	CPU	(108)
GalaxyDock3	Licence Required	Incorporation of side chain flexibility	✗	CPU	(109)
Glide (HTVS)	Paid Academic Licencing	Less exhaustive sampling and scoring	✗	CPU	(110)
Glide (SP)	Paid Academic Licencing	More exhaustive sampling (relative to Glide HTVS) & multiple levels of accuracy	✗	CPU	(110)
Glide (XP)	Paid Academic Licencing	Most accurate sampling and scoring for Glide & different levels of accuracy	✗	CPU	(110)
Gnina	Open Source	Deep learning scoring of complex	✓	GPU/CPU scoring	(68)
GOLD	Licence Required	Genetic algorithm based docking; Covalent docking	✗	CPU	(111)
GWOVina	Open Source	Grey wolf optimization	✗	CPU	(112)
HDOCK	Academic Licence	Protein-RNA/DNA docking	✗	CPU	(113)
HSYMDOCK	Academic Licence	Protein-protein/rna docking with Cn and Dn symmetry	✗	CPU	(114)
idock	Open Source	AutoDock Vina derivative with improved speed	✗	CPU	(115)
iGEMDOCK	Open Source	-	✗	CPU	(116)
LeDock	Freeware for academic use	Excellent Sampling power (comparable to GOLD)	✗	CPU	(117)
LigandFit	Open Source	Identification of active site regions for docking	✗	CPU	(118)
LightDock	Open Source	Protein-protein docking	✗	CPU	(119)
MDock	Licence Required	Ensemble docking for multiple protein conformations	✗	CPU	(120)
MCDock	Open Source	Use of Monte Carlo simulation technique	✗	CPU	(121)
Molegro	Licence Required	-	✗	CPU	(122)
MpsDockzn	Licence Required	Zinc metalloprotein-specific scoring function	✗	CPU	(123)
PSOVina	Open Source	Improved pose sampling and speed	✗	CPU	(124)
PLANTS	Open Source	Use of ant colony optimization	✗	CPU	(125)
PI-LZerD	Open Source	Protein-protein docking	✗	CPU	(126)
PIPER	Open Source	Use of pairwise interaction potentials for protein-protein docking	✗	CPU	(127)
QuickVina 2	Open Source	Extremely fast	✗	CPU	(90)
QuickVina-W	Open Source	Improved speed for blind docking	✗	CPU	(128)
QVina2-GPU	Open Source	GPU-accelerated version of QVina2	✗	GPU	(58)
QVina2-W-GPU	Open Source	GPU-accelerated version of QVina2-W	✗	GPU	(58)
rDock	Open Source	Support for RNA targets	✗	CPU	(129)
Rosetta Ligand	Licence Required	-	✗	CPU	(130)
RLDOCK	Open Source	RNA-ligand docking	✗	CPU	(131)
SEED	Open Source	Fragment-based docking	✗	CPU	(132)
SMINA	Open Source	User specific scoring functions	✗	CPU	(133)
TANKBind	Open Source	Fast neural network based complex prediction	✓	CPU/GPU	(12)
VinaCarb	Open Source	Improved glycosidic modelling	✗	CPU	(134)
VinaXB16	Open Source	Improved halogen bonding scoring function	✗	CPU	(135)
Vina-GPU	Open Source	GPU-accelerated version of Vina	✗	GPU	(58)
Vina-GPU-2.0	Open Source	GPU-accelerated version of Vina-2.0	✗	GPU	(58)

Extended Data Table 4. Supported docking programs. The collection of docking programs supported within AFVS and AF Unity that perform both pose prediction and scoring.

Docking program	License model	Special features	DL-based	Hardware support	Reference
AD4 Scoring	Open Source	Improved over AutoDock 3 forcefield	✗	CPU	(136)
ASP Scoring	Licence required	Atom-atom distance potential	✗	CPU	(111)
Bump Filter Scoring	Free for academics	Consideration of molecule overlap with receptor atoms	✗	CPU	(106)
ChemScore	Licence required	Improved speed compared to GoldScore	✗	CPU	(111)
Contact Scoring	Free for academics	Incorporates the scoring performed with the DISTMAP program	✗	CPU	(106)
Continuous Scoring	Free for academics	No grid calculations	✗	CPU	(106)
ChemPLP	Licence required	Slightly faster than ChemScore and up to four times faster than GoldScore	✗	CPU	(111)
DOCK3.5 Scoring	Free for academics	Incorporation of solvation effects	✗	CPU	(106)
DeepAffinity	Open Source	Use of convolutional neural networks	✓	CPU/GPU	(70)
DeepBindRG	Open Source	Use of recurrent and convolutional neural networks	✓	CPU/GPU	(71)
Footprint Scoring	Free for academics	Calculates intermolecular hydrogen bonds and footprint comparisons, in addition to standard interactions	✗	CPU	(106)
Gnina Scoring	Open Source	Convolutional neural network scoring	✓	CPU/GPU	(68)
GOLD Scoring	Licence required	Use of the chemscore function	✗	CPU	(111)
GoldScore	Licence required	Original scoring function provided with GOLD	✗	CPU	(111)
Grid-Based Scoring	Free for academics	Based on the non-bonded terms of the molecular mechanic force field	✗	CPU	(106)
Hawkins GB/SA Scoring	Free for academics	Use of pairwise GB solvation model	✗	CPU	(137)
NNScore 2.0	Open Source	Use of a fully-connected neural networks	✓	CPU/GPU	(69)
MM/PBSA	Open Source	Calculation of end-state free energies of molecules in solution	✗	CPU	(74)
MM/GBSA	Open source	Improved scoring accuracy	✗	CPU	(74)
MultiGrid FPS Scoring	Free for academics	Pair-wise interaction energies are computed over multiple grids rather than in Cartesian space	✗	CPU	(106)
OpenBPM	Open Source	Improved scoring accuracy	✗	CPU	(72)
PLP Scoring	Open source	Pairwise additive scoring function	✗	CPU	(125)
PLP95 Scoring	Open source	Origin of PLP	✗	CPU	(125)
PB/SA Scoring	Free for academics	Application of the ZAP electrostatics toolkit from OpenEye	✗	CPU	(106)
Pharmacophore Matching Similarity Scoring	Free for academics	Calculates the level of pharmacophore overlap between a reference molecule and a candidate molecule	✗	CPU	(106)
Rosetta Scoring	Open Source	Improved scoring for side-chain flexibility	✗	CPU	(130)
RF-Score-VS	Open Source	Random Forest-based scoring function	✗	CPU	(138)
SASA Scoring	Free for academics	Calculates the percent exposure of a ligand	✗	CPU	(139)
SEED Scoring	Open Source	Continuum dielectric approximation of electrostatic contribution	✗	CPU	(132)
SMINA Scoring	Open Source	Improved AutoDock Vina scoring	✗	CPU	(133)
Vinardo Scoring	Open Source	Improved Vina scoring	✗	CPU	(140)
Vina Scoring	Open Source	Improved scoring compared to AD4	✗	CPU	(101)
Zou GB/SA Scoring	Free for academics	Fast algorithm, pairwise free energy model, for ligand binding affinity calculations	✗	CPU	(139)

Extended Data Table 5. Supported Protein-Ligand Scoring Functions. The collection of docking programs supported within AFVS and AF Unity that perform both pose prediction and scoring.

Docking program	License model	Special features	DL-based	Hardware support	Reference
DiffDock	Open Source	Blind docking pose prediction	✓	GPU/CPU	(66)
EquiBind	Open Source	Blind docking pose prediction	✓	GPU/CPU	(67)
EnzyDock	Open Source	Identifies probable binding modes of multiple ligand states in an enzyme active site	✗	CPU	(141)
FRED	Academic Licence	Improved pose sampling	✗	CPU	(142)

Extended Data Table 6. Supported pose prediction methods. The collection of docking programs supported within AFVS and AF Unity that perform both pose prediction and scoring.

Target Protein	Structure ID	Structure Source	Target Site	Target Site Type
KRAS ^{G12V}	4TQ9	Protein Data Bank (RCSB)	Active site	Enzymatic site
CB1	6KPG	Protein Data Bank (RCSB)	Primary ligand pocket	Orthosteric site
PTPRB	2H04	Protein Data Bank (RCSB)	Active site	Enzymatic site
JNK3	7S1N	Protein Data Bank (RCSB)	Active site	Enzymatic site
TLR8	6ZJZ	Protein Data Bank (RCSB)	Dimerization site	Protein-protein interaction
PIK3CG	2CHZ	Protein Data Bank (RCSB)	Active site	Enzymatic site
NTR1	4GRV	Protein Data Bank (RCSB)	Primary ligand pocket	Orthosteric site
MED15-KIX	2GUT	Protein Data Bank (RCSB)	SREBP interaction site	Protein-protein interaction
KEAP1	5FNQ	Protein Data Bank (RCSB)	NRF2 interaction site	Protein-protein interaction
Raptor	AFQ8N122_F1	AlphaFold Structure Database	TOS motif	Protein-protein interaction

Extended Data Table 7. ATG-VS Benchmark Systems. The target proteins and binding sites used in the benchmark studies comparing ATG-VSs with standard ULVs.