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

Domiziana Cecchini[1,+], AkshatKumar Nigam[2,3,+], Ming Tang[4,5,6,+], Joana Reis[7,+], Matt Koop[8], Andrea Gottinger[1], Callum Robert Nicoll[1], Yao Wang[9], Abhilash Jayaraj[10,7], Süleyman Selim Çınaroğlu[11,12], Ricarda Törner[7,13], Yehor Malets[14], Minko Gehev[15], Krishna M. Padmanabha Das[10,6], Kelly Churion[6], Jongwan Kim[6], Nidhin Thomas[6], Yong Li[16], Hyuk-Soo Seo[7], Sirano Dhe-Paganon[7], Christopher Secker[17,18], Mohammad Haddadnia[19,10,7], Alexander Hasson[20,21], Minkai Li[22], Abhishek Kumar[23,6], Roni Levin-Konigsberg[2], Eun-Bee Choi[24], Geoffrey I. Shapiro[24], Huel Cox 3rd[7], Luke Sebastian[7], Chelsea Braithwaite[7], Puspalata Bashyal[7], Dmytro S. Radchenko[25], Aditya Kumar[26], Lei Yang[16], Pierre-Yves Aquilanti[8,27], Henry Gabb[28], Amr Alhossary[29,30], Gerhard Wagner[10], Alán Aspuru-Guzik[31,32,33,34], Yurii S. Moroz[25,35,36], Charalampos G. Kalodimos[6], Konstantin Fackeldey[17,26], John D. Schuetz[9], Andrea Mattevi[1,*], Haribabu Arthanari[10,7,*], Christoph Gorgulla[6,10,37,*]

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 protocols. 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 Q_1, we validated these hits experimentally and determined the 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 10^{60} (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,

[1]Department of Biology and Biotechnology, University of Pavia, Italy. [2]Klyne, San Francisco, USA. [3]Department of Computer Science, Stanford University, Stanford, USA. [4]The University of Queensland, Frazer Institute, Faculty of Medicine at the Translational Research Institute Australia, Brisbane, Australia. [5]Queensland University of Technology, School of Biomedical Sciences, Centre for Genomics and Personalised Health at the Translational Research Institute Australia, Brisbane, Australia. [6]Department of Structural Biology, St. Jude Children's Research Hospital, Memphis, Tennessee, USA. [7]Department of Cancer Biology, Dana-Farber Cancer Institute, Boston, USA. [8]Amazon Web Services, Seattle, USA. [9]Department of Pharmacy and Pharmaceutical Sciences, St. Jude Children's Research Hospital, Memphis, Tennessee, USA. [10]Department of Biological Chemistry and Molecular Pharmacology, Harvard Medical School, Harvard University, Boston, USA. [11]Department of Biochemistry, University of Oxford, UK. [12]Department of Bioengineering, Ege University, Bornova, 35040, Türkiye. [13]Department of Chemistry, University of Zurich, Switzerland. [14]V. P. Kukhar Institute of Bioorganic Chemistry and Petrochemistry, National Academy of Science of Ukraine, Kyiv, Ukraine. [15]Google, Mountain View, USA. [16]Department of Chemical Biology and Therapeutics, St. Jude Children's Research Hospital, Memphis, Tennessee, USA. [17]Zuse Institute Berlin (ZIB), Berlin, Germany. [18]Max Delbrück Center for Molecular Medicine, Berlin, Germany. [19]Department of Molecular Genetics, University of Toronto, Toronto, Canada. [20]Department of Oncology, University of Oxford, Oxford, UK. [21]Mathematical Institute, University of Oxford, Oxford, UK. [22]Harvard College, Cambridge, USA. [23]Department of Biosciences and Bioengineering, Indian Institute of Technology Guwahati, India. [24]Department of Medical Oncology, Dana-Farber Cancer Institute, Boston, MA, USA. [25]Enamine Ltd, Kyiv, Ukraine. [26]Institute for Mathematics, Technical University Berlin, Berlin, Germany. [27]NVIDIA, Santa Clara, CA, USA. [28]Intel, Santa Clara, USA. [29]Wesleyan University, Middletown, USA. [30]Pulse for Integrated Solutions GmbH, Unterschleißheim, Bavaria, Germany. [31]Department of Computer Science, University of Toronto, Toronto, Ontario, Canada. [32]Vector Institute for Artificial Intelligence, Toronto, Ontario, Canada. [33]Department of Chemistry, University of Toronto, Toronto, Ontario, Canada. [34]Canadian Institute for Advanced Research, Toronto, Ontario, Canada. [35]Chemspace LLC, Kyiv, Ukraine. [36]Taras Shevchenko National University of Kyiv, Kyiv, Ukraine. [37]Department of Physics, Faculty of Arts and Sciences, Harvard University, Cambridge, USA.

[+]These authors contributed equally.

[*]email: christoph.gorgulla@stjude.org, andrea.mattevi@unipv.it, hari@hms.harvard.edu

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

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

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

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

AdaptiveFlow

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

AdaptiveFlow introduces major enhancements across three core components: the AdaptiveFlow Ligand Preparation (AFLP) module for efficient preprocessing of massive libraries; the AdaptiveFlow for Virtual Screening (AFVS) engine, which supports over 1,500 docking protocols; and the AdaptiveFlow Unity (AFU) module, which integrates these components into a unified, modular workflow adaptable to both CPU and GPU-based infrastructures (overview in Fig. 1 and in Extended Data Table 1). These modules have been engineered to support both traditional and emerging use cases, while optimizing for large-scale parallel execution on cloud and on-premise infrastructures. AdaptiveFlow is natively compatible with cloud services such

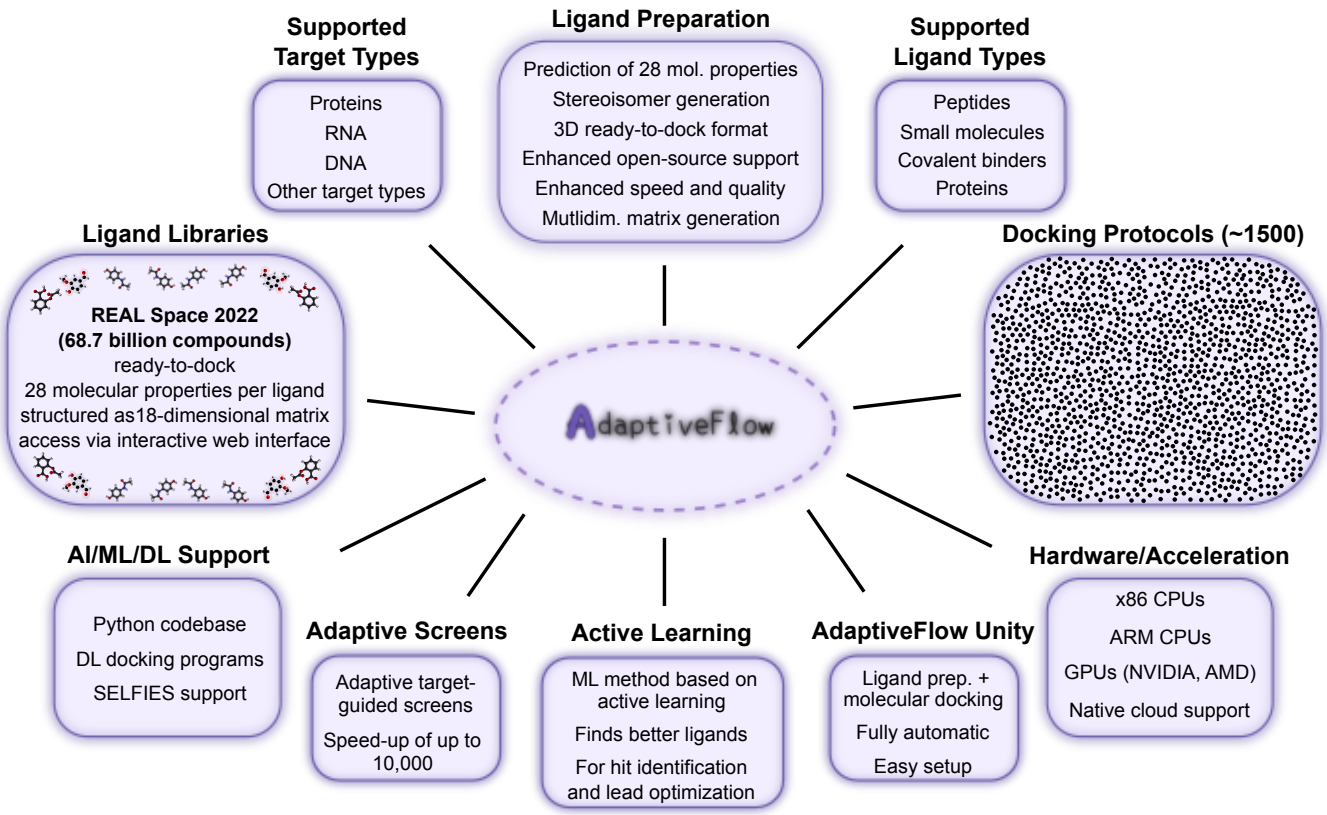

Fig. 1. Overview of features of AdaptiveFlow. AdaptiveFlow introduces several key advances that significantly expand the capabilities of ULVS, including, i) curation of the Enamine REAL Space library, comprising 69 billion molecules in a ready-to-dock format compatible in multiple chemical file formats, ii) integration of an extensive suite of approximately 1,500 docking protocols (defined as specific combinations of sampling algorithms and scoring functions), 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 protocols. 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.

as AWS (Extended Data Fig. 1), where it demonstrates near-linear scaling up to 5.6 million virtual CPUs (Extended Data Fig. 2), enabling efficient exploration of ultra-large libraries. The use of spot (preemptible) instances further reduces the costs.

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

The AFVS module enables large-scale docking campaigns using over 1,500 distinct docking protocols. We define a docking protocol as a specific combination of a sampling algorithm (pose prediction) and a scoring function, allowing for a vast number of unique workflows assembled from more than 40 individual docking programs. AFVS supports a wide range of ligand and target types, including peptides, DNA, and RNA, and includes recent advances in machine learning-based docking such as DiffDock and TANKBind (12, 13), which have shown promise in speed and accuracy, especially for blind or flexible docking scenarios (Extended Data Table 4, 5, 6; Supplementary Table 1, 2; Supplementary Fig. 4).

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

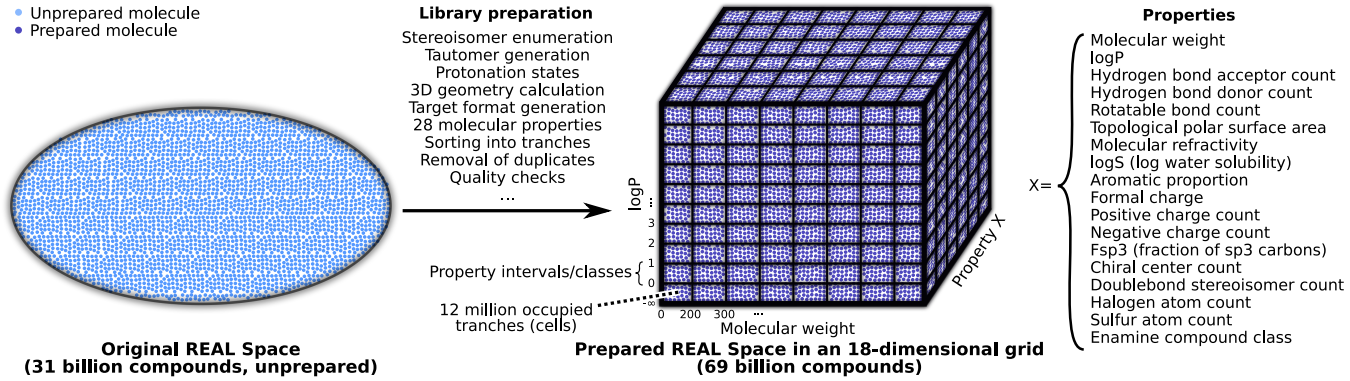

Fig. 2. AdaptiveFlow version of the Enamine REAL Space. The initial version of the REAL Space (enumerated release, 2022q1–2) comprised approximately 31 billion molecules represented as SMILES strings in an unprocessed format, e.g. without full stereoisomer enumeration or structural standardization. A 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-VSs) that are introduced in this work.

Enamine REAL Space

Virtual screening performance scales significantly with library size, because larger screening sizes tend to yield both higher potencies and improved hit rates (i.e., the number of confirmed hits divided by the number of tested compounds) (7, 9–11). 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. 99% of the compounds in the library conform to 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 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

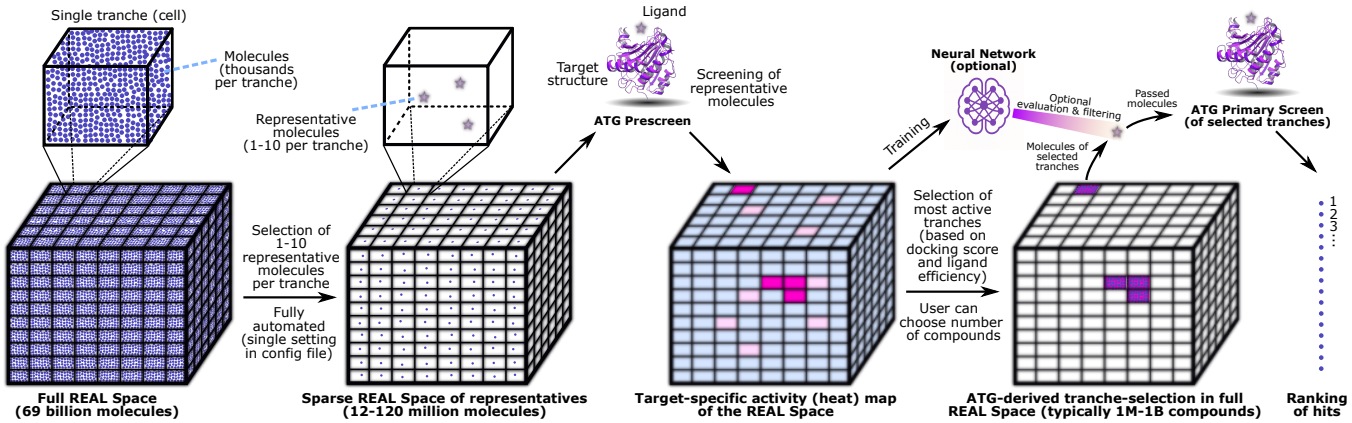

Fig. 3. Conceptual overview of Adaptive Target-Guided Virtual Screens (ATG-VSs). The REAL Space, divided into an 18-dimensional grid, populates over 12 million individual tranches (cells) represented in the figure schematically as a 3-dimensional matrix and each of these tranches has on average 5600 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.

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

Adaptive Ultra-Large Virtual Screenings

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

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

1. **Library Decomposition.** On average, a tranche in the current library has approximately 5600 molecules, and for each tranche up to ten representative ligands are selected for a pilot screen. The number of representatives per tranche (1-10) can be selected by the user. Since there are approximately 12 million tranches (that contain molecules from the REAL Space), if one representative per tranche is selected, the total number of molecules screened in the prescreen would be approximately 12 million for the entire library. If more than one representative molecule, per tranche, is chosen by the user then these additional molecules are automatically selected to be diverse from each other within the given tranche.

2. **Screening Representative Molecules (ATG Prescreen).** Instead of screening the entire library, a prescreen of only the selected representative ligands is performed. In addition, instead of screening representative ligands from all the 12 million tranches, users can apply an optional dynamic filtering procedure that corresponds to the 18-dimensional properties of the ligands. Here, users can manually apply criteria that the molecules/tranches need to satisfy, e.g. based on the properties of the targeted site, which can significantly reduce the number of dockings required in the prescreen. After

the prescreen, the sub-space (tranches) of the ligand library that have the best docking scores among the representative ligands are selected for the next screening stage.

3. **Selection of Active Tranches.** AdaptiveFlow automatically selects the most promising tranches based on docking scores of the ATG Prescreen. The total number of ligands to be screened in the next stage can be specified by the user.

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

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

AdaptiveFlow is a highly flexible platform that supports funnel-type screening strategies, either as a standalone method or in combination with ATG-VS. Users can perform iterative filtering by screening a subset of the library (or the full library, if computationally feasible) to generate an initial ranking, then re-screening the top candidates with increased exhaustiveness or different scoring functions. This allows for an arbitrary number of screening stages to be deployed. Additionally, users can manually select library subsets based on molecular properties or generate smaller custom libraries via the AFLP module to facilitate multi-stage workflows.

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

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

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

(ii) Impact of machine learning (test–based benchmarks): To further evaluate the benefit of incorporating an ML classifier into the ATG workflow, we conducted experiments using a sub-sampled version of the REAL Space comprising 5 million

compounds. We compared three approaches: (1) a random screen of 1 million compounds ("1M Random" in Fig. 4); (2) the ATG-VS method, which involved decomposing the library based on 18 molecular properties, performing a prescreen, and selecting compounds for a primary screen, with a total of 100,000 docking evaluations across both stages ("100K ATG"); and (3) an enhanced ATG-VS approach in which a machine learning model, trained on docking scores and Morgan fingerprints from the prescreen, was used to prioritize compounds for the primary screen ("100K ATG+AL"). It should be noted that in the case of ATG and ATG+AL scenarios, the total number of docking evaluations, including the prescreen, was 100K, compared to a million in the random screen. As shown in Fig. 4,the ML-augmented ATG-VS consistently improved the enrichment of high-affinity virtual hits relative to the standard ATG-VS. Though both versions of the ATG screen sampled only 100K molecules, they outperformed the random screen sampling of 1 million molecules.

While ATG-based screens consistently outperformed random screening, the extent of improvement varied depending on the biophysical properties of the target. Proteins with complex, feature-rich binding pockets, particularly those combining extensive hydrophobic surfaces with well-defined polar subpockets, derived the greatest benefit from ATG-based prioritization and ML enhancement. Taken together, these findings underscore the scalability, efficiency, and adaptability of ATG-VS. By intelligently narrowing the search space and optionally integrating machine learning for compound prioritization, AdaptiveFlow enables high-performance virtual screening at the scale of tens of billions of molecules.

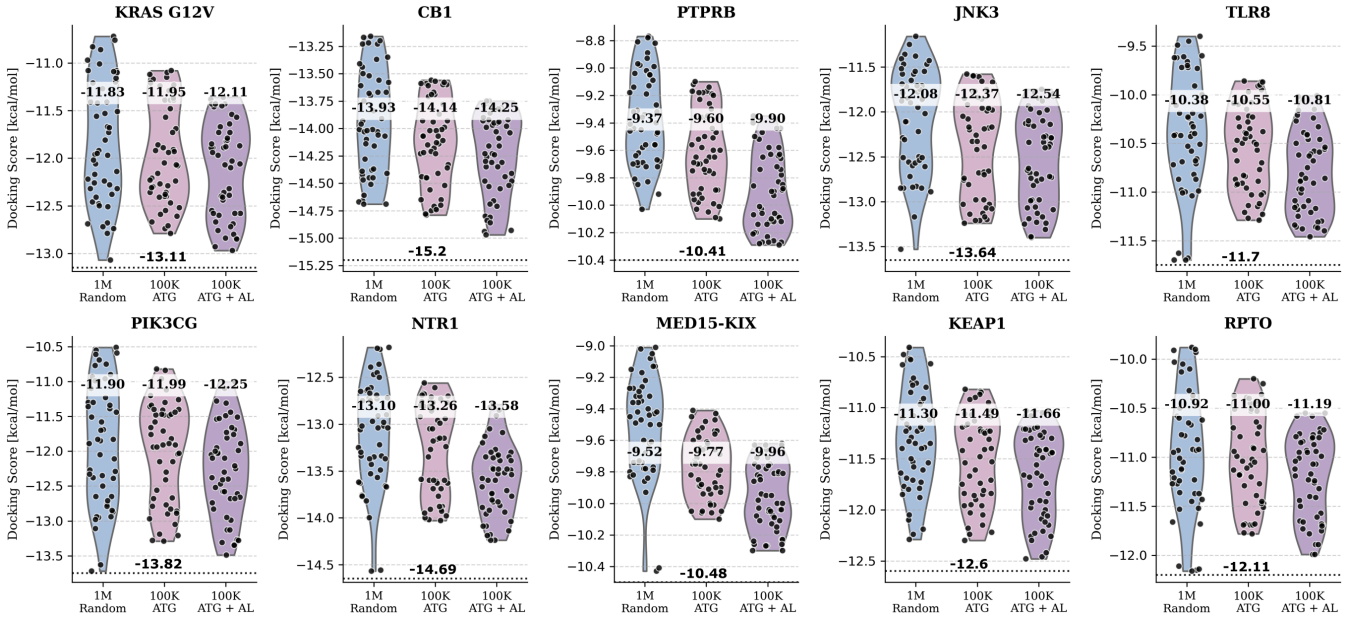

Fig. 4. Benchmarks Involving Adaptive Target-Guided (ATG) Primary Virtual Screens and Standard ULVSs. Violin plots of the docking scores (the more negative, the better) of the top 50 virtual screening hits obtained using three approaches: (i) standard ultra-large virtual screening (ULVS) of 1 million 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. The horizontal dashed lines indicate the best docking score obtained from a brute-force screen of the entire library, serving as a reference for the absolute best possible hit (5-million-compound subset). 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.

To evaluate the sensitivity of our binning strategy to different sampling algorithms and scoring functions, we repeated the benchmarks using Smina (employing the Vinardo scoring function) and GWOVina (incorporating flexible side-chain docking). The physicochemical property-based binning remains effective independent of the docking protocol used. As shown in Supplementary Fig. 11, the ATG-VS approach consistently identified high-scoring compounds with speedups and enrichment factors comparable to those observed with QuickVina 2. Specifically, the average predicted binding affinity of the top 50 molecules improved relative to standard random screens, while the computational costs remained dramatically lower. This demonstrates that the descriptor-based binning correlates well with binding potential across docking protocols of varying sophistication.

Machine Learning and Cloud Support

AdaptiveFlow introduces expanded capabilities designed to support artificial intelligence (AI) and machine learning (ML) workflows across molecular biology and drug discovery. In addition to integrating state-of-the-art ML-based docking programs, including EquiBind (30), DiffDock (13), and TANKBind (12), the platform offers several features aimed at facilitating the development and application of ML models. The AFLP module computes a wide range of physicochemical properties commonly used as features or optimization targets in ML pipelines, including the Quantitative Estimate of Druglikeness (QED) (31) and octanol–water partition coefficient (logP) (32). AFLP also supports the generation of SELFIES (Self-Referencing

Embedded Strings) molecular representations (33, 34), which are widely adopted in generative modeling due to their syntactic robustness and ability to encode valid molecules across arbitrary token sequences (35) (Extended Data Table 2, 3, 3, 4, 5, 6, 7; Supplementary Table 1, 2). Additionally, AFVS introduces numerous deep learning-based docking approaches by modularly assembling scoring functions and pose prediction tools into comprehensive docking workflows. To further streamline ML-driven workflows the AFU module provides a fully integrated pipeline for ligand preparation and docking. This unified interface simplifies data generation for model training and enables seamless incorporation of docking scores as objective functions in generative or optimization-based ML models (36). With support for approximately 1,500 docking protocols accessible through a single command-line interface, AFU eliminates the need for users to individually configure and learn multiple docking programs, thus significantly lowering the barrier to entry for ML and drug discovery communities alike. Together, these features position AdaptiveFlow as a versatile platform not only for applying ML methods to virtual screening but also for developing new ML architectures and pipelines for structure-based drug discovery (Extended Data Fig. 11, 12).

In parallel with its support for machine learning, AdaptiveFlow is optimized for efficient deployment in cloud computing environments. The platform runs natively on the popular workload manager Slurm (37), widely supported across providers like Google Cloud and Microsoft Azure, and AWS Batch (38). Users provide a SLURM submission template (e.g., template1.slurm.sh) in the configuration. AdaptiveFlow populates this template with the necessary variables (job arrays, partitions, memory) at runtime. This decouples the workflow logic from the specific cluster configuration. Users can adapt the submission script to their local requirements (e.g., specific headers, modules, or partition names) without modifying the core Python codebase. Further, the size of computing jobs (both size in cores and length of time) can be configured based on site-specific cluster configurations.

In this study, we demonstrate near-linear scalability on AWS using over 5.6 million virtual CPUs, setting a new global benchmark for parallelization in cloud-based drug discovery. This level of scalability enables the completion of ultra-large virtual screens involving billions of compounds within hours, timelines that would otherwise require weeks or months on traditional on-premise compute clusters. By combining cloud-native performance with ML-native features, AdaptiveFlow serves as a high-performance and accessible foundation for next-generation AI-driven drug discovery.

Experimental validation of AdaptiveFlow

To experimentally validate the capabilities and performance of AdaptiveFlow, we selected two therapeutically significant target proteins, ferroptosis suppressor protein 1 (FSP1) and poly(ADP-ribose) polymerase 1 (PARP1). FSP1 is a recently characterized oxidoreductase that plays a central role in protecting cells from ferroptosis, an iron-dependent form of regulated cell death triggered by membrane lipid peroxidation (39–42). FSP1 maintains pools of reduced coenzyme Q_{10} and vitamin K hydroquinone using NAD(P)H, operating independently of GPX4 to suppress oxidative membrane damage. Inhibiting FSP1 disrupts this defense, sensitizing cells, especially redox-imbalanced cancer cells (43), to ferroptosis. Despite its importance, only a few FSP1 inhibitors have been described and their binding to FSP1 has been structurally characterized (2, 44–46), making it an ideal target for exploring novel inhibitor space using AdaptiveFlow. We also targeted PARP1, a well-established target involved in DNA single-strand break repair, for benchmarking AdaptiveFlow on a mechanistically distinct, clinically validated enzyme class.

Structural Studies of FSP1. To streamline hit discovery and structural studies, we employed ancestral sequence reconstruction, a method known to yield highly stable proteins while preserving functional and structural fidelity to modern homologs (47, 48). Using this approach, we reconstructed the tetrapod ancestor of human FSP1, which shares 76% sequence identity and retains all active-site residues (Supplementary Fig. 12, 13). The ancestral FSP1 variant expressed efficiently in *E. coli* ($\sim$100 mg/L culture), bound FAD, exhibited high thermostability ($T_{\mathrm{m}} = 62$ °C; Supplementary Fig. 13B), and retained full enzymatic activity (Supplementary Table 3; Supplementary Fig. 14 and 15). Crucially, it yielded high-quality, well-diffracting crystals (Supplementary Table 4), enabling detailed structural analyses.

We determined high-resolution crystal structures of FSP1 bound to FAD, to FAD and NAD^+, and in complex with FAD, NAD^+ and coenzyme Q_1 (Fig. 5; Supplementary Fig. 16a; Supplementary Table 4). The overall structure of the ancestral protein is very similar to that of the human enzyme (Supplementary Fig. 16b). NAD^+ triggers a 2 Å shift of the flavin ring and its binding mode is fully consistent with a hydride transfer mechanism for flavin reduction as the nicotinamide C4 atom is located right in front of the isoalloxazine N5 (Fig. 5B). Coenzyme Q_1 binds between two aromatic side chains of Phe354 and Y290, with its quinone moiety at the edge of the pyrimidine ring of the flavin in a canyon leading from the protein surface (Fig. 5C-D). Specifically, quinone oxygen engages the N(3)H group of the flavin through a hydrogen bond. This geometry and the proximity between the two molecules can enable the facile electron transfer from the flavin to coenzyme Q. Comparison between the FSP1-FAD-NAD^+ and FSP1-FAD-NAD^+-coenzyme Q_1 complexes shows that NAD^+ binding is unaffected by the presence of coenzyme Q. Thus, it is demonstrated that coenzyme Q and $NAD(P)^+$ can simultaneously bind to the protein and this structural attribute perfectly aligns with the ternary-complex mechanism evidenced by the kinetics data (Supplementary Fig. 15). The structure also suggests that there is ample space for the oxygen to access the flavin N5 locus where its reaction with the flavin is known to take place. This feature explains why FSP1 can also operate as a NAD(P)H oxidase (49). This

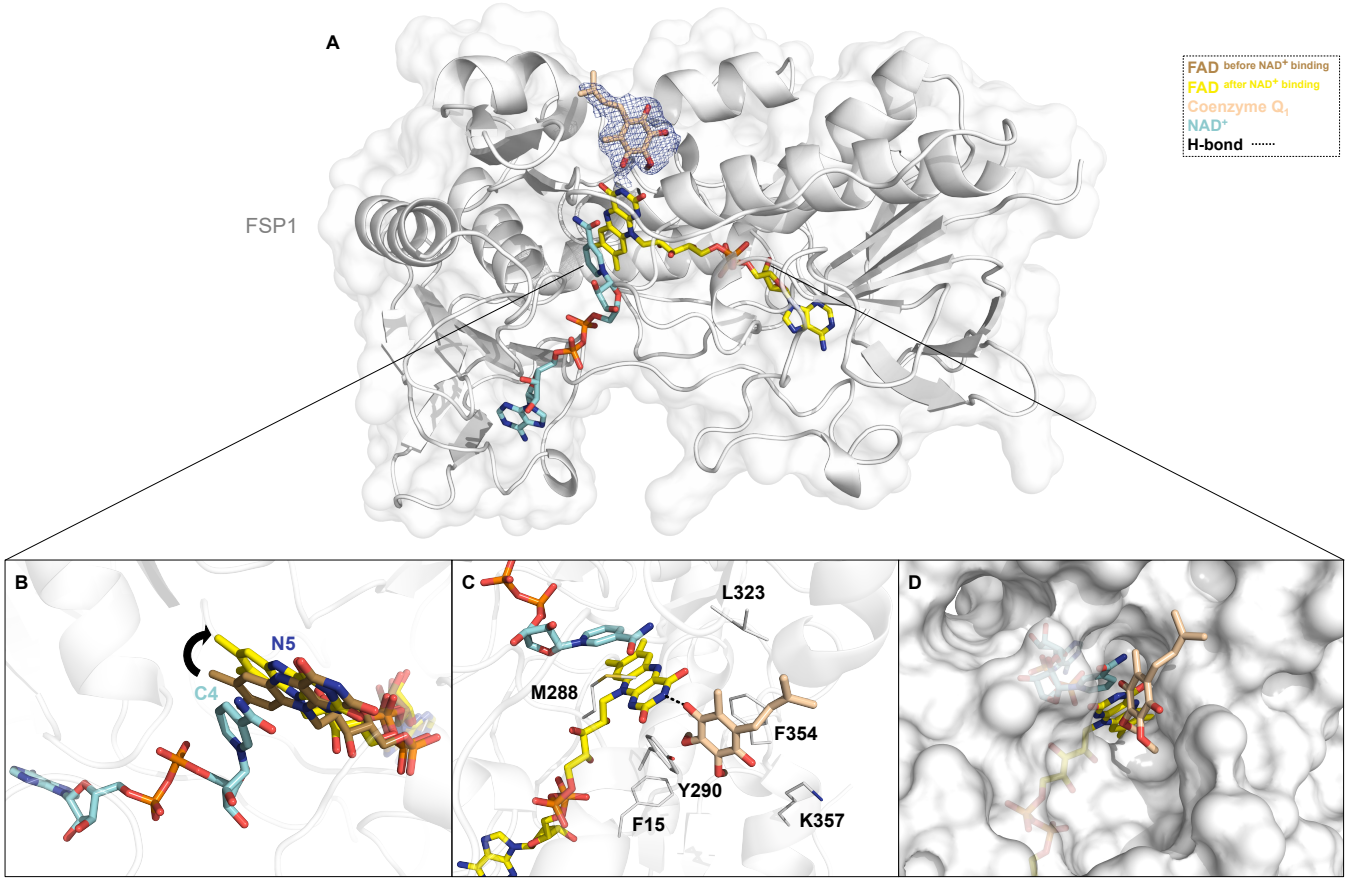

Fig. 5. Co-Crystal Structures of FSP1 in Complex with its Cofactors and Substrate. We determined three crystal structures of FSP1: (i) bound to FAD alone, (ii) bound to FAD and NAD^+, and (iii) bound to FAD, NAD^+, and coenzyme Q_1. (A) Overall structure of the fully assembled FSP1 complex with FAD (yellow carbons), NAD^+ (cyan), and coenzyme Q_1 (beige). The polder omit map for coenzyme Q_1 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 (50, 51) (see also Supplementary Fig. 16A). (C) In the ternary complex, coenzyme Q_1 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 Q_1 occupy distinct but converging channels that meet at the FAD active site, illustrating the architecture of the substrate-binding cavity.

detailed experimental visualization of the NAD(P)H and quinone binding modes and their interactions provided the foundation for inhibitor discovery through in silico screenings targeting the substrate binding sites.

Targeting FSP1 with AdaptiveFlow. To assess the effectiveness of AdaptiveFlow, we screened FSP1 inhibitors targeting its coenzyme Q site using AdaptiveFlow. Our high-resolution crystal structure of FSP1 in complex with FAD and NAD^+ (PDB code: 9IFT) was used as receptor. Maestro from Schrödinger (52) was applied to prepare the receptor's structure for screening, by removing solvent molecules, assigning missing hydrogen bonds, correcting protonation states at pH 7.4, and assigning correct bond orders. AutoDockTools (v.1.5.7) (53) was utilized to convert the prepared structure from the PDB format to the PDBQT format. The coenzyme Q site of FSP1 (with FAD and NAD^+ bound to the corresponding sites) was targeted. In the ATG Prescreen, 1 representative compound was used per tranche (resulting in 12 million dockings), and based on these results 10 million compounds were chosen for the primary screen using the ATG-VS method (without ML feature). Basic filters were applied to the top hit candidates, such as logP less than 5, LogS greater than -5, Topological Polar Surface Area (TPSA) smaller than 140 $Å^2$, number of stereo centers not greater than 3, and compounds with problematic functional groups and predicted toxicity (mutagenic compounds, compounds with reproductive effects, tumorigenic compounds) were excluded. In total, 42 compounds from 37 different clusters with zero alerts suggested by SwissADME (54) were ordered from Enamine. 33 compounds were synthesized and experimentally tested.

Experimental Validation: Inhibitors extend from the coenzyme Q site to the NAD site. We assessed the potential binding and inhibitory activities of the synthesized 33 virtual hits from AdaptiveFlow, using thermal shift and NADH-depletion assays. Two of the hits elicited stabilizing shifts of at least 1.5 °C in the protein melting temperature and inhibited FSP1 oxidoreductase activity by at least 50% at a 10 µM concentration. These two compounds were found to potently inhibit coenzyme Q-reduction activity (Table 1; Supplementary Fig. 17). Their K_i values are comparable to, or better than, those reported for previously described FSP1 inhibitors (0.283 µM for afi-FSP1-1 and 0.777 µM for afi-FSP1-2; Supplementary Table 5; Supplementary Fig. 18, 19) (2, 44, 45). Critically, a cellular thermal shift assay (CETSA; (55)) performed on FSP1-expressing

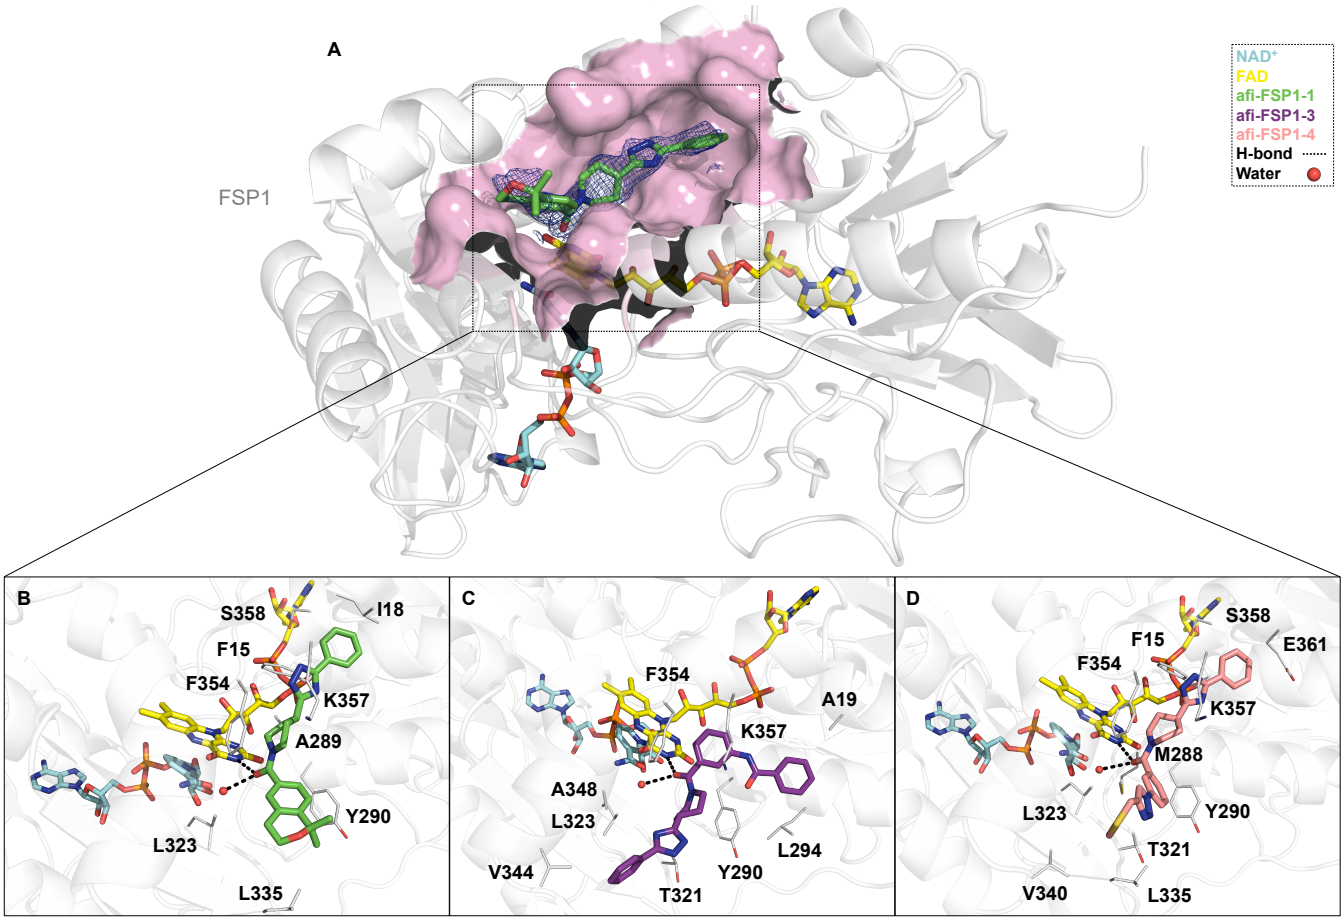

Fig. 6. Co-Crystal Structures of Inhibitors binding to FSP1. (A) Overall structure of FSP1 in complex with inhibitor afi-FSP1-1. The polder omit map for the inhibitor is contoured at the 3.0 σ level. (B-D) Close-up views of the binding site with afi-FSP1-1 (B, green), afi-FSP1-3 (C, violet), and afi-FSP1-4 (D, pink) bound in the region normally occupied by coenzyme Q_1, adjacent to FAD (yellow) and NAD^+ (cyan). All inhibitors displace coenzyme Q_1 and interact directly with the catalytic cavity.

HEK293 cells shows that afi-FSP1-1 and afi-FSP1-2 effectively engage human FSP1 in cells, as evidenced by a thermal shift of 2.8 °C and 3.1 °C, respectively (Supplementary Fig. 20A-J). The activity of afi-FSP1-1 was tested against a T-cell acute lymphoblastic leukemia cell line (Jurkat) that expresses FSP1 (Supplementary Fig. 20M) and is known to undergo ferroptosis (56). Cells proved to be sensitive to the combination of afi-FSP1-1 with 0.5 μM erastin, a ferroptosis inducer. The activity of afi-FSP1-1 matched that shown by the widely used iFSP1(2). When used as a single agent, afi-FSP1-1 showed no cell toxicity at concentrations up to 10 μM. The pro-ferroptotic effect induced by 25 µM inhibitor was completely rescued by liproxstatin-1, a ferroptosis inhibitor. (Supplementary Fig. 20N-P). We determined the crystal structure of the most potent compound (afi-FSP1-1) bound to FSP1-FAD-NAD^+ which reveals that the inhibitor occupies the coenzyme Q binding site in a mode consistent with the predicted docking pose (Fig. 6A-B; Supplementary Fig. 16a). 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. 21). 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 Q_1. Similar features characterize the binding compound afi-FSP1-2 whose bulkier structure enables even more extended interaction with the protein (Supplementary Fig. 21, 22). 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. 23). 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 (Supplementary Fig. 17C-F; Supplementary Fig. 20). Notably, these compounds share an aromatic-amide-saturated heterocycle core as displayed also by afi-FSP1-1 and 2

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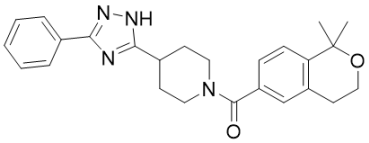	$0.283^b \pm 0.05$	0.098
afi-FSP1-2	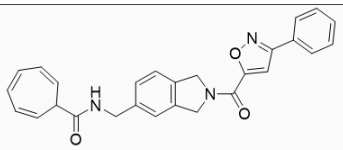	0.777 ± 0.22	0.266
afi-FSP1-3	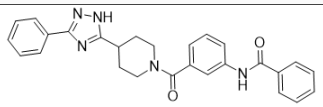	0.520 ± 0.11	0.184
afi-FSP1-4	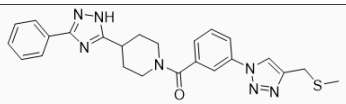	0.284 ± 0.07	0.198
afi-FSP1-5	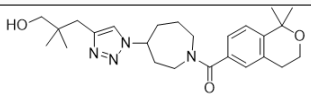	1.00 ± 0.32	n.d.
afi-FSP1-6	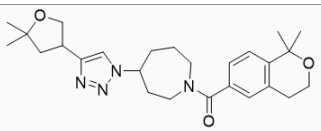	3.94 ± 0.71	n.d.

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 Q_1, and 0.1–0.2 μM FSP1 (Supplementary Fig 17). K_d values were derived from independent binding experiments based on surface plasmon resonance (SPR, Supplementary Fig. 19). 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. Supplementary Table 5 compares the K_i values of the afi-FSP1 compounds with those of known FSP1 inhibitors whereas the IC_{50} values measured using the indirect resazurin assay, commonly used in previous literature (2, 44, 45), are listed in supplementary Table 9.

(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 bond 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. 21, 22). 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, U-shaped conformation laying on Tyr290. These binding modes are validated by mutagenesis experiments (Supplementary Table 6, 7; Supplementary Fig. 21, 22) and correlate with the measured K_i values. Compounds afi-FSP1-1 and afi-FSP1-4 exhibit practically identical K_is, whereas compound afi-FSP1-3 displays a twofold decrease in potency. Nonetheless, Absorption, Distribution, Metabolism, and Excretion (ADME) assays indicated that afi-FSP-3 is the most promising candidate for further medicinal chemistry optimization (see Supplementary Information Section H for details). Besides being the best inhibitors, compounds afi-FSP1-1, afi-FSP1-2, and afi-FSP1-3 exhibit a 10-12-fold reduced efficacy against the NADH oxidase activity of FSP1, while compound afi-FSP1-4 exhibited no efficacy at all (Supplementary Fig. 25; Supplementary Table 8). The weaker inhibition reflects the retention of NAD(P)H binding in the inhibited protein as shown by crystal structures and indicates that reactivity with oxygen is only mildly or not affected by inhibitor binding in the coenzyme Q site. We posit that preferentially inhibiting the quinone-reducing activity while sparing the reactivity with oxygen, might be pharmacologically beneficial. Inhibitors of this type may convert the enzyme from being anti-ferroptotic (coenzyme Q reducing activity) to being pro-ferroptotic (uninhibited lipid-damaging ROS production). Future work is needed to explore and validate this concept in FSP1 pharmacology.

Applying AdaptiveFlow to PARP1. To further validate the effectiveness of the AdaptiveFlow platform, we selected human Poly(ADP-ribose) polymerase 1 (PARP1) as a benchmark target due to its well-characterized structure, critical role in DNA repair, and relevance in cancer therapy. Following the ATG-prescreen and a primary screen of 100 million molecules, 160 promising candidates were synthesized for further experimental verification, as detailed in the Methods section. A two-concentration high-throughput enzymatic screen identified seven PARP1 inhibitors, four of which (iParp1, iParp2, iParp3, and iParp4) showed sub-250 nM IC50 values and were selected for further characterization. Protein NMR confirmed direct binding to the active site of PARP1, with chemical shift perturbations supporting the predicted docking poses. X-ray crystallography

revealed the 2.05 Å structure of the PARP1 catalytic domain bound to a hydrolyzed form of iParp1, validating its binding mode and key interactions. Clonogenic assays in BRCA1-deficient TNBC cells demonstrated selective cytotoxicity for iParp1 and iParp2, though reduced cellular activity, especially for iParp1, was attributed to poor membrane permeability due to hydrolysis. Remarkably, AdaptiveFlow identified PARP1 inhibitors with submicromolar to low nanomolar potency, including iParp1, which exhibited an IC_{50} of 8.8 nM, comparable to the FDA-approved drug Olaparib. This result is particularly notable given that our hits were derived from a purely virtual screen without any prior knowledge of known inhibitors or medicinal chemistry optimization of the hit. Structural similarity analyses confirmed that several hits shared key scaffolds with known PARP1 inhibitors, while others showed different scaffold variations. Details can be found in the Supplementary Information Section I. 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. 27 , 29, 30, 31, 33, 32; Supplementary Table 13.

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, 57, 58), 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 (59, 60). 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 target central to DNA damage repair. Our ultra-large screen yielded seven inhibitors, including iParp1, which displayed an IC_{50} comparable to that of the FDA-approved drug olaparib (61). These results underscore the power of AdaptiveFlow in rapidly identifying potent hits across diverse protein classes and demonstrate the advantage of adaptive screening over brute-force enumeration. Notably, while our previous identification of nanomolar KEAP1–NRF2 inhibitors required an unbiased screen of 1.3 billion molecules, the ATG-guided screen for FSP1 achieved comparable success with just 22 million docking instances. Although the effectiveness of chemical space focusing is target-dependent, our benchmarking and experimental studies highlight the potential of ATG-VS to tailor the search space to specific binding pockets, significantly enhancing screening efficiency. As chemical space continues to expand, focusing on target-specific subspaces will no longer be optional, but essential.

With the infrastructure now in place, adding molecules to the library and directing them to their respective tranches is a computational process that scales linearly. As new tranches, beyond the 12 million currently populated are filled, only representatives from these newly populated tranches need to be incorporated into the prescreen set. The open-source nature of AdaptiveFlow enables seamless adaptation, extension, and development of the platform by both academic and industry users. With the increasing availability of high-resolution structural data, enabled by techniques such as cryo-EM and de novo structure prediction, combined with site-specific functional insights from gene editing technologies, in silico screening is poised to deliver chemical tools capable of modulating precise biomolecular functions. As modalities like molecular glues and targeted protein degraders mature, the ability to identify small molecules that bind specific protein surfaces with functional consequences is becoming increasingly valuable. AdaptiveFlow's modular architecture supports over 1,500 docking protocols, encompassing combinations of scoring functions, sampling algorithms, and ML-based pose predictors such as DiffDock and TANKBind. The platform also natively supports SELFIES representations, physicochemical property calculations, and active learning–driven workflows, making it fully compatible with generative models and reinforcement learning pipelines (62). As generative models continue to advance, platforms like AdaptiveFlow will be essential for rapidly evaluating designed molecules against biologically relevant structural targets. In conclusion, AdaptiveFlow offers a unified, open-source, and massively scalable solution for ultra-large virtual screening. It provides seamless integration of conventional and AI-driven docking protocols, access to the largest ready-to-dock chemical library, and adaptive strategies for efficient navigation of trillion-scale chemical space.

With demonstrated success across diverse therapeutic targets, AdaptiveFlow stands as a next-generation infrastructure poised to accelerate the design and discovery of future therapeutics.

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, Milene 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(7784):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 Castañón, 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 Balius, 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 McCorvy, Grant C Glatfelter, 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 σ2 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, Chinzorig 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) carbonate 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, Dmitriy 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 chlorooxoacetate 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 Viniychuk, 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 Hurmach, 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.
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.
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.
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.
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.
27. Anna Tomberg and Jonas Boström. 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.
28. Frank Lovering, Jack Bikker, and Christine Humblet. Escape from flatland: increasing saturation as an approach to improving clinical success. Journal of medicinal chemistry, 52(21):6752–6756, 2009.
29. Frank Lovering. Escape from flatland 2: complexity and promiscuity. MedChemComm, 4(3):515–519, 2013.
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.
31. G. Richard J. Bickerton, Gaia Valeria Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L. Hopkins. Quantifying the chemical beauty of drugs. Nature chemistry, 4 2:90–8, 2012.
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.
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.
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.
35. Daniel Flam-Shepherd, Kevin Zhu, and Alán Aspuru-Guzik. Language models can learn complex molecular distributions. Nature Communications, 13(1):1–10, 2022.
36. AkshatKumar Nigam, Robert Pollice, Gary Tom, Kjell Jorner, John Willes, 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.
37. SLURM. https://slurm.schedmd.com, Accessed: 2018-12-22.
38. AWS Batch. https://aws.amazon.com/batch/, Accessed: 2022-12-31.
39. Ashley R Brown, Tal Hirschhorn, and Brent R Stockwell. Ferroptosis—disease perils and therapeutic promise. Science, 386(6724):848–849, 2024.
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.
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.
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.
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.
44. Joseph M Hendricks, Cody E Doubravsky, Eddie Wehri, Zhipeng Li, Melissa A Roberts, Kirandeep K Deol, Mike Lange, Irene Lasheras-Otero, Jeremiah D Momper, 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.

45. Toshitaka Nakamura, Eikan Mishima, Naoya Yamada, André Santos Dias Mourão, Dietrich Trümbach, 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.
46. Sitao Zhang, Amalia H. Megarioti, Joseph M. Hendricks, Junshu Zhou, Qingxiang Sun, Da Jia, and James A. Olzmann. Cocrystal structure reveals the mechanism of fsp1 inhibition by fsen1. Proceedings of the National Academy of Sciences, 122(22):e2505197122, 2025. doi: 10.1073/pnas.2505197122.
47. 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.
48. 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.
49. 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.
50. Sitao 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.
51. 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.
52. Schrödinger, LLC. Maestro, 2024. Schrödinger Release 2024-1.
53. GM Morris, R Huey, and AJ Olson. Using autodock for ligand-receptor docking. curren t protocols in bioinformatics. Unitas, 8(14):5, 2008.
54. Antoine Daina, Olivier Michielin, 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.
55. 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.
56. Mariluz Soula, Ross A Weber, Omkar Zilka, Hanan Alwaseem, Konnor La, Frederick Yen, Henrik Molina, Javier Garcia-Bermudez, Derek A Pratt, and Kıvanç Birsoy. Metabolic determinants of cancer cell sensitivity to canonical ferroptosis inducers. Nature Chemical Biology, 16(12):1351–1360, 2020.
57. 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.
58. AkshatKumar Nigam, Matthew FD Hurley, Fengling Li, Eva Konkoľová, Martin Klíma, Jana Trylčová, Robert Pollice, Süleyman Selim Çinaroğlu, 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.
59. 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.
60. 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.
61. Giorgio Valabrega, Giulia Scotto, 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.
62. Benjamin Sanchez-Lengeling and Alán Aspuru-Guzik. Inverse molecular design using machine learning: Generative models for matter engineering. Science, 361(6400):360–365, 2018.
63. 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.
64. 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.
65. Greg Landrum et al. Rdkit: Open-source cheminformatics, 2006.
66. 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.
67. 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.
68. 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.
69. 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.
70. Andrew T 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.
71. 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.
72. 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.
73. 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.
74. 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.
75. 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.
76. 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.
77. 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 ranking of protein–ligand binding poses. Journal of Chemical Information and Modeling, 2022.
78. 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 metabolon unveils the molecular determinants of coenzyme q biosynthesis. Nature catalysis, 7(2):148–160, 2024.
79. 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.
80. Kazutaka Katoh, John Rozewicki, and Kazunori D Yamada. Mafft online service: multiple sequence alignment, interactive sequence choice and visualization. Briefings in bioinformatics, 20(4): 1160–1166, 2019.
81. Alexandros Stamatakis. Raxml version 8: a tool for phylogenetic analysis and post-analysis of large phylogenies. Bioinformatics, 30(9):1312–1313, 2014.
82. Fréderic Lemoine, J-B Domelevo Entfellner, Eduan Wilkinson, Damien Correia, M Dávila Felipe, Tulio De Oliveira, and Olivier Gascuel. Renewing felsenstein's phylogenetic bootstrap in the era of big data. Nature, 556(7702):452–456, 2018.
83. Maria Laura Mascotti. Resurrecting enzymes by ancestral sequence reconstruction. In Enzyme Engineering: Methods and Protocols, pages 111–136. Springer, 2021.
84. Ziheng Yang. Paml 4: phylogenetic analysis by maximum likelihood. Molecular biology and evolution, 24(8):1586–1591, 2007.
85. Geeta N Eick, Jamie T Bridgham, Douglas P Anderson, Michael J Harms, and Joseph W Thornton. Robustness of reconstructed ancestral protein functions to statistical uncertainty. Molecular biology and evolution, 34(2):247–261, 2016.
86. Xavier Robert and Patrice Gouet. Deciphering key features in protein structures with the new endscript server. Nucleic acids research, 42(W1):W320–W324, 2014.
87. 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 binding. The FEBS journal, 276(10):2833–2840, 2009.
88. Wolfgang Kabsch. xds. Acta Crystallographica Section D: Biological Crystallography, 66(2):125–132, 2010.
89. Jon Agirre, Mihaela Atanasova, Haroldas Bagdonas, Charles B Ballard, Arnaud Baslé, James Beilsten-Edmands, Rafael J Borges, David G Brown, J Javier Burgos-Mármol, John M Berrisford, et al. The ccp4 suite: integrative software for macromolecular crystallography. Biological Crystallography, 79(6):449–461, 2023.
90. Ana Casañal, Bernhard Lohkamp, and Paul Emsley. Current developments in coot for macromolecular model building of electron cryo-microscopy and crystallographic data. Protein Science, 29 (4):1055–1064, 2020.
91. 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 chimerax: Structure visualization for researchers, educators, and developers. Protein science, 30(1):70–82, 2021.
92. 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, 2015.
93. 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 and modeling, 55(2):460–473, 2015.
94. BioSolveIT GmbH. infinisee version 6.2.0, 2024.
95. BioSolveIT GmbH. Seesar, 2024.
96. 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, G. Giambaşu, T. Giese, M.K. Gilson, H. Gohlke, 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, 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,

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. York, C. Zhao, Q. Zhu, and P.A. Kollman. Amber 2024, 2024.

97. 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. Gohlke, S. Izadi, K. Kasavajhala, M.C. Kaymak, E. King, T. Kurtzman, 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, 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. J. Chem. Inf. Model., 63:6183–6191, 2023. doi: 10.1021/acs.jcim.3c01153.
98. Ryne C. Johnston, Kun Yao, Zachary Kaplan, Monica Chelliah, Karl Leswing, Sean Seekins, Shawn Watts, David Calkins, Jackson Chief Elk, Steven V. Jerome, Matthew P. Repasky, and John C. 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: 37023332.
99. 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): 6946–6956, 2024. doi: 10.1021/acs.jctc.4c00328. PMID: 38832803.
100. 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 receptor flexibility. Journal of computational chemistry, 30(16):2785–2791, 2009.
101. 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 information and modeling, 54(8):2371–2379, 2014.
102. 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 Information and Modeling, 61(8):3891–3898, 2021.
103. 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 chemistry, 31(2):455–461, 2010.
104. 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. PLoS computational biology, 11(12):e1004586, 2015.
105. Junkai Ji, Jin Zhou, Zhangfan Yang, Qiuzhen Lin, Jianqiang Li, and Carlos A Coello Coello. Autodock koto: A gradient boosting differential evolution for molecular docking. IEEE Transactions on Evolutionary Computation, 2022.
106. 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.
107. 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 scoring. Journal of chemical information and modeling, 54(7):1932–1940, 2014.
108. William J Allen, Trent E Balius, 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 current docking performance. Journal of computational chemistry, 36(15):1132–1156, 2015.
109. Bernd Kramer, Matthias Rarey, and Thomas Lengauer. Evaluation of the flexx incremental construction algorithm for protein–ligand docking. Proteins: Structure, Function, and Bioinformatics, 37 (2):228–241, 1999.
110. 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.
111. Woong-Hee Shin and Chaok Seok. Galaxydock: protein–ligand docking with flexible protein side-chains. Journal of chemical information and modeling, 52(12):3225–3232, 2012.
112. 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. enrichment factors in database screening. Journal of medicinal chemistry, 47(7):1750–1759, 2004.
113. 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, 52(4):609–623, 2003.
114. 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.
115. Yumeng Yan, Huanyu Tao, Jiahua He, and Sheng-You Huang. The hdock server for integrated protein–protein docking. Nature protocols, 15(5):1829–1852, 2020.
116. Yumeng 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 (W1):W423–W431, 2018.
117. 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 Bioinformatics and Computational Biology (CIBCB), pages 77–84. IEEE, 2012.
118. 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. BMC bioinformatics, 12(1):1–11, 2011.
119. 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.
120. 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 of Molecular Graphics and Modelling, 21(4):289–307, 2003.
121. 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 docking. Bioinformatics, 34(1):49–55, 2018.
122. 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.
123. Ming Liu and Shaomeng Wang. Mcdock: a monte carlo simulation approach to the molecular docking problem. Journal of computer-aided molecular design, 13(5):435–451, 1999.
124. Gabriela Bitencourt-Ferreira and Walter Filgueira de Azevedo. Molegro virtual docker for docking. In Docking screens for drug discovery, pages 149–167. Springer, 2019.
125. 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 sampling and iteration optimization strategy. Journal of Chemical Information and Modeling, 55(4):833–847, 2015.
126. 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.
127. 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.
128. Bin Li and Daisuke Kihara. Protein docking prediction using predicted protein-protein interface. BMC bioinformatics, 13(1):1–17, 2012.
129. 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.
130. 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.
131. Sergio Ruiz-Carmona, Daniel Alvarez-Garcia, Nicolas Foloppe, A Beatriz Garmendia-Doval, Szilveszter Juhos, Peter Schmidtke, 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.
132. Gordon Lemmon and Jens Meiler. Rosetta ligand docking with flexible xml protocols. In Computational Drug Discovery and Design, pages 143–155. Springer, 2012.
133. Li-Zhen Sun, Yangwei Jiang, Yuanzhe Zhou, and Shi-Jie Chen. Rldock: a new method for predicting rna–ligand interactions. Journal of chemical theory and computation, 16(11):7173–7183, 2020.
134. 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.
135. 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.
136. 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.
137. Mathew R Koebel, Grant Schmadeke, Richard G Posner, and Suman Sirimulla. Autodock vinaxb: implementation of xbsf, new empirical halogen bond scoring function, into autodock vina. Journal of cheminformatics, 8(1):1–8, 2016.
138. 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.
139. 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.
140. Maciej Wójcikowski, Pedro J Ballester, and Pawel Siedlecki. Performance of machine-learning scoring functions in structure-based virtual screening. Scientific Reports, 7(1):1–10, 2017.
141. Hao-Yang Liu, Irwin D Kuntz, and Xiaoqin Zou. Pairwise gb/sa scoring function for structure-based drug design. The Journal of Physical Chemistry B, 108(17):5453–5462, 2004.
142. 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.
143. Susanta Das, Mor Shimshi, Keren Raz, Neta Nitoker Eliaz, Anil Ranu Mhashal, Tamar Ansbacher, 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.
144. Mark McGann. Fred pose prediction and virtual screening accuracy. Journal of chemical information and modeling, 51(3):578–596, 2011.
145. 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.

146. Hannah Farmer, Nuala McCabe, Christopher J Lord, Andrew NJ 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.
147. 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.
148. 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.
149. Thomas Helleday. The underlying mechanism for the parp and brca synthetic lethality: clearing up the misunderstandings. Molecular oncology, 5(4):387–393, 2011.
150. Junko Murai, Shar-yin N 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.
151. 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):5330–5357, 2019.
152. Fons Elstrodt, Antoinette Hollestelle, Jord HA Nagel, Michael Gorin, Marijke Wasielewski, Ans Van Den Ouweland, 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.
153. 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.
154. 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.
155. Schrödinger LLC New York NY. Schrödinger release 2024-1: Maestro. https://newsite.schrodinger.com/platform/products/maestro/, 2024.
156. 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.
157. 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.
158. John S Delaney. Esol: estimating aqueous solubility directly from molecular structure. Journal of chemical information and computer sciences, 44(3):1000–1005, 2004.
159. 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.
160. Wenxiu 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.
161. 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.
162. 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.
163. 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.
164. 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.
165. 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.
166. 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 Laue. The ccpn data model for nmr spectroscopy: development of a software pipeline. Proteins: structure, function, and bioinformatics, 59(4):687–696, 2005.
167. 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.
168. Paul D Adams, Pavel V Afonine, Gábor Bunkóczi, 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.
169. Paul Emsley and Kevin Cowtan. Coot: model-building tools for molecular graphics. Acta crystallographica section D: biological crystallography, 60(12):2126–2132, 2004.

Methods

AdaptiveFlow

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

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

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

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

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

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

- Scheduling and Orchestration: AWS Batch is a cloud-native scheduler and orchestrator for containerized jobs. As part of running a job, AdaptiveFlow can make use of AWS Batch to automatically scale cloud infrastructure up and down as needed and run all of the individual subjobs. AdaptiveFlow can also automatically generate the Docker container used for processing, eliminating the need to understand containerization or other infrastructure.

- Object Storage (Amazon S3): In addition to the usage of standard cluster filesystems, AdaptiveFlow also supports the usage of object storage for both ingesting of input datasets and the output of the resulting data from the processing.

- Spot Capacity Awareness and Spot Resilience: Spot capacity is unused computing capacity that cloud providers make available at significantly discounted prices compared to normal on-demand access. In exchange for a lower price, a workload running on a spot instance may be reclaimed (stopped) with a short warning. AdaptiveFlow is designed to take advantage of this capacity model by using shorter running subjobs (approximately 30 minutes) and automatically re-running jobs that have prematurely exited through integration with AWS Batch. Additionally, through integration with AWS Batch, AdaptiveFlow is able to select instances that are less likely to be interrupted through the "spot capacity aware" scheduling method.

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

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

Enamine REAL Space

The REAL Space (enumerated, version 2022q1-2) was prepared using AFVS. JChem Suite was used for neutralizing the molecules, stereoisomer enumeration, tautomer generation, protonation state prediction, 3D coordinate calculation, and the calculation of multiple molecular properties, JChem Suite 18.20.0, ChemAxon (https://www.chemaxon.com). Open Babel (version 2.3.2) was used as backup program for protonation state prediction, 3D coordinate calculation (63), 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 overview of the processing steps and preparation methods used is shown in Supplementary Table 2.

Multiple Formats. AdaptiveFlow provides the REAL Space library in ready-to-dock 3D formats including PDB, PDBQT, MOL2, and SDF. These formats are compatible with a wide range of docking programs for small molecules. The total size of the ready-to-dock library is approximately 50 terabytes in compressed form for each of these formats, totaling around 200 terabytes in compressed form (uncompressed, the library is around 2 petabytes). In addition to these ready-to-dock formats, the REAL Space library is also provided in the SELFIES (33), SMILES (64), and Parquet formats. These formats are useful for postprocessing the results of virtual screening (e.g., for easy representation of each molecule), and they can also be used as standalone libraries for other purposes such as in the AI/ML community, where SMILES and SELFIES are the standard molecular input file formats. The Parquet format of the REAL Space only stores the properties of the molecules, which can be useful for fast library-wide queries.

AdaptiveFlow for Ligand Preparation (AFLP)

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

Preparation into Ready-To-Dock Format. The primary task of AFLP is to prepare ligands into a ready-to-dock format. To accomplish this, AFLP desalts the original ligands (in case they are salts), neutralizes them, generates stereoisomers and tautomers, predicts the protonation states, computes the minimum-potential energy conformation of the ligand (3D coordinates), optionally validates the generated 3D conformers, and converts them into the desired target formats. In total over 100 output formats are supported. A detailed overview of the processing steps can be seen in Supplementary Fig. 7. The support input formats are SMILES, SELFIES, as well as amino acid sequences. AFLP uses external tools such as ChemAxon's JChem, Open Babel, and RDKit (65), to carry out the preparation steps. For many of the preparation steps, the user can choose which external tool should be used. Details on which tools are available for each processing step can be found in Supplementary Table 3.

Molecular Properties Calculation. AFLP enables researchers to calculate a variety of physicochemical properties and molecular descriptors for the molecules in the input ligand library, in addition to its primary task of preparing them for molecular docking. By calculating properties such as QED (Quantitative Estimate of Druglikeness), logP, molecular weight, and others (a full list of supported properties can be found in Supplementary Table 2), researchers can more effectively select and prioritize ligands, which can help reduce costs and the time required for screens using AFVS. AFLP uses mostly external tools to calculate the molecular properties, such as ChemAxon's JChem (https://chemaxon.com/calculators-and-predictors, Open Babel (63), and RDKit (for details see Supplementary Table 3). By combining the results of AFLP with those of AFVS, researchers can gain valuable insights into the structural and chemical characteristics of their ligands relative to specific target proteins, which can help them design and optimize new molecules with improved affinity and specificity.

Automatic Tranche Assignments. In AdaptiveFlow, ligand libraries are organized into tranches. Each tranche corresponds to a property of the ligands and is divided into multiple intervals or categories. For example, the molecular weight of the ligands 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, 400-500, and 500-∞ daltons. Multiple tranches allow for the partitioning of a ligand library into a multidimensional grid or table, with one dimension corresponding to each tranche type (ligand property).

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

To characterize how the Enamine REAL Space library populates the 18-dimensional property grid used by AdaptiveFlow, we quantified tranche occupancy (the number of molecules assigned to each grid cell) and summarized it as a histogram (Supplementary Fig. 34). We counted molecules per tranche across the full library. The resulting occupancy counts span

several orders of magnitude, so we report the distribution using logarithmically spaced occupancy ranges (1, 2–10, 11–100, 101–1K, 1K–10K, 10K–100K, 100K–1M, 1M–10M, and >10M molecules per tranche). The histogram displays the number of tranches in each occupancy range (y-axis), highlighting that only a small fraction of the theoretical grid cells are populated and that populated tranches vary widely in size (from single-molecule tranches to highly populated tranches exceeding 10 million molecules), with an average of 5600 molecules per populated tranche.

Support of SELFIES. SELF-referencing Embedded Strings (SELFIES) (33, 34) has become one of the most popular line notation formats for small molecules within machine-learning communities. The primary reason is that SELFIES are fully robust, meaning that every SELFIES string corresponds to a valid molecule, which is not the case for SMILES. Thus generative ML models cannot create invalid molecules when using SELFIES. AFLP now supports SELFIES as an input format for the molecules in addition to SMILES. Additionally, the entire REAL Space is free to be downloaded in the SELFIES format, enabling the creation of powerful datasets.

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

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

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

Enhanced Quality Checks. AFLP checks the validity of each ligand that it prepares in the ready-to-dock format by calculating the potential energy of the ligand in 3D format using Open Babel. This allows AFLP to automatically remove ligands that have an excessively high potential energy, which can be an indication that the compound was damaged or corrupted during the preparation process. Additionally, AFLP supports optional plausibility checks on generated 3D conformers using PoseBusters (66). This helps to ensure the integrity and accuracy of the prepared ligand library.

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

AdaptiveFlow for Virtual Screening (AFVS)

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

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

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

To evaluate the effectiveness of ATG-VS, we performed two types of benchmark studies: (1) full-scale production benchmarks using the entire Enamine REAL Space, and (2) smaller-scale "test set-based" benchmarks designed for rapid prototyping and detailed comparison across different ATG configurations. In the test set-based benchmarks, we randomly selected synthetic reactions from the Enamine space to create a compound subset totaling approximately 5 million molecules. For each of the 10 protein targets tested (listed in Supplementary Table 7), a different random subset was used to ensure unbiased evaluation across targets and screening scenarios.

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

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

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

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

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

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

AdaptiveFlow Unity

AdaptiveFlow Unity (AFU) is the all-in-one version of AdaptiveFlow, combining both ligand preparation and virtual screening into a single streamlined workflow. It is the third module besides AFLP and AFVS. This new module adds many conveniences for workflows where both ligand preparation and virtual screenings have to be carried out in tandem. Notably, AFU is designed to be run as a package, supporting approximately 1500 docking protocols that can be run independently of job schedulers (see Supplementary Data Table 1). A conceptual overview of the workflow of AFU can be found in Supplementary Fig. 11, and a detailed UML workflow diagram is shown in Supplementary Fig. 5. Specifically, as input, a user specifies the molecule to be docked (as a SMILES, SELFIES, or amino-acid sequence), the protein file, the docking site, and the docking program to be run. Based on the user’s choices, error checks in the form of missing docking-choice-specific files are performed, and calculations are run. As a result of a successful calculation, the docking score and pose are returned. AFU can be run in standalone mode,

where it is run via the command line and specifies the options via a configuration file (an example configuration file is shown in Supplementary Data Fig. M). This mode is for example useful for scientists in the drug discovery communities who want to carry out molecular dockings in the most convenient and simple way, without requiring computational expertise, while having a large number of docking protocols available. In addition, AFU can be run in API mode to directly interface with other programs and code, which can for instance be useful for the machine learning communities that develop new methods for drug discovery.

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

Joint Features of AFVS and AFU

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

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

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

Enhanced Performance. The QuickVina 2 molecular docking tool is used to accelerate AutoDock Vina with no loss of accuracy (67). QuickVina 2 execution was profiled in an effort to improve performance even further. To establish baseline performance, the open-source GitHub repository was cloned and the qvina_1buffer branch was compiled using the GNU C++ compiler (version 9.4), the Boost C++ libraries (version 1.79, available from `https://www.boost.org`), and the default parameters (see `https://qvina.github.io/compilingQvina2.html` for the build instructions). A single docking experiment of Avanafil binding to unsolvated 5FNQ was run through the Intel VTune Profiler to look for performance bottlenecks. Unfortunately, the execution profile was flat. The most time-consuming function accounted for less than 20% of the total run-time. Tuning the source code would give rapidly diminishing returns at the risk of affecting accuracy, so compiler-level optimization was the best option to improve performance. QuickVina 2 was rebuilt using the Intel oneAPI DPC++/C++ Compiler (version 2022.2.0.20220730) and aggressive optimization: automatic vectorization, interprocedural optimization, and relaxed floating-point constraints. Profile-guided optimization was also tried, but it did not improve performance. The final executable was 18% faster than the default builds with no loss of accuracy. The new binaries are provided on the QuickVina and AF homepages.

GPU Support. Several of the docking programs support graphics processing units (GPUs), which can strongly reduce computation time and cost. For pose prediction, these include softwares: AutoDock-GPU (59), QVina2-GPU (60), QVina2-W-GPU (60), Vina-GPU (60). Additionally, deep-learning methods such as TANKBind (12), DiffDock (68), EquiBind (69) use GPU-compatible neural networks for direct pose prediction of several ligands in parallel. Additionally, Gnina (70), NNScore 2.0 (71), DeepAffinity (72) and DeepBindRG (73) use GPU-compatible deep learning models for scoring protein-ligand complexes. AdaptiveFlow supports GPU compatibility of all these software. The relative speedups of these methods for single calculations can be compared within Fig. 12.

Molecular Dynamics-Based Methods. Molecular docking is a computationally fast way to predict binding poses and binding-free energies, but it has sometimes trouble distinguishing between compounds with similar binding affinities. Molecular docking algorithms often prioritize speed over accuracy by using fast pose generation and approximate scoring functions to save time, especially when screening extensive collections of compounds against a protein target. This trade-off reduces the effectiveness of the algorithms. More precise algorithms based on molecular dynamics simulations are sometimes used downstream of docking in drug discovery processes to improve the number of accurate predictions. The combination of molecular docking and molecular dynamics can yield results near experimentally determined crystal structures (74, 75). AFVS includes methods for selecting and scoring poses based on molecular dynamics simulations, such as MM/PBSA and MM/GBSA, for

evaluating docking poses and predicting binding affinities (76). MM/PBSA and MM/GBSA are popular for binding free energy prediction because they are more accurate than molecular docking and less computationally intensive than alchemical methods. AdaptiveFlow also supports Binding Pose Metadynamics (BPMD), which uses a stability score based on the resistance of the ligand to perturbation away from the initial pose to rerank predicted binding poses from docking (74, 77). The OpenBPMD algorithm can accurately predict binding poses within 2 Å RMSD of the crystallographic complex structure for many systems and efficiently rerank docked poses. The MM/PBSA, MM/GBSA, and BPMD methods can be used within AdaptiveFlow as rescoring and pose refinement methods.

Enhanced Quality Checks. The docking pose of each ligand screened by AFU is checked for corruption by evaluating the coordinates and computing the potential energy of the docked compound with Open Babel. Any corrupted compounds are automatically removed. AFVS supports optional plausibility checks on generated docking poses of Vina-based docking protocols using PoseBusters, an open-source tool, which checks for chemical as well as intramolecular ligand and intermolecular validity of ligand-protein complexes (66).

Experimental Validation Studies Involving FSP1

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

Chemicals. All chemicals and proteins employed for binding and enzymatic assays were purchased from Sigma Aldrich, except for 2,3,5-trimethylcyclohexa-2,5-diene-1,4-dione (α-tocopherolquinone hydrophilic head) purchased from Synthonix and menaquinone (Vitamin K2) from Supelco Analytical. FSP1 inhibitor candidates were purchased from Enamine; iFSP1 from Medchemexpress; viFSP1 and FSEN1 from Cayman Chemical. Inhibitors were analyzed by NMR (supplementary data). Samples were prepared by dissolving the compounds in deuterated DMSO (d6-DMSO) to 1 mM final concentration. Spectra were acquired with 100 scans on a Bruker Avance III 400 MHz spectrometer and processed using TopSpin version 4.5.0.

Phylogenetic Inference and Ancestral Sequence Reconstruction. Phylogenetic analysis and ancestral sequence reconstruction were carried out following the methods described in (78). Briefly, human FSP1 sequence (UniprotID: Q9BRQ8) was employed as query for homology searches with BLASTP. Datasets were constructed vetting all chordate classes according to TimeTree (79) by mining at least two species with fully sequenced genomes. Multiple sequence alignments were constructed in MAFFT v7 (80) and manually trimmed for single sequence insertions/terminal extensions. Neighbour-joining guide trees were constructed in MEGA v10.2 to assess the quality of datasets under construction. Once the working multiple-sequence alignments were obtained, best-fit substitution models and gamma distribution values (α) were calculated using ProtTest (data for each protein phylogeny are shown in Supplementary Fig. 12B). Maximum likelihood phylogenies were inferred using RaxML v8.2.10 (HPC-PTHREADS module (81)), using rapid bootstrap analysis and searching for the best-scoring ML tree, with 500 bootstraps replicates and the given best-fit model under gamma distribution. When required, a species tree was used to constrain the phylogeny, this was constructed using TimeTree (timetree.org). Once the phylogenies were inferred, the bootstrap values were subjected to transfer bootstrap expectation values using BOOSTER online (82). Figtree v1.4.2 was employed for analysing and visualizing the trees. Ancestral sequence reconstruction was performed employing PAML v4.9a (CODEML module) as marginal reconstruction, using the phylogenies obtained previously, empirical amino acid substitution model (model=2) and JTT substitution matrix, four gamma categories and re-estimation of gamma shape parameter (83, 84). The distribution of the posterior probabilities for each of the ancestral states was analysed at the node corresponding to the tetrapod ancestor in the phylogeny. Sites that displayed posterior probabilities <0.8 were considered ambiguously reconstructed when alternative states displayed posterior probabilities >0.2 (85). The length of the targeted nodes was treated by Fitch's parsimony. To visualize the degree of sequence conservation across tetrapod organisms, multiple sequence alignments were generated using ESPript 3 (86).

Construct design and expression of FSP1. The gene encoding for human full-length FSP1-twinSTREP-eGFP was cloned into pcDNA3 (Genscript). HEK293-EBNA1-6E cells were used for the transient transfection. Cells were maintained in suspension with FreeStyle medium (Invitrogen). The day before transfection, cells were seeded at a concentration of 0.5×10^5 cells/mL. The day after (1×10^6 cells/mL), cells were transiently transfected using polyethyleneimine with a ratio of 1:3 DNA:polyethyleneimine. After 48 h, cells were harvested by centrifugation at 1,200 g for 5 min at 20 °C. Trypan Blue exclusion assay (Trypan Blue solution w/v, Sigma) was used to evaluate cell viability. The genes encoding for His_6-SUMO-ancestral FSP1, His_6-SUMO-ancestral L323A FSP1, His_6-SUMO-ancestral F15A FSP1, His_6-SUMO-ancestral F354A FSP1, lacking the Gly-2 myristoylation site, were cloned in a pET-24a(+) vector with kanamycin resistance (Genscript). Plasmids were then transformed by heat shock into *E. coli* BL21(DE3) cells (25 s at 42 °C). Cells from a single colony were picked and pre-inoculated into 100 mL LB broth containing 100 µg mL^{-1} kanamycin and grown overnight at 37 °C. 10 mL of pre-cultures were inoculated into 1 liter of Terrific Broth and grown at 37 °C, 200 r.p.m. for 3 h until the optical density (OD_{600}) reached 0.5-0.7. Protein expression was induced with isopropyl β-D-1-thiogalactopyranoside (0.1 mM final) and incubated overnight at 30 °C while shaking 200 r.p.m.. Cells were harvested by centrifugation (5,000 g, 15 min, 4 °C), flash frozen in liquid nitrogen, and stored at - 20 °C.

Protein purification. Cells (ca. 29 g) were re-suspended 1:5 (grams:milliters) in Buffer A (50 mM HEPES pH 7.2, 250 mM NaCl, 10% (v/v) glycerol) supplemented with 1 mM FAD and protease inhibitors (1 mM phenylmethylsulphonyl fluoride, 10 µM leupeptin, 10 µM pepstatin) and 5 µg DNAse I /g cells. Cell resuspensions were stirred at 4 °C for 25 min before cell lysis, which was performed with a high-pressure homogeniser (Emulsiflex c-3, ATA Scientific) in three cycles. Cell lysates 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). The resulting supernatant was filtered through a 0.45 µm filter and loaded on a gravity column containing Nickel Sepharose HP IMAC resin (Cytiva) pre-equilibrated with Buffer A. The crude extract was passed onto the column twice. The resin was then washed with at least five column volumes of Buffer A. The wash up of the resin was performed adding 5 column volumes of Buffer A supplemented with increasing concentrations of imidazole, 5 to 30 mM. The protein was then eluted with at least ten column volumes of Buffer B (Buffer A supplemented with 300 mM imidazole). The purified protein sample was concentrated using an Amicon Ultra (Merck) with a 30 kDa cut-off up to 2.5 mL. The imidazole was then removed with a PD-10 batch desalting column (Cytiva) pre-equilibrated with Buffer A. To cleave the His_6-SUMO tag, the desalted sample was concentrated down to 800 µL and then added of a His_6-tagged SUMO protease (1.2 mg mL^{-1}) to a volume ratio of 100:1 and incubated overnight at 4 °C on a rotating wheel. The sample was then loaded onto an Äkta Pure system (Cytiva) equipped with a Ni HiTrap-HP 5 mL column (Cytiva) pre-equilibrated in Buffer A and proteins were eluted out using an imidazole concentration gradient of 2% Buffer B corresponding to 6 mM imidazole. The elution of the tag-less protein was monitored with a multiwavelength detector following the absorbance at 280 nm and 450 nm. The sample was loaded onto an Äkta Pure system (Cytiva) equipped with a Superdex 200 10/300 GL (Cytiva) pre-equilibrated in 100 mM NaCl, 50 mM HEPES pH 7.2. The purified protein was then concentrated down to about 1 mM, measuring the concentration using a NanoDrop ND-100 UV/Vis spectrophotometer (Thermo Scientific) based on the absorption at 280 nm and at 458 nm (ε280 nm predicted by ProtParam tool - Expasy). Purity of the sample was evaluated by SDS-PAGE. The purified protein was finally flash frozen in liquid nitrogen and stored at -80 °C.

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

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

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

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

Nano Differential Scanning Fluorimetry. Thermo-stability analyses were carried out using a TychoTMNT.6 system (NanoTemper GmbH). The protein was diluted to 1 mg mL^{-1} in Buffer A and pre-incubated with ligands at room temperature for 15 minutes before experiments. Melting curves were obtained following the intrinsic fluorescence of tryptophan and tyrosine residues (emission 350 nm and 330 nm, respectively) applying a temperature gradient from 35 °C to 95 °C for a time range of three minutes. Data were analysed as F_{350}/F_{330} ratio and derived to determine the inflection temperature (T_m). All measurements were performed in duplicate. The following condition were tested: FSP1 in presence of 1 mM FAD, NAD^+ or 1 mM CoQ_1 or 1 mM α-tocopherolquinone hydrophilic head or 1 mM vitamin K3 or 10 µM inhibitors or 5% DMSO v/v as a control.

ThermoFAD. Experiments were performed with 20 µM FSP1 after 10 min incubation with inhibitors (10 µM) or 5% v/v DMSO in HEPES buffer (50 mM, pH 7.2) (87). The temperature gradient was set to 30–90 °C with fluorescence detection at a constant rate 0.1 °C/s at 450 ± 30 nm excitation and 535 ± 30 nm emission (BioRad MiniOpticon Real-Time PCR System). Data were normalized and fitted with sigmoidal non-linear regression using GraphPad Prism 9.

Cellular thermal shift assays. Inhibitor afi-FSP1-1: HEK293-EBNA1-6E cells were seeded at a density of 3.0×10^7 cells/mL and incubated with either 5% v/v DMSO or 10 µM of 1 for 1 h at 37 °C in 2 ml of culture medium. Cells were collected by centrifugation at 1,200 × g for 5 min at room temperature, washed twice with PBS, and resuspended at the initial cell concentration. Aliquots of 50 µL were distributed into PCR strip tubes. Samples were subjected to a heat treatment at different temperatures (ranging from 38.7 to 86.2 °C) using a Bio-Rad T100 Thermal Cycler. Following heat exposure, cells were lysed by adding 1% Triton X-100, protease inhibitors (see Protein purification section), and DNase I (5 µg per g of cells). Lysis was performed for 1 h at 4 °C on a rotation wheel. Lysates were centrifuged at 20,000 × g for 20 min at 4 °C, and 15 µL of each supernatant was loaded onto Bio-Rad Mini-PROTEAN TGX SDS-PAGE gels. Proteins were transferred to 0.2 µm PVDF membranes (Bio-Rad) using the Trans-Blot Turbo system. Membranes were exposed with an anti-AIFM2 antibody (clone 1I15 ZooMAb®, rabbit monoclonal, Sigma-Aldrich), followed by an HRP-conjugated anti-rabbit secondary antibody (Millipore). Chemiluminescent signals were detected with the Bio-Rad ECL detection kit on a ChemiDoc MP imaging system. For data analysis, band intensities at each temperature were quantified and normalized by setting the strongest and weakest signals to 100% and 0%, respectively. Melting curves were fitted using a Boltzmann sigmoidal model in GraphPad Prism 9. Aggregation temperatures were calculated from two independent biological replicates, and results are reported as single replicates.

Inhibitors afi-FSP1-2, 3, 4: 1.5×10^7 HEK293-EBNA1-6E cells/mL expressing human FSP1-twinSTREP-eGFP were treated with 1 µM compound 2, or 10 µM compound 3, or 10 µM compound 4 or 5% v/v DMSO following the above-described protocol. The heat treatment ranged from 31.6 to 78.1 °C. PVDF membranes containing proteins were incubated with Precision Protein™ Strep-Tactin® HRP-conjugate (BioRad). The detection and data analysis were described before.

Immunoblot analysis. Jurkat T-lymphoid cell line was obtained from ATCC and cultured in RPMI 1640 with 10% FBS, 100 units/mL penicillin and streptomycin. The cells were lysed using sonication in M-PER buffer (cat:78501 from Thermo Scientific) and 50 µg of total protein per lane was loaded on the gel. FSP1 antibody is a Rabbit-derived cat. # HY-P80679 from MCE; GPX4 antibody is mouse-derived cat. # 67763-1-1g from Protein tech; GAPDH antibody (6C5) is mouse-derived, cat. # AM4300 from Invitrogen

Cell viability assays. Trypan blue dye exclusion was used to determine cell viability. Logarithmic phase cells were seeded at 1×10^5 cells/mL in a 12-well plate. Cells were then treated with either DMSO (control) or the test compounds solubilized in DMSO, gently mixed and incubated for 24 hours under culture conditions as described. Viability was assessed by removing an aliquot (10 µL) of the cell suspension and mixing with an equal volume of 0.4% Trypan Blue solution. Live cells were then enumerated by light microscopy. Viability was expressed as a percentage of the DMSO control number.

Statistical Analysis and Reproducibility. Significance levels are indicated in legends of each figure as follows: *P< 0.05, **P< 0.01, ***P<0 .001, ****P< 0.0001 (Student's t test). Tests used to derive significance are indicated in legends. Individual data points are indicated in graphs. In the latter case all replicates are independent and not multiple measurements taken from a single sample. Analysis of significance was done using unpaired Student's t-test, two-sided.

Crystallisation and structural determination of FSP1. Crystallisation was performed using 13 g/mL in 50 mM HEPES pH 7.2, 1 mM FAD, 1 mM NAD^+ with or without 1 mM coenzyme Q_1. Small crystals or crystalline precipitate appeared in a few days under several conditions of the Mopheus screening kit (Molecular Dimensions) using the vapour diffusion sitting drop 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, 0.06 M $MgCl_2$, 0.06 M $CaCl_2$ at 20 °C with a ratio of 1:1 in a sitting drop. After harvesting, they were smashed to generate small crystals, and used for microseeding. Specifically, the drops were composed of 0.15 µL protein, 0.05 µL seeds, and 0.2 µL reservoir from the following commercial screens.

- FSP1-FAD structure; 0.06 M $MgCl_2$, 0.06 M $CaCl_2$, 0.1 M HEPES/MOPS (acid) pH 7.5; 20% v/v PEG500*MME, 10% PEG20000.
- FSP1-FAD-NAD^+ structure; 0.06 M $MgCl_2$, 0.06 $CaCl_2$, 0.1 M HEPES/MOPS (acid) pH 7.5, 12.5% PEG1000, 12.5% PEG3350, 12.5% MPD.
- FSP1-FAD-NAD^+-coenzyme Q_1 structure; 0.06 M $MgCl_2$, 0.06 $CaCl_2$, 0.1 M imidazole/MES monohydrate (acid) pH 6.5, 20.0% PEG500* MME, 10% PEG20000. Crystals were then soaked with 20 mM coenzyme Q_1 added to the crystallisation drop.
- FSP1-FAD-NAD^+-afi-FSP1-1; 0.06 M $MgCl_2$, $CaCl_2$, 0.1M imidazole/MES monohydrate (acid) pH 6.5, 12.5% MPD, 12.5% PEG1000, 12.5% PEG3350. Crystals were then soaked with 1 mM afi-FSP1-1 added to the crystallisation drop.
- FSP1-FAD-NAD^+-afi-FSP1-3; 0.12 M diethylene glycol, 0.12 M triethylene glycol, 0.12 tetraethylene glycol, 0.12 pentaethylene glycol, 0.1 M tris (base)/bicine pH 8.5, 12.5% MPD, 12.5% PEG1000, 12.5% PEG3350. Crystals were then soaked with 10 mM afi-FSP1-3 added to the crystallisation drop.

- FSP1-FAD-NAD$^+$-afi-FSP1-4; 0.06 M $MgCl_2$, 0.06 M $CaCl_2$, 0.1 M imidazole/MES monohydrate (acid) pH 6.5, 12.5% MPD, 12.5% PEG1000, 12.5% PEG3350. Crystals were then soaked with 10 mM afi-FSP1-4 added to the crystallisation drop.

- FSP1(L323A)-FAD-NAD$^+$-afi-FSP1-2; 0.06 M $MgCl_2$, 0.06 M $CaCl_2$, 0.1 M HEPES/MOPS (acid) pH 7.5, 20% ethylene glycol, 20% PEG8000. Crystals were then soaked with 1 mM afi-FSP1-2 added to the crystallisation drop.

Electron densities from co-crystallisations with previously reported inhibitors (Supp. Table 5) did not show any ligand binding. Data were collected at the European Synchrotron Radiation Facility (Grenoble, France) and processed with the XDS and CCP4 packages (88, 89). The phase problem was solved by molecular replacement using an Alpha Fold model and CCP4 programs. Model building and refinement were conducted using COOT and Refmac5 (90). Figures were then generated using ChimeraX (91), PyMOL (DeLano Scientific; www.pymol.org).

SPR Binding Assays. SPR binding assays were carried out using the Biacore 8K SPR system (Cytiva) with the running buffer (50 mM HEPES, pH 7.2, 150 mM NaCl, 0.05% Tween-20, and 2% DMSO) under controlled conditions at 25 °C. The His-Sumo–FSP1 protein was immobilized on a preconditioned and activated Series S NTA sensor chip (Cytiva). Preconditioning involved treatment with 350 mM EDTA, followed by activation with 0.5 mM $NiCl_2$. To enhance stability and minimize signal drift, the chip was further activated using EDC (1-ethyl-3-(3-dimethylaminopropyl) carbodiimide) and NHS (N-hydroxysuccinimide) before coupling the protein to the surface. The protein was immobilized to a response level of 7500 RU.

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

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

Virtual Screens. The AdaptiveFlow open-source drug-discovery platform was used to carry out the virtual screen. The size of the docking box was set to $14 \times 18 \times 16$ Å. The FSP1-FAD-NAD$^+$ complex structure was held rigid during the virtual screens (ATG prescreen and primary screen). QuickVina 2 (92) using a hybrid empirical and knowledge-based scoring function was used as the docking program. The exhaustiveness value was set to 1. The number of replicas was set to 1. The virtual screen was run in Amazon Web Services (AWS). Postprocessing of the results including clustering with the similarity limit set to 0.8 was done with DataWarrior (v6.0.0) (93). In detail: For each target, the virtual library was docked and the top ∼2,000 compounds by docking score were retained for further processing. Compounds flagged by DataWarrior "Nasty Functions" and PAINS alerts were removed to exclude known assay-interfering or chemically problematic substructures (e.g., strong Michael acceptors, redox-active quinones, rhodanines, catechols/hydroquinones, azo dyes). The remaining compounds were filtered by simple physicochemical criteria, requiring predicted aqueous solubility $\geq 10\ \mu\mathrm{M}$, overall aromaticity $\leq 75\%$, and fraction of sp^3 carbons ≥ 0.1. This filtered set was then clustered by 2D similarity, and from each cluster one or more representatives were chosen, prioritizing the best docking scores, to yield a final diverse set of ∼80 compounds per target for experimental testing. Apart from these automated substructure and physicochemical filters and the docking scores, no additional interaction-based or other structural criteria were applied.

Discovery of Inhibitor Analogs. Analog Hunter and Scaffold Hopper modules in the infiniSee (94) module from BioSolveIT were used to screen analogs of compound afi-FSP1-1. Top-ranking 5000 compounds from each module were exported to SeeSAR (95) from BioSolveIT and docked with the FlexX algorithm in the general docking mode to the coenzyme Q binding site of the FSP1/FAD/NAD/afi-FSP1-1 complex structure (PDB code: 9IFU) with compound afi-FSP1-1 removed. The docked poses were scored with the HYDE scoring function. Top-ranking 168 compounds with a similar binding pose as compound afi-FSP1-1 were tested with 100-ns-long molecular dynamics (MD) simulations with the PMEMD module (CUDA version) of the AMBER 22 package and protein.ff19SB force field for protein (96). The Molecular Mechanics/Generalized Born Surface Area (MM/GBSA) binding free energies between protein and the ligands that remained in the binding pocket were calculated by the MMPBSA.py script in AmberTools 23 (97). The top-ranking 20 analogs with the lowest MM/GBSA binding free energies were ordered from Enamine for testing.

Data Availability

The ready-to-dock Enamine REAL Space can be explored online via an interactive web interface available on the AdaptiveFlow homepage `http://adaptive-flow.ai/`. The data can be accessed directly via an AWS S3 bucket provided by AWS Open Data (`https://aws.amazon.com/opendata`). 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 AdaptiveFlow homepage https://adaptive-flow.ai/real-space-2022q12 (free registration required).

Code Availability

The source code of AdaptiveFlow is released under the GNU GPL v2.0, a free and open-source license. The code can be accessed via GitHub. The AdaptiveFlow GitHub repos can be found at https://github.com/QuantumAI4Bio, which contains multiple repositories. The AFLP and AFVS modules of AdaptiveFlow can be found in AFLP (https://github.com/QuantumAI4Bio/AdaptiveFlow-LP) and the AFVS (https://github.com/QuantumAI4Bio/AdaptiveFlow-VS) GitHub repositories, respectively. The AFU module (both single machine and parallelized versions) can be found in the GitHub repo (https://github.com/QuantumAI4Bio/AdaptiveFlow-Unity).

Acknowledgments

The authors would like to thank ChemAxon for a free academic license for JChem. AK.N. thanks Anshul Kundaje and Michael Bassik for their support. The authors thank AWS for research credits for the AWS Cloud and thank Nelson Gonzales, Stephen Litster, Mayank Thakkar, Robert Wang, Arun Subramaniyan, and Christine Tsien Silvers for technical and administrative support within AWS. The authors would like to thank Akanksha Bilani and Caroline Rhoades from Intel for their support, and Brigitte Klein for proofreading the manuscript. AK.N. acknowledges funding from the Bio-X Stanford Interdisciplinary Graduate Fellowship (SGIF). C.S. and K.F. acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany´s Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689). H.A. gratefully acknowledges support from the NIH/NIGMS (GM136859 and GM158220), a grant from Pivotal Life Sciences, and a gift from J. Goldberg through Dana-Farber. A.M acknowledges funding from the European Research Council (ERC Advanced Grant, MetaQ no. 101094471), and by the Associazione Italiana per la Ricerca sul Cancro (AIRC) Investigator Grant (no. 28754). This work was supported by ALSAC, the fundraising and awareness organization of St. Jude Children's Research Hospital, and the Blue Sky Kinase Project at St. Jude Children's Research Hospital.

Competing Interests

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

1254 # Extended Data

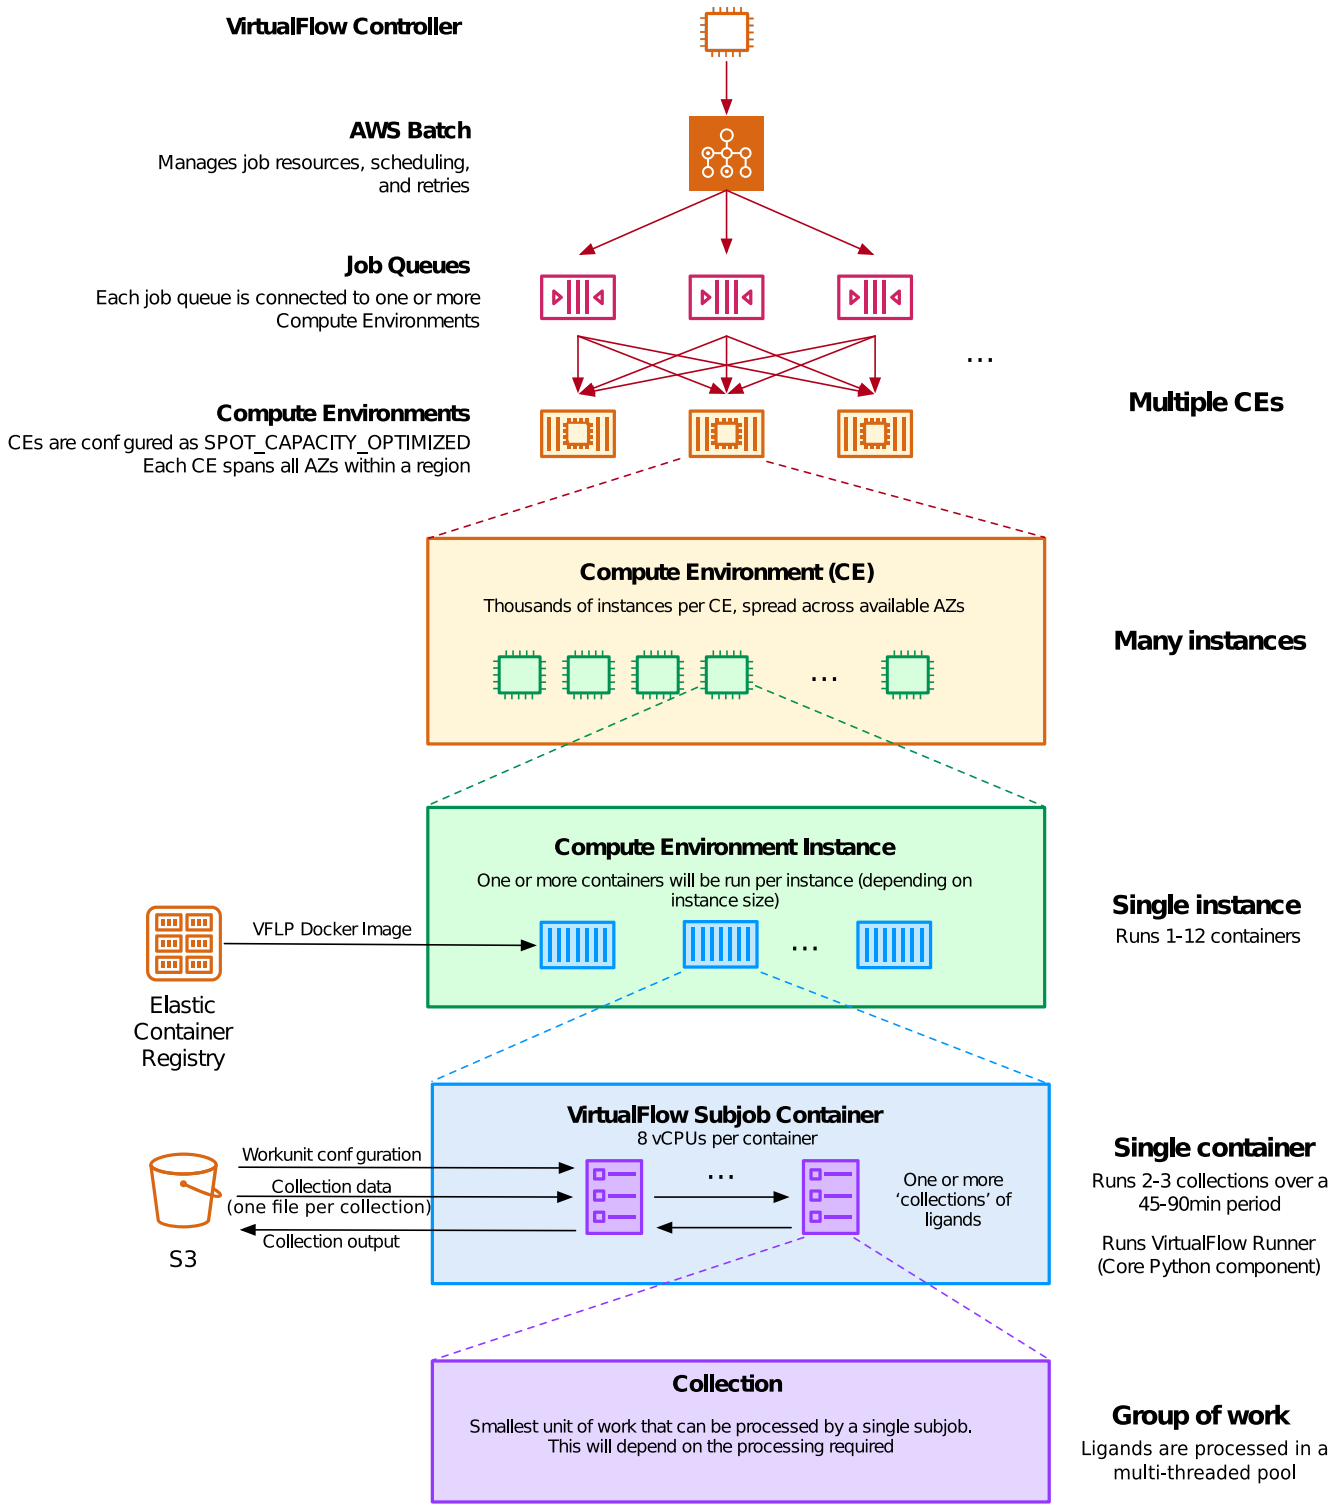

Extended Data Fig. 1. Compute Stack deployed by AdaptiveFlow in AWS. The login node is used by the user to control the workflow in AWS via the workload manager AWS Batch. AWS Batch has multiple Job Queues to which jobs are submitted by AdaptiveFlow, and each job queue deploys multiple Compute Environments 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).

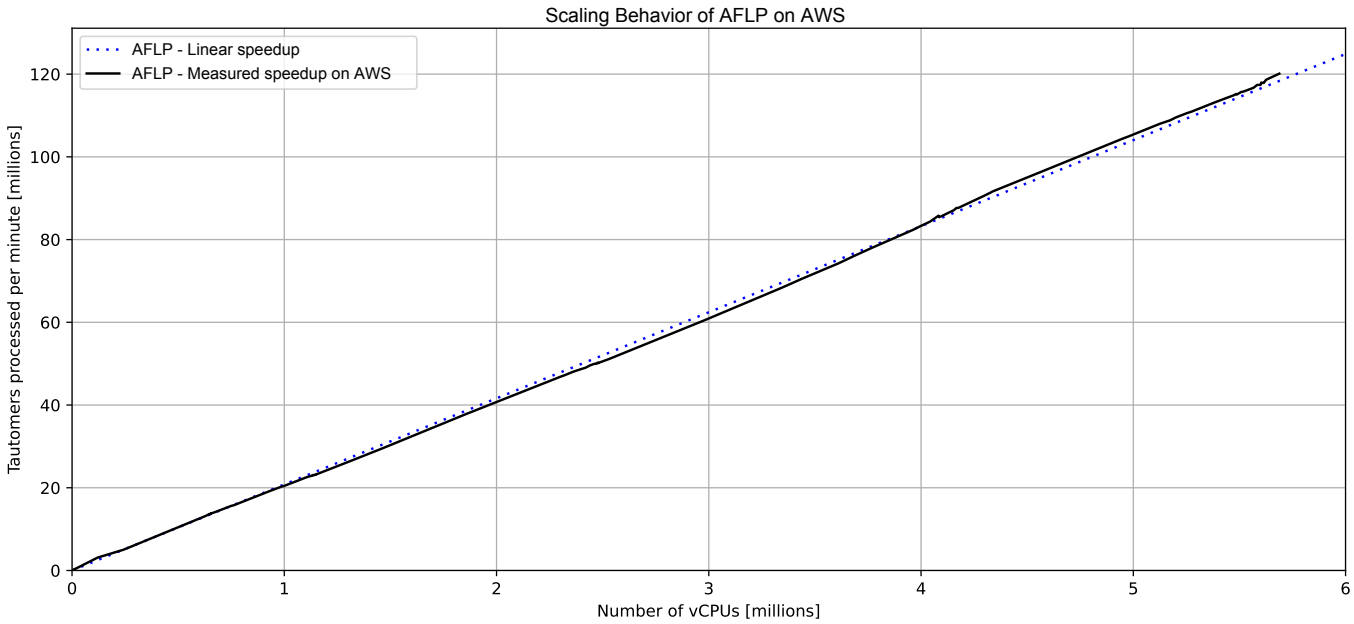

Extended Data Fig. 2. Scaling behavior of AFLP of on AWS. The scaling behavior is perfectly linear even when using extremely large numbers of CPUs. Up to 5.6 million vCPUs were used in parallel across multiple AWS regions. As can be seen in the plot, the scaling behavior is perfectly linear, implying 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**	**200 < X ≤ 300**	**300 < X ≤ 400**	**400 < X ≤ 450**	**X > 450**	
		3.8E+6	564.9E+6	19.8E+9	48.4E+9	0.0E+0	
logP	2	**X ≤ 0**	**0 < X ≤ 2**	**2 < X ≤ 3**	**3 < X ≤ 4**	**4 < X ≤ 5**	**X > 5**
		12.5E+9	30.5E+9	14.2E+9	8.0E+9	2.8E+9	706.3E+6
HBA count	3	**X ≤ 2**	**2 < X ≤ 4**	**4 < X ≤ 7**	**7 < X ≤ 10**	**X > 10**	
		270.9E+6	5.1E+9	31.0E+9	25.5E+9	6.8E+9	
HBD count	4	**X = 0**	**0 < X ≤ 2**	**2 < X ≤ 4**	**4 < X ≤ 5**	**X > 5**	
		6.0E+9	44.5E+9	16.8E+9	1.2E+9	277.2E+6	
Rotatable bond count	5	**X = 0**	**0 < X ≤ 2**	**2 < X ≤ 4**	**4 < X ≤ 7**	**7 < X ≤ 10**	**X > 10**
		2.0E+6	966.3E+6	11.6E+9	35.2E+9	17.6E+9	3.3E+9
TPSA	6	**X ≤ 40**	**40 < X ≤ 80**	**80 < X ≤ 110**	**110 < X ≤ 140**	**X > 140**	
		377.6E+6	24.0E+9	28.9E+9	12.6E+9	2.9E+9	
Molecular refractivity	7	**X ≤ 40**	**40 < X ≤ 80**	**80 < X ≤ 130**	**X > 130**		
		79.2E+3	540.6E+6	63.6E+9	4.6E+9		
logS	8	**X ≤ -6**	**-6 < X ≤ -5**	**-5 < X ≤ -4**	**-4 < X ≤ -3**	**X > -3**	
		4.9E+9	10.7E+9	17.1E+9	16.8E+9	19.2E+9	
Aromatic proportion	9	**X = 0**	**0 < X ≤ 0.25**	**0.25 < X ≤ 0.5**	**0.25 < X ≤ 0.75**	**0.75 < X ≤ 1.0**	
		6.8E+9	25.9E+9	31.9E+9	4.0E+9	49.7E+6	
Formal charge	10	**X ≤ -2**	**X = -1**	**X = 0**	**X = 1**	**X ≥ 2**	
		158.3E+6	3.1E+9	59.9E+9	5.4E+9	140.9E+6	
Positive charge count	11	**X = 0**	**X = 1**	**X ≥ 2**			
		60.1E+9	8.2E+9	341.7E+6			
Negative charge count	12	**X = 0**	**X = 1**	**X ≥ 2**			
		62.5E+9	5.7E+9	484.1E+6			
Fsp3	13	**0 ≤ X ≤ 0.2**	**0.2 < X ≤ 0.4**	**0.4 < X ≤ 0.6**	**0.6 < X ≤ 0.8**	**0.8 < X ≤ 1.0**	
		647.0E+6	11.9E+9	30.8E+9	20.7E+9	4.7E+9	
Chiral center count	14	**X = 0**	**X = 1**	**X ≥ 2**			
		6.1E+9	18.7E+9	43.9E+9			
Doublebond stereoisomer count	15	**X = 0**	**X ≥ 1**				
		59.6E+9	9.1E+9				
Halogen atom count	16	**X = 0**	**X ≥ 1**				
		45.7E+9	23.0E+9				
Sulfur atom count	17	**X = 0**	**X ≥ 1**				
		54.2E+9	14.5E+9				
Enamine compound class	18	**X = S**	**X = M**				
		66.9E+9	1.8E+9				

Color Scale

0 100M 50B

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).

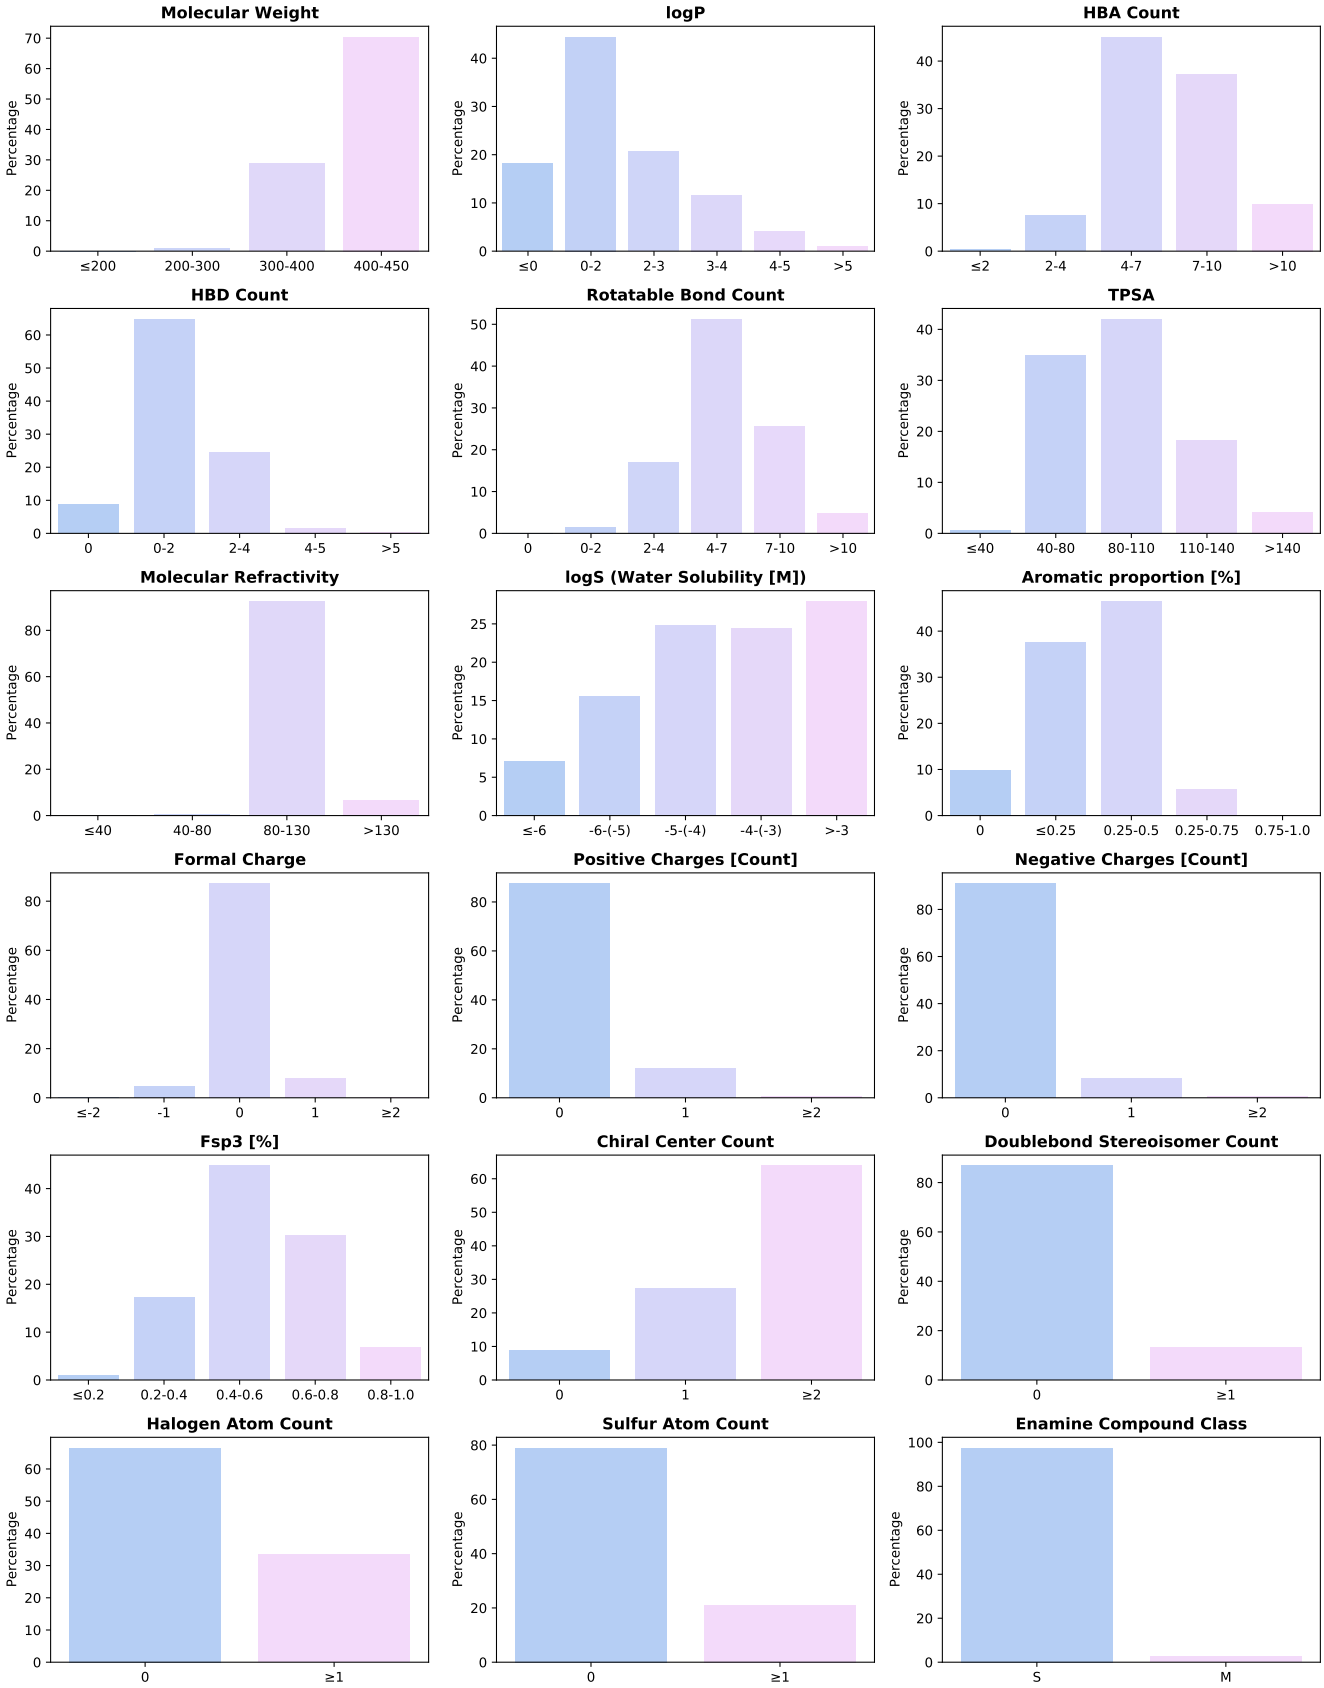

Extended Data Fig. 4. Overview of the distribution of the molecular properties of the Enamine REAL Space (version 2022q1-2). Each subplot shows the distribution of the molecular property associated with each dimension of the 18-dimensional grid into which the REAL Space was partitioned.

MW

SlogP

	0 - 200	200 - 300	300 - 400	400 - ∞	Row sums:
-∞ - 0	1.03M	121M	4.39B	7.95B	12.5B
0 - 2	2.04M	265M	9.34B	20.9B	30.5B
2 - 3	513K	104M	3.59B	10.5B	14.2B
3 - 4	157K	53.6M	1.75B	6.18B	7.99B
4 - 5	17.2K	18.4M	552M	2.28B	2.85B
5 - ∞	253	3.99M	149M	553M	706M
Column sums:	3.76M	565M	19.8B	48.4B	

Compounds selected: 68.7B **from** 68.7B **Tranches selected:** 12.3M **from** 12.3M

Molecular Weight: 0 to ∞
0 200 300 400 ∞

Formal Charge: -∞ to ∞
-∞ -2 -1 0 1 ∞

Partition Coefficient (SlogP): -∞ to ∞
-∞ 0 2 3 4 5 ∞

Positive Charge Count: 0 to ∞
0 1 ∞

Hydrogen Bond Acceptors: 0 to ∞
0 2 4 7 10 ∞

Negative Charge Count: 0 to ∞
0 1 ∞

Hydrogen Bond Donors: 0 to ∞
0 2 4 5 ∞

Fraction of sp^3 Carbon Atoms, %: 0 to 100
0 20 40 60 80 100

Rotatable Bonds: 0 to ∞
0 2 4 7 10 ∞

Chiral Center Count: 0 to ∞
0 1 ∞

Topological Polar Surface Area: 0 to ∞
0 40 80 110 140 ∞

Doublebond Stereoisomer Count: 0 to ∞
0 ∞

Molecular Refractivity: 0 to ∞
0 40 80 130 ∞

Halogen Atom Count: 0 to ∞
0 ∞

Solubility Coefficient: -∞ to ∞
-∞ -6 -5 -4 -3 ∞

Sulfur Atom Count: 0 to ∞
0 ∞

Aromatic Proportion, %: 0 to 100
0 25 50 75 100

Enamine Compound Class: S to M
S M

Download method for tranches: wget Download

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.

	≤200	300	400	450	Row sums:
0	1M	120.8M	4.4B	8B	12.5B
2	2M	264.6M	9.3B	20.9B	30.5B
3	512.9K	103.6M	3.6B	10.5B	14.2B
4	157.4K	53.6M	1.7B	6.2B	8B
5	17.2K	18.4M	552.5M	2.3B	2.8B
>5	253	4M	149.2M	553.1M	706.3M
Column sums:	3.8M	564.9M	19.8B	48.4B	

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.

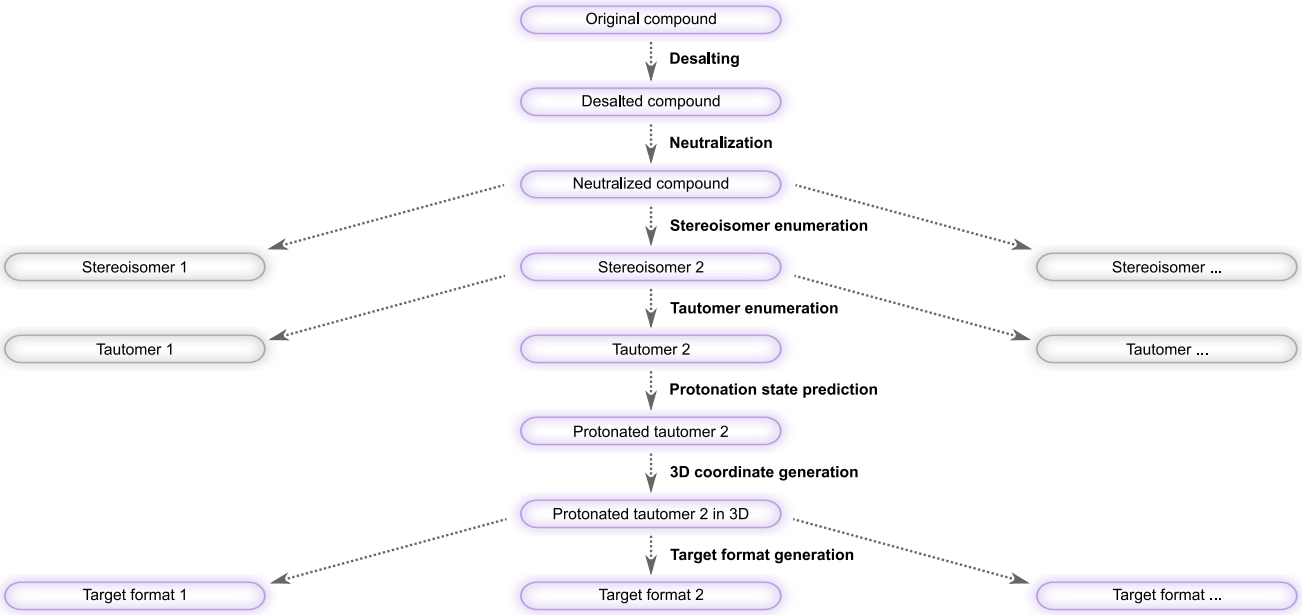

Extended Data Fig. 7. AFLP - Overview of the Ligand Preparation Procedure. Ligands are initially in a format in line notation (e.g. SMILES, SELFIES, or amino acid sequences). The preparation steps include desalting, neutralization, stereoisomer enumeration, tautomer enumeration, protonation state prediction, 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
logP	**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
HBA count	**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
HBD count	**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
TPSA	**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
	E	X > 140	-8.10076	1	-8.08286	1	-8.32545	1
Molecular refractivity	**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
logS	**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
Aromatic proportion	**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
Formal charge	**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
Negative charge count	**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
Fsp3	**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
	E	0.8 < X ≤ 1.0	-7.02762	5	-6.99253	5	-7.26721	5
Chiral center count	**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
Doublebond stereoisomer count	**A**	X = 0	-7.7409	1	-7.72163	1	-7.92645	1
	B	X ≥ 1	-7.58906	2	-7.56916	2	-7.81019	2
Halogen atom count	**A**	X = 0	-7.6665	2	-7.65526	2	-7.88305	2
	B	X ≥ 1	-7.7349	1	-7.70825	1	-7.90392	1
Sulfur atom count	**A**	X = 0	-7.79078	1	-7.77764	1	-7.98746	1
	B	X ≥ 1	-7.5662	2	-7.53497	2	-7.7627	2
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.

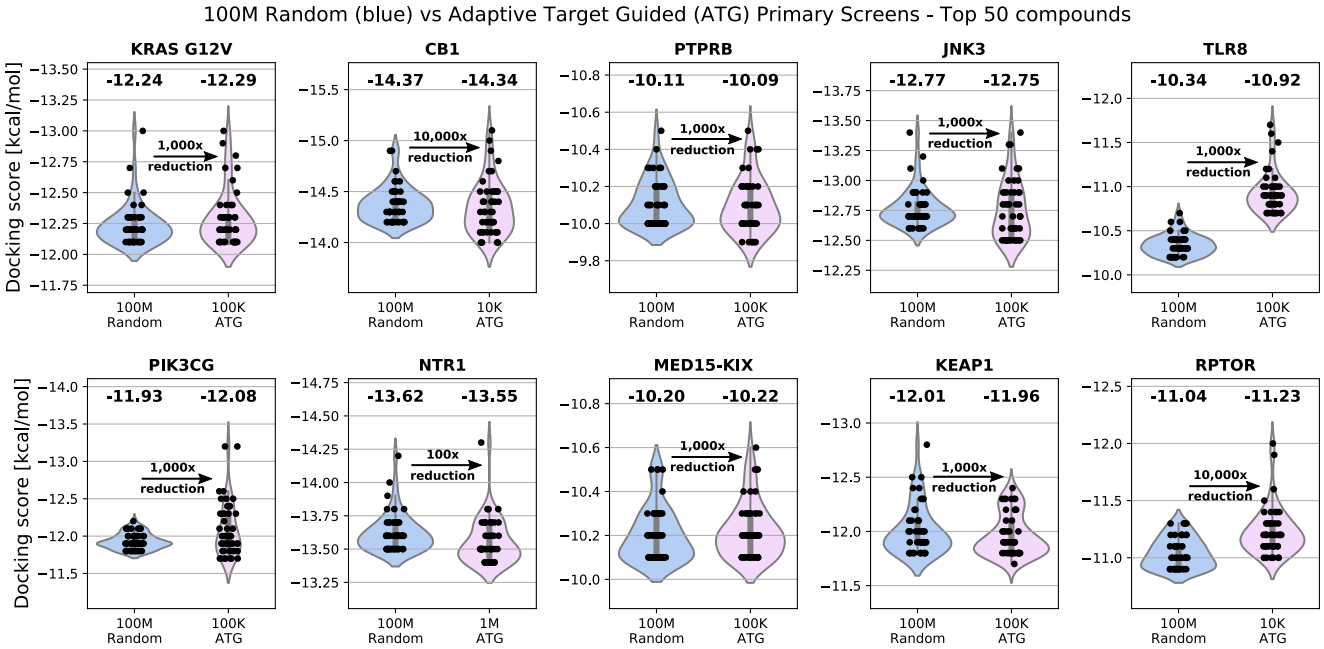

Extended Data Fig. 9. Benchmarking large-scale virtual screening strategies. Violin plots compare the distribution of docking scores for the top 50 molecules identified by two approaches: (1) a standard ultra-large virtual screen of 100 million randomly selected compounds, and (2) the Adaptive Target-Guided (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.

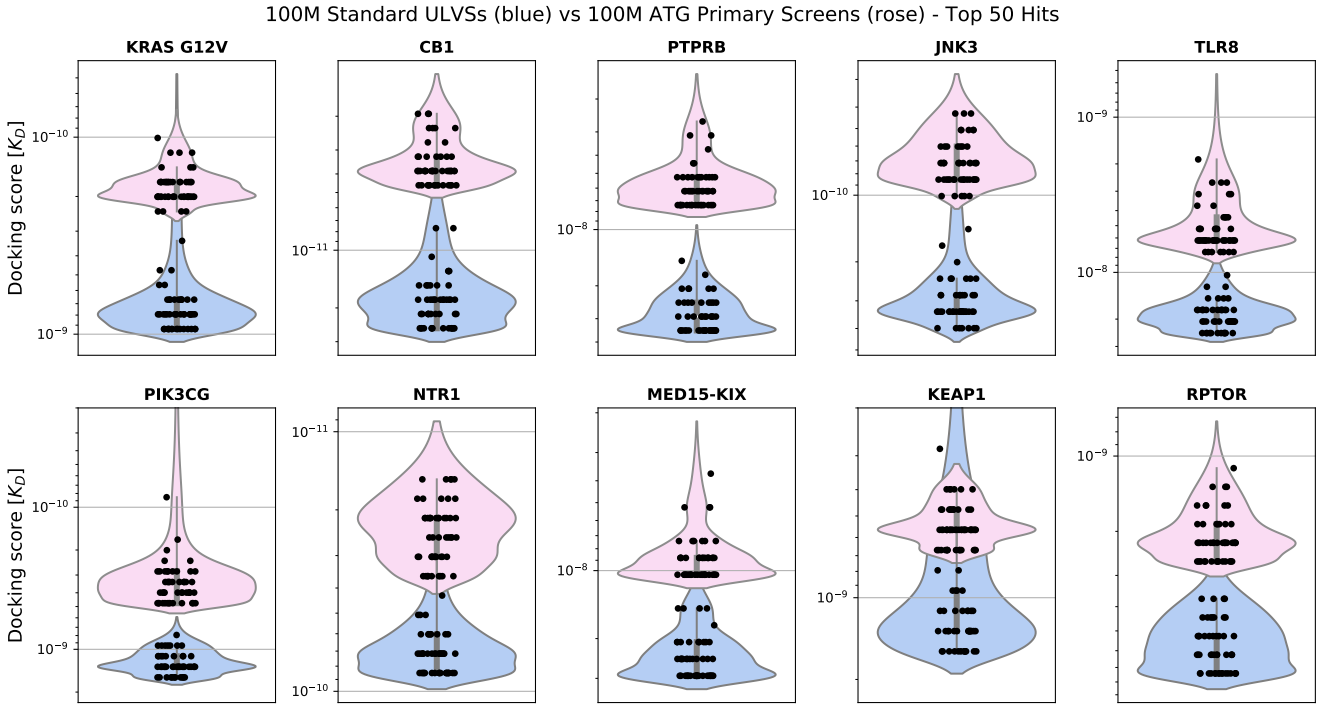

Extended Data Fig. 10. Comparison of Adaptive Target-Guided (ATG) Primary Virtual Screens and standard ULVS. 100 million compounds were screened with standard ULVS of randomly selected compounds, as well as with ATGs where the 100 million compounds to be screened were selected based on the ATG prescreen. 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 overlayed. The docking scores of the top 50 compounds of the adaptive screens are considerably better for all target proteins than for the standard ULVSs.

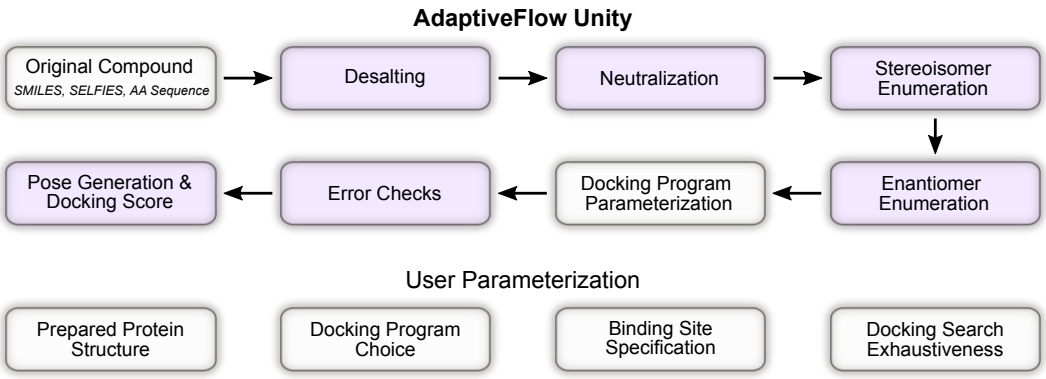

Extended Data Fig. 11. Conceptual Overview of AdaptiveFlow Unity (AFU). AFU takes as input a user-specified string representation of a molecule and processes the provided molecule by desalting, neutralizing, and generating stereoisomer (with RDKit (65)). Subsequently, based on the docking program choice, all ligands are converted into 3D conformers (with OpenBabel (63)) 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.

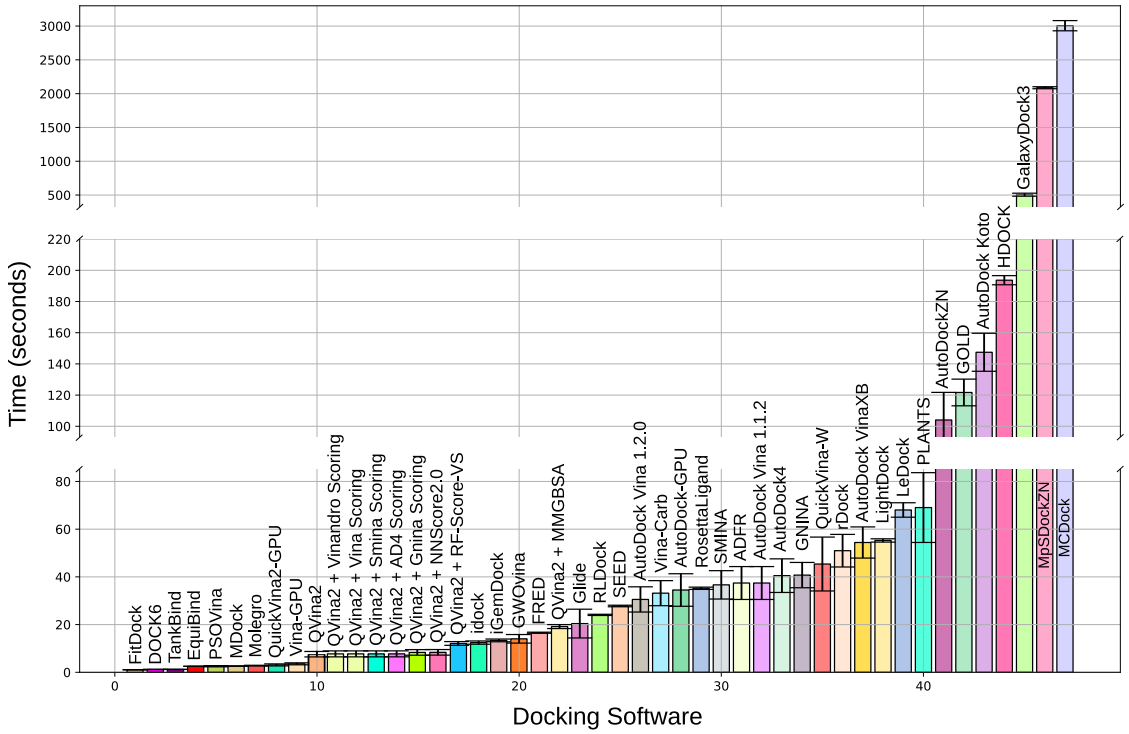

Extended Data Fig. 12. Docking Time Estimates. Timing estimates for running a single calculation averaged over 25 independent runs for all docking programs. The calculations were run using AFU.

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-VSs)	available
Interactive web interface	yes (18 dimensions)
Accessability	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 protocols	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 protocols)
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 coun	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 sp3 carbon atoms (fsp3)	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 (98), QupKake (99) 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	(59)
AutoDock4	Open Source	Improved modelling of waters and macrocycles	✗	CPU	(100)
AutoDockZN	Open Source	Special support for zinc-coordinating ligands	✗	CPU	(101)
AutoDock Vina 1.2.0	Open Source	Python support; macromolecule support	✗	CPU	(102)
AutoDock Vina 1.1.2	Open Source	Improved optimization, scoring and multithreading	✗	CPU	(103)
AutoDockFR (ADFR)	Open Source	Improved receptor side-chain flexibility modelling	✗	CPU	(104)
AutoDock Koto	Open Source	Improved sampling	✗	CPU	(105)
ATPdock	Open Source	ATP-specific protein-ligand docking	✗	CPU	(106)
CovDock	Licence Required	Covalent docking	✗	CPU	(107)
DOCK 6.0	Free for academics	RNA-ligand, DNA-ligand, protein-protein docking	✗	CPU	(108)
FlexX	Licence Required	Covalent docking. Ensemble docking	✗	CPU	(109)
FitDock	Academic Licence	Incorporation of template fitting for docking	✗	CPU	(110)
GalaxyDock3	Licence Required	Incorporation of side chain flexibility	✗	CPU	(111)
Glide (HTVS)	Paid Academic Licencing	Less exhaustive sampling and scoring	✗	CPU	(112)
Glide (SP)	Paid Academic Licencing	More exhaustive sampling (relative to Glide HTVS) & multiple levels of accuracy	✗	CPU	(112)
Glide (XP)	Paid Academic Licencing	Most accurate sampling and scoring for Glide & different levels of accuracy	✗	CPU	(112)
Gnina	Open Source	Deep learning scoring of complex	✓	GPU/CPU scoring	(70)
GOLD	Licence Required	Genetic algorithm based docking; Covalent docking	✗	CPU	(113)
GWOVina	Open Source	Grey wolf optimization	✗	CPU	(114)
HDOCK	Academic Licence	Protein–RNA/DNA docking	✗	CPU	(115)
HSYMDOCK	Academic Licence	Protein–protein/rna docking with Cn and Dn symmetry	✗	CPU	(116)
idock	Open Source	AutoDock Vina derivative with improved speed	✗	CPU	(117)
iGEMDOCK	Open Source	-	✗	CPU	(118)
LeDock	Freeware for academic use	Excellent Sampling power (comparable to GOLD)	✗	CPU	(119)
LigandFit	Open Source	Identification of active site regions for docking	✗	CPU	(120)
LightDock	Open Source	Protein-protein docking	✗	CPU	(121)
MDock	Licence Required	Ensemble docking for multiple protein conformations	✗	CPU	(122)
MCDock	Open Source	Use of Monte Carlo simulation technique	✗	CPU	(123)
Molegro	Licence Required	-	✗	CPU	(124)
MpsDockzn	Licence Required	Zinc metalloprotein-specific scoring function	✗	CPU	(125)
PSOVina	Open Source	Improved pose sampling and speed	✗	CPU	(126)
PLANTS	Open Source	Use of ant colony optimization	✗	CPU	(127)
PI-LZerD	Open Source	Protein-protein docking	✗	CPU	(128)
PIPER	Open Source	Use of pairwise interaction potentials for protein-protein docking	✗	CPU	(129)
QuickVina 2	Open Source	Extremely fast	✗	CPU	(92)
QuickVina-W	Open Source	Improved speed for blind docking	✗	CPU	(130)
QVina2-GPU	Open Source	GPU-accelerated version of QVina2	✗	GPU	(60)
QVina2-W-GPU	Open Source	GPU-accelerated version of QVina2-W	✗	GPU	(60)
rDock	Open Source	Support for RNA targets	✗	CPU	(131)
Rosetta Ligand	Licence Required	-	✗	CPU	(132)
RLDOCK	Open Source	RNA-ligand docking	✗	CPU	(133)
SEED	Open Source	Fragment-based docking	✗	CPU	(134)
SMINA	Open Source	User specific scoring functions	✗	CPU	(135)
TANKBind	Open Source	Fast neural network based complex prediction	✓	CPU/GPU	(12)
VinaCarb	Open Source	Improved glycosidic modelling	✗	CPU	(136)
VinaXB16	Open Source	Improved halogen bonding scoring function	✗	CPU	(137)
Vina-GPU	Open Source	GPU-accelerated version of Vina	✗	GPU	(60)
Vina-GPU-2.0	Open Source	GPU-accelerated version of Vina-2.0	✗	GPU	(60)

Extended Data Table 4. Supported docking programs. The collection of unique underlying docking programs supported within AFVS and AF Unity. These programs serve as the foundational components for the 1,500 supported docking protocols (specific combinations of pose prediction and scoring algorithms) enabled by the platform.

Docking program	License model	Special features	DL-based	Hardware support	Reference
AD4 Scoring	Open Source	Improved over AutoDock 3 forcefield	✗	CPU	(138)
ASP Scoring	Licence required	Atom-atom distance potential	✗	CPU	(113)
Bump Filter Scoring	Free for academics	Consideration of molecule overlap with receptor atoms	✗	CPU	(108)
ChemScore	Licence required	Improved speed compared to GoldScore	✗	CPU	(113)
Contact Scoring	Free for academics	Incorporates the scoring performed with the DISTMAP program	✗	CPU	(108)
Continuous Scoring	Free for academics	No grid calculations	✗	CPU	(108)
ChemPLP	Licence required	Slightly faster than ChemScore and up to four times faster than GoldScore	✗	CPU	(113)
DOCK3.5 Scoring	Free for academics	Incorporation of solvation effects	✗	CPU	(108)
DeepAffinity	Open Source	Use of convolutional neural networks	✓	CPU/GPU	(72)
DeepBindRG	Open Source	Use of recurrent and convolutional neural networks	✓	CPU/GPU	(73)
Footprint Scoring	Free for academics	Calculates intermolecular hydrogen bonds and footprint comparisons, in addition to standard interactions	✗	CPU	(108)
Gnina Scoring	Open Source	Convolutional neural network scoring	✓	CPU/GPU	(70)
GOLD Scoring	Licence required	Use of the chemscore function	✗	CPU	(113)
GoldScore	Licence required	Original scoring function provided with GOLD	✗	CPU	(113)
Grid-Based Scoring	Free for academics	Based on the non-bonded terms of the molecular mechanic force field	✗	CPU	(108)
Hawkins GB/SA Scoring	Free for academics	Use of pairwise GB solvation model	✗	CPU	(139)
NNScore 2.0	Open Source	Use of a fully-connected neural networks	✓	CPU/GPU	(71)
MM/PBSA	Open Source	Calculation of end-state free energies of molecules in solution	✗	CPU	(76)
MM/GBSA	Open source	Improved scoring accuracy	✗	CPU	(76)
MultiGrid FPS Scoring	Free for academics	Pair-wise interaction energies are computed over multiple grids rather than in Cartesian space	✗	CPU	(108)
OpenBPMD	Open Source	Improved scoring accuracy	✗	CPU	(74)
PLP Scoring	Open source	Pairwise additive scoring function	✗	CPU	(127)
PLP95 Scoring	Open source	Origin of PLP	✗	CPU	(127)
PB/SA Scoring	Free for academics	Application of the ZAP electrostatics toolkit from OpenEye	✗	CPU	(108)
Pharmacophore Matching Similarity Scoring	Free for academics	Calculates the level of pharmacophore overlap between a reference molecule and a candidate molecule	✗	CPU	(108)
Rosetta Scoring	Open Source	Improved scoring for side-chain flexibility	✗	CPU	(132)
RF-Score-VS	Open Source	Random Forest-based scoring function	✗	CPU	(140)
SASA Scoring	Free for academics	Calculates the percent exposure of a ligand	✗	CPU	(141)
SEED Scoring	Open Source	Continuum dielectric approximation of electrostatic contribution	✗	CPU	(134)
SMINA Scoring	Open Source	Improved AutoDock Vina scoring	✗	CPU	(135)
Vinardo Scoring	Open Source	Improved Vina scoring	✗	CPU	(142)
Vina Scoring	Open Source	Improved scoring compared to AD4	✗	CPU	(103)
Zou GB/SA Scoring	Free for academics	Fast algorithm, pairwise free energy model, for ligand binding affinity calculations	✗	CPU	(141)

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	(68)
EquiBind	Open Source	Blind docking pose prediction	✓	GPU/CPU	(69)
EnzyDock	Open Source	Identifies probable binding modes of multiple ligand states in an enzyme active site	✗	CPU	(143)
FRED	Academic Licence	Improved pose sampling	✗	CPU	(144)

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 ULVSs.