

Quantum Computing-Enhanced Algorithm Unveils Novel Inhibitors for KRAS

Mohammad Ghazi Vakili,^{1,2} Christoph Gorgulla,^{3,4} AkshatKumar Nigam,⁵ Dmitry Bezrukov,⁶ Daniel Varoli,⁷ Alex Aliper,⁶ Daniil Polykovsky,⁶ Krishna M. Padmanabha Das,^{8,9} Jamie Snider,¹⁰ Anna Lyakisheva,¹⁰ Ardalan Hosseini Mansob,^{10,11} Zhong Yao,¹⁰ Lela Bitar,^{10,12} Eugene Radchenko,⁶ Xiao Ding,⁶ Jinxin Liu,⁶ Fanye Meng,⁶ Feng Ren,⁶ Yudong Cao,¹³ Igor Stagljjar,^{10,14,11,15} Alán Aspuru-Guzik,^{1,2,16,17,18,19,*} and Alex Zhavoronkov^{6,†}

¹*Department of Computer Science, University of Toronto, Canada*

²*Department of Chemistry, University of Toronto, Canada*

³*Department of Structural Biology, St. Jude Children’s Research Hospital, USA*

⁴*Department of Physics, Harvard University, USA*

⁵*Department of Computer Science, Stanford University*

⁶*Insilico Medicine AI Limited, Abu Dhabi, UAE*

⁷*Zapata AI, 100 Federal St., Boston, MA 02110*

⁸*Department of Cancer Biology, Dana-Farber Cancer Institute, Boston, MA, USA*

⁹*Department of Biological Chemistry and Molecular Pharmacology,*

Harvard Medical School, Harvard University, Boston, MA, USA

¹⁰*Donnelly Centre, Temerty, Faculty of Medicine, University of Toronto, ON, Canada*

¹¹*Department of Molecular Genetics, University of Toronto, Ontario, Canada*

¹²*Department for Lung Diseases Jordanovac, Clinical Hospital Centre Zagreb, University of Zagreb, Croatia*

¹³*Zapata AI, 100 Federal St., Boston, MA, USA*

¹⁴*Department of Biochemistry, University of Toronto, Ontario, Canada*

¹⁵*Mediterranean Institute for Life Sciences, Split, Croatia*

¹⁶*Department of Chemical Engineering and Applied Chemistry, University of Toronto, Canada*

¹⁷*Department of Materials Science and Engineering, University of Toronto, Canada*

¹⁸*Vector Institute for Artificial Intelligence, Toronto, Canada*

¹⁹*Lebovic Fellow, Canadian Institute for Advanced Research (CIFAR), Toronto, Ontario, Canada*

Abstract The discovery of small molecules with therapeutic potential is a long-standing challenge in chemistry and biology. Researchers have increasingly leveraged novel computational techniques to streamline the drug development process to increase hit rates and reduce the costs associated with bringing a drug to market. To this end, we introduce a quantum-classical generative model that seamlessly integrates the computational power of quantum algorithms trained on a 16-qubit IBM quantum computer with the established reliability of classical methods for designing small molecules. Our hybrid generative model was applied to designing new KRAS inhibitors, a crucial target in cancer therapy. We synthesized 15 promising molecules during our investigation and subjected them to experimental testing to assess their ability to engage with the target. Notably, among these candidates, two molecules, ISM061-018-2 and ISM061-22, each featuring unique scaffolds, stood out by demonstrating effective engagement with KRAS. ISM061-018-2 was identified as a broad-spectrum KRAS inhibitor, exhibiting a binding affinity to KRAS-G12D at $1.4\mu M$. Concurrently, ISM061-22 exhibited specific mutant selectivity, displaying heightened activity against KRAS G12R and Q61H mutants. To our knowledge, this work shows for the first time the use of a quantum-generative model to yield experimentally confirmed biological hits, showcasing the practical potential of quantum-assisted drug discovery to produce viable therapeutics. Importantly, comparative analysis with existing classical generative models indicates that integrating quantum computing enhances distribution learning from established datasets, suggesting a potential advantage for quantum generative models over their classical counterparts. Moreover, our findings reveal that the efficacy of distribution learning correlates with the number of qubits utilized, underlining the scalability potential of quantum computing resources. Overall, we anticipate our results to be a stepping stone towards developing more advanced quantum generative models in drug discovery.

I. INTRODUCTION

Drug discovery is a multifaceted process involving various stages, including identifying, developing, and rigorously testing novel molecules intended to combat a spectrum of diseases [1]. A drug discovery campaign

typically spans a decade to fifteen years or more and commands a financial commitment that often exceeds \$2.5 billion [2] during clinical trials. Significantly, these substantial investments do not guarantee success; when a drug development cycle fails, it represents a financial setback with the potential loss of the entire capital investment [3]. Consequently, the pharmaceutical industry continually seeks innovative and cutting-edge technologies to integrate into their workflows, aiming to enhance their prospects for successful market entry.

* Correspondence to: aspuru@utoronto.ca

† Correspondence to: alex@insilicomedicine.com

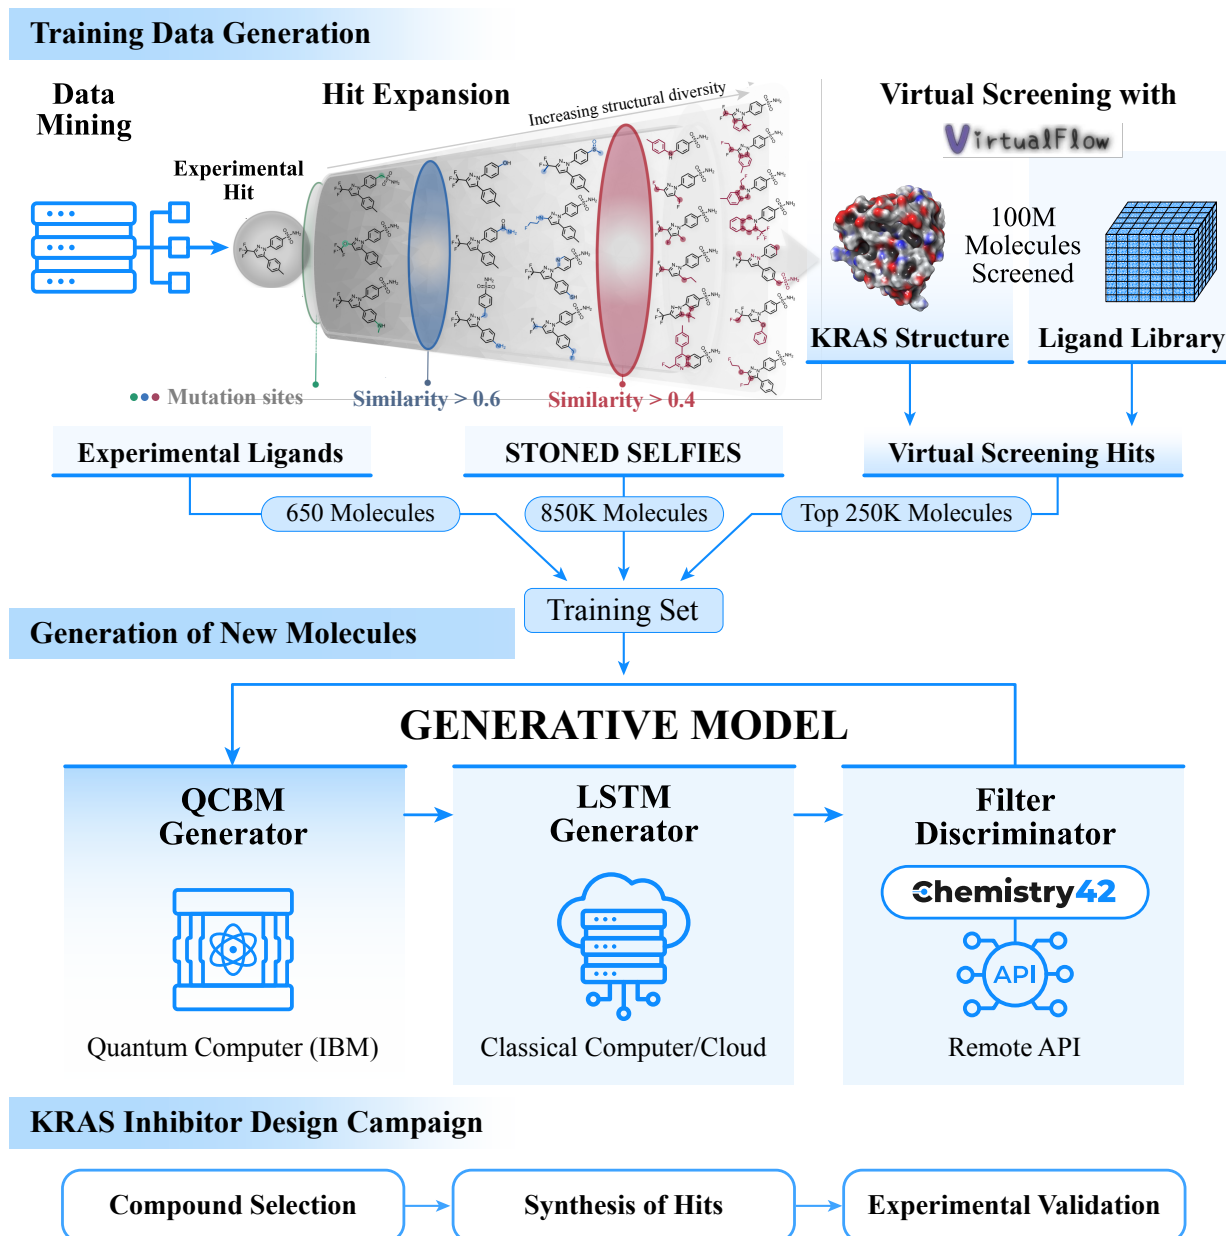


FIG. 1. Schematic Representation of the Hybrid Quantum-Classical Framework for KRAS Ligand Development. The initial phase concentrates on compiling a dataset for model training. A curated set of 650 experimentally verified inhibitors targeting the KRAS protein is extracted from the literature. By applying the STONED-SELFIES algorithm, analogs for each identified compound are derived, yielding an expanded collection of around 850,000 compounds. This dataset is further enhanced by the addition of the top 250,000 candidates, identified via a virtual screening process using the REAL ligand library against the KRAS protein, culminating in a dataset of over 1 million molecules for training our generative model. Upon completing the training of our model, new molecules targeting KRAS are created employing both a classical LSTM network and a Quantum Circuit Born Machine (QCBM) as the underlying generative frameworks. The LSTM network processes sequential data encapsulating the chemical structures of ligands, while QCBM, trained based on the quality of LSTM-generated samples, creates complex, high-dimensional probability distributions. The combined workflow utilizes Chemistry42 as a reward function to incentivize the creation of structurally diverse and synthesizable molecules.

The drug discovery journey commences with identifying a critical target, usually a protein or enzyme integral to a disease’s pathophysiology [4]. Following this step, researchers employ many techniques, notably virtual screening [5–7], to design and rigorously assess potential drug candidates creatively. These candidates are meticulously evaluated for their proficiency in engaging with and modulating the target, propelling the pursuit of therapeutic innovations [8]. Concurrently, generative modelling is emerging as a transformative technology in molecule design [9–12]. Generative models utilize machine learning techniques to comprehend the underlying distribution of atoms and bonds in a specified dataset. Subsequently, these models are employed to construct molecules with predefined properties, a process known as inverse molecular design [13–15]. A promising aspect of these models is their ability to navigate the vast chemical space, proposing interesting molecules within the challenging realm of 10^{60} drug-like molecules [16].

A transition from the intricate landscape of traditional drug discovery to the realm of advanced computational techniques underscores the industry’s adaptation to innovative methodologies. Within this evolution, quantum machine learning has garnered attention, particularly in enhancing generative models. Hibat-Allah et al. [17] introduced a framework that juxtaposes the performance of quantum and classical generative models, focusing on the practical quantum advantage and the potential superiority of Quantum Circuit Born Machines (QCBM) over their conventional counterparts. Their research highlights the generalization capabilities of QCBMs, notably in generating novel, valid samples that extend beyond the training dataset. These models adhere to the target distribution and address training challenges such as barren plateaus, effectively integrating tensor networks alongside QCBMs [18, 19]. Despite the promising strides, quantum information processing confronts inherent limitations, notably in data loading and trainability within quantum circuits. These challenges have prompted a shift towards hybrid algorithms, amalgamating the strengths of both quantum and classical machine learning paradigms. A notable contribution in this regard is from Manuel et al. [20], who demonstrated the enhanced exploration of target space facilitated by integrating a multi-basis QCBM into a classical framework, termed the Quantum Circuit Associative Adversarial Network (QC-AAN). This hybrid model not only augments model training beyond the capabilities of classical hardware but also transcends the constraints imposed by the size of the quantum prior, signifying a significant leap in the field of quantum-enhanced machine learning. Furthermore, Zeng et al. [21] introduced a quantum-classical hybrid framework (conditional QCBM and a classical model) for image generation tasks, utilizing quantum circuits that encode conditional information via additional qubits.

Here, we propose a novel quantum-classical generative model specifically designed to overcome qubit limitations while combining the best of classical and quantum computational methods. Our hybrid model is engineered to generate realistic ligands to design compounds for targeted proteins. A prime focus of our endeavour is the KRAS protein, a target notorious for its intricate complexity and historical resistance in drug discovery ventures [22–24]. The comprehensive workflow of our approach, spanning from data generation to experimental validation, is illustrated in Figure 1. Notably, our model capitalizes on established classical computational pipelines. We employ the STONED-SELFIES algorithm [25] for its proficiency in exploring and interpolating the vast chemical space. Additionally, we incorporate Chemistry42 [26] as an external reward signal, a strategic move to enhance the realistic design of molecules. This integration ensures that our generated ligands exhibit novel and drug-like properties and are highly relevant for the targeted protein. To validate the predictions of our model, we synthesized a series of 15 candidate compounds. These engineered compounds underwent extensive experimental evaluation to ascertain their efficacy in targeting KRAS. One candidate, distinguished by a novel chemotype, manifested as a pan-KRAS inhibitor. We believe that this study lays foundational groundwork for the integration of quantum algorithms in the field of drug discovery. This work underscores the potential and viability of quantum-enhanced methodologies in drug discovery by achieving what we understand to be the first experimental hit attributed to a quantum algorithm.

II. RESULTS AND DISCUSSION

Our methodology encompasses a comprehensive workflow, extending from data preparation to experimental validation, as delineated in Figure 1. This workflow is structured into three pivotal stages:

(1) Generation of Training Data: We initiate the process by constructing a robust dataset for training our generative model to target the KRAS protein. The foundation of this dataset is approximately 650 experimentally confirmed KRAS inhibitors, compiled through an extensive literature review [28–31]. Acknowledging the necessity of a more expansive dataset to develop a model for ligand design effectively, we adopted a two-pronged approach: virtual screening and local chemical space exploration. In the virtual screening phase, we employed Virtual Flow 2.0 [6] to screen 100 million molecules, utilizing Enamine’s REAL library [32] in conjunction with molecular docking techniques. The top 250k compounds from this screen, exhibiting the lowest docking scores, were subsequently integrated into our dataset. Complementing this, the local chemical space exploration was conducted using the STONED-SELFIES algorithm [33], which was applied

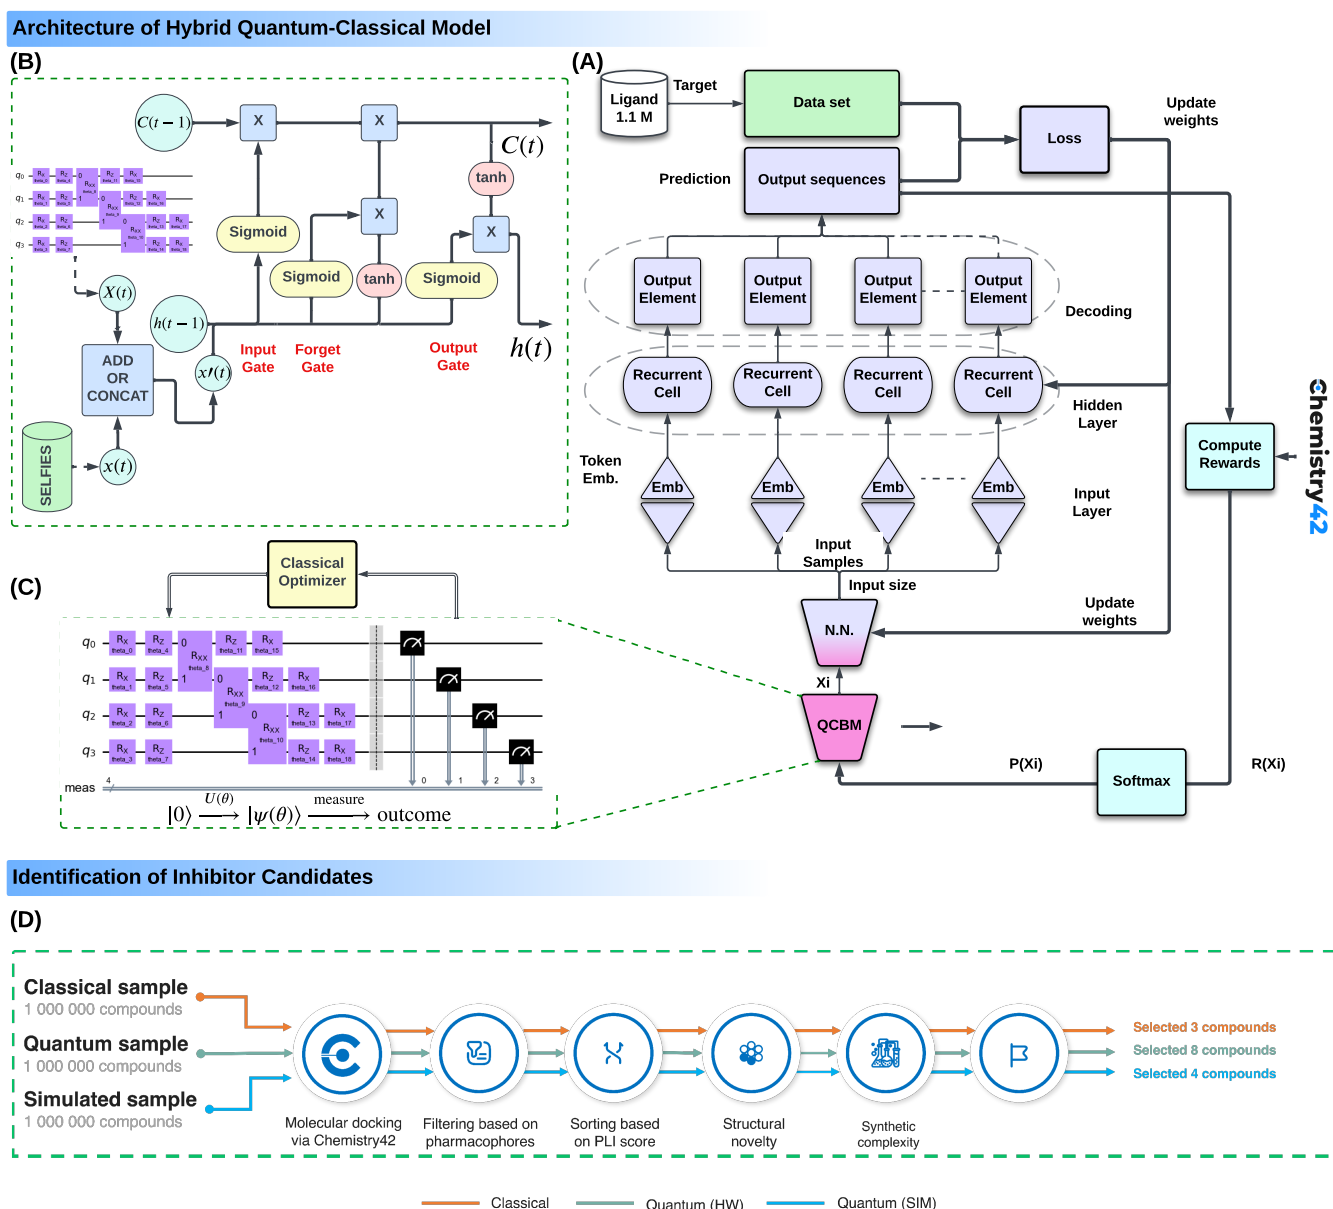


FIG. 2. Quantum-Enhanced Generative Model for Drug Discovery Applications. (A) Hybrid model combining a Quantum Circuit Born Machine (QCBM) with Long Short-Term Memory (LSTM). This model iteratively trains using prior samples from quantum hardware. (B) Integration method of prior samples into the LSTM architecture. Molecular information (in SELFIES encoding) and quantum data are merged by addition or concatenation. The resultant samples, $X'(t)$, are then input to the LSTM cell. (C) Quantum prior component described as a QCBM, generating samples from quantum hardware each training epoch and trains with a reward value, $P(x) = \text{Softmax}(R(x))$, calculated using Chemistry42 or a local filter. (D) Process of experimental sample selection: 1 million compounds are sampled from each model—classical samples (via vanilla LSTM), quantum samples (QCBM on quantum hardware), and simulated samples (quantum simulation on classical hardware). These samples undergo evaluation by Chemistry42, filtering out compounds unsuitable for pharmacological purposes and ranking the remaining compounds by their docking score (PLI score). Subsequently, 15 novel compounds were selected for synthesis.

to the 650 experimentally derived hits. This algorithm distinctively introduces random point mutations into the SELFIES representations [34–36] of the molecules, thereby generating novel compounds that maintain a structural resemblance to the starting point. The result-

ing derivatives were filtered based on synthesizability (see Methods A 1 a and A 1 b for details), culminating in the addition of 850 thousand molecules to our training set.


Tartarus Benchmarking

(A)

	1SYH			6Y2F			4LDE		
	ΔE_{1SYH}		SR	ΔE_{6Y2F}		SR	ΔE_{4LDE}		SR
	QUICKVINA2	SMINA		QUICKVINA2	SMINA		QUICKVINA2	SMINA	
NATIVE DOCKING	-10.2	-10.5	100.0%	-4.9	-5.3	0.0%	-11.6	-12.1	100.0%
DATASET	-9.9	-10.2	100.0%	-8.2	-8.2	100.0%	-12.2	-13.1	100.0%
SMILES-VAE	-10.0 \pm 0.7	-10.4 \pm 0.6	12.3%	-8.3 \pm 0.6	-8.9 \pm 0.8	13.1%	-10.7 \pm 0.2	-11.1 \pm 0.4	12.6%
SELFIES-VAE	-10.4 \pm 0.3	-10.9 \pm 0.3	34.8%	-9.6 \pm 0.5	-10.1 \pm 0.4	38.9%	-11.1 \pm 0.4	-11.9 \pm 0.2	38.9%
MoFlow	-10.6 \pm 0.4	-11.0 \pm 0.3	35.5%	-9.8 \pm 0.4	-10.6 \pm 0.3	35.9%	-12.1 \pm 0.4	-13.0 \pm 0.3	36.2%
SMILES-LSTM-HC	-10.6 \pm 0.6	-11.1 \pm 0.4	71.7%	-10.0 \pm 0.6	-10.4 \pm 0.7	70.1%	-12.4 \pm 0.3	-13.3 \pm 0.4	73.9%
SELFIES-LSTM-HC	-10.8 \pm 0.5	-11.2 \pm 0.2	73.2%	-10.4 \pm 0.4	-10.6 \pm 0.6	71.5%	-12.7 \pm 0.3	-13.6 \pm 0.5	75.6%
REINVENT	-11.8 \pm 0.4	-12.1 \pm 0.2	77.8%	-11.1 \pm 0.3	-11.4 \pm 0.3	76.8%	-12.8 \pm 0.2	-13.7 \pm 0.5	76.8%
GB-GA	-11.6 \pm 0.5	-12.0 \pm 0.2	72.6%	-10.9 \pm 0.2	-11.0 \pm 0.2	73.9%	-12.9 \pm 0.1	-13.8 \pm 0.4	71.4%
JANUS	-11.7 \pm 0.4	-11.9 \pm 0.2	68.4%	-11.3 \pm 0.3	-11.9 \pm 0.4	70.4%	-12.8 \pm 0.2	-13.6 \pm 0.5	65.3%
QCBM-LSTM	-10.2 \pm 0.3	-10.4 \pm 0.3	90.85%	-8.6 \pm 0.4	-9.1 \pm 0.3	88.19%	-13.0 \pm 0.2	-12.9 \pm 0.2	79.2%

Classical vs. Quantum Modelling

(B)



(C)

Model	SR % local	QV 2.0	DF %	UF %	SR % Ch42	R Ch42	SA Ch42	PLI Ch42
QCBM-LSTM (SIM)	63.56 \pm 0.31	-12.7 \pm 0.35	86.36	99.66	9.95	2.54	5.69	-11.64
QCBM-LSTM (HW)	62.62 \pm 0.50	-12.8 \pm 0.28	86.94	99.96	9.41	2.55	5.75	-11.75
MQCBM-LSTM (SIM)	60.94 \pm 0.38	-12.1 \pm 0.22	86.61	99.74	9.85	2.54	5.70	-11.95
LSTM	52.3 \pm 0.50	-12.2 \pm 0.97	87.70	99.95	8.24	2.57	5.75	-11.73

(D)

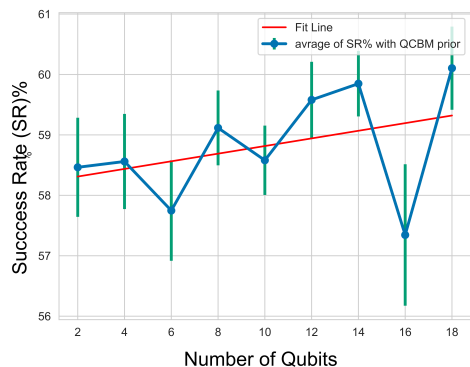


FIG. 3. **Comparative Benchmarking of Quantum and Classical Ligand Design Methods.** (A) Evaluation of the proposed model against classical counterparts using the Tartarus benchmark suite [27] for ligand design across three protein targets: 1SYH, 6Y2F, 4LDE, with models trained on a subset of the DTP Open Compound Collection. Displayed metrics show both the average and the variability (mean \pm standard deviation) of the optimal objective values for the targets, compiled from five individual experiments. 'Dataset' refers to the molecule with the highest performance in the training dataset, whereas 'Native Docking' indicates the initial ligands within their crystallographic structures. The notation ΔE_X signifies the docking score relative to the protein target designated by X. SR stands for the success ratio, indicating the percentage of molecules that meet the predefined structural benchmarks. (B) Comparative analysis of our hybrid approaches with varied priors. The performance of the Quantum Circuit Born Machine (QCBM) was assessed using both a quantum simulator (Sim) and a hardware backend (HW), and contrasted with a Multi-bases QCBM (MQCBM) operating solely on a quantum simulator (SIM), as well as an LSTM model devoid of quantum priors (representing a fully classical architecture). We calculated the number of generated molecules that met a series of synthesizability and stability criteria as stipulated by the Tartarus benchmarking platform (referred to as Local Filters) and by Chemistry42 (referred to as Chemistry42 Filters). (C) Comparative analysis of prior sampling techniques in producing high-docking molecules, as assessed by QuickVina2 and Chemistry42. This comparison delineates the Success Rate (SR %) of molecules meeting Tartarus filter criteria, the uniqueness of generated ligands (Unique Fraction, UF %), and the Structural Diversity Fraction (DF %) of the generated ligands, across various methods. the success rates (SR% Ch42) of molecules meeting Chemistry42's filter criteria, the top reward values (R Ch42) assigned to molecules by Chemistry42, the synthetic accessibility score (SA Ch42) of drug-like molecules, and the highest PLI Ch42 scores found in the generation. The PLI score is measured in kcal/mol, with more negative values indicating better scores. (D) Success rate of generating molecules that meet Tartarus's filter criteria as a function of the number of qubits used in modeling priors for the QCBM.

(2) Generation of New Molecules: Our approach is structured around the integration of three primary components: a) the Quantum Circuit Born Machine (QCBM), b) the classical Long Short Term Memory (LSTM) model, and c) Chemistry42 for AI-driven validation, as shown in Figure 2. The QCBM generator [17] employs a 16-qubit IBM quantum processor, utilizing quantum circuits to model complex data distributions. The integration method of quantum priors into the LSTM architecture, as shown in Figure 2B, involves merging molecular information encoded in SELFIES and quantum data by addition or concatenation to form samples, $X'(t)$, which are then input into the LSTM cell. The quantum component, depicted in Figure 2C, is a QCBM that generates samples from quantum hardware each training epoch and is trained with a reward value, $P(x) = \text{Softmax}(R(x))$, calculated using Chemistry42 or a local filter. This cyclical sampling, training, and validation process forms a loop aimed at continually improving the generated molecular structures for targeting KRAS.

(3) Experimental validation: The process of selecting experimental sample candidates is illustrated in Figure 2D. After training our model, we sampled 1 million compounds from each prior model listed in Figure 2D. These samples underwent evaluation by Chemistry42, filtering out unsuitable compounds for pharmacological purposes and ranking the remaining compounds by their docking score (PLI score). Subsequently, 15 novel compounds were selected for synthesis and underwent Surface Plasmon Resonance (SPR) and cell-based assay experiments.

Before initiating our campaign to design new KRAS inhibitors, we sought to compare our hybrid quantum-classical approach with established classical algorithms. Two principal questions guided our evaluation: First, does integrating quantum methodologies contribute to generating high-quality molecules featuring potentially strong target-specific docking scores? To address this, we evaluated our approach against the Tartarus benchmarking suite [27], specifically designed for drug discovery tasks on three distinct protein targets. Second, we examined the influence of the quantum prior, replacing it with a classical counterpart and increasing the number of qubits in our quantum model to determine whether there is a proportional relationship between the number of qubits and the quality of the generated molecules.

A. Computational Benchmarks - Classical vs Quantum Models

1. Tartarus Benchmark

We employed the Tartarus platform [27] to benchmark our proposed QCBM-LSTM methodology against an

array of classical state-of-the-art models, including REINVENT [37], SMILES-VAE [38], SELFIES-VAE [36], MoFlow [39], SMILES-LSTM-HC [40, 41], SELFIES-LSTM-HC, GB-GA [42], and JANUS [43]. The study focused on three protein targets selected from the Tartarus dataset: 1SYH, an ionotropic glutamate receptor associated with neurological and psychiatric disorders such as Alzheimer’s, Parkinson’s, and epilepsy [44]; 6Y2F, the main protease of the SARS-CoV-2 virus, crucial for its RNA translation [45]; and 4LDE, the β_2 -adrenoceptor GPCR, a cell membrane-spanning receptor that binds to adrenaline, a hormone implicated in muscle relaxation and bronchodilation [46]. For each target, we had a dual objective: to generate novel molecules that exhibit strong binding affinity to the specified proteins, as determined by active sites assigned by Tartarus, and to minimize the docking score using QuickVina2 [47]. Additionally, these molecules were required to pass a comprehensive set of filters designed to eliminate reactive, unsynthesizable, or unstable groups, thereby streamlining the drug discovery process. The top-performing molecules, post-filtering, were subject to a refined re-scoring using a more precise scoring function provided by SMINA [48], at an increased level of exhaustiveness.

We conducted experiments utilizing the QCBM with 16 qubits as a quantum prior and the LSTM as a classical model. The local filter from the Tartarus paper served as the reward function to train the QCBM. As recommended by Tartarus, our models were trained on a subset of 150,000 molecules from the Developmental Therapeutics Program (DTP) Open Compound Collection [49, 50], referred to as DATASET in Figure 3, Table (A). Notably, all 150,000 structures underwent a rigorous screening process using structural filters to eliminate reactive, unsynthesizable, or unstable groups. As such, generative models adept at capturing the distribution of the provided molecule set would exhibit a correspondingly high success rate in generating novel molecules without structural violations. Our observations indicate that only a few generative models demonstrate a high success rate. However, QCBM-LSTM model is very strong in producing a substantial number of high-quality samples that successfully meet the filter criteria, as evidenced by the elevated success rate (SR) depicted in Figure 3, Table (A). Consequently, we believe that the incorporation of a quantum prior leads to improved distribution matching. We further benchmark the influence of a classical/quantum prior in the subsequent section. Moreover, our analysis reveals that for the 4LDE target, our model generates the highest-scoring molecules relative to other generative models. While the docking scores for the remaining two targets are not as high as those produced by classical algorithms, we speculate that incorporating a docking-score-based reward, in conjunction with the filter success rate, could potentially improve our results.

2. Benchmarking of Prior distributions

To evaluate the impact of prior selection on the quality of the molecules generated by our model, we trained four distinct model variants, each incorporating different priors (refer to Figure 3(B)). Specifically, we examined a Quantum Circuit Born Machine (QCBM) prior and implemented it on both a quantum simulator (Sim) and a hardware backend (HW), in contrast with a Multi-bases QCBM (MQCBM) operating exclusively on a quantum simulator (Sim), and a classical LSTM model devoid of quantum priors. These models were tasked with designing KRAS inhibitors, utilizing a meticulously curated dataset of over one million molecules (see Figure 1). Figure 3(B) showcases the optimal results obtained following a comprehensive optimization of the corresponding architectures using Optuna [51] (detailed in Methods A3). We assessed the quality of the generated molecules employing two distinct sets of criteria: one derived from Tartarus [27], termed the "Local Filter," and a more stringent set provided by Chemistry42, termed the "Chemistry42 Filter." In both assessments, we observed that incorporating a quantum prior enhances the success rate, as gauged by the proportion of molecules satisfying the criteria set by the two filters. Furthermore, utilizing the top model from each prior category, we sampled 5,000 molecules that successfully met the filter criteria and examined their respective docking scores (as presented in Figure 3, Table (C)). Intriguingly, these molecules displayed comparably high docking scores as determined by QuickVina2 (denoted as QV 2.0 in the Table) and the protein-ligand interaction score (PLI), as evaluated by Chemistry42 (noted as PLI Ch42 in the Table). Additionally, the synthesized molecules demonstrated consistent metrics across various parameters, including Diversity Fraction (DF%), Uniqueness Fraction (UF%), Chemistry42 Reward (CH42 R), and the Chemistry42 Synthetic Accessibility Score (Chemistry42 SA) [26].

Encouraged by our observation that quantum priors enhance molecule quality, we further investigated the influence of the number of qubits used in modelling priors on the quality of generated molecules, as shown in Figure 3(D). Specifically, we analyzed the percentage of 5,000 uniquely generated random molecules that satisfied a series of local filters. Interestingly, our findings reveal that the success rate correlates roughly linearly with the number of qubits employed in modelling the prior, indicating a direct relationship between the complexity of the quantum model and the effectiveness in generating high-quality molecules. This trend underscores the potential of increasing qubit numbers in quantum models to improve molecular design outcomes systematically.

B. KRAS Inhibitor Design Campaign

A critical aspect of our work involved validating the most promising compounds identified by our hybrid quantum-classical model, specifically targeting KRAS proteins. This validation process is crucial for substantiating our computational predictions and bridging the gap between theoretical models and their practical, real-world applications. Through the testing of these selected compounds in a campaign aimed at KRAS proteins, we aim to showcase the potential practicality of our proposed generative model, thereby emphasizing the real-world relevance of our findings in the field of drug discovery.

1. Chemistry42 Post-Screening and Selection of Promising Candidate Structures for Synthesis

Our study employed the Chemistry42 platform's structure-based drug design (SBDD) screening workflow to evaluate the generated structures and select promising ligand structures for KRAS inhibition. This comprehensive workflow includes a series of filters and scoring modules designed to efficiently identify structures with favourable drug-likeness, synthetic accessibility, and target interactions, as detailed in several steps (see Figure 2D).

The initial phase of screening involved the application of various 2D and 3D filters:

- 2D structural filters evaluating simple structural and compositional parameters, including hydrogen bond donors, oxygen atoms, aromatic atom fraction, and rotatable bonds.
- 2D property filters assessing estimated compound properties like molecular weight, lipophilicity, topological polar surface area, and molecular flexibility.
- Medicinal chemistry filters identifying undesirable or problematic structural fragments.
- A synthetic accessibility filter based on the Retrosynthesis-Related (ReRSA) model [52].
- 3D pharmacophore hypothesis derived from the X-ray structure of the complexed inhibitor molecule (PDB: 7EW9) [22], indicating common structural features of known KRAS inhibitors (Figure 6).
- Pocket-ligand interaction (PLI) score, based on molecular docking, to assess binding features of KRAS inhibitors (Figure 6).
- Calculation of an integrated reward value based on the weighted scores.

For the final selection of candidate structures for synthesis, we imposed more stringent criteria to increase the likelihood of identifying compounds with the desired biological activity. These criteria, more restrictive than those used during the screening evaluation, included:

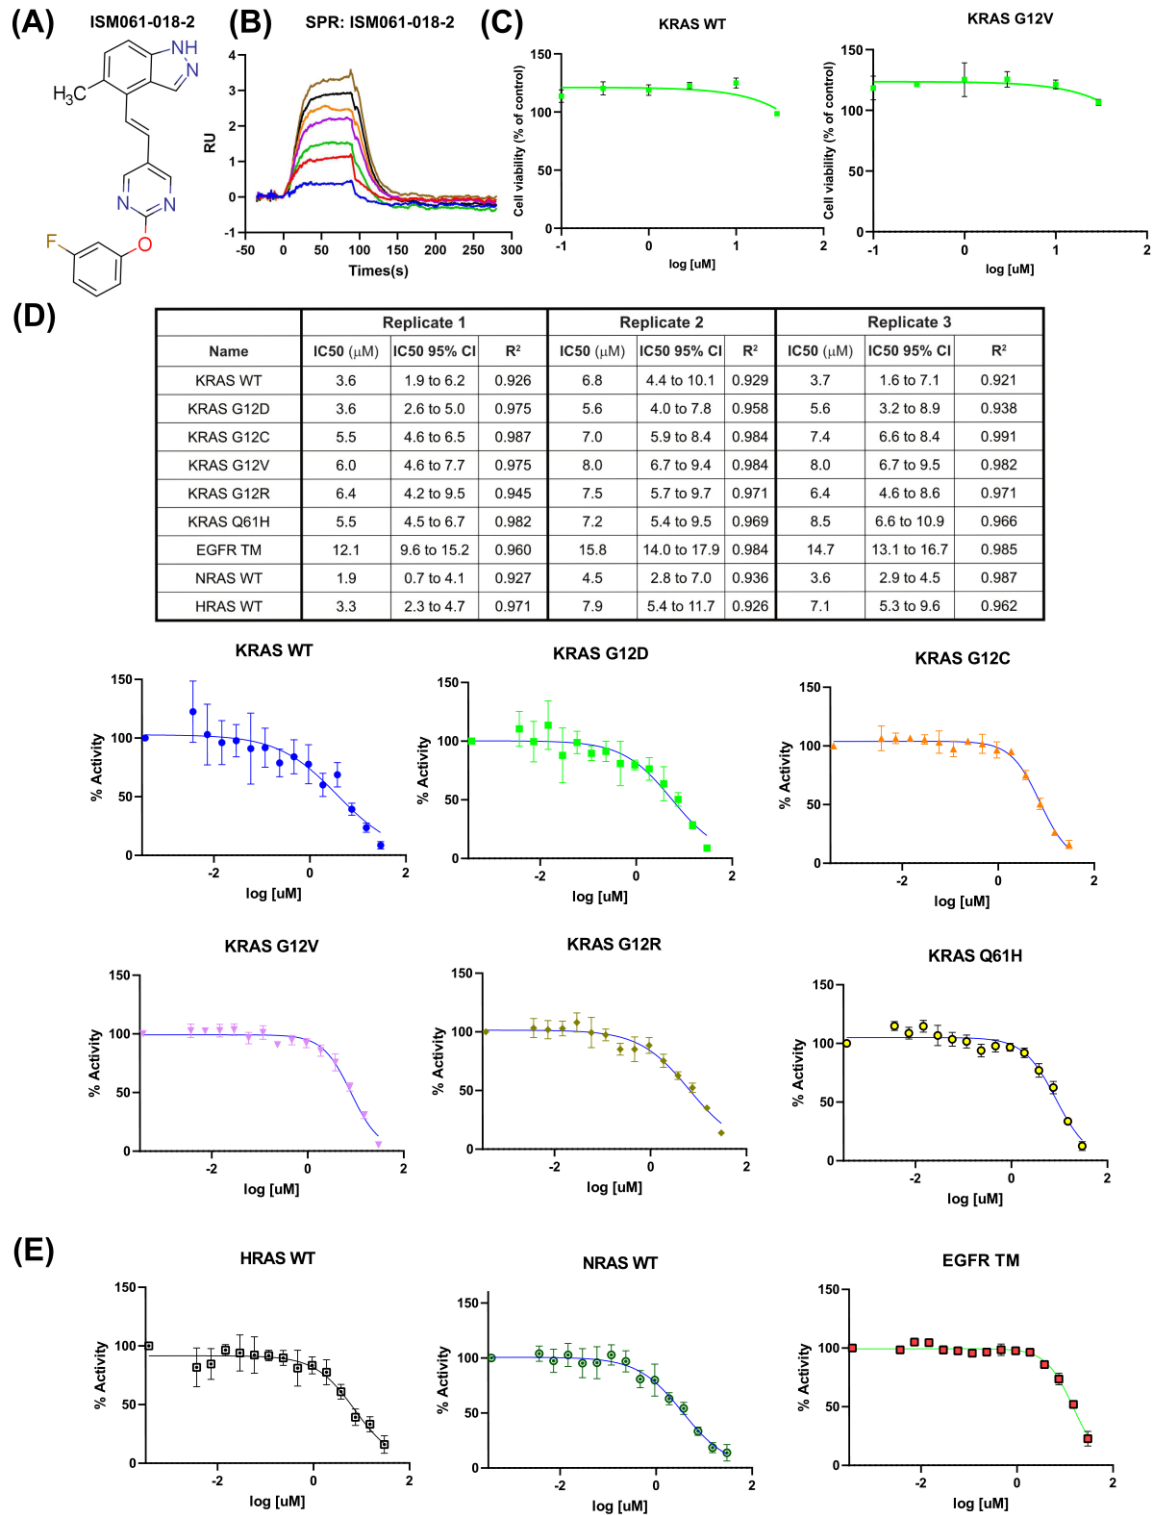


FIG. 4. Pharmacological Characterization of Compound ISM061-018-2 Through Surface Plasmon Resonance and Cellular Activity Assays. (A) Chemical structure of ISM061-018-2. (B) Surface Plasmon Resonance (SPR) sensorgram illustrating the binding kinetics of ISM061-018-2 with various KRAS proteins. (C) Results of Cell-Titer-Glo viability assays illustrating the impact of the compound on cellular proliferation across a concentration range from 123 nM to 30 μ M. Reported values represent the mean of three technical replicates, with standard deviation (S.D.) indicated. (D, Table) A compendium of IC₅₀ values derived from MaMTH-DS dose-response assays, conducted in biological triplicate, evaluating a range of RAS protein baits interactions with the RAF1 prey partner. Investigated RAS members include the wild-type forms of KRAS, HRAS, and NRAS, alongside five oncogenic KRAS mutants of clinical significance. The interaction between EGFR and the SHCI adapter was additionally examined as an off-target control. We provide 95% confidence intervals and R-squared values to verify the accuracy of the curve fitting. (D,E): Dose-response curves from MaMTH-DS assays graphing the modulation of activity of various KRAS proteins, NRAS, HRAS, and EGFR, in response to increasing concentrations of ISM061-018-2 (from 4 nM to 30 μ M), plotted on a logarithmic scale. The curves displayed represent one set from three biological replicates. Each point denotes the mean of three to four technical replicates, with S.D. provided. Curve fitting was executed in GraphPad Prism as delineated in the Methods section. These profiles underscore the compound's differential potency against distinct targets, shedding light on its pharmacological spectrum.

- Successful passage through all Chemistry42 filters.
- An integrated reward value greater than 0.7.
- A Protein-Ligand Interaction (PLI) score less than -8 kcal/mol.
- A pharmacophore match score exceeding 0.7.
- A Synthetic Accessibility (ReRSA) score below 5.

At the stage of final selection and determination of structures for subsequent experimental verification, the resulting sets of molecules at the last step were first clustered by chemical similarity, and 15–20 clusters were determined for each generation. Then, within each cluster, a ranking was carried out based on the ReRSA estimate and the Chemistry42 Protein-Ligand Interaction (PLI) score. Medical chemistry experts analyzed the resulting sets of 100–150 molecules, which were selected based primarily on chemical novelty, structural complexity and potentially undesirable chemical functionality.

2. Experimental Evaluation of Generated Compounds

From the pool of identified structures, we synthesized and characterized 15 compounds. Detailed methodologies of this process are elaborated in the Supplementary Materials¹ (Section S1). We showcased the molecular structures of the two most promising compounds (named ISM061-018-2 and ISM061-22) in Figures 4(A) and 5(A). Each synthesized compound underwent a rigorous two-phase evaluation: their binding affinities were determined using Surface Plasmon Resonance (SPR), and their biological efficacies were gauged through cell-based assays. Notably, the compound ISM061-018-2, engineered through our hybrid quantum model and illustrated in Figure 4, demonstrated a substantial binding affinity to KRAS G12D, registered at $1.4\mu\text{M}$. To delve deeper into this molecule’s effectiveness across a spectrum of KRAS mutations, we commenced an extensive series of tests using a cell-based assay. Specifically, we evaluated the molecule’s performance in a biological context, employing a commercial cell viability assay (CellTiter-Glo; Promega) in conjunction with MaMTH-DS, an advanced split-ubiquitin-based platform for the real-time detection of small molecules targeting specific cellular interactions [53–59].

Further, the biological efficacy of ISM061-018-2 was rigorously tested. Interestingly, it demonstrated no detrimental impact on the viability of HEK293 cells, even when expressing either KRAS WT or KRAS G12V

bait in MaMTH-DS format, and subjected to concentrations as high as $30\mu\text{M}$ for 18–20 hours (Figure 4C). Subsequent testing using MaMTH-DS across a spectrum of cell lines expressing various KRAS “baits” (both WT and five clinically significant oncogenic mutants) in combination with RAF1 “prey” (a recognized KRAS effector) revealed a dose-responsive inhibition of interactions, with IC50s in the micromolar range (Figure 4D). The compound’s activity was not specific to mutants, as it targeted both WT and mutant interactions with similar efficacy. It also showed comparable effectiveness in disrupting the interactions of WT NRAS and HRAS “baits” with RAF1 “prey” (Figure 4D,E). However, it was notably less potent, by a factor of 2–3 (based on IC50 comparisons), against our control interaction comprising EGFR triple mutant “bait” and SHCI adapter “prey”, an interaction crucial in signalling but functioning upstream of RAS in the pathway (Figure 4D,E). These results collectively hint at the potential pan-RAS activity of ISM061-018-2, albeit with indications of some off-target effects.

ISM061-22, as illustrated in Figure 5A, also stood out as a compound of promise, particularly due to its selectivity towards certain KRAS mutants. This compound demonstrated a mild impact on cellular viability at high concentrations after 18–20 hours of exposure, as shown in Figure 5C. Mirroring the performance of ISM061-018-2, ISM061-22 revealed dose-responsive inhibition of the interaction between KRAS “bait” and RAF1 “prey” within the micromolar range. Distinctively, it exhibited a markedly greater effect on mutant KRAS compared to WT, with the degree of influence ranging from 2–8 fold, depending on the specific mutant. Notably, KRAS G12R and Q61H displayed the highest sensitivity, as depicted in Figure 5D,E. Diverging from ISM061-018-2, ISM061-22 did not show binding to KRAS G12D. The compound also demonstrated activity against WT HRAS and NRAS, although it was less potent against HRAS. Furthermore, ISM061-22’s interaction with the EGFR control displayed a unique inhibition profile, stabilizing at 50% maximal activity instead of decreasing to zero, as typically observed in the RAS/RAF interaction curves (Figure 5D,E). This distinct pattern suggests an alternative mode of action for ISM061-22, potentially indicative of partial non-specific activity or a feedback mechanism influencing the EGFR signaling pathway.

In essence, these live-cell experimental observations underscore the robustness of our approach, effectively identifying small molecule candidates with biological activity. This underlines the potential of our methodology to address and surmount the complexities inherent in targeting clinically challenging biomolecules.

¹ <https://github.com/aspuru-guzik-group/quantum-generative-models>

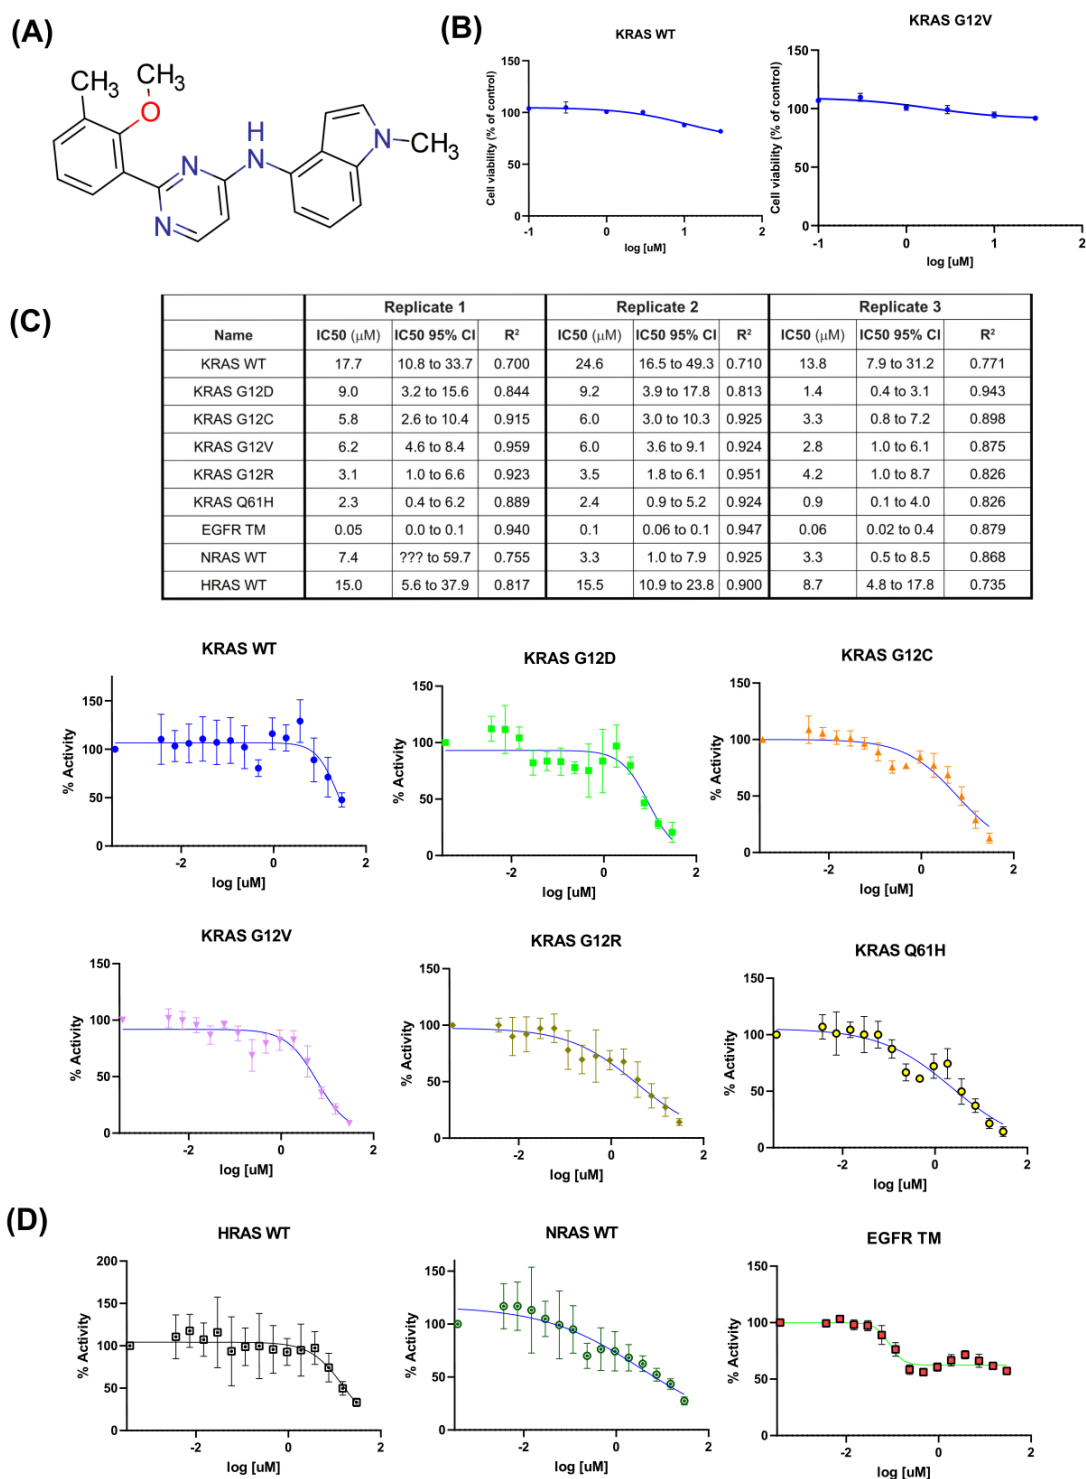


FIG. 5. Pharmacological Evaluation of Compound ISM061-22 Against KRAS Variants and Other Related Proteins. (A) Chemical structure of ISM061-22. (B) Results from Cell-Titer-Glo viability assays depicting the effect of the compound on cellular proliferation over a concentration range from 123 nM to 30 μM . The values represent the mean of three technical replicates, with standard deviation (S.D.) indicated. (C, Table) Summary of IC₅₀ values derived from MaMTH-DS dose-response assays conducted in biological triplicate, against various RAS protein baits' interactions with the RAF1 prey partner. Tested RAS members include wild-type KRAS, HRAS, and NRAS, as well as five clinically significant oncogenic KRAS mutants. The EGFR's interaction with the SHC1 adapter was also examined as an off-target control. The 95% confidence intervals and R-squared values are reported to confirm the precision of the curve fitting. (C,D) Dose-response curves from MaMTH-DS assays, illustrating the modulation of activity of different KRAS proteins, NRAS, HRAS, and EGFR, by increasing concentrations of ISM061-22 (from 4 nM to 30 μM), presented on a logarithmic scale. The curves represent one instance from three biological replicates. Each data point is the average of three to four technical replicates, with S.D. presented. Curve fitting procedures were executed in GraphPad Prism, as outlined in the Methods section. The data emphasizes the compound's nuanced effectiveness against various protein targets, illuminating its potential therapeutic value.

III. CONCLUSION

We introduce a hybrid quantum-classical algorithm, meticulously crafted for near-term quantum computers, aiming to discover novel ligands for specific molecular targets. Our method ingeniously integrates a Quantum Circuit Born Machine (QCBM) as a prior distribution with a classical Long Short-Term Memory (LSTM) network, enhancing this amalgamation with a reward function designed to foster the generation of drug-like small molecules. We benchmarked our approach through computational evaluations across two distinct sets of tasks. Initially, we engaged in a comparative analysis against classical generative models, focusing on the design of potent small molecule binders for three different proteins, employing the Tartarus benchmarking suite. The results highlight our model’s ability to generate high-quality molecules, registering the highest success rate across the tasks, albeit with a marginally lower docking score when juxtaposed with classical approaches. In addition, we examined the impact of the number of qubits on the modelling of the prior distribution and observed a trend that suggests a roughly linear correlation between the modelling success rate of generating high-quality molecules and the number of qubits involved. These findings suggest a potential advantage of the quantum-enhanced model in navigating the synthesizable chemical space with greater efficacy than its traditional classical counterparts while simultaneously yielding molecules that exhibit strong binding affinities.

Subsequently, we deployed our model in a practical scenario to design inhibitors targeting the notoriously challenging cancer protein, KRAS. In this endeavour, we experimentally synthesized and evaluated 15 promising ligands conceived by hybrid and fully classical models. Our empirical findings substantiated the inhibitory properties of two molecules, ISM061-018-2 and ISM061-22, on KRAS, affirming our models’ proficiency in generating viable new ligands for complex drug targets. ISM061-018-2 was identified as a binder to KRAS-G12D, exhibiting a binding affinity of $1.4\mu M$, and it demonstrated pan-KRAS inhibition. Furthermore, ISM061-018-2 exhibited specificity towards certain mutants, showing pronounced activity, particularly against KRAS G12R and Q61H. Both molecules were derived from our hybrid methodology, indicating their superior efficacy to the fully classical model. In addition, both compounds introduced novel chemotypes distinct from existing KRAS inhibitors, illustrating our hybrid model’s capacity to effectively navigate and explore diverse regions within chemical space.

While the results showcased are promising, they do not conclusively establish a ‘quantum advantage’, defined as achieving results beyond the reach of classical methods within a reasonable timeframe. The modest count of

16 qubits in our hybrid algorithm permits simulation on classical platforms, suggesting that state-of-the-art classical algorithms might match or even exceed the efficacy of our quantum-assisted approach. Hence, a critical future research direction involves comprehensively assessing our hybrid quantum-classical algorithm’s performance compared to its classical equivalents. Essential factors for this comparative study include analyzing the scalability relative to qubit quantity, the intricacies of qubit types and their interconnections, the influence of quantum noise and errors, and how the algorithm measures up against top-tier classical algorithms in terms of success rates and other crucial metrics like the docking scores of ligands.

Our research indicates that current near-term quantum hardware can already be harnessed for practical drug discovery applications, mitigating the need to wait for fully fault-tolerant quantum computers, which may be a decade from fruition. Moreover, since our algorithm uses only 16 qubits within the realm of classical simulation, it shows how quantum computing can catalyze the development of more efficient algorithms for classical hardware. In conclusion, we have introduced a hybrid quantum-classical algorithm surpassing its classical performance counterpart. The modest qubit count used, absent any error correction and with limited connectivity, hints at the potential of more advanced quantum computers alongside better quantum-classical algorithms for future drug discovery. With an anticipated increase in qubits, improved fidelity, error correction capabilities, and enhanced connectivity, the prospects for quantum computing in drug discovery are a new frontier of computational and experimental science.

IV. ACKNOWLEDGEMENTS

We want to thank Alejandro Perdomo-Ortiz, Marta Mauri, Brian Dellabetta, Vladimir Vargas-Calderón, Austin Cheng, Jacob Miller, Mohsen Bagherimehrab, and Kouhei Nakaji for their valuable discussion and support in this research. Additionally, we are thankful to the Defense Advanced Research Projects Agency (DARPA) for their funding (grant number HR0011-23-3-0017), which significantly supports our scientific endeavours. A.K.N. acknowledges funding from the Bio-X Stanford Interdisciplinary Graduate Fellowship (SGIF). A.A.-G. thanks Anders G. Frøseth for his generous support. A.A.-G. also acknowledges the generous support of Natural Resources Canada and the Canada 150 Research Chairs program. This research was undertaken thanks in part to funding provided to the Acceleration Consortium of the University of Toronto from the Canada First Research Excellence Fund.

- [1] Elina Petrova. Innovation in the pharmaceutical industry: The process of drug discovery and development. In *Innovation and Marketing in the Pharmaceutical Industry: Emerging Practices, Research, and Policies*, pages 19–81. Springer, 2013.
- [2] Joseph A DiMasi, Henry G Grabowski, and Ronald W Hansen. Innovation in the pharmaceutical industry: new estimates of r&d costs. *Journal of health economics*, 47:20–33, 2016.
- [3] Duxin Sun, Wei Gao, Hongxiang Hu, and Simon Zhou. Why 90% of clinical drug development fails and how to improve it? *Acta Pharmaceutica Sinica B*, 12(7):3049–3062, 2022.
- [4] James P Hughes, Stephen Rees, S Barrett Kalindjian, and Karen L Philpott. Principles of early drug discovery. *British journal of pharmacology*, 162(6):1239–1249, 2011.
- [5] Christoph Gorgulla, Andras Boeszoermyeni, 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, et al. An open-source drug discovery platform enables ultra-large virtual screens. *Nature*, 580(7805):663–668, 2020.
- [6] Christoph Gorgulla, AkshatKumar Nigam, Matt Koop, Suleyman S Cinaroglu, Christopher Secker, Mohammad Haddadnia, Abhishek Kumar, Yehor Malets, Alexander Hasson, Roni Levin-Konigsberg, et al. Virtualflow 2.0-the next generation drug discovery platform enabling adaptive screens of 69 billion molecules. *bioRxiv*, pages 2023–04, 2023.
- [7] AkshatKumar Nigam, Matthew FD Hurley, Fengling Li, Eva Konkořová, 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.
- [8] Steven M Paul, Daniel S Mytelka, Christopher T Dunwiddie, Charles C Persinger, Bernard H Munos, Stacy R Lindborg, and Aaron L Schacht. How to improve r&d productivity: the pharmaceutical industry’s grand challenge. *Nature reviews Drug discovery*, 9(3):203–214, 2010.
- [9] Alex Zhavoronkov, Yan A Ivanenkov, Alex Aliper, Mark S Veselov, Vladimir A Aladinskiy, Anastasiya V Aladinskaya, Victor A Terentiev, Daniil A Polykovskiy, Maksim D Kuznetsov, Arip Asadulaev, et al. Deep learning enables rapid identification of potent ddr1 kinase inhibitors. *Nature biotechnology*, 37(9):1038–1040, 2019.
- [10] Jonathan M Stokes, Kevin Yang, Kyle Swanson, Wengong Jin, Andres Cubillos-Ruiz, Nina M Donghia, Craig R MacNair, Shawn French, Lindsey A Carfrae, Zohar Bloom-Ackermann, et al. A deep learning approach to antibiotic discovery. *Cell*, 180(4):688–702, 2020.
- [11] Feng Ren, Xiao Ding, Min Zheng, Mikhail Korzinkin, Xin Cai, Wei Zhu, Alexey Mantsyzov, Alex Aliper, Vladimir Aladinskiy, Zhongying Cao, et al. Alphafold accelerates artificial intelligence powered drug discovery: efficient discovery of a novel cdk20 small molecule inhibitor. *Chemical Science*, 14(6):1443–1452, 2023.
- [12] AkshatKumar Nigam, Robert Pollice, Pascal Friederich, and Alán Aspuru-Guzik. Artificial design of organic emitters via a genetic algorithm enhanced by a deep neural network. *Chemical Science*, 2024.
- [13] Benjamin Sanchez-Lengeling and Alán Aspuru-Guzik. Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400):360–365, 2018.
- [14] Robert Pollice, Gabriel dos Passos Gomes, Matteo Aldeghi, Riley J Hickman, Mario Krenn, Cyrille Lavigne, Michael Lindner-D’Addario, AkshatKumar Nigam, Cher Tian Ser, Zhenpeng Yao, et al. Data-driven strategies for accelerated materials design. *Accounts of Chemical Research*, 54(4):849–860, 2021.
- [15] AkshatKumar Nigam, Robert Pollice, Matthew FD Hurley, Riley J Hickman, Matteo Aldeghi, Naruki Yoshikawa, Seyone Chithrananda, Vincent A Voelz, and Alán Aspuru-Guzik. Assigning confidence to molecular property prediction. *Expert opinion on drug discovery*, 16(9):1009–1023, 2021.
- [16] 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.
- [17] Mohamed Hibat-Allah, Marta Mauri, Juan Carrasquilla, and Alejandro Perdomo-Ortiz. A Framework for Demonstrating Practical Quantum Advantage: Racing Quantum against Classical Generative Models, March 2023. arXiv:2303.15626 [cond-mat, physics:quant-ph].
- [18] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature communications*, 9(1):4812, 2018.
- [19] Kaitlin Gili, Mohamed Hibat-Allah, Marta Mauri, Chris Ballance, and Alejandro Perdomo-Ortiz. Do quantum circuit Born machines generalize? *Quantum Science and Technology*, 8(3):035021, July 2023.
- [20] Manuel S. Rudolph, Ntwali Bashige Toussaint, Amara Katarwa, Sonika Johri, Borja Peropadre, and Alejandro Perdomo-Ortiz. Generation of High-Resolution Handwritten Digits with an Ion-Trap Quantum Computer. *Physical Review X*, 12(3):031010, July 2022. Publisher: American Physical Society.
- [21] Qing-Wei Zeng, Hong-Ying Ge, Chen Gong, and Nan-Run Zhou. Conditional quantum circuit Born machine based on a hybrid quantum-classical framework. *Physica A: Statistical Mechanics and its Applications*, 618:128693, May 2023.
- [22] RCSB Protein Data Bank. RCSB PDB - 7EW9: GDP-bound KRAS G12D in complex with TH-Z816.
- [23] Xiaolun Wang, Shelley Allen, James F. Blake, Vickie Bowcut, David M. Briere, Andrew Calinisan, Joshua R. Dahlke, Jay B. Fell, John P. Fischer, Robin J. Gunn, Jill Hallin, Jade Laguer, J. David Lawson, James Medwid, Brad Newhouse, Phong Nguyen, Jacob M. O’Leary, Peter Olson, Spencer Pajk, Lisa Rahbaek, Mareli Rodriguez, Christopher R. Smith, Tony P. Tang, Nicole C. Thomas, Darin Vanderpool, Guy P. Vigers, James G. Christensen, and Matthew A. Marx. Identification of MRTX1133, a Noncovalent, Potent, and Selective KRASG12D Inhibitor. *Journal of Medicinal Chemistry*, 65(4):3123–3133, February 2022. Publisher: American Chemical Society.

- [24] Zhongwei Mao, Hongying Xiao, Panpan Shen, Yu Yang, Jing Xue, Yunyun Yang, Yanguo Shang, Lilan Zhang, Xin Li, Yuying Zhang, Yanan Du, Chun-Chi Chen, Rey-Ting Guo, and Yonghui Zhang. KRAS(G12D) can be targeted by potent inhibitors via formation of salt bridge. *Cell Discovery*, 8(1):1–14, January 2022. Number: 1 Publisher: Nature Publishing Group.
- [25] AkshatKumar Nigam, Robert Pollice, Mario Krenn, Gabriel dos Passos Gomes, and Alan Aspuru-Guzik. Beyond Generative Models: Superfast Traversal, Optimization, Novelty, Exploration and Discovery (STONED) Algorithm for Molecules using SELFIES, January 2021.
- [26] Yan A. Ivanenkov, Daniil Polykovskiy, Dmitry Bezrukov, Bogdan Zagribelnyy, Vladimir Aladinskiy, Petrina Kamyra, Alex Aliper, Feng Ren, and Alex Zhavoronkov. Chemistry42: An AI-Driven Platform for Molecular Design and Optimization. *Journal of Chemical Information and Modeling*, 63(3):695–701, February 2023. Publisher: American Chemical Society.
- [27] AkshatKumar Nigam, Robert Pollice, Gary Tom, Kjell Jorner, Luca A. Thiede, Anshul Kundaje, and Alan Aspuru-Guzik. Tartarus: A Benchmarking Platform for Realistic And Practical Inverse Molecular Design, July 2023. arXiv:2209.12487 [cs].
- [28] Albert K Kwan, Gary A Piazza, Adam B Keeton, and Caio A Leite. The path to the clinic: a comprehensive review on direct kras12c inhibitors. *Journal of Experimental & Clinical Cancer Research*, 41(1):1–23, 2022.
- [29] Kaushal Parikh, Giuseppe Banna, Stephen V Liu, Alex Friedlaender, Aakash Desai, Vivek Subbiah, and Alfredo Addeo. Drugging kras: current perspectives and state-of-art review. *Journal of Hematology & Oncology*, 15(1):152, 2022.
- [30] Tarapong Srisongkram, Patcharapa Khamtang, and Natthida Weerapreeyakul. Prediction of kras12c inhibitors using conjoint fingerprint and machine learning-based qsar models. *Journal of Molecular Graphics and Modelling*, 122:108466, 2023.
- [31] Misako Nagasaka, Yiwei Li, Ammar Sukari, Sai-Hong Ignatius Ou, Mohammed Najeeb Al-Hallak, and Asfar S Azmi. Kras g12c game of thrones, which direct kras inhibitor will claim the iron throne? *Cancer treatment reviews*, 84:101974, 2020.
- [32] 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), 2020.
- [33] AkshatKumar Nigam, Robert Pollice, Mario Krenn, Gabriel dos Passos Gomes, and Alan Aspuru-Guzik. Beyond generative models: superfast traversal, optimization, novelty, exploration and discovery (stoned) algorithm for molecules using selfies. *Chemical science*, 12(20):7079–7090, 2021.
- [34] 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, 2020.
- [35] Alston Lo, Robert Pollice, AkshatKumar Nigam, Andrew D White, Mario Krenn, and Alán Aspuru-Guzik. Recent advances in the self-referencing embedding strings (selfies) library. *arXiv preprint arXiv:2302.03620*, 2023.
- [36] 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), 2022.
- [37] Marcus Olivecrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics*, 9(1):1–14, 2017.
- [38] Rafael Gómez-Bombarelli, Jennifer N. Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D. Hirzel, Ryan P. Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2):268–276, Jan 2018.
- [39] Chengxi Zang and Fei Wang. Moflow: an invertible flow model for generating molecular graphs. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 617–626, 2020.
- [40] Marwin HS Segler, Thierry Kogej, Christian Tyrchan, and Mark P Waller. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS central science*, 4(1):120–131, 2018.
- [41] Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher. Guacamol: benchmarking models for de novo molecular design. *Journal of chemical information and modeling*, 59(3):1096–1108, 2019.
- [42] Jan H Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. *Chemical science*, 10(12):3567–3572, 2019.
- [43] AkshatKumar Nigam, Robert Pollice, and Alán Aspuru-Guzik. Parallel tempered genetic algorithm guided by deep neural networks for inverse molecular design. *Digital Discovery*, 1(4):390–404, 2022.
- [44] Anne Frandsen, Darryl S. Pickering, Bente Vestergaard, Christina Kasper, Bettina Bryde Nielsen, Jeremy R. Greenwood, Giuseppe Campiani, Caterina Fattorusso, Michael Gajhede, Arne Schousboe, and Jette Sandholm Kastруп. Tyr702 Is an Important Determinant of Agonist Binding and Domain Closure of the Ligand-Binding Core of GluR2. *Molecular Pharmacology*, 67(3):703–713, March 2005. Publisher: American Society for Pharmacology and Experimental Therapeutics Section: ORIGINAL ARTICLE.
- [45] Linlin Zhang, Daizong Lin, Xinyuanyuan Sun, Ute Curth, Christian Drostén, Lucie Sauerhering, Stephan Becker, Katharina Rox, and Rolf Hilgenfeld. Crystal structure of SARS-CoV-2 main protease provides a basis for design of improved -ketoamide inhibitors. *Science*, 368(6489):409–412, April 2020. Publisher: American Association for the Advancement of Science.
- [46] Aaron M. Ring, Aashish Manglik, Andrew C. Kruse, Michael D. Enos, William I. Weis, K. Christopher Garcia, and Brian K. Kobilka. Adrenaline-activated structure of 2-adrenoceptor stabilized by an engineered nanobody. *Nature*, 502(7472):575–579, October 2013. Number: 7472 Publisher: Nature Publishing Group.
- [47] 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.

- [48] 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.
- [49] Johannes H. Voigt, Bruno Bienfait, Shaomeng Wang, and Marc C. Nicklaus. Comparison of the nci open database with seven large chemical structural databases. *Journal of Chemical Information and Computer Sciences*, 41(3):702–712, 2001.
- [50] Wolf-Dietrich Ihlenfeldt, Johannes H. Voigt, Bruno Bienfait, Frank Oellien, and Marc C. Nicklaus. Enhanced cactvs browser of the open nci database. *Journal of Chemical Information and Computer Sciences*, 42(1):46–57, 2002.
- [51] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631, 2019.
- [52] Bogdan ZAGRIBELNYY, Evgeny Olegovich Putin, Sergei Andreevich FEDORCHENKO, Yan A. IVANENKOV, and Aleksandr Zavoronkovs. Retrosynthesis-related synthetic accessibility, November 2021.
- [53] Punit Saraon, Jamie Snider, Yannis Kalaidzidis, Leanne E. Wybenga-Groot, Konstantin Weiss, Ankit Rai, Nikolina Radulovich, Luka Drecun, Nika Vučković, Adriana Vučetić, Victoria Wong, Brigitte Thériault, Nhu-An Pham, Jin H. Park, Alessandro Datti, Jenny Wang, Shivanthy Pathmanathan, Farzaneh Aboualizadeh, Anna Lyakisheva, Zhong Yao, Yuhui Wang, Babu Joseph, Ahmed Aman, Michael F. Moran, Michael Prakesch, Gennady Poda, Richard Marcellus, David Uehling, Miroslav Samaržija, Marko Jakopović, Ming-Sound Tsao, Frances A. Shepherd, Adrian Sacher, Natasha Leighl, Anna Akhmanova, Rima Al-Awar, Marino Zerial, and Igor Stagljjar. A drug discovery platform to identify compounds that inhibit EGFR triple mutants. *Nature Chemical Biology*, 16(5):577–586, May 2020.
- [54] Igor Stagljjar, Chantal Korostensky, Nils Johnsson, and Stephan Te Heesen. A genetic system based on split-ubiquitin for the analysis of interactions between membrane proteins in vivo. *Proceedings of the National Academy of Sciences*, 95(9):5187–5192, 1998.
- [55] Julia Petschnigg, Bella Groisman, Max Kotlyar, Mikko Taipale, Yong Zheng, Christoph F Kurat, Azin Sayad, J Rafael Sierra, Mojca Mattiazzi Usaj, Jamie Snider, et al. The mammalian-membrane two-hybrid assay (mamth) for probing membrane-protein interactions in human cells. *Nature methods*, 11(5):585–592, 2014.
- [56] Punit Saraon, Jamie Snider, Wiebke Schormann, Ankit Rai, Nikolina Radulovich, Maria Sánchez-Osuna, Jasmin Coulombe-Huntington, Caroline Huard, Mohammed Mohammed, Evelyne Lima-Fernandes, et al. Chemical genetics screen identifies copb2 tool compounds that alters er stress response and induces rtk dysregulation in lung cancer cells. *Journal of Molecular Biology*, 433(23):167294, 2021.
- [57] Wiep Scheper, Safia Thaminy, Sanja Kais, Igor Stagljjar, and Karin Romisch. Coordination of n-glycosylation and protein translocation across the endoplasmic reticulum membrane by sss1 protein. *Journal of Biological Chemistry*, 278(39):37998–38003, 2003.
- [58] Abba Benleulmi-Chaachoua, Lina Chen, Kate Sokolina, Victoria Wong, Igor Jurisica, Michel Boris Emerit, Michèle Darmon, Almudena Espin, Igor Stagljjar, Petra Tafelmeyer, et al. Protein interactome mining defines melatonin mt 1 receptors as integral component of presynaptic protein complexes of neurons. *Journal of pineal research*, 60(1):95–108, 2016.
- [59] João P Lopes, Xavier Morató, Carolina Souza, Cindy Pinhal, Nuno J Machado, Paula M Canas, Henrique B Silva, Igor Stagljjar, Jorge Gandía, Víctor Fernández-Dueñas, et al. The role of parkinson’s disease-associated receptor gpr 37 in the hippocampus: functional interplay with the adenosinergic system. *Journal of Neurochemistry*, 134(1):135–146, 2015.
- [60] Christoph Gorgulla, AkshatKumar Nigam, Matt Koop, Süleyman Selim Çınaroğlu, Christopher Secker, Mohammad Haddadnia, Abhishek Kumar, Yehor Malets, Alexander Hasson, Minkai Li, Ming Tang, Roni Levin-Konigsberg, Dmitry Radchenko, Aditya Kumar, Minko Gehev, Pierre-Yves Aquilanti, Henry Gabb, Amr Alhosary, Gerhard Wagner, Alán Aspuru-Guzik, Yurii S. Moroz, Konstantin Fackeldey, and Haribabu Arthanari. VirtualFlow 2.0 - The Next Generation Drug Discovery Platform Enabling Adaptive Screens of 69 Billion Molecules, April 2023. Pages: 2023.04.25.537981 Section: New Results.
- [61] Matthew E Welsch, Anna Kaplan, Jennifer M Chambers, Michael E Stokes, Pieter H Bos, Arie Zask, Yan Zhang, Marta Sanchez-Martin, Michael A Badgley, Christine S Huang, et al. Multivalent small-molecule pan-ras inhibitors. *Cell*, 168(5):878–889, 2017.
- [62] Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, November 1997.
- [63] F.A. Gers, J. Schmidhuber, and F. Cummins. Learning to forget: continual prediction with LSTM. In *1999 Ninth International Conference on Artificial Neural Networks ICANN 99. (Conf. Publ. No. 470)*, volume 2, pages 850–855 vol.2, September 1999. ISSN: 0537-9989.
- [64] F.A. Gers and J. Schmidhuber. Recurrent nets that time and count. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, volume 3, pages 189–194 vol.3, July 2000. ISSN: 1098-7576.
- [65] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization, January 2017. arXiv:1412.6980 [cs].
- [66] P.J. Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, October 1990. Conference Name: Proceedings of the IEEE.
- [67] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, January 2014.




FIG. 6. Depiction of the Pharmacophore Model for KRAS Inhibitors Based on the Co-crystallized Ligand Structure Analyzed with Chemistry42 (PDB: 7EW9 [22]). Essential pharmacophoric elements are delineated: a blue sphere illustrates a ring system that contributes to structural integrity, a green sphere marks a hydrophobic moiety pivotal for binding affinity, and a cyan sphere signifies a hydrogen bond donor, integral for specificity and interaction with the KRAS protein. The protein structure is displayed on the right, with the pharmacophore interactions within the KRAS Switch 2 Binding Site illustrated on the left.

TABLE I. Summary of binding kinetics and affinity parameters for reference compounds interacting with the KRAS G12D mutation. The table presents the on-rate (K_{on}), off-rate (K_{off}), and dissociation constants (K_d) in both kinetic and steady-state measurements, providing a comprehensive overview of the binding dynamics of MRTX1133 and BI-2852 compounds with the target protein.

Compound	$K_{on} (M^{-1}s^{-1})$	$K_{off} (s^{-1})$	Kinetic $K_d (M)$	Steady state $K_d (M)$
MRTX1133	1.26×10^{10}	0.0732	5.83×10^{-12}	—
BI-2852	6.13×10^6	0.0101	1.65×10^{-9}	1.41×10^{-8}

Appendix A: Methods

This section explains the methods and workflow incorporated in our proposed approach, offering a comprehensive understanding of the mechanisms used in our study. Figure 1 illustrates the workflow we employed in our study.

1. Data Acquisition and Pre-processing

Our preliminary dataset, sourced from Insilico Medicine², included approximately 650 data points. These were selectively collated from existing literature, specifically targeting the KRAS G12D mutation (refer to Figure 1). Given the dataset's limited size, we opted to expand it to improve the robustness of our model

² <https://insilico.com/>

TABLE II. **Dissociation constants (K_d) for a series of newly generated compounds, illustrating their binding affinities to the KRAS G12D protein.** Five of the twelve compounds evaluated demonstrated discernible binding affinities, with one notable compound achieving binding affinity in the single-digit micromolar range. Contrasting with the reference inhibitors, these compounds exhibit a ‘fast off’ kinetic profile, which may hold significance for their pharmacodynamic properties. While their affinities are considered weak for immediate therapeutic application, these compounds, synthesized from a small and structurally diverse training set, represent promising scaffolds for further optimization in pursuit of novel KRAS protein inhibitors.

Compound K_d (μM)		Compound K_d (μM)		Compound K_d (μM)		Compound K_d (μM)	
ISM061-6	39.1	ISM061-18	19.2	ISM061-14	N/A	ISM061-22	N/A
ISM061-11	N/A	ISM061-18-2	1.4	ISM061-15	N/A	ISM061-4	N/A
ISM061-13	N/A	ISM061-21	56.9	ISM061-16	N/A	ISM061-24-2	16

Note: N/A – No Activity observed.

during training.

a. STONED-SELFIES

We utilized the STONED-SELFIES [25] algorithm, available at <https://github.com/aspuru-guzik-group/stoned-selfies>, to mine our initial set of 650 molecules. For a given molecule in SMILES format, we first randomized the string using RDKit³. These randomized strings were then converted into SELFIES. Each SELFIES string underwent mutations—in the form of character deletions, replacements, and additions—up to 500 times. Subsequently, the synthesizability and stability of the mutated strings were assessed using Chemistry42. We generated 850k molecules, which served as the training set for our generative models.

b. Virtual Screening Process

VirtualFlow 2.0 [60] was used to identify additional molecules predicted to bind to KRAS G12D. The adaptive target-guided (ATG) method performed the virtual screening in two stages. In the first stage, the ATG pre-screen was performed, in which a spare version of the 69 billion REAL Space from Enamine (version 2022q12) was screened. In the second stage, the most potent tranches of ligands were screened in full, amounting to 100 million ligands. The docking program used was QuickVina 2 [47], with exhaustiveness set to 1 in both stages of the screen. The screen was carried out in the AWS (Amazon Web Services) cloud computing platform. The protein structure used in the screen is the PDB structure 5US4 [61], which was prepared before the virtual screen with Schrödinger’s Protein Preparation Wizard (addition of hydrogens, protonation state prediction). The docking box was of size 14x14x20 Angstrom.

2. Quantum Assisted Algorithm

As Figure 1 shows, Our quantum-assisted model is a hybrid algorithm composed of both quantum and classical generative components. The quantum generative model utilizes a Quantum Circuit Born Machine (QCBM) model, while the classical component utilizes a Long Short Term Memory (LSTM) model. Figure 2 illustrates the flowchart of our proposed generative model. Within this model, we utilized Chemistry42 and a local filter to validate sample generation at each step, which was then employed to train the QCBM model. The QCBM model, a quantum circuit model, was executed on a quantum processing unit. Subsequently, samples from the trained QCBM were fed into the LSTM model, which generated sequences based on these samples. The reward value for each sample was computed at every step using Local filter until epoch 20, and after we selected Chemistry42. This reward value was then used to train our quantum generative model. During the first epoch, no rewards were available, so the algorithm sampled from the untrained QCBM model, designated as X_i . From the second epoch onward, rewards were computed, allowing us to calculate the **Softmax** of the rewards for each X_i , where $i \in [1, N]$. The corresponding pseudocode can be found in Algorithm A.1.

a. QCBM Model

The Quantum Circuit Born Machine (QCBM) model represents a quantum variational generative model, necessitating a classical optimizer to train its parameters. The total count of these parameters is computed with the number of qubits and layers defined in the model (refer to Figure 7). Upon specifying the parameters, denoted as θ_n , we obtain a quantum state $|\psi(\theta)\rangle$. Here, each θ_n exerts an impact on the wave function, expressed as $\psi(\theta)$. To optimize these parameters θ , the model is initially configured with randomly assigned parameters $|\psi(\theta)\rangle$. These parameters are subsequently calculated throughout the training process. The training of the

³ <https://www.rdkit.org/>

QCBM model involves minimizing the Exact Negative Log-Likelihood (Exact NLL) loss function.

b. Classical Model: LSTM Model

Long Short-Term Memory (LSTM) networks (see Figure 1) are employed for a classical part of this architecture. LSTM is simple and has a good record of learning the string pattern in natural language processing for a long time. LSTM networks are specialized Recurrent Neural Networks (RNN) capable of learning long-term dependencies in sequence data [62]. They are particularly useful in applications where the context from earlier parts of the sequence is needed to interpret later parts, such as in natural language processing, time-series forecasting, and more [63]. The LSTM architecture consists of a chain of repeating modules called cells. Each cell contains three gates that control the flow of information:

1. **Forget Gate:** This gate decides what information from the cell state should be thrown away or kept. It takes the output of the previous LSTM cell and the current input and passes them through a sigmoid function, outputting a number between 0 and 1 for each number in the cell state. A 0 means “completely forget this” and a 1 means “completely keep this.”
2. **Input Gate:** This gate updates the cell state with new information. It has two parts: a sigmoid layer called the “input gate layer” and a hyperbolic tangent layer. The sigmoid layer decides what values to update, and the hyperbolic tangent layer creates a vector of new candidate values that could be added to the state.
3. **Output Gate:** This gate decides the next hidden state. The hidden state contains information on previous inputs. The hidden state is used to calculate the output of the LSTM and the next hidden state.

The following equations can describe the LSTM’s operations:

$$\text{Forget Gate: } f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \quad (\text{A1})$$

$$\text{Input Gate: } i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \quad (\text{A2})$$

$$\text{Candidate Values: } \tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \quad (\text{A3})$$

$$\text{Update Cell State: } C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t \quad (\text{A4})$$

$$\text{Output Gate: } o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \quad (\text{A5})$$

$$\text{Update Hidden State: } h_t = o_t \cdot \tanh(C_t) \quad (\text{A6})$$

Where σ is the sigmoid activation function, W and b are the weight matrices and bias vectors for each gate, and x_t is the input at time t [64].

Training an LSTM involves optimizing the network’s weights and biases to minimize a specific loss function.

This is typically accomplished using gradient-based optimization algorithms such as stochastic gradient descent (SGD) or Adam [65]. The backpropagation through time (BPTT) algorithm is employed to compute the gradients relative to the loss function, considering the sequential nature of the data [66]. The networks are trained using the Adam Optimizer with the Negative Log Likelihood Loss function, and to mitigate overfitting, regularization techniques like dropout are implemented [67]. The Negative Log-Likelihood Loss for a single data point is given by:

$$L(y, \hat{y}) = -\log(\hat{y}_y) \quad (\text{A7})$$

where y is the true class label, and \hat{y}_y is the predicted probability for the true class label y .

The loss for a batch of data is the mean of the individual losses for each data point in the batch:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \log(\hat{y}_{y_i}) \quad (\text{A8})$$

Where N is the number of data points in the batch, y_i is the true class label for the i -th data point, and \hat{y}_{y_i} is the predicted probability for the true class label of the i -th data point.

In the hyperparameter tuning process, we utilized Optuna⁴, an optimization framework, to adjust parameters such as the number of hidden dimensions, embedding dimensions, and layers within the model. The model presented in this research integrates a deep learning architecture. This architecture is designed to incorporate prior information (samples) into the generative process. Additionally, the model employs Chemistry42 feedback in conjunction with Quantum Circuit Born Machines (QCBM), aiming to enhance its generative accuracy. Figure 2(B) illustrates the proposed architecture at a cell level. The prior samples are combined with Input samples $x'_t = X(i) \# x_t$ in LSTM cell; This combination has two methods: adding and concatenating samples. LSTM’s operations will be updated with

$$\begin{aligned} \text{Prior Sampling } x'_t &= X(i) \# x_t \\ \text{OR } x'_t &= X(i) + x_t \end{aligned} \quad (\text{A9})$$

$$\text{Forget Gate: } f_t = \sigma(W_f \cdot [h_{t-1}, x'_t] + b_f) \quad (\text{A10})$$

$$\text{Input Gate: } i_t = \sigma(W_i \cdot [h_{t-1}, x'_t] + b_i) \quad (\text{A11})$$

$$\text{Candidate Values: } \tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x'_t] + b_C) \quad (\text{A12})$$

$$\text{Update Cell State: } C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t \quad (\text{A13})$$

$$\text{Output Gate: } o_t = \sigma(W_o \cdot [h_{t-1}, x'_t] + b_o) \quad (\text{A14})$$

$$\text{Update Hidden State: } h_t = o_t \cdot \tanh(C_t) \quad (\text{A15})$$

⁴ <https://optuna.org/>

To generate samples, the process begins with sampling from the prior, followed by the LSTM network processing these prior samples to generate compounds representations. The compounds will be validated through the Chemistry42 platform, specifically tailored to assess ligand quality for the KRAS G12D mutation. This methodology offers designing ligands targeted at specific proteins. Moreover, the LSTM model is a classical approach for learning ligand structures and constructing a latent ligand space. The Quantum Circuit Born Machine (QCBM) functions as a prior, guiding the LSTM in the generation of novel ligand samples. This procedure is subjected to an iterative process to enhance the quality of ligands, which is evaluated using the Chemistry42 platform.

c. Quantum Generative Model: QCBM Model

The QCBM is a variational quantum algorithm that utilizes the foundational principles of quantum mechanics, particularly the Born rule, to generate complex and diverse data samples. The core of our QCBM model is a parameterized quantum state $|\psi(\theta)\rangle$, where θ denotes the parameters, or ansatz, of our quantum circuit. As per the Born rule, given a measurement basis, which is commonly the computational basis in our case, the probability of observing a specific outcome $|x\rangle$ is expressed as $|\langle x|\psi(\theta)\rangle|^2$.

Training a Quantum Circuit Born Machine (QCBM) involves optimizing the parameters of the quantum circuit to produce a probability distribution that closely approximates the target distribution (Probability that is computed by chemistry Reward values. This process is fundamentally iterative, where the quantum circuit parameters, denoted as θ , are adjusted in each step to reduce the discrepancy between the generated and target distributions. A classical optimization algorithm recommends adjusting parameters, which operates based on the feedback received from the evaluation of the circuit’s output. At each iteration, the quantum circuit is sampled to produce a set of states. These states are then compared against the target distribution, and the difference between them informs the direction and magnitude of parameter adjustments in the quantum circuit. This iterative process continues until the distribution generated by the QCBM closely aligns with the target distribution or until a predefined convergence criterion is met.

In the context of QCBM training, the Exact Negative Log-Likelihood (Exact NLL) functions as the primary loss function, providing a quantitative measure of the difference between the distributions. The Exact NLL for a QCBM is the negative sum of the logarithms of the probabilities the quantum circuit assigns to the states in the training dataset. Mathematically, this is represented as $\text{NLL}(\theta) = -\sum_{x \in D} \log p_\theta(x)$, where D is the set of data points, and $p_\theta(x)$ is the probability of observing

state x under the current parameters θ of the quantum circuit. Minimizing the NLL involves adjusting θ such that the quantum circuit’s output distribution increasingly resembles the empirical distribution of the data. This optimization is typically carried out using gradient-based methods or other heuristic techniques suited to the quantum computing context. In our project, we used COBYLA for our optimizer. As the NLL decreases, the fidelity of the QCBM in modeling the target distribution correspondingly increases, indicating successful training of the quantum model.

Figure 7 shows the QCBM architecture and illustrates its associated ansatz. We used linear topology for our project. Our QCBM model is built with 16 qubits and 4 layers, and we had 96 parameters to optimize in total. The initial probability of the samples, $P(X(i))$, is computed based on the rewards returned by the Chemistry42 model. These reward-based probabilities are then passed through a `Softmax` function to ensure they are normalized and fall within the range of 0 to 1. The resulting values serve as the "true" probabilities of the samples and are used as the target values during the model’s training process.

3. Benchmark Setup

Our benchmark uses both classical and quantum hardware. Our classical computational setup is based on a cluster equipped with GPU nodes. This cluster consists of two GPU nodes, each with specific features. These features include two AMD EPYC 7V13 64-core Processors, resulting in a total of 128 CPU cores per node. In addition, each node is equipped with 512 GB of RAM. The nodes also contain eight AMD Instinct™ MI100 GPUs, each with a GPU RAM of 32GB. For the classical training, we utilized four of these GPUs in parallel (i.e., one GPU node). Furthermore, we used a Nvidia GPU (RTX3090Ti) to facilitate our classical quantum simulations. For the quantum hardware setup, we employed the Guadalupe quantum system, equipped with 16 Qubits and a Falcon r4P processor type. Our Quantum Circuit Born Machine (QCBM) model, accompanied by an error correction circuit, was executed on this quantum processor.

Regarding software, we utilized several packages provided by Zapata AI under the Qml core agreements⁵. We implemented our Variational Quantum Circuit and classical LSTM model using the Qml Core Python package. We used the STONED-SELFIES and VirtualFlow 2.0 packages to prepare a diverse dataset. Additionally, we employed RDkit and Insilico APIs to compute the reward value and conduct some post-processing analyses. The QCBM model underwent a training regimen spanning 30

⁵ <https://docs.orquestra.io>

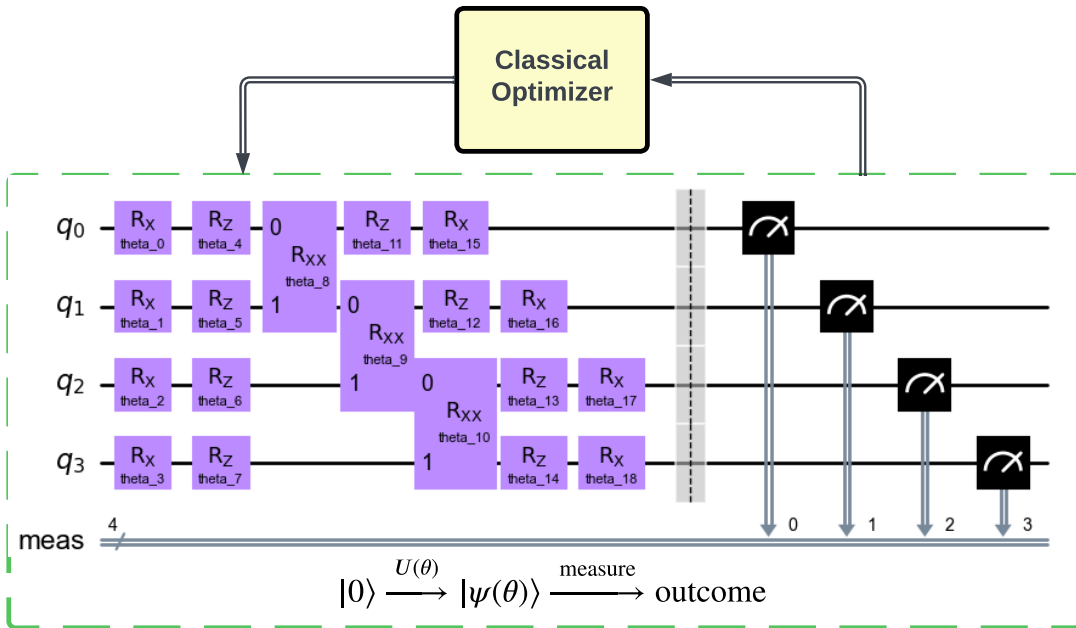


FIG. 7. **Schematic representation of the Quantum Circuit Born Machine (QCBM)** implemented in our numerical experiments, illustrating a variational quantum circuit with a configuration of three layers and four qubits. In practice, our numerical experiments utilized a system with 16 qubits. The depicted quantum gates, including parameterized rotations (Rx, Rz) and entangling CNOT gates, are orchestrated to evolve the initial state $|0\rangle$ into a complex quantum state $|\psi(\theta)\rangle$. The outcome is measured, and the resulting data are used by the classical optimizer to iteratively refine the parameters θ , thus leading the circuit towards an optimal solution for ligand generation.

epochs. In contrast, the LSTM model was trained over a total of 40 epochs.

We utilized the Optuna platform to optimize the hyperparameters in the benchmarking. We ran Optuna tuning for 100 trials for each model, to determine the optimal number of QCBM layers, LSTM layers, and embedding dimensions. Additionally, we tuned the sampling temperature, which defines the balance between determinism and stochasticity in the model, particularly between the prior input and the LSTM output.

4. SPR Conditions

A Biacore 8K system was used for all experiments. For preliminary compound screening, N-terminal biotinylated KRASG12D protein (synthesized by VIVA Biotech (Shanghai) Ltd, purity $\geq 95\%$) was captured on a Sensor Chip SA (GE Healthcare) at a density of about 2000RU. Protein immobilization was done using 1x HBS-EP+, 2mM TCEP, 2% DMSO as a running buffer. Protein was injected for 70s at a flow rate of $5 \mu\text{L}/\text{min}$. The protein concentration was $5 \mu\text{g}/\text{mL}$. We performed an initial screening of compounds prepared samples by serial 2-fold dilutions from $200 \mu\text{M}$ to $0.39 \mu\text{M}$ in 1x HBS-EP+, 2mM TCEP, 2% DMSO. Samples were injected for 60s at a flow rate of $30 \mu\text{L}/\text{min}$ and dissociation time was 180s. A Biacore 8K machine was used to carry out the SPR

experiments and subsequent data analysis.

5. MaMTH-DS Dose-Response Assays

MaMTH-DS FLP HEK293 reporter cell lines[53] stably expressing KRAS (WT or mutant), HRAS, NRAS or EGFR triple mutant L858R/T790M/C797S bait alongside RAF1 (for RAS baits) or SHC1 (for EGFR) preys were seeded into 384-well white-walled, flat-bottomed, tissue-culture treated microplates (Greiner #781098) at a concentration of 100,000 /mL ($50 \mu\text{L}$ total volume/well) in DMEM/10% FBS/1% Pen-Strep media. Seeding was performed using a MultiFlo-FX multi-mode liquid dispenser (BioTek). Plates were left at room temperature for 30-60 minutes following seeding before transfer to a Heracell 150i incubator (Thermo) and growth at $37^\circ\text{C}/5\% \text{CO}_2$ for 3 hours. After growth $10 \mu\text{L}$ of DMEM/10%FBS/1% Pen-Strep supplemented with $3 \mu\text{g}/\text{mL}$ Tetracycline (to induce bait/prey expression; BioShop, TET701) and $60 \text{ng}/\text{mL}$ EGF (to stimulate RAS signaling; Sigma #E9644) was added to each well via multichannel pipette. As appropriate, 6X concentration of drug (or DMSO only) was also included in the media, with all lower concentrations produced via serial dilution starting from the highest concentration solution. Plates were then grown overnight (18 – 20 hours) at $37^\circ\text{C}/5\% \text{CO}_2$. Luciferase assay was performed the next

day using 10 μ L of 20 μ M native coelenterazine substrate (Nanolight #303) per well. Luminescence was measured using a Clariostar plate reader (BMG Labtech) with a Gain of 3200-3800 and a 1 second integration time. All data analysis was performed using Microsoft Excel and GraphPad Prism. Curve fits were performed in Prism (non-linear regression, log(inhibitor) vs. response, variable slope (4 parameters) with bottom constrained to zero, except for EGFR-SHC1 in the presence of ISM061-22, for which no constraint was applied).

6. Cell Viability Assay

MaMTH-DS FLP HEK293 reporter cell lines[53] stably expressing KRAS (WT or G12V mutant) bait alongside RAF1 prey were seeded into 96-well white-walled, μ CLEAR[®] flat-bottomed, tissue-culture treated plates (Greiner #655098) at 40 000 cells/well in DMEM/10% FBS/1% Pen-Strep media (60 μ L total volume/well). Seeding was performed using a MultiFlo-FX multi-mode liquid dispenser (BioTek). Plates were left at room temperature for 30-60 minutes following seeding before

transfer to a Heracell 150i incubator (Thermo) and growth at 37°C / 5% CO₂ for 3 hours. After growth 30 μ L of 3X concentration of the drug (or DMSO only) in DMEM/10%FBS/1% Pen-Strep media were added to wells, with all lower concentrations produced via serial dilution starting from the highest concentration solution (final drug concentration 30 μ M to 123nM). 37°C/5% CO₂. Plates were then grown overnight (18-20 hours) at 37°C/5% CO₂. Effect of the drug on cell viability was assessed by CellTiter-Glo[®] Luminescent Cell Viability Assay from Promega (#G7570). 90 μ L of the CellTiter-Glo[®] reagent were added directly into each well following 30 min equilibration of the plate at room temperature. Contents of the wells were mixed on an orbital shaker for 2 min, and plates were then incubated at room temperature for 10 minutes to stabilize luminescent signal. Luminescence was measured using a Clariostar plate reader (BMG Labtech) with a Gain of 3600 and a 1 second integration time. Values represent the mean \pm S.D. of three replicates for each tested drug concentration. All data analysis was performed using Microsoft Excel and GraphPad Prism.

Algorithm A.1. **overview of Quantum-Assisted Drug Discovery using an LSTM framework.** This pseudocode details the iterative process, starting with the initialization of the LSTM and QCBM models, followed by the generation and validation of new compounds. Valid compounds are subjected to a reward calculation and probability assessment, which in turn inform the subsequent training of the QCBM (detailed in Algorithm A.2). This cycle continues until convergence, illustrating the dynamic interplay between quantum predictions and LSTM-generated compounds underpinned by the Chemistry42 evaluation.

```

1: Initialize: LSTM, QCBM, filter (Chemistry42)
2: Generate initial samples  $X_i$  from QCBM
3: while not converged do
4:   Train LSTM with  $X_i$ 
5:   LSTM generates a new compound from the current samples  $X_i$ 
6:   Validate the new compound with the filter
7:   if new compound is valid then
8:     Compute rewards for the new compound
9:     Compute probabilities  $P(X_i)$  for each new compound
10:    Train QCBM with  $X_i$  and  $P(X_i)$ 
11:    Generate new  $X_i$  from QCBM
12:   end if
13: end while

```

Algorithm A.2. **pseudocode outlining the training regimen for the Quantum Circuit Born Machine (QCBM) model.** This process delineates the iterative optimization of the QCBM parameters.

```

1: Initialize: QCBM model with a certain number of qubits and layers
2: Set: Parameterized quantum state  $|\psi(\theta)\rangle$ 
3: while not converged do
4:   Compute exact negative log-likelihood (exact NLL) loss function
5:   Compute gradient of exact NLL with respect to  $\theta$ 
6:   Adjust parameters using an optimizer
7:   Validate the sample and compute its reward value
8:   if sample is valid then
9:     Compute rewards for the sample
10:    Adjust probabilities  $P(X_i)$  based on the rewards
11:    Train QCBM model with adjusted probabilities
12:   end if
13: end while

```