



Quantum-computing-enhanced algorithm unveils potential KRAS inhibitors

Received: 7 May 2024

Accepted: 6 December 2024

Published online: 22 January 2025

Check for updates

Mohammad Ghazi Vakili^{1,2}, Christoph Gorgulla^{1,3,4}✉, Jamie Snider⁵, Akshat Kumar Nigam⁶✉, Dmitry Bezrukov⁷, Daniel Varoli⁸, Alex Aliper⁷, Daniil Polykovsky⁹, Krishna M. Padmanabha Das^{10,11}, Huel Cox III¹¹, Anna Lyakisheva⁵, Ardalan Hosseini Mansob^{10,12}, Zhong Yao⁵, Lela Bitar^{5,13}, Danielle Tahoulas^{10,14}, Dora Čerina^{10,15}, Eugene Radchenko^{10,7}, Xiao Ding⁷, Jinxin Liu⁷, Fanye Meng⁷, Feng Ren^{10,7}, Yudong Cao¹⁶, Igor Stagljar^{10,5,12,14,17}✉, Alán Aspuru-Guzik^{10,1,2,18,19,20,21}✉ & Alex Zhavoronkov^{10,7}✉

We introduce a quantum–classical generative model for small-molecule design, specifically targeting KRAS inhibitors for cancer therapy. We apply the method to design, select and synthesize 15 proposed molecules that could notably engage with KRAS for cancer therapy, with two holding promise for future development as inhibitors. This work showcases the potential of quantum computing to generate experimentally validated hits that compare favorably against classical models.

Drug discovery is a multifaceted and resource-intensive process encompassing the discovery, development and comprehensive testing of new molecules. Typically extending over a decade and incurring substantial costs, the pharmaceutical industry faces substantial financial risks^{1–3}. The pressing need for efficiency and innovation in drug discovery has led to integrating advanced computational tools into traditional pharmaceutical research methodologies. Concurrently, generative modeling has emerged as a transformative technology in molecule design^{4–7}. Generative models use machine learning techniques to understand the underlying distribution of atoms and bonds in a specified dataset, which are then used to construct molecules with predefined properties, a process known as inverse molecular design^{8–10}. These models excel at exploring the vast drug-like chemical space (~10⁶⁰ molecules), identifying potential molecules efficiently¹¹.

By merging the advancements in quantum machine learning with traditional drug discovery, the industry is shifting toward innovative computational strategies¹². While purely classical algorithms have made notable strides in drug discovery, hybrid classical–quantum approaches offer unique advantages because of the ability of quantum circuit Born machines (QCBMs) to leverage quantum effects such as superposition and entanglement^{13,14}. The introduction of QCBMs illustrates this advancement, offering a generative model that can outperform classical ones in certain aspects. QCBMs are quantum generative models that leverage quantum circuits to learn complex probability distributions, enabling them to generate new samples that resemble the training data¹⁵. Hibat-Allah et al.¹⁶ demonstrated the superior generalization capabilities of QCBMs, which can produce cardinality distributions beyond the training set and overcome

¹Department of Computer Science, University of Toronto, Toronto, Ontario, Canada. ²Department of Chemistry, University of Toronto, Toronto, Ontario, Canada. ³Department of Structural Biology, St. Jude Children's Research Hospital, Memphis, TN, USA. ⁴Department of Physics, Harvard University, Cambridge, MA, USA. ⁵Donnelly Centre, Temerty Faculty of Medicine, University of Toronto, Toronto, Ontario, Canada. ⁶Department of Computer Science, Stanford University, Stanford, CA, USA. ⁷Insilico Medicine AI Limited, Abu Dhabi, UAE. ⁸Zapata AI, Boston, MA, USA. ⁹Insilico Medicine Canada, Inc., Montreal, Québec, Canada. ¹⁰Department of Cancer Biology, Dana-Farber Cancer Institute, Boston, MA, USA. ¹¹Department of Biological Chemistry and Molecular Pharmacology, Harvard Medical School, Harvard University, Boston, MA, USA. ¹²Department of Molecular Genetics, University of Toronto, Toronto, Ontario, Canada. ¹³Department for Lung Diseases Jordanovac, Clinical Hospital Centre Zagreb, University of Zagreb, Zagreb, Croatia. ¹⁴Donnelly Centre, Department of Biochemistry, Temerty Faculty of Medicine, University of Toronto, Toronto, Ontario, Canada. ¹⁵Department of Oncology, University Hospital Center Split, School of Medicine, University of Split, Split, Croatia. ¹⁶AQI, Inc., Boston, MA, USA. ¹⁷Mediterranean Institute for Life Sciences (MedILS), School of Medicine, University of Split, Split, Croatia. ¹⁸Department of Chemical Engineering and Applied Chemistry, University of Toronto, Toronto, Ontario, Canada. ¹⁹Department of Materials Science and Engineering, University of Toronto, Toronto, Ontario, Canada. ²⁰Vector Institute for Artificial Intelligence, Toronto, Ontario, Canada. ²¹Fellow, Canadian Institute for Advanced Research (CIFAR), Toronto, Ontario, Canada.

✉ e-mail: christoph.gorgulla@stjude.org; akshat98@stanford.edu; igor.stagljar@utoronto.ca; aspuru@utoronto.ca; alex@insilicomedicine.com

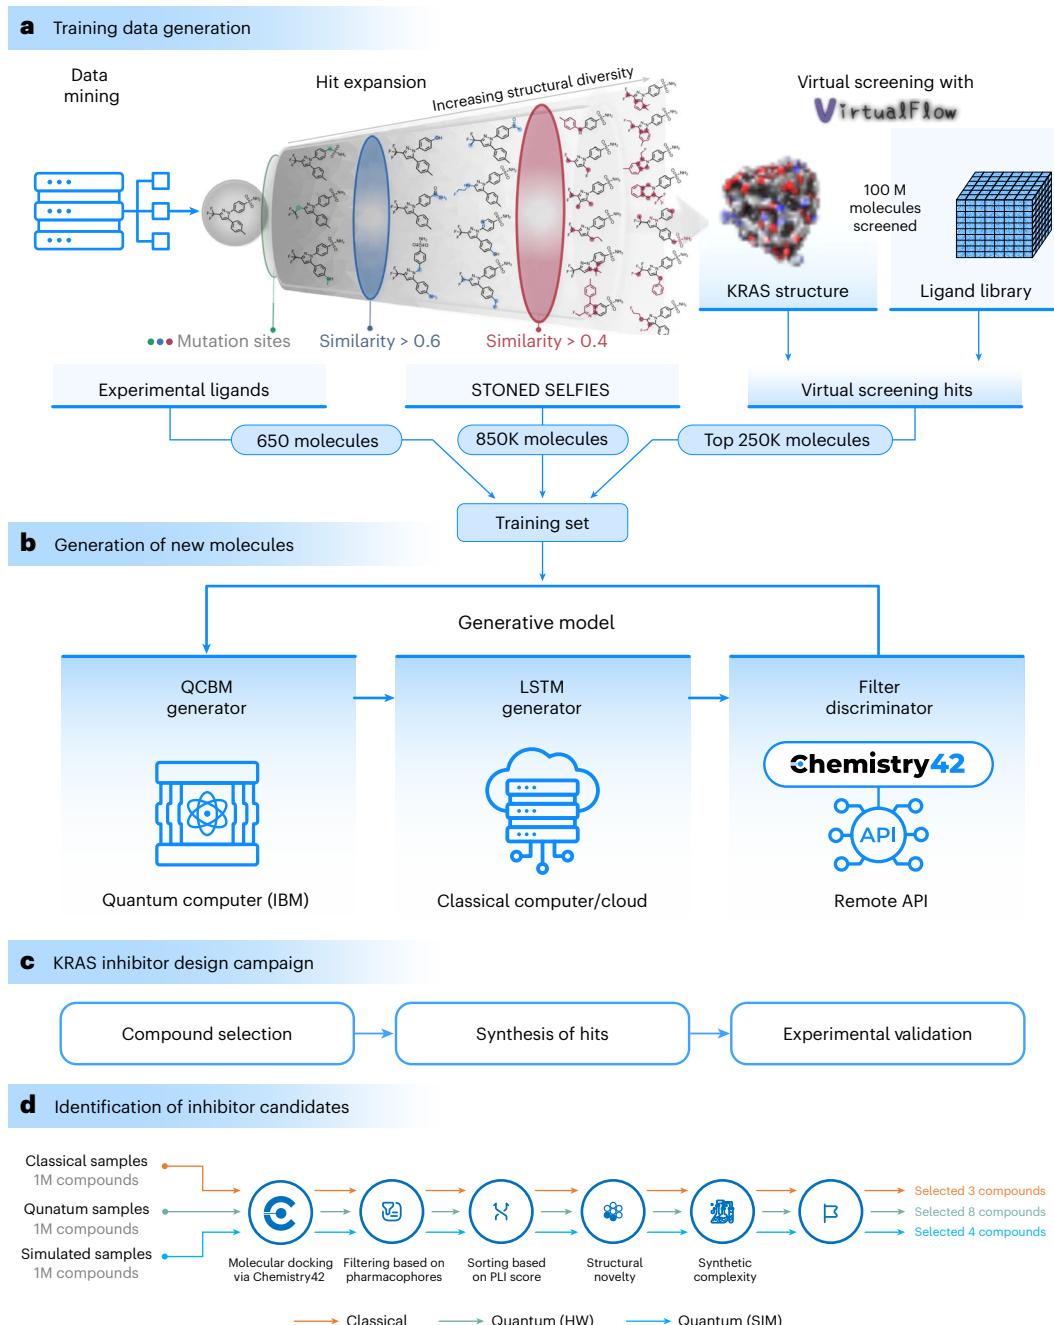


Fig. 1 | Schematic representation of the hybrid quantum-classical framework for KRAS ligand development. **a**, The initial phase involved assembling a training dataset, starting with 650 experimentally verified KRAS inhibitors sourced from the literature. Using the STONED-SELFIES algorithm, analogs of these inhibitors were generated, expanding the dataset to approximately 850,000 compounds. This was further augmented by adding 250,000 top candidates from a virtual screening of the REAL ligand library against KRAS, creating a total dataset of over 1 million molecules. **b**, In the generation phase, the dataset was used to train our generative model, consisting of both a classical LSTM network and a QCBM. The LSTM network processed sequential data of

chemical structures, while the QCBM, trained on the output from the LSTM, generated complex, high-dimensional probability distributions. This workflow incorporated Chemistry42 as a reward function to encourage the production of structurally diverse and synthesizable molecules. **c**, Workflow for KRAS inhibitor design, detailing the process from computational compound selection to laboratory synthesis and experimental validation. **d**, A total of 1 million compounds (classical samples from the LSTM, quantum samples from QCBM on quantum hardware and simulated quantum samples on classical hardware) were evaluated by Chemistry42 to filter out unsuitable candidates and rank the rest by their PLI scores. Finally, 15 promising compounds were selected for synthesis.

learning challenges such as barren plateaus. The integration of tensor networks furthers their effectiveness^{17,18}. However, the challenges in quantum data processing and circuit trainability have led to the development of hybrid algorithms that harness the strengths of both quantum and classical machine learning. Manuel et al.¹⁹ showcased how a hybrid quantum circuit generative adversarial network can enhance target space exploration, overcoming limitations of

classical systems and quantum prior size. Additionally, Zeng et al.²⁰ presented a quantum-classical hybrid for image generation, highlighting the ability to encode conditions into quantum circuits with extra qubits.

Here, we propose a hybrid quantum-classical model that addresses qubit limitations and combines quantum and classical approaches to generate compounds targeting the KRAS protein,

which is known for its intricate complexity and historical resistance to drug discovery efforts^{21,22}. As illustrated in Fig. 1, our workflow was structured into three pivotal stages. The first stage was the generation of training data, beginning with creating a dataset comprising approximately 650 known KRAS inhibitors from the literature. To enhance this dataset, we used VirtualFlow 2.0 to screen 100 million molecules from the Enamine REAL library, selecting the top 250,000 with the best docking scores²³. Additionally, we used the STONED (superfast traversal, optimization, novelty, exploration and discovery) algorithm on the SELFIES (self-referencing embedded strings) molecular representation of known inhibitors, generating structurally similar compounds. After applying synthesizability filtering, this process added 850,000 molecules to our training set²⁴. Furthermore, we used Chemistry42 to validate the molecules generated in this workflow²⁵. We merged data from various sources to compile a single dataset containing 1.1 million data points, all used for training our generative model. The next stage was the generation of new molecules, which combined (1) the QCBM¹⁶ using a 16-qubit processor to generate a prior distribution; (2) a long short-term memory (LSTM) network as the classical model; and (3) Chemistry42 for validation (Fig. 1). The quantum component, depicted in Extended Data Fig. 3c, was a QCBM that generated samples from quantum hardware in every training epoch and was trained with a reward value, $P(x) = \text{softmax}(R(x))$ calculated using Chemistry42 or a local filter. We note that, in practice, the reward function can be tailored to a user's specific criteria, including accounting for off-target effects by evaluating multiple docking scenarios. This recurrent sampling, training and validation process formed a cycle that continuously improved the generated molecular structures targeting the KRAS protein. The final step was experimental validation. From our trained models, we sampled 1 million compounds using three different models (selection workflow in Fig. 1). We used Chemistry42 to screen these samples for pharmacological viability and ranked them on the basis of their docking scores (protein–ligand interaction (PLI) score). The top 15 candidates were synthesized and tested using surface plasmon resonance (SPR) and cell-based assays. This study highlights the promising intersection of quantum computing and drug discovery, marking a quantum computer's first experimental hit.

Before initiating our campaign to design additional KRAS inhibitors, we aimed to compare our hybrid quantum–classical approach to established classical algorithms. We benchmarked our model, targeting three different proteins, against the state of the art using the Tartarus benchmarking suite for drug discovery²⁶. The results indicate that, while few classical models achieved high success rates and docking scores, QCBM–LSTM excelled by generating numerous high-quality samples. It achieved a high success rate and a docking score comparable to the best, as demonstrated in Supplementary Table 1. As illustrated in Extended Data Fig. 1a and Supplementary Table 2, we evaluated the impact of incorporating a quantum prior (QCBM/multibase (M)QCBM–LSTM) versus a fully classical model (vanilla LSTM). We observed that a quantum prior enhanced the success rate, as measured by the proportion of molecules meeting the criteria set by the same filters. Our observations revealed that the use of QCBM–LSTM (Supplementary Table 1) offered a 21.5% improvement in passing filters that assessed the synthesizability and stability of the generated molecules, indicating a higher quality of generated structures. Additionally, we analyzed how the size of the prior (number of qubits) affected sample quality. We found that more qubits improved the success rates for molecule generation, suggesting that larger quantum models could enhance molecular design. Our findings suggest that the success rate correlates approximately linearly with the number of qubits, as illustrated in Extended Data Fig. 1c.

We believe that the improvement in our hybrid classical–quantum approach stemmed from the quantum effects, such as superposition and entanglement, which allow QCBMs to explore and represent

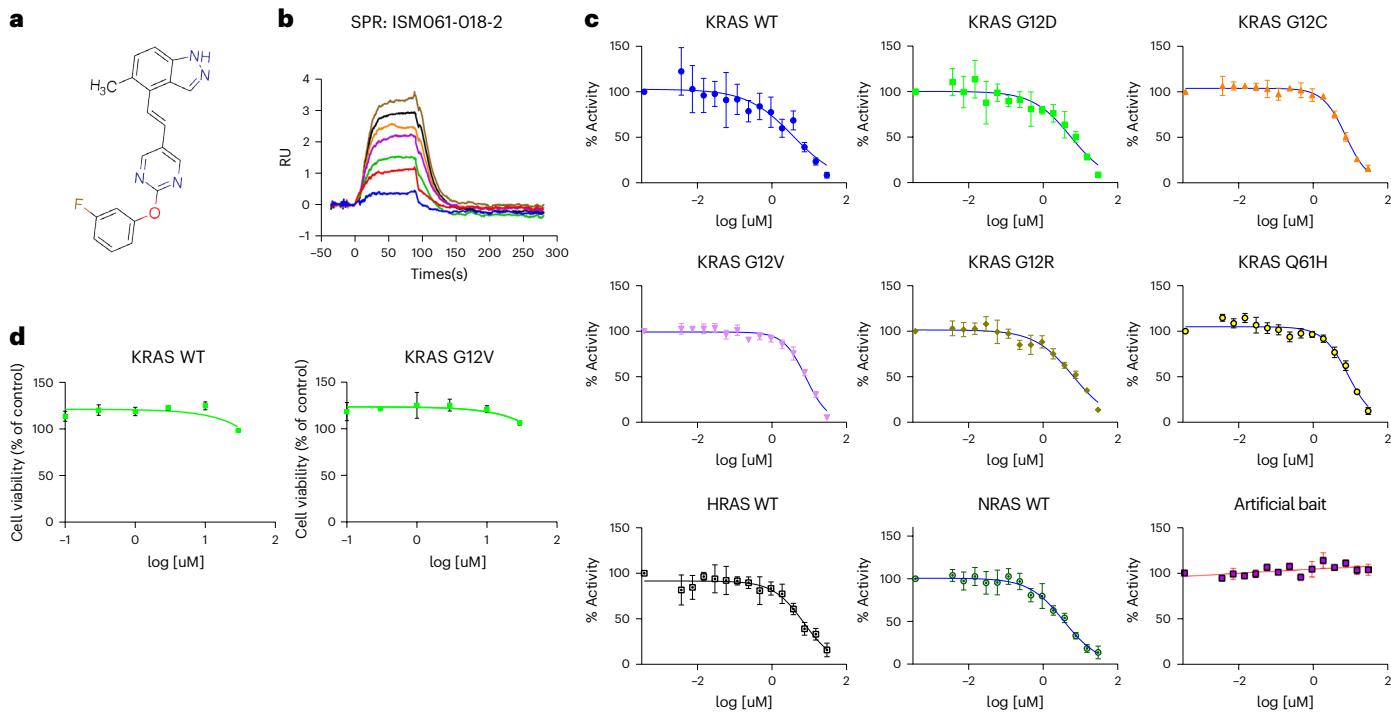
complex, high-dimensional probability distributions more efficiently than classical models²⁷. Entanglement enables the creation of correlations between qubits, capturing intricate dependencies within the prior distribution. This capability is particularly advantageous in generative models, as it allows for a more accurate representation of the underlying distributions in complex datasets, leading to the improved generation of molecular structures. Additionally, these quantum properties help to escape barren plateaus more effectively during optimization—a common challenge in machine learning models¹⁷. Consequently, this resulted in a more comprehensive exploration of the solution space, potentially uncovering molecular structures that our naive LSTM could not generate within the same time frame.

Furthermore, our KRAS inhibitor campaign validated compounds derived from both the hybrid quantum–classical and vanilla approaches using the Chemistry42 platform's structure-based drug design workflow. Within our study, 15 compounds were selected after filtering generated compounds (Fig. 1) using Chemistry42 and synthesized for experimental analysis. The two compounds demonstrating the greatest promise, ISM061-018-2 and ISM061-022, were characterized (structures in Fig. 2a,e, respectively). Their assessment involved a two-stage process; the SPR determined their binding affinities and cell-based assays were subsequently used to gauge biological efficacy. The compound ISM061-018-2, engineered through our hybrid quantum model (Fig. 2b and Supplementary Table 6), demonstrated substantial binding affinity to KRAS-G12D, registered at 1.4 μM . To delve deeper into this molecule's effectiveness across a spectrum of KRAS mutants, we commenced an extensive series of tests using a cell-based assay. Specifically, we evaluated the molecule's performance in a biological context using a commercial cell viability assay (CellTiter-Glo, Promega) in conjunction with MaMTH-DS (mammalian membrane two-hybrid drug screening), an advanced split-ubiquitin-based platform for the real-time detection of small molecules targeting specific cellular interactions^{28–35}.

The biological activity of ISM061-018-2 was rigorously tested. Importantly, it demonstrated no detrimental impact on the viability of HEK293 cells, even when expressing KRAS wild type (WT) or KRAS-G12V bait in MaMTH-DS format and being subjected to concentrations as high as 30 μM for 18–20 h, showing that the compound did not possess any general, nonspecific toxicity (Fig. 2d). Subsequent testing using MaMTH-DS across a spectrum of cell lines expressing various KRAS baits (WT and five clinically important oncogenic mutants) in combination with Raf1 'prey' (a recognized KRAS effector) revealed a dose-responsive inhibition of interactions, with half-maximal inhibitory concentration (IC_{50}) values in the micromolar range (Fig. 2c and Supplementary Table 3). 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 (Fig. 2c). However, it had no effect on the interaction of a completely unrelated artificial bait–prey control pair (consisting of membrane-anchored ALFA tag bait and nanobody ALFA prey)³⁶, supporting the biological specificity of the interaction (Fig. 2c). Collectively, these results support a potential pan-Ras activity of ISM061-018-2.

ISM061-022, as illustrated in Fig. 2e, also stood out as a compound of promise, particularly because of its selectivity toward certain KRAS mutants. Within our in vitro examination, ISM061-022 demonstrated a concentration-dependent inhibition of KRAS interactions with IC_{50} values in the micromolar range (Fig. 2f and Supplementary Table 4), while manifesting only a mild, general impact on cell viability at higher concentrations over an 18–20-h exposure (Fig. 2g). The observed inhibition mirrored that of ISM061-018-2 yet displayed enhanced selectivity toward certain KRAS mutants, particularly KRAS-G12R and KRAS-Q61H, which were most receptive to the compound's action (Fig. 2f). Diverging from ISM061-018-2, ISM061-022 did not show binding to KRAS-G12D

ISM061-018-2



ISM061-022

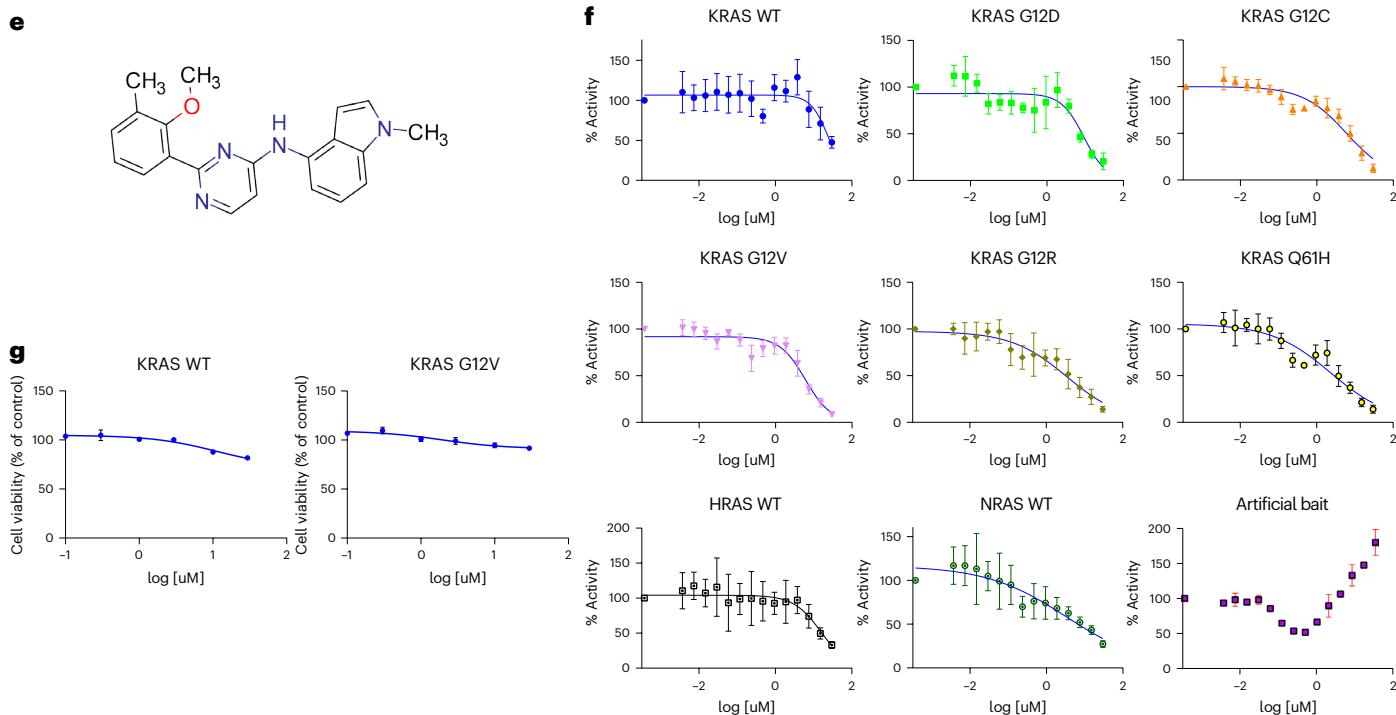


Fig. 2 | Pharmacological characterization of compounds ISM061-018-2 and ISM061-022 through SPR and cellular activity assays. **a,e**, Chemical structures of ISM061-018-2 (**a**) and ISM061-022 (**e**). **b**, SPR sensorgrams illustrating binding effects of ISM061-018-2 on KRAS protein. **c,f**, MaMTH-DS dose-response curves illustrate the binding kinetics and effects of ISM061-018-2 (**c**) and ISM061-022 (**f**) on various KRAS proteins, as well as NRAS, HRAS and artificial bait, with activities measured across concentrations from 4 nM to 30 μ M. Data points are presented as averages from $n = 4$ technical replicates with error bars showing the s.d. Each graph is representative of $n = 3$ biological replicates, performed under the same conditions. Analyses were carried out using GraphPad Prism. **d,g**, Results from CellTiter-Glo viability assays measuring the impact of ISM061-018-2 (**d**) and ISM061-022 (**g**) on cellular proliferation over a concentration range from 123 nM to 30 μ M, demonstrating that they do not display general, nonspecific toxicity. Data points are presented as averages from $n = 3$ technical replicates with error bars showing the s.d. Each graph is representative of $n = 3$ biological replicates performed under the same conditions.

s.d. Each graph is representative of $n = 3$ biological replicates, performed under the same conditions. Analyses were carried out using GraphPad Prism. **d,g**, Results from CellTiter-Glo viability assays measuring the impact of ISM061-018-2 (**d**) and ISM061-022 (**g**) on cellular proliferation over a concentration range from 123 nM to 30 μ M, demonstrating that they do not display general, nonspecific toxicity. Data points are presented as averages from $n = 3$ technical replicates with error bars showing the s.d. Each graph is representative of $n = 3$ biological replicates performed under the same conditions.

according to SPR. The compound also demonstrated activity against WTHRAS and NRAS, although it was less potent against HRAS. Notably, the compound had an unusual effect on our artificial control interaction pair, leading to a distinct dip to approximately 50% residual activity at concentrations of 250–500 nM, before climbing again and ultimately leading to mild enhancement of the interaction at higher micromolar concentrations (Fig. 2f). This distinct pattern suggests an alternative mode of action for ISM061-022, revealing at least some degree of nonspecific activity, although partial specificity for mutant KRAS protein does still appear to be in evidence. Further investigation into the mechanism of action of this molecule will be the focus of future study.

Overall, these live-cell experimental observations underscore the robustness of our approach, effectively identifying small-molecule candidates with biological activity. While these molecules are certainly at an early stage and in need of further refinement, our results underline the potential of our methodology to address and surmount the complexities inherent in targeting clinically challenging biomolecules. To further characterize ISM061-018-2's interaction with KRAS-G12D, we used protein-detected nuclear magnetic resonance (NMR) spectroscopy. ¹⁵N-labeled KRAS-G12D (residues 1–169) protein was used to record ¹H–¹⁵N transverse relaxation optimized spectroscopy (TROSY) heteronuclear single quantum coherence (HSQC) spectra in the presence and absence of ISM061-018-2 under guanosine diphosphate (GDP)-loaded conditions (Extended Data Fig. 2a). Substantial chemical shift perturbations (CSPs) were observed, indicating fast exchange between free and bound states, with some residues showing intensity loss because of exchange broadening. Mapping the weighted CSP values onto the Protein Data Bank (PDB) structure of GDP-bound KRAS-G12D (Extended Data Fig. 2b) revealed notable perturbations in the P-loop, α 2 helix, switch II region and nearby β 4 sheet residues, suggesting a binding mode within the switch II pocket. Additional CSPs in the α 1 helix, switch I and α 4 regions suggest conformational changes upon ligand binding. Although direct interaction with D12 (ref. 21), a residue critical for selectivity in inhibitors such as MRTX1133, was not confirmed, CSPs in the P-loop residues (G10 and A11) imply possible specific interactions. These NMR results strongly support ISM061-018-2's binding to the switch II pocket of KRAS-G12D, further corroborating its proposed mechanism of action. However, cocrystallization studies are necessary to fully elucidate the binding residues. In summary, current quantum hardware is proving crucial for practical drug discovery applications, preempting the need for fully fault-tolerant quantum systems in the future. Quantum algorithms have the potential to design small molecules that better align with the distribution of the provided training set, achieving outcomes that are challenging for classical methods because of their ability to more effectively explore and represent complex, high-dimensional probability distributions. Using just 16 qubits, our algorithm demonstrates how quantum computing can enhance classical computation, leading to more efficient algorithms. Our hybrid quantum–classical model surpasses classical models in performance and matches up well against the state of the art according to the Tartarus benchmark. We applied this model to design inhibitors targeting the complex KRAS protein, synthesizing and evaluating 15 promising ligands. Among these, ISM061-018-2 and ISM061-022 stood out, showing inhibitory effects on KRAS variants with ISM061-018-2 displaying a high binding affinity of 1.4 μ M and broad activity across KRAS mutants. These results, derived from our hybrid approach, indicate superior efficacy to fully classical models and introduce potential chemotypes that expand the chemical space explored in KRAS inhibitor design. While these findings are encouraging, they stop short of definitively proving a ‘quantum advantage’, achieving results unattainable by classical methods within a reasonable time frame. Future work will focus on exploring transformer-based generative models, coupled with increasing the number of qubits, to further enhance the quality and diversity of generated molecules.

This synergistic approach could potentially lead to the discovery of therapeutics that might be overlooked by other methods, while also reducing the preclinical drug discovery phase from several years to just a few months.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41587-024-02526-3>.

References

1. Petrova, E. 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* (eds Ding, M. et al.) 19–82 (Springer, 2014).
2. DiMasi, J. A., Grabowski, H. G. & Hansen, R. W. Innovation in the pharmaceutical industry: new estimates of R&D costs. *J. Health Econ.* **47**, 20–33 (2016).
3. Sun, D., Gao, W., Hu, H. & Zhou, S. Why 90% of clinical drug development fails and how to improve it? *Acta Pharm. Sin. B* **12**, 3049–3062 (2022).
4. Zhavoronkov, A. et al. Deep learning enables rapid identification of potent ddr1 kinase inhibitors. *Nat. Biotechnol.* **37**, 1038–1040 (2019).
5. Stokes, J. et al. A deep learning approach to antibiotic discovery. *Cell* **180**, 688–702 (2020).
6. Ren, F. et al. AlphaFold accelerates artificial intelligence powered drug discovery: efficient discovery of a novel CDK20 small molecule inhibitor. *Chem. Sci.* **14**, 1443–1452 (2023).
7. Nigam, A. K., Pollice, R., Friederich, P. & Aspuru-Guzik, A. Artificial design of organic emitters via a genetic algorithm enhanced by a deep neural network. *Chem. Sci.* **15**, 2618–2639 (2024).
8. Sánchez-Lengeling, B. & Aspuru Guzik, A. Inverse molecular design using machine learning: generative models for matter engineering. *Science* **361**, 360–365 (2018).
9. Pollice, R. et al. Data-driven strategies for accelerated materials design. *Acc. Chem. Res.* **54**, 849–860 (2021).
10. Nigam, A. K. et al. Assigning confidence to molecular property prediction. *Exp. Opin. Drug Discov.* **16**, 935–949 (2021).
11. Bohacek, R. S., McMurtin, C. & Guida, W. C. The art and practice of structure-based drug design: a molecular modeling perspective. *Med. Res. Rev.* **16**, 3–50 (1996).
12. Santagati, R. et al. Drug design on quantum computers. *Nat. Phys.* **20**, 549–557 (2024).
13. Li, J., Topaloglu, R. O. & Ghosh, S. Quantum generative models for small molecule drug discovery. *IEEE Trans. Quantum Eng.* **2**, 1–8 (2021).
14. Batra, K. et al. Quantum machine learning algorithms for drug discovery applications. *J. Chem. Inf. Model.* **61**, 2641–2647 (2021).
15. Benedetti, M., Lloyd, E., Sack, S. & Fiorentini, M. Parameterized quantum circuits as machine learning models. *Quantum Sci. Technol.* **4**, 043001 (2019).
16. Hibat-Allah, M., Mauri, M., Carrasquilla, J. & Perdomo-Ortiz, A. A framework for demonstrating practical quantum advantage: racing quantum against classical generative models. *Commun. Phys.* **7**, 68 (2024).
17. McClean, J. R., Boixo, S., Smelyanskiy, V., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812 (2018).
18. Gili, K., Hibat-Allah, M., Mauri, M., Ballance, C. J. & Perdomo-Ortiz, A. Do quantum circuit born machines generalize? *Quantum Sci. Technol.* **8**, 015021 (2023).

19. Rudolph, M. S. et al. Generation of high-resolution handwritten digits with an ion-trap quantum computer. *Phys. Rev. X* **12**, 031010 (2022).
20. Zeng, Q., Ge, H.-Y., Chen, G. & Zhou, N. Conditional quantum circuit born machine based on a hybrid quantum-classical framework. *Physica A Stat. Mech. Appl.* **608**, 128693 (2023).
21. Wang, X. et al. Identification of MRTX1133, a noncovalent, potent, and selective KRas^{G12D} inhibitor. *J. Med. Chem.* **65**, 3123–3133 (2022).
22. Mao, Z. et al. KRas(G12D) can be targeted by potent inhibitors via formation of salt bridge. *Cell Discov.* **8**, 5 (2022).
23. Gorgulla, C. et al. Virtualflow 2.0—the next generation drug discovery platform enabling adaptive screens of 69 billion molecules. Preprint at *bioRxiv* <https://doi.org/10.1101/2023.04.25.537981> (2023).
24. Nigam, A. K., Pollice, R., Krenn, M., Gomes, G. & Guzik, A. A. Beyond generative models: superfast traversal, optimization, novelty, exploration and discovery (STONED) algorithm for molecules using SELFIES. *Chem. Sci.* **12**, 7079–7090 (2021).
25. Ivanenkov, Y. A. et al. Chemistry42: an AI-driven platform for molecular design and optimization. *J. Chem. Inf. Model.* **63**, 695–701 (2023).
26. Nigam, A. et al. Tartarus: A benchmarking platform for realistic and practical inverse molecular design. In *Advances in Neural Information Processing Systems* **36**, 3263–3306 (2023).
27. Schuld, M., Sinayskiy, I. & Petruccione, F. An introduction to quantum machine learning. *Quantum Mach. Intell.* **3**, 28 (2021).
28. Saraon, P. et al. A drug discovery platform to identify compounds that inhibit EGFR triple mutants. *Nat. Chem. Biol.* **16**, 577–586 (2020).
29. Stagljar, I., Korostensky, C., Johnsson, N. & te Heesen, S. A genetic system based on split-ubiquitin for the analysis of interactions between membrane proteins in vivo. *Proc. Natl Acad. Sci. USA* **95**, 5187–5192 (1998).
30. Petschnigg, J. et al. The mammalian-membrane two-hybrid assay (MAMTH) for probing membrane–protein interactions in human cells. *Nat. Methods* **11**, 585–592 (2014).
31. Saraon, P. et al. Chemical genetics screen identifies COPB2 tool compounds that alters ER stress response and induces RTK dysregulation in lung cancer cells. *J. Mol. Biol.* **433**, 167294 (2021).
32. Scheper, W., Thaminy, S., Kais, S., Stagljar, I. & Römisch, K. Coordination of N-glycosylation and protein translocation across the endoplasmic reticulum membrane by Sss1 protein. *J. Biol. Chem.* **278**, 37998–38003 (2003).
33. Benleulmi-Chaachoua, A. et al. Protein interactome mining defines melatonin MT1 receptors as integral component of presynaptic protein complexes of neurons. *J. Pineal Res.* **60**, 95–108 (2016).
34. Lopes, J. P. et al. The role of Parkinson’s disease-associated receptor GPR37 in the hippocampus: functional interplay with the adenosinergic system. *J. Neurochem.* **134**, 135–146 (2015).
35. Grozavu, I. et al. D154Q mutation does not alter KRas dimerization. *J. Mol. Biol.* **434**, 167543 (2022).
36. Götzke, H. et al. The ALFA-tag is a highly versatile tool for nanobody-based bioscience applications. *Nat. Commun.* **10**, 4403 (2019).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025

Methods

Benchmark setup

Our benchmark used both classical and quantum hardware. Our classical computational setup was based on a cluster equipped with graphics processing unit (GPU) nodes. This cluster consisted of two GPU nodes, each with specific features. These features included two AMD EPYC 7V13 64-core processors, resulting in a total of 128 central processing unit cores per node. In addition, each node was equipped with 512 GB of random-access memory (RAM). The nodes also contain eight AMD Instinct MI100 GPUs, each with a GPU RAM of 32 GB. For the classical training, we used four of these GPUs in parallel (that is, one GPU node). Furthermore, we used an Nvidia GPU (RTX3090Ti) to facilitate our classical–quantum simulations. For the quantum hardware setup, we used the Guadalupe quantum system, equipped with 16 qubits and a Falcon r4P processor type. Our QCBM model, accompanied by an error correction circuit, was executed on this quantum processor.

Regarding software, we used several packages provided by Zapata AI under the Qml core agreements. 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 used RDkit and Insilico APIs to compute the reward value and conduct some postprocessing analyses. The QCBM model underwent a training regimen spanning 30 epochs. In contrast, the LSTM model was trained over a total of 40 epochs.

We used 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, number of 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.

Computational benchmarks: classical versus quantum models
Tartarus benchmark. We used the Tartarus platform²⁶ to benchmark our proposed QCBM–LSTM methodology against an array of classical state-of-the-art models, including REINVENT³⁷, SMILES–VAE³⁸, SELFIES–VAE³⁹, MoFlow⁴⁰, SMILES–LSTM–HC^{41,42}, SELFIES–LSTM–HC, GB–GA⁴³ and JANUS⁴⁴. The study focused on three protein targets selected from the Tartarus dataset: (1) PDB 1SYH, an ionotropic glutamate receptor associated with neurological and psychiatric disorders such as Alzheimer disease, Parkinson disease and epilepsy⁴⁵; (2) PDB 6Y2F, the main protease of severe acute respiratory syndrome coronavirus 2, crucial for its RNA translation⁴⁶; and (3) PDB 4LDE, the β_2 -adrenoceptor G-protein-coupled receptor, a cell-membrane-spanning receptor that binds to adrenaline, a hormone implicated in muscle relaxation and bronchodilation⁴⁷. 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 QuickVina 2 (ref. ⁴⁸). 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, after filtering, were subjected to a refined rescoring using a more precise scoring function provided by SMINA⁴⁹, at an increased level of exhaustiveness.

We conducted experiments using 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 open compound collection^{50,51}, referred to as DATASET in Supplementary Table 1. 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 indicated that only a few generative models demonstrated a high success rate. However, the QCBM–LSTM model was very strong in producing a substantial number of high-quality samples that successfully meet the filter criteria, as evidenced by the elevated success rate depicted in Supplementary Table 1. 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 revealed that, for the PDB 4LDE target, our model generated the highest-scoring molecules relative to other generative models. While the docking scores for the remaining two targets were 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.

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 (Extended Data Fig. 1b). Specifically, we examined a QCBM prior and implemented it on both a quantum simulator and a hardware backend, in contrast with an MQCBM operating exclusively on a quantum simulator and a classical LSTM model devoid of quantum priors. These models were tasked with designing KRAS inhibitors, using a meticulously curated dataset of over 1 million molecules (Fig. 1). Extended Data Figure 1b showcases the optimal results obtained following a comprehensive optimization of the corresponding architectures using Optuna⁵². We assessed the quality of the generated molecules using two distinct sets of criteria: one derived from Tartarus²⁶, 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 enhanced the success rate, as gauged by the proportion of molecules satisfying the criteria set by the two filters. Furthermore, using the top model from each prior category, we sampled 5,000 molecules that successfully met the filter criteria and examined their respective docking scores (Supplementary Table 2). Intriguingly, these molecules displayed comparably high docking scores as determined by QuickVina 2 and the PLI score, as evaluated by Chemistry42. Additionally, the synthesized molecules demonstrated consistent metrics across various parameters, including the diversity fraction, uniqueness fraction, Chemistry42 reward and Chemistry42 synthetic accessibility score²⁵.

Encouraged by our observation that quantum priors enhance molecule quality, we further investigated the influence of the number of qubits used in modeling priors on the quality of generated molecules (Extended Data Fig. 1d). Specifically, we analyzed the percentage of 5,000 uniquely generated random molecules that satisfied a series of local filters. Interestingly, our findings revealed that the success rate correlated roughly linearly with the number of qubits used in modeling 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.

SPR conditions

A Biacore 8K system was used for all experiments. For preliminary compound screening, N-terminal biotinylated KRAS-G12D protein (synthesized by VIVA Biotech; purity $\geq 95\%$) was captured on a sensor chip SA (GE Healthcare) at a density of about 2,000 RU. Protein immobilization was conducted using 1× HBS-EP+, 2 mM TCEP and 2% DMSO as a running buffer. Protein was injected for 70 s at a flow rate of 5 $\mu\text{l min}^{-1}$. The protein concentration was 5 $\mu\text{g ml}^{-1}$. We performed an initial screening of compounds prepared samples by serial twofold

dilutions from 200 μ M to 0.39 μ M in 1× HBS-EP+, 2 mM TCEP and 2% DMSO. Samples were injected for 60 s at a flow rate of 30 μ l min⁻¹ and dissociation time of 180 s. A Biacore 8K machine was used to carry out the SPR experiments and subsequent data analysis.

MaMTH-DS dose–response assays

MaMTH-DS FLP HEK293 reporter cell lines²⁸ stably expressing KRAS (WT or mutant), HRAS, NRAS or artificial membrane-anchored ALFA tag bait alongside Raf1 (for Ras baits) or nanobody ALFA (for artificial bait) preys were seeded into 384-well white-walled, flat-bottomed, tissue-culture-treated microplates (Greiner, 781098) at a concentration of 100,000 cells per ml (50 μ l total volume per well) in DMEM, 10% FBS and 1% penicillin–streptomycin. Seeding was performed using a MultiFlo-FX multimode liquid dispenser (BioTek). Plates were left at room temperature for 30–60 min following seeding before transfer to a Heracell 150i incubator (Thermo Fisher Scientific) and growth at 37 °C in 5% CO₂ for 3 h. After growth, 10 μ l of DMEM, 10% FBS and 1% penicillin–streptomycin supplemented with 3 μ g ml⁻¹ tetracycline (to induce bait and prey expression; BioShop, TET701) and 60 ng ml⁻¹ epidermal growth factor (to stimulate Ras signaling; Sigma, E9644) was added to each well using a multichannel pipette. As appropriate, a 6× concentration of drug (or DMSO only) was also included in the medium, with all lower concentrations produced by serial dilution starting from the highest concentration solution. Plates were then grown overnight (18–20 h) at 37 °C in 5% CO₂. A 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 3,200–3,800 and a 1-s integration time. All data analysis was performed using Microsoft Excel and GraphPad Prism. Curve fits were performed in Prism (nonlinear regression) using log(inhibitor) versus response curves, with a variable slope (four parameters) bottom-constrained to zero.

Cell viability assay

MaMTH-DS FLP HEK293 reporter cell lines²⁸ 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 per well in DMEM, 10% FBS and 1% penicillin–streptomycin (60 μ l total volume per well). Seeding was performed using a MultiFlo-FX multimode liquid dispenser (BioTek). Plates were left at room temperature for 30–60 min following seeding before transfer to a Heracell 150i incubator (Thermo Fisher Scientific) and growth at 37 °C in 5% CO₂ for 3 h. After growth, 30 μ l of a 3× concentration of the drug (or DMSO only) in DMEM, 10% FBS and 1% penicillin–streptomycin was added to wells, with all lower concentrations produced by serial dilution starting from the highest concentration solution (final drug concentration: 30 μ M to 123 nM). Plates were then grown overnight (18–20 h) at 37 °C in 5% CO₂. The effect of the drug on cell viability was assessed by the CellTiter-Glo luminescent cell viability assay from Promega (G7570). Briefly, 90 μ l of the CellTiter-Glo reagent was 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 min to stabilize the luminescence signal. Luminescence was measured using a Clariostar plate reader (BMG Labtech) with a gain of 3,600 and a 1-s 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.

Protein purification and NMR: sample preparation

A construct encoding N-His–TEV (tobacco etch virus)–KRAS-G12D was transformed into BL21(DE3) cells. The cells were grown in minimal medium containing 1 g l⁻¹ [¹⁵N]NH₄Cl and 4 g l⁻¹ glucose to an optical density at 600 nm of ~0.7–1.0 and induced with 1 mM IPTG for 16 h at 16 °C. Cells were pelleted and resuspended in lysis buffer (20 mM

sodium phosphate pH 8.0, 500 mM NaCl, 10 mM imidazole, 1 mM 2-mercaptoethanol and 5% (v/v) glycerol) containing PMSF and benzimidine. Protein was purified over a His trap column (Cytiva) following standard Ni-affinity protocols (wash with 10 mM and 20 mM imidazole and then elute with 300 mM imidazole) and the His-tag was removed by TEV cleavage. KRAS-G12D was further purified by size-exclusion chromatography using a Superdex 75 Increase column (10/300 GL) with PBS buffer (137 mM NaCl, 2.7 mM KCl, 10 mM Na₂HPO₄ and 2 mM KH₂PO₄, pH 7.4) containing 10 mM EDTA. The samples were prepared by mixing 30 μ M of protein with 200 μ M ISM018-2 in the presence of 2% DMSO-*d*₆, 5 mM GDP and 15 mM MgCl₂ in the buffer. The purification protocols were adapted from previously published protocols for KRAS-G12C (ref. 53) and KRAS-G12D (ref. 54).

NMR experiments and analysis

¹H–¹⁵N TROSY HSQC experiments were recorded at 298 K on an 800-MHz Bruker spectrometer equipped with an AVANCE III console and a cryogenically cooled probe. The acquired spectra were processed using NMRPipe⁵⁵ and were analyzed using CCPNMR 2.0 (ref. 56). The chemical shifts were transferred from Biological Magnetic Resonance Bank 27719 (ref. 54). The CSPs on both ¹H and ¹⁵N dimensions were used to calculate the weighted CSP⁵⁷ and these values were then plotted onto the PDB structure using PyMol to map the regions that took part in ligand binding or underwent conformational changes upon ligand binding.

Experimental validation methods

Chemistry42 methods: after screening and selection of promising candidate structures for synthesis. Our selection process relied on a tiered system of the following criteria:

- Structural and compositional parameter filters, such as hydrogen bond donor count and aromatic atom fraction.
- Property evaluation for molecular weight, lipophilicity and other physicochemical traits.
- Medicinal chemistry filters to exclude problematic structural motifs.
- A synthetic accessibility assessment based on the ReRSA model⁵⁸.
- Three-dimensional pharmacophoric analysis in reference to an X-ray cocrystal structure (PDB 7EW9) (Extended Data Fig. 4).
- PLI scoring to estimate binding efficiency.
- An overall reward calculation integrating the above scores

For synthesis candidacy, we applied more demanding conditions, ensuring compounds met the following enhanced benchmarks:

- Clearance of all Chemistry42 filters.
- An aggregate reward value exceeding 0.7.
- A PLI score indicative of strong KRAS binding (<~8 kcal mol⁻¹).
- A pharmacophore match score above 0.7.
- Favorable synthetic accessibility with a ReRSA score under 5.

After screening, molecules were clustered and ranked within clusters by chemical similarity, allowing for expert analysis to further prioritize on the basis of novelty and structural intricacy, culminating in a selection of 100–150 molecules for potential synthesis and subsequent examination.

Experimental evaluation methods 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 Information. The molecular structures of the two most promising compounds (ISM061-018-2 and ISM061-022) are showcased in Fig. 2a,e. Each synthesized compound underwent a rigorous two-phase evaluation; their binding affinities were determined using SPR and their biological efficacies were gauged through cell-based

assays. Notably, compound ISM061-018-2, engineered through our hybrid quantum model (Supplementary Table 3), demonstrated a substantial binding affinity to KRAS-G12D, registered at 1.4 μ M. To delve deeper into this molecule's effectiveness across a spectrum of KRAS mutants, we commenced an extensive series of tests using a cell-based assay. Specifically, we evaluated the molecule's performance in a biological context using 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^{28–34}.

Furthermore, the biological activity of ISM061-018-2 was rigorously tested. Importantly, 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 being subjected to concentrations as high as 30 μ M for 18–20 h, showing that the compound did not possess any general, nonspecific toxicity (Fig. 2d). Subsequent testing using MaMTH-DS across a spectrum of cell lines expressing various KRAS baits (WT and five clinically important oncogenic mutants) in combination with Raf1 prey (a recognized KRAS effector) revealed a dose-responsive inhibition of interactions, with IC_{50} values in the micromolar range (Supplementary Table 3). 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 (Fig. 2c). However, it had no effect on the interaction of a completely unrelated artificial bait–prey control pair (consisting of membrane-anchored ALFA tag bait and nanobody ALFA prey)³⁶, supporting the biological specificity of the interaction (Fig. 2c). Collectively, these results support a potential pan-Ras activity of ISM061-018-2.

ISM061-022 (Fig. 2e) also stood out as a compound of promise, particularly because of its selectivity toward certain KRAS mutants. Within our *in vitro* examination, ISM061-022 demonstrated a concentration-dependent inhibition of KRAS interactions (Fig. 2f), while manifesting only a mild, general impact on cell viability at higher concentrations over an 18–20-h exposure (Fig. 2g). The observed inhibition mirrored that of ISM061-018-2 yet displayed enhanced selectivity toward certain KRAS mutants, particularly KRAS-G12R and KRAS-Q61H, which were most receptive to the compound's action (Fig. 2f and Supplementary Table 4). Diverging from ISM061-018-2, ISM061-022 did not show binding to KRAS-G12D. The compound also demonstrated activity against WT HRAS and NRAS, although it was less potent against HRAS. Notably, the compound had an unusual effect on our artificial control interaction pair, leading to a distinct dip to approximately 50% residual activity at concentrations of 250–500 nM, before climbing again and ultimately leading to mild enhancement of the interaction at higher micromolar concentrations (Fig. 2f). This distinct pattern suggests an alternative mode of action for ISM061-022, revealing at least some degree of nonspecific activity, although partial specificity for mutant KRAS protein does still appear to be in evidence. Further investigation into the mechanism of action of this molecule will be the focus of future study.

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.

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 used in our study.

Data acquisition and preprocessing

Our preliminary dataset, sourced from Insilico Medicine, included approximately 650 data points. These were selectively collated from existing literature, specifically targeting the KRAS-G12D mutant (Fig. 1).

Given the dataset's limited size, we opted to expand it to improve the robustness of our model during training.

STONED–SELFIES. We used the STONED–SELFIES²⁴ algorithm (<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 RDKit. 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 850,000 molecules, which served as the training set for our generative models.

Virtual screening process. VirtualFlow 2.0 (ref. 23) 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 prescreen 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 (ref. 48), with exhaustiveness set to 1 in both stages of the screen. The screen was carried out in the Amazon Web Services cloud computing platform. The protein structure used in the screen was PDB 5US4 (ref. 59), which was prepared before the virtual screen with Schrödinger's protein preparation wizard (addition of hydrogens and protonation state prediction). The size of the docking box was $14 \times 14 \times 20 \text{ \AA}^3$.

Quantum-assisted algorithm

As shown in Fig. 1, our quantum-assisted model was a hybrid algorithm composed of both quantum and classical generative components. The quantum generative model used a QCBM model while the classical component used an LSTM model. Extended Data Fig. 3 illustrates the flowchart of our proposed generative model.

Within this model, we used Chemistry42 and a local filter to validate sample generation at each step, which was then used 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 on the basis of these samples. The reward value for each sample was computed at every step using the local filter until epoch 20, after which we selected Chemistry42. This reward value was then used to train our quantum generative model. During the first epoch, no rewards were available; hence, 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 1.

Quantum-computing-enhanced workflow

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

(1) Generation of training data. We initiated the process by constructing a robust dataset for training our generative model to target the KRAS protein. The foundation of this dataset was approximately 650 experimentally confirmed KRAS inhibitors, compiled through an extensive literature review^{60–63}. 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 used VirtualFlow 2.0 (ref. 23) to screen 100 million molecules, applying Enamine's REAL library⁶⁴ in conjunction with molecular docking techniques. The top 250,000 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 algorithm²⁴, which was applied to the 650 experimentally derived hits. This algorithm distinctively introduces random point substitutions into the SELFIES representations^{39,65,66} of the molecules, thereby generating novel compounds that maintain a structural resemblance to the starting point. The resulting derivatives were filtered on the basis of synthesizability, culminating in the addition of 850,000 molecules to our training set.

(2) Generation of new molecules. Our approach was structured around the integration of three primary components: (1) the QCBM; (2) the classical LSTM model; and (3) Chemistry42 for artificial-intelligence-driven validation, as shown in Extended Data Fig. 3. The QCBM generator¹⁶ used a 16-qubit IBM quantum processor with quantum circuits to model complex data distributions. The integration method of quantum priors into the LSTM architecture (Supplementary Fig. 4) involved merging molecular information encoded in SELFIES and quantum data by addition or concatenation to form samples, $X'(t)$, which were then input into the LSTM cell. The quantum component (Supplementary Fig. 4) was a QCBM that generated samples from quantum hardware each training epoch and was 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 formed 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 at the bottom of Fig. 1. After training our model, we sampled 1 million compounds from each prior model listed at the bottom of Fig. 1. 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 SPR and cell-based assay experiments.

Algorithm 1. Algorithm 1 provides an 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 (Algorithm 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

```

QCBM model. The QCBM model represents a quantum variational generative model, necessitating a classical optimizer to train its parameters. The total count of these parameters was computed with the number of qubits and layers defined in the model (Extended Data Fig. 5). 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 QCBM model involves minimizing the exact negative log-likelihood (exact NLL) loss function.

Classical model: LSTM model. LSTM networks (Fig. 1) were used for the 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 capable of learning long-term dependencies in sequence data⁶⁷. 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 and time-series forecasting⁶⁸. The LSTM architecture consists of a chain of repeating modules called cells. Each cell contains three gates that control the flow of information:

- 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, where 0 means 'completely forget this' and 1 means 'completely keep this'.
- 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.
- 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:

$$\text{Forget gate: } f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + \mathbf{b}_f) \quad (1)$$

$$\text{Input gate: } i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + \mathbf{b}_i) \quad (2)$$

$$\text{Candidate values: } \tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + \mathbf{b}_C) \quad (3)$$

$$\text{Update cell state: } C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t \quad (4)$$

$$\text{Output gate: } o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + \mathbf{b}_o) \quad (5)$$

$$\text{Update hidden state: } h_t = o_t \cdot \tanh(C_t) \quad (6)$$

Here, σ is the sigmoid activation function, W and \mathbf{b} are the weight matrices and bias vectors for each gate and x_t is the input at time t ⁶⁹.

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 or Adam⁷⁰. The backpropagation through time algorithm was used to compute the gradients relative to the loss function, considering the sequential nature of the data⁷¹. The networks were trained using the Adam optimizer with the NLL loss function; to mitigate overfitting, regularization techniques such as dropout were implemented⁷². The NLL loss for a single data point is given by

$$L(y, \hat{y}) = -\log(\hat{y}_y) \quad (7)$$

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:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \log(\hat{y}_{y_i}) \quad (8)$$

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 used Optuna, 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 integrated a deep learning architecture. This architecture was designed to incorporate prior information (samples) into the generative process. Additionally, the model used Chemistry42 feedback in conjunction with QCBMs, aimed at enhancing its generative accuracy. Supplementary Figure 4 illustrates the proposed architecture at a cell level. The prior samples were combined with input samples $x'_t = X(i) + x_t$ in the LSTM cell. This combination constituted two methods: adding and concatenating samples. The LSTM's operations were updated with the following operations:

$$\text{Prior sampling } x'_t = X(i) + x_t \text{ OR } x'_t = X(i) + x_t \quad (9)$$

$$\text{Forget gate: } f_t = \sigma(W_f \cdot [h_{t-1}, x'_t] + \mathbf{b}_f) \quad (10)$$

$$\text{Input gate: } i_t = \sigma(W_i \cdot [h_{t-1}, x'_t] + \mathbf{b}_i) \quad (11)$$

$$\text{Candidate values: } \tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x'_t] + \mathbf{b}_C) \quad (12)$$

$$\text{Update cell state: } C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t \quad (13)$$

$$\text{Output gate: } o_t = \sigma(W_o \cdot [h_{t-1}, x'_t] + \mathbf{b}_o) \quad (14)$$

$$\text{Update hidden state: } h_t = o_t \cdot \tanh(C_t) \quad (15)$$

To generate samples, the process began with sampling from the prior, followed by the LSTM network processing these prior samples to generate compounds representations. The compounds were validated through the Chemistry42 platform, specifically tailored to assess ligand quality for the KRAS-G12D mutant. This methodology allows 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 QCBM functions as a prior, guiding the LSTM in the generation of novel ligand samples. The procedure was subjected to an iterative process to enhance the quality of ligands, which was evaluated using the Chemistry42 platform.

Quantum generative model: QCBM model. The QCBM is a variational quantum algorithm that uses 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 QCBM involves optimizing the parameters of the quantum circuit to produce a probability distribution that closely approximates the target distribution (probability computed by Chemistry42 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 on the basis of 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 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 that 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.

Extended Data Fig. 5 shows the QCBM architecture and illustrates its associated ansatz. We used linear topology for our project. Our QCBM model was built with 16 qubits and four layers and we had 96 parameters to optimize in total. The initial probability of the samples, $P(X(i))$, was computed on the basis of the rewards returned by the Chemistry42 model. These reward-based probabilities were then passed through a softmax function to ensure that they were normalized and fell within the range of 0–1. The resulting values served as the 'true' probabilities of the samples and were used as the target values during the model's training process.

Algorithm 2. Algorithm 2 provides the pseudocode outlining the training regimen for the 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 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  $P(X_i)$  on the basis of the rewards
11:    Train QCBM model with adjusted probabilities
12:   end If
13: end while

```

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The datasets generated and analyzed during the current study, along with the generative model, are available via Zenodo at <https://doi.org/10.5281/zenodo.11137638> (ref. 73).

Code availability

The full code for running our generative model is available via Zenodo at <https://doi.org/10.5281/zenodo.11137638> (ref. 73).

References

37. Olivecrona, M., Blaschke, T., Engkvist, O. & Chen, H. Molecular de-novo design through deep reinforcement learning. *J. Cheminform.* **9**, 48 (2017).
38. Bombarelli, R. G. et al. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Cent. Sci.* **4**, 268–276 (2018).
39. Krenn, M. et al. SELFIES and the future of molecular string representations. *Patterns* **3**, 100588 (2022).
40. Zang, C. & Wang, F. Moflow: an invertible flow model for generating molecular graphs. In *Proc. of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (eds Gupta, R. & Liu, Y.) 617–626 (ACM, 2020).
41. Segler, M., Kogej, T., Tyrchan, C. & Waller, M. P. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS Cent. Sci.* **4**, 120–131 (2018).
42. Brown, N., Fiscato, M., Segler, M. & Vaucher, A. C. Guacamol: benchmarking models for de novo molecular design. *J. Chem. Inf. Model.* **59**, 1096–1108 (2019).
43. Jensen, J. H. A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space. *Chem. Sci.* **10**, 3567–3572 (2019).
44. Nigam, A. K., Pollice, R. & Guzik, A. A. Parallel tempered genetic algorithm guided by deep neural networks for inverse molecular design. *Digit. Discov.* **1**, 390–404 (2022).
45. Frandsen, A. et al. Tyr702 is an important determinant of agonist binding and domain closure of the ligand-binding core of GLUR2. *Mol. Pharmacol.* **67**, 703–713 (2005).
46. Zhang, L. et al. Crystal structure of SARS-CoV-2 main protease provides a basis for design of improved α -ketoamide inhibitors. *Science* **368**, 409–412 (2020).
47. Ring, A. M. et al. Adrenaline-activated structure of β_2 -adrenoceptor stabilized by an engineered nanobody. *Nature* **502**, 575–579 (2013).
48. Alhossary, A., Handoko, S. D., Mu, Y. & Kwoh, C. K. Fast, accurate, and reliable molecular docking with QuickVina 2. *Bioinformatics* **31**, 2214–2216 (2015).
49. Koes, D. R., Baumgartner, M. P. & Camacho, C. J. Lessons learned in empirical scoring with smina from the CSAR 2011 benchmarking exercise. *J. Chem. Inf. Model.* **53**, 1893–1904 (2013).
50. Voigt, J. H., Bienfait, B., Wang, S. & Nicklaus, M. C. Comparison of the NCI open database with seven large chemical structural databases. *J. Chem. Inf. Comput. Sci.* **41**, 702–712 (2001).
51. Ihlenfeldt, W. D., Voigt, J. H., Bienfait, B., Oellien, F. & Nicklaus, M. C. Enhanced CACTVS browser of the open NCI database. *J. Chem. Inf. Comput. Sci.* **42**, 46–57 (2002).
52. Akiba, T., Sano, S., Yanase, T., Ohta, T. & Koyama, M. Optuna: a next-generation hyperparameter optimization framework. In *Proc. of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (eds Teredesai, A. & Kumar, V.) 2623–2631 (ACM, 2019).
53. Lim, S. M. et al. Therapeutic targeting of oncogenic K-Ras by a covalent catalytic site inhibitor. *Angew. Chem. Int. Ed. Engl.* **53**, 199–204 (2014).
54. Pálfy, G., Vida, I. & Perczel, A. ^1H , ^{15}N backbone assignment and comparative analysis of the wild type and G12C, G12D, G12V mutants of K-Ras bound to GDP at physiological pH. *Biomol. NMR Assign.* **14**, 1–7 (2020).
55. Delaglio, F. et al. NMRPipe: a multidimensional spectral processing system based on unix pipes. *J. Biomol. NMR* **6**, 277–293 (1995).
56. Vranken, W. et al. The CCPN data model for NMR spectroscopy: development of a software pipeline. *Proteins* **59**, 687–696 (2005).
57. Williamson, M. P. Using chemical shift perturbation to characterise ligand binding. *Prog. Nucl. Magn. Reson. Spectrosc.* **73**, 1–16 (2013).
58. Zagribelnyy, B., Putin, E. O., Fedorchenco, S. A., Ivanenkov, Y. A. & Zavoronkovs, A. Retrosynthesis-related synthetic accessibility. WO patent 2021/229454 A1 (2021).
59. Welsch, M. et al. Multivalent small-molecule pan-RAS inhibitors. *Cell* **168**, 878–889 (2017).
60. Kwan, A. K., Piazza, G. A., Keeton, A. B. & Leite, C. A. The path to the clinic: a comprehensive review on direct KRas $^{\text{G12C}}$ inhibitors. *J. Exp. Clin. Cancer Res.* **41**, 27 (2022).
61. Parikh, K. et al. Drugging KRas: current perspectives and state-of-art review. *J. Hematol. Oncol.* **15**, 152 (2022).
62. Srisongkram, T., Khamtang, P. & Weerapreeyakul, N. Prediction of KRas $^{\text{G12C}}$ inhibitors using conjoint fingerprint and machine learning-based QSAR models. *J. Mol. Graph. Model.* **120**, 108197 (2023).
63. Nagasaka, M. et al. KRas G12C Game of Yhrone, which direct KRas inhibitor will claim the iron throne? *Cancer Treat. Rev.* **84**, 101974 (2020).
64. Grygorenko, O. O. et al. Generating multibillion chemical space of readily accessible screening compounds. *iScience* **23**, 101681 (2020).
65. Krenn, M., Häse, F., Nigam, A. K. & Friederich, P. Self-referencing embedded strings (SELFIES): a 100% robust molecular string representation. *Mach. Learn. Sci. Technol.* **1**, 045024 (2020).
66. Lo, A. et al. Recent advances in the self-referencing embedded strings (SELFIES) library. *Digit. Discov.* **2**, 897–908 (2023).
67. Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural Comput.* **9**, 1735–1780 (1997).
68. Gers, F. A., Schmidhuber, J. & Cummins, F. Learning to forget: continual prediction with LSTM. In *Proc. of the 9th International Conference on Artificial Neural Networks* (eds Sandoval, F. et al.) 850–855 (Springer, 1999).
69. Gers, F. A. & Schmidhuber, J. Recurrent nets that time and count. In *Proc. of the IEEE-INNS-ENNS International Joint Conference on Neural Networks* (eds Amari, S.-I. et al.) 189–194 (IEEE, 2000).
70. Kingma, D. P. & Ba, J. Adam: a method for stochastic optimization. In *Proc. of the 3rd International Conference on Learning Representations* (ed. Yue, X.-G.) 1–13 (ACM, 2015).
71. Werbos, P. J. Backpropagation through time: what it does and how to do it. *Proc. IEEE* **78**, 1550–1560 (1990).
72. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**, 1929–1958 (2014).
73. Vakili, M. G. Quantum computing-enhanced algorithm unveils novel inhibitors for KRas. Zenodo <https://doi.org/10.5281/zenodo.11137638> (2024).

Acknowledgements

We want to thank A. Perdomo-Ortiz, M. Mauri, B. Dellabetta, V. Vargas-Calderón, A. Cheng, J. Miller, M. Bagherimehrab and K. Nakaji for their valuable discussion and support in this research. We thank H. Arthanari for providing us with his NMR resources. Additionally, we are thankful to the Defense Advanced Research Projects Agency for their funding (grant number HR0011-23-3-0017), which greatly supports our scientific endeavors. A.K.N. acknowledges funding from the Bio-X Stanford Interdisciplinary Graduate Fellowship. A.A.-G. thanks A. 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. C.G. was funded through the Blue Sky Kinase Project at St. Jude Children's Research Hospital. The research in the I.S. lab is supported by the grants from the Canadian Institute for Health Research, Genome Canada, Ontario Research Fund, Ontario

Genomics Institute, Cystic Fibrosis Canada and Toronto Innovation Acceleration Partners. This research was undertaken thanks partly to funding provided to the Acceleration Consortium of the University of Toronto from the Canada First Research Excellence Fund.

Author contributions

M.G.V. led the project, developed the quantum machine learning algorithm, benchmarked the model and contributed substantially to writing the manuscript. C.G. led the virtual screening and NMR experiments. A.K.N. worked on the genetic algorithm and generated the initial data. The Insilico team focused on synthesizing the compounds and performing the SPR assays. J.S. led the cell-based assays conducted in the I.S. lab. The remaining authors worked in parallel to validate the findings, discuss the results and contribute to the preparation and analysis of the final results.

Competing interests

A.A.G. serves as the chief visionary officer and is a board member of Kebotix. A.Z., A.A., D.B., D.P., X.D., J.L., E.R., F.R. and F.M. are affiliated with Insilico Medicine. Y.C. and D.V. are affiliated with Zapata AI. C.G.

is a cofounder and consultant of Virtual Discovery and Quantum Therapeutics. The remaining authors declare no competing interests.

Additional information

Extended data is available for this paper at <https://doi.org/10.1038/s41587-024-02526-3>.

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41587-024-02526-3>.

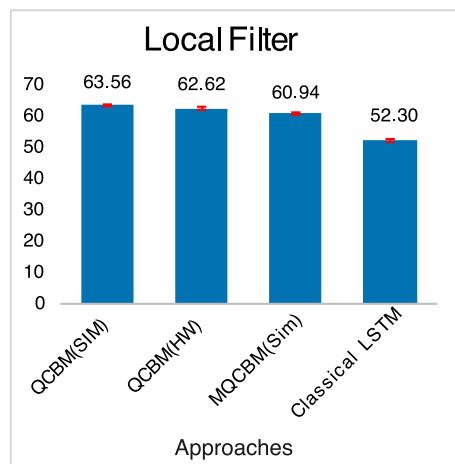
Correspondence and requests for materials should be addressed to Christoph Gorgulla, AkshatKumar Nigam, Igor Stagljar, Alán Aspuru-Guzik or Alex Zhavoronkov.

Peer review information *Nature Biotechnology* thanks Su-Yang Xu and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

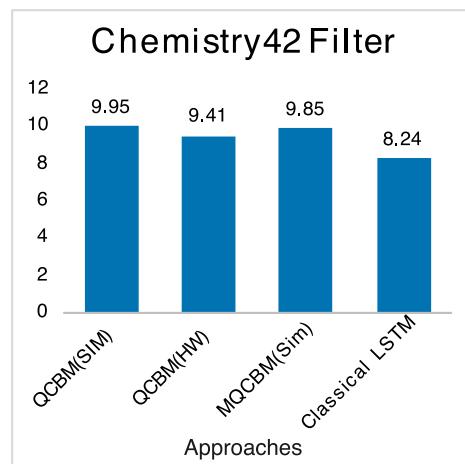
Reprints and permissions information is available at www.nature.com/reprints.

Clasical vs. Quantum Modelling

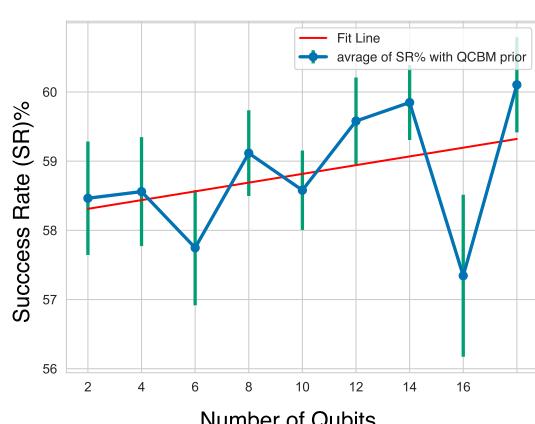
(A)



(B)



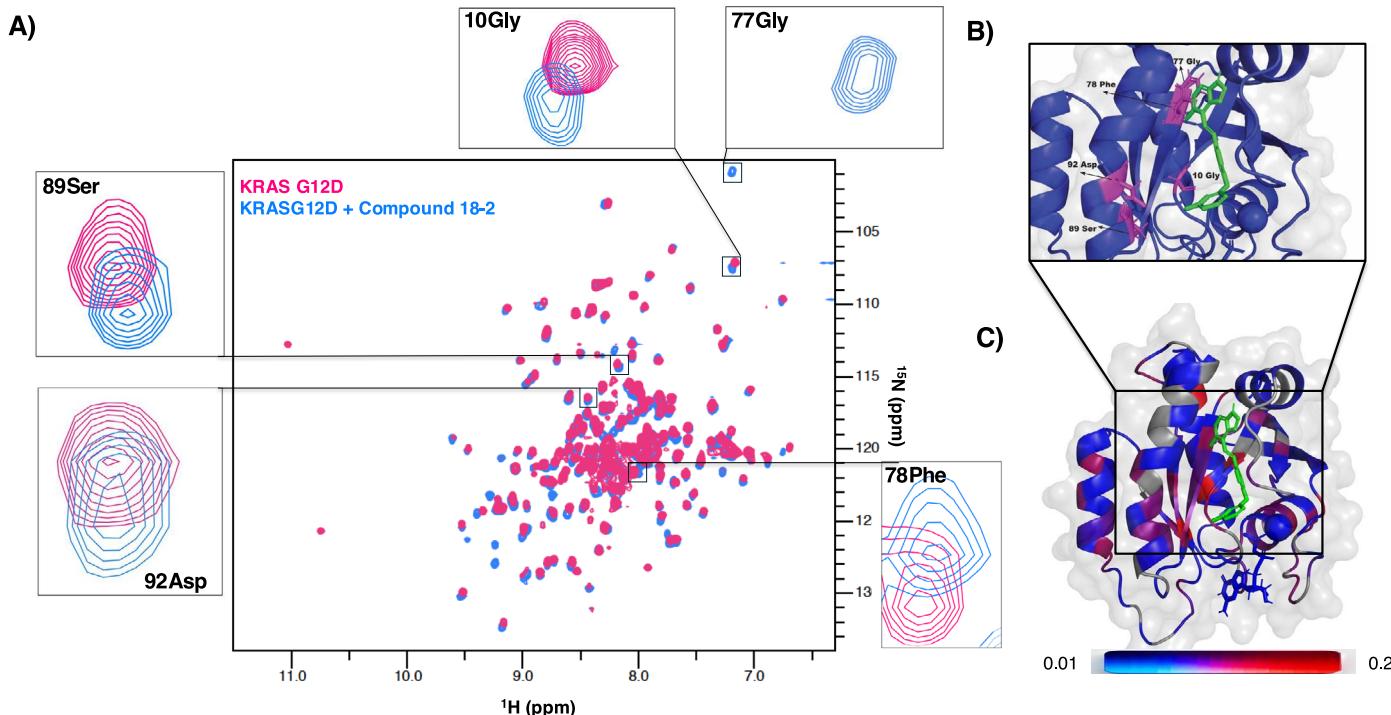
(C)



Number of Qubits

Extended Data Fig. 1 | Comparative Benchmarking of Quantum and Classical Ligand Design Methods. (A, 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)²⁶ and by Chemistry42 (referred to as Chemistry42 Filters). To generate Figure A, we repeated the experiments five times, sampling n=1,000 compounds in each repetition and applying the filter. The reported values represent the mean for each data point, with error bars indicating the standard deviation across the repetitions (Mean \pm Std.). more detailed is reported in Supplementary Information Table S3.2. (C) 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. We repeated the experiments five times, sampling n=5,000 compounds in each repetition and applying the filter. The reported values represent the mean for each data point, with error bars indicating the standard deviation across the repetitions.

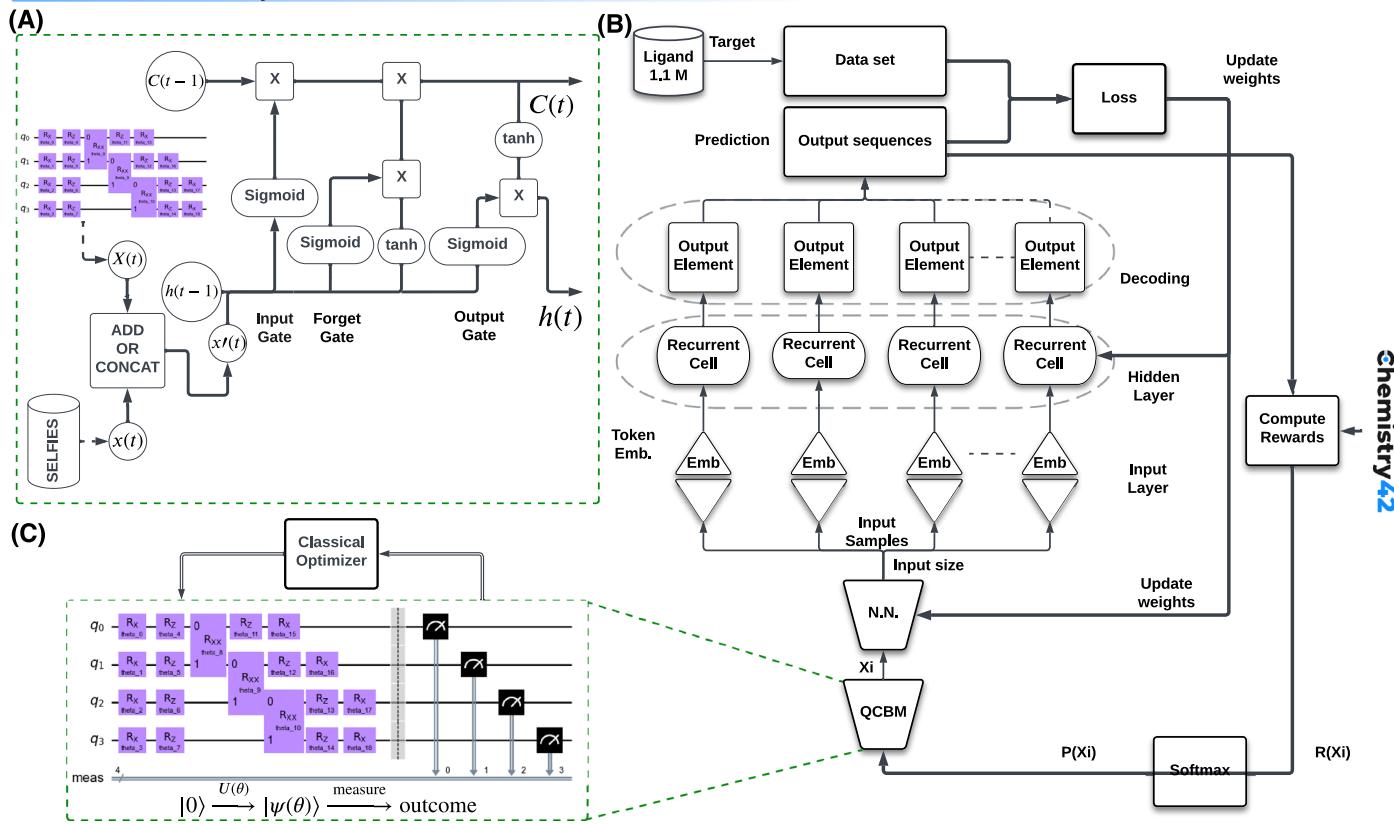
approaches with varied priors. The performance of the Quantum Circuit Born Machine (QCBM) was assessed using both a quantum simulator (Sim) and a hardware backend (HW), and contrasted with a Multi-bases QCBM (MQCBM) operating solely on a quantum simulator (Sim), as well as an LSTM model devoid of quantum 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)²⁶ and by Chemistry42 (referred to as Chemistry42 Filters). To generate Figure A, we repeated the experiments five times, sampling n=1,000 compounds in each repetition and applying the filter. The reported values represent the mean for each data point, with error bars indicating the standard deviation across the repetitions (Mean \pm Std.). more detailed is reported in Supplementary Information Table S3.2. (C) 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. We repeated the experiments five times, sampling n=5,000 compounds in each repetition and applying the filter. The reported values represent the mean for each data point, with error bars indicating the standard deviation across the repetitions.



Extended Data Fig. 2 | Protein-detected NMR experiments. Protein-detected NMR experiments elucidate the binding mode of Compound 18-2 to KRAS-G12D. (A) Significant chemical shift perturbations (CSPs) and intensity changes due to chemical exchange are observed upon binding to the compound. Residues lining the Switch-II pocket that show significant CSPs are highlighted in the inset, including Gly10 (from the P-loop), Phe78 and Gly77 (from the α -2 helix adjacent to the Switch-II region), Ser89, and Asp92 (from the β -4 sheet). (B) The CSP values

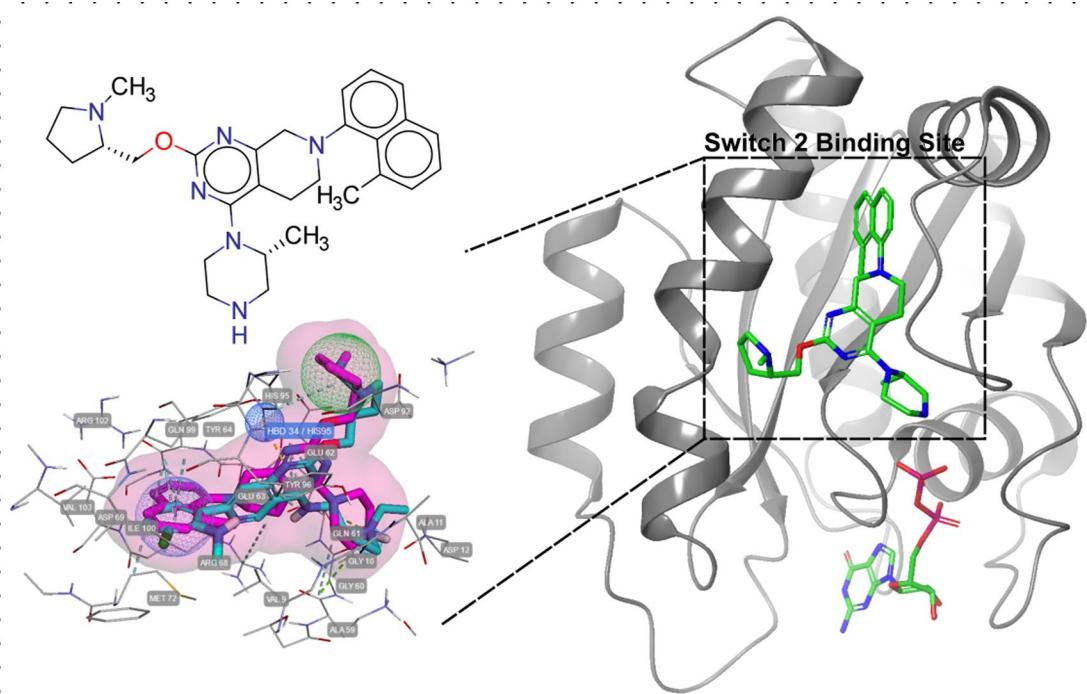
are mapped onto the G12D structure and displayed in a color gradient. Residues that could not be assigned are shown in grey. The majority of significant CSPs are observed near the Switch-II pocket; however, CSPs are also noted in the α -1 helix, Switch-I, and α -4 regions, possibly indicating conformational changes upon binding to the compound. (C) Zoomed-in region of the protein, highlighting residues from the Switch-II pocket that exhibit CSPs.

Architecture of Hybrid Quantum-Classical Model



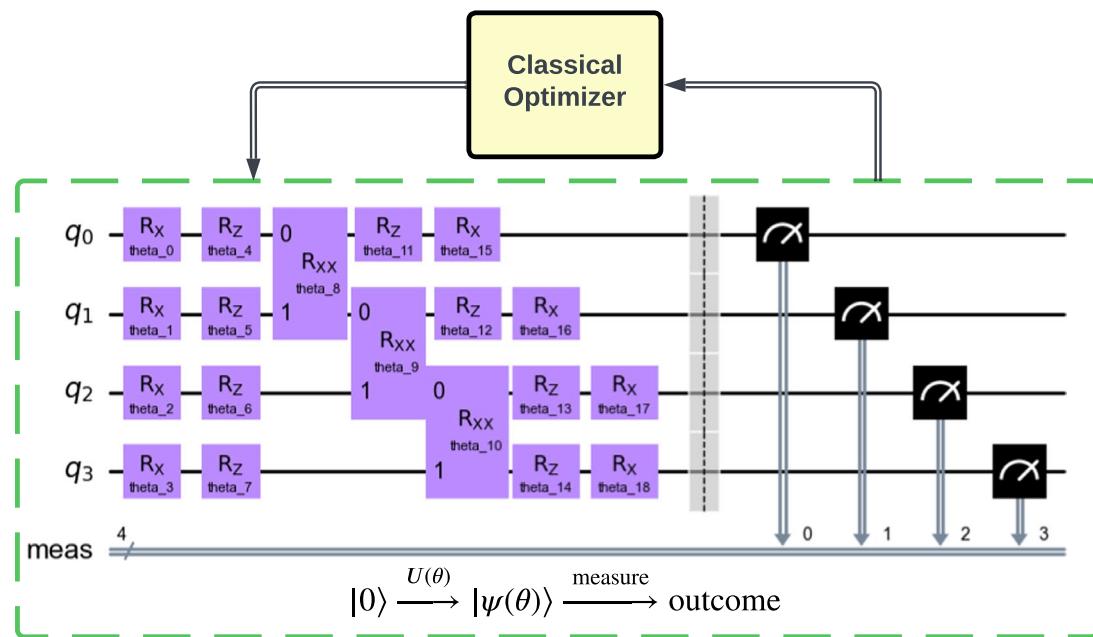
Extended Data Fig. 3 | Quantum-Enhanced Generative Model for Drug Discovery Applications. (B) Hybrid model combining a Quantum Circuit Born Machine (QCBM) with Long Short-Term Memory (LSTM). This model iteratively trains using prior samples from quantum hardware. (A) Integration method of prior samples 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.



Extended Data Fig. 4 | Pharmacophore Model Depiction. Pharmacophore Model Depiction for KRAS Inhibitor TH-Z816 Based on the Co-crystallized Ligand Structure Analyzed with Chemistry42 (PDB: [7EW9](https://doi.org/10.1101/2023.09.04.553312)). This figure illustrates the key pharmacophoric features identified from the ligand structure of TH-Z816 bound to KRAS G12D. A blue sphere represents a critical ring system that

supports structural integrity, a green sphere highlights a hydrophobic moiety essential for binding affinity, and a cyan sphere indicates a hydrogen bond donor that contributes to interaction specificity with the KRAS protein. The protein structure is shown on the right, while the pharmacophore interactions within the KRAS Switch-II binding pocket are detailed on the left.



Extended Data Fig. 5 | Quantum Circuit Born Machine (QCBM) Model.
 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 (R_x , R_z) 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.

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection	The data used in this research were generated using VirtualFlow 2.0 to produce ligand candidates, followed by validation and scoring using the Chemistry42 platform. Detailed information about the dataset preparation is provided in the manuscript.
Data analysis	Data analysis was conducted using custom Python scripts, with the workflow described in detail in the manuscript. Statistical computations, performance evaluations, and visualization were performed using standard Python libraries. The scripts are available in the GitHub repository and have been cited in the paper for accessibility and reproducibility.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

Zenodo Repository: <https://zenodo.org/records/11137638>

Research involving human participants, their data, or biological material

Policy information about studies with [human participants or human data](#). See also policy information about [sex, gender \(identity/presentation\), and sexual orientation](#) and [race, ethnicity and racism](#).

Reporting on sex and gender

N/A

Reporting on race, ethnicity, or other socially relevant groupings

N/A

Population characteristics

N/A

Recruitment

N/A

Ethics oversight

N/A

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences

Behavioural & social sciences

Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size

The sample size for experimental validation was determined based on standard practices in biological assays and the availability of synthesized compounds. For Surface Plasmon Resonance (SPR) experiments, serial dilutions of the compounds were tested across a concentration range of 0.39 μ M to 200 μ M, ensuring sufficient data points for reliable binding affinity measurements. Cell-based assays were conducted with three technical replicates for each concentration, ensuring statistical robustness in viability and inhibition data. The chosen sample sizes are sufficient to detect meaningful differences in activity while minimizing resource use.

Data exclusions

No data were excluded from the analysis unless experiments failed due to technical issues such as improper compound dissolution or equipment malfunction. Excluded data points and their rationale are noted in the raw dataset and supplementary materials. The exclusion criteria were predefined to ensure data integrity and reliability.

Replication

All experimental findings were independently replicated to confirm reproducibility. SPR experiments were repeated for each compound to validate binding kinetics, and cell-based assays were conducted in triplicate to ensure consistent results. Results were verified to align across replicates, confirming the robustness of the findings.

Randomization

Compounds were assigned to experimental groups based on their order in the synthesis workflow and tested across the full range of concentrations in a randomized sequence to minimize systematic errors. Cell lines and assay plates were randomized during preparation to avoid any unintended bias.

Blinding

Blinding was not implemented during the data collection or analysis phases, as compound testing was conducted in a systematic and automated manner using established protocols. The lack of subjective decision-making in the experimental workflow made blinding unnecessary for this study.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
<input checked="" type="checkbox"/>	Antibodies
<input type="checkbox"/>	Eukaryotic cell lines
<input checked="" type="checkbox"/>	Palaeontology and archaeology
<input checked="" type="checkbox"/>	Animals and other organisms
<input checked="" type="checkbox"/>	Clinical data
<input checked="" type="checkbox"/>	Dual use research of concern
<input checked="" type="checkbox"/>	Plants

Methods

n/a	Involved in the study
<input checked="" type="checkbox"/>	ChIP-seq
<input checked="" type="checkbox"/>	Flow cytometry
<input checked="" type="checkbox"/>	MRI-based neuroimaging

Eukaryotic cell lines

Policy information about [cell lines and Sex and Gender in Research](#)

Cell line source(s)

HEK293 (T-REx) cell lines used in this study were obtained from Thermo Fisher (Cat# R71007). Cells were cultured in DMEM supplemented with 10% fetal bovine serum and 1% penicillin-streptomycin under standard conditions (37°C, 5% CO₂).

Authentication

HEK293 cells used in this study were authenticated by STR analysis at the Toronto Hospital for Sick Children authentication facility.

Mycoplasma contamination

HEK293 cells used in this study are routinely tested for mycoplasma contamination using PCR-based mycoplasma detection. Tests performed before experimentation confirmed that the cell lines are negative for mycoplasma contamination.

Commonly misidentified lines
(See [ICLAC](#) register)

The HEK293 cell line is not commonly misidentified. It is a well-characterized human embryonic kidney cell line widely recognized and validated for biological research.

Plants

Seed stocks

NA

Novel plant genotypes

NA

Authentication

NA