



Cite this: DOI: 10.1039/d6dd00012f

# A property–agnostic framework for scalable molecular inverse design *via* quantum annealing

Yuki Deguchi \* and Masato Taki

Technologies for designing molecules with desired properties have the potential to drive innovation across a wide range of fields. Molecular inverse design typically involves three key tasks: chemical latent space representation, property prediction, and molecule generation. While deep learning models trained on large molecule–property datasets can address all three tasks within a unified framework, they often require substantial property-specific retraining when targeting new molecular properties, limiting scalability. In this work, we propose an algorithmic framework that integrates machine learning and quantum annealing by explicitly decoupling chemical latent space representation, property prediction, and molecule generation. By freezing deep molecular representations learned from large datasets and separating molecule generation from deep learning, the proposed method enables inverse design for new target properties at a low training cost, requiring only lightweight model adaptation. Using quantum annealing, 98% of the generated molecules were novel and exhibited properties close to the desired targets, indicating efficient exploration beyond the training distribution. Moreover, the molecular generation rate was approximately six times faster than that of classical optimization algorithms. These results demonstrate that modularizing molecular inverse design into complementary learning and optimization components provides a scalable and effective alternative to monolithic deep learning-based approaches.

Received 12th January 2026  
Accepted 20th May 2026

DOI: 10.1039/d6dd00012f

rsc.li/digitaldiscovery

## Introduction

The discovery of novel molecules with desired properties is a crucial task across a wide range of scientific and industrial domains. For example, in drug discovery, deep generative models have been used to design and experimentally validate novel kinase inhibitors<sup>1</sup> and nuclear receptor agonists.<sup>2</sup> In materials science, machine learning has been applied to the design of organic photovoltaic materials with desired electronic properties.<sup>3,4</sup> Data-driven molecular design has also been extended to agrochemistry for the discovery of novel herbicides and insecticides.<sup>5</sup> However, the rational design of molecules across these domains remains challenging due to the vastness of the chemical space. Even when restricted to molecules composed solely of H, C, N, O, and S with up to 30 atoms, the size of the accessible chemical space has been estimated to reach approximately  $10^{60}$  possible structures.<sup>6</sup> Nevertheless, exhaustive or screening-based exploration of the chemical space remains computationally infeasible. Consequently, goal-directed molecular inverse design has emerged as a promising paradigm for accelerating innovation across these domains.

Molecular inverse design based on desired properties is commonly referred to as inverse Quantitative Structure–Activity Relationship (inverse QSAR). Inverse QSAR typically involves three essential tasks: chemical latent space representation, property prediction, and molecule generation. In recent years, deep learning-based generative models have attracted considerable attention, as they can address these tasks within a unified framework while achieving high performance.<sup>7</sup> Representative approaches include recurrent neural network (RNN)-based,<sup>8–11</sup> variational autoencoder (VAE)-based,<sup>12,13</sup> generative adversarial network (GAN)-based,<sup>14,15</sup> diffusion-based,<sup>16,17</sup> and Transformer-based models.<sup>4,18,19</sup>

In all cases, candidate molecules are generated by sampling from learned chemical representations conditioned on target property objectives. Despite methodological differences among these approaches, extending inverse molecular design to new target properties remains challenging. End-to-end generative models typically require substantial retraining when the target property changes,<sup>9,20</sup> while partially decoupled or semi-supervised approaches still depend on property-specific predictors and sufficient labeled data for each new objective.<sup>19,21</sup> In practice, many molecular properties cannot be reliably estimated *in silico*, and experimental acquisition of property labels remains costly and time-consuming, making it difficult to obtain sufficient data for scalable adaptation.<sup>22</sup> As a result, the applicability of existing deep learning-based inverse

Graduate School of Artificial Intelligence and Science, Rikkyo University, Tokyo, Japan.  
E-mail: hijoo.guchi@gmail.com; taki\_m@rikkyo.ac.jp



QSAR methods across diverse and previously unseen property objectives remains fundamentally limited.

Recently, molecular inverse design methods that integrate deep learning with additional optimization or simulation techniques have been actively explored. Representative examples include genetic algorithms for the design of organic emitters with inverted singlet–triplet gaps,<sup>23,24</sup> quantum annealing for inverse design of molecules with targeted drug-likeness,<sup>21,25</sup> and docking simulations for structure-based generation of molecules with high binding affinity to specific protein targets.<sup>26</sup> These hybrid approaches relax the reliance on purely end-to-end deep learning models and provide greater flexibility in the optimization process. However, in most existing hybrid frameworks, the optimization objective remains tightly coupled to task-dependent evaluation models or simulation pipelines, such as property-specific surrogate models used to evaluate fitness in genetic algorithms<sup>24,27</sup> or computationally expensive physics-based docking calculations required for each target protein.<sup>26</sup> As a result, adapting these methods to new molecular properties often incurs substantial computational and data costs, such as retraining surrogate models or reconfiguring simulation pipelines for each new objective, limiting their scalability. These observations highlight the need for a molecular inverse design framework that can adapt to new target properties without reconfiguring the entire optimization pipeline.


In this study, we propose a property-agnostic molecular inverse design framework that integrates machine learning and quantum annealing (Fig. 1). To address computational limitations in existing approaches, the framework explicitly decouples chemical latent space representation, property prediction, and molecule generation. A pretrained deep neural network is used

to provide a general molecular representation that is shared across tasks and kept fixed. When designing molecules for a new target property, only a lightweight task-specific prediction head is adapted, without reconfiguring the overall optimization pipeline or retraining deep generative models. Because the prediction head has a simple functional form and operates on fixed molecular descriptors, the inverse design problem can be formulated as a quadratic unconstrained binary optimization (QUBO) problem. This QUBO formulation is central to the proposed framework, as it transforms the molecular inverse design problem into a well-defined combinatorial optimization that can be solved by a variety of solvers. Among them, quantum annealing is particularly well suited, as its solution-searching dynamics, driven by quantum tunneling effects, can efficiently navigate the rugged optimization landscapes that arise when multiple property constraints must be satisfied simultaneously.<sup>28</sup> Solving this QUBO *via* quantum annealing enables the efficient identification of molecular descriptors satisfying the desired property objectives, which are subsequently decoded into molecular structures. Unlike prior approaches that focus on high-accuracy inverse estimation for specific properties, the proposed framework enables scalable adaptation to new objectives while maintaining high levels of molecular novelty, synthesizability, and diversity, comparable to or exceeding those reported in previous studies.<sup>19,21</sup>

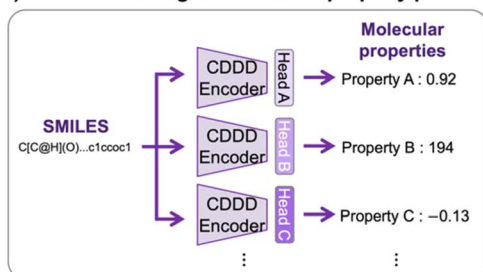
The major contributions of this study are summarized as follows:

(1) We propose a property-agnostic framework for molecular inverse design that integrates machine learning and quantum annealing through feature-based transfer learning. By fixing a learned molecular representation and adapting only a lightweight task-specific prediction head for each target property,

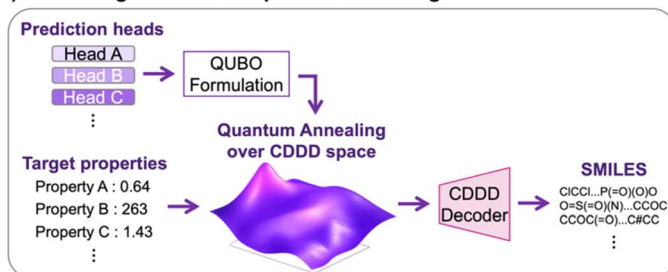
### a) Chemical latent space representation



### b) Transfer learning for molecular property prediction



### c) Molecule generation via quantum annealing



**Fig. 1** Overview of the proposed molecular inverse design framework. (a) Molecular structures are encoded into a chemical latent space using the pretrained Continuous and Data-driven Molecular Descriptor (CDDD) encoder and decoded back into Simplified Molecular Input Line Entry System (SMILES). (b) Lightweight task-specific prediction heads are adapted to the CDDD representations, enabling transfer learning for molecular property prediction. (c) Given target property values, the inverse optimization problem on the CDDD space is formulated as a quadratic unconstrained binary optimization (QUBO) and solved by quantum annealing to generate molecules satisfying the desired properties.



the proposed framework enables efficient inverse design without retraining deep generative models.

(2) We formulate molecular inverse design as a QUBO problem defined on pretrained molecular descriptors, allowing direct optimization by quantum annealing. This formulation decouples molecule generation from property-specific deep generative model training while preserving alignment with target property objectives.

(3) We demonstrate that the proposed framework generates molecules that closely match desired property values while achieving high novelty, synthesizability, and diversity. In addition, quantum annealing enables a molecular generation rate approximately six times faster than that of classical optimization algorithms.

## Results

### Overview of the molecular inverse design workflow

In this study, we propose a molecular inverse design workflow that integrates feature-based transfer learning with quantum annealing (QA), a metaheuristic optimization method that exploits quantum mechanical effects to solve combinatorial optimization problems.<sup>25</sup> A pretrained deep neural network is used to extract a general molecular representation, which is kept fixed across tasks, while molecular properties are modeled using lightweight task-specific prediction heads adapted for each target property. QA is then employed as an optimization backend to solve the inverse design problem defined on the molecular descriptor space. As illustrated in Fig. 1a, molecules are represented as strings using the Simplified Molecular Input Line Entry System (SMILES).<sup>29–31</sup> These SMILES strings are encoded into 512-dimensional continuous vectors using a pretrained RNN-based autoencoder from the Continuous and Data-driven Molecular Descriptor (CDDD) framework,<sup>32</sup> which learns continuous molecular representations by translating between equivalent chemical notations of a large corpus of molecular structures. The pretrained CDDD encoder is used without retraining to obtain fixed molecular descriptors defining a chemical latent space. As shown in Fig. 1b, molecular property prediction is performed by adapting lightweight task-specific prediction heads operating on the fixed CDDD representation. Because these prediction heads have simple functional forms and operate on fixed molecular descriptors, the inverse problem of identifying descriptor values that satisfy desired property objectives can be formulated as a QUBO problem and efficiently solved using QA. Finally, as illustrated in Fig. 1c, the molecular descriptors obtained through this inverse optimization process are decoded back into molecular structures using the pretrained CDDD decoder.<sup>32</sup> Through this workflow, molecules satisfying an arbitrary number of specified property constraints can be generated in a scalable and property-agnostic manner.

### Capability of molecular property prediction

In the proposed framework, molecular property prediction is formulated as a lightweight task-specific adaptation on top of a fixed deep molecular representation, rather than as end-to-

end deep learning training. Specifically, molecular properties are predicted by adapting linear prediction heads operating on pretrained CDDD representations. State-of-the-art deep learning models can achieve high predictive performance when large amounts of high-quality labeled data are available. However, such models typically require substantial retraining for each new target property, which limits scalability. In contrast, the proposed transfer learning strategy requires significantly less computation and training data, enabling efficient adaptation to new molecular properties. In this study, CDDD is used as a fixed molecular representation for property prediction. Leveraging this pretrained representation allows the lightweight prediction heads to achieve strong predictive performance despite their simplicity.

Table 1 reports the predictive performance of the proposed approach alongside two representative deep learning models, Structure–Property Multi-Modal foundation model (SPMM)<sup>19</sup> and Ajagekar's model.<sup>21</sup> The evaluation datasets differ across the three methods. The proposed method was evaluated on 100 000 molecules from ZINC-22,<sup>33</sup> PubChem<sup>34</sup> and ChEMBL<sup>35</sup> using a scaffold-based split, which ensures that no molecular scaffolds are shared between the training and test sets. In contrast, Ajagekar's model was evaluated on a 1000-molecule test set drawn from a 12 000-molecule ZINC 12 (ref. 36) subset using a random split,<sup>21,37</sup> and SPMM was evaluated on 1000 unseen molecules from ZINC 15.<sup>19,38</sup> Although the specific test molecules used in the prior studies are not publicly available, all three methods draw from ZINC-derived databases and predict the same RDKit-computed molecular properties, and structural overlap among the datasets is expected. Notably, the proposed method employs a larger and more diverse evaluation set with a stricter splitting strategy than either baseline. Despite these differences in evaluation conditions, the proposed method achieves competitive or superior performance in terms of the coefficient of determination ( $R^2$ ), root mean squared error (RMSE) and mean absolute error (MAE), while using substantially fewer trainable parameters. These results suggest that feature-based transfer learning on fixed deep molecular representations provides an effective and scalable approach to molecular property prediction.

To verify the effectiveness of the CDDD representation itself, a comparison against extended-connectivity fingerprints (ECFP; radius 2 and 2048 bits),<sup>39</sup> one of the most widely used molecular representations in cheminformatics, was conducted under otherwise identical conditions. The CDDD-based models outperformed the ECFP baseline across all properties, with a mean  $R^2$  of  $0.955 \pm 0.060$  compared to  $0.862 \pm 0.074$ , confirming that the pretrained CDDD representation captures richer molecular information than standard fixed-length fingerprints.

### Selectivity, efficiency, and scalability of molecule generation

Two experiments were conducted to evaluate the selectivity, efficiency, and scalability of the proposed molecule generation framework (Table 2). In these experiments, three molecular properties commonly used as indicators of drug-likeness were targeted:<sup>40–42</sup> Quantitative Estimate of Drug-likeness (QED),<sup>43</sup>



**Table 1** Molecular property predictive performance of our models, Ajagekar's model<sup>21</sup> and SPMM.<sup>19</sup> ECFP (radius 2 and 2048 bits) are included as a baseline to benchmark the effectiveness of the CDDD representation. The evaluated properties include commonly used indicators of drug-likeness in drug discovery.<sup>40–42</sup> (QED: Quantitative Estimate of Drug-likeness,<sup>43</sup> SAS: Synthetic Accessibility Score,<sup>44</sup> log *P*: Wildman–Crippen partition coefficient,<sup>45</sup> TPSA: Topological Polar Surface Area<sup>46</sup>)

Target property	Our models			Ajagekar's	SPMM	ECFP			
	MAE	RMSE	<i>R</i> <sup>2</sup>	MAE	RMSE	<i>R</i> <sup>2</sup>	MAE	RMSE	<i>R</i> <sup>2</sup>
QED	<b>0.040</b>	0.053	0.828	0.10	<b>0.047</b>	<b>0.896</b>	0.055	0.070	0.694
SAS	<b>0.167</b>	<b>0.214</b>	<b>0.882</b>	0.66	—	—	0.180	0.235	0.857
Log <i>P</i>	<b>0.092</b>	<b>0.120</b>	<b>0.991</b>	1.27	0.410	0.901	0.429	0.547	0.821
Molecular weight	<b>4.682</b>	<b>6.344</b>	<b>0.967</b>	—	12.666	0.953	13.575	18.098	0.941
TPSA	<b>1.200</b>	<b>1.565</b>	<b>0.996</b>	—	3.104	0.985	6.938	8.996	0.861
Number of hydrogen bond acceptors	<b>0.093</b>	<b>0.118</b>	<b>0.995</b>	—	0.557	0.871	0.480	0.618	0.856
Number of hydrogen bond donors	<b>0.073</b>	0.103	<b>0.991</b>	—	<b>0.094</b>	0.990	0.254	0.345	0.897
Number of rotatable bonds	<b>0.324</b>	<b>0.418</b>	<b>0.966</b>	—	0.758	0.868	0.535	0.700	0.904
Number of aromatic rings	<b>0.075</b>	<b>0.114</b>	<b>0.982</b>	—	0.336	0.848	0.172	0.226	0.931

**Table 2** Summary of experimental conditions. Experiments (a) and (b) evaluate multi-objective generation with QA and SA as a classical baseline. Experiments (c) and (d) evaluate single-objective controllability in opposite directions

Experiment	Annealing method	Target conditions	Number of generations
(a)	QA	QED > 0.7 & SAS < 3	1000
(a)'	SA	QED > 0.7 & SAS < 3	1000
(b)	QA	QED > 0.7 & log <i>P</i> < 1	1000
(b)'	SA	QED > 0.7 & log <i>P</i> < 1	1000
(c)	QA	QED > 0.7	100
(d)	QA	QED < 0.7	100

Synthetic Accessibility Score (SAS),<sup>44</sup> and Wildman–Crippen partition coefficient (log *P*).<sup>45</sup>

(a) Generation of molecules with high drug-likeness and high synthetic accessibility (QED > 0.7 and SAS < 3).

(b) Generation of molecules with high drug-likeness and reduced lipophilicity (QED > 0.7 and log *P* < 1).

In both experiments, molecules were generated by solving the QUBO-based inverse optimization problem *via* QA under the

specified property constraints. The generation was performed within the CDDD latent space learned from organic compounds with molecular weights in the range of 12–600. Such combinations of property targets are of practical importance in drug discovery,<sup>40,44</sup> yet are typically sparsely represented in available chemical datasets, making them challenging targets for data-driven generation. Under each condition, 1000 molecules were generated using both QA and classical simulated annealing (SA),<sup>47</sup> with SA serving as a widely used classical baseline optimization method for comparison.

Table 3 reports the annealing error relative to the target property values, the deviation between predicted and target properties, as well as the validity, uniqueness, and novelty of the generated molecules. The extremely low annealing errors (MAE = 0.000) indicate that the underlying optimization problems were solved accurately. This result is expected by design. Because the prediction heads are linear functions of the CDDD descriptors, the inverse optimization reduces to a quadratic problem that QA can solve to near-exact optimality. The annealing MAE therefore reflects the precision of the QUBO solver, not the accuracy of the generated molecules' actual properties. The latter is captured by the MAE (actual value) column, which quantifies the deviation between the RDKit-

**Table 3** Annealing conditions and generated results. Validity, uniqueness, and novelty are defined as the proportions of chemically valid, structurally distinct, and previously unreported molecules among the generated samples, respectively

	Annealing	Target	Target	MAE	MAE	Validity	Uniqueness	Novelty <sup>a</sup>
	Method	Property	Value	Annealing	Actual value			
(a)	QA	QED	0.85	0.000	0.107	54.1%	99.1%	98.0%
		SAS	2.5	0.000	0.883			
(a)'	SA	QED	0.85	0.000	0.184	56.0%	99.2%	95.2%
		SAS	2.5	0.000	1.057			
(b)	QA	QED	0.85	0.000	0.112	51.9%	98.6%	97.8%
		Log <i>P</i>	0.0	0.000	1.171			
(b)'	SA	QED	0.85	0.000	0.213	35.9%	97.0%	93.0%
		Log <i>P</i>	0.0	0.000	1.976			

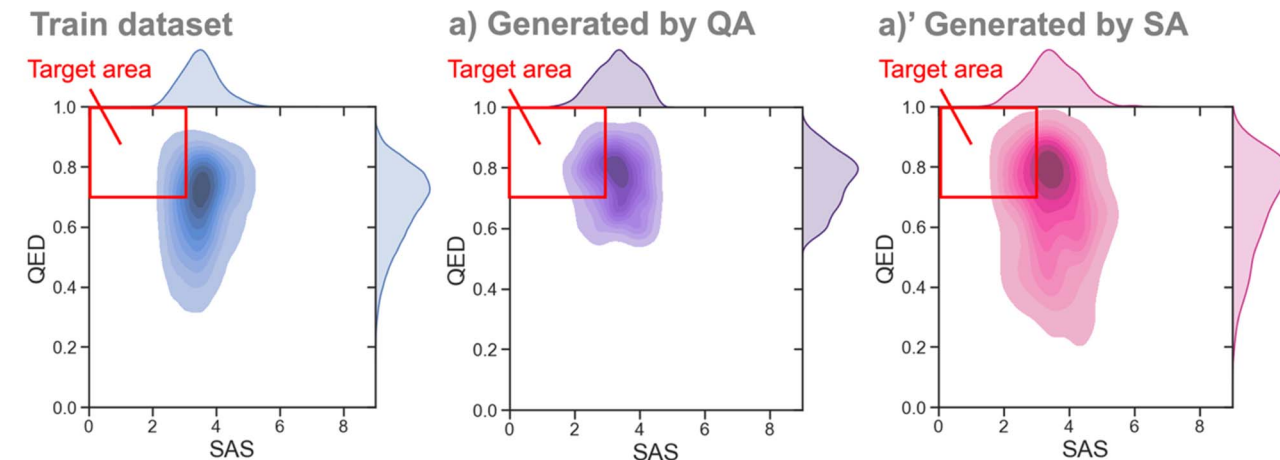
<sup>a</sup> Novelty was only calculated for valid molecules.



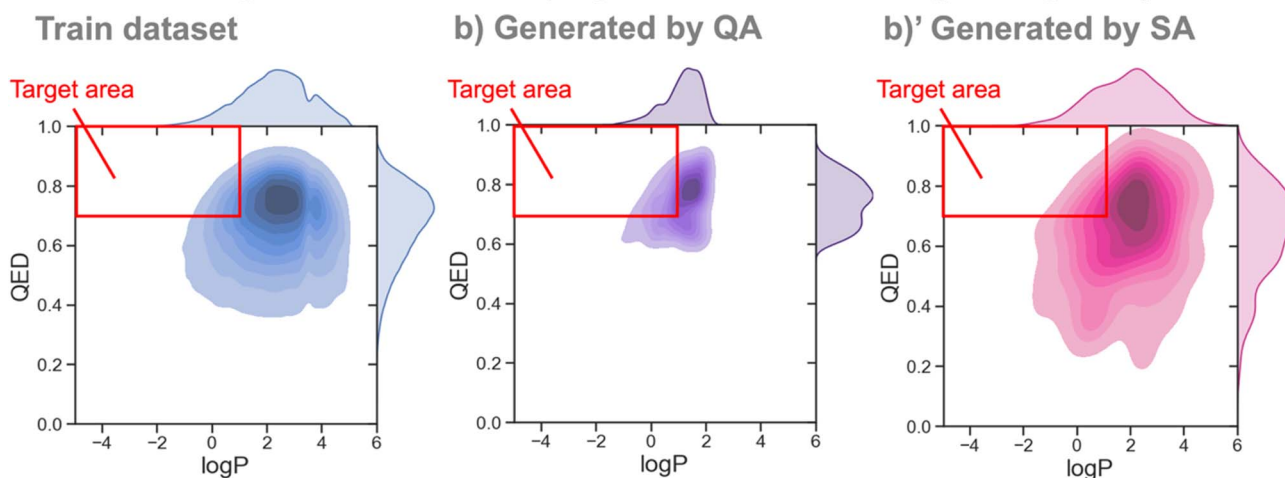
computed<sup>48</sup> properties of the decoded molecules and the target values. The nonzero actual MAE arises primarily from the approximation introduced during CDDD decoding, where the optimized descriptor vector is mapped back to a discrete molecular structure.

Fig. 2 shows the QED-SAS and QED-log *P* distributions of the generated molecules, estimated using kernel density estimation (KDE).<sup>49,50</sup> Table 4 summarizes the proportions of molecules satisfying the target property criteria in both the training dataset and the generated samples. In both experiments, molecules

## • QED and SAS distributions (High QED and low SAS targeted)



## • QED and log *P* distributions (High QED and low log *P* targeted)



## • Examples of generated molecular structures




Fig. 2 Distributions of the target properties of the training dataset and the generated molecules, and examples of generated molecular structures.



**Table 4** Proportion of the training dataset and the generated molecules that match the target property values. Experiments (a) and (a)' targeted high QED and low SAS. Experiments (b) and (b)' targeted high QED and low log  $P$

Data source	QED > 0.7	QED > 0.8	QED > 0.9	SAS < 3.0	SAS < 2.5	SAS < 2.0	QED > 0.7 & SAS < 3.0
(a) Generated by QA	<b>70.2%</b>	<b>34.4%</b>	<b>3.3%</b>	<b>32.0%</b>	<b>11.8%</b>	<b>2.4%</b>	<b>24.4%</b>
(a)' Generated by SA	52.1%	24.5%	2.0%	25.2%	9.5%	1.8%	15.4%
Training dataset	46.3%	14.1%	0.8%	20.2%	4.1%	0.1%	7.9%
Data source	QED > 0.7	QED > 0.8	QED > 0.9	Log $P$ < 1.0	Log $P$ < 0.0	Log $P$ < -1.0	QED > 0.7 & log $P$ < 1.0
(b) Generated by QA	<b>66.3%</b>	<b>25.0%</b>	<b>0.2%</b>	<b>35.8%</b>	7.3%	0.8%	<b>18.5%</b>
(b)' Generated by SA	40.4%	14.5%	<b>0.8%</b>	26.2%	<b>8.6%</b>	<b>2.2%</b>	5.0%
Training dataset	46.3%	14.1%	0.8%	16.9%	5.1%	1.2%	4.7%

generated by QA and SA exhibit property distributions that are substantially more concentrated within the target regions than those of the training dataset, demonstrating effective and selective inverse design.

Notably, while QA and SA exhibit comparable performance in terms of molecular validity and novelty, QA consistently achieves a substantially higher success rate in generating molecules that simultaneously satisfy multiple target property constraints. As shown in Table 4, the proportion of molecules meeting both QED and SAS or QED and log  $P$  criteria is significantly higher for QA than for SA or the training dataset. This indicates that QA provides enhanced selectivity in multi-objective molecular inverse design, rather than merely increasing diversity or random exploration. Such behavior is particularly relevant in constrained optimization settings, where feasible solutions occupy narrow and sparsely distributed regions of the search space, and the ability of QA to navigate complex and rugged optimization landscapes enables more effective identification of solutions satisfying multiple competing objectives. Further analysis of the distribution of generated molecules in the chemical latent space, discussed in the following section, suggests that this enhanced selectivity is accompanied by controlled exploration beyond the training distribution.

In terms of computational efficiency, QA achieved a mean end-to-end computation time of  $9.55 \pm 0.63$  s per molecule in experiment a), whereas SA required  $54.65 \pm 0.45$  s, corresponding to an approximately 5.7-fold speedup. Notably, the actual computation time of the quantum annealer itself was only  $82 \pm 11$  ms. The remaining overhead arises primarily from classical components of the hybrid workflow, including application programming interface (API) communication, solver orchestration, and neural network-based decoding of molecular representations. This indicates that the observed latency reflects implementation-level constraints rather than fundamental limitations of QA.

Overall, these results demonstrate that the proposed framework enables selective and efficient generation of molecules with specified property profiles. Importantly, within this architecture, adapting inverse molecular design to a new target property requires only the lightweight adaptation of a task-specific prediction head on a fixed molecular representation,

without retraining deep generative models. As a result, the proposed method achieves high selectivity, computational efficiency, and scalability across a broad range of molecular design objectives.

### Novelty, synthesizability and diversity of generated molecules

The novelty, synthesizability, and diversity of generated molecules are critical factors in molecular inverse design. Low novelty reduces experimental efficiency, while poor synthesizability hinders downstream experimental validation and practical application. In addition, insufficient diversity may lead to biased molecular structures and limited chemical exploration.


In this study, validity is defined as the proportion of generated SMILES that can be successfully parsed into chemically valid molecular structures by RDKit.<sup>48</sup> Uniqueness refers to the proportion of valid molecules with distinct canonical SMILES. Novelty is defined as the proportion of valid molecules whose canonical SMILES do not appear in the reference dataset of approximately 22 billion molecules. Table 3 reports the novelty and uniqueness of molecules generated under each experimental condition. 98.0% of the molecules generated by QA in the high QED and low SAS setting were novel, and the uniqueness reached 99.1%. These results indicate that the proposed framework does not rely on brute-force search within the training dataset, but instead samples molecular candidates from a well-generalized chemical latent space.

To examine whether the novelty arises from entirely new core structures or from modifications of known molecular frameworks, we decomposed each generated molecule into its core ring system and linkers (Murcko scaffolds<sup>51</sup>) and compared them against the reference dataset. Of the unique scaffolds identified, 64.0% were absent from the reference set, indicating that the framework generates molecules with entirely new core structures. The remaining 36.0% shared scaffolds with the reference set, yet these molecules were still novel at the full-structure level. Analysis of the modification patterns on these known scaffolds revealed that the modifications extend well beyond simple single-site substitutions: on average, 8.8 heavy atoms were added beyond the scaffold, with 93.8% of molecules containing more than three additional heavy atoms. Furthermore, 85.1% of these molecules introduced atom types not

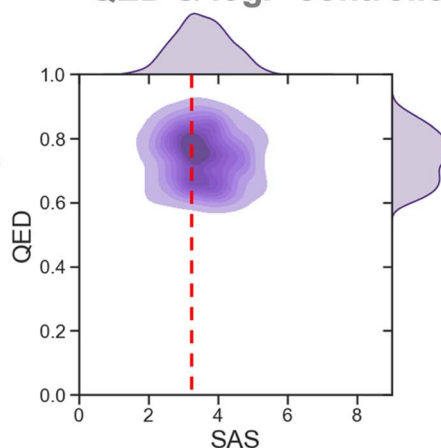


## • SAS and QED distributions

1) Ajagekar *et al.*  
QED controlled




b) Ours  
QED & log*P* controlled



## • SAS and log*P* distributions

2) Ajagekar *et al.*  
log*P* controlled



b) Ours  
QED & log*P* controlled




Fig. 3 Distributions of SAS and the target properties of the molecules generated by our method and Ajagekar's model.<sup>21</sup> The red line represents the mean SAS.

present in their scaffold, such as halogens and heteroatoms. These findings suggest that the novelty of generated molecules arises from two complementary mechanisms: the generation of entirely new core structures through interpolation and extrapolation in the continuous CDDD latent space, and non-trivial multi-site functional group diversification on known frameworks.

A trade-off between molecular novelty and validity was observed, with the validity of generated molecules being approximately 50%. This behaviour reflects a deliberate design choice. While many likelihood-based generative models commonly adopted in molecular design are optimized to generate syntactically valid molecules close to the training distribution, the proposed optimization-driven framework actively explores out-of-distribution regions of the learned chemical latent space. As a result, highly novel and diverse


molecular structures can be identified, albeit at the cost of reduced validity. Importantly, this limitation is not fundamental, and validity is expected to improve through the incorporation of more expressive generative backbones, such as transformer-based architectures, or more robust molecular string representations such as Self-referencing Embedded Strings (SELFIES).<sup>52</sup>

Fig. 3 illustrates the QED-SAS and log*P*-SAS distributions estimated using KDE for molecules generated in experiment (b) and by Ajagekar's model. Table 5 summarizes the mean SAS values for each method. In experiment (a), SAS was explicitly constrained to low values, resulting in a large fraction of molecules with high synthetic accessibility. In experiment (b), only QED and log*P* were controlled, while SAS was not directly optimized. Nevertheless, the mean SAS value of the generated molecules ( $3.555 \pm 0.787$ ) closely approached that of the



**Table 5** Mean SAS of the training dataset and the molecules generated by our method and Ajagekar's model.<sup>21</sup> SAS ranges from 1 (easy to synthesize) to 10 (difficult to synthesize).<sup>44</sup>

Data source	Target property	Mean SAS
(b) Generated by our method	QED & log <i>P</i>	3.555 ± 0.787
(1) Generated by Ajagekar's model	QED	6.516 ± 0.931
(2) Generated by Ajagekar's model	Log <i>P</i>	7.023 ± 0.832
Training dataset	—	3.522 ± 0.622



**Fig. 4** CDDD reduced by UMAP of the training dataset and the generated molecules.

training dataset ( $3.522 \pm 0.622$ ). In contrast, when controlling log *P*, Ajagekar's model produced molecules with a substantially higher mean SAS value of  $7.023 \pm 0.832$ , indicating increased synthetic difficulty in this setting. Notably, the proposed framework generates molecules with lower SAS (*i.e.*, improved synthetic accessibility) than Ajagekar's model, even without explicitly optimizing SAS. This advantage arises from performing optimization within a fixed chemical latent space learned from chemically valid molecules, which naturally constrains the search to structurally plausible regions. As a result, the method satisfies target property objectives while avoiding excessive

structural complexity, leading to improved synthetic accessibility relative to previous studies.

Fig. 4 visualizes the distribution of molecular representations using Uniform Manifold Approximation and Projection (UMAP).<sup>53,54</sup> Most molecules generated by QA and SA were sampled from regions overlapping with the training distribution, while a subset of molecules extended into previously unexplored regions of the chemical latent space. Notably, molecules with unconventional structural motifs, such as phosphite and tetrahydro-1,2,5-triazepine, were observed. This diversity may be influenced by thermodynamic noise<sup>55</sup> and solution-searching behaviour associated with quantum tunnelling effects in QA.<sup>28</sup> Together, these observations indicate that the proposed method enables broad and diverse exploration of chemical space beyond the training distribution.

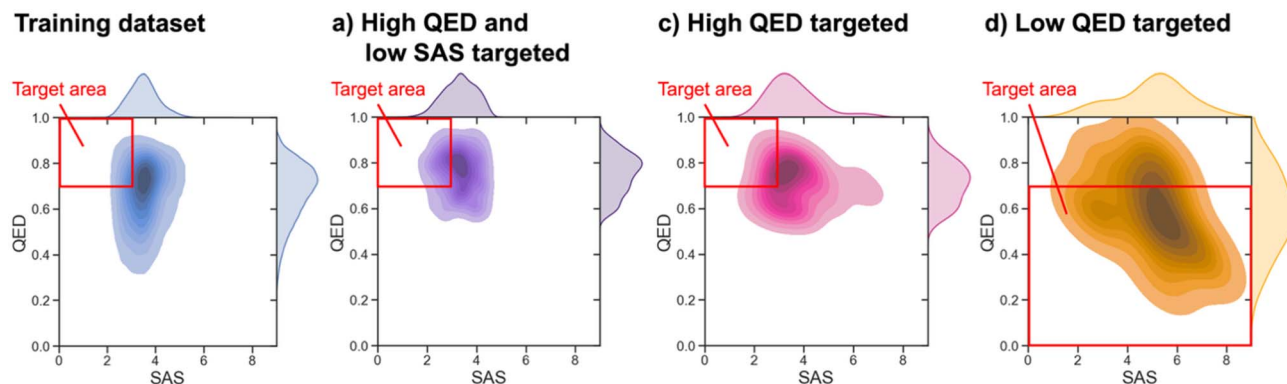
Overall, these results demonstrate that the proposed framework generates molecules that are highly novel, synthetically accessible, and structurally diverse. Given that molecular design under specific objectives often remains constrained to known chemical scaffolds and their analogues,<sup>56</sup> the ability to systematically explore out-of-distribution chemical space represents a key advantage of the proposed approach over conventional human-centered molecular design strategies.

## Performance study

### Multiple molecular properties control

Experiment a) demonstrated that molecules with high QED and low SAS can be generated by explicitly controlling multiple molecular properties simultaneously. To further assess whether the proposed framework genuinely controls multiple properties, rather than merely biasing generation toward high-QED molecules, experiment (c) was conducted in which only QED was controlled. The target value was set to QED > 0.7, and 100 molecules were generated using QA.

Fig. 5(a) shows the QED-SAS distribution of molecules generated in experiment (a), while Fig. 5(c) shows the distribution obtained in experiment (c). In both experiments, the generated molecules exhibit higher QED values than those in the training dataset. Quantitatively, experiment a) achieved



**Fig. 5** Distributions of QED and SAS of the training dataset and generated molecules: (a) high QED and low SAS targeted, (c) high QED targeted and (d) low QED targeted.



a mean QED of  $0.753 \pm 0.089$  with a corresponding mean SAS of  $3.307 \pm 0.646$ . In contrast, when only QED was controlled in experiment (c), the generated molecules showed a lower mean QED of  $0.709 \pm 0.087$  and a higher mean SAS of  $3.558 \pm 0.970$ . These results indicate that simultaneous control of QED and SAS in experiment (a) systematically shifts the generated molecules toward lower synthetic accessibility scores while maintaining higher drug-likeness, compared with single-objective optimization. This demonstrates that the proposed framework enables explicit and simultaneous control of multiple molecular properties, rather than implicitly favoring a single dominant objective.

### Direction of control

Experiment (c) focused on generating molecules with high QED. However, because molecules with high QED are more prevalent in the training dataset than those with low QED, it is possible that such generation reflects the underlying data distribution rather than true controllability. Although high QED is generally desirable, the optimal direction of control for a given molecular property depends on the intended application. To verify whether the proposed framework supports bidirectional control of molecular properties, experiment (d) was conducted with the objective of generating molecules with low QED. The target value was set to  $\text{QED} < 0.7$ , and 100 molecules were generated using QA.

Fig. 5(d) illustrates the QED-SAS distribution of molecules generated in experiment (d). While molecules generated in experiment (c) exhibit higher QED values than those in the training dataset, molecules generated in experiment (d) are clearly shifted toward lower QED values relative to experiment (c). These results demonstrate that the proposed framework enables selective and bidirectional control of molecular properties, rather than passively reproducing trends present in the training data.

## Discussion

The proposed framework demonstrates a scalable and flexible approach to molecular inverse design by decoupling chemical latent space representation, property prediction, and molecule generation into complementary components. Unlike many existing inverse design methods that prioritize high-accuracy optimization of a small number of predefined molecular properties, the present work focuses on improving task scalability by enabling efficient adaptation to new target properties through feature-based transfer learning and optimization-driven molecule generation.

A key characteristic of the proposed framework is the deliberate separation of deep representation learning from downstream property optimization. By fixing a pretrained molecular representation and adapting only a lightweight task-specific prediction head, the framework significantly reduces the computational and data requirements associated with introducing new molecular objectives. This design choice allows inverse molecular design to be formulated as

a constrained optimization problem over a learned chemical latent space, which can be efficiently solved using quantum annealing. As demonstrated in the Results section, this approach enables the simultaneous control of multiple molecular properties while maintaining strong selectivity toward target property regions.

The observed trade-off between molecular novelty and validity warrants careful discussion. The validity of generated molecules is lower than that reported for likelihood-based generative models optimized for syntactic correctness and sampling efficiency, such as transformer-based or reinforcement learning-driven SMILES generators. However, this difference reflects a fundamental distinction in design objectives rather than a limitation of the proposed method. Likelihood-based generative models are typically trained to maximize the probability of generating molecules close to the training distribution, resulting in high validity but limited exploration beyond known chemical space. In contrast, the proposed optimization-driven framework explicitly targets constrained inverse problems and actively explores out-of-distribution regions of the learned chemical latent space. This design choice leads to substantially higher molecular novelty and diversity, albeit at the cost of reduced validity. However, this trade-off is not fundamental and can be mitigated through architectural and representational improvements.

Computational efficiency represents another important consideration. Although the end-to-end generation time per molecule is on the order of several seconds, the actual quantum annealing time required to solve the underlying QUBO problem is on the order of milliseconds. The dominant sources of latency arise from implementation-level factors, including hybrid solver orchestration, API communication overhead, and neural network-based decoding of molecular representations. These observations indicate that the current generation speed is not a fundamental limitation of the framework itself, but rather reflects the present software and hardware integration. Future improvements in quantum hardware accessibility, solver parallelization, and decoder optimization are expected to substantially reduce end-to-end generation times. Moreover, classical alternatives such as simulated or momentum annealing accelerated on graphics processing unit (GPU) may serve as practical complements when quantum hardware availability is constrained.<sup>57,58</sup>

An important advantage of quantum annealing observed in this study is its enhanced selectivity in constrained multi-objective optimization. While quantum annealing and simulated annealing achieve comparable levels of molecular validity and novelty, quantum annealing consistently yields a higher proportion of molecules that simultaneously satisfy multiple target property constraints. This behaviour is particularly valuable in molecular inverse design problems where feasible solutions occupy narrow and sparsely distributed regions of the search space. The solution-searching dynamics of quantum annealing, influenced by tunneling effects and thermodynamic noise, appear well suited to navigating such rugged optimization landscapes, thereby improving the efficiency of constraint satisfaction rather than merely increasing random exploration.



Overall, the proposed framework provides a complementary alternative to monolithic deep learning-based molecular generators. Rather than competing with likelihood-based models in terms of raw generation speed or syntactic validity, the present approach is designed to support scalable, property-agnostic inverse design under complex and evolving constraints. This capability is particularly relevant in early-stage molecular discovery, where exploration beyond known chemical scaffolds and rapid adaptation to new design objectives are essential. By modularizing molecular inverse design into representation learning, lightweight transfer learning, and optimization-based generation, the proposed framework offers a promising foundation for future hybrid approaches that combine machine learning and advanced optimization technologies.

## Methods

### Chemical latent space representation

In this study, we adopt Continuous and Data-Driven Molecular Descriptors (CDDD), which are obtained using a pretrained RNN-based autoencoder,<sup>32</sup> as the chemical latent space representation. CDDD encodes molecular structures, provided as SMILES strings, into continuous 512-dimensional vectors with values in the range of  $[-1, +1]$ . This continuous representation is well suited for optimization-based molecular inverse design, as it enables smooth navigation of chemical latent space and facilitates inverse mapping from descriptors to molecular structures.

A key advantage of CDDD is its reversibility, as molecular descriptors can be decoded back into valid SMILES strings using the corresponding pretrained decoder. This property is essential for the proposed framework, which relies on inverse optimization of molecular descriptors followed by molecule reconstruction. In contrast to task-specific or end-to-end trained generative models, the CDDD encoder–decoder is kept fixed throughout this study, and no retraining is performed. As a result, the learned chemical representation can be reused across different molecular property objectives, forming a stable foundation for feature-based transfer learning. All CDDD representations in this study were computed using the pretrained models released by Winter *et al.*<sup>32</sup>

### Property prediction

Lightweight task-specific prediction heads operating on fixed CDDD representations were adapted for each target molecular property. In this study, these prediction heads were implemented as linear models, which were selected due to their scalability, ease of training, and suitability for inverse optimization. Owing to their simple functional form, the inverse estimation of molecular descriptors can be formulated as a quadratic unconstrained binary optimization (QUBO) problem, enabling molecule generation *via* quantum annealing (QA).

Training data were collected from ZINC-22,<sup>33</sup> PubChem,<sup>34</sup> and ChEMBL.<sup>35</sup> ZINC-22 (37 223 934 261 molecules), PubChem (115 347 688 molecules), and ChEMBL (2 399 743 molecules)

datasets available as of 16 October 2023 were combined and preprocessed to meet the requirements for CDDD calculation.<sup>32</sup>

(1) Ionic bonds were cleaved and the largest molecular fragments were retained.

(2) Stereochemical information was removed, molecules were converted to canonical SMILES using RDKit,<sup>48</sup> and duplicates were removed.

(3) Molecules containing atoms other than H, B, C, N, O, F, P, S, Cl, Br, and I were removed.

(4) Molecules with  $\log P$  outside the range of  $[-7, 5]$ , molecular weight outside  $[12\ 600]$ , or fewer than three heavy atoms were removed.

Of the initial 37 341 681 692 molecules, approximately 15.3 billion (40.9%) were removed during preprocessing, yielding a final dataset of 22 080 110 355 molecules, which includes all 72 million molecules used to train the CDDD autoencoder.<sup>32</sup> The vast majority of the removal (approximately 14.6 billion, 39.0%) occurred in step 2, while steps 3 and 4 each removed comparatively small fractions of approximately 0.3% and 1.6%, respectively. It should be noted that the source databases used in this study are subject to known coverage biases. ZINC-22 primarily catalogues commercially available and make-on-demand compounds, covering only a small fraction of drug-like chemical space even at the scale of tens of billions of molecules.<sup>33</sup> More generally, widely used molecular datasets have been shown to lack uniform coverage of biomolecular structures, and models trained on such datasets may have limited predictive power for underrepresented regions of chemical space.<sup>59</sup> While no explicit bias mitigation techniques were applied in this study, the optimization-driven generation *via* quantum annealing enables exploration of out-of-distribution regions beyond the training data, partially alleviating the impact of these biases. Nevertheless, the CDDD latent space inevitably reflects the biases of its training data, and incorporating more diverse molecular sources remains a direction for future work.

From this dataset, 1 million molecules were used for training and 0.1 million molecules for testing the task-specific prediction heads using a scaffold-based split. A separate validation set was not employed, as the prediction heads were intentionally designed to have a closed-form solution with no hyperparameters to tune. This simplicity is not incidental but a deliberate design choice that enables the inverse design problem to be directly formulated as a QUBO problem, which is central to the proposed framework. The full dataset was used for evaluating the novelty of generated molecules.

All transfer learning procedures for adapting the task-specific prediction heads were performed using an NVIDIA A100 GPU (40 GB memory).

### Molecule generation

To improve scalability across diverse molecular property objectives, molecular property prediction is decoupled from deep generative modeling and formulated as a lightweight linear prediction head operating on a fixed deep molecular representation. While this feature-based transfer learning strategy



enables efficient adaptation to new target properties, it makes molecule generation non-trivial when using conventional sampling-based approaches. To address this challenge, we propose a molecule generation algorithm inspired by optimization-based inverse design methods, including genetic algorithms,<sup>24,60,61</sup> gradient-based optimization,<sup>62</sup> and QA.<sup>21</sup> In the proposed framework, molecule generation proceeds as follows:

(1) Molecular descriptors in the CDDD space that satisfy the desired property constraints are obtained by solving the inverse optimization problem defined by the linear prediction heads using QA.

(2) The obtained molecular descriptors are decoded into SMILES strings using the pretrained CDDD decoder.

QA is a stochastic optimization paradigm inspired by simulated annealing (SA), which explores the ground state of an Ising or QUBO model by introducing quantum fluctuations.<sup>63</sup> Previous studies have shown that QA can outperform SA on optimization landscapes characterized by sharp and deep local minima, owing to quantum tunneling effects.<sup>28</sup> These properties make QA particularly suitable for inverse molecular design problems, where feasible solutions satisfying multiple constraints often occupy narrow and fragmented regions of the search space.

In this study, QA was accessed *via* QA hardware provided by D-Wave Systems, Inc. through an API interface.<sup>64</sup> Current quantum annealers remain limited in terms of effective problem size due to hardware connectivity constraints. For example, although the Advantage\_system6.4 supports 5612 physical qubits, problems involving dense second-order interactions can only be mapped to approximately 180 logical variables.<sup>65</sup> This limitation arises because multiple physical qubits are required to represent a single logical variable and its interactions.<sup>66</sup>

To overcome this restriction, a hybrid binary quadratic model solver was employed (hybrid\_binary\_quadratic\_model\_version2p), which combines classical optimization with QA and can handle up to one million variables. This solver is conceptually related to qbsolv,<sup>67</sup> where large optimization problems are decomposed into smaller subproblems that are iteratively optimized using QA and classical heuristics such as SA and Tabu search.<sup>68-70</sup> While the hybrid solver introduces additional classical overhead compared to pure QA execution, it enables the practical application of

function is defined by a set of linear prediction heads corresponding to the target molecular properties. Because QA-based solvers operate on binary variables, the inverse optimization problem is formulated as a QUBO problem. Optimization algorithms based on the Ising model, including QA and SA, require the objective function to be expressed in terms of binary variables, as the system energy (Hamiltonian) is defined using spin variables  $\sigma_i \in (-1, +1)$ . An equivalent formulation is given by the QUBO model, in which the spin variables  $\sigma_i$  are replaced by binary variables  $q_i \in (0, +1)$ . The equivalence between the Ising and QUBO formulations has been well established for a wide range of NP-hard optimization problems.<sup>71,72</sup> Accordingly, the inverse molecular design problem is represented using the QUBO formulation, which provides a natural and intuitive representation compatible with QA-based optimization.

Let  $\mathbf{d}$  be the value of  $l$  desired molecular properties,  $\mathbf{W} = (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_l)^T$  be the weight of  $l$  linear prediction heads corresponding to the desired molecular properties,  $\mathbf{b}$  be the intercept of  $l$  linear prediction heads, and  $\mathbf{x} \in [-1, +1]$  be the value of the  $m = 512$  dimensional CDDD to be obtained and find CDDD that minimizes the following evaluation function.

$$L(\mathbf{x}) = \sum_{g=1}^l \left( \sum_{h=1}^m W_{gh} x_h + c_g \right)^2 \quad (1)$$

where

$$\mathbf{c} = \mathbf{b} - \mathbf{d}$$

Although each element of CDDD is a continuous value within the range of  $-1$  to  $+1$ , the QUBO model requires the use of binary variables  $\mathbf{q} \in \{0, +1\}$ . Therefore, each element  $x \in [-1, +1]$  of the CDDD is approximated using  $n$  binary variables  $\mathbf{q}$  as follows. In this study,  $n = 8$  was employed to balance computational complexity and accuracy.

$$x = \sum_{i=1}^n 2^{1-i} q_i - 1 \quad (2)$$

From eqn (1) and (2), the evaluation function is expressed as follows:

$$\begin{aligned} L(\mathbf{q}) = & \sum_{g=1}^l \left( \sum_{h=1}^m W_{gh} \left( \sum_{i=1}^n 2^{1-i} q_{f(h,i)} - 1 \right) + c_g \right)^2 = \sum_{g=1}^l \sum_{h,i=1}^m \sum_{j,k=1}^n W_{gh} W_{gi} 2^{1-j} q_{f(h,j)} 2^{1-k} q_{f(i,k)} - 2 \sum_{g=1}^l \sum_{h,i=1}^m \sum_{j=1}^n W_{gh} W_{gi} 2^{1-j} q_{f(h,j)} \\ & + 2 \sum_{g=1}^l \sum_{h=1}^m \sum_{j=1}^n c_g W_{gh} 2^{1-j} q_{f(h,j)} + \sum_{g=1}^l \sum_{h,i=1}^m W_{gh} W_{gi} - 2 \sum_{g=1}^l \sum_{h=1}^m c_g W_{gh} + \sum_{g=1}^l c_g^2 \end{aligned} \quad (3)$$

quantum-assisted optimization to high-dimensional molecular inverse design problems.

Within this framework, molecule generation is reduced to an optimization problem over the CDDD space, where the objective

where

$$f(y, z) = n(y - 1) + z$$



As the additive constant in eqn (3) can be disregarded, the Hamiltonian for the inverse estimation problem of a molecule with the desired properties can be expressed using the upper triangular matrix  $Q$  as follows:

$$\mathcal{H} = \mathbf{q}^{\text{ST}} Q \mathbf{q} \quad (4)$$

where

$$Q_{ij} = \begin{cases} \sum_{g=1}^i f(g, i) \left( f(g, i) - 2 \sum_{h=1}^m W_{gh} + 2c_g \right) & \text{if } i = j \\ \sum_{g=1}^i f(g, i) f(g, j) & \text{if } i < j \\ 0 & \text{if } i > j \end{cases}$$

$$f(y, z) = W_{ye(z)} 2^{(1-z) \bmod n}$$

$$e(z) = \frac{(z - z \bmod n)}{n} + 1$$

$\mathbf{q}$  that minimizes eqn (4) is obtained using QA, and the desired CDDD is obtained by converting  $\mathbf{q}$  to  $\mathbf{x}$  according to eqn (2). Finally, the pretrained CDDD decoder converts the CDDD into SMILES to generate molecules.

Due to thermodynamic noise and environmental fluctuations, QA does not guarantee reproducibility of individual solutions, and the probability of reaching the ground state varies across runs.<sup>55</sup> In this study, this stochasticity is intentionally exploited. Multiple annealing runs are performed under identical conditions, and each run yields a distinct molecular descriptor. Consequently, a single annealing run corresponds to the generation of one candidate molecule. In contrast, although SA can be made reproducible by fixing random seeds, no such constraints were imposed here in order to promote molecular diversity.

## Author contributions

Y. D.: conceptualization, methodology, software, validation, formal analysis, investigation, data curation, visualization, writing – original draft, writing – review & editing, funding acquisition. M. T.: supervision, writing – review & editing, funding acquisition.

## Conflicts of interest

The authors declare no competing interests.

## Data availability

This study was carried out using publicly available molecular data from ZINC-22 (<https://files.docking.org/zinc22>), PubChem (<https://pubchem.ncbi.nlm.nih.gov>), and ChEMBL (<https://www.ebi.ac.uk/chembl>). Data on the molecules and properties

generated in this study are available at [https://github.com/hijooguchi/qa\\_molecular\\_inverse\\_design](https://github.com/hijooguchi/qa_molecular_inverse_design) with DOI: <https://doi.org/10.5281/zenodo.20039825>.

Code availability: The source code for this study is available at [https://github.com/hijooguchi/qa\\_molecular\\_inverse\\_design](https://github.com/hijooguchi/qa_molecular_inverse_design) with DOI: <https://doi.org/10.5281/zenodo.20039825>.

## Acknowledgements

This work was supported by The Public Foundation of Chubu Science and Technology Center.

## References

- 1 A. Zhavoronkov, Y. A. Ivanenkov, A. Aliper, M. S. Veselov, V. A. Aladinskiy, A. V. Aladinskaya, V. A. Terentiev, D. A. Polykovskiy, M. D. Kuznetsov, A. Asadulaev, Y. Volkov, A. Zholus, R. R. Shayakhmetov, A. Zhebrak, L. I. Minaeva, B. A. Zagribelnyy, L. H. Lee, R. Soll, D. Madge, L. Xing, T. Guo and A. Aspuru-Guzik, *Nat. Biotechnol.*, 2019, **37**, 1038–1040.
- 2 K. Atz, L. Cotos, C. Isert, M. Håkansson, D. Focht, M. Hilleke, D. F. Nippa, M. Iff, J. Ledergerber, C. C. G. Schiebriock, V. Romeo, J. A. Hiss, D. Merk, P. Schneider, B. Kuhn, U. Grether and G. Schneider, *Nat. Commun.*, 2024, **15**, 3408.
- 3 W. Sun, Y. Zheng, K. Yang, Q. Zhang, A. A. Shah, Z. Wu, Y. Sun, L. Feng, D. Chen, Z. Xiao, S. Lu, Y. Li and K. Sun, *Sci. Adv.*, 2019, **5**, eaay4275.
- 4 P. Yoo, D. Bhowmik, K. Mehta, P. Zhang, F. Liu, M. L. Pasini and S. Irle, *Sci. Rep.*, 2023, **13**, 20031.
- 5 Y. Nakayama, S. Morishita, H. Doi, T. Hirano and H. Kaneko, *ACS Omega*, 2024, **9**, 18488–18494.
- 6 R. S. Bohacek, C. McMartin and W. C. Guida, *Med. Res. Rev.*, 1996, **16**, 3–50.
- 7 B. Sanchez-Lengeling and A. Aspuru-Guzik, *Science*, 2018, **361**, 360–365.
- 8 D. E. Rumelhart, G. E. Hinton and R. J. Williams, *Nature*, 1986, **323**, 533–536.
- 9 A. Gupta, A. T. Müller, B. J. H. Huisman, J. A. Fuchs, P. Schneider and G. Schneider, *Mol. Inf.*, 2018, **37**, 1700111.
- 10 M. H. S. Segler, T. Kogej, C. Tyrchan and M. P. Waller, *ACS Cent. Sci.*, 2018, **4**, 120–131.
- 11 K. Kim, S. Kang, J. Yoo, Y. Kwon, Y. Nam, D. Lee, I. Kim, Y.-S. Choi, Y. Jung, S. Kim, W.-J. Son, J. Son, H. S. Lee, S. Kim, J. Shin and S. Hwang, *npj Comput. Mater.*, 2018, **4**, 67.
- 12 D. P. Kingma and M. Welling, in *Proceedings of the 2nd international conference on learning representations*, 2014.
- 13 S. Kang and K. Cho, *J. Chem. Inf. Model.*, 2019, **59**, 43–52.
- 14 I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville and Y. Bengio, in *Proceedings of the 27th international conference on neural information processing systems*, 2014, vol. 2, pp. 2672–2680.
- 15 N. De Cao and T. Kipf, *arXiv*, 2018, preprint, arXiv:1805.11973 DOI: [10.48550/arXiv.1805.11973](https://doi.org/10.48550/arXiv.1805.11973).
- 16 J. Ho, A. Jain and P. Abbeel, in *Proceedings of the 34th international conference on neural information processing systems*, 2020.



- 17 T. Weiss, E. M. Yanes, S. Chakraborty, L. Cosmo, A. M. Bronstein and R. Gershoni-Poranne, *Nat. Comput. Sci.*, 2023, **3**, 873–882.
- 18 A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser and I. Polosukhin, in *Proceedings of the 31st international conference on neural information processing systems*, 2017, pp. 6000–6010.
- 19 J. Chang and J. C. Ye, *Nat. Commun.*, 2024, **15**, 2323.
- 20 C. Bilodeau, W. Jin, T. Jaakkola, R. Barzilay and K. F. Jensen, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.*, 2022, **12**, e1608.
- 21 A. Ajagekar and F. You, *npj Comput. Mater.*, 2023, **9**, 143.
- 22 B. Dou, Z. Zhu, E. Merkurjev, L. Ke, L. Chen, J. Jiang, Y. Zhu, J. Liu, B. Zhang and G.-W. Wei, *Chem. Rev.*, 2023, **123**, 8736–8780.
- 23 J. H. Holland, *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*, U Michigan Press, 1975.
- 24 A. Nigam, R. Pollice, P. Friederich and A. Aspuru-Guzik, *Chem. Sci.*, 2024, **15**, 2618–2639.
- 25 T. Kadowaki and H. Nishimori, *Phys. Rev.*, 1998, **58**, 5355–5363.
- 26 B. Ma, K. Terayama, S. Matsumoto, Y. Isaka, Y. Sasakura, H. Iwata, M. Araki and Y. Okuno, *J. Chem. Inf. Model.*, 2021, **61**, 3304–3313.
- 27 H. Kneiding and D. Balcells, *ChemRxiv*, 2024.
- 28 V. S. Denchev, S. Boixo, S. V. Isakov, N. Ding, R. Babbush, V. Smelyanskiy, J. Martinis and H. Neven, *Phys. Rev. X*, 2016, **6**, 031015.
- 29 D. Weininger, *J. Chem. Inf. Comput. Sci.*, 1988, **28**, 31–36.
- 30 D. Weininger, A. Weininger and J. L. Weininger, *J. Chem. Inf. Comput. Sci.*, 1989, **29**, 97–101.
- 31 D. Weininger, *J. Chem. Inf. Comput. Sci.*, 1990, **30**, 237–243.
- 32 R. Winter, F. Montanari, F. Noé and D.-A. Clevert, *Chem. Sci.*, 2019, **10**, 1692–1701.
- 33 B. I. Tingle, K. G. Tang, M. Castanon, J. J. Gutierrez, M. Khurelbaatar, C. Dandarchuluun, Y. S. Moroz and J. J. Irwin, *J. Chem. Inf. Model.*, 2023, **63**, 1166–1176.
- 34 S. Kim, J. Chen, T. Cheng, A. Gindulyte, J. He, S. He, Q. Li, B. A. Shoemaker, P. A. Thiessen, B. Yu, L. Zaslavsky, J. Zhang and E. E. Bolton, *Nucleic Acids Res.*, 2023, **51**, D1373–D1380.
- 35 D. Mendez, A. Gaulton, A. P. Bento, J. Chambers, M. De Veij, E. F. María, P. Magariños, J. F. Mosquera, P. Mutowo, M. Nowotka Michałand Gordillo-Marañón, F. Hunter, L. Junco, G. Mugumbate, M. Rodriguez-Lopez, F. Atkinson, N. Bosc, C. J. Radoux, A. Segura-Cabrera, A. Hersey and A. R. Leach, *Nucleic Acids Res.*, 2019, **47**, D930–D940.
- 36 J. J. Irwin, T. Sterling, M. M. Mysinger, E. S. Bolstad and R. G. Coleman, *J. Chem. Inf. Model.*, 2012, **52**, 1757–1768.
- 37 V. P. Dwivedi, C. K. Joshi, A. T. Luu, T. Laurent, Y. Bengio and X. Bresson, *J. Mach. Learn. Res.*, 2023, **24**, 1–48.
- 38 T. Sterling and J. J. Irwin, *J. Chem. Inf. Model.*, 2015, **55**, 2324–2337.
- 39 D. Rogers and M. Hahn, *J. Chem. Inf. Model.*, 2010, **50**, 742–754.
- 40 C. A. Lipinski, F. Lombardo, B. W. Dominy and P. J. Feeney, *Adv. Drug Delivery Rev.*, 1997, **23**, 3–25.
- 41 D. F. Veber, S. R. Johnson, H.-Y. Cheng, B. R. Smith, K. W. Ward and K. D. Kopple, *J. Med. Chem.*, 2002, **45**, 2615–2623.
- 42 T. J. Ritchie and S. J. F. Macdonald, *Drug Discov. Today*, 2009, **14**, 1011–1020.
- 43 G. R. Bickerton, G. V. Paolini, J. Besnard, S. Muresan and A. L. Hopkins, *Nat. Chem.*, 2012, **4**, 90–98.
- 44 P. Ertl and A. Schuffenhauer, *J. Cheminform.*
- 45 S. A. Wildman and G. M. Crippen, *J. Chem. Inf. Comput. Sci.*, 1999, **39**, 868–873.
- 46 P. Ertl, B. Rohde and P. Selzer, *J. Med. Chem.*, 2000, **43**, 3714–3717.
- 47 S. Kirkpatrick, C. D. Jr. Gelatt and M. P. Vecchi, *Science*, 1983, **210**, 671–680.
- 48 *RDKit documentation*, 2024, <https://www.rdkit.org/docs/index.html>.
- 49 M. Rosenblatt, *Ann. Math. Statist.*, 1956, **27**, 832–837.
- 50 E. Parzen, *Ann. Math. Statist.*, 1962, **33**, 1065–1076.
- 51 G. W. Bemis and M. A. Murcko, *J. Med. Chem.*, 1996, **39**, 2887–2893.
- 52 M. Krenn, F. Häse, A. Nigam, P. Friederich and A. Aspuru-Guzik, *Mach. Learn.: Sci. Technol.*, 2020, **1**, 4.
- 53 L. McInnes, J. Healy and J. Melville, *J. Open Source Softw.*, 2018, **3**, 861.
- 54 F. Trozzi, X. Wang and P. Tao, *J. Phys. Chem. B*, 2021, **125**, 5022–5034.
- 55 N. G. Dickson, M. W. Johnson, M. H. Amin, R. Harris, F. Altomare, A. J. Berkley, P. Bunyk, J. Cai, E. M. Chapple, P. Chavez, F. Cioata, T. Cirip, P. deBuen, M. Drew-Brook, C. Enderud, S. Gildert, F. Hamze, J. P. Hilton, E. Hoskinson, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Lanting, T. Mahon, R. Neufeld, T. Oh, I. Perminov, C. Petroff, A. Przybysz, C. Rich, P. Spear, A. Tcaciuc, M. C. Thom, E. Tolkacheva, S. Uchaikin, J. Wang, A. B. Wilson, Z. Merali and G. Rose, *Nat. Commun.*, 2013, **4**, 1903.
- 56 N. D. Austin, N. V. Sahinidis and D. W. Trahan, *Chem. Eng. Res. Des.*, 2016, **116**, 2–26.
- 57 Y. Han, S. Roy and K. Chakraborty, in *2011 12th international symposium on quality electronic design*, 2011, pp. 1–7.
- 58 T. Okuyama, T. Sonobe, K. Kawarabayashi and M. Yamaoka, *Phys. Rev. E*, 2019, **100**, 012111.
- 59 F. Kretschmer, J. Seipp, M. Ludwig, G. W. Klau and S. Böcker, *Nat. Commun.*, 2025, **16**, 554.
- 60 P. Willett, *Trends Biotechnol.*, 1995, **13**, 516–521.
- 61 R. Tom, S. Gao, Y. Yang, K. Zhao, I. Bier, E. A. Buchanan, A. Zaykov, Z. Havlas, J. Michl and N. Marom, *Chem. Mater.*, 2023, **35**, 1373–1386.
- 62 R. A. Vargas-Hernández, K. Jorner, R. Pollice and A. Aspuru-Guzik, *J. Chem. Phys.*, 2023, **158**, 104801.
- 63 W. Lenz, *Eur. Phys. J. A*, 1920, **21**, 613–615.



- 64 M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson and G. Rose, *Nature*, 2011, **473**, 194–198.
- 65 Inc. D-Wave Systems, 2024.
- 66 V. Choi, *Quantum Inf. Process.*, 2008, **7**, 193–209.
- 67 M. Booth, S. P. Reinhardt and A. Roy, D-Wave Systems, Inc., Technical Report, 2017.
- 68 F. Glover, *Comput. Oper. Res.*, 1986, **13**, 533–549.
- 69 F. Glover, *ORSA J. Comput.*, 1989, **1**, 190–206.
- 70 F. Glover, *ORSA J. Comput.*, 1990, **2**, 4–32.
- 71 A. Lucas, *Front. Phys.*, 2014, **2**, 5.
- 72 F. Glover, G. Kochenberger, R. Hennig and Y. Du, *Ann. Oper. Res.*, 2022, **314**, 141–183.

