Thinking like a structural biologist: A pocket-based 3D molecule generative model fueled by electron density

*Lvwei Wang [†1], Rong Bai[†1], Xiaoxuan Shi[†1], Wei Zhang[2], Yinuo Cui[1], Xiaoman Wang[1], Cheng Wang[1], Haoyu Chang[1], Yingsheng Zhang[1], Jielong Zhou[1], Wenbiao Zhou[*1], and Bo Huang[*1]*

Affiliation:

1. Beijing StoneWise Technology Co Ltd., Haidian Street #15, Haidian District, Beijing 100080, China

2. Innovation Center For Pathogen Research, Guangzhou Laboratory, Guangzhou 510320, China

[†]Equal contributors

*Corresponding author

Abstract

We report for the first time the use of experimental electron density (ED) as training data for the generation of drug-like three-dimensional molecules based on the structure of a target protein pocket. Similar to a structural biologist building molecules based on their ED, our model functions with two main components: a generative adversarial network (GAN) to generate the ligand ED in the input pocket and an ED interpretation module for molecule generation. The model was tested on three targets including kinase, protease, and nuclear receptor, and evaluated with a reference dataset composed of over 8,000 compounds that have their activities reported in the literature. The

evaluation examined the chemical validity, chemical space distribution-based diversity, and similarity with reference active compounds concerning molecular structure and pocket-binding mode. Our model can reproduce classical active compounds and can also generate novel molecules with similar binding modes as active compounds, making it a promising tool to generate a library for high-throughput virtual screening. Our model is available as an online service to academic users via https://edmg.stonewise.cn/#/create.

Key words: experimental electron density, artificial intelligence, molecular generative models, de novo drug design

Introduction

Molecular generative models using the three-dimensional (3D) information of target pockets have garnered increasing attention in the field of de novo drug design[1,2]. The process of drug design is generally perceived as an inherently multi-constrained optimization process. The major constraints include complementarity between the ligand and protein, regarding multiple aspects, such as shape and non-covalent interactions (NCI), and requirements of the ligand itself, including synthesizability and low strain energy binding conformation. Therefore, efficiently maintaining these constraints during training of the molecular generative model is subject to intensive discussions. Some representative attempts include the following: 1) using autoregressive algorithms or introducing conditional tokens in the training process for models generating molecules in a sequential manner[3]; 2) leveraging generative adversarial networks (GAN) and reinforcement learning to reflect the desired bias in the output distribution[4,5]; 3) employing Bayesian optimization for the search of appropriate regions in the latent space from trained models

such as variational autoencoder (VAE)[6]. These approaches all possess delicately designed architecture and have accomplished great achievements by working in an end-to-end manner in the fields of machine translation, games, and image processing. However, when applied to molecule generation for drug design, additional challenges arise, including an increasing number of constraints and the lack of data. As a result, the application of end-to-end design that intends to satisfy all the constraints at one stroke by using massive amounts of training data is not as effective as in other fields. Therefore, researchers pursue solutions in two directions: 1) reducing the constraints by generating non-3D molecules, in which the molecules are represented as strings (i.e. SMILES) or graphs[7-9]; 2) expanding the dataset, such as by generating ligand-protein complex data using docking approaches[10]. Although these attempts are inspiring, a groundbreaking strategy that can pour additional experimental data in the models is still absent. The current data expansion approaches are based on calculations heavily relying on computer-aided drug design (CADD) theory. This promotes artificial intelligence (AI) models learning from well-established rules instead of "real world" information.

To find a solution to the challenges posed by complicated constraints and insufficient data, an appropriate molecule representation, which can simplify the introduction of massive amounts of experiment data to the training process and more comprehensively reflect the nature of constraints than the traditionally used molecular representations, is crucial. There are primarily two traditional 3D molecule representations used for the AI model: the node-edge based representation[11,12] and dot-cloud-based representation such as Gaussian filtering[13] and van der Waals radius[14]. Compared to the node-edge based representation in which one atom is located at one voxel, the dot clouds are continuously smoothing and more compatible with convolutional neural networks (CNN). However, regarding reflection of the physical and chemical properties of a molecule, such as

intermolecular noncovalent interactions, bond order, and electrostatics, none of these traditional representations are ideal. Therefore, multiple channels are required to complementarily describe bond orders, elements, and NCI-related properties, such as hydrogen bond donor or hydrogen bond acceptor. Consequently, using these molecular representations does not alleviate the lack of data but rather aggravates the problem by creating a data sparsity problem for some channels (e.g., Br, I for elements, and sigma hole for NCI). In contrast, electron density (ED) is ideal for representing both the physical and chemical properties of a molecule: the ED intensity isosurface describes the shape of a molecule; electron localization function (ELF)[15] describes the bond properties; and ED topology, such as saddle points, indicates the NCI between molecules[16]. As a result, there is no need to involve multiple channels, thereby avoiding the data sparsity problem. Moreover, ED is also the form in which massive amounts of experimental data are stored: there are over 120,000 high quality ED maps accumulated globally in the Protein Data Bank (PDB) over the past 60 years. Only part of the information in these experimental EDs is used for the determination of atom coordinates now widely applied by AI models, whereas other information reflecting NCI and solvent distribution remains untapped[16].

Based on the forementioned reasons, we propose the use of ED as the molecule representation in this work. In so doing, we are able to train our model with experimental ED. In addition, the design of our model is inspired by the procedures performed by structural biologists. This includes building molecules in the pocket, where the shape and key atoms are first defined based on their EDs, and stereochemistry constraints are then applied to improve conformation. We trained our model to learn constraints and generate molecules in a step-by-step manner. Initially, a GAN was used to take the ED of a pocket as input to learn pocket-ligand complementarity and generate the ligand ED. Subsequently, an ED interpretation module powered by vector quantized variational

autoencoder (VQ-VAE2) [17], PixelCNN[17,18], and V-Net[19] was used to learn constraints on ligand validity and interpret the generated ED into molecules. Experimental EDs were used to train the GAN, and quantum mechanics (QM) based computational ED as well as force field based molecular conformations were used to train the ED interpretation module. A rule-based fragment assembler was also designed as part of such a module to ensure the validity of the generated molecules. We evaluated our model using classical CADD indicators, such as quantitative estimate of drug-likeness (QED)[20] and synthetic accessibility score (SAS)[21], to test the molecular validity. Additionally, we employed a set of more intuitive indicators to test whether our model could generate novel active compounds: 1) the reproduction of classical active compounds; and 2) the generation of novel molecules sharing a similar binding mode with the classical active compounds. Regarding diversity, chemical space distribution was used for measurement.

Results

Model Design

The model architecture (Fig. 1) was designed under the principle of learning constraints step-by-step: less abundant experimental ED data were used to learn the complementarity between ligands and pockets, and relatively abundant computational data were used to learn the constraints on molecule validity. Specifically, ED was used as a molecular representation to convert the molecules from discrete symbols to a continuous numerical representation. Subsequently, a GAN (Fig. 1a) was designed for the generation of ligand ED (filler ED) and trained with 27,006 ligand-protein complex data to learn how the ligand ED could best complement the pocket ED. However, such a generated filler ED may not exist as a valid molecule, necessitating the VQVAE2-PixelCNN-VNet component for interpretation of the filler ED to valid and diversified molecules.

By reconstructing 40 million semiempirical quantum mechanical (QM) level EDs (details in Methods section), the VQ-VAE2[17] (Fig. 1b) was trained to learn two discrete latent spaces (top-level and bottom-level). Next, an autoregressive sampling process powered by PixelCNN was conducted in the latent spaces using generated ligand ED as a condition to retain the complementarity learned by the GAN. After conditional sampling, the filler ED was converted into a group of reconstructed EDs. Then, the V-Net[19] component takes the reconstructed ED as input to determine how to position the atoms to best maintain the shape of such ED and the type of atom required to best reflect the strength of the ED. The output of the V-Net component is called the map skeleton. A fragment library and a rule-based fragment substitution method were applied on map skeletons to create the final molecules.

Benefitting from the fact that the ED intrinsically contains the information, such as shape, element, bond properties, and interactions, it is unnecessary to include extra channels in our model. To assure the efficiency of retaining constraints during training of the whole model, we increased the weight of the region where the ligand and pocket show non-covalent interactions (NCI) during the GAN training.

Figure 1. Model Architecture. a) The GAN for generating filler ED based on pocket ED. b) The training and inference process of VQVAE2-PixelCNN-VNet for generating molecules using filler

ED as input. c) Illustration of evaluation framework. The blue part indicates the requirements for the validity test. The first level of the orange part indicates that the model can generate molecules that are previously reported active compounds; the second level indicates that the generated molecules should be more similar to high-activity compounds than low-activity compounds (i.e. a tendency to reproduce molecules with high activity). The green part represents the design meant to test whether there are active molecules with novel structures. All molecules with a low similarity (less than 0.5 Tanimoto similarity by using ECFP4) to the literature-reported active compounds were selected and first tested using the docking score and then tested using the binding mode by analyzing the NCI fingerprint. ED, electron density; GAN, generative adversarial network; NCI, non-covalent interactions; QED, quantitative estimate of drug-likeness; SAS, synthetic accessibility score; VQ-VAE, vector quantized variational autoencoder.

To demonstrate the functionality of our model, we applied the model to a kinase target, HPK1. The complex structure of HPK1 in DFG-in conformation binding with a reference ligand (PDB:7KAC; Fig. 2a) was used as a starting point. The pocket was defined as all the residues within 5 Å from the reference ligand. Using experimental ED (2Fo-Fc map contoured at 0.8 σ) at a resolution of 2.5 Å for the pocket as input (Fig. 2b), the filler ED was generated as shown in Fig. 2c. Interestingly, the generated filler ED replaced the region originally occupied by water molecules and covered the unoccupied cavities (indicated with red arrows in Fig. 2c). Then, the generated filler ED was converted into several reconstructed EDs (Fig. 2d). For each reconstructed ED, a number of map skeletons were generated and then interpreted into fragments. It is not a surprise to observe that the classical hinge binding groups could be well fitted into the reconstructed EDs for a kinase target (Fig. 2e) Then, the fragments were selectively connected to create different molecules (Fig. 2f and 2g).

Figure 2. ED based 3D molecule generation for HPK1. a) Binding pocket and reference ligand (PDB code 7KAC). Experimental ED (2Fo-Fc map at 1.2 s contour level) for the pocket and ligand are shown as a blue mesh. b) Pocket ED (2Fo-Fc map at 0.8 s contour level) with ligand removed. c) Generated filler ED. For the rainbow color scheme, red indicates a strong ED intensity, and blue indicates a weak ED intensity. The extension of generated ED to the region occupied originally by water molecules and cavities originally unoccupied are indicated by red arrows. d) Reconstructed ED generated based on filler ED. e) Map skeleton and hinge binding fragments interpreted from the map skeleton. Map skeletons are shown as white lines. f) Examples of generated molecules and their map skeleton aligned with reconstructed ED. G) List of examples of generated molecules. ED, electron density.

Model evaluation

As shown in Fig. 1c, our molecule generative model was evaluated from three perspectives: 1) the ability to generate valid molecules in terms of QED and SAS while maintaining a reasonable diversity; 2) the ability to reproduce classical active compounds, defined as generation of molecules with over 0.5 Tanimoto similarity[22] with the reference active compounds; 3) the ability to generate novel active compounds, defined as generation of molecules with less than 0.5 Tanimoto similarity with the reference active compounds but possessing a similar binding mode.

We selected three targets including HPK1 (PDB: 7KAC), Covid19-3CL (PDB:7VU6), and VDR (PDB: 1S19), to test the performance of our model. They were selected as representatives of kinase, protease, and nuclear receptors. Regarding references (Table S1), 6334, 1101, and 757 compounds with reported activity were used for HPK1, 3CL, and VDR, respectively (supplementary material

1). In addition, a 3D molecule generative model[23] reported in NIPS2021 was used as a comparison benchmark.

Table S1. Reference Compounds

Target	# of reference compounds			
	Active	Medium	Not Active	Total
HPK1	3847	2319	168	6334
Covid19-3CL	222	248	631	1101
VDR	329	67	361	757

Note: Active=less than 100 nM potency (Ki, Kd, EC50, or IC50); Medium=100 nM to 1,000 nM potency; Not active=over 1,000 nM potency

Validity Test

The most fundamental test for molecules generated by AI models is the validity test evaluating the QED and SAS. We generated 10,000 molecules for each of the four targets using our model and the benchmark model and compared the QED and SAS of the generated molecules. To make the comparison results easy to understand, we also calculated the QED and SAS of the reference compounds and used them as positive controls. As shown in Table 1, although both our model and the benchmark model performed well for QED, our model outperformed the benchmark model on SAS.

Table 1. Evaluation of molecular validity

Metric for validity test		Target								
		HPK1			Covid19-3CL			VDR		
		Ours	BM[a]	Ref.[b]	Ours	BM[a]	Ref.[b]	Ours	BM[a]	Ref.[b]
QED	Avg.	0.54	0.55	0.47	0.40	0.62	0.32	0.55	0.48	0.49
	Med.	0.53	0.56	0.46	0.39	0.64	0.29	0.53	0.47	0.48
SAS	Avg.	3.0	5.5	3.6	3.2	4.5	3.2	2.9	5.3	3.0

	Med.	2.9	5.5	3.6	3.1	4.5	2.8	2.9	5.3	2.8

Notes: 10,000 molecules are respectively generated by our model and benchmark model; [a]BM refers to a model[23] reported in NIPS 2021 and used as benchmark here; [b]Reference compounds (Table S1); About the color scheme: all the values of reference compounds have their cells colored yellow. For generated molecules, if their values are greater than or equal to that of the reference, then their cells are colored green; otherwise, the cells are colored red. Specifically, for QED, the higher is better; for SAS, the lower is better

Table S2. Diversity of molecules generated by our model

Diversity		Target								
		HPK1			3CL			VDR		
		Ours	BM[a]	Ref.[b]	Ours	BM[a]	Ref.[b]	Ours	BM[a]	Ref.[b]
Diversity[a]	Avg.	0.86	0.89	0.82	0.86	0.90	0.68	0.86	0.86	0.64
	Med.	0.87	0.90	0.85	0.86	0.91	0.84	0.87	0.86	0.54

Note: [a]Diversity=1- (pairwise Tanimoto similarities over ECFP4 fingerprints among the target set of molecules). The number of reference molecules subjected to the analysis is the same as that of Table 1. For each target, 10,000 molecules are respectively generated and subjected to the analysis.

Chemical Space Distribution

To understand the diversity of the molecules generated by our model, we calculated the Tanimoto-similarity-based diversity. As shown in Table S2, the molecules generated by the benchmark model exhibit higher diversity than those generated by our model. However, such general diversity is not suitable for the evaluation of molecules generated with constraints, because well-functioning constraints may reduce the general diversity. To perform a comprehensive evaluation, we referred to the small-molecule-universe (SMU) chemical space[24], a self-organizing-

map (SOM)-supported chemical space used to describe the distribution of molecules with molecular weights under 500. To simplify the comparison results, samples from PubChem were used to represent molecules lacking target-specific constraints, and active reference compounds are used to represent molecules with tight constraints. As shown in Fig. 3, we observed that our model achieved a better balance between the diversity and constraints than the benchmark model (i.e. the molecule generated by our model is concentrated towards the region where the active compounds are positioned, whereas the molecules of the benchmark model are distracted to some other regions and thus appear more diverse).

Molecular Distribution in SMU Chemical Space

(a)

PubChem
(N=96 million)

PubChem
(Randomly selected 1 million)

PubChem
(Randomly selected 10,000)

(b)

Active HPK1 inhibitors
(N=3,847)

Generated Molecules for HPK1
using our model
(N=27,000)

Generated Molecules for HPK1
using benchmark model
(N=27,000)

(c)

Active 3CL inhibitors
(N=222)

Generated Molecules for 3CL
using our model
(N=12,700)

Generated Molecules for 3CL
using benchmark model
(N=12,700)

(d) Active VDR agonists
(N=329)

Generated Molecules for VDR
using our model
(N=11,000)

Generated Molecules for VDR
using benchmark model
(N=11,000)

Figure 3. Chemical space distribution of generated molecules and references. A 120 × 120 SOM was created using SMU[24] compounds. The color indicates the number of molecules on a logarithmic scale. a) The distribution of PubChem with different numbers of randomly selected samples. b) The distribution of HPK1 reference active compounds and molecules generated using different models. c) The distribution of Covid19-3CL reference active compounds and molecules generated using different models. d) The distribution of VDR reference active compounds and molecules generated using different models. The number of molecules generated by our model is adjusted to match the number that can be generated by the benchmark model within a reasonable time frame. The SMILES of the molecules and their positions in the chemical space are provided in Supplementary material 3. SOM, self-organizing-map; SMU, small-molecule-universe.

Reproduction of classical active compounds Besides testing the validity of the generated molecules, we also needed to test whether the molecular generative model can reproduce already known active compounds. To this end, 1 million molecules were generated for each of the three targets and compared to reference compounds by measuring the Tanimoto similarity of the ECFP4 fingerprint. In addition, generating a molecule with over 0.5 Tanimoto similarity against a reference compound was considered a successful reproduction of this reference compound. Our model successfully reproduced reference compounds for all three targets (Supplementary material 2). Taking HPK1 as an example, some reference compounds and their generated counterparts are listed in Fig. 4. Furthermore, when we sorted reference compounds by their activity, we observed that our model tended to generate molecules with higher similarity to active and medium-active compounds, compared with inactive ones, for all three targets (Table 2). Apart from the similarity trend, our model also reproduced more active and medium-active compounds than inactive ones

for HPK1 and Covid19-3CL. Regarding the comparison with the benchmark model which failed to generate million-level molecules within a reasonable time frame, our molecules that were subjected to the comparison were randomly selected form the 1 million previously generated to match the capacity of the benchmark model. As shown in Table S3, although neither of the two models reproduced active reference compounds under the condition that only tens of thousands of molecules were generated, our model still provided molecules with better Tanimoto similarity to the references than the benchmark model did.

Figure 4. Examples of generated molecules that are similar to reference compounds for HPK1.

Table 2. Test of reproducing classical active compounds for HPK1

Metric for classical active compounds reproduction test		HPK1			Covid19-3CL			VDR		
		Active	Medium	Not active	Active	Medium	Not active	Active	Medium	Not active
# of reference compounds		3847	2319	168	222	248	631	329	67	361
Our Model	# of molecules generated	1 million			1 million			1 million		
	Max. of Tanimoto similarity to ref. cpd.	0.76	0.72	0.51	0.58	0.73	0.61	0.57	0.44	0.55
	# of reference molecule reproduced[a]	58	53	1	3	15	11	1	0	7
	# of generated molecules with over 0.5 Tanimoto similarity to ref. cpd.	30	20	2	14	158	42	1	0	10

Note: [a]*reproduced: for a reference compound, if a generated molecule has over 0.5 Tanimoto similarity with it for ECFP4, then this reference compound is considered to have been reproduced. The SMILES of the generated molecules and their similar reference compounds are provided in Supplementary material 2.*

Table S3. Comparison with Benchmark Model in Reproducing Reference Compounds

Metric for classical active compounds reproduction test		HPK1			Covid19-3CL			VDR		
		Active	Medium	Not active	Active	Medium	Not active	Active	Medium	Not active
# of reference compounds		3,847	2,319	168	222	248	631	329	67	361
Our Model	# of molecules generated	27,000			12,700			11,000		
	Max. of Tanimoto similarity to ref. cpd.	0.49	0.47	0.44	0.38	0.50	0.43	0.42	0.42	0.50
Benchmark Model	# of molecules generated	27,000			12,700			11,000		
	Max. of Tanimoto similarity to ref. cpd.	0.36	0.37	0.35	0.34	0.47	0.52	0.31	0.33	0.4

Generation of novel active compounds If a molecule generative model can only generate classical active compounds; it is not attractive to researchers searching for new drugs. To test our model's ability to generate active compounds with novel structures, we searched the generated library for the molecules exhibiting a <0.5 Tanimoto similarity with reference compounds while sharing a similar binding mode.

Taking HPK1 as an example, 2,769 HPK1 reference compounds with reported Ki or Kd were selected and docked into the 7KAC pocket by using Glide SP[25,26] to provide a background for binding mode analysis. This is because the binding mode is more theoretically related to Ki and Kd than to EC50 or IC50. However, as shown in Fig. 5a, the glide score could not efficiently distinguish the reference compounds with different activities, which implies the need for a more powerful descriptor of binding mode. Thus, we referred to NCI fingerprints depicted using the independent gradient model (IGM) method[27]. The IGM method calculates NCIs by analyzing the topological properties of EDs, and it can provide a full spectrum of NCIs, whereas the traditional rule-based method only provides a short list of classical NCIs. Compound clustering based on the NCI fingerprint was conducted using T-distributed stochastic neighbor embedding (t-SNE)[28]. As shown in Fig. 5b, the clustering results aligned quite well with the activity: compounds sorted into one cluster tended to possess similar activities despite the low discriminative capability for samples with very high activity (separating compounds below 1 nM from those below 10 nM).

To analyze the binding mode of molecules generated for HPK1, we randomly selected 10,000 generated molecules with <0.5 Tanimoto similarity against the reference compounds and then scored them using Glide by only employing the minimization and scoring functions. This operation was conducted for both our model and the benchmark model. Similarly to the reference compounds,

the Glide score was unable to efficiently distinguish between the performances of our model and the benchmark model (Fig. 5a). Therefore, we randomly selected 200 molecules from each of the two models and conducted NCI fingerprint clustering. As shown in Fig. 5b, the molecules generated by the benchmark model tended to be more concentrated and farther from the cluster of active compounds than those generated by our model. Molecules at different positions are listed in Fig. 5c and 5d, and their NCI fingerprints are provided in Fig. S1. The hinge region-related NCIs, which are considered crucial for active compounds, were exhibited for molecules generated close to the cluster of active compounds. These NCIs were weak for molecules far from the clusters, such as #2-BM and #3-BM. Although there were cases where the molecule from our model was close to the benchmark model molecule in the NCI fingerprint map, and both of them were close to the cluster of active compounds, our molecules are still easily distinguished for their superior synthetic accessibility.

Figure 5. Binding mode analysis of the generated molecules for HPK1. a) Glide Score distribution

for active reference compounds and generated molecules. b) Results of t-SNE clustering using

IGM calculated NCIs as features. c) Binding mode of selected molecules generated by our model. d) Binding mode of selected molecules generated by benchmark model. For panel c and d, NCI regions are indicated with dots colored using the rainbow scheme, in which blue indicates weak interactions and red indicates strong interactions. Details of docking score, clustering analysis, and NCI fingerprints of selected molecules are provided in supplementary material 4, 5, and 6, respectively. IGM, independent gradient model; NCI, non-covalent interaction; t-SNE, t-distributed stochastic neighbor embedding.

Figure S1. NCI fingerprint of selected molecules. Hinge region related NCIs (i.e. E92 or C94 involved) are indicated with red box. More details are provided in supplementary material 6.

Discussion

To solve the fundamental problem of target compound identification for structure-based drug design, we used experimental ED to train an AI 3D molecule-generative model for the first time. The core of our approach can be understood from three perspectives: data, molecule representation, and model architecture. Since experimental EDs contain experimentally observed time-averaged conformational change and solvent information that are critical for ligand binding, the introduction of these ED data to AI molecular generation provides a promising solution to the challenge of insufficient data. In addition, we highlighted the value of representing molecules as ED; such representations are not only compatible with the usage of experimental ED data during model training but also avoid the data sparsity problem caused by using multiple channels. Furthermore, we proposed a procedure allowing for step-by-step constraint learning, just as structural biologists build small molecules in the pocket. In summary, we presented a novel solution for AI-based 3D molecule generation.

One possible limitation to our approach is associated with a scenario in which the model is used without experimental ED data as input. Such a scenario could occur when the generative model is used for a pocket provided by a molecular dynamic approach. For our model, although it is trained using experimental ED data as input, it also works when using a calculated Fourier-synthesis based ED map as input. Specifically, one just needs to add reasonable B factors (e.g., 20 with a Gaussian perturbation) for each atom and then calculate the ED map using the Fourier-synthesis.

To further improve the performance of our model, we are considering the use of multi resolution ED. For the training of our current model, all pockets were represented as ED at a resolution of 2.5 Å, and therefore, some of the PDB entries with a resolution higher than 2.5 Å had their data at high resolution shells unused. One possible method for utilizing these valuable data is to create

multiple channels for different resolutions. However, one must strike a balance between creating multiple channels and avoiding the data sparsity problem for high resolution channels receiving PDB entries with low resolution data.

Methods

1. Molecule Training

We designed a two-step generation scheme using ED to represent molecules. For the first step, we used 27,006 complex data to focus on the learning of the complementarity between protein pockets and ligand molecules. For the second step, we trained the model to generate valid molecules under constraints contained in the generated ED in the first step. Decoupling the two steps makes it possible to use millions of QM and force-field-generated computational data for model training in the second step. ED was featured using 3D CNN, and data augmentation was used to achieve rotational invariance.

1.1 Ligand ED Generative Model

Similar to the image generative model (such as pix2pix[29]) used for image inpainting, our model was designed to generate ligand ED (also referred as filler ED) for pocket ED while maintaining complementarity in terms of shape and NCI. The model was built on the GAN framework with the architecture shown in Fig. 1a.

To prepare data for model training, the experimental EDs for pockets were generated using Phenix[30] and experimental ED coefficients downloaded from PDB. These pocket EDs were used as input features for the AI model. The computational ED for ligands were generated based on their coordinates using xtb[31] at the GNF2-xTB level. Since the ED intensity within the nuclear region is much stronger than that within the region where the NCIs take place and

where the shape of the molecule is defined (i.e., the isosurface at 0.03 e/Å³), the range of values were compressed using logarithms and then used as labels for the AI model.

The pocket EDs were input into a V-Net 3D image generation network, which acts as a generator. In addition to the generator, two discriminators, called complex and ligand discriminators, were designed. The complex discriminator was responsible for examining the level of complementarity between the ligand and the pocket, and the ligand discriminator was responsible for verifying ED validity from the perspectives of gradient and connectivity.

The protein-ligand pairs used to train the filler ED generative model were from the PDBbind database[32]. In total, 27,006 ligand-protein pairs were extracted from 12,905 complexes. The number of the ligand-protein pairs was larger than that of the complexes because we included drug-like molecules as well as other binders with molecular weight less than 600, such as ATP, ADP, and sugars. To avoid the appearance of similar pockets in both the training and the testing sets, training-testing split was done by referring to the pocket classification from previous studies focusing on the 1D and 3D pocket similarity[10,33]. The key NCIs identified using a previously reported method[16] were also included in the model during training with the purpose of making the network more attentive on the NCI related regions. Smooth L1 loss was performed separately to strengthen the loss of the NCI related regions.

The GAN and its loss were expressed using the following equations:

$$G^* = argmin_G max_D \left(L_{cGAN}(G, D_{complex}) + \alpha L_{GAN}(G, D_{ligand}) \right) + \lambda L(G) \qquad (1)$$

$$L_{cGAN}(G, D_{complex}) = \mathbb{E}_{x,y}[log\ D_{complex}(x,y)] + \mathbb{E}_{x,y}\left[log\left(1 - D_{complex}(x, G(x))\right)\right] \quad (2)$$

$$L_{GAN}(G, D_{ligand}) = \mathbb{E}_{x,y}[log\ D_{ligand}(x,y)] + \mathbb{E}_{x,y}\left[log\left(1 - D_{ligand}(G(x))\right)\right] \qquad (3)$$

$$L(G) = \quad \beta L_{gen} + \gamma L_{ligand} + \varepsilon L_{NCI} + \delta(L_{tv1} + L_{tv2}) \qquad (4)$$

where $D_{complex}$ represents the complex discriminator, and D_{ligand} represents the ligand discriminator. $L(G)$ indicates the regression loss of GAN, x represents the input pocket, y represents the ground truth ligand, and $\alpha, \lambda, \beta, \gamma, \varepsilon, \delta$ are parameters to be learned. L_{tv1} and L_{tv2} are used to measure the similarity of two EDs from the perspectives of the first and second orders derivative of intensity, respectively.

The components of $L(G)$ are defined as follows:

$$\text{SmoothL1} = \begin{cases} 0.5x^2, \text{if } |x| < 1 \\ |x| - 0.5, \quad \text{otherwise} \end{cases} \tag{5}$$

$$L_{gen} = \frac{1}{N_I} \Sigma_{i\in I} \text{SmoothL1}(G(x)_i - y_i) \tag{6}$$

where I indicates the entire region, and N_I indicates the number of elements in I.

$$L_{ligand} = \frac{1}{N_{Gt}} \Sigma_{i\in Gt} \text{SmoothL1}(G(x)_i - y_i) \tag{7}$$

where Gt indicates the region covered by the ground truth, and N_{Gt} indicates the number of elements in Gt

$$L_{NCI} = \frac{1}{N_{NCI}} \Sigma_{i\in NCI} \text{SmoothL1}(G(x)_i - y_i) \tag{8a}$$

where NCI indicates the region covered by NCI, and N_{NCI} indicates number of elements in NCI

$$L_{tv1} = \frac{1}{N_I} \Sigma_{i\in I} \text{SmoothL1}(\widehat{tv1}_i - tv1_i) \tag{8b}$$

$$L_{tv2} = \frac{1}{N_I} \Sigma_{i\in I} \text{SmoothL1}(\widehat{tv2}_i - tv2_i) \tag{8c}$$

$$tv1 = \left(I_{(i+1,j,k)\in I} - I_{(i,j,k)\in I}\right) \oplus \left(I_{(i,j+1,k)\in I} - I_{(i,j,k)\in I}\right) \oplus \left(I_{(i,j,k+1)\in I} - I_{(i,j,k)\in I}\right) \tag{8d}$$

$$tv2 = \left(tv1_{(i+1,j,k)\in I} - tv1_{(i,j,k)\in I}\right) \oplus \left(tv1_{(i,j+1,k)\in I} - tv1_{(i,j,k)\in I}\right) \oplus \left(tv1_{(i,j,k+1)\in I} - tv1_{(i,j,k)\in I}\right) \tag{8e}$$

where $\oplus$ indicates the operator of concatenate.

1.2 Ligand Structure Generative Model

The model was designed to infer the atom type and coordinates of small molecules using generated filler ED as input. A PixelCNN[17,18] was trained to do the sampling with the generated

filler ED as a condition in the latent space constructed using a VQ-VAE2 framework. A group of reconstructed EDs that are more similar to a molecule were generated in the sampling process. Next, the reconstructed ED was submitted to a V-Net framework for the detection of key points. This step is similar to that of key point detection in human skeletons in the field of image processing but differs in that it faces severe complications associated with the flexible number of points, diverse types of points, and complicated topology.

VQ-VAE2 was used to compress the ED into a discrete latent space and learn a codebook in which EDs are represented by a series of embeddings[17]. Each pixel in the latent space can be described by an integer in the codebook and the integer can be logically considered as a 3D molecular fragment. To retain the constraints in the input ED while ensuring that the output ED can match a valid molecule, a separate autoregressive prior (PixelCNN)[17,18] was taught to sample the latent space with input ED as a given condition. Two codebooks were learned during training: the top-level codebook focusing on the extraction of general profile information (such as shape) and the bottom-level codebook focusing on the detailed information (such as local conformations). Regarding sampling with the two codebooks, the input ED was used as a condition to sample the top-level codebook, and the top-level encoding was used as the condition to sample the bottom-level codebook. Thus, the relationship between the conditional ED (i.e., input ED) and the generated ED was decoupled.

The input ED value at each pixel was normalized to a range between -1 and 1 using tanh plus batch normalization in the encoder before it proceeded to the codebook for sampling. This improves the codebook encoding utilization, and thereby increases the amount of information represented through encoding.

The data used to train the VQVAE2-PixelCNN-VNet model for inferring atoms from generated ligand ED were obtained from PubChem. Two million molecules with a QED exceeding 0.3 were extracted. For each molecule, 20 conformers were generated using ConfGen[34]. EDs of the 40 million conformers were generated using xtb[31] and then used as training data for the VQ-VAE2. The V-Net training, which was used to infer the atom location, relied on the above 40 million EDs and the conformers used to generate these EDs, which were combined into conformer-ED pairs. The V-Net output was comprised of connected atoms and called a map skeleton.

EDs were voxelized into a discrete 0.5Å cubic grid with a side size of 24Å. There were represented as a tensor $p \in \mathcal{R}^{1 \times 48 \times 48 \times 48}$, where channel size was 1, expressing the intensity of ED.

The VQ-VAE2 loss is shown in Eq (9), where x is the training instance, D is the decoder of the VQ-VAE2 and e is the encoder. The reconstruction loss is shown in Eq (10). It combines the $L_{SoftDiceLoss}$ as well as $L_{SmoothL1}$. To generate better ED reconstructed shapes, we adopted $L_{SoftDiceLoss}$ as shown in Eq (11) for segmenting ED and the background where y is the target and $\hat{y}$ is the prediction. The target label of the pixel with the ED value was set to 1, and the rest of the background part was set to 0. Further, $L_{SmoothL1}$ was obtained using Eq (12); similarly, y is the target and $\hat{y}$ is the prediction.

$$L(x, D(e)) = L_{Reconstruction}(x, D(e)) + \|sg[E(x) - e]\|_2^2 + \beta\|sg[e] - E(x)\|_2^2 \qquad (9)$$

$$L_{Reconstruction}(x, D(x)) = L_{SoftDiceLoss}(x, D(x)) + \alpha L_{SmoothL1}(x, D(x)) \qquad (10)$$

$$L_{SoftDiceLoss}(\hat{y}, y) = 1 - \frac{2\sum_{i \in I-1} \hat{y}_i y_i}{\sum_{i \in I-1} \hat{y_i}^2 + \sum_{i \in I-1} y_i^2} \qquad (11)$$

$$L_{SmoothL1}(\hat{y}, y) = \frac{1}{N_I}\sum_{i \in I} SmoothL1(\hat{y}_i - y_i) \qquad (12)$$

PixelCNN optimized the negative log-likelihood of the training data *to maximize the probability* $p(x|h) = \prod_{i=1}^{HWD} p(x_i|x_1, \ldots, x_{i-1}, h)$, where H represents height, W represents width, and D represents depth of the 3D cube where h is the condition. Top-level and bottom-level prior networks were modeled with 6×6×6 and 12×12×12 latent variables, respectively. Additionally, the condition of the top-level PixelCNN was the previously generated filler ED. The bottom-level network was conditioned on the top-level prior.

Respecting the V-Net output, atoms from molecules were characterized as one-hot representations in each cell of the 3D grid. The ground truth of the atom types was represented as a tensor, where 10 channels were denoted as the PAD (no atoms), C, N, O, S, P, F, Cl, Br, and I, respectively, in this work. Then, V-Net utilized the cross-entropy loss to classify atom types in each voxel $x \in \mathcal{R}^3$.

1.3 Map skeleton substitution and fragment assembling

The V-Net generated map skeletons that were already positioned in the pocket. A map skeleton was then divided into two fragments, and each fragment was replaced by molecular fragments with similar shapes from the conformation library. Briefly, the library entries with a root mean square distance less than 1.5 Å from the fragment of the map skeleton were used to substitute that fragment.

The molecules from PubChem were matched with 35 reaction templates[35] according to the substructure superposition supported by RDKit[36], and then cut into two fragments at the matching site. If a molecule matched multiple reaction templates, it was still cut into two fragments each time, but the cutting site was different. In this way, over 270,000 fragments with labeled cutting

sites were obtained. Then OpenBabel[37] was used to generate a maximum of 20 conformations for each fragment, and it finally produced a library with 3.5 million 3D conformations.

The map skeleton was cut into several fragment pairs in the following manner: each time the whole molecule was cleaved at one acyclic single bond, two fragments were generated; next time, the whole molecule was cleaved at another acyclic single bond; the process was repeated until all acyclic single bonds had been cut. For every fragment pair, the above-established 3.5 million 3D conformation library was searched for entries that were similar to the fragments. Next, the two groups of selected entries were assembled together in an enumerative way to make a list of target molecules. Subsequently, several filters were applied to remove the unqualified molecules. These filters include the following:

1) Collides with pocket: the distance from the pocket heavy atom is less than 2.5 Å.

2) Collides between the fragments: the distance from the non-bonded heavy atom is less than 2.0 Å.

3) Stability and synthetic accessibility filters and drug-like filter reported by Virshup et al.[24]

2. Model Evaluation

The data we used for chemical space construction consisted of 8.8 million molecules from SMU[24].

The evaluation data were collected from 144 publications including research papers and patents.

3. Chemical space construction

Chemical space was constructed mainly based on the method described in a previous study[24].

Specifically, Moreau−Broto autocorrelation descriptors were used for molecule description.

These descriptors encode topological correlations between the atomic properties of a molecule into a fixed-length vector.

$$AC(d,p) = \sum_{i \leq j} p_i p_j \delta(d_{ij} - d) \qquad (13)$$

$$\delta(d_{ij} - d) = \begin{cases} 1, if\ d_{ij} = d \\ 0, otherwise \end{cases} \qquad (14)$$

where d_{ij} is the number of bonds separating atoms i and j, while p_i is the value of atomic property p on atom i. In our work, p includes four properties: carbon-scaled atomic mass, carbon-scaled atomic van der Waals volume, carbon-scaled atomic Sanderson electronegativity, and carbon-scaled atomic polarizability. The values of d ranged from 0 to 7. As a result, a molecule was finally described as a vector with 32 dimensions. The autocorrelation descriptors were calculated using PyBioMed[38].

Up to 8.8 million molecules from SMU[24] were converted into 32D vectors and used for the construction of the chemical space by using the SOM method[39]. The SOM size was 120 ×120. The SOM was implemented using MiniSom library[40].

4. NCI fingerprint analysis

Docking was implemented by using Glide[25,26]. IGM-based NCI analysis was implemented using Multiwfn[41]. The atom-pair-based NCI list output by Multwfn was further annotated with the Mol2/Sybyl atom types for both of the atoms using OpenBabel[42] and PyBel[37] packages, and then submitted to Scikit-learn[43] for t-SNE[28] clustering analysis.

5. Benchmark Model

The source code of benchmark model is downloaded from its official website: https://github.com/ luost26/3D-Generative-SBDD

6. Software for Figures and Tables

The structure and ED figures were made using Pymol.

AUTHOR INFORMATION

Author Contributions

Bo Huang and Wenbiao Zhou conceived the idea. Bo Huang provided instructions for experimental ED processing, molecular representations and model evaluation. Wenbiao Zhou provided instructions on AI models. Lvwei Wang developed the GAN module. Rong Bai developed the VQVAE2-PixelCNN-VNet module. Xiaoxuan Shi developed the fragment assembling code and conducted the evaluation. Wei Zhang and Yinuo Cui prepared evaluation data. Cheng Wang supported chemical space distribution analysis. Haoyu Chang supported GAN training. Xiaoman Wang supported website development. Yingsheng Zhang and Jielong Zhou provided instructions on general model architecture.

DATA AVAILABILITY

Reference compounds extracted from literatures, data used for chemical space distribution analysis and NCI fingerprint analysis, and results of reference compound reproduction are available as supplementary materials.

CODE AVAILABILITY

Partial codes are available upon request.

Our model is available as online service to academic users via https://edmg.stonewise.cn/#/create.

COMPETING INTERESTS

The authors declare no competing interests.

ACKNOWLEDGMENT

All the authors received funding from StoneWise Ltd.; B.H. and W.Z. also received funding for cloud computing from Beijing Municipal Science & Technology Commission project (No. Z211100003521001).

REFERENCES

Walters, W. P. & Barzilay, R. Applications of Deep Learning in Molecule Generation and Molecular Property Prediction. *Acc Chem Res* **54**, 263-270, doi:10.1021/acs.accounts.0c00699 (2021).
Yang, N. *et al.* Molecule Generation for Drug Design: a Graph Learning Perspective. arXiv:2202.09212 (2022). <https://ui.adsabs.harvard.edu/abs/2022arXiv220209212Y>.
Shi, C. *et al.* GraphAF: a Flow-based Autoregressive Model for Molecular Graph Generation. arXiv:2001.09382 (2020). <https://ui.adsabs.harvard.edu/abs/2020arXiv200109382S>.
Olivecrona, M., Blaschke, T., Engkvist, O. & Chen, H. Molecular de-novo design through deep reinforcement learning. *J Cheminform* **9**, 48, doi:10.1186/s13321-017-0235-x (2017).
De Cao, N. & Kipf, T. J. a. e.-p. MolGAN: An implicit generative model for small molecular graphs. arXiv:1805.11973 (2018). <https://ui.adsabs.harvard.edu/abs/2018arXiv180511973D>.
Lim, J., Ryu, S., Kim, J. W. & Kim, W. Y. Molecular generative model based on conditional variational autoencoder for de novo molecular design. *J Cheminform* **10**, 31, doi:10.1186/s13321-018-0286-7 (2018).
Xu, M., Ran, T. & Chen, H. De Novo Molecule Design Through the Molecular Generative Model Conditioned by 3D Information of Protein Binding Sites. *J Chem Inf Model* **61**, 3240-3254, doi:10.1021/acs.jcim.0c01494 (2021).
Li, Y. *et al.* DeepScaffold: A Comprehensive Tool for Scaffold-Based De Novo Drug Discovery Using Deep Learning. *J Chem Inf Model* **60**, 77-91, doi:10.1021/acs.jcim.9b00727 (2020).
Mahmood, O., Mansimov, E., Bonneau, R. & Cho, K. Masked graph modeling for molecule generation. *Nat Commun* **12**, 3156, doi:10.1038/s41467-021-23415-2 (2021).
Francoeur, P. G. *et al.* Three-Dimensional Convolutional Neural Networks and a Cross-Docked Data Set for Structure-Based Drug Design. *J Chem Inf Model* **60**, 4200-4215, doi:10.1021/acs.jcim.0c00411 (2020).
Deng, D. *et al.* XGraphBoost: Extracting Graph Neural Network-Based Features for a Better Prediction of Molecular Properties. *J Chem Inf Model* **61**, 2697-2705, doi:10.1021/acs.jcim.0c01489 (2021).
Gligorijevic, V. *et al.* Structure-based protein function prediction using graph convolutional networks. *Nat Commun* **12**, 3168, doi:10.1038/s41467-021-23303-9 (2021).
Kuzminykh, D. *et al.* 3D Molecular Representations Based on the Wave Transform for Convolutional Neural Networks. *Mol Pharm* **15**, 4378-4385, doi:10.1021/acs.molpharmaceut.7b01134 (2018).
Skalic, M., Jimenez, J., Sabbadin, D. & De Fabritiis, G. Shape-Based Generative Modeling for de Novo Drug Design. *J Chem Inf Model* **59**, 1205-1214, doi:10.1021/acs.jcim.8b00706 (2019).
Koumpouras, K. & Larsson, J. A. Distinguishing between chemical bonding and physical binding using electron localization function (ELF). *J Phys Condens Matter* **32**, 315502, doi:10.1088/1361-648X/ab7fd8 (2020).
Ding, K. *et al.* Observing Noncovalent Interactions in Experimental Electron Density for Macromolecular Systems: A Novel

Perspective for Protein-Ligand Interaction Research. *J Chem Inf Model*, doi:10.1021/acs.jcim.1c01406 (2022).

Razavi, A., van den Oord, A. & Vinyals, O. J. a. e.-p. Generating Diverse High-Fidelity Images with VQ-VAE-2. arXiv:1906.00446 (2019). <https://ui.adsabs.harvard.edu/abs/2019arXiv190600446R>.

Pombo, G., Gray, R., Varsavsky, T., Ashburner, J. & Nachev, P. J. a. e.-p. Bayesian Volumetric Autoregressive generative models for better semisupervised learning. arXiv:1907.11559 (2019). <https://ui.adsabs.harvard.edu/abs/2019arXiv190711559P>.

Milletari, F., Navab, N. & Ahmadi, S.-A. J. a. e.-p. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. arXiv:1606.04797 (2016). <https://ui.adsabs.harvard.edu/abs/2016arXiv160604797M>.

Bickerton, G. R., Paolini, G. V., Besnard, J., Muresan, S. & Hopkins, A. L. Quantifying the chemical beauty of drugs. *Nat Chem* **4**, 90-98, doi:10.1038/nchem.1243 (2012).

Ertl, P. & Schuffenhauer, A. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *J Cheminform* **1**, 8, doi:10.1186/1758-2946-1-8 (2009).

Bajusz, D., Racz, A. & Heberger, K. Why is Tanimoto index an appropriate choice for fingerprint-based similarity calculations? *J Cheminform* **7**, 20, doi:10.1186/s13321-015-0069-3 (2015).

Shitong Luo, J. G., Jianzhu Ma, Jian Peng. A 3D Generative Model for Structure-Based Drug Design. *NeurIPS 2021* (2021).

Virshup, A. M., Contreras-Garcia, J., Wipf, P., Yang, W. & Beratan, D. N. Stochastic voyages into uncharted chemical space produce a representative library of all possible drug-like compounds. *J Am Chem Soc* **135**, 7296-7303, doi:10.1021/ja401184g (2013).

Friesner, R. A. *et al.* Extra precision glide: docking and scoring incorporating a model of hydrophobic enclosure for protein-ligand complexes. *J Med Chem* **49**, 6177-6196, doi:10.1021/jm051256o (2006).

Schrödinger Release 2021-4. **Glide, Schrödinger, LLC, New York, NY, 2021.**

Contreras-Garcia, J. *et al.* NCIPLOT: a program for plotting non-covalent interaction regions. *J Chem Theory Comput* **7**, 625-632, doi:10.1021/ct100641a (2011).

Belkina, A. C. *et al.* Automated optimized parameters for T-distributed stochastic neighbor embedding improve visualization and analysis of large datasets. *Nat Commun* **10**, 5415, doi:10.1038/s41467-019-13055-y (2019).

Isola, P., Zhu, J.-Y., Zhou, T. & Efros, A. A. J. a. e.-p. Image-to-Image Translation with Conditional Adversarial Networks. arXiv:1611.07004 (2016). <https://ui.adsabs.harvard.edu/abs/2016arXiv161107004I>.

Adams, P. D. *et al.* Recent developments in the PHENIX software for automated crystallographic structure determination. *J Synchrotron Radiat* **11**, 53-55, doi:10.1107/s0909049503024130 (2004).

Bannwarth, C., Ehlert, S. & Grimme, S. GFN2-xTB-An Accurate and Broadly Parametrized Self-Consistent Tight-Binding Quantum Chemical Method with Multipole Electrostatics and Density-Dependent Dispersion Contributions. *J Chem Theory Comput* **15**, 1652-1671, doi:10.1021/acs.jctc.8b01176 (2019).

Wang, R., Fang, X., Lu, Y. & Wang, S. The PDBbind database: collection of binding affinities for protein-ligand complexes with known three-dimensional structures. *J Med Chem* **47**, 2977-2980, doi:10.1021/jm030580l (2004).

Bhagavat, R., Sankar, S., Srinivasan, N. & Chandra, N. An Augmented Pocketome: Detection and Analysis of Small-Molecule Binding Pockets in Proteins of Known 3D Structure. *Structure* **26**, 499-512 e492, doi:10.1016/j.str.2018.02.001 (2018).

Watts, K. S. *et al.* ConfGen: a conformational search method for efficient generation of bioactive conformers. *J Chem Inf Model* **50**, 534-546, doi:10.1021/ci100015j (2010).

Konze, K. D. *et al.* Reaction-Based Enumeration, Active Learning, and Free Energy Calculations To Rapidly Explore Synthetically Tractable Chemical Space and Optimize Potency of Cyclin-Dependent Kinase 2 Inhibitors. *J Chem Inf Model* **59**, 3782-3793, doi:10.1021/acs.jcim.9b00367 (2019).

RDKit. RDKit: Open-source cheminformatics. *http://www.rdkit.org* (2021).

O'Boyle, N. M., Morley, C. & Hutchison, G. R. Pybel: a Python wrapper for the OpenBabel cheminformatics toolkit. *Chem Cent J* **2**, 5, doi:10.1186/1752-153X-2-5 (2008).

Dong, J. *et al.* PyBioMed: a python library for various molecular representations of chemicals, proteins and DNAs and their interactions. *J Cheminform* **10**, 16, doi:10.1186/s13321-018-0270-2 (2018).

Takacs, G., Sandor, M., Szalai, Z., Kiss, R. & Balogh, G. T. Analysis of the uncharted, druglike property space by self-organizing maps. *Mol Divers*, doi:10.1007/s11030-021-10343-y (2021).

Vettigli, G. MiniSom: minimalistic and NumPy-based implementation of the Self Organizing Map. (2018).

Lu, T. & Chen, F. Multiwfn: a multifunctional wavefunction analyzer. *J Comput Chem* **33**, 580-592, doi:10.1002/jcc.22885 (2012).

O'Boyle, N. M. *et al.* Open Babel: An open chemical toolbox. *J Cheminform* **3**, 33, doi:10.1186/1758-2946-3-33 (2011).

Pedregosa, F. a. V., G. and Gramfort, A. and Michel, V. and Thirion, B. and Grisel, O. and Blondel, M. and Prettenhofer, P. and Weiss, R. and Dubourg, V. and Vanderplas, J. and Passos, A. and Cournapeau, D. and Brucher, M. and Perrot, M. and Duchesnay, E. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12**, 2825-2830 (2011).