

# Deep learning enables rapid identification of potent DDR1 kinase inhibitors

Alex Zhavoronkov<sup>1\*</sup>, Yan A. Ivanenkov<sup>1</sup>, Alex Aliper<sup>1</sup>, Mark S. Veselov<sup>1</sup>, Vladimir A. Aladinskiy<sup>1</sup>, Anastasiya V. Aladinskaya<sup>1</sup>, Victor A. Terentiev<sup>1</sup>, Daniil A. Polykovskiy<sup>1</sup>, Maksim D. Kuznetsov<sup>1</sup>, Arip Asadulaev<sup>1</sup>, Yuri Volkov<sup>1</sup>, Artem Zholus<sup>1</sup>, Rim R. Shayakhmetov<sup>1</sup>, Alexander Zhebrak<sup>1</sup>, Lidiya I. Minaeva<sup>1</sup>, Bogdan A. Zagribelnyy<sup>1</sup>, Lennart H. Lee<sup>2</sup>, Richard Soll<sup>2</sup>, David Madge<sup>2</sup>, Li Xing<sup>2</sup>, Tao Guo<sup>2</sup> and Alán Aspuru-Guzik<sup>3,4,5,6</sup>

**We have developed a deep generative model, generative tensorial reinforcement learning (GENTRL), for de novo small-molecule design. GENTRL optimizes synthetic feasibility, novelty, and biological activity. We used GENTRL to discover potent inhibitors of discoidin domain receptor 1 (DDR1), a kinase target implicated in fibrosis and other diseases, in 21 days. Four compounds were active in biochemical assays, and two were validated in cell-based assays. One lead candidate was tested and demonstrated favorable pharmacokinetics in mice.**

Drug discovery is resource intensive, and involves typical timelines of 10–20 years and costs that range from US\$0.5 billion to US\$2.6 billion<sup>1,2</sup>. Artificial intelligence promises to accelerate this process and reduce costs by facilitating the rapid identification of compounds<sup>3,4</sup>. Deep generative models are machine learning techniques that use neural networks to produce new data objects. These techniques can generate objects with certain properties, such as activity against a given target, that make them well suited for the discovery of drug candidates. However, few examples of generative drug design have achieved experimental validation involving synthesis of novel compounds for in vitro and in vivo investigation<sup>5–16</sup>.

Discoidin domain receptor 1 (DDR1) is a collagen-activated pro-inflammatory receptor tyrosine kinase that is expressed in epithelial cells and involved in fibrosis<sup>17</sup>. However, it is not clear whether DDR1 directly regulates fibrotic processes, such as myofibroblast activation and collagen deposition, or earlier inflammatory events that are associated with reduced macrophage infiltration. Since 2013, at least eight chemotypes have been published as selective DDR1 (or DDR1 and DDR2) small-molecule inhibitors (Supplementary Table 1). Recently, a series of highly selective, spiro-indoline-based DDR1 inhibitors were shown to have potential therapeutic efficacy against renal fibrosis in a *Col4a3*<sup>-/-</sup> mice model of Alport syndrome<sup>18</sup>. A wider diversity of DDR1 inhibitors would therefore enable further basic understanding and therapeutic intervention.

We developed generative tensorial reinforcement learning (GENTRL), a machine learning approach for de novo drug design. GENTRL prioritizes the synthetic feasibility of a compound, its effectiveness against a given biological target, and how distinct it is from other molecules in the literature and patent space. In this work, GENTRL was used to rapidly design novel compounds that are active against DDR1 kinase. Six of these compounds, each complying with Lipinski's rules<sup>1</sup>, were designed, synthesized, and

experimentally tested in 46 days, which demonstrates the potential of this approach to provide rapid and effective molecular design (Fig. 1a).


To create GENTRL, we combined reinforcement learning, variational inference, and tensor decompositions into a generative two-step machine learning algorithm (Supplementary Fig. 1)<sup>19</sup>. First, we learned a mapping of chemical space, a set of discrete molecular graphs, to a continuous space of 50 dimensions. We parameterized the structure of the learned manifold in the tensor train format to use partially known properties. Our auto-encoder-based model compresses the space of structures onto a distribution that parameterizes the latent space in a high-dimensional lattice with an exponentially large number of multidimensional Gaussians in its nodes. This parameterization ties latent codes and properties, and works with missing values without their explicit input. In the second step, we explored this space with reinforcement learning to discover new compounds.

GENTRL uses three distinct self-organizing maps (SOMs) as reward functions: the trending SOM, the general kinase SOM, and the specific kinase SOM. The trending SOM is a Kohonen-based reward function that scores compound novelty using the application priority date of structures that have been disclosed in patents. Neurons that are abundantly populated with novel chemical entities reward the generative model. The general kinase SOM is a Kohonen map that distinguishes kinase inhibitors from other classes of molecules. The specific kinase SOM isolates DDR1 inhibitors from the total pool of kinase-targeted molecules. GENTRL prioritizes the structures it generates by using these three SOMs in sequence.

We used six data sets to build the model: (1) a large set of molecules derived from a ZINC data set, (2) known DDR1 kinase inhibitors, (3) common kinase inhibitors (positive set), (4) molecules that act on non-kinase targets (negative set), (5) patent data for biologically active molecules that have been claimed by pharmaceutical companies, and (6) three-dimensional (3D) structures for DDR1 inhibitors (Supplementary Table 1). Data sets were preprocessed to exclude gross outliers and to reduce the number of compounds that contained similar structures (see Methods).

We started to train GENTRL (pretraining) on a filtered ZINC database (data set 1, described earlier), and then continued training using the DDR1 and common kinase inhibitors (data set 2 and data set 3). We then launched the reinforcement learning stage with the reward described earlier. We obtained an initial output of 30,000 structures (Supplementary Data Set), which were then

<sup>1</sup>Insilico Medicine Hong Kong Ltd, Pak Shek Kok, New Territories, Hong Kong. <sup>2</sup>WuXi AppTec Co., Ltd, Shanghai, China. <sup>3</sup>Department of Chemistry, University of Toronto, Toronto, Ontario, Canada. <sup>4</sup>Department of Computer Science, University of Toronto, Toronto, Ontario, Canada. <sup>5</sup>Vector Institute for Artificial Intelligence, Toronto, Ontario, Canada. <sup>6</sup>Canadian Institute for Advanced Research, Toronto, Ontario, Canada. \*e-mail: alex@insilico.com



**Fig. 1 | GENTRL model design, workflow, and nanomolar hits.** **a**, The general workflow and timeline for the design of lead candidates using GENTRL. IP, intellectual property. **b**, Representative examples of generated hits structures compared to the parent DDR1 kinase inhibitor. **c**, Generated compounds with the highest inhibition activity against human DDR1 kinase.

automatically filtered to remove molecules bearing structural alerts or reactive groups, and the resulting chemical space was reduced by clustering and diversity sorting (Supplementary Table 2). We then evaluated structures using (1) the general and specific kinase SOMs, and (2) pharmacophore modeling on the basis of crystal structures of compounds in complex with DDR1 (Supplementary Figs. 2 and 3). On the basis of the values of molecular descriptors and root-mean-square deviation (RMSD) calculated in two previous steps (steps 6 and 7), we used Sammon mapping to assess the distribution of the remaining structures.

To narrow our focus to a smaller set of molecules for analysis, we randomly selected 40 structures that smoothly covered the resulting chemical space and distribution of RMSD values (Supplementary Fig. 4 and Supplementary Table 3). Of the 40 selected structures, 39 were likely to fall outside the scope of any published patents or applications (Supplementary Table 4). Six of these were chosen for experimental validation on the basis of synthetic accessibility. Of note, our approach led to several examples of nontrivial potentially bioisosteric replacements and topological modifications (Fig. 1b).

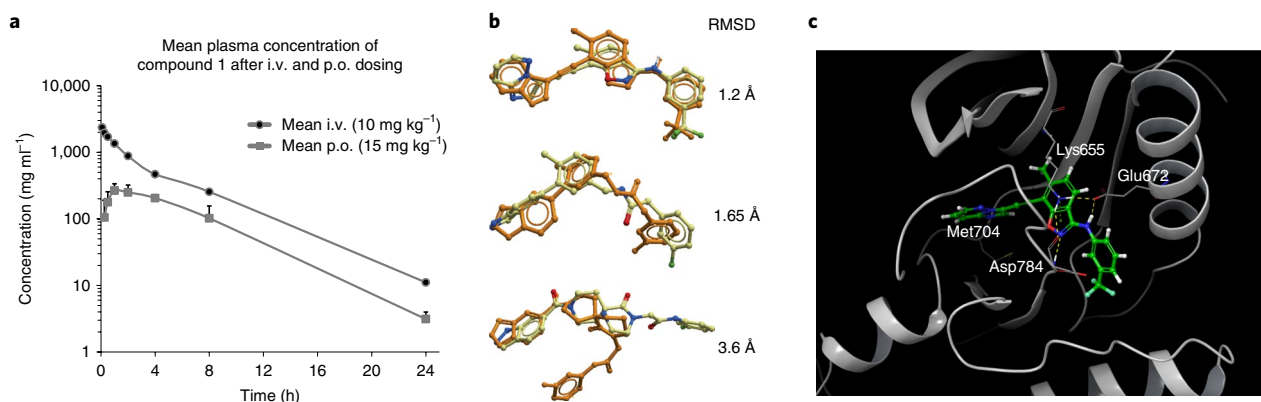
By day 23 after target selection, we had identified six lead candidates, and by day 35, these molecules had been successfully synthesized (Fig. 1c). They were then tested for *in vitro* inhibitory activity in an enzymatic kinase assay (Supplementary Fig. 5). Compounds 1 and 2 strongly inhibited DDR1 activity (half-maximum inhibitory concentration ( $IC_{50}$ ) values of 10 and 21 nM, respectively), compounds 3 and 4 demonstrated moderate potency ( $IC_{50}$  values of 1  $\mu$ M and 278 nM, respectively), and compounds 5 and 6 were inactive. Compounds 1 and 2 both exhibited selectivity towards DDR1

over DDR2 (Fig. 1c). Furthermore, compound 1 exhibited a relatively high selectivity index compared to those of 44 diverse kinases (Supplementary Fig. 6).

Next, we investigated the DDR1 inhibitory activity of compound 1 and compound 2 as measured by autophosphorylation in U2OS cells. The compounds showed  $IC_{50}$  values of 10.3 and 5.8 nM, respectively (Supplementary Fig. 7). Both molecules inhibited the induction of fibrotic markers  $\alpha$ -actin and CCN2 in MRC-5 lung fibroblasts (Supplementary Fig. 8). These molecules also inhibited the expression of collagen (a hallmark of fibrosis) in LX-2 hepatic stellate cells, with compound 1 showing potent activity at 13 nM (Supplementary Fig. 9).

We then performed *in vitro* microsomal stability studies to characterize the metabolic stability of compounds 1 and 2 in human, rat, mouse, and dog liver microsomes. Compounds 1 and 2 had half-life and clearance values that were similar to or more favorable than those of routinely used control molecules (Supplementary Table 5). Compound 2 was also found to be very stable in buffer conditions (Supplementary Table 6). Neither compound strongly inhibited cytochrome P450, and both compounds showed favorable physicochemical properties, including satisfying Lipinski's rules (Supplementary Tables 7 and 8).

Finally, we tested compound 1 in a rodent model. Compound 1 was delivered to mice intravenously (i.v.) (10 mg  $kg^{-1}$ ) and orally (p.o., 15 mg  $kg^{-1}$ ). The two administrations resulted in similar half-lives,  $\sim$ 3.5 h (Fig. 2a and Supplementary Tables 9 and 10). I.v. administration conferred a peak plasma concentration of 2,357 ng  $ml^{-1}$  on initial delivery, whereas p.o. administration resulted in a lower maximum of 266 ng  $ml^{-1}$ , which peaked 1 h after delivery.



**Fig. 2 | Pharmacokinetic characterization and structural basis of hit activity.** **a**, Plasma concentrations of compound 1 in mouse pharmacokinetic study at doses of 10 and 15 mg kg<sup>-1</sup> for i.v. and p.o. treatment, respectively. Measure of center is mean; error bars are s.d.;  $n=3$  biologically independent animals used for each route of administration. **b**, The rigid alignment of a conformation that best fit the pharmacophore hypothesis and a conformation predicted by quantum mechanical calculations. Superpositions are presented for compound 1, compound 3, and compound 5. Orange, quantum mechanical calculation; yellow, pharmacophore modeling. **c**, The putative binding mode of compound 1 ( $IC_{50}=10$  nM) in DDR1 kinase (PDB code: 3Z0S) derived from docking simulations. The receptor is shown in gray; compound 1 is shown as sticks and balls, and key receptor residues that are involved in ligand binding are shown as sticks. Hydrogen bonds are shown as yellow dashed lines.

Quantum mechanical analysis was used to explore the mechanistic basis of the activity of compound 1. The predicted conformation of compound 1 according to pharmacophore modeling was very similar to the conformation predicted to be preferred and stable by quantum mechanical calculations (Fig. 2b). We proposed a 'lock and key' entropy-driven binding mechanism between compound 1 and DDR1, and further characterized this binding via molecular docking. The putative binding mode suggests a type II inhibition mechanism (Fig. 2c). In summary, compound 1 forms multiple hydrogen bonds and has favorable charge and hydrophobic interactions with the active site residues of DDR1 kinase. The complementarity of compound 1 to the ATP site may help to explain its inhibitory activity against DDR1.

Despite reasonable microsomal stability and pharmacokinetic properties, the compounds that have been identified here may require further optimization in terms of selectivity, specificity, and other medicinal chemistry properties.

In this work, we designed, synthesized, and experimentally validated molecules targeting DDR1 kinase in less than 2 months and for a fraction of the cost associated with a traditional drug discovery approach<sup>1</sup>. This illustrates the utility of our deep generative model for the successful, rapid design of compounds that are synthetically feasible, active against a target of interest, and potentially innovative with respect to existing intellectual properties. We anticipate that this technology will be improved further as a useful tool to identify drug candidates.

### Online content

Any methods, additional references, Nature Research reporting summaries, source data, statements of code and data availability and associated accession codes are available at <https://doi.org/10.1038/s41587-019-0224-x>.

Received: 1 November 2018; Accepted: 12 July 2019;  
Published online: 2 September 2019

### References

- Paul, S. M. et al. *Nat. Rev. Drug Discov.* **9**, 203–214 (2010).
- Avorn, J. N. *Engl. J. Med.* **372**, 1877–1879 (2015).
- Goodfellow, I. et al. Generative adversarial nets. in *Advances in Neural Information Processing Systems* 2672–2680 (2014).
- Mamoshina, P. et al. *Mol. Pharm.* **13**, 1445–1454 (2016).
- Sanchez-Lengeling, B. & Aspuru-Guzik, A. *Science* **361**, 360–365 (2018).
- Kadurin, A. et al. *Oncotarget* **8**, 10883–10890 (2016).

- Kadurin, A. et al. *Mol. Pharm.* **14**, 3098–3104 (2017).
- Gómez-Bombarelli, R. et al. *ACS Cent. Sci.* **4**, 268–276 (2018).
- Putin, E. et al. *Mol. Pharm.* **15**, 4386–4397 (2018).
- Putin, E. et al. *J. Chem. Inf. Model.* **58**, 1194–1204 (2018).
- Harel, S. & Radinsky, K. *Mol. Pharm.* **15**, 4406–4416 (2018).
- Polykovskiy, D. et al. *Mol. Pharm.* **15**, 4398–4405 (2018).
- Kuzminykh, D. et al. *Mol. Pharm.* **15**, 4378–4385 (2018).
- Segler, M. H. S. et al. *Nature* **555**, 604–610 (2018).
- Merk, D. et al. *Mol. Inform.* **37**, 1–2 (2018).
- Merk, D. et al. *Commun. Chem.* **1.1**, 68 (2018).
- Moll, S. et al. *Biochim. Biophys. Acta Mol. Cell Res.* <https://doi.org/10.1016/j.bbamcr.2019.04.004> (2019).
- Richter, H. et al. *ACS Chem. Biol.* **14**, 37–49 (2019).
- Elton, D. C. et al. *Mol. Syst. Des. Eng.* **4**, 828–849 (2019).

### Acknowledgements

The authors thank T. Oprea (University of New Mexico School of Medicine) for the valuable contributions, review, and assessment of the novelty of the intellectual property generated by GENTRL. The authors would like to thank NVIDIA Corporation and M. Berger for providing early access to the graphics processing equipment used for deep learning applications by Insilico Medicine. The authors acknowledge T. Lu, L. Duan, Y. Hu, and the WuXi AppTec chemistry team for providing chemical synthesis of the presented compounds. The authors thank S. Djuric, whose valuable comments informed further experiments.

### Author contributions

A. Zhavoronkov, Y.A.I., and A.A. led the project, designed and planned the experiments, and wrote the manuscript. M.S.V., V.A.A., A.V.A., and V.A.T. planned and performed computational chemistry experiments. D.A.P., M.D.K., A. Zholus, A.A., Y.V., R.R.S., and A. Zhebrak developed and implemented the GENTRL. L.I.M. curated chemical synthesis, and B.A.Z. collected and prepared the data. L.H.L., R.S., D.M., L.X., and T.G. helped write the manuscript. A.A.-G. provided manuscript and methodological feedback.

### Competing interests

A. Zhavoronkov, Y.A.I., A. Aliper, M.S.V., V.A.A., A.V.A., V.A.T., D.A.P., M.D.K., A. Zholus, A. Asadulaev, Y.V., A. Zhebrak, R.R.S., L.I.M., and B.A.Z. work for Insilico Medicine, a commercial artificial intelligence company. L.H.L., R.S., D.M., L.X., and T.G. work for WuXi AppTec, a commercial research organization. A.A.-G. is a cofounder and board member of, and consultant for, Kebotix, an artificial intelligence-driven molecular discovery company and a member of the science advisory board of Insilico Medicine.

### Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41587-019-0224-x>.

Reprints and permissions information is available at [www.nature.com/reprints](http://www.nature.com/reprints).

Correspondence and requests for materials should be addressed to A.Z.

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

© The Author(s), under exclusive licence to Springer Nature America, Inc. 2019

## Methods

**Pretraining data set.** For the pretraining procedure, we have prepared a data set of structures using the Clean Leads set from the ZINC database<sup>20</sup> and proprietary databases from our partners. We have removed structures containing atoms other than carbon, nitrogen, oxygen, sulfur, fluorine, chlorine, bromine, and hydrogen. Routine medicinal chemistry filters were applied to exclude compounds with potentially toxic and reactive groups.

**Kinase inhibitors and 'negative' data set.** The data set of molecules that actively inhibit and do not inhibit various kinases was prepared using the data available in the Integrity and ChEMBL databases.

**Compounds from patent records by priority date.** The Integrity database was used to collect the data set of structures claimed as new drug substances in patent records from 1950 to the present day by the top ten pharmaceutical companies (as ranked by market capitalization in 2017 according to <https://www.globaldata.com>). The final data set contained 17,000 records.

**Model.** Our generative pipeline was created using the GENTRL model, a variational auto-encoder with a rich prior distribution in the latent space (Supplementary Code and Supplementary Fig. 1). We used tensor decomposition to encode the relationships between molecular structures and their properties, and trained a model in a semisupervised fashion without imputing unknown biochemical properties of molecules.

The tensor-train decomposition<sup>21</sup> approximates high-dimensional tensors using a relatively small number of parameters. A joint distribution  $p(r_1, r_2, \dots, r_n)$  of discrete random variables  $r_i \in \{0, \dots, N_i - 1\}$  can be represented as elements of  $n$ -dimensional tensor:

$$p(r_1, r_2, \dots, r_n) = \frac{1}{Z} \mathbf{1}_m \cdot \prod_{i=1}^n Q_i[r_i] \cdot \mathbf{1}_m^T$$

where tensors  $Q_i \in \mathbb{R}_+^{N_i \times m \times m}$  are cores,  $\mathbf{1}_m$  is a vector of ones, and  $Z$  is a normalizing constant. With larger core sizes, the flexibility of the model improves, although the number of parameters grows quadratically with core size  $m$ . In tensor train, we can efficiently marginalize the distribution with respect to any variable, as follows:

$$p(r_1, \dots, r_{k-1}, r_{k+1}, \dots, r_n) = \frac{1}{Z} \mathbf{1}_m \cdot \left( \prod_{i=1}^k Q_i[r_i] \right) \cdot \tilde{Q}_k \cdot \left( \prod_{i=k+1}^n Q_i[r_i] \right) \cdot \mathbf{1}_m^T$$

where  $\tilde{Q}_k = \sum_{r_k} Q_k[r_k]$  can be computed efficiently. With marginal distributions, we can compute the conditional distributions and sample using a chain rule. The normalizing constant  $Z$  is given by

$$Z = \mathbf{1}_m \cdot \prod_{i=1}^n \tilde{Q}_i \cdot \mathbf{1}_m^T$$

As generative auto-encoders use continuous latent codes, we use continuous tensor-train representation. For simplicity of notation, assume that latent codes  $z$  are continuous and properties  $y$  are discrete. We approximate distributions  $p_\psi(z_i)$  as mixtures of Gaussians with component index  $s$ . The joint distribution on  $z$  and  $y$  is

$$p_\psi(z, y) = \sum_{s_1, \dots, s_d} p_\psi(z, y, s) = \sum_{s_1, \dots, s_d} P[s, y] \cdot p_\psi(z | y, s)$$

For conditional distribution  $p_\psi(z|y, s)$ , we select a fully factorized Gaussian that does not depend on  $y$ :

$$p_\psi(z|y, s) = p_\psi(z|s) = \prod_{k=1}^d \mathcal{N}(z_k | \mu_{k, s_k}, \sigma_{k, s_k}^2)$$

The tunable parameters  $\psi$  of the distribution  $p_\psi$  are tensor-train cores  $Q_i$ , means  $\mu_{k, s_k}$ , and variances  $\sigma_{k, s_k}^2$  of the Gaussian components. We store tensor  $P[s, y]$  in a tensor-train format. The resulting distribution becomes

$$p_\psi(z, y) = \sum_{s_1, \dots, s_d} P[s, y] \cdot \prod_{k=1}^d \mathcal{N}(z_k | \mu_{k, s_k}, \sigma_{k, s_k}^2)$$

Our model is a variational auto-encoder with a prior distribution  $p_\psi(z, y)$ , encoder  $q_\phi$ , and a decoder  $p_\theta$ . Consider a training example  $(x, y_{ob})$ , where  $x$  is a molecule and  $y_{ob}$  are its known properties. The lower bound on a log-marginal likelihood (also known as the evidence lower bound) for our model is

$$L(\theta, \phi, \psi) = \mathbb{E}_{q_\phi(z|x, y_{ob})} (\log p_\theta(x|z, y_{ob}) + \log p_\psi(y_{ob}|z)) - \mathcal{KL}(q_\phi(z|x, y_{ob}) || p_\psi(z|y_{ob}))$$

As the molecule determines its properties, we assume that  $q_\phi(z|x, y_{ob}) = q_\phi(z|x)$ . We also assume that  $p_\theta(x|z, y_{ob}) = p_\theta(x|z)$ , indicating that an object is fully defined by its latent code. The resulting evidence lower bound is

$$\begin{aligned} L(\theta, \phi, \psi) &= \mathbb{E}_{q_\phi(z|x)} (\log p_\theta(x|z) + \log p_\psi(y_{ob}|z)) \\ &\quad - \mathcal{KL}(q_\phi(z|x) || p_\psi(z|y_{ob})) \\ &\approx \frac{1}{l} \sum_{i=1}^l \left[ \log p_\theta(x|z_i) + \log p_\psi(y_{ob}|z_i) - \log \frac{q_\phi(z_i|x)}{p_\psi(z_i|y_{ob})} \right] \end{aligned}$$

where  $z_i \sim q_\phi(z|x)$ . For the proposed joint distribution  $p_\psi(z, y)$ , we can compute the density of the posterior distribution on the latent codes, given observed properties  $p_\psi(z|y_{ob})$ , analytically.

By maximizing the evidence lower bound, we trained an auto-encoder and a prior distribution on three data sets described above (pretraining, kinase and patent data sets): we sampled molecules in a simplified molecular input line entry system (SMILES) format from the data set along with their properties, including MCE-18, pIC<sub>50</sub> (negative common logarithm of IC<sub>50</sub>) and a binary feature that indicates whether a molecule passed medicinal chemistry filters (MCFs). We trained this model and obtained a mapping from the chemical space to the latent codes. This mapping was aware of the relationship between molecules and their biochemical properties.

In the next stage of training, we fine-tuned the model to preferentially generate DDR1 kinase inhibitors. We used reinforcement learning to expand the latent manifold towards discovering novel inhibitors with reward functions (general kinase SOM, specific kinase SOM, and trending SOM), which are described in the next section. We used the REINFORCE<sup>22</sup> algorithm (also known as a log-derivative trick) to directly optimize the model:

$$\max_\psi \mathbb{E}_{z \sim p_\psi(z)} R(z), \quad R(z) = \mathbb{E}_{x \sim p_\theta(x|z)} [R_{\text{general}}(x) + R_{\text{specific}}(x) + R_{\text{trending}}(x)]$$

$$\nabla_\psi \mathbb{E}_{z \sim p_\psi(z)} R(z) = \mathbb{E}_{z \sim p_\psi(z)} \nabla_\psi \log p_\psi(z) \cdot R(z)$$

We reduced the variance of the gradient using a standard variance reduction technique called a 'baseline'. The rewards for each molecule in a batch are calculated and averaged, and the average reward is then subtracted from each individual reward:

$$\nabla_\psi \mathbb{E}_{z \sim p_\psi(z)} R(z) \approx \frac{1}{N} \sum_{i=1}^N \nabla_\psi \log p_\psi(z_i) \left[ R(z_i) - \frac{1}{N} \sum_{j=1}^N R(z_j) \right]$$

To preserve the mapping of the chemical space, we fixed the parameters of the encoder and decoder, and trained only the manifold distribution  $p_\psi(z)$ . We combined exploration and exploitation approaches. For exploration, we sampled  $z^{\text{explore}} \sim \mathcal{N}(\mu, (2\sigma)^2)$  outside from the currently explored latent space, where  $\mu$  and  $\sigma^2$  are the mean and variance of  $p_\psi(z)$  for all dimensions. If the reward  $R(z^{\text{explore}})$  for a newly discovered area was high, the latent manifold expanded toward it (Supplementary Fig. 1).

The comparison of generative chemistry models is very important for the advancement of this emerging field, and there are several benchmarking platforms in development<sup>12,23</sup>. We successfully compared the performance of GENTRL with previous approaches, including objective-reinforced generative adversarial networks (ORGAN)<sup>24,25</sup>, reinforced adversarial neural computer (RANC)<sup>10</sup>, and adversarial threshold neural computer (ATNC)<sup>9</sup>. Training details are provided in the Supplementary Note.

**Reward function.** A reward function was developed on the basis of the Kohonen self-organizing maps (SOM)<sup>26</sup> (Supplementary Fig. 3). This algorithm was introduced by Teuvo Kohonen as a unique unsupervised machine-learning dimensionality reduction technique. It can effectively reproduce an intrinsic topology and patterns hidden in the input chemical space in a faithful and unbiased fashion. The input chemical space is usually described in terms of molecular descriptors (input vector), and the output typically includes a 2D or 3D feature map for convenient visual inspection. An ensemble of three SOMs was used as a reward function: the first SOM (general kinase SOM,  $R_{\text{general}}$ ) was trained to predict the activity of compounds against kinases, the second SOM (specific kinase SOM,  $R_{\text{specific}}$ ) was developed to select compounds located in neurons associated with DDR1 inhibitors within the whole kinase map, and the last SOM (trending SOM,  $R_{\text{trending}}$ ) was trained to assess the novelty of chemical structures in terms of the current trends in medicinal chemistry. During learning, the generative model was rewarded when the generated structures were classified as molecules acting on kinases, positioned in neurons attributed to DDR1 inhibitor. The model was also rewarded for generating novel structures.

**Pharmacophore hypotheses.** On the basis of X-ray data available in the Protein Data Bank (PDB) database (PDB codes 3ZOS, 4BKJ, 4CKR, 5BVN, 5BVO, 5FDP,

5FDX, and 6GWR), we have developed three pharmacophore models describing DDR1 inhibitors. To obtain the superposition of the ligands, 3D alignment of the complexes was carried out. These three-, four- and five-centered pharmacophore hypotheses contain key features that are responsible for binding to the active site of DDR1 kinase, including a hydrogen bond acceptor at the hinge region, an aromatic or hydrophobic linker, and a hydrophobic center in the pocket located in proximity to the DFG motif. For detailed information on pharmacophore features and distances, see Supplementary Fig. 2.

**Nonlinear Sammon mapping.** To make the final selection, we used a Sammon-based mapping technique<sup>27</sup>. The main goal of this algorithm lies in the approximation of local geometric and topological relationships hidden in the input chemical space on a visually intelligible 2D or 3D plot. The fundamental idea of this method is to substantially reduce the high dimensionality of the initial data set into the low-dimensional feature space, and, in this aspect, it resembles an SOM approach with multidimensional scaling. However, in contrast to other algorithms, a classical Sammon-based method allows scientists to construct a projection that reflects global topographic relationships as pair-wise distances between all of the objects within the whole space of input vector samples. Structures that successfully passed all of the selection procedures described earlier were used as an input chemical space. For mapping, we used the same set of molecular descriptors that was applied for specific kinase SOM and added RMSD values obtained during pharmacophore modeling as additional inputs. Euclidean distances were used as a similarity metric. The stress threshold was 0.01, the interaction number was 300, the optimization step was 0.3 and the structural similarity factor was 0.5. The resulting map (Supplementary Fig. 4) demonstrates that structures are normally distributed within the Sammon plot.

**Molecule generation and selection procedure.** Using our model, we generated 30,000 unique valid structures by sampling latent codes from the learned manifold  $p_v(z)$  and sampling structures from the decoder distribution  $p_r(x|z)$ . To select the batch of molecules for synthesis and biological studies, we developed a prioritization pipeline (for examples of rejected molecules, see Supplementary Fig. 10). At the initial step, the data set was reduced to 12,147 compounds using the following molecular descriptor thresholds:  $-2 < \log P < 7$ ,  $250 < MW < 750$ ,  $HBA + HBD < 10$ ,  $TPSA < 150$ , and  $NRB < 10$ . After that, 150 in-house MCFs were applied to remove potentially toxic structures and compounds containing reactive and undesirable groups. These include substrates for 1,4-addition (Michael-bearing moieties) and other electrophilic species (for example, *para*- or *ortho*-halogen-substituted pyridines, 2-halogen-substituted furans and thiophenes, alkyl halides, and aldehydes and anhydrides), disulfides, isatins, barbiturates, strained heterocycles, fused polyaromatic systems, detergents, hydroxamic acids and diazo-compounds, peroxides, unstable fragments, and sulfonyl ester derivatives. In addition, we used more trivial filtering rules that excluded the following:  $< 2$  NO<sub>2</sub> groups,  $< 3$  Cl,  $< 2$  Br,  $< 6$  F, and  $< 5$  aromatic rings, and undesired atoms, such as silicon, cobalt or phosphorus. This reduced the number of structures spread within the entire chemical space to drug-like molecules without structural alerts. This procedure resulted in 7,912 structures. A clustering analysis was then performed using Tanimoto similarity as a metric and standard Morgan fingerprints implemented in the RDKit package. All compounds that satisfied a 0.6 similarity threshold were assigned to the same cluster, with a minimum value of five structures per cluster. Inside each cluster, the compounds were sorted according to their internal dissimilarity coefficient to output the top five items with the maximum diversity in structure. As a result, the data set was reduced to 5,542 molecules. Then, we performed a similarity search using vendors' collections (MolPort (<https://www.molport.com>) and ZINC<sup>18</sup>) and removed a further 900 compounds with similarity  $> 0.5$  to increase the novelty of the generated structures. General kinase SOM and specific kinase SOM were used to prioritize the compounds by their potential activity against DDR1 kinase. Out of 2,570 molecules classified as kinase inhibitors by general kinase SOM, 1,951 molecules were classified as DDR1 inhibitors by specific kinase SOM and were used for pharmacophore-based virtual screening. For every molecule, ten conformations were generated and minimized using RDKit's implementation of the universal force field<sup>28</sup>. Using the developed hypotheses, the screening procedure was carried out, resulting in a set of RMSD values for 848 molecules matching at least one pharmacophore hypothesis. On the basis of Sammon mapping, we uniformly selected 20 molecules from ellipses corresponding to four- and five-centered pharmacophores (Supplementary Table 3 and Supplementary Fig. 4). Forty molecules were selected for synthesis and subsequent biological evaluation.

**Ab initio calculation details.** We carried out first-principles calculations on the lowest conformer as predicted with the universal-force-field methodology presented earlier. Geometry optimization was performed using a local correlated coupled-cluster method that included single and double excitations (LCCSD) with the 6-31++G basis set. Final energies were calculated at the LCCSD(T) level of theory. The localized Pipek–Mezey procedure was used to obtain the initial molecular orbitals.

**Docking simulations.** Molecular modeling was performed in the Maestro suite (<https://www.schrodinger.com>). PDB structure 3ZOS was preprocessed and energy

minimized using the Prep module. The binding site grid was generated around the ATP binding site with 20 Å buffer dimensions. Docking poses were generated by extra-precision (XP) Glide runs using the optimized ligand structure. The final model was selected on the basis of its docking score of  $-15$  kcal mol<sup>-1</sup>, which is lowest among all of the obtained models.

**In vitro activity assays.** The activity of the molecules against human DDR1 and human DDR2 kinases was assessed using KinaseProfiler (Eurofins Scientific).

**Cell-culture activity assay.** To measure autophosphorylation, the gene encoding human DDR1b with a hemagglutinin tag was cloned into pCMV Tet-On vector (Clontech), and stable inducible cell lines established in U2OS were used for the IC<sub>50</sub> test. *DDR1* expression was induced for 48 h before DDR1 activation by rat tail collagen I (Sigma 11179179001). The cells were detached with trypsinization and transferred to a 15 ml tube. Then after pretreatment with the compound for 0.5 h, the cells were treated with compounds in the presence of 10 μg ml<sup>-1</sup> rat tail collagen I for 1.5 h at 37 °C.

**Cell-culture fibrosis assay.** MRC-5 or human hepatic LX-2 cells were grown in reduced serum medium and treated with compounds for 30 minutes. Subsequently, the cells were stimulated with 10 ng ml<sup>-1</sup> or 4 ng ml<sup>-1</sup> TGF-β (R&D Systems, 240-B-002) for 48 or 72 h. The cells were lysed in radioimmunoprecipitation assay buffer and cell lysate of each sample was loaded onto a Wes automated western blot system (ProteinSimple, a Bio-Techne brand).

**Cytochrome inhibition.** Water used in the assay and analysis was purified by ELGA Lab purification systems. Potassium phosphate buffer (PB, concentration of 100 mM) and MgCl<sub>2</sub> (concentration of 33 mM) were used. Test compounds (compound 1 and compound 2) and standard inhibitors (α-naphthoflavone, sulfaphenazole, (+)-*N*-3-benzylirivanol, quinidine, and ketoconazole) working solutions (100×) were prepared. Microsomes were taken out of a freezer ( $-80$  °C) to thaw on ice, labeled with the date, and returned to the freezer immediately after use. Next, 20 μl of the substrate solutions was added to corresponding wells, 20 μl PB was added to blank wells, and 2 μl of the test compounds and positive control working solution was added to corresponding wells. We then prepared a working solution of human liver microsomes (HLM), and 158 μl of the HLM working solution was added to all wells of the incubation plate. The plate was prewarmed for approximately 10 minutes in a water bath at 37 °C. Then, reduced nicotinamide adenine dinucleotide phosphate (NADPH) cofactor solution was prepared and 20 μl NADPH cofactor was added to all incubation wells. The solution was mixed and incubated for 10 minutes in a water bath at 37 °C. At this point, the reaction was terminated by adding 400 μl cold stop solution (200 ng ml<sup>-1</sup> tolbutamide and 200 ng ml<sup>-1</sup> labetalol in acetonitrile (ACN)). The samples were centrifuged at 4,000 r.p.m. for 20 minutes to precipitate protein. Then, 200 μl supernatant was transferred to 100 μl HPLC water and shaken for 10 minutes. Xlfit was used to plot the per cent of vehicle control versus the test compound concentrations, and for nonlinear regression analysis of the data. IC<sub>50</sub> values were determined using three- or four-parameter logistic equation. IC<sub>50</sub> values were reported as  $> 50$  μM when per cent inhibition at the highest concentration (50 μM) was less than 50%.

**Microsomal stability.** The microsomal stability of compound 2 was assessed as follows: working solutions of compound 2 and control compounds (testosterone, diclofenac, and propafenone) were prepared. The appropriate amount of NADPH powder (β-nicotinamide adenine dinucleotide phosphate reduced form, tetrasodium salt, NADPH-4Na, catalog no. 00616; Chem-Impex International) was weighed and diluted into MgCl<sub>2</sub> (10 mM) solution (working solution concentration, 10 units ml<sup>-1</sup>; final concentration in reaction system, 1 unit ml<sup>-1</sup>). The appropriate concentration of microsome working solutions (human: HLM, catalog no. 452117, Corning; SD rat: RLM, catalog no. R1000, Xenotech; CD-1 mouse: MLM, catalog no. M1000, Xenotech; Beagle dog: DLM, catalog no. D1000, Xenotech) was prepared with 100 mM PB. Cold ACN, including 100 ng ml<sup>-1</sup> tolbutamide and 100 ng ml<sup>-1</sup> labetalol as internal standard (IS), was used for the stop solution. Compound or control working solution (10 μl per well) was added to all plates (T0, T5, T10, T20, T30, T60, and NCF60), except the matrix blank. Dispensed microsome solution (80 μl per well) was added to every plate by Apricot and the mixture of microsome solution and compound was incubated at 37 °C for approximately 10 minutes. After prewarming, dispensed NADPH regenerating system (10 μl per well) was added to every plate by Apricot to start a reaction. The solution was then incubated at 37 °C. Stop solution (300 μl per well, 4 °C) was then added to terminate the reaction. The sampling plates were shaken for approximately 10 minutes. The samples were centrifuged at 4,000 r.p.m. for 20 minutes at 4 °C. While centrifuging, new 8 × 96-well plates were loaded with 300 μl HPLC water, and then 100 μl supernatant was transferred and mixed for liquid chromatography–tandem mass spectrometry (LC/MS/MS).

**Buffer stability.** The stability of compound 2 was assessed in phosphate buffer (pH 7.0 and 7.4). Test compounds (at 10 μM) were incubated at 25 °C with 50 mM phosphate buffer (pH 7.4), 8 mM MOPS (pH 7.0), and 0.2 mM EDTA (pH 7.0). Duplicate samples were used. Time samples (0, 120, 240, 360, and 1,440 minutes)

were removed and immediately mixed with cold 50% aqueous ACN solution containing IS. Curcumin was used as positive control in this assay at neutral–basic condition. The samples were analyzed by LC/MS/MS, and the disappearance of the test compound was assessed on the basis of peak area ratios of the analyte and IS (no standard curve).

**Pharmacokinetic studies.** A study was permitted by the Institutional Animal Care and Use Committee, Shanghai Site (IACUC-SH, WuXi AppTec (Shanghai) Co., Ltd.). The pharmacokinetic profiling of compound 1 was performed on male C57BL/6 mice (7–9 weeks old). Then, we performed i.v. (10 mg kg<sup>-1</sup>) and p.o. (15 mg kg<sup>-1</sup>) administration of compound 1. Each group consisted of three mice. *N*-Methyl-2-pyrrolidone:polyethylene glycol 400:H<sub>2</sub>O = 1:7:2 solution was used as a vehicle at 5 and 3 ml kg<sup>-1</sup> for i.v. and PO, respectively. All blood samples (approximately 25 µl blood per time point) were transferred into prechilled commercial K2-EDTA tubes, and then placed on wet ice. The blood samples were immediately processed for plasma by centrifugation at approximately 4°C, 3,200g for 10 minutes. The plasma was transferred into one pre-labeled polypropylene microcentrifuge tube, quick frozen over dry ice, and kept at -70 ± 10°C until LC/MS/MS analysis. Plasma concentration versus time data was analyzed by non-compartmental approaches using the Phoenix WinNonlin 6.3 software program.

**Statistics and reproducibility.** The sample sizes can be found in the figures and tables or corresponding legends. For microsomal stability experiments, *R*<sup>2</sup> values were calculated. The number of samples for each experiment can be found in the footnote to Supplementary Table 4. All western blot experiments were performed at least twice with similar results.

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

### Data availability

All data are available in the main text or the supplementary materials.

### Code availability

The code for the GENTRL model is available at <http://github.com/insilicomedicine/gentrl> and in Supplementary Code.

### References

20. Irwin, J. J. et al. *J. Chem. Inf. Model.* **52**, 1757–1768 (2012).
21. Oseledets, I. V. *SIAM J. Sci. Comput.* **33**, 2295–2317 (2011).
22. Williams, R. J. *Mach. Learn.* **8**, 229–256 (1992).
23. Brown, N. et al. *J. Chem. Inf. Model.* **59**, 1096–1108 (2018).
24. Guimaraes, G. L. et al. Objective-Reinforced Generative Adversarial Networks (ORGAN) for sequence generation models. Preprint at <https://arxiv.org/abs/1705.10843> (2017).
25. Sanchez-Lengeling, B. et al. Optimizing distributions over molecular space. An Objective-Reinforced Generative Adversarial Network for Inverse-design Chemistry (ORGANIC). Preprint at [https://chemrxiv.org/articles/ORGANIC\\_1\\_pdf/5309668](https://chemrxiv.org/articles/ORGANIC_1_pdf/5309668) (2017).
26. Ritter, H. & Kohonen, T. *Biol. Cybern.* **61**, 241–254 (1989).
27. Sammon, J. W. *IEEE Trans. Comput.* **C-18**, 401–409 (1969).
28. Rappe, A. K. *J. Am. Chem. Soc.* **114**, 10024–10035 (1992).

## Reporting Summary

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

### Statistics

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

n/a Confirmed

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

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

### Software and code

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

Data collection

The code for the GENerative adversarial Tensorial Reinforcement Learning (GENTRL) system is provided with the manuscript. It will be freely available at GitHub upon the acceptance of the manuscript.

Data analysis

Maestro Release 2018-3; Phoenix WinNonlin 6.3; custom python scripts using PyTorch 0.4.1 and RDKit 2018.03.4 libraries

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

### Data

Policy information about [availability of data](#)

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

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

The 30,000 structures generated by GENTRL for the DDR1 kinase are available in supplementary materials.

### Field-specific reporting

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

- Life sciences       Behavioural & social sciences       Ecological, evolutionary & environmental sciences

## Life sciences study design

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

Sample size	No applicable.
Data exclusions	Data outside 3xSD were excluded.
Replication	Results are verified by independent biological experiments.
Randomization	Not applicable.
Blinding	Not applicable.

## Reporting for specific materials, systems and methods

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

### Materials & experimental systems

n/a	Involved in the study
<input type="checkbox"/>	<input checked="" type="checkbox"/> Antibodies
<input type="checkbox"/>	<input checked="" type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology
<input type="checkbox"/>	<input checked="" type="checkbox"/> Animals and other organisms
<input checked="" type="checkbox"/>	<input type="checkbox"/> Human research participants
<input checked="" type="checkbox"/>	<input type="checkbox"/> Clinical data

### Methods

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

## Antibodies

Antibodies used	Mouse anti- human COL1A1 antibody (3G3) (Santa Cruz, sc-293182); mouse anti-human CTGF antibody (E-5) (Santa Cruz, sc-365970); mouse anti human $\alpha$ -smooth muscle actin antibody (SPM332) (Santa Cruz, sc-56499); mouse anti-human GLYCERALDEHYDE-3-PDH, (Merck-Millipore, Merck-MAB374); rabbit anti-Phospho-DDR1 (Tyr513) (E1N8F) (Cell Signaling, 14531S); mouse anti-HA antibody (Sigma, H9658-200UL); goat Anti-Mouse IgG H&L (HRP) (Abcam, ab205719); goat Anti-Rabbit IgG H&L (HRP) (Abcam, ab205718)
Validation	The antibodies were validated by the manufacturers for Western blot.

## Eukaryotic cell lines

Policy information about [cell lines](#)

Cell line source(s)	MRC-5 was from ATCC (CCL-171); LX-2 was from Merck Millipore (SCC064); U-2 OS-DDR1 stable cell line was made at RSD Biology Department, WuXi Apptec, Inc. from U-2 OS parent cell line from ATCC® (HTB96™);
Authentication	The cells were authenticated by STR.
Mycoplasma contamination	All cell lines used have been tested free of mycoplasma contamination.
Commonly misidentified lines (See <a href="#">ICLAC</a> register)	None.

## Animals and other organisms

Policy information about [studies involving animals](#); [ARRIVE guidelines](#) recommended for reporting animal research

Laboratory animals	Mouse, C57BL/6, male, 7- 9 weeks
Wild animals	Study did not involve wild animals
Field-collected samples	Following arrival at WuXi AppTec animals were assessed as to their general health by a member of the veterinary staff or other authorized personnel. Animals were acclimated for at least 3 days (upon arrival at WuXi AppTec) before being placed on study.



Animals were group housed during acclimation and individually housed during the study. The animal room environment was controlled (target conditions: temperature 18 to 26°C, relative humidity 30 to 70%, 12 hours artificial light and 12 hours dark). Temperature and relative humidity was monitored daily. The animals were overnight fasted. They had access to Certified Rodent Diet ad libitum 4 hr post dose. The lot number and specifications of each lot used were archived at WuXi AppTec. Water was autoclaved before provided to the animals ad libitum. Periodic analyses of the water was performed and the results archived at WuXi AppTec. There are no known contaminants in the diet or water that, at the levels of detection, is expected to interfere with the purpose, conduct or outcome of the study.

#### Ethics oversight

Study was permitted by IACUC (Institutional Animal Care and Use Committee).

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