Deep Generative Models for Graphs

VAEs, GANs, and Reinforcement Learning for *de novo* drug discovery

Nicola De Cao

The University of Edinburgh
and University of Amsterdam

September 24, 2019

The drug design processes

Which problem do we want to solve?

- **Drug design** is the process of finding new drugs

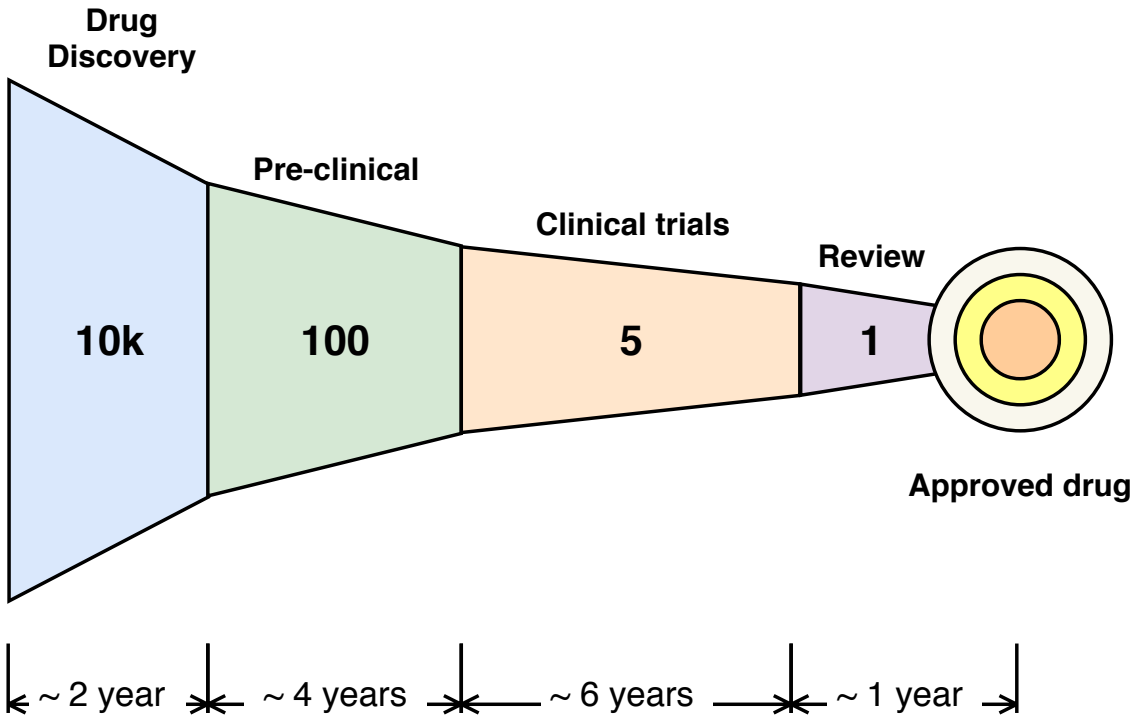

- The first step is **Drug Discovery**
  - screening large compound libraries
  - designing of **new unknown molecules** (*de novo*)

Our contributions

How others proposed to study the problem?

- Generating **SMILES** representations [Gómez-Bombarelli et al., 2016]
- Generating **labeled graphs** [Simonovsky and Komodakis, 2018]

How do we study the problem?

- Using **labeled graphs**
- **Likelihood-based vs. likelihood-free** methods (VAE vs. GAN)
- Biasing the process using **reinforcement learning**

Background

Variational Auto-Encoders

Likelihood-based generative process [Kingma and Welling, 2013]

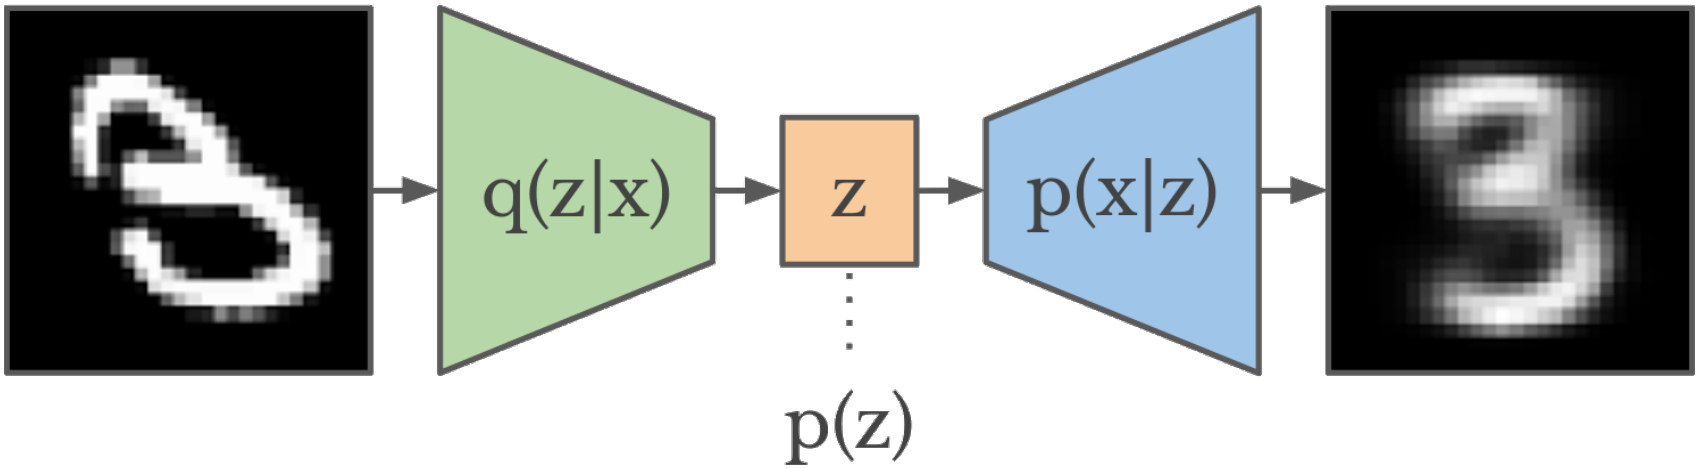

Image credit [Hafner, 2018]

Training VAEs I

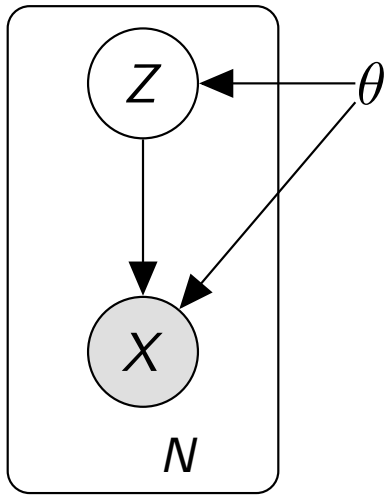

Figure: Graphical model of a simple VAE.

Trained to **maximize the log-evidence**:

$$\log p_\theta(\mathbf{x}^{(i)}) = \log \int p_\theta(\mathbf{x}^{(i)}, \mathbf{z})\, d\mathbf{z}$$

Training VAEs II

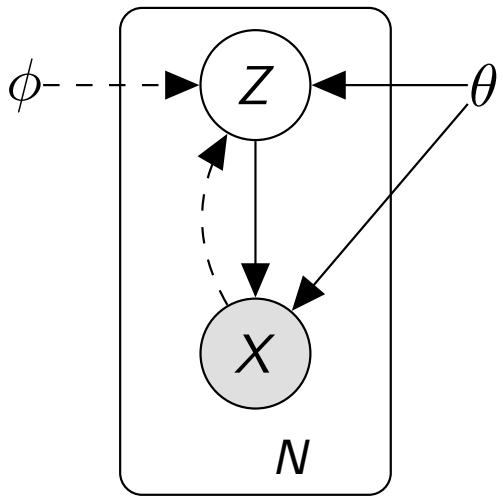

Figure: Graphical model of a simple VAE.

Evidence Lower Bound (ELBO)

Optimizing a **lower bound** of the loss makes the problem feasible:

$$\log p_\theta(\mathbf{x}^{(i)}) \geq \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}^{(i)})}\left[\log p_\theta(\mathbf{x}^{(i)}|\mathbf{z})\right] - D_{KL}\left[\, q_\phi(\mathbf{z}|\mathbf{x}^{(i)}) \parallel p_\theta(\mathbf{z}) \,\right]$$

Generative Adversarial Networks

Likelihood-free generative process [Goodfellow et al., 2014]

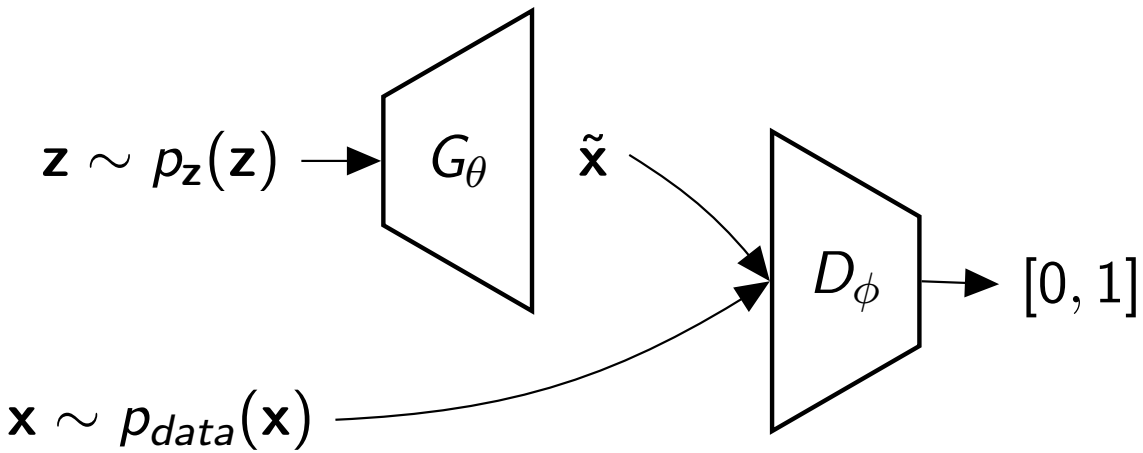

Figure: Schema of **GAN architecture**.

Training GANs

Originally proposed loss:

$$\min_{\theta} \max_{\phi} \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log D_{\phi}(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D_{\phi}(G_{\theta}(\mathbf{z}))]$$

but there are **better alternatives**:

- f-GAN [Nowozin et al., 2016];
- WGAN [Arjovsky et al., 2017];
- WGAN-GP [Gulrajani et al., 2017] (used in this work).

Reinforcement Learning

Can we train a **non-differentiable** generative process?

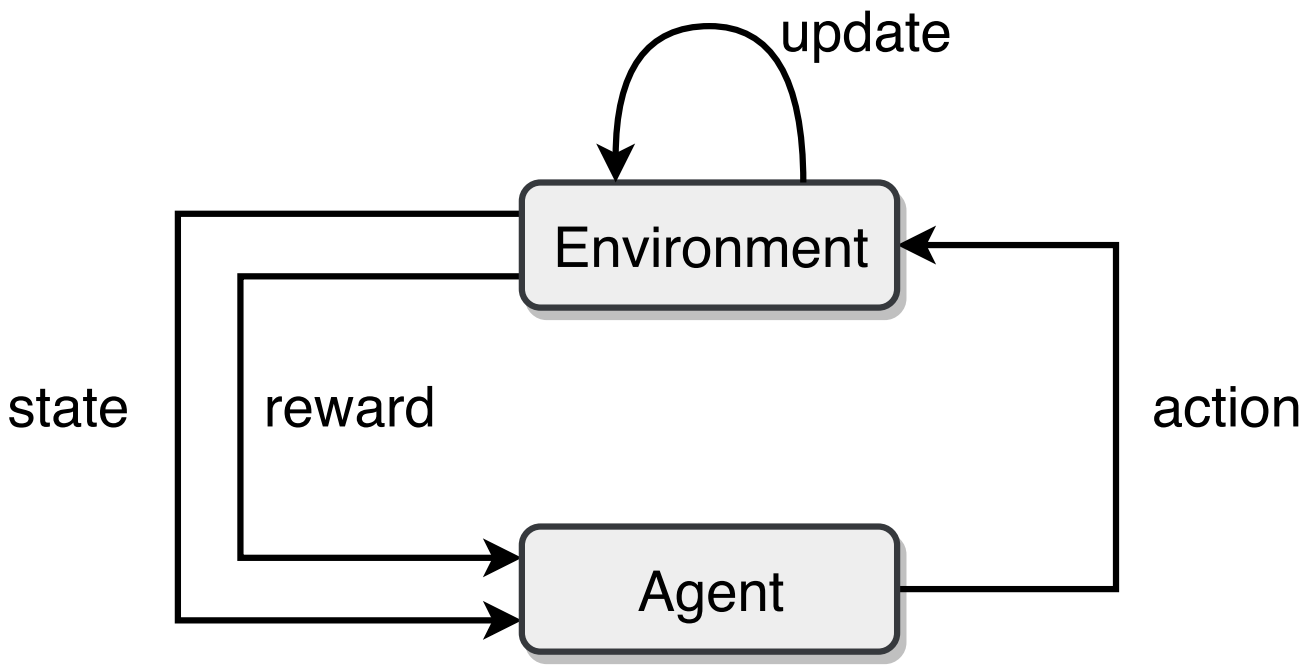

Figure: **Reinforcement Learning** loop schema.

Deep Deterministic Policy Gradient (DDPG) I

How to learn a RL policy?

Deep deterministic policy gradient [Lillicrap et al., 2016]

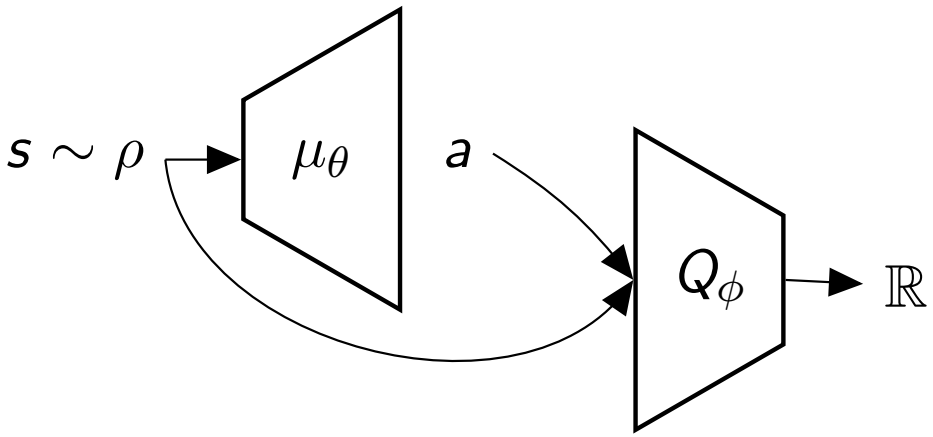

Deep Deterministic Policy Gradient (DDPG) II

Off-policy deterministic policy gradient

Update θ according to [Silver et al., 2014]:

$$\mathcal{L}(\mu_\theta) = \mathbb{E}_{s\sim\rho, a=\mu_\theta(s)}\left[Q_\phi(s,a)\right]$$
$$\nabla_\theta \mathcal{L}(\mu_\theta) = \mathbb{E}_{s\sim\rho, a=\mu_\theta(s)}\left[\nabla_a Q_\phi(s,a)\ \nabla_\theta \mu_\theta(s)\right]$$

and ϕ with:

$$\mathcal{L}(Q_\phi) = \mathbb{E}_{s\sim\rho, a=\mu_\theta(s)}\left[\|Q_\phi(s,a) - r(s,a)\|^2\right]$$

Models

Vectorial representation of graphs I

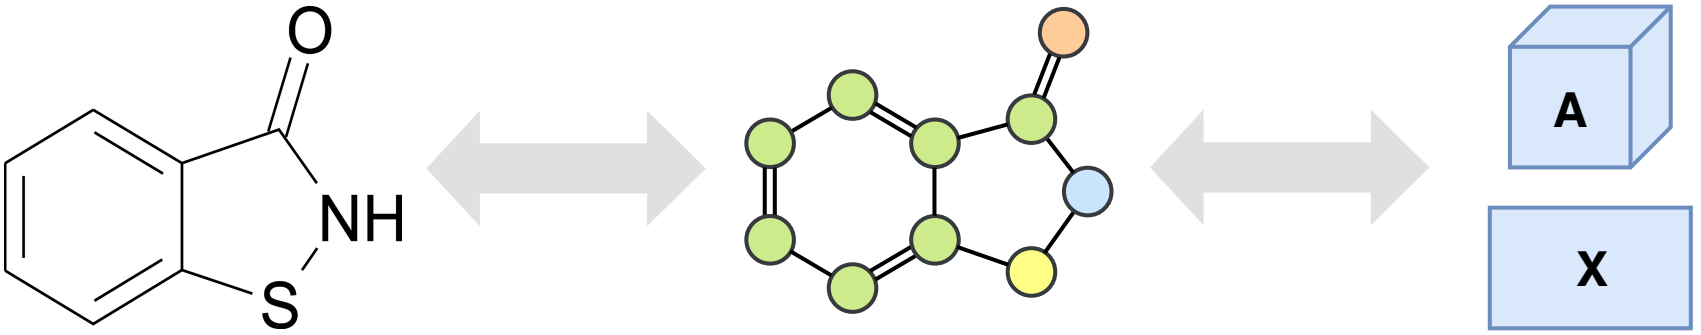

Figure: The **molecule** (a) is represented as an **labeled graph** (b) which can be encoded into an **adjacency tensor A** and an **annotation matrix X**.

Vectorial representation of graphs II

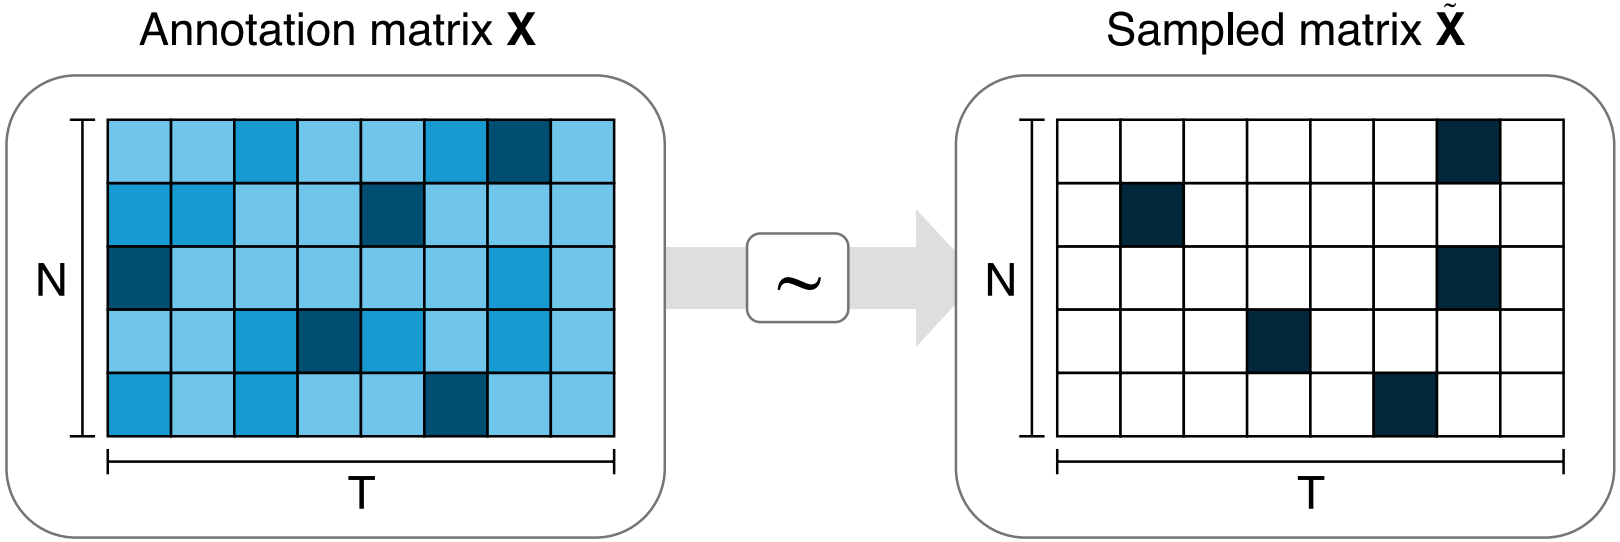

Vectorial representation of graphs III

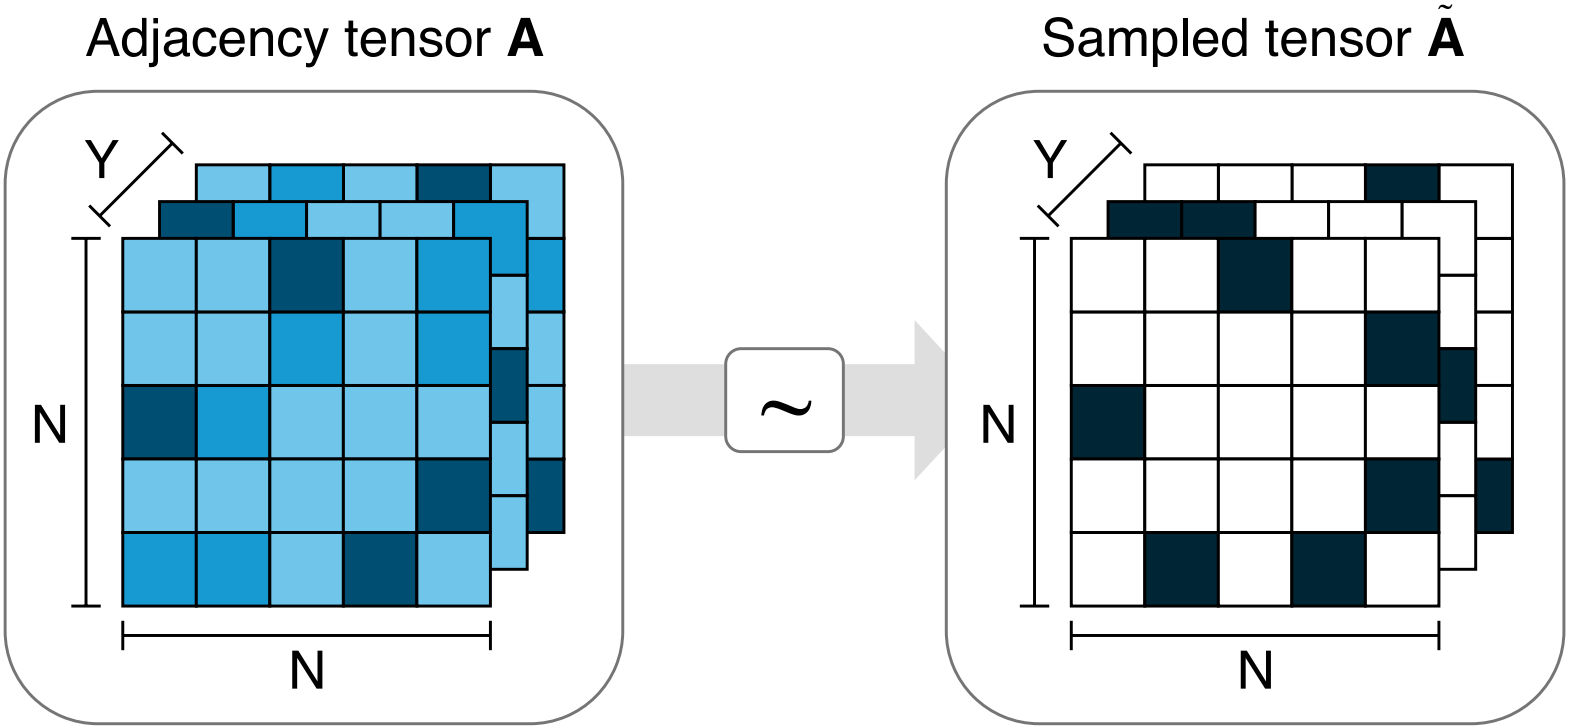

Graph Convolutional Networks I

Technically:

$$(f * g)(t) = \int_{\mathbb{R}^n} f(\tau) g(t - \tau) d\tau$$

but in practice we do **discrete** convolutions:

$$(f * g)[n] = \sum_{m \in \mathbb{S}} f(m) g(n - m)$$

Graph Convolutional Networks II

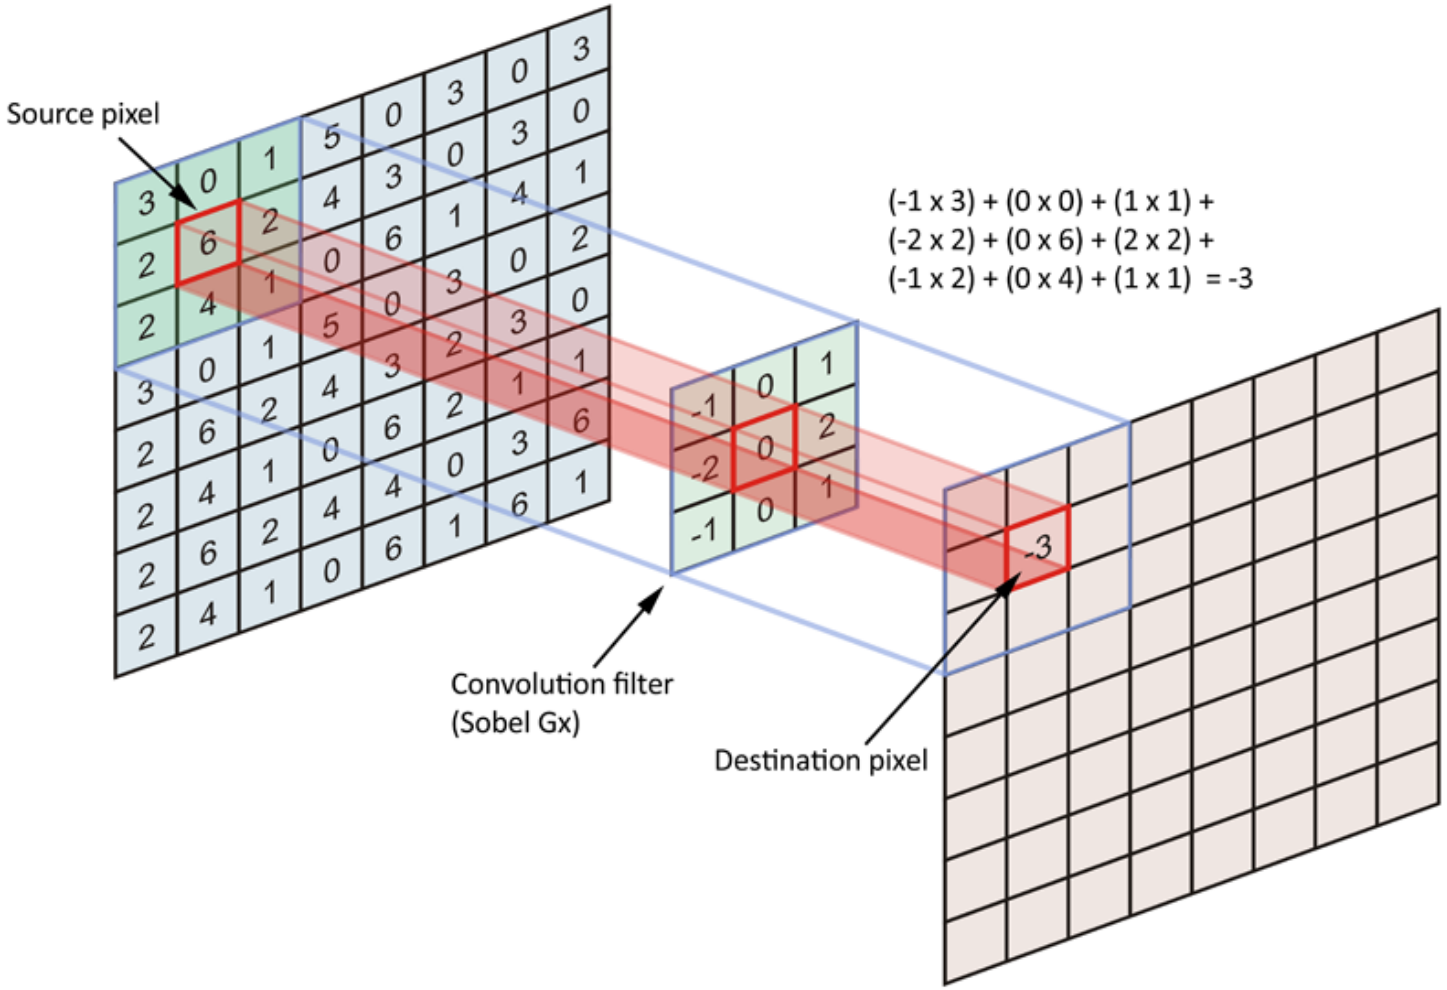

Figure: Graph convolution on an image.

Graph Convolutional Networks III

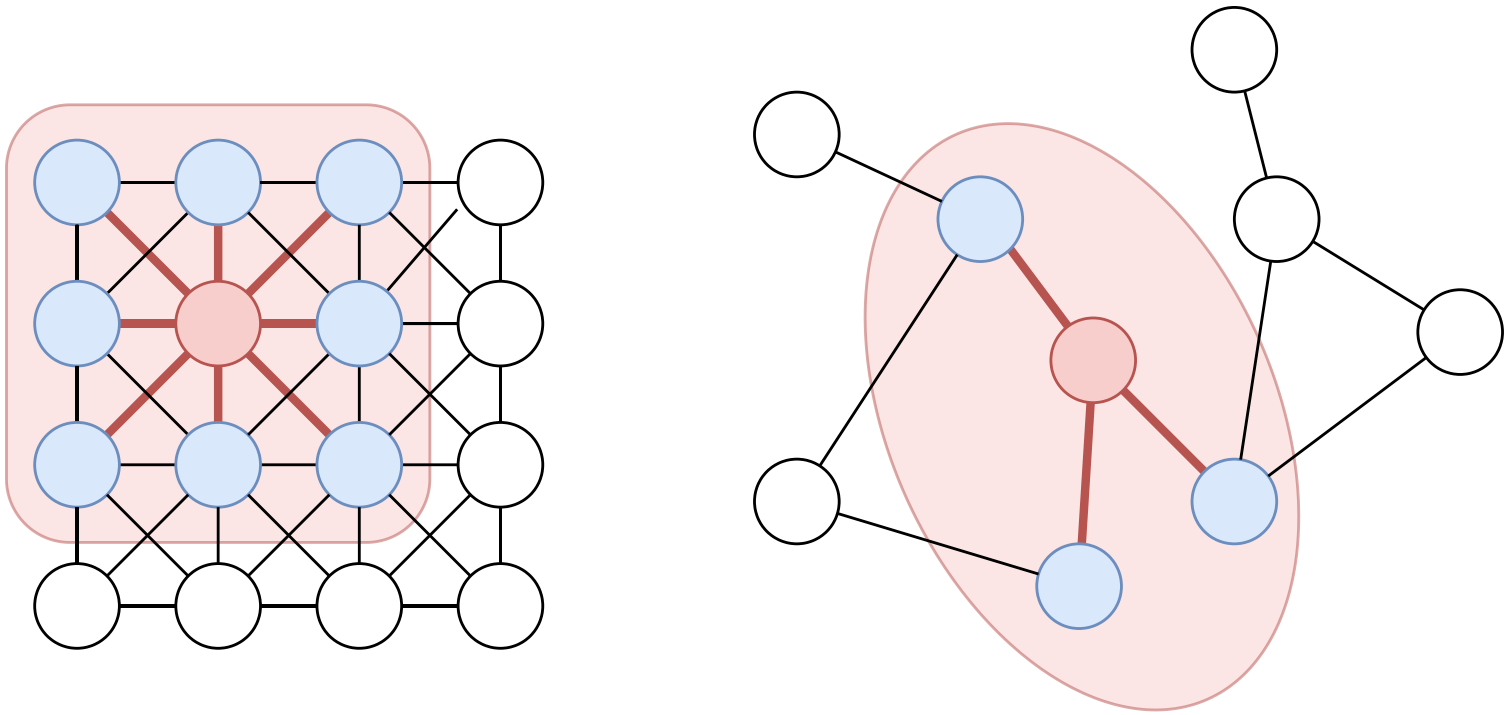

- On images the **topology is regular** and neighbours are **pixels**
- On graphs the **topology is arbitrary** and neighbours are **nodes**

Graph Convolutional Networks in practice

Edge-type-conditioned convolutions based on Relational-GCN [Schlichtkrull et al., 2018]:

$$\mathbf{h}_i^{(\ell+1)} = \tanh\left(f_s^{(\ell)}(\mathbf{h}_i^{(\ell)}, \mathbf{x}_i) + \sum_{j=1}^{N} \sum_{y=1}^{|\mathcal{X}_\mathcal{E}|} \frac{1}{|\mathcal{N}_i|} \mathbf{A}_{ijy} \, f_y^{(\ell)}(\mathbf{h}_j^{(\ell)}, \mathbf{x}_j) \right) ,$$

and following [Li et al., 2016], we define a graph level representation vector as

$$\mathbf{h}_\mathcal{G} = \tanh\left(\sum_{i=1}^{N} \sigma\left(f_g\left(\mathbf{h}_i^{(L)}, \mathbf{x}_i \right) \right) \odot \tanh\left(f_r\left(\mathbf{h}_i^{(L)}, \mathbf{x}_i \right) \right) \right) .$$

Molecular graph VAE

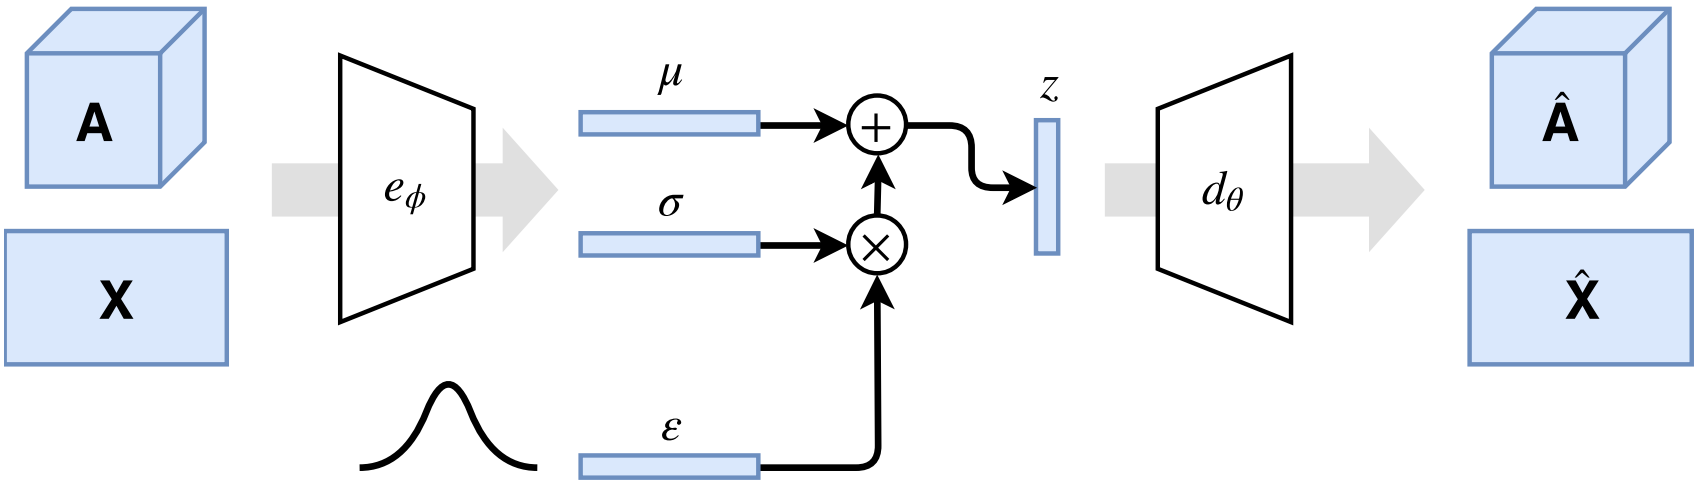

The reconstruction loss is a sum of two **categorical cross entropy** losses.

Molecular graph GAN

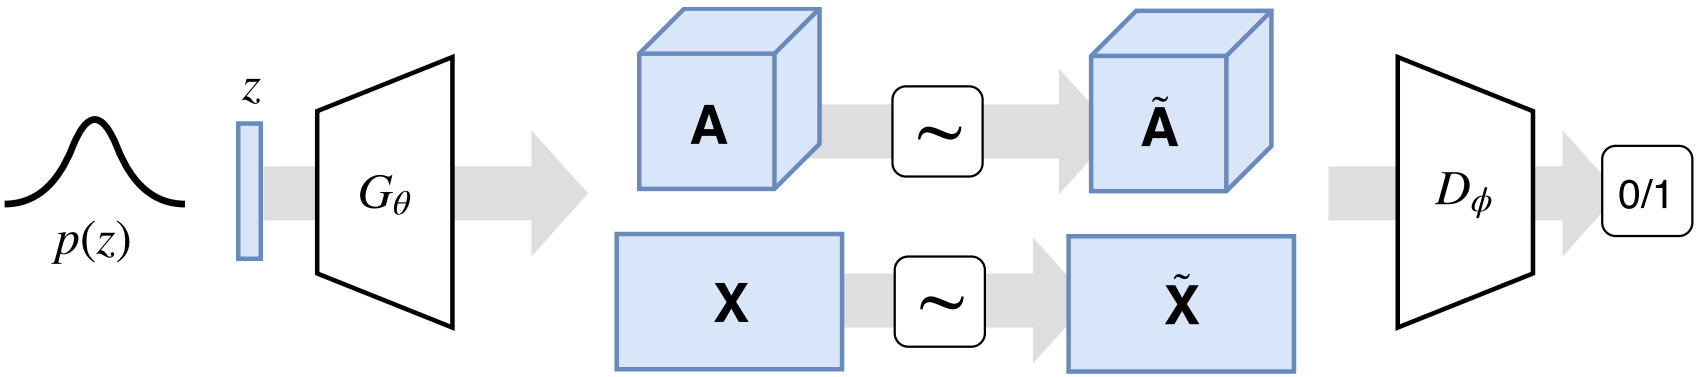

From generator to discriminator with **differentiable sampling** [Jang et al., 2017].

Molecular graph GAN with RL

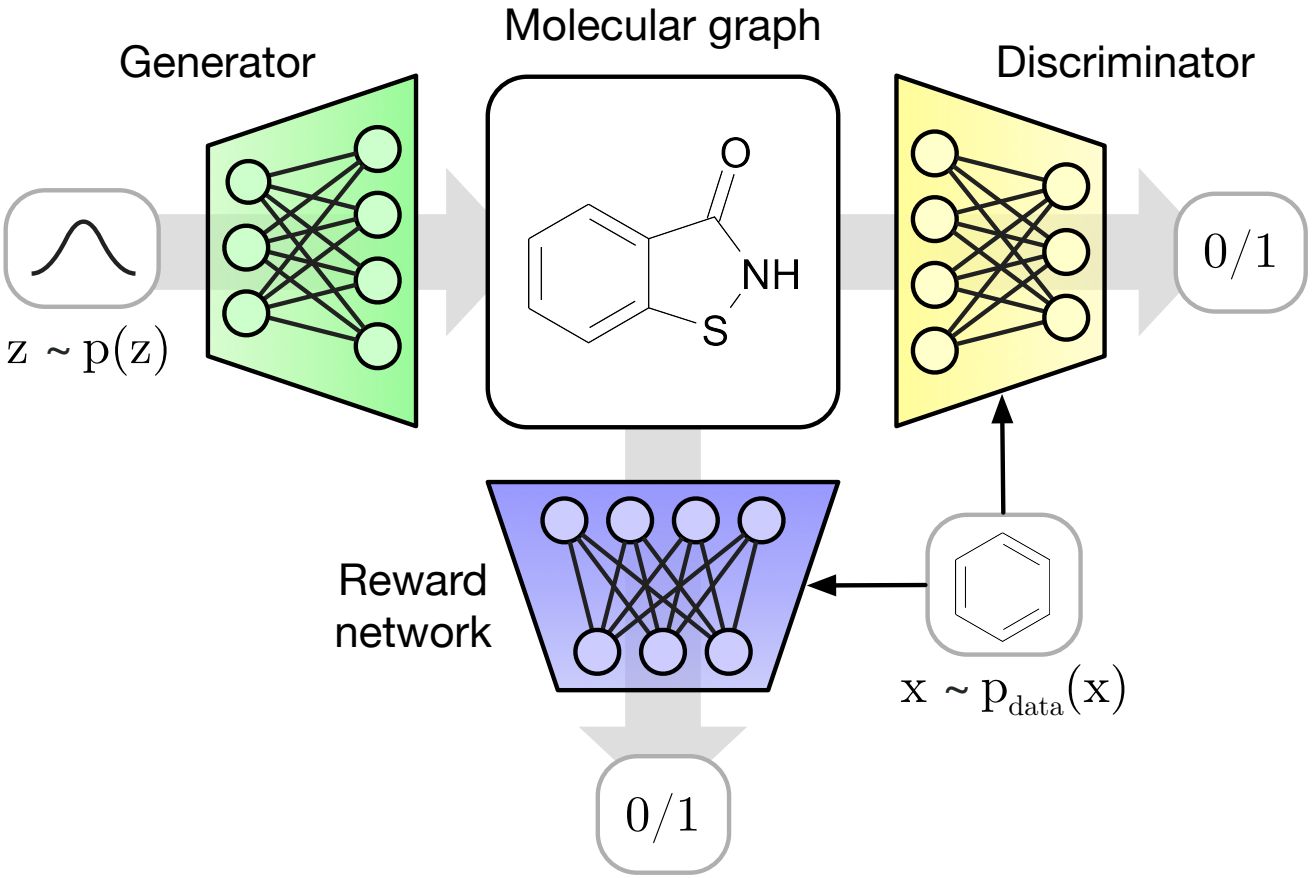

Architecture from our previous work **MolGAN** [De Cao and Kipf, 2018]

Experiments

Experiments

Which questions we would like to answer?

- **likelihood-based vs. likelihood-free** (VAEs vs. GANs)
- the **effect of RL** towards chemical objectives
- is generating a **graph better than a SMILES** representation?

Experiments on QM9 [Ramakrishnan et al., 2014].

VAE advantages and disadvantages

VAEs train an **encoder!**

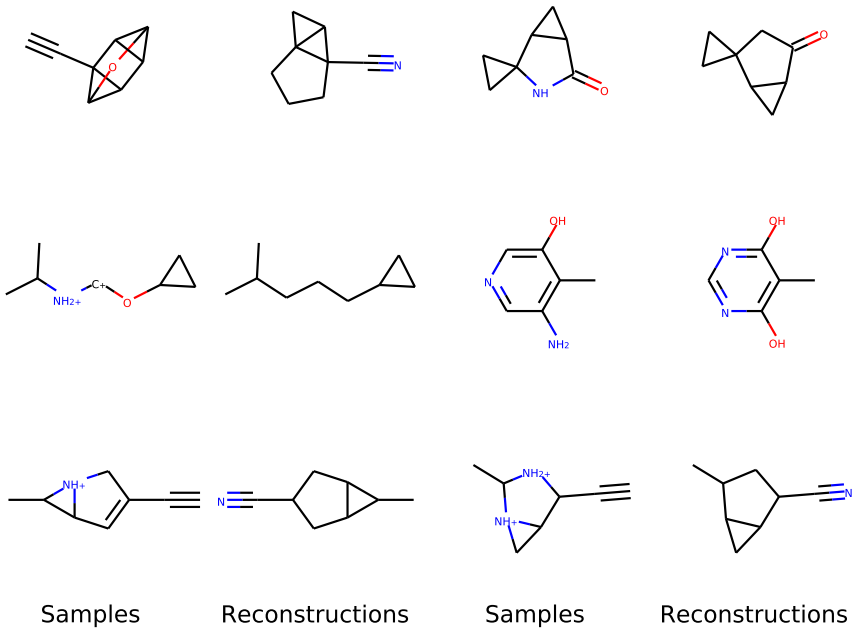

- **VAE objective**: reconstruction loss and divergence
- **RL objective**: sampled molecules should maximize a score

There is a mismatch between these two!

Synthetic accessibility score (SAS) distributions I

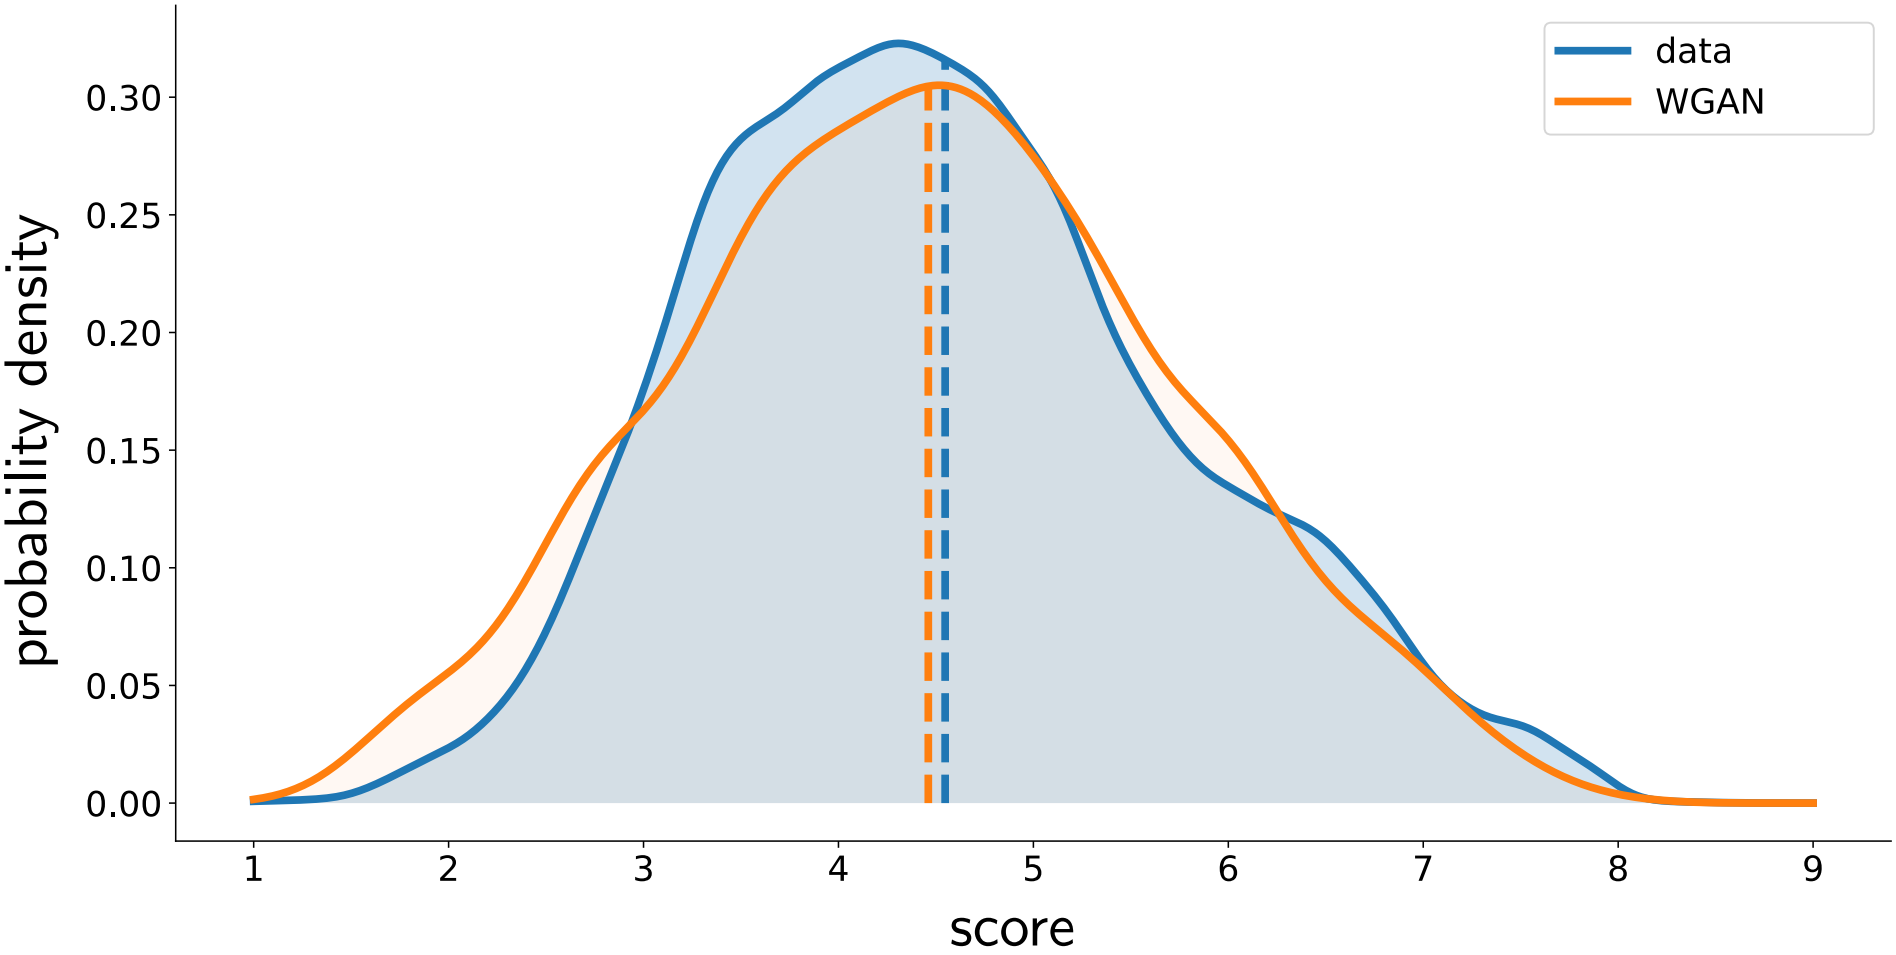

Figure: WGAN matches the data distribution of the synthetic accessibility score [Ertl and Schuffenhauer, 2009].

Synthetic accessibility score (SAS) distributions II

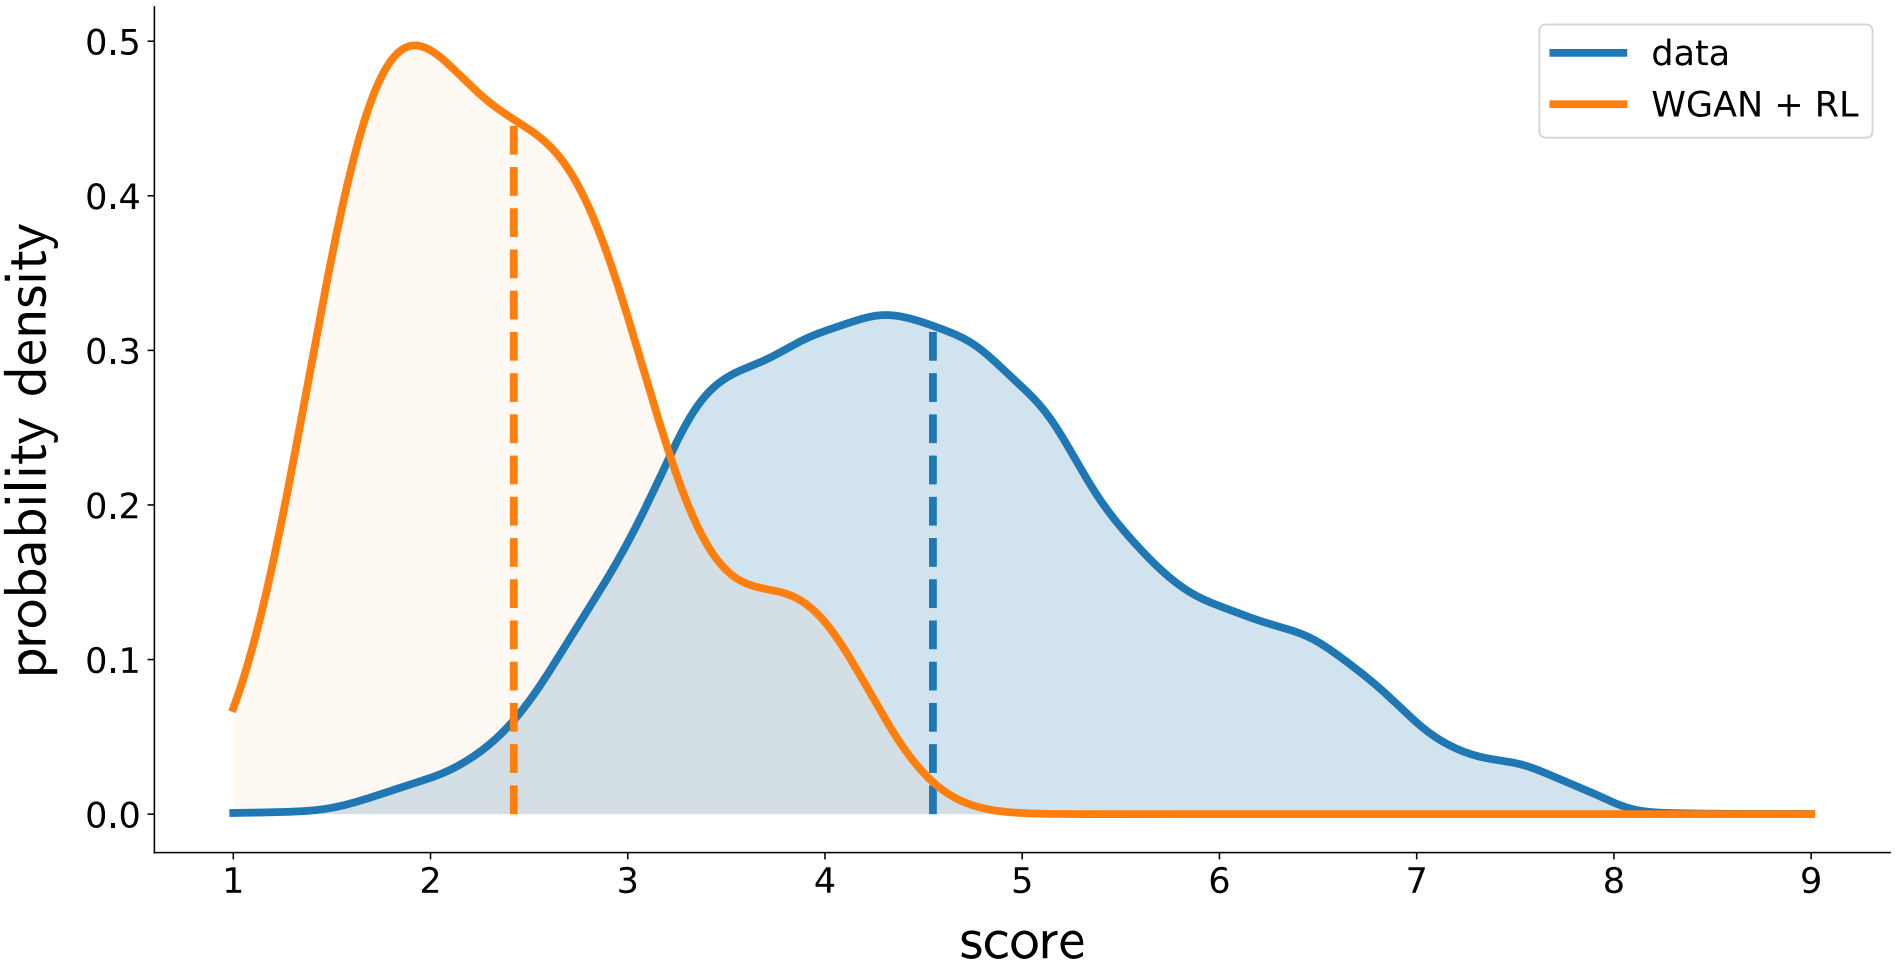

Figure: WGAN in combination with RL push the distribution of the synthetic accessibility score (SAS) to be as low as possible.

Trade-off between WGAN and RL

Method	validity	uniqueness	QED*
$\lambda = 0.0$ (full RL)	**100.00**	3.16	**0.61**
$\lambda = 0.125$	**100.00**	7.21	**0.61**
$\lambda = 0.25$	99.80	10.16	**0.61**
$\lambda = 0.375$	99.90	11.11	0.60
$\lambda = 0.5$	99.40	31.29	0.56
$\lambda = 0.625$	97.20	49.69	0.51
$\lambda = 0.75$	93.70	64.35	0.51
$\lambda = 0.875$	89.40	**69.69**	0.50
$\lambda = 1.0$ (no RL)	90.10	63.91	0.50

Table: WGAN and RL objectives trade-off. *QED is the quantitative estimate of drug-likeness [Bickerton et al., 2012].

Evolution of the QED during training

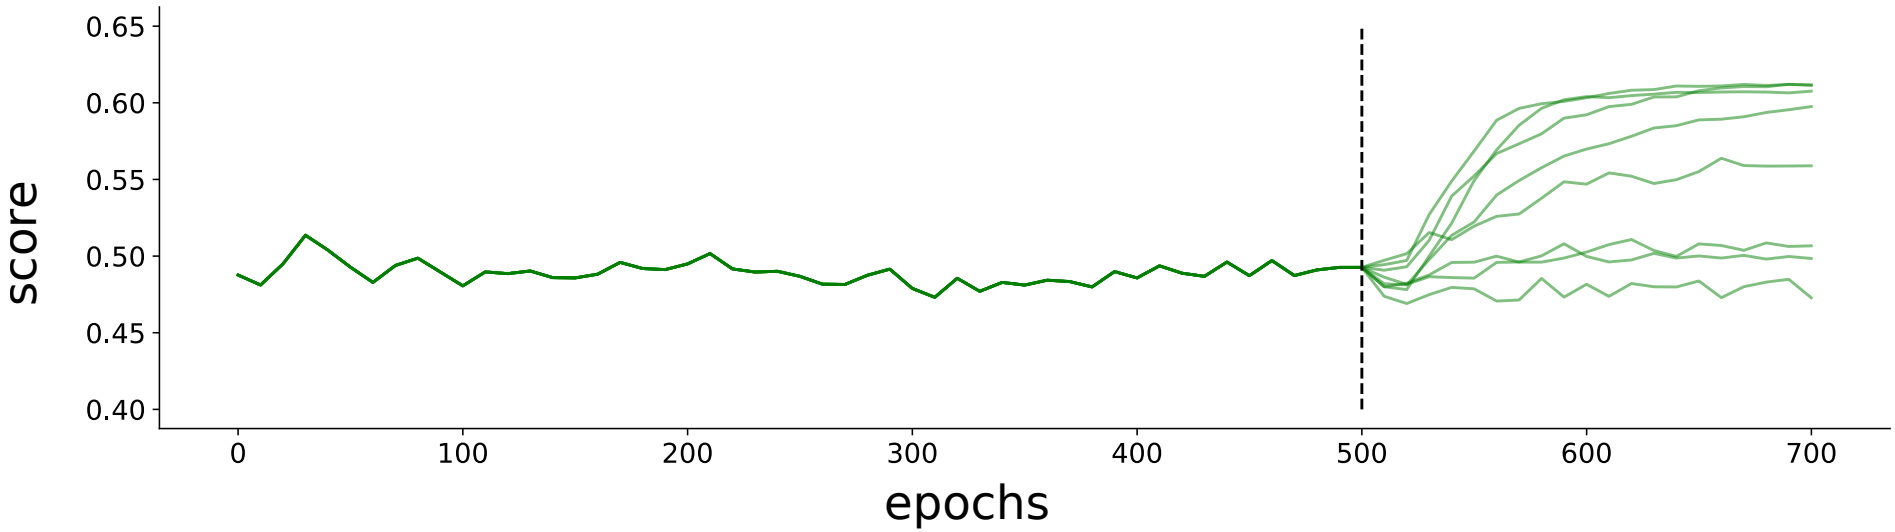

Figure: Evolution of the QED during training with different λ values.

Comparison with VAE based methods

Method	validity	uniqueness	novelty
CharacterVAE	10.3	67.5	90.0
GrammarVAE	60.2	9.3	80.9
GraphVAE	55.7	76.0	61.6
GraphVAE/imp	56.2	42.0	75.8
GraphVAE NoGM	81.0	24.1	61.0
Our VAE	61.5	**97.6**	69.1
Our VAE with RL	**89.1**	11.1	**92.3**
Our WGAN	**89.2**	26.5	55.7
Our WGAN with RL	**99.6**	14.5	**97.7**

Table: Baseline results from [Gómez-Bombarelli et al., 2016, Kusner et al., 2017, Simonovsky and Komodakis, 2018]

Comparison with a GAN based method

Method	validity	SAS	time (h)
ORGAN	96.5	0.83	8.7
OR(W)GAN	97.6	0.75	9.6
Naive RL	97.7	0.83	10.6
Our VAE with RL	89.6	0.71	**0.09**
Our VAE with RL (full QM9)	94.0	**0.86**	**2.2**
Our WGAN with RL	**100.0**	0.70	**0.15**
Our WGAN with RL (full QM9)	**99.8**	**0.92**	**3.3**

Table: Baseline results from ORGAN [Guimaraes et al., 2017].

Best QED samples

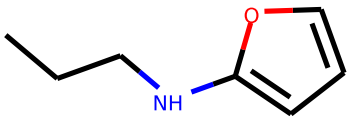

QED=0.67

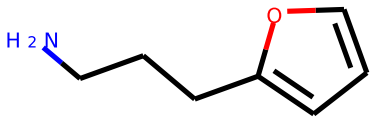

QED=0.66

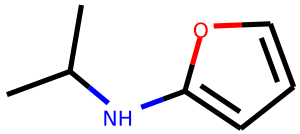

QED=0.65

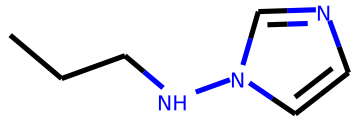

QED=0.65

(a) Best VAE with RL molecules.

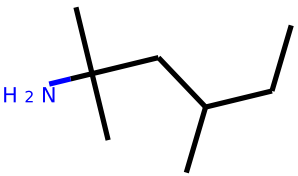

QED=0.62

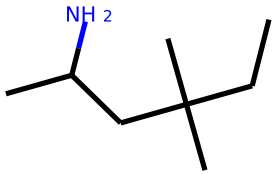

QED=0.62

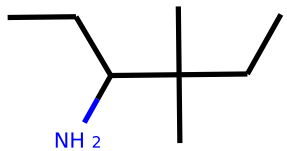

QED=0.62

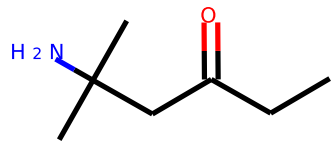

QED=0.62

(b) Best WGAN with RL molecules.

Figure: Top four molecules with QED scores.

Conclusion and future work

Conclusion

Considering experimental, we identify these further **contributions**:

- recurrent SMILES generation is more **computational expensive**
- likelihood-based models are **difficult to be optimized** with RL

... but keeping in mind and those **limitations**:

- we experimented using compounds of **at most 9 atoms**
- models are susceptible to **mode collapse**

Future work

We identify four principal directions for future work:

- address **mode collapse** [Srivastava et al., 2017]
- combine **variational** approaches with **adversarial** learning to benefit from both approaches [Mescheder et al., 2017, Rosca et al., 2017]
- train our models on **ChEMBL** [Gaulton et al., 2011]
- more **realistic reward** functions [Li et al., 2018]

References I


Arjovsky, M., Chintala, S., and Bottou, L. (2017).
Wasserstein generative adversarial networks.
In *International Conference on Machine Learning*, pages 214–223.


Bickerton, G. R., Paolini, G. V., Besnard, J., Muresan, S., and Hopkins, A. L. (2012).
Quantifying the chemical beauty of drugs.
Nature chemistry, 4(2):90.


De Cao, N. and Kipf, T. (2018).
MolGAN: An implicit generative model for small molecular graphs.
ICML 2018 workshop on Theoretical Foundations and Applications of Deep Generative Models.

References II


Ertl, P. and Schuffenhauer, A. (2009).
Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions.
Journal of cheminformatics, 1(1):8.


Gaulton, A., Bellis, L. J., Bento, A. P., Chambers, J., Davies, M., Hersey, A., Light, Y., McGlinchey, S., Michalovich, D., Al-Lazikani, B., et al. (2011).
ChEMBL: a large-scale bioactivity database for drug discovery.
Nucleic acids research, 40(D1):D1100–D1107.

References III

Gómez-Bombarelli, R., Wei, J. N., Duvenaud, D., Hernández-Lobato, J. M., Sánchez-Lengeling, B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T. D., Adams, R. P., and Aspuru-Guzik, A. (2016).
Automatic chemical design using a data-driven continuous representation of molecules.
ACS Central Science, 4(2):268–276.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014).
Generative adversarial nets.
In *Advances in neural information processing systems*, pages 2672–2680.

References IV


Guimaraes, G. L., Sanchez-Lengeling, B., Farias, P. L. C., and Aspuru-Guzik, A. (2017).
Objective-reinforced generative adversarial networks (ORGAN) for sequence generation models.
arXiv preprint arXiv:1705.10843.


Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. (2017).
Improved training of wasserstein gans.
In *Advances in Neural Information Processing Systems*, pages 5769–5779.


Hafner, D. (2018).
Building variational auto-encoders in tensorflow.
Blog post.

References V


Jang, E., Gu, S., and Poole, B. (2017).
Categorical reparameterization with gumbel-softmax.
International Conference on Learning Representations.


Kingma, D. P. and Welling, M. (2013).
Auto-encoding variational bayes.
International Conference on Learning Representations.


Kusner, M. J., Paige, B., and Hernández-Lobato, J. M. (2017).
Grammar variational autoencoder.
In *International Conference on Machine Learning*, pages 1945–1954.


Li, Y., Tarlow, D., Brockschmidt, M., and Zemel, R. (2016).
Gated graph sequence neural networks.
International Conference on Learning Representations.

References VI


Li, Y., Zhang, L., and Liu, Z. (2018).
Multi-objective de novo drug design with conditional graph generative model.
arXiv preprint arXiv:1801.07299.


Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. (2016).
Continuous control with deep reinforcement learning.
International Conference on Learning Representations.


Mescheder, L., Nowozin, S., and Geiger, A. (2017).
Adversarial variational bayes: Unifying variational autoencoders and generative adversarial networks.
International Conference on Machine Learning.

References VII


Nowozin, S., Cseke, B., and Tomioka, R. (2016).
f-gan: Training generative neural samplers using variational divergence minimization.
In *Advances in Neural Information Processing Systems*, pages 271–279.


Ramakrishnan, R., Dral, P. O., Rupp, M., and Von Lilienfeld, O. A. (2014).
Quantum chemistry structures and properties of 134 kilo molecules.
Scientific data, 1:140022.


Rosca, M., Lakshminarayanan, B., Warde-Farley, D., and Mohamed, S. (2017).
Variational approaches for auto-encoding generative adversarial networks.
arXiv preprint arXiv:1706.04987.

References VIII


Schlichtkrull, M., Kipf, T. N., Bloem, P., van den Berg, R., Titov, I., and Welling, M. (2018).
Modeling relational data with graph convolutional networks.
In *European Semantic Web Conference*, pages 593–607. Springer.


Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. (2014).
Deterministic policy gradient algorithms.
In *International Conference on Machine Learning*.


Simonovsky, M. and Komodakis, N. (2018).
GraphVAE: Towards generation of small graphs using variational autoencoders.
arXiv preprint arXiv:1802.03480.

References IX


Srivastava, A., Valkoz, L., Russell, C., Gutmann, M. U., and Sutton, C. (2017).
VEEGAN: Reducing mode collapse in GANs using implicit variational learning.
In *Advances in Neural Information Processing Systems*, pages 3308–3318.