Unified Guidance for Geometry-Conditioned Molecular Generation

Sirine Ayadi[*1,2] **Leon Hetzel**[*1,2,3] **Johanna Sommer**[*1,2]

Fabian Theis[1,2,3] **Stephan Günnemann**[1,2]

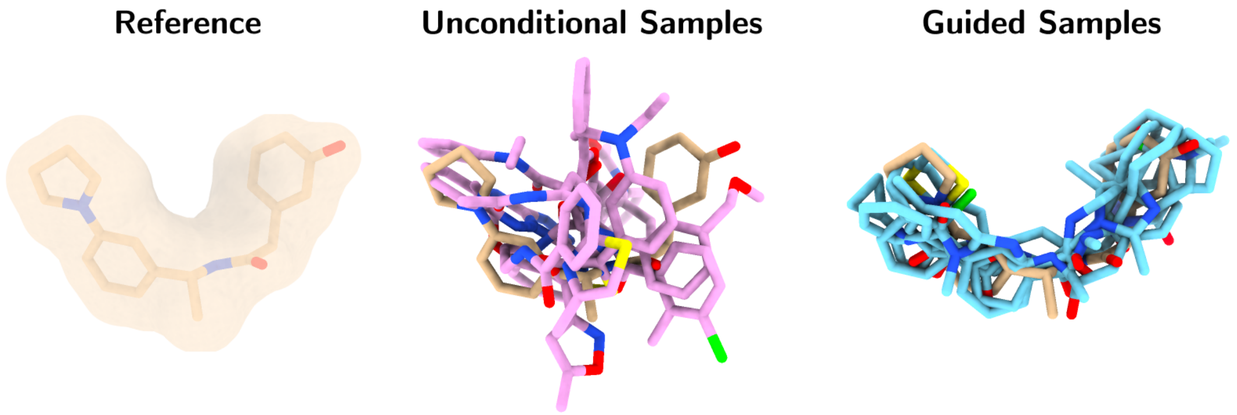

https://www.cs.cit.tum.de/daml/uniguide/

Diffusion models for drug design

Equivariant (unconditional) diffusion model for 3D molecule generation

- coordinates: $\mathbf{x}$
- atom types: $\mathbf{h}$
- configuration: $\mathbf{z}$

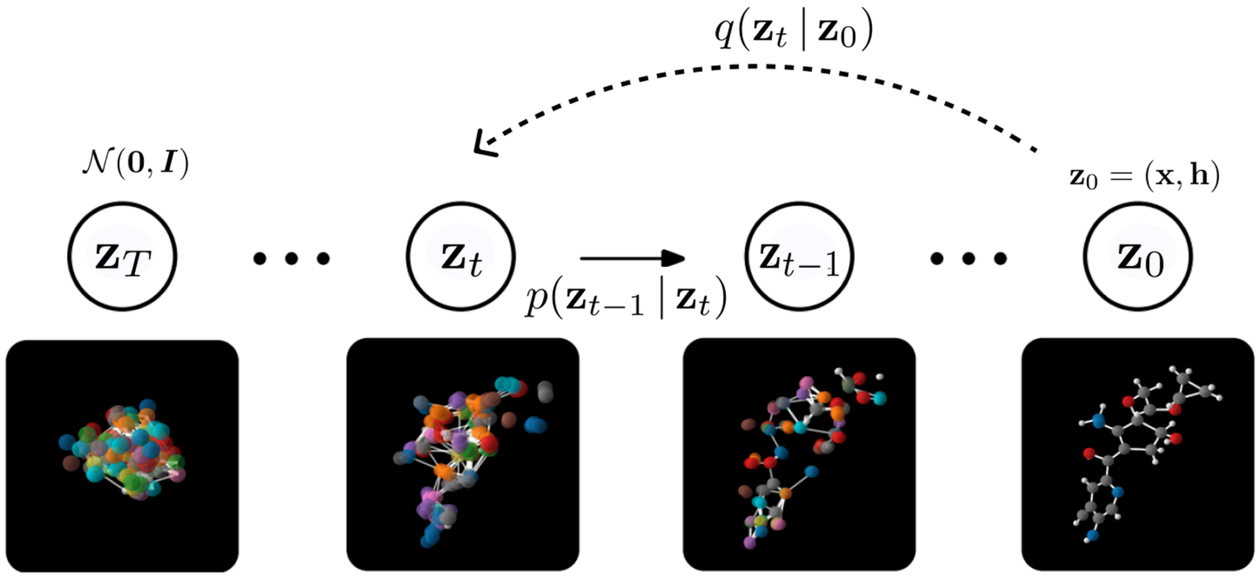

Hoogeboom, E., Satorras, V. G., Vignac, C., & Welling, M. (2022, June). Equivariant diffusion for molecule generation in 3D.

Diffusion models for drug design

Extending unconditional diffusion models for molecules to various conditioning tasks

Structure-Based Drug Design

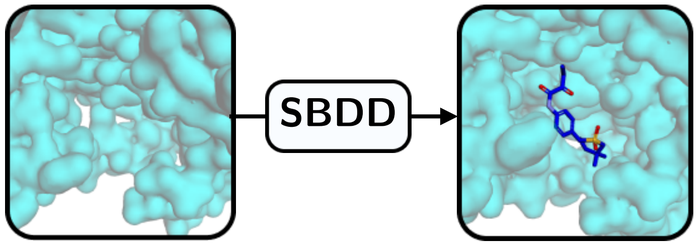

Protein

Generate a molecule that **binds to the protein**

Fragment-Based Drug Design

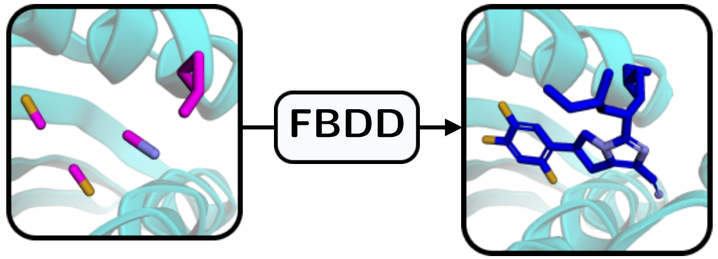

Protein & Fragments

Generate a molecule that **contains the fragments & binds to the protein**

Ligand-Based Drug Design

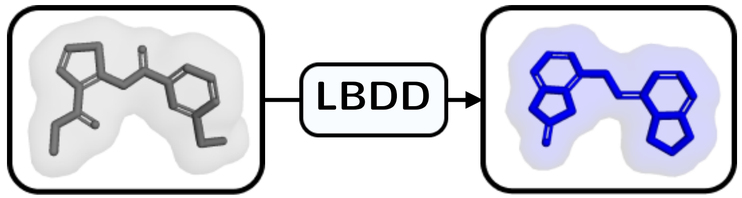

Molecular Surface $\mathcal{M}_{\text{ref}}$

Generate a novel molecule that **satisfies the shape**

Controlling diffusion models

Previous diffusion-based methods are limited to a single drug design task

Previous Work

Conditional Diffusion Models

- ✘ Require conditional training per task
- ✘ Require task-specific data
- ✘ Extra training might be required to find a suitable condition representation

Guiding Diffusion Models

- ✓ Does not require extra training
- ✘ Requires external model for guidance
- ✘ External model has to be trained on noisy states of the diffusion process

Self-guiding diffusion models

Incorporating the condition c to the unconditional model $\epsilon_\theta(\mathbf{z}_t, t)$ during inference

Bayes' rule

$$p_\theta(\mathbf{z}_t \mid \boldsymbol{c}) = \frac{p_\theta(\mathbf{z}_t)\, p_\theta(\boldsymbol{c} \mid \mathbf{z}_t)}{p(\boldsymbol{c})}$$

Score function

$$\nabla_{\mathbf{z}_t} \log p_\theta(\mathbf{z}_t \mid \boldsymbol{c}) = \underbrace{\nabla_{\mathbf{z}_t} \log p_\theta(\mathbf{z}_t)}_{-\frac{\boldsymbol{\epsilon}_\theta(\mathbf{z}_t, t)}{\sqrt{1-\bar{\alpha}_t}}} + S\, \nabla_{\mathbf{z}_t} \log p_\theta(\boldsymbol{c} \mid \mathbf{z}_t)$$

Guided diffusion model

$$\hat{\boldsymbol{\epsilon}}_\theta(\mathbf{z}_t, t, \boldsymbol{c}) = \boldsymbol{\epsilon}_\theta(\mathbf{z}_t, t) - \sqrt{1-\bar{\alpha}_t}\, S\, \nabla_{\mathbf{z}_t} \log p_\theta(\boldsymbol{c} \mid \mathbf{z}_t)$$

Self-guiding diffusion models

Incorporating the condition c to the unconditional model $\epsilon_\theta(\mathbf{z}_t, t)$ during inference

Assumption

$$p_\theta(\boldsymbol{c} \,|\, \mathbf{z}_t) = \mathcal{N}\big(\boldsymbol{c} \,|\, \boldsymbol{f}_\theta(\mathbf{z}_t, t), \boldsymbol{I}\big)$$

Use clean data point approximation

$$\boldsymbol{f}_\theta(\mathbf{z}_t, t) = \frac{\mathbf{z}_t - \sqrt{1-\bar{\alpha}_t}\,\boldsymbol{\epsilon}_\theta(\mathbf{z}_t, t)}{\sqrt{\bar{\alpha}_t}} =: \hat{\mathbf{z}}_0$$

Guiding signal reduces to squared error

$$\nabla_{\mathbf{z}_t} \log p_\theta(\boldsymbol{c} \,|\, \mathbf{z}_t) = -\frac{1}{2} \nabla_{\mathbf{z}_t} \Big\| \qquad \Big\|^2$$

The condition must lie in the configuration space

Self-guided diffusion model

$$\hat{\boldsymbol{\epsilon}}_\theta(\mathbf{z}_t, t, \boldsymbol{c}) = \boldsymbol{\epsilon}_\theta(\mathbf{z}_t, t) + \lambda_t\, S\, \nabla_{\mathbf{z}_t} \big\| \hat{\mathbf{z}}_0 - \boldsymbol{c} \big\|_2^2$$

UniGuide: A unified guidance framework

The condition map transforms arbitrary source conditions to suitable guidance targets

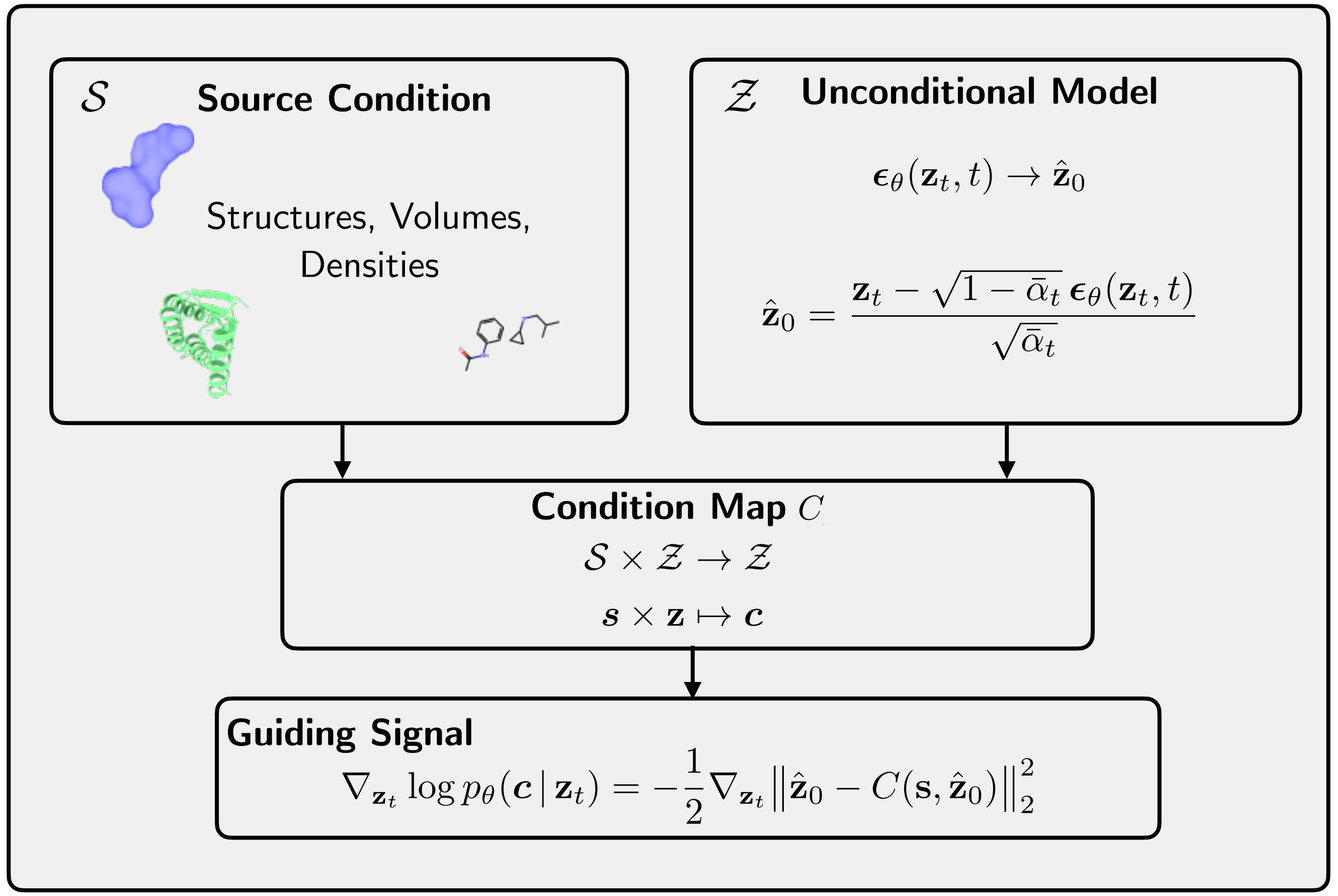

UniGuide

- ✓ No re-training required
- ✓ Does not break equivariance
- ✓ Works in a low-data regime
- ✓ Readily adaptable to various drug design tasks
- ✓ No external models

UniGuide: A unified guidance framework

The condition map transforms arbitrary source conditions to suitable guidance targets

Surface Condition Map (LBDD)

Sample points $\mathbf{y}$ from the surface

$$C_{\partial V} : \mathbb{R}^{K\times 3} \times \mathbb{R}^{N\times 3} \to \mathbb{R}^{N\times 3}$$
$$\mathbf{y} \times \hat{\mathbf{x}}_0 \mapsto \boldsymbol{c}_{\mathbf{x}} \quad ,$$

Compute projection on surface from closest neighbours

$$\bar{\boldsymbol{y}}_i = \frac{1}{k} \sum_{j \in \mathcal{N}_{\hat{\boldsymbol{x}}_i}} R_{\hat{\mathbf{x}}_0} \boldsymbol{y}_j \quad ,$$

$$\text{with} \quad \mathcal{N}_{\hat{\boldsymbol{x}}_i} = \arg \min_{I \subset \{1,\dots,K\}, |I|=k} \sum_{j \in I} \left\| R_{\hat{\mathbf{x}}_0} \boldsymbol{y}_j - \hat{\boldsymbol{x}}_i \right\|_2$$

Compute the (point-wise) target condition

$$\boldsymbol{c}_{\mathbf{x},i} = \begin{cases} \bar{\boldsymbol{y}}_i + \frac{\alpha}{d}(\bar{\boldsymbol{y}}_i - \hat{\boldsymbol{x}}_i) \,, & \text{if } \hat{\boldsymbol{x}}_i \text{ outside } V \\ \bar{\boldsymbol{y}}_i - \frac{\alpha}{d}(\bar{\boldsymbol{y}}_i - \hat{\boldsymbol{x}}_i) \,, & \text{if } \hat{\boldsymbol{x}}_i \text{ inside } V \wedge d < \alpha \\ \hat{\boldsymbol{x}}_i \,, & \text{otherwise} \,, \end{cases}$$

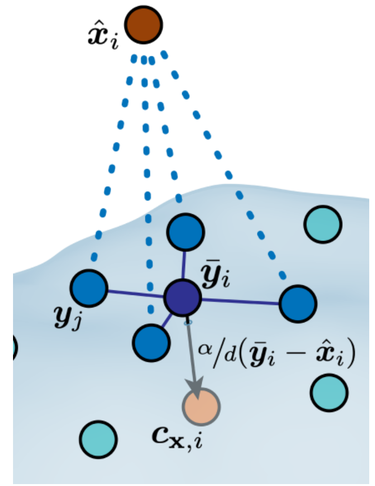

Figure: Surface Condition Map $C_{\partial V}$: For every atom $\boldsymbol{x}_i$, the closest surface points $\boldsymbol{y}_j$ are computed.

UniGuide: A unified guidance framework

The condition map transforms arbitrary source conditions to suitable guidance targets

Surface Condition Map (LBDD)

Sample points $\mathbf{y}$ from the surface

$$C_{\partial V} : \mathbb{R}^{K\times 3} \times \mathbb{R}^{N\times 3} \to \mathbb{R}^{N\times 3}$$
$$\mathbf{y} \times \hat{\mathbf{x}}_0 \mapsto \boldsymbol{c}_{\mathbf{x}} \quad ,$$

Compute projection on surface from closest neighbours

$$\bar{\boldsymbol{y}}_i = \frac{1}{k} \sum_{j \in \mathcal{N}_{\hat{\boldsymbol{x}}_i}} R_{\hat{\mathbf{x}}_0} \boldsymbol{y}_j \quad ,$$

$$\text{with} \quad \mathcal{N}_{\hat{\boldsymbol{x}}_i} = \arg \min_{I \subset \{1,\dots,K\}, |I|=k} \sum_{j \in I} \left\| R_{\hat{\mathbf{x}}_0} \boldsymbol{y}_j - \hat{\boldsymbol{x}}_i \right\|_2$$

Compute the (point-wise) target condition

$$\boldsymbol{c}_{\mathbf{x},i} = \begin{cases} \bar{\boldsymbol{y}}_i + \frac{\alpha}{d}(\bar{\boldsymbol{y}}_i - \hat{\boldsymbol{x}}_i) \,, & \text{if } \hat{\boldsymbol{x}}_i \text{ outside } V \\ \bar{\boldsymbol{y}}_i - \frac{\alpha}{d}(\bar{\boldsymbol{y}}_i - \hat{\boldsymbol{x}}_i) \,, & \text{if } \hat{\boldsymbol{x}}_i \text{ inside } V \wedge d < \alpha \\ \hat{\boldsymbol{x}}_i \,, & \text{otherwise} \,, \end{cases}$$

Special Case $\mathcal{S} = \mathcal{Z}$

Ensure equivariant update

$$C_{\mathcal{Z}} : \quad \mathbb{R}^{m\times(3+d)} \times \mathbb{R}^{m\times(3+d)} \to \mathbb{R}^{m\times(3+d)}$$
$$\tilde{\mathbf{z}} \times \hat{\mathbf{z}}_0^{\mathcal{A}} \mapsto T_{\hat{\mathbf{z}}_0^{\mathcal{A}}} \tilde{\mathbf{z}}$$

SBDD

$$C_{\mathcal{Z}}(\tilde{\mathbf{z}}^{\mathcal{P}}, \hat{\mathbf{z}}_t^{\mathcal{P}}) = T_{\hat{\mathbf{z}}_t^{\mathcal{P}}} \tilde{\mathbf{z}}^{\mathcal{P}}$$

FBDD

$$C_{\mathcal{Z}}(\tilde{\mathbf{z}}^{\mathcal{A}}, \hat{\mathbf{z}}_t^{\mathcal{A}}) = T_{\hat{\mathbf{z}}_t^{\mathcal{A}}} \tilde{\mathbf{z}}^{\mathcal{A}}$$

$$\mathcal{A} = \mathcal{P} \cup \mathcal{F}$$

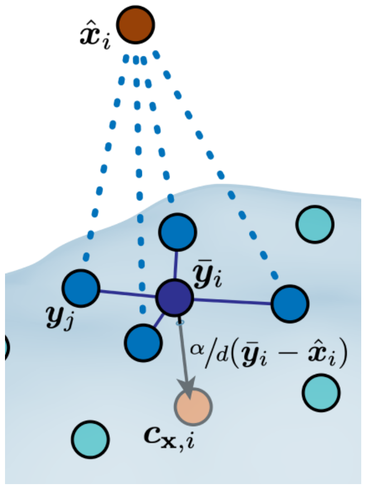

Figure: Surface Condition Map $C_{\partial V}$: For every atom $\boldsymbol{x}_i$, the closest surface points $\boldsymbol{y}_j$ are computed.

Results: Ligand-Based Drug Design

Guiding an off-the-shelf EDM model for LBDD shows superior performance, even to Virtual Screening.

	SimS (↑)	MaxSimS (↑)	SimG (↓)	**Ratio (↑)**
Virtual Screening	0.729	0.807	0.226	3.226
ShapeMol	0.677	0.797	**0.239**	2.834
ShapeMol+g	0.744	0.849	0.242	3.074
UniGuide (ShapeMol [U])	0.726	0.827	0.248	2.927
UniGuide (ShapeMol)	**0.760**	**0.857**	0.240	**3.167**
UniGuide (EDM)	0.749	0.860	0.212	3.536

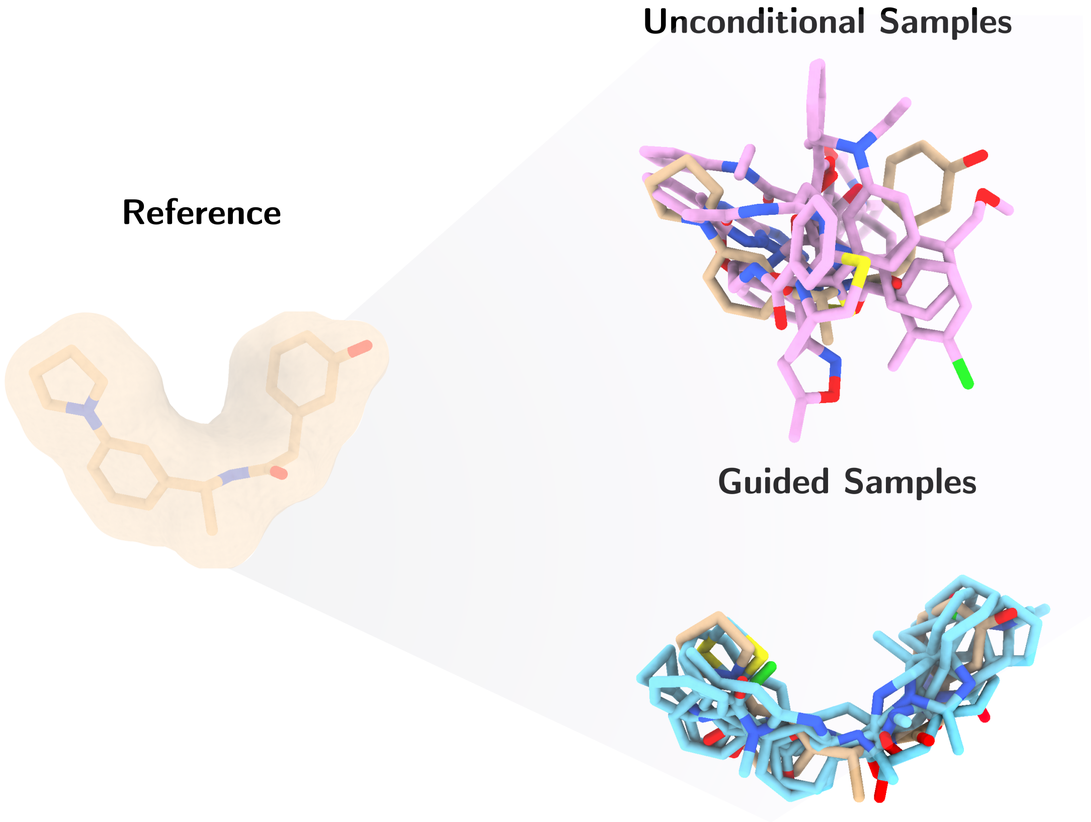

Results: Structure-Based Drug Design

UniGuide outperforms different diffusion-based conditioning approaches.

Quantitative Comparison on CrossDocked. We highlight in **bold** the best approach given the same backbone and underline the best approach overall.

	Vina Score (↓)	Vina Min (↓)	Vina Dock (↓)	QED (↑)
Test set	−6.362	−6.707	−7.450	0.48
DecompDiff	−4.750	−6.170	–	–
TargetDiff	−5.466	−6.643	−7.802	0.48
DiffSBDD-cond	−3.684	−4.670	−6.941	0.47
DiffSBDD	−4.097	−6.306	−7.889	**0.57**
UniGuide	**−5.103**	**−6.610**	**−7.921**	**0.57**

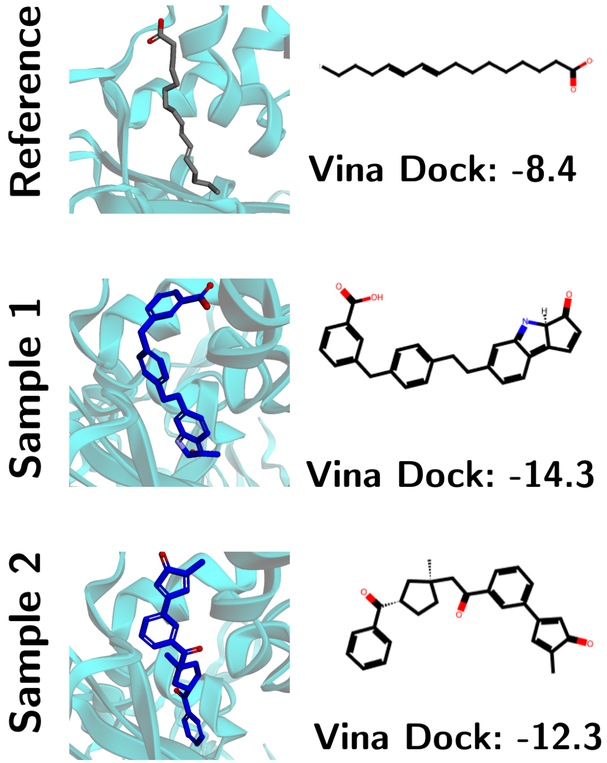

Results: Fragment-Based Drug Design

UniGuide is readily applicable to various FBDD settings in a unified fashion

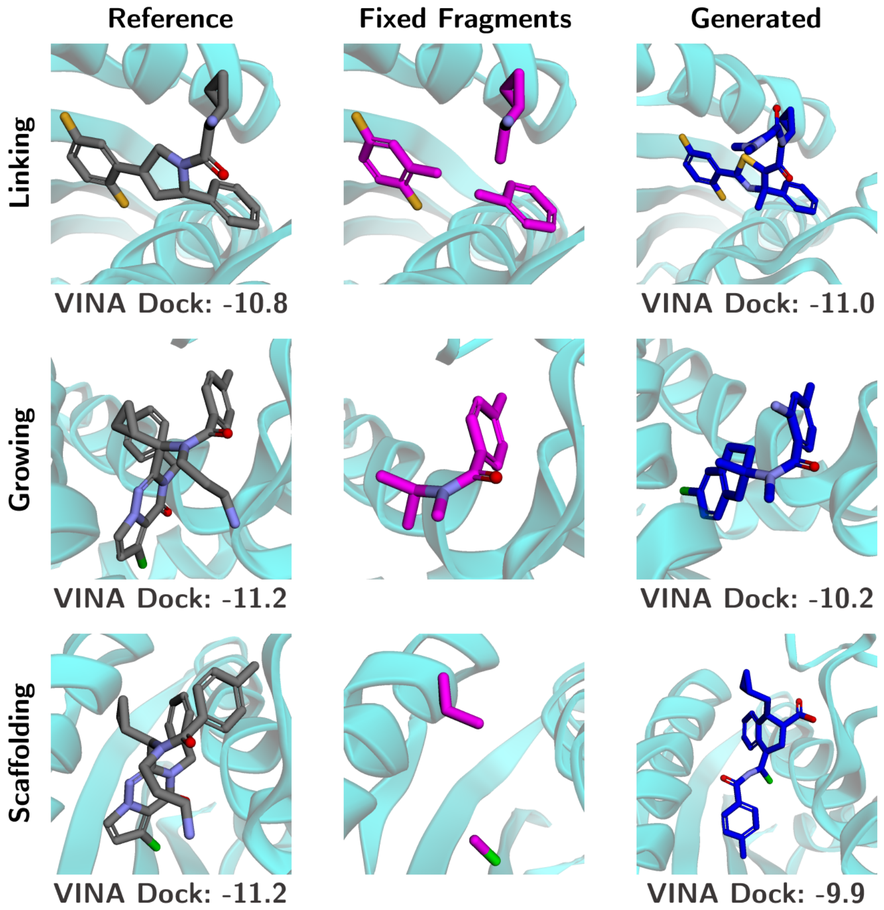

Unified Guidance for Geometry-Conditioned Molecular Generation

Sirine Ayadi[*1,2] **Leon Hetzel**[*1,2,3] **Johanna Sommer**[*1,2]

Fabian Theis[1,2,3] **Stephan Günnemann**[1,2]

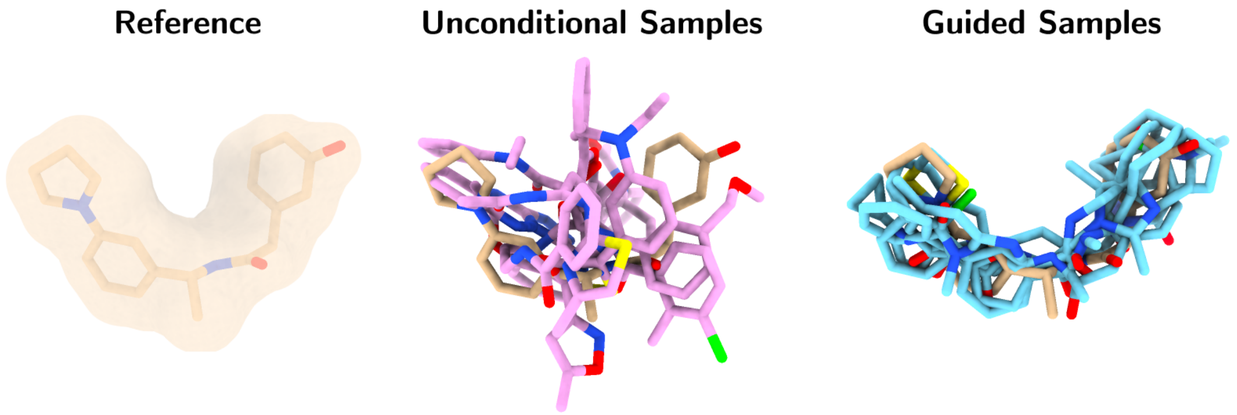

https://www.cs.cit.tum.de/daml/uniguide/