Symmetry-Preserving Diffusion Models via Target Symmetrization

Vinh Tong, Yun Ye, Dung Trung Hoang, Anji Liu, Guy Van den Broeck, Mathias Niepert
trung.hoang@unibe.ch

We propose a **loss function** that enables **diffusion models** to respect **data symmetries** without **architectural constraints** or **data augmentation**, ensuring **stable, low-variance training**.

Introduction

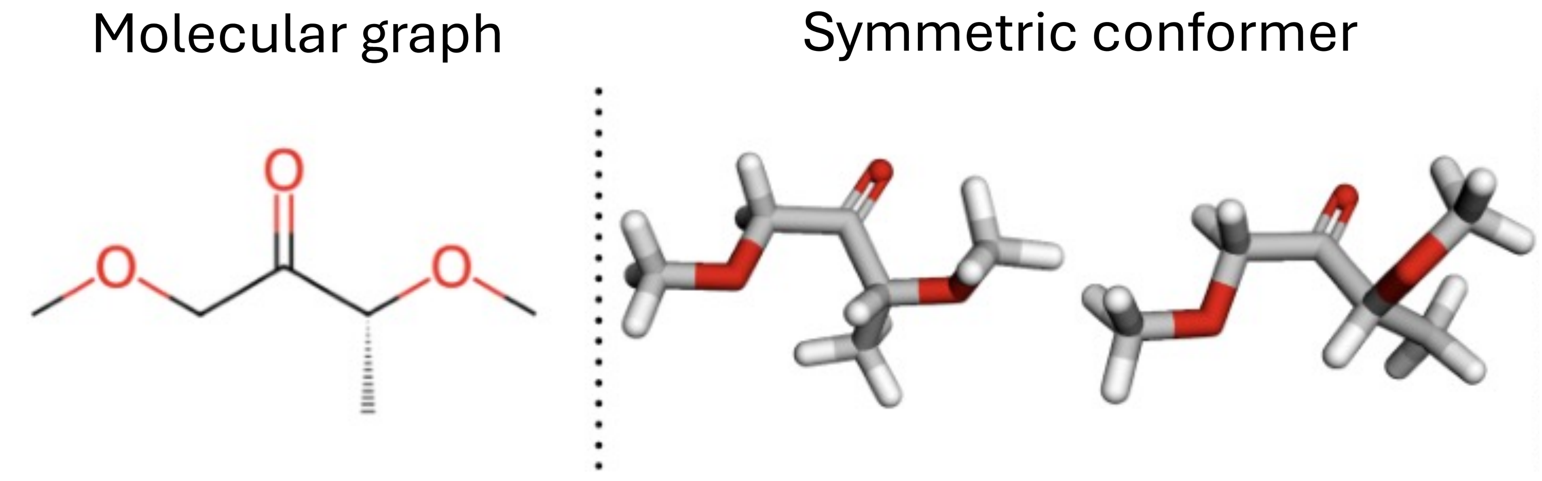

- **Diffusion models** struggle with symmetries (e.g., rotations, reflections) since the **empirical dist.** is not necessarily **G-invariant**.
- Existing methods:
  - **Equivariant denoisers** → complex to design.

$$\mathcal{L}_t(\phi) = E_{(x_0,x_t)\sim\hat{q}_t(x_0,x_t)}\left[\omega(t)\|\phi(x_t,t)-x_0\|^2\right]$$
$$\text{s.t}\quad \phi(g\circ x_t,t) = g\circ\phi(x_t,t)\quad \forall x_t\in\Omega \text{ and } \forall g\in G$$

  - **Data augmentation** to approximate G-invariant dist. → expensive.

$$\mathcal{L}^G_t(\phi) = E_{(x_0,x_t)\sim\hat{q}^G(x_0,x_t)}\left[\omega(t)\|\phi(x_t,t)-x_0\|^2\right]$$

- Our idea: **Symmetrized loss** → simple, stable, symmetry-aware.

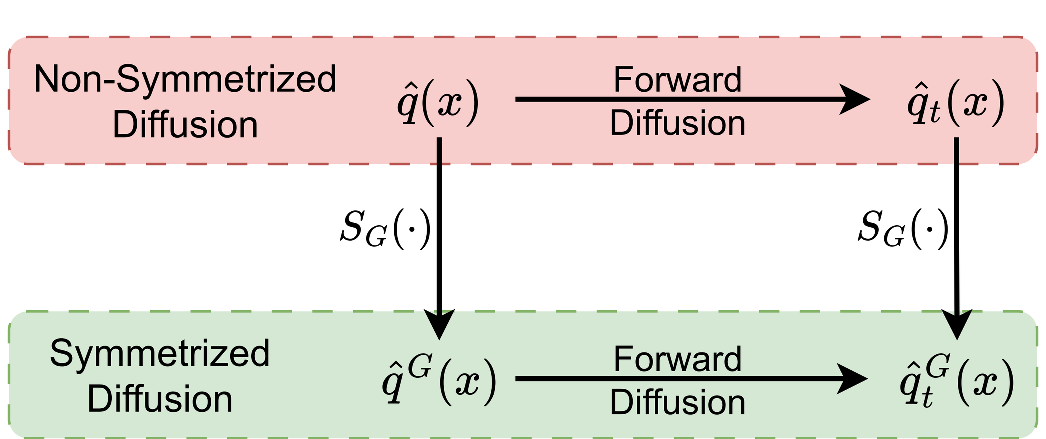

Method

$$\tilde{\mathcal{L}}^G_t(\phi) = E_{(x_0,x_t)\sim\hat{q}_t(x_0,x_t)}\left[\omega(t)\|\phi(x_t,t)-\phi^*_G(x_t,t)\|^2\right]$$
$$\text{s.t}\quad \phi^*_G(x_t,t) = \int_\Omega x_0\hat{q}^G_t(x_0\mid x_t)\,\mathrm{d}x_0$$

- The minimizer is **provably G-equivariant.**
- **Reduced gradient variance** → stable training.

Algorithm 1 Approximation of $\phi^*_G(x_t,t)$

1: Draw M samples $x_0^m\sim\hat{q}(x_0)$, K group elements $g_k\sim\mu_G$
2: **for each** x_0^m, g_k **do**
3: $\quad x_0^{mk}\leftarrow g_k\circ x_0^m,\quad R_{mk}\leftarrow\frac{-\|x_t-\alpha_t x_0^{mk}\|^2}{2\sigma_t^2}$
4: **end for**
5: $w_{mk}\leftarrow\frac{\exp(R_{mk})}{\sum_{j=1}^{MK}\exp(R_j)},\quad \phi^*_G(x_t,t)\leftarrow\sum_{j=1}^{MK}w_jx_0^j$

ICLR

Results

Reflection Group in 1D

Generate ±1 from a single training sample using reflection symmetry.

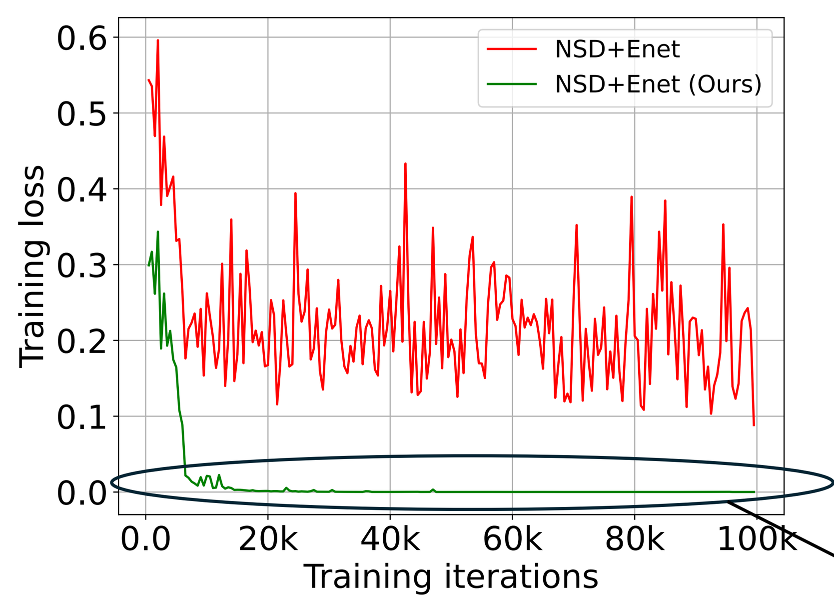

Models	RMSD (x10^{-3})	W2 Distance (x10^{-3})
NSD + Enet	8.392	0.014
NSD + Enet (Ours)	2.049	0.002

Our method converges faster with lower variance.

Rotation and Permutation Groups in 2D

Generate a 4-point 2D point cloud, ensuring equivariance to rotation and permutation.

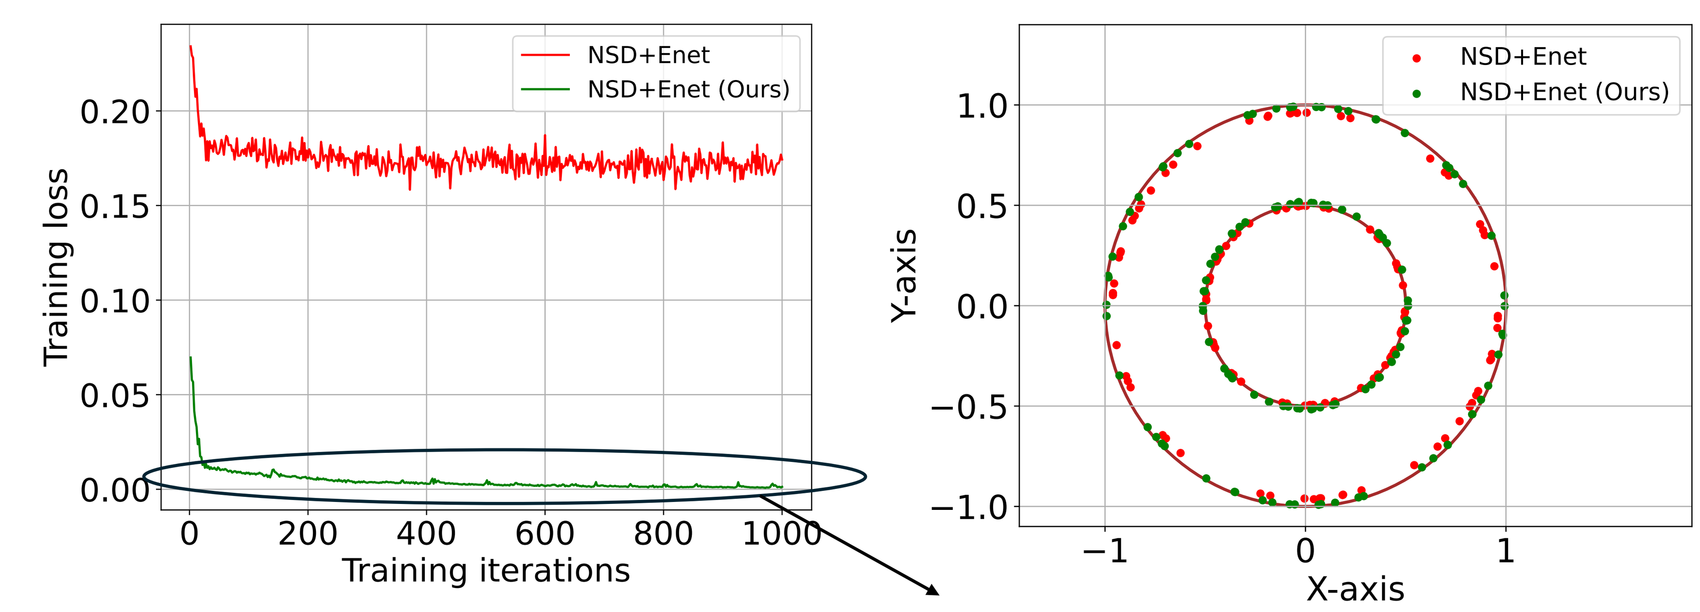

Our method converges faster with lower variance.

Molecular Conformer Prediction

Models	Recall				Precision			
	Coverage ↑		AMR ↓		Coverage ↑		AMR ↓	
	M	Mdn	M	Mdn	M	Mdn	M	Mdn
ETDiff	80.7	**100.0**	0.092	0.038	71.4	75.0	0.162	0.093
ETDiff + Ours	81.0	100.0	0.089	0.037	72.2	80.0	0.151	0.077
ETFlow	83.8	**100.0**	0.082	0.032	77.7	90.0	0.130	0049
ETFlow + Ours	**84.7**	100.0	**0.079**	**0.028**	**79.1**	**92.5**	**0.122**	**0.046**

Conclusion

Our method injects equivariance via a simple loss modification.

- No architectural changes.
- Minimal overhead.
- Strong theoretical backing.
- Better generalization on symmetric data.

Universität Stuttgart
imprs-is
UCLA
UNIVERSITÄT BERN