\Orchestrating a brighter world

Universität Stuttgart

LoGDIFF: Logical Guidance for the Exact Composition of Diffusion Models

ICML 2026

Francesco Alesiani* [1], Jonathan Warrell* [2], Tanja Bien* [3],
Henrik Christiansen[1], Matheus Ferraz[4], Mathias Niepert[1 3]

*Equal Contribution [1]NEC Laboratories Europe [2]NEC Laboratories America [3]University of Stuttgart [4]NEC OncoImmunity

LoGDiff: Exact Boolean Calculus

For a Boolean formula φ, we derive an exact calculus for computing the **logical score**, $s_t(\varphi, \boldsymbol{x}) = \nabla_{\boldsymbol{x}} \log p_t(\varphi \mid \boldsymbol{x})$, where mixing coefficients are **dynamic** and depend on the **time-varying posterior probability** $\widehat{\pi}(\varphi)$.

Operator	**Posterior Prob.** $\widehat{\pi}(\varphi)$	**Logical Score** $s_t(\varphi, \boldsymbol{x})$
Atom c	$p_t(c \mid \boldsymbol{x})$	$\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x} \mid c) - \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x})$
Negation $\neg\psi$	$1 - \widehat{\pi}(\psi)$	$-\frac{\widehat{\pi}(\psi)}{1-\widehat{\pi}(\psi)}\, s_t(\psi, \boldsymbol{x})$
Conjunction $\psi \wedge \chi$	$\widehat{\pi}(\psi)\widehat{\pi}(\chi)$	$s_t(\psi, \boldsymbol{x}) + s_t(\chi, \boldsymbol{x})$
Disjunction (ME)	$\widehat{\pi}(\psi) + \widehat{\pi}(\chi)$	$\frac{\widehat{\pi}(\psi) s_t(\psi, \boldsymbol{x}) + \widehat{\pi}(\chi) s_t(\chi, \boldsymbol{x})}{\widehat{\pi}(\psi) + \widehat{\pi}(\chi)}$

Sufficient Condition: The formula admits a circuit representation where conjunctions are conditionally independent (CI) and disjunctions are mutually exclusive (ME) or CI.

Theoretical Analysis

- We introduce two types of query that meet our sufficiency condition:
    - **Categorical-CI queries**: example digits, shapes or target proteins
    - **Taxonomy queries** (i.e. nested queries) : example protein secondary structure classification
- We show that our recursive rules are **representationally complete** (any Boolean formula may be compiled into a valid circuit) for the following types of distributions:
    - atoms are sets of **mutually CI categorical variables**
    - **pairs of distinct atoms** are either ME or nested (i.e atoms form a taxonomy)
    - **nested atoms at terminal time** $\implies$ representational completeness at all times

Practical Implementation: Hybrid Guidance

Hybrid Strategy:

- Scores (s_t) come from *Classifier-Free Guidance*.
- Weights ($\widehat{\pi}$) come from *Noise-aware Classifiers*.

Adaptive Repulsive Guidance:

- Goal: Avoid class confusion (e.g., $A \wedge \neg B$).
- Push away from competing class B only when B is locally probable.
- Improves FID on ImageNet.

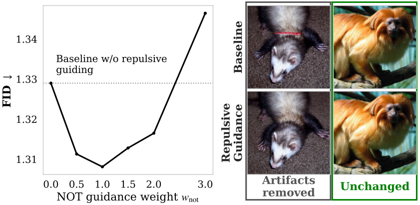

Repulsive Guidance (RG) improves ImageNet FID.

Results: Complex Logical Queries

Benchmark (CMNIST/Shapes3D):

- Tested recursive queries up to complexity $N = 5$.
- LoGDiff maintains **high conformity and high diversity** (Joint Entropy)

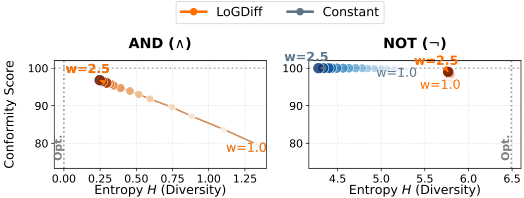

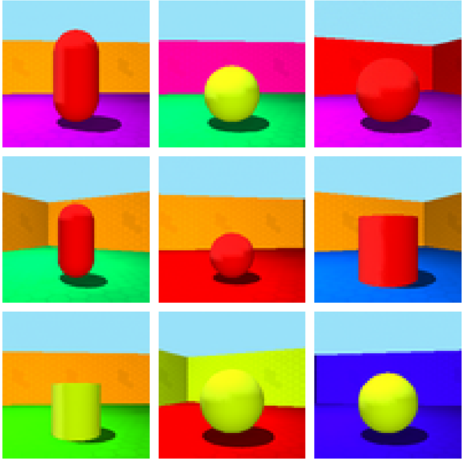

Diverse samples for (*OrangeWall* ∨ *Sphere*) ∧ (*Red* ∨ *Yellow*).

Application: Selective Drug Design

Designing ligands that bind to target A but **not** target B.

- **Dual Binding:** $\varphi = A \land B$.
- **Selective Binding:** $\varphi = A \land \neg B$.

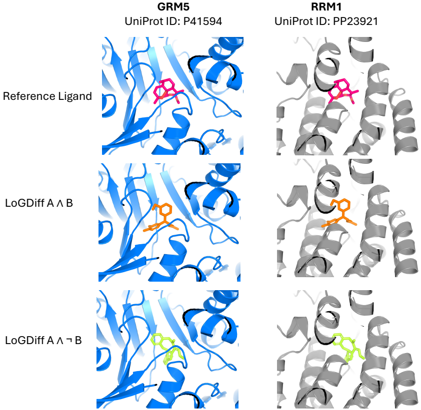

- LoGDiff improves the $\Delta\Delta G$ in docking scores (On-target vs. Off-target).
- Outperforms specialized baselines (DualDiff).

Conclusion

- LoGDiff provides an **exact Boolean calculus** for diffusion guidance.
- Translates logical structure into **dynamic mixture coefficients**.
- Scalable to complex, nested queries ($N > 5$).
- Broad applicability: From image steering to selective protein-ligand binding.

Code and paper available at: