Fuzzy Logic Composition of Diffusion Models

Peter Blohm
Aalto University
peter.blohm@aalto.fi

Vikas Garg
YaiYai Ltd and Aalto University
vgarg@csail.mit.edu

Abstract

We investigate the composition of diffusion models as combinatorial operations on their underlying probability distributions. Theory from fuzzy logic allows us to define a class of score-composition operators that follow DeMorgan's law. We reframe Mixture of Experts in this context to obtain the class of *Dombi* operators, which provide a parametrised tradeoff between smoothness and precise combinatorial reasoning. We show that score-composition with Products of Experts (PoE) can lead to unstable behaviour and implicit temperature rescaling. In preliminary experiments on SAT problems in continuous settings, we show that Dombi composition enables combinatorial reasoning and does not suffer from the bias of PoE.

1 Introduction

We investigate compositions of diffusion models from the viewpoint of fuzzy set theory and fuzzy logic. We propose a procedure to derive sets of well-behaved composition operators, and among them, propose *Dombi* operators, which extend the commonly used mixture of experts [Jordan and Jacobs, 1994] (MoE) to allow modelling of conjunction, disjunction, and explicit negation.

The high fidelity and flexibility of diffusion models have led to widespread use. As they, in essence, just sample from probability distributions, recent research efforts focus on using ensembles of diffusion models to encode constraints or enable precise steering of the diffusion process [Ho and Salimans, 2021, Du et al., 2020, Liu et al., 2022, Garipov et al., 2023, Du and Kaelbling, 2024, Skreta et al., 2025a,b, Hasan et al., 2025, Thornton et al., 2025]. In the context of energy-based models (EBMs), compositional generation has long been explored [Jordan and Jacobs, 1994, Hinton, 1999, 2002, Du et al., 2020] via algebraic manipulation of learned energy functions. In this prior work, different connectives to imitate unions and intersections have been proposed, although products and mixtures of experts (PoE, MoE) are most commonly used.

We show that existing operators, and especially the used *sets* of operators, do not carry the well-understood and favorable properties of fuzzy set operators. These missing properties make them ill-equipped to deal with more complex compositions of models or to encode model constraints without introducing biases. Specifically, in Appendix A, we provide a definition for a general class of DeMorgan dual operators for score models, and show that operators based on the Dombi t-norm generalize the classical mixture and can be tuned to approximate idempotent and distributive behavior. In Section 4, we motivate the advantages of these operators and illustrate pathologies that are avoided. In Section 5, we show as proof of concept that ensembles of models can be composed to allow combinatorial reasoning and compare the Dombi operators against PoE.

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: DiffCoALG.

2 Background & Related Work

2.1 Fuzzy Logic

Fuzzy logic relaxes classical logic from a binary domain to the reals in $[0, 1]$. We follow the definitions and notation from Klement et al. [2013] for the following concepts. We define a *t-norm*, a generalization of conjunction or intersection operations as a function $T : [0, 1]^2 \to [0, 1]$ which is commutative, associative, monotonously increasing, and fulfills the boundary condition $\forall x \in [0, 1] : T(x, 1) = x$. Under the standard negation $N(x) = 1 - x$, we can define the *dual t-conorm* $S : [0, 1]^2 \to [0, 1]$, the corresponding disjunction, via DeMorgan's law as

$$S(x, y) = N(T(N(x), N(y))).$$

T-norms that are *strict*, i.e., continuous and strictly increasing, can be *generated* [Dombi, 1982, Klement et al., 2013] by a continuous, strictly decreasing function $f : [0, 1] \to [0, \infty]$ with $f(1) = 0$, as so-called *additive generator*, i.e., $T(x, y) \coloneqq f^{-1}\left(f(x) + f(y)\right)$. For this work, the parametrised Dombi t-norm is the most important representative, generated by $f_\lambda(x) = (\frac{1}{x} - 1)^{-\lambda}$, it is defined as

$$T_\lambda(x, y) = \frac{1}{1 + \left(\left(\frac{1}{x} - 1\right)^\lambda + \left(\frac{1}{y} - 1\right)^\lambda\right)^{1/\lambda}}. \quad (1)$$

A favorable property of the Dombi t-norm, is that $\lim_{\lambda\to\infty} T_\lambda = T_M = \min$. The minimum t-norm T_M together with $S_M = \max$ forms the *only* DeMorgan dual that is continuous, idempotent with $T_M(x, x) = x$ and that is distributive with $T_M(x, S_M(y, z)) = S_M(T_M(x, y), T_M(x, z))$ and is the basis for Gödel logic [Klement et al., 2013].

2.2 Score-Based Models

We want to generate samples $\mathbf{x} \in \mathcal{X}$, following some unknown distribution represented by a score-model. A score model parametrised by θ allows us to draw samples according to the density function

$$p_\theta(\mathbf{x}) = \frac{\exp(-E_\theta(\mathbf{x}))}{Z_\theta}.$$

Here $E_\theta(\mathbf{x})$ (the energy) represents the unnormalized negative log-likelihood, and Z_θ is the (unknown) normalizing constant

$$Z_\theta = \int \exp(-E_\theta(\mathbf{x}))d\mathbf{x}.$$

Diffusion models then learn the *score function* $s_\theta(\mathbf{x}) = \nabla_\mathbf{x} \log p_\theta(\mathbf{x}) = -\nabla_\mathbf{x} E_\theta(\mathbf{x})$. Our target distribution, however, is not static: we first consider a family of *perturbed distributions* $\{p_t(\mathbf{x})\}_{t\in[0,1]}$, defined by a forward *diffusion process*. Starting from the data distribution $p_0(\mathbf{x})$, we gradually add Gaussian noise according to the Itô SDE

$$d\mathbf{x}_t = f(\mathbf{x}_t, t)\, dt + g(t)\, d\mathbf{w}_t, \qquad \mathbf{x}_0 \sim p_0,$$

where $f(t)$ is the drift coefficient, $g(t)$ the diffusion coefficient, and $\mathbf{w}_t$ a standard Wiener process [Ho et al., 2020, Song et al., 2021]. This defines a Markov chain with marginal densities $p_t(\mathbf{x})$ that interpolate between the data distribution ($t = 0$) and an easy-to-sample prior such as $\mathcal{N}(0, \mathbf{I})$ ($t = 1$). This construction is linear: applying the noising operator N_t to a mixture of distributions is equivalent to mixing the individually perturbed distributions, $N_t\Big(\sum_i \lambda_i p^{(i)}\Big) = \sum_i \lambda_i N_t\big(p^{(i)}\big)$. To sample, we simulate the backward process

$$d\mathbf{x}_t = \left[f(\mathbf{x}_t, t) - g^2(t)\nabla_\mathbf{x} \log p_t(\mathbf{x}_t)\right] dt + g(t)d\overline{\mathbf{w}}.$$

As the focus of this work is not on the dynamics of this process, we write our results for $t = 0$ unless explicitly mentioned otherwise.

2.3 Composition of Score Fields

A rapidly growing body of work exists on the compositions, mixtures, and products of EBMs, flow, and diffusion models. We explicitly focus on *training-free* mixtures of score functions in diffusion. Prior work [Du et al., 2020, Skreta et al., 2025b] mainly bases composition on probabilistic operations on the underlying distributions. As the interpretation of these operations is often logical or set-theoretic, we will use the symbols $\{\vee, \wedge, \neg\}$ to denote them, for both probability densities and their scores. In score-based modelling, conjunctions are then represented by products

$$p_1(\mathbf{x}) \wedge_\times p_2(\mathbf{x}) \coloneqq p_1(\mathbf{x})p_2(\mathbf{x}) \Longrightarrow s_1(\mathbf{x}) \wedge_\times s_2(\mathbf{x}) = s_1(\mathbf{x}) + s_2(\mathbf{x}) \tag{2}$$

and disjunctions by mixtures

$$p_1(\mathbf{x}) \vee_+ p_2(\mathbf{x}) \coloneqq \frac{1}{2}p_1(\mathbf{x}) + \frac{1}{2}p_2(\mathbf{x}) \Longrightarrow \tag{3}$$

$$s_1(\mathbf{x}) \vee_+ s_2(\mathbf{x}) \;=\; \sum_{i\in\{1,2\}} \frac{\exp(-E_i(\mathbf{x}))}{\sum_{j\in\{1,2\}}\exp(-E_j(\mathbf{x}))} s_i(\mathbf{x}), \tag{4}$$

where the required energies, i.e., the negative log-likelihoods, can be estimated during inference with Itô's lemma [Karczewski et al., 2025, Skreta et al., 2025b]. It is to note that for $t \neq 0$, nonlinear compositions do not commute with the noising operator, i.e., $N_t(p_1) \vee_+ N_t(p_2) = N_t(p_1 \vee_+ p_2)$ but $N_t(p_1) \wedge_\times N_t(p_2) \neq N_t(p_1 \wedge_+ p_2)$. This means that naive composition of perturbed score models leads to a bias that can be corrected with methods like sequential Monte Carlo [Skreta et al., 2025a, Thornton et al., 2025].

To *avoid* samples from certain distributions, EBM's and score models are often negated *relative* to others [Vedantam et al., 2018, Garipov et al., 2023, Skreta et al., 2025a], as also done in classifier-free guidance [Ho and Salimans, 2021] (CFG), e.g., for $\gamma \in \mathbb{R}_{>0}$ $p_1(\mathbf{x})/p_2(\mathbf{x})^\gamma$. This and similar operations make implicit use of the reciprocal as pseudo-inverse $\tilde{\neg}\, p_2(\mathbf{x}) = 1/p_2(\mathbf{x})$, but to the best of our knowledge, explicit negations in score-models are not often explored. Two noteworthy exceptions from product-based conjunctions are Garipov et al. [2023], who model conjunctions with the *harmonic mean*

$$p_1(\mathbf{x}) \wedge_H p_2(\mathbf{x}) = \frac{p_1(\mathbf{x})p_2(\mathbf{x})}{p_1(\mathbf{x}) + p_2(\mathbf{x})},$$

but do not define disjunctions, and Skreta et al. [2025b], who reweigh individual scores at each time (even negatively) to approach a region of equal density in all distributions.

3 Dombi Density Operators

We can extend the definition of T-norm and conorm pairs to obtain DeMorgan dual density and score operators. Appendix A describes the exact requirements to generate a set of operators. As a special class we propose and investigate the DeMorgan operators generated by $f_\lambda(x) = \left(\frac{1}{x} - 1\right)^{-\lambda}$ and $\phi_c(x) = \frac{x}{x+c}$ for $\lambda, c \in \mathbb{R}_{\geq 0}$. This choice of f not only recovers the Dombi t-norm, but ϕ_c expresses negation with *reference* to c. This helps stabilize the behaviour of the operator in practice, and shows that our formalism is fully compatible with Garipov et al. [2023].

Definition 3.1 (Dombi Operators). *Choose $\lambda \in R_{\geq 0}$ and a fixed continuous function $c : \mathcal{X} \to \mathbb{R}_{\geq 0}$. With $f_\lambda(x) = \left(\frac{1}{x} - 1\right)^{-\lambda}$ and $\phi_{c(\mathbf{x})}(x) = \frac{1}{1+\frac{c(\mathbf{x})}{x}}$ for $\lambda \in \mathbb{R}_{>0}$, where $s_c = \nabla_\mathbf{x} \log c$. Then the* Dombi operators *are defined as*

$$\neg_c p(\mathbf{x}) = \frac{c(\mathbf{x})^2}{p(\mathbf{x})} \Longrightarrow \neg_c s(\mathbf{x}) = 2s_c(\mathbf{x}) - s(\mathbf{x}) \tag{5}$$

$$p_1(\mathbf{x}) \wedge_\lambda p_2(\mathbf{x}) = \frac{p_1(\mathbf{x})p_2(\mathbf{x})}{\left(p_1(\mathbf{x})^\lambda + p_2(\mathbf{x})^\lambda\right)^{1/\lambda}} \Longrightarrow s_1(\mathbf{x}) \wedge_\lambda s_2(\mathbf{x}) = \frac{\sum_{i\in\{1,2\}}\exp(\lambda E_i(\mathbf{x}))s_i(\mathbf{x})}{\sum_{i\in\{1,2\}}\exp(\lambda E_i(\mathbf{x}))} \tag{6}$$

$$p_1(\mathbf{x}) \vee_\lambda p_2(\mathbf{x}) = \left(p_1(\mathbf{x})^\lambda + p_2(\mathbf{x})^\lambda\right)^{1/\lambda} \Longrightarrow s_1(\mathbf{x}) \vee_\lambda s_2(\mathbf{x}) = \frac{\sum_{i\in\{1,2\}}\exp(-\lambda E_i(\mathbf{x}))s_i(\mathbf{x})}{\sum_{i\in\{1,2\}}\exp(-\lambda E_i(\mathbf{x}))} \tag{7}$$

This definition bears multiple remarkable properties. First, we can note that $p_1(\mathbf{x}) \vee_1 p_2(\mathbf{x}) = p_1(\mathbf{x}) + p_2(\mathbf{x})$ and $p_1(\mathbf{x}) \wedge_1 p_2(\mathbf{x}) = \frac{p_1(\mathbf{x})p_2(\mathbf{x})}{p_1(\mathbf{x})+p_2(\mathbf{x})}$, the harmonic mean operator used by Garipov et al. [2023] Further, using a probability distribution as a reference allows us to recover the contrast operator of Garipov et al. [2023] as well:

$$p(\mathbf{x}) \wedge \neg_{p(\mathbf{x})} q(\mathbf{x}) = \frac{p(\mathbf{x})^2}{\frac{p(\mathbf{x})^2}{p(\mathbf{x})} + q(\mathbf{x})} = \frac{p(\mathbf{x})^2}{p(\mathbf{x}) + q(\mathbf{x})} = q(\mathbf{x}) \,\mathbin{\text{◐}}\, p(\mathbf{x}).$$

Besides the connection to prior work, the parameter λ from the Dombi operators naturally appears as inverse temperature in the score composition. For $\lambda \to \infty$, the Dombi operators recover the exact $\{\min, \max\}$ lattice, along with its distributive and idempotent behavior. For finite λ, the following simple bounds can be used to quantify biases in density compositions.

This means that for high values of λ, we can recover an exact lattice structure, and with it, obtain idempotent and distributive laws. For a given value of λ, the maximal difference between the Dombi operators and the $\min$ / $\max$ functions can be easily bounded as an additive term in log-likelihood:

Proposition 3.2. *Let $\wedge_\lambda, \vee_\lambda$ be the Dombi density operators. Then it holds that*

$$\forall x, y \in \mathbb{R}_{\geq 0} : \frac{\min\{x, y\}}{2^{1/\lambda}} \leq x \wedge_\lambda y \leq \min\{x, y\} \tag{8}$$

$$\forall x, y \in \mathbb{R}_{\geq 0} : \max\{x, y\} \leq x \vee_\lambda y \leq \max\{x, y\} 2^{1/\lambda} \tag{9}$$

Proposition 3.2 gives us a constant additive error-bound on energies, which we can easily propagate through formulas to gauge the maximal density bias we can introduce when applying distributive laws.

Corollary 3.3 (Idempotency Bias)**.** *Let $\wedge_\lambda, \vee_\lambda$ be the Dombi density operators. From the proof of Proposition 3.2 it follows that*

$$\forall x \in \mathbb{R}_{\geq 0} : x \vee_\lambda x = 2^{1/\lambda} x \tag{10}$$

$$\forall x \in \mathbb{R}_{\geq 0} : x \wedge_\lambda x = 2^{-1/\lambda} x \tag{11}$$

Corollary 3.4 (Distributivity Bias)**.** *Let $\wedge_\lambda, \vee_\lambda$ be the Dombi density operators. From Proposition 3.2 it follows that formulas distribute with bounded bias, i.e.,*

$$\forall x, y, z \in \mathbb{R}_{\geq 0} : x \vee_\lambda (y \wedge_\lambda z) \in ((x \vee_\lambda y) \wedge_\lambda (x \vee_\lambda z)) 2^{\pm 2/\lambda} \tag{12}$$

$$\forall x, y, z \in \mathbb{R}_{\geq 0} : x \wedge_\lambda (y \vee_\lambda z) \in ((x \wedge_\lambda y) \vee_\lambda (x \wedge_\lambda z)) 2^{\pm 2/\lambda} \tag{13}$$

4 Properties of Combinatorial Operators

In the previous section, we introduced DeMorgan density and score operators. As a specific example, we demonstrated that the Dombi operators enable a tradeoff between smoothness, distributivity, and idempotence. While the smoothed lattice formed by these operators facilitates a combinatorial interpretation of composition, it sacrifices a clear probabilistic interpretation. In this section, we provide an overview of how DeMorgan duality, idempotency, and distributivity facilitate ensemble settings to motivate this trade-off. For each of these properties, we illustrate how its absence leads to failure modes, or "intuition pitfalls."

4.1 Bounded Negation

While negation is a well-defined concept in fuzzy logic, it cannot be defined as a pointwise operator on probability densities. As a proxy for a proper negation, reciprocals of distributions are commonly used to discourage the generation of points from their high-density regions. Using reciprocals carelessly is problematic, even if, in terms of scores, only a sign flip is needed. In settings with unbounded support, for a given density function $p(\mathbf{x})$, in general $1/p(\mathbf{x})$ blows up at the tails and is non-normalizable, i.e., $\int 1/p(\mathbf{x}) d\mathbf{x} = \infty$. Garipov et al. [2023] observe that the harmonic mean can still utilize "negated" distributions safely, as the result of their contrast operator ◐ is always bounded. Similarly, for the Dombi operators, we can use reciprocals of distributions safely as 1. Proposition 3.2 guarantees that

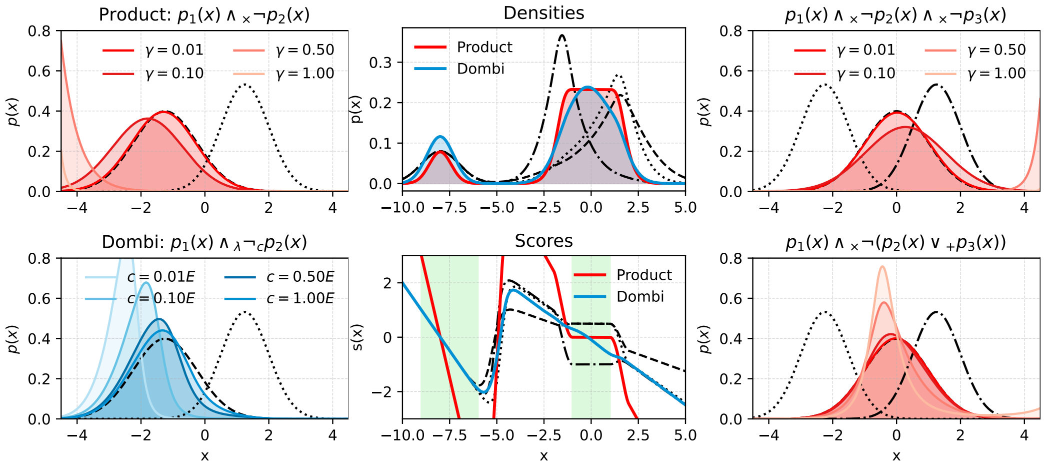

(a) Negation Stability (b) Temperature Rescaling (c) No DeMorgan Duality

Figure 1: Failure Modes of PoE/MoE composition in combinatorial settings.
Figure 1a illustrates that direct fractions such as $p_1(\mathbf{x})/p_2(\mathbf{x})^\gamma$ can lead to unstable behaviour, compared to Dombi operators with constant reference c as fraction of $\mathbb{E}p_1(\mathbf{x})$.
Figure 1b shows an intersection, where the product can lead to locally underscaled *and* overscaled temperatures simultaneously (green), in contrast to the convex score combinations of Dombi operators.
Figure 1c shows that logically equivalent formulas result in different PoE/MoE compositions.

any nonnegative function $q(\mathbf{x})$, under conjunction with a density $p(\mathbf{x})$, yields a normalizeable density. 2. If $q(\mathbf{x})$ models a *subpopulation* of some distribution $p(\mathbf{x})$, we can assume $\frac{p(\mathbf{x})^2}{q(\mathbf{x})}$ has a finite upper bound, so referenced negation is stable in isolation here. 3. Under finite support, the reciprocal of any density is finite; we can use, e.g., a uniform density over the support to reference negation. Figure 1a illustrates the formula $p(\mathbf{x}) \wedge \neg q(\mathbf{x})$ with different operators to compare their parameter sensitivity.

4.2 Implicit Temperature Rescaling

In the context of probability densities, idempotent behavior ensures that scores remain stable. In the case of Dombi operators, we are guaranteed a convex combination of scores. This means that the effective temperature of our composition is bounded above by the temperatures of the individual models. For higher λ, the composite score will converge to one of the components, which also prevents *reductions* in effective temperature.

If our operators do not respect idempotency, we might end up with locally temperature-scaled behavior, which is difficult to control. For example for $p \approx q \approx r$ it holds that $p \wedge_\times q \wedge_\times r \approx p^3$, while $p \wedge_\lambda q \wedge_\lambda r \approx p$, for high enough λ, as per Proposition 3.2. This difference in behaviors is illustrated in Figure 1b.

4.3 Equivalent Formula Rewriting

The DeMorgan laws allow us to freely move negations across disjunctions and conjunctions while not changing the result of our score composition. Formulas can then be expressed in negation normal form, ensuring that negations are applied only to individual score models and preventing the arbitrary and numerically challenging nesting of negations. If we can further assume distributivity (approximately, Corollary 3.4), we can apply factoring to reduce the size of compositions. Similarly, we can rewrite it in disjunctive normal form to represent compositions as mixtures, which speeds up sampling.

Finally, we argue that DeMorgan's Law prevents unintuitive behavior. If we want to sample from p, but steer away from distributions q and r, DeMorgan duality can prevent modelling mistakes (Figure 1c):

$$p \wedge_\times \tilde{\neg} q \wedge_\times \tilde{\neg} r = p \wedge_\times \tilde{\neg} (p \wedge_\times q) = \frac{p}{qr} \neq p \wedge_\times \tilde{\neg} (p \vee_+ q) = \frac{p}{q+r}.$$

5 Experiments

5.1 Setup

We illustrate the capability of Dombi compositions to adhere to combinatorial constraints by sampling uniformly from satisfying variable assignments of propositional formulas. For a formula with k propositional variables P_i, for $i \in [1, k]$, we set up our diffusion ensemble as follows: In $\mathbb{R}^k$, we place 2^k Guassian modes, one for each possible variable assignment. Then, in our ensemble, each of k score models simulates one propositional variable. For $i \in [1, k]$, we have access to s_i, which defines a denoising process to a uniform mixture of the 2^{k-1} Gaussian modes, where the P_i is true. Additionally, a reference model defines a denoising process uniformly to *all* 2^k Gaussian modes. For $k = 2$, this setup is visualized in Figure 2a.

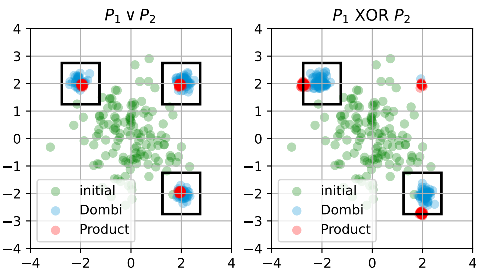

(a) SAT experiment in $\mathbb{R}^2$.

Formula	Method	$k=2$ η_{corr}	$k=2$ η_{all}	$k=2$ PPL↑	$k=5$ η_{corr}	$k=5$ η_{all}	$k=5$ PPL↑	$k=10$ η_{corr}	$k=10$ η_{all}	$k=10$ PPL↑
Maj_k	Dombi	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**0.98**
	Prod	**1.00**	**1.00**	0.82	**1.00**	**1.00**	0.09	0.00	0.00	0.00
XOR_k	Dombi	0.97	**1.00**	**1.00**	**0.95**	**1.00**	**1.00**	**0.91**	**1.00**	**0.98**
	Prod	**1.00**	0.00	**1.00**	0.00	0.00	0.00	0.00	0.00	0.00
OneHot_k	Dombi	0.97	**1.00**	**1.00**	**0.59**	**1.00**	**1.00**	**0.09**	**1.00**	**0.99**
	Prod	**1.00**	0.00	**1.00**	0.00	0.00	0.00	0.00	0.00	0.00

(b) Overview of SAT experiment for three formulas.

Figure 2: Figure 2a shows the SAT experiment in $\mathbb{R}^2$, with squares corresponding to satisfying assignments. The corresponding numerical overview for $k \in \{2, 5, 10\}$ in Figure 2b. Best are bold.

Our objective is then to use score-composition to uniformly sample from all satisfying variable assignments. We repeat this setup for the Dombi operators, as well as PoE/MoE composition for three formulas for $k \in [1, 10]$, and report mode coverage, uniformity, and stability of the composition. For further details of the experimental setup and detailed experimental results, refer to Appendix C.

5.2 Results

Figure 2a shows samples of formulas in $\mathbb{R}^2$. An overview of the experimental results is provided in Figure 2b. We can see multiple shortcomings of products in our experimental results. On the negation-free Maj_k, PoE drastically reduces the per-mode variance, as seen in Figure 2a, drops most of the modes for $k = 5$, and completely breaks down for $k = 10$. In contrast, the dombi Operators do not drop modes and maintain a close-to-uniform distribution over modes in high dimensions. For XOR_k and OneHot_k, PoE breaks down for $k = 2$ already, due to the negated literals. In Figure 2a, the modes of the PoE sample appear drastically biased by the negated clause. Somewhat surprisingly, the Dombi composition can sample comparatively well from the exponentially sized XOR_{10}, and struggles much more for OneHot, which is comprised of many purely negative clauses.

6 Discussion and Limitations

This work investigates composition methods for score-based models that improve *combinatorial* reasoning. It advocates for a clearer distinction between probabilistic and combinatorial arguments, highlighting the shortcomings of probabilistic operators in combinatorial settings. The Dombi operators in particular re-ground MoE in Fuzzy logic and allow us to obtain guarantees about their behavior in the context of temperature scaling.

For the practical evaluation of this work in realistic settings, more and larger-scale experiments are required. While toy experiments on SAT reflect differences in the behaviour of Dombi composition and PoE, practical differences in suboptimal settings, and in complex models need to be explored further. Promising directions for future research include highly constrained neurosymbolic diffusion tasks, such as Sudoku [Avdeyev et al., 2023], where ensembled models have much simpler learning tasks. Similarly, the exploration of unlearning or *shielded generation* [Kirchhof et al., 2025] might be an application where the combinatorial approach be favorable to probabilistic methods.

7 Acknowledgements

This work was supported by the Finnish Center for Artificial Intelligence (FCAI) under Flagship R5 (award 15011052). VG also acknowledges the support from Saab-WASP (grant 411025), the Academy of Finland (grant 342077), and the Jane and Aatos Erkko Foundation (grant 7001703).

References

Pavel Avdeyev, Chenlai Shi, Yuhao Tan, Kseniia Dudnyk, and Jian Zhou. Dirichlet diffusion score model for biological sequence generation. In *International Conference on Machine Learning*, pages 1276–1301. PMLR, 2023.

J. Dombi. A general class of fuzzy operators, the demorgan class of fuzzy operators and fuzziness measures induced by fuzzy operators. *Fuzzy Sets and Systems*, 8(2):149–163, August 1982. ISSN 0165-0114. doi: 10.1016/0165-0114(82)90005-7.

Yilun Du and Leslie Pack Kaelbling. Compositional generative modeling: A single model is not all you need. *CoRR*, abs/2402.01103, 2024. URL https://doi.org/10.48550/arXiv.2402.01103.

Yilun Du, Shuang Li, and Igor Mordatch. Compositional visual generation with energy based models. *Advances in Neural Information Processing Systems*, 33:6637–6647, 2020.

Timur Garipov, Sebastiaan De Peuter, Ge Yang, Vikas Garg, Samuel Kaski, and Tommi Jaakkola. Compositional Sculpting of Iterative Generative Processes. *Advances in neural information processing systems*, 36:12665–12702, 2023.

Mohsin Hasan, Marta Skreta, Alan Aspuru-Guzik, Yoshua Bengio, and Kirill Neklyudov. Discrete feynman-kac correctors. In *2nd AI for Math Workshop @ ICML 2025*, 2025. URL https://openreview.net/forum?id=CafvO7fIdq.

Geoffrey E Hinton. Products of experts. In *Proceedings of the Ninth International Conference on Artificial Neural Networks*, volume 1, pages 1–6. IET, 1999.

Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8):1771–1800, 2002.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. URL https://openreview.net/forum?id=qw8AKxfYbI.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 6840–6851, 2020.

Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6(2):181–214, 1994.

Rafal Karczewski, Markus Heinonen, and Vikas Garg. Diffusion models as cartoonists: The curious case of high density regions. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=RiS2cxpENN.

Michael Kirchhof, James Thornton, Louis Béthune, Pierre Ablin, Eugene Ndiaye, and Marco Cuturi. Shielded Diffusion: Generating Novel and Diverse Images using Sparse Repellency, May 2025.

Erich Peter Klement, Radko Mesiar, and Endre Pap. *Triangular norms*, volume 8. Springer Science & Business Media, 2013.

Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B. Tenenbaum. Compositional Visual Generation with Composable Diffusion Models. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 423–439, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-19790-1. doi: 10.1007/978-3-031-19790-1_26.

Marta Skreta, Tara Akhound-Sadegh, Viktor Ohanesian, Roberto Bondesan, Alán Aspuru-Guzik, Arnaud Doucet, Rob Brekelmans, Alexander Tong, and Kirill Neklyudov. Feynman-Kac Correctors in Diffusion: Annealing, Guidance, and Product of Experts, June 2025a.

Marta Skreta, Lazar Atanackovic, Avishek Joey Bose, Alexander Tong, and Kirill Neklyudov. The Superposition of Diffusion Models Using the Itô Density Estimator, February 2025b.

Yang Song, Jascha Sohl-Dickstein, Durk P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2021.

James Thornton, Louis Bethune, Ruixiang Zhang, Arwen Bradley, Preetum Nakkiran, and Shuangfei Zhai. Composition and Control with Distilled Energy Diffusion Models and Sequential Monte Carlo, February 2025.

Ramakrishna Vedantam, Ian Fischer, Jonathan Huang, and Kevin Murphy. Generative models of visually grounded imagination. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=HkCsm6lRb.

A DeMorgan Dual Fuzzy Logic Operators

In this section, we define the class of DeMorgan dual density and score operators, and investigate one example, the Dombi operators, in detail. We show that they generalize probabilistic mixtures and the harmonic mean, and discuss methods to stabilize explicitly used negations with these operators. We first extend the definition of fuzzy logic operators to the domain of probability densities.

Definition A.1 (DeMorgan Density Operators)**.** *Let $\phi : [0, \infty] \to [0, 1]$ be an order-isomorphism and $f : [0, 1] \to [0, \infty]$ be a continuous, strictly decreasing function with $f(0) = \infty$. For $g = f \circ \phi$, we define*

$$\neg p(\mathbf{x}) := \phi^{-1}(1 - \phi(p(\mathbf{x})) \tag{14}$$

$$p_1(\mathbf{x}) \wedge p_2(\mathbf{x}) := g^{-1}(g(p_1(\mathbf{x}) + g(p_1(\mathbf{x})) \tag{15}$$

$$p_1(\mathbf{x}) \vee p_2(\mathbf{x}) := \neg(\neg p(\mathbf{x}) \wedge \neg p(\mathbf{x})) \tag{16}$$

As f, by definition, generates a strict t-norm, the resulting set of operators is well-behaved, as we illustrate in Section 4. For differentiable f and ϕ, the application to scores follows directly:

Lemma A.1 (DeMorgan score calculus)**.** *Let ϕ and f be fully differentiable functions that generate the DeMorgan density operators $\{\wedge, \vee, \neg\}$. Then with $g = f \circ \phi$, $h : x \mapsto f(1 - \phi(x))$, $w(x) := x\, g'(x)$ and $\overline{w}(x) := x\, h'(x)$ the corresponding operations on the energies and scores are defined as*

$$\neg s(\mathbf{x}) = -\frac{\phi'(p(\mathbf{x}))p(\mathbf{x})}{\phi'(\neg p(\mathbf{x}))\neg p(\mathbf{x})}\, s(\mathbf{x}) \tag{17}$$

$$s_1(\mathbf{x}) \wedge s_2(\mathbf{x}) = \frac{w(p_1(\mathbf{x}))s_1(\mathbf{x}) + w(p_2(\mathbf{x}))s_2(\mathbf{x})}{w(p_1(\mathbf{x}) \wedge p_2(\mathbf{x}))} \tag{18}$$

$$s_1(\mathbf{x}) \vee s_2(\mathbf{x}) = \frac{\overline{w}(p_1(\mathbf{x}))s_1(\mathbf{x}) + \overline{w}(p_2(\mathbf{x}))s_2(\mathbf{x})}{\overline{w}(p_1(\mathbf{x}) \vee p_2(\mathbf{x}))}. \tag{19}$$

Proof. See Appendix B. □

B Proofs

Lemma A.1 (DeMorgan score calculus)**.** *Let ϕ and f be fully differentiable functions that generate the DeMorgan density operators $\{\wedge, \vee, \neg\}$. Then with $g = f \circ \phi$, $h : x \mapsto f(1 - \phi(x))$, $w(x) := x\, g'(x)$*

and $\overline{w}(x) := x\,h'(x)$ *the corresponding operations on the energies and scores are defined as*

$$\neg s(\mathbf{x}) = -\frac{\phi'(p(\mathbf{x}))p(\mathbf{x})}{\phi'(\neg p(\mathbf{x}))\neg p(\mathbf{x})}\, s(\mathbf{x}) \tag{17}$$

$$s_1(\mathbf{x}) \wedge s_2(\mathbf{x}) = \frac{w(p_1(\mathbf{x}))s_1(\mathbf{x}) + w(p_2(\mathbf{x}))s_2(\mathbf{x})}{w(p_1(\mathbf{x}) \wedge p_2(\mathbf{x}))} \tag{18}$$

$$s_1(\mathbf{x}) \vee s_2(\mathbf{x}) = \frac{\overline{w}(p_1(\mathbf{x}))s_1(\mathbf{x}) + \overline{w}(p_2(\mathbf{x}))s_2(\mathbf{x})}{\overline{w}(p_1(\mathbf{x}) \vee p_2(\mathbf{x}))}. \tag{19}$$

Proof of Lemma A.1. $\neg$

$$\begin{aligned}\neg s_1(\mathbf{x}) &= \nabla_{\mathbf{x}} \log \neg p(\mathbf{x}) && (20)\\ &= \frac{\nabla_{\mathbf{x}} \neg p(\mathbf{x})}{\neg p(\mathbf{x})} && (21)\\ &= \frac{\nabla_{\mathbf{x}} \phi^{-1}(1-\phi(p(\mathbf{x})))}{\phi^{-1}(1-\phi(p(\mathbf{x})))} && (22)\\ &= \frac{\nabla_{\mathbf{x}}(1-\phi(p(\mathbf{x})))}{\phi'(\phi^{-1}(1-\phi(p(\mathbf{x}))))\,\phi^{-1}(1-\phi(p(\mathbf{x})))} && (23)\\ &= \frac{-\phi'(p(\mathbf{x}))p(\mathbf{x})}{\phi'(\phi^{-1}(1-\phi(p(\mathbf{x}))))\,\phi^{-1}(1-\phi(p(\mathbf{x})))}\, s(\mathbf{x}) && (24)\\ &= \frac{-\phi'(p(\mathbf{x}))p(\mathbf{x})}{\phi'(\neg p(\mathbf{x}))\,\neg p(\mathbf{x})}\, s(\mathbf{x}) && (25)\end{aligned}$$

$\wedge$

$$\begin{aligned}s_1(\mathbf{x}) \wedge s_2(\mathbf{x}) &= \nabla_{\mathbf{x}} \log(p_1(\mathbf{x}) \wedge p_2(\mathbf{x})) && (26)\\ &= \frac{\nabla_{\mathbf{x}}(p_1(\mathbf{x}) \wedge p_2(\mathbf{x}))}{p_1(\mathbf{x}) \wedge p_2(\mathbf{x})} && (27)\\ &= \frac{\nabla_{\mathbf{x}} g^{-1}\big(g(p_1(\mathbf{x})) + g(p_2(\mathbf{x}))\big)}{p_1(\mathbf{x}) \wedge p_2(\mathbf{x})} && (28)\\ &= \frac{g'(p_1(\mathbf{x}))p_1(\mathbf{x})s_1(\mathbf{x}) + g'(p_2(\mathbf{x}))p_2(\mathbf{x})s_2(\mathbf{x})}{g'(p_1(\mathbf{x}) \wedge p_2(\mathbf{x}))\,\big(p_1(\mathbf{x}) \wedge p_2(\mathbf{x})\big)} && (29)\end{aligned}$$

$\vee$ Symmetric derivation with h instead of g.

We note that, if we can relate the ratios of the weights, we can give upper *and* lower bounds on the norm of the scores of compositions. □

Proposition 3.2. *Let* $\wedge_\lambda, \vee_\lambda$ *be the Dombi density operators. Then it holds that*

$$\forall x, y \in \mathbb{R}_{\geq 0} : \frac{\min\{x, y\}}{2^{1/\lambda}} \leq x \wedge_\lambda y \leq \min\{x, y\} \tag{8}$$

$$\forall x, y \in \mathbb{R}_{\geq 0} : \max\{x, y\} \leq x \vee_\lambda y \leq \max\{x, y\} 2^{1/\lambda} \tag{9}$$

Proof. We show the case for $p \vee_\lambda q = \big(p^\lambda + q^\lambda\big)^{1/\lambda}$ first. The definition of $\vee_\lambda$ is equivalent to that of a P-norm over two components. We have the standard inequality (w.l.o.g. for $p \geq q$)

$$p \vee_\lambda q = \big(p^\lambda + q^\lambda\big)^{1/\lambda} \leq \big(2p^\lambda\big)^{1/\lambda} = 2^{1/\lambda} \max\{p, q\} \tag{30}$$

The lower bound similarly follows from

$$p \vee_\lambda q = \big(p^\lambda + q^\lambda\big)^{1/\lambda} \geq \big(p^\lambda\big)^{1/\lambda} = \max\{p, q\} \tag{31}$$

For $\wedge_\lambda$, we can use DeMorgan to obtain the symmetric bounds. We can note that the upper bound is tight for $p = q$ and the lower bound is tight for $q = 0$. □

C Experiments

C.1 SAT Formulas

We use three different propositional formulas: majority, xor, and one-hot. The formulations of these formulas are designed to test different aspects of the score composition.

Majority We define the formula over k variables as

$$\mathrm{Maj}_k(P_1, \ldots, P_k) \equiv \bigwedge_{\substack{S \subseteq \{P_1, \ldots, P_k\} \\ |S| = \lceil k/2 \rceil}} \bigvee_{P \in S} .$$

This formula is negation-free, but might lead to mode dropping for variable assignments with fewer positive variables.

One-Hot We define a formula where exactly one variable has to be true as

$$\mathrm{OneHot}_k(P_1, \ldots, P_k) \equiv \Big(\bigvee_{i=1}^{k} P_i \Big) \wedge \Big(\bigwedge_{1 \leq i < j \leq k} (\neg P_i \vee \neg P_j) \Big).$$

It is only quadratic in the length of the variables, but it contains many clauses without positive literals, requiring precise handling of explicit negation.

Exclusive Or We define xor as a parity function over k variables as

$$\mathrm{XOR}_k(P_1, \ldots, P_k) \equiv \bigwedge_{\substack{v \in \{0,1\}^k \\ \sum_i v_i \equiv 0 \pmod 2}} \bigvee_{i=1}^{k} \big(v_i ? \neg P_i : \, P_i \big).$$

This formula can only be expressed in exponential length with 2^{k-1} clauses, which explicitly exclude one assignment with even parity.

C.2 Score Model Setup

We translate each of the 2^k propositional variable assignments to a Gaussian mode in $\mathbb{R}^k$ as

$$p(\mathbf{x}) = \frac{1}{2^k} \sum_{v \in \{0,1\}^k} \mathcal{N}_k(\mathbf{x} | 4v - 2, \sigma^2).$$

We then define "directional" diffusion models

$$\forall i \in [1, k] : p_i(\mathbf{x}) = \frac{1}{2^{k-1}} \sum_{\substack{v \in \{0,1\}^k \\ v_i = 1}} \mathcal{N}_k(\mathbf{x} | 4v - 2, \sigma^2).$$

In this setup, each distribution plays the role of one propositional variable. The distributions p_i can then be composed to mirror a propositional formula, with the goal that particles converge only to modes that correspond to satisfying variable assignments. We use p as an additional stabilizing model to guide particles to any location that corresponds to an assignment.

As these models are mixtures of Gaussians, we derive optimal scores and energy functions from the standard Gaussian to our distributions in closed form.

We then model each type of formula for $k \in [1, 10]$ as direct composition and simulate 2^{14} particles over 100 denoising steps.

For each mode, when then check a L_∞ bounding box around its mean of sidelength 3σ and consider all particles within that radius to be valid assignments.

In Figure 2b we show the most important metrics: η_{corr}, the fraction of particles within bounding boxes of satisfying modes, η_{all}, the fraction of particles converging to any mode. Additionally, we measure the normalized perplexity in the particle distributions across as PPL. In this experiment, PPL measures mode uniformity, where a higher number indicates more uniform samples from satisfying modes of the formula. In a formula with K satisfying variable assignments, for a batch of n particles, with $n\eta_{\text{corr}}$ particles within satisfying modes, we denote the fraction of particles within the bounding box of the *assignment index* $i \in [1, K]$ as η_i with $\sum_i \eta_i = \eta_{\text{corr}}$. We then calculate PPL for mode confusion as

$$\text{PPL} = 2^{\left(-\sum_{i=1}^{K} \frac{\eta_i}{\eta_{\text{corr}}} \log_2 \frac{\eta_i}{\eta_{\text{corr}}}\right)}/K.$$