DISCRETE DIFFUSION PRIORS FOR RAIN-FIELD RECONSTRUCTION FROM COMMERCIAL MICROWAVE LINKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Reconstructing rain fields from commercial microwave links (CMLs) requires inferring spatially distributed intensities from sparse, path-integrated attenuation measurements. Rainfall intermittency and spatial dependence make the prior central to this inverse problem. We introduce discrete diffusion priors for rain fields: masked and uniform diffusion models trained on fields encoded with an explicit dry class, which preserves rainfall occurrence exactly, and a codebook learned on positive intensities. We formulate CML reconstruction under three observation models (midpoint and extended virtual rain gauges, and a path-integrated likelihood) and benchmark conditioning methods: masked in-filling, CML adaptations of G2D2 and GILC, and D-MGPS, a new discrete variant of midpoint guidance. On the OpenMRG dataset, masked in-filling reaches an RMSE comparable to that of continuous diffusion priors at one to two orders of magnitude lower runtime and attains the highest wet/dry accuracy. With midpoint gauges, its 95% intervals are much narrower than kriging intervals and cover 0.94 of all locations, but only 0.42 of wet locations. Likelihood-guided discrete samplers, including D-MGPS, are currently less accurate than in-filling.

1 INTRODUCTION

Spatial estimates of rainfall support hydrological modeling, flood warning, and water-resource management. Rain gauges, weather radar, and satellites provide complementary observations with different spatial coverage and sources of uncertainty. Commercial microwave links (CMLs) offer an additional source of near-ground information by reusing cellular network infrastructure (Messer et al., 2006; Overeem et al., 2016). Rain attenuates the transmitted signal, so each link constrains an integral of a nonlinear function of rain intensity along its propagation path. Recovering a gridded field from a sparse, irregular network is therefore an underdetermined inverse problem: distinct spatial configurations can produce similar link measurements. A reconstruction method must combine these observations with a prior on rainfall occurrence, intensity, and spatial dependence.

Rainfall is intermittent and spatially structured: at each location, its distribution has a probability mass at zero and a continuous component over positive intensities. Classical stochastic models capture these features, for example, with transformed and censored latent Gaussian processes (Bárdossy & Plate, 1992; Ailliot et al., 2009). For CML reconstruction, interpolation methods provide simple estimates from virtual rain gauges, while stochastic methods generate fields subject to point observations and link-averaged rain-rate constraints (Eshel et al., 2021; Haese et al., 2017). An alternative is to learn the prior from reference rain fields and combine it with an explicit observation model.

Diffusion models provide such learned priors and support inference-time conditioning for inverse problems (Chung et al., 2023). Recent work has applied continuous diffusion priors to Bayesian rain-field reconstruction from CMLs (Moufad et al., 2026). Representing dry locations then requires care: an absolutely continuous output distribution assigns zero probability to exactly zero intensity, so an additional mechanism such as censoring is needed. A dedicated dry class, combined with a finite codebook for strictly positive intensities, provides an explicit alternative. Spatial dependence must still be learned, and quantizing intensities introduces a measurable approximation error.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

We study discrete diffusion priors on rain fields encoded by separating dry from wet cells and quantizing only the positive intensities, with levels learned by Lloyd's algorithm; the dry class preserves rainfall occurrence exactly in the encoded data. We consider three observation models: midpoint virtual rain gauges (VRGs), extended VRGs (E-VRGs) that assign each link's intensity to every cell it crosses, and a path-integrated Gaussian likelihood evaluated after dequantization. The first two yield categorical constraints, which masked diffusion enforces by in-filling; the third requires likelihood guidance. These procedures approximate the posterior, so the calibration of their uncertainty estimates must be assessed empirically.

Our contributions are threefold. **1** We introduce a discrete rain-field prior with an occurrence-preserving encoding, trained with masked (MDM) and uniform (UDM) diffusion (Austin et al., 2021). **2** We formulate CML reconstruction with these priors and benchmark conditioning methods: masked in-filling, CML adaptations of G2D2 (Murata et al., 2025) and GILC (Dou et al., 2026), and D-MGPS, a new discrete variant of midpoint guidance posterior sampling (MGPS) (Moufad et al., 2025) that fits a product distribution at an intermediate diffusion time and samples the next state through an exact categorical bridge. **3** We evaluate these methods on OpenMRG against meteorological and continuous-diffusion baselines in terms of reconstruction error, wet/dry accuracy, runtime, and (against kriging and MGPS) interval coverage. Masked in-filling attains the highest wet/dry accuracy and an RMSE comparable to that of continuous diffusion at one to two orders of magnitude lower runtime. At the 95% level, VRG in-filling reaches near-nominal coverage over all locations with intervals much narrower than those of ordinary kriging, but its coverage drops to 0.42 on wet locations. At the 50% and 75% levels, its intervals overcover across all locations and undercover on wet locations.

2 Discrete Diffusion Models for Rain Fields

This section presents our first contribution: discrete diffusion priors (masked and uniform) on rain fields encoded with an explicit dry class and a codebook learned on positive intensities. A rain field $X : \Omega \to \mathbb{R}_+ = [0, \infty)$ maps each location $\mathbf{s}$ of a spatial domain $\Omega \subset \mathbb{R}^2$ to a rain intensity $X(\mathbf{s})$ in mm/h. On an $H \times W$ grid with locations $\mathbf{s}_1, \dots, \mathbf{s}_L$ ($L = HW$), the flattened field is $x = (x^1, \dots, x^L) \in \mathbb{R}_+^L$ with $x^\ell = X(\mathbf{s}_\ell)$.

Rain-field quantization. The codebook $\mathcal{C} = \{a_0, \dots, a_{K-1}\}$ has $K \geq 2$ classes: the dry class $a_0 = 0$ and $K-1$ positive levels $0 < a_1 < \cdots < a_{K-1}$ for wet cells. The positive levels are chosen to minimize the mean squared quantization error of strictly positive intensities; the dry class takes no part in this fit (Appendix D):

$$(a_1^\star, \dots, a_{K-1}^\star) \in \operatorname*{arg\,min}_{a_1 < \cdots < a_{K-1}} \mathbb{E}\left[\min_{1 \leq k < K} |R - a_k|^2 \,\middle|\, R > 0 \right].$$

Here R is the intensity at a random location, and Lloyd iterations (MacQueen, 1967; Lloyd, 1982) minimize the empirical counterpart of this objective locally. Once the codebook is learned, the quantizer maps $x = (x^1, \dots, x^L)$ to class indices at each cell according to

$$z = \mathcal{Q}(x) \in \{0, \dots, K-1\}^L, \qquad z^\ell = \begin{cases} 0, & \text{if } x^\ell = 0, \\ \arg\min_{1 \leq k < K} |x^\ell - a_k|, & \text{if } x^\ell > 0. \end{cases}$$

Conversely, the dequantizer is a lookup table that maps class indices back to rain intensities as

$$\widetilde{x} = \mathcal{D}(z) \in \mathbb{R}_+^L, \qquad \widetilde{x}^\ell = a_{z^\ell}.$$

Hence $\widetilde{x}^\ell = 0$ if and only if $x^\ell = 0$, so the wet–dry pattern is preserved, including arbitrarily weak rainfall. Figure 1 shows that the quantization RMSE decreases as K increases.

Discrete diffusion models. We model the distribution p_0 of quantized fields $z_0 = \mathcal{Q}(x)$ with discrete diffusion (Austin et al., 2021), identifying class k with the one-hot vector e_k. For a nonincreasing schedule with $\alpha_0 = 1$, $0 < \alpha_t < 1$ for $0 < t < T$, and $\alpha_T = 0$, set $\alpha_{t|s} = \alpha_t / \alpha_s$ for $0 \leq s < t \leq T$. The forward process corrupts locations independently:

$$q_{t|s}(z_t \mid z_s) = \prod_{\ell=1}^{L} \mathrm{Cat}\big(z_t^\ell; \alpha_{t|s} z_s^\ell + (1 - \alpha_{t|s})\bar{q}\big).$$

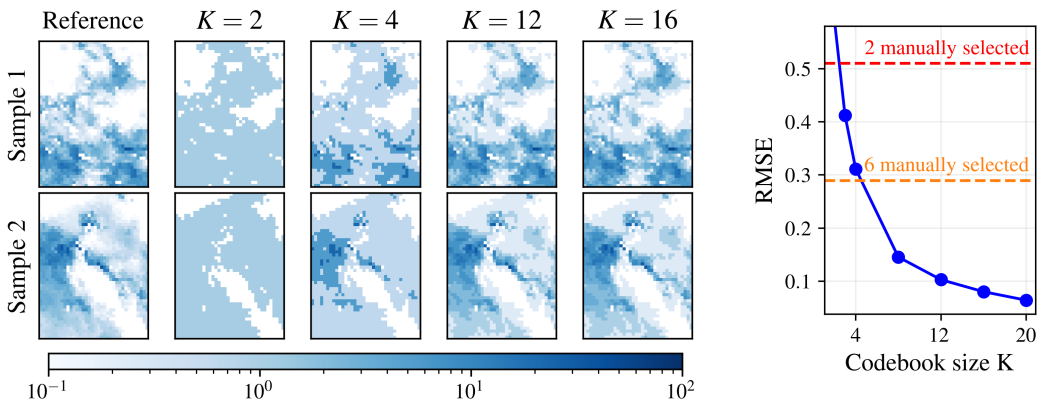

Figure 1: Effect of rain-intensity codebook size. **Left:** Reference rain fields and their dequantized versions using $K \in \{2, 4, 12, 16\}$ learned codebooks. Each row is one rain field. **Right:** Quantization RMSE across 1,000 rain fields as a function of codebook size. Dashed lines show the RMSE of manually specified codebooks with $K = 2$ and $K = 6$ using quantiles of intensities.

Uniform diffusion (UDM) uses $\bar{q} = \mathbf{1}/K$, whereas masked diffusion (MDM) adds an absorbing class `M` with $\bar{q} = e_{\mathtt{M}}$ and extends clean vectors by a zero mask coordinate (Sahoo et al., 2024; Shi et al., 2024). A denoiser predicts clean-class probabilities $P_\ell^\theta(z_t, t)$ at each cell ℓ from the complete corrupted field. It defines the factorized approximation $p_{0|t}^\theta(z_0 \mid z_t) = \prod_\ell \mathrm{Cat}(z_0^\ell; P_\ell^\theta(z_t, t))$ and the learned reverse transition

$$p_{s|t}^\theta(z_s \mid z_t) = \sum_{z_0} q_{s|0,t}(z_s \mid z_0, z_t)\, p_{0|t}^\theta(z_0 \mid z_t),$$

where the bridge $q_{s|0,t}$, the law of z_s given z_0 and z_t under the forward process, is available in closed form and factorizes across cells; see Appendix E. Hence $p_{s|t}^\theta(\cdot \mid z_t)$ factorizes across cells and can be sampled exactly, while successive transitions can still produce spatially dependent fields. In masked diffusion, the denoiser probability of an unmasked cell's class is set to one, so the cell remains unchanged; in uniform diffusion, cells can change class throughout reverse sampling.

Training minimizes a weighted denoising cross-entropy which regresses the corrupted field z_t to its clean counterpart z_0; see Appendix E.4. Sampling draws independent terminal tokens from $\bar{q}$ and applies the learned transitions along a decreasing time grid. We denote the resulting prior by p_0^θ and dequantize its samples with $\mathcal{D}$; Figure 2 shows samples from both priors.

3 Posterior Sampling from CMLs

Problem formulation. CMLs measure signal attenuation between transmitters and receivers. After isolating the rain-induced component, we model the noiseless attenuation along link $\mathcal{L}_i$ as $y_i = \kappa_i \int_{\mathcal{L}_i} X(\mathbf{s})^{b_i}\, \mathrm{d}\ell$, which integrates the specific-attenuation power law (International Telecommunication Union, 2005) with coefficients $\kappa_i > 0$ and $b_i > 0$ that depend on the link characteristics and are assumed constant along the link. The operator is nonlinear when $b_i \neq 1$, and sparse link measurements do not uniquely determine the field. Let $z_0 \sim p_0^\theta$ and $x_0 = \mathcal{D}(z_0)$. Given a likelihood $p(y \mid z_0)$, we define the reward $r_y(z_0) = \log p(y \mid z_0)$, up to an additive constant independent of z_0. By Bayes' rule, the target posterior is

$$\pi_0^\theta(z_0 \mid y) \propto p_0^\theta(z_0) \exp(r_y(z_0)).$$

It is a posterior over quantized fields under the chosen observation model: reconstruction proceeds by sampling quantized fields from $\pi_0^\theta(\cdot \mid y)$ and dequantizing them.

Observation models. A midpoint virtual rain gauge (VRG) represents each link by an intensity at its center (Eshel et al., 2021; Zhang et al., 2023). For a link of length $d_i = |\mathcal{L}_i|$, assuming constant rain along the link gives $R_i^{\mathrm{VRG}} = (y_i/(\kappa_i d_i))^{1/b_i}$; for spatially varying rain, this is an effective

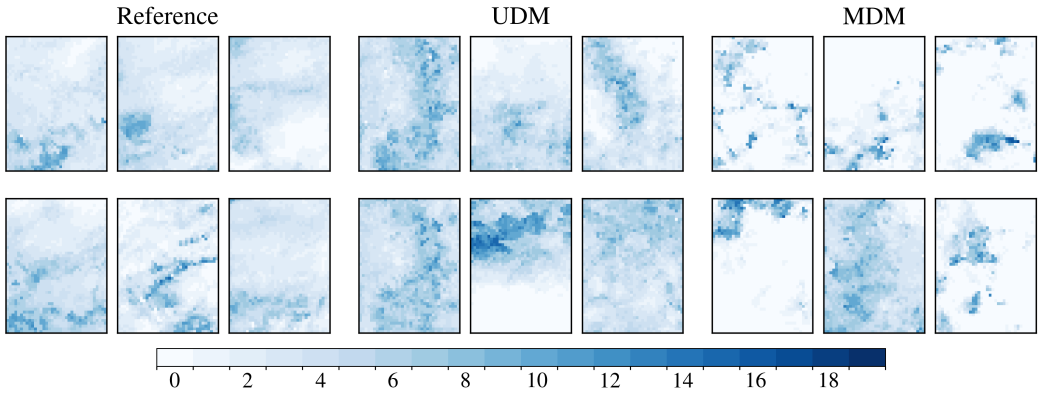

Figure 2: Quantized reference rain fields (left) and samples from the uniform (center) and masked (right) diffusion priors, trained with codebook size $K = 20$; the quantitative experiments use $K \in \{4, 8, 20\}$. The colorbar indicates class indices (class 0 is the dry class).

intensity that generally differs from the arithmetic path average when $b_i \neq 1$. Each midpoint is assigned to its nearest grid cell and its intensity is quantized. With $I \subseteq \{1, \dots, L\}$ the observed cells and y_{VRG} their classes, the noiseless likelihood is $p(y_{\mathrm{VRG}} \mid z_0) = \mathbf{1}\{z_0^I = y_{\mathrm{VRG}}\}$, provided that no cell receives conflicting labels; otherwise, an aggregation rule or a noisy likelihood is required. An extended VRG (E-VRG) assigns the class $y^{i,\ell}_{\mathrm{EVRG}} = \mathcal{Q}(R_i^{\mathrm{VRG}})$ to every cell $\ell \in J_i$ traversed by link i, with likelihood $p(y_{\mathrm{EVRG}} \mid z_0) = \prod_i \prod_{\ell \in J_i} \mathbf{1}\{z_0^\ell = y^{i,\ell}_{\mathrm{EVRG}}\}$. It retains the spatial extent of each path but assumes constant rain along it. Both models produce pseudo-observations: their hard constraints do not enforce the path-integrated attenuation, and E-VRGs impose stronger assumptions than VRGs without necessarily being more accurate.

The *discretized path-integrated operator* instead treats rain as piecewise constant on the grid: $[\mathcal{M}_{\mathrm{CML}}(x_0)]_i = \kappa_i \sum_{\ell=1}^{L} \Delta_\ell^i (x_0^\ell)^{b_i}$, where Δ_ℓ^i is the intersection length of link $\mathcal{L}_i$ with cell ℓ, which can be computed with the Siddon algorithm (Siddon, 1985); for links inside the grid, $\sum_\ell \Delta_\ell^i = d_i$. For m links, we assume independent Gaussian errors, $y = \mathcal{M}_{\mathrm{CML}}(\mathcal{D}(z_0)) + \Sigma\varepsilon$ with $\varepsilon \sim \mathcal{N}(0, I_m)$ and fixed $\Sigma = \mathrm{diag}(\sigma_1, \dots, \sigma_m)$, $\sigma_i > 0$, which gives

$$p(y \mid z_0) \propto \exp\left\{-\frac{1}{2} \left\| y - \mathcal{M}_{\mathrm{CML}}(\mathcal{D}(z_0)) \right\|^2_{\Sigma^{-2}}\right\}, \tag{1}$$

with $\|v\|^2_{\Sigma^{-2}} = v^\top \Sigma^{-2} v$. This model retains the path geometry and allows rain to vary along a link. The Gaussian error model is a modeling choice: the covariance Σ^2 does not automatically account for quantization or preprocessing errors.

Conditioning methods. We benchmark three types of conditioning methods for the discrete priors, all of which keep the pretrained model fixed: (a) in-filling, (b) CML adaptations of G2D2 and GILC, and (c) D-MGPS, a new method described below. (a) *VRG and E-VRG in-filling.* With a masked diffusion prior, compatible noiseless VRG or E-VRG constraints require no likelihood guidance. We initialize the observed cells with their prescribed tokens, mask all other cells, start the reverse process at the time whose masking probability matches the fraction of masked cells, and reveal the missing tokens while keeping the observations fixed. This enforces the categorical constraints exactly but does not by itself guarantee exact sampling from $\pi_0^\theta(\cdot \mid y)$.

(b) *G2D2 and GILC.* The path-integrated likelihood does not prescribe individual tokens and requires explicit guidance. We adapt two zero-shot guidance methods for discrete diffusion to the reward r_y obtained from Equation (1). At reverse time t, G2D2 (Murata et al., 2025) repeatedly updates a distribution that factorizes over clean tokens to minimize the sum of the expected negative log-likelihood under this distribution and a KL penalty relative to the denoiser prediction $p^\theta_{0|t}(\cdot \mid z_t)$; Appendix G derives its CML adaptation and explains the scope of its theoretical error bound. GILC (Dou et al., 2026) averages the reward under the denoiser's predicted clean-field distribution, corrects the predicted class scores (logits) once along the gradient of this average, and

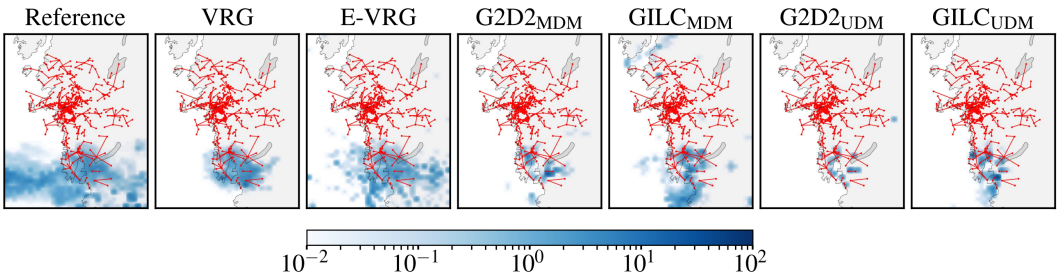

Figure 3: One OpenMRG test field: reference (left) and sample reconstructions by six discrete methods (VRG and E-VRG in-filling with MDM; G2D2 and GILC with MDM and UDM). Red segments denote the CML network. Rain rate (mm/h) on a logarithmic color scale.

then samples the next reverse state. For the Gaussian likelihood, this average and its gradient can be computed exactly (Appendices F and H); Algorithm 7 gives the masked sampler.

(c) D-MGPS. We also introduce discrete midpoint guidance posterior sampling (D-MGPS), which adapts the variational midpoint construction of MGPS (Moufad et al., 2025) to categorical states. Whereas G2D2 fits a product distribution over clean fields, D-MGPS fits one over the state at an adjustable intermediate time, which may be the clean endpoint, and then samples the next state from an exact categorical bridge. For a reverse step $t \to s$, choose $0 \le u \le s < t$ and let $v = (v^1, \dots, v^L)$ be a value of Z_u, the grid state at time u; in masked diffusion, v^ℓ can be the mask when $u > 0$. Let $P^t_{\ell k} = P^\theta_{\ell k}(z_t, t)$ be the denoiser probability of clean class k at cell ℓ given the current field z_t; in masked diffusion, a revealed cell keeps its class with probability one. Since the denoiser prediction factorizes over cells given z_t and the forward bridge acts cellwise, the reference distribution also factorizes:

$$\rho_{u|t}(v \mid z_t) = \sum_{z_0} q_{u|0,t}(v \mid z_0, z_t)\, p^\theta_{0|t}(z_0 \mid z_t) = \prod_{\ell=1}^{L} \rho^\ell_{u|t}(v^\ell \mid z_t), \tag{2}$$

with $\rho^\ell_{u|t}(v^\ell \mid z_t) = \sum_k P^t_{\ell k}\, q_{u|0,t}(v^\ell \mid e_k, z^\ell_t)$, where $q_{u|0,t}$ denotes the single-cell bridge and zero-weight terms are omitted. Each factor still depends on the whole field z_t through P^t; Appendix I.2 gives the calculation and distinguishes this reference from the exact prior conditional. The reference is computed once per reverse step and kept fixed during the fit. For $u > 0$, the reward is averaged under a new denoiser prediction at (v, u), $\mathcal{R}_u(v) = \mathbb{E}_{p^\theta_{0|u}(\cdot|v)}[r_y(Z_0)]$; at $u = 0$, $\mathcal{R}_0(v) = r_y(v)$. With a guidance strength $\lambda_t \ge 0$, we fit a product distribution $q_\phi(v) = \prod_\ell V^\phi_\ell(v^\ell)$, parameterized by logits ϕ (Appendix I.3), by minimizing

$$\mathcal{F}_{u,t}(\phi) = \mathrm{KL}\big(q_\phi \,\big\|\, \rho_{u|t}(\cdot \mid z_t)\big) - \lambda_t \mathbb{E}_{q_\phi}[\mathcal{R}_u(Z_u)], \tag{3}$$

which is equivalent to minimizing $\mathrm{KL}(q_\phi \,\|\, \nu_{u|t})$, where $\nu_{u|t} \propto \rho_{u|t}(\cdot \mid z_t) \exp(\lambda_t \mathcal{R}_u)$. After fitting, we draw $v \sim q_\phi$ using new random draws rather than optimization samples, and then $z_s \sim q_{s|u,t}(\cdot \mid v, z_t)$ from the exact categorical bridge, which returns v when $u = s$; Appendix I derives the underlying midpoint identity.

For uniform diffusion, each factor of q_ϕ is supported on the classes with positive reference probability. For masked diffusion, revealed tokens remain fixed; at a masked cell, we fix the midpoint mask probability to $(1-\alpha_u)/(1-\alpha_t)$ and optimize only the class probabilities conditional on unmasking, which preserves the unguided probability of revealing the cell at each reverse step while letting the observations affect its class. We initialize $q_\phi = \rho_{u|t}(\cdot \mid z_t)$, so that $\lambda_t = 0$ recovers the unguided reverse step for every u; sampling starts from $z_T \sim \bar{q}^{\otimes L}$ and ends with $s = u = 0$ and a clean field. For the Gaussian likelihood, $\mathcal{R}_u(v)$ has a closed form for each v (Appendix F), and the KL term is exact. For $u > 0$, the gradient of the reward term is estimated with Gumbel–Softmax or straight-through relaxations through the midpoint denoiser, or with a score-function (REINFORCE) estimator that needs no input derivatives; Appendix F.2 gives a common derivation and algorithms for G2D2, GILC, and D-MGPS, including the relaxed reward extensions and their bias. For $u = 0$, an analytic alternative minimizes the objective exactly over one cell's class probabilities at a time;

these coordinate updates cannot increase the objective (Appendix G.3). Algorithm 8 specifies the complete sampler with either fit.

Like G2D2 and GILC, D-MGPS is not guaranteed to sample exactly from $\pi_0^\theta(\cdot \mid y)$, although its bridge transitions and reward evaluations are exact. It replaces the exact prior conditional of Z_u by the denoiser-based reference, restricts q_ϕ to a product family, and uses finitely many optimization steps; for $u > 0$, it also replaces the exact midpoint weight, the likelihood $p(y \mid Z_0)$ averaged over the clean-field distribution conditional on $Z_u = v$, by the exponential of the expected log-likelihood under the denoiser prediction at (v, u) (Appendices I.1–I.2). At $u = 0$, D-MGPS uses the same CML objective and updates as our G2D2 adaptation, and the complete step agrees with the Markov variant of G2D2 when all fitting choices agree (Appendix G). Starting from the reference at $u = 0$, a single gradient step on the logits recovers our analytic CML version of GILC, because the KL gradient vanishes at this initialization.

4 Related Work

CML rain-field reconstruction. CMLs support rain-rate retrieval and countrywide rainfall mapping (Leijnse et al., 2007; Overeem et al., 2013; Chwala & Kunstmann, 2019). Existing methods combine attenuation measurements with spatial assumptions via virtual-gauge interpolation (Eshel et al., 2021; Zhang et al., 2023), iterative reconstruction (Goldshtein et al., 2009), or conditional stochastic fields (Haese et al., 2017; Blettner et al., 2022). Stochastic methods already represent spatial uncertainty, so ensemble reconstruction is not specific to diffusion models. Recent work uses continuous diffusion priors with explicit CML likelihoods (Moufad et al., 2026) or trains a diffusion model conditioned on CML observations (Ji et al., 2026). Like the former, we apply the observation model only at inference time, but with a categorical prior and an occurrence-preserving encoding.

Discrete diffusion. Discrete diffusion uses categorical transition matrices (Austin et al., 2021), including uniform ones (Hoogeboom et al., 2021) and ones with an absorbing mask state (Sahoo et al., 2024; Shi et al., 2024). Other formulations use continuous-time Markov chains (Campbell et al., 2022); Lou et al. (2024) learn ratios of state probabilities as discrete scores, and Zhao et al. (2025) relate continuous- and discrete-time formulations. Uniform diffusion can change a class at later reverse steps, whereas masked diffusion keeps a revealed class fixed.

Guidance for discrete diffusion. Continuous-diffusion posterior samplers include diffusion posterior sampling (DPS) (Chung et al., 2023), MGPS (Moufad et al., 2025), and the twisted diffusion sampler (TDS) (Wu et al., 2023). Several discrete guidance methods require a trained predictor, conditional model, or guidance network (Nisonoff et al., 2025; Schiff et al., 2025; Wan et al., 2026) and therefore require training for each new measurement model. Zero-shot alternatives condition a fixed prior: G2D2 fits a clean-token distribution trading off an expected log-likelihood against a KL penalty to the denoiser prediction (Murata et al., 2025), and GILC corrects the denoiser logits along the gradient of an expected reward (Dou et al., 2026), here the observation log-likelihood. D-MGPS, which we introduce and evaluate alongside both, adapts the MGPS midpoint construction to categorical states. All three apply the observation at each reverse step through a product distribution, which cannot represent the dependence between cells induced by a path-integrated observation. Appendix J extends this review.

5 Experiments

Experimental setting. We use OpenMRG (Andersson et al., 2022; Fencl et al., 2024), which provides synchronized measurements from 364 bidirectional CMLs, weather radar, and rain gauges in Gothenburg, Sweden, from June to August 2015. Following common practice (Michelson & Koistinen, 2000), the reference fields are radar-derived rain-rate fields adjusted with rain-gauge measurements. After preprocessing, the dataset contains $26{,}495$ fields of size $H \times W = 48 \times 36$, which we split chronologically into training (80%) and test (20%) sets. We simulate CML observations from the reference fields on the OpenMRG network, following the protocol of Moufad et al. (2026). Appendix A details the preprocessing, architectures, and training.

Table 1: Rain-field reconstruction by discrete-diffusion ($K = 20$, 100 reverse steps), continuous-diffusion, and meteorological methods: means over test problems with 95% confidence intervals, following Moufad et al. (2026). Cum. rain-rate diff.: sum over grid cells of reconstructed minus reference rain rate (mm/h); negative values indicate underestimation. The fraction of wet locations in the test set is 20%. Runtime: seconds per reconstruction; diffusion methods generate 10 reconstructions per batch, and the batch runtime is divided by 10. Arrows indicate preferred values; dark and light shading mark the best and second-best values.

	Method	RMSE ↓	Wet/dry accuracy ↑	PCC ↑	Cum. rain-rate diff. → 0	Runtime (s) ↓
Masked diffusion	VRG in-filling	0.84 ± 0.09	0.89 ± 0.01	0.26 ± 0.02	-135.96 ± 27.60	0.12
	E-VRG in-filling	0.85 ± 0.09	0.90 ± 0.01	0.24 ± 0.01	-136.16 ± 23.80	0.14
	G2D2	0.94 ± 0.09	0.83 ± 0.01	0.15 ± 0.01	-490.96 ± 69.63	0.66
	GILC	1.14 ± 0.10	0.79 ± 0.01	0.10 ± 0.01	-177.26 ± 59.35	0.33
	D-MGPS	1.01 ± 0.10	0.81 ± 0.01	0.07 ± 0.01	-606.89 ± 86.49	2.11
Uniform diffusion	G2D2	0.95 ± 0.09	0.82 ± 0.01	0.14 ± 0.01	-423.81 ± 65.22	0.65
	GILC	0.93 ± 0.09	0.86 ± 0.01	0.15 ± 0.01	-364.01 ± 46.59	0.30
	D-MGPS	0.99 ± 0.10	0.81 ± 0.01	0.06 ± 0.01	-531.26 ± 77.47	2.51
Continuous diffusion	MGPS	0.86 ± 0.08	0.62 ± 0.01	0.24 ± 0.01	42.86 ± 34.65	5.22
	TDS	0.84 ± 0.09	0.86 ± 0.01	0.21 ± 0.02	-216.44 ± 35.89	18.49
Meteorological	IDW	0.90 ± 0.09	0.38 ± 0.01	0.21 ± 0.01	-144.30 ± 32.26	0.01
	GMZ	0.92 ± 0.09	0.34 ± 0.01	0.21 ± 0.01	-65.87 ± 29.26	1.01
	OK	0.87 ± 0.09	0.22 ± 0.02	0.23 ± 0.02	158.66 ± 38.67	0.27

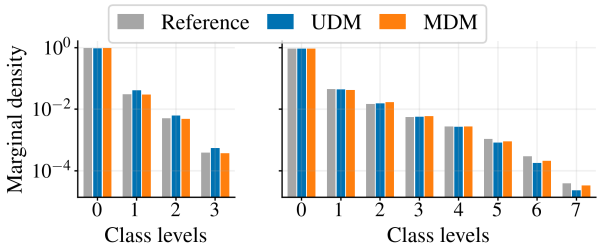

Codebook size	Prior	Sampling steps 20	50	100	400	800
$K = 4$	MDM	0.88	0.75	0.66	0.58	0.55
	UDM	0.60	0.58	0.57	0.58	0.59
$K = 8$	MDM	0.94	0.83	0.75	0.64	0.62
	UDM	0.64	0.62	0.60	0.61	0.61
$K = 20$	MDM	0.95	0.86	0.78	0.67	0.65
	UDM	0.69	0.65	0.64	0.65	0.64

Figure 4: Discrete diffusion priors versus reference data. **Left**: marginal class frequencies (log scale) of reference fields and of samples from uniform (UDM) and masked (MDM) diffusion priors, for $K = 4$ (left panel) and $K = 8$ (right panel). **Right**: C2ST accuracy for codebook sizes $K \in \{4, 8, 20\}$ and numbers of sampling steps from 20 to 800; values closer to 0.5 mean generated fields are harder to distinguish from reference fields.

Discrete diffusion priors. We train MDM and UDM priors with the same U-Net architecture (Ronneberger et al., 2015) and training configuration (Appendix A.2); each model has approximately one million parameters and trains in approximately 30 minutes on a single NVIDIA L40S GPU. Unless stated otherwise, the codebook size is $K = 20$. We assess prior fidelity with the classifier two-sample test (C2ST) (Lopez-Paz & Oquab, 2017): we pool generated and reference fields into a labeled dataset, divide it into five folds, train a classifier on four folds, evaluate it on the held-out fold, and report the mean accuracy over folds. An accuracy near 0.5 indicates that the classifier cannot distinguish generated fields from reference fields. For $K \in \{4, 8\}$, both priors closely match the reference class frequencies, including the large mass of the dry class, with small deviations in the wet classes, largest for UDM (Figure 4, left). For $K = 20$, UDM reaches C2ST accuracies of 0.64–0.69 for all step counts, whereas MDM accuracy decreases from 0.95 at 20 steps to 0.65 at 800 steps (Figure 4, right). UDM is thus much closer to 0.5 with few sampling steps, and the two priors become comparable at 400–800 steps.

Rain-field reconstruction with CMLs. We compare discrete-diffusion methods with meteorological and continuous-diffusion baselines. The meteorological baselines are inverse-distance weighting (IDW), the Goldshtein–Messer–Zinevich method (GMZ) (Goldshtein et al., 2009), and ordinary kriging (OK) (Eshel et al., 2021); the continuous-diffusion baselines are MGPS (Moufad et al., 2025) and the twisted diffusion sampler (TDS) (Wu et al., 2023). The discrete methods are VRG and E-VRG in-filling with the MDM prior, and three samplers guided by the path-integrated likelihood, run with both priors: our CML adaptations of G2D2 (Murata et al., 2025) and GILC (Dou

Table 2: Empirical marginal coverage and mean interval width (mm/h) at nominal levels of 50%, 75%, and 95%, over 1,000 test problems with 500 reconstructions each. All-location results include dry locations; wet-location results include only locations with positive reference rain rate. Dark and light shading mark the best and second-best values in each column, including ties. Coverage is ranked by absolute deviation from the nominal level; smaller interval width is preferred. D-MGPS, TDS, IDW, and GMZ are not reported.

		All locations						Wet locations					
		Marginal coverage			Interval width ↓			Marginal coverage			Interval width ↓		
	Method	50%	75%	95%	50%	75%	95%	50%	75%	95%	50%	75%	95%
Masked diffusion	VRG in-filling	0.85	0.89	0.94	0.17	0.33	0.81	0.13	0.23	0.42	0.32	0.61	1.45
	E-VRG in-filling	0.84	0.87	0.91	0.16	0.30	0.72	0.11	0.19	0.35	0.29	0.54	1.24
	G2D2	0.82	0.83	0.88	0.07	0.14	0.39	0.03	0.05	0.22	0.14	0.27	0.72
	GILC	0.83	0.88	0.97	0.16	0.44	2.08	0.06	0.32	0.84	0.27	0.70	3.04
Uniform diffusion	G2D2	0.82	0.84	0.94	0.06	0.14	0.74	0.02	0.06	0.69	0.10	0.23	0.98
	GILC	0.83	0.85	0.90	0.14	0.25	0.58	0.05	0.08	0.23	0.24	0.46	1.12
Continuous diffusion	MGPS	0.72	0.88	0.95	0.29	0.50	1.00	0.27	0.49	0.77	0.38	0.66	1.28
Meteorological	OK	0.39	0.57	0.72	2.31	3.95	6.72	0.47	0.61	0.73	2.24	3.82	6.51

et al., 2026), and D-MGPS. All discrete methods use $K = 20$ and, unless stated otherwise, 100 reverse steps (Appendix B). We compare the reconstructions (dequantized for discrete methods) with the reference fields using RMSE, wet/dry accuracy (the fraction of grid cells whose wet/dry status matches the reference), the Pearson correlation coefficient (PCC), and the cumulative rain-rate difference (the sum over grid cells of the reconstructed minus the reference rain rate, in mm/h; negative values indicate underestimation), as defined in Appendix C. For each problem, we average 10 generated samples before computing the metrics. For IDW and GMZ, we use their single reconstructions; for OK, we use its posterior mean.

Results. Table 1 reports the results, and Figure 3 shows discrete reconstructions of one test field. VRG in-filling has the lowest mean RMSE (0.84, tied with TDS), followed by E-VRG in-filling (0.85); the differences from continuous diffusion (0.84–0.86) and from the meteorological baselines (0.87–0.92) are within the 95% confidence intervals. VRG in-filling also has the highest mean PCC (0.26, versus 0.24 for E-VRG in-filling and MGPS). The two in-filling variants have the highest wet/dry accuracy (0.89–0.90), compared with 0.86 for TDS, 0.62 for MGPS, and 0.22–0.38 for the meteorological baselines. The dry class lets the discrete methods represent dry areas explicitly, which this metric rewards; among continuous methods, TDS nevertheless reaches 0.86, above G2D2 (0.82–0.83) and GILC with MDM (0.79). All discrete methods underestimate total rain, as do TDS, IDW, and GMZ, whereas MGPS and OK overestimate it; the underestimation is larger for GILC (-177 and -364) and much larger for G2D2 (-424 and -491) than for in-filling (-136). G2D2 and GILC are also less accurate than in-filling, most clearly in PCC (0.10–0.15; RMSE 0.93–1.14); UDM improves GILC (RMSE 0.93 versus 1.14), whereas G2D2 performs similarly with both priors. D-MGPS currently performs below G2D2 and GILC (Table 1): its RMSE (0.99 with UDM, 1.01 with MDM) exceeds that of all other methods except GILC (MDM), its PCC (0.06–0.07) is the lowest, and it underestimates total rain the most (-531 and -607). Its wet/dry accuracy (0.81 with both priors) exceeds only that of GILC (MDM) among the discrete methods.

Uncertainty quantification. For each of 1,000 test problems, we draw 500 reconstructions and form location-wise intervals at nominal levels of 50%, 75%, and 95%. Table 2 reports empirical marginal coverage (the fraction of locations whose interval contains the reference rain rate, averaged over problems) and mean interval width, both over all locations and restricted to locations with positive reference rain rate (Appendix C). The OK baseline provides posterior means and standard deviations. For a consistent comparison with the other methods, we draw 500 samples from its Gaussian process posterior. Over all locations, VRG in-filling and G2D2 (UDM) attain 0.94 coverage at the 95% level with mean widths of 0.81 and 0.74 mm/h, respectively; MGPS attains 0.95 with width 1.00 mm/h. In contrast, wet-location coverage at this level is 0.42, 0.69, and 0.77, respectively, with wider intervals (1.45, 0.98, and 1.28 mm/h). Across the discrete methods, wet-location coverage at 95% ranges from 0.22 for G2D2 (MDM) to 0.84 for GILC (MDM). At the 50% and 75% levels, all discrete methods overcover when averaged over all locations (0.82–0.89), but undercover on wet locations (0.02–0.32). Dry reference locations can inflate all-location coverage: an interval

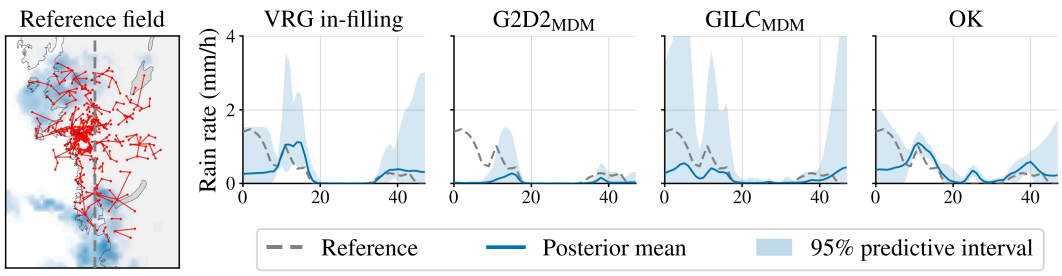

Figure 5: Uncertainty along a central vertical slice of one test field. **Left**: reference field, CML network (red segments), and slice location (vertical dashed line). **Other panels**: reference rain rate (dashed), posterior mean (solid), and 95% predictive interval (shaded) along the slice for VRG in-filling, G2D2, GILC, and OK. The x-axis is the grid index along the slice.

concentrated at zero contains a zero reference at every nominal level. OK undercovers at 95% both over all locations (0.72) and on wet locations (0.73), despite mean widths of 6.72 and 6.51 mm/h, respectively. Figure 5 shows posterior means and 95% intervals along a slice of one test field.

Computational cost. The last column of Table 1 reports the runtime per reconstruction. Diffusion methods generate 10 reconstructions per problem in one batch, and we report the batch runtime divided by 10 for comparison with the meteorological methods, which produce a single reconstruction. All runtimes were measured on the same NVIDIA L40S GPU. VRG in-filling takes 0.12 s per reconstruction, compared with 5.22 s for MGPS and 18.49 s for TDS. Per reconstruction, in-filling, G2D2, and GILC (0.12–0.66 s) are within the runtime range of the meteorological baselines (0.01–1.01 s), whereas D-MGPS takes 2.11–2.51 s, still less than MGPS and TDS.

Sensitivity analyses. *Priors.* The C2ST results for $K = 4$ and $K = 8$ follow the same pattern as for $K = 20$ (Figure 4, right): UDM accuracy changes little with the number of sampling steps and is much closer to 0.5 than MDM accuracy up to 100 steps, whereas MDM improves steadily and becomes comparable at 400–800 steps; for $K = 4$, MDM is even slightly closer to 0.5 at 800 steps (0.55 versus 0.59). For both priors, C2ST accuracy increases with K, which is consistent with finer codebooks being harder to learn; since each K defines a different discretization, this comparison is only indicative. *Reconstruction.* In-filling is insensitive to the number of reverse steps (Table 5), with RMSE 0.84–0.86 and cumulative rain-rate differences between -148 and -134; the RMSE of G2D2 and D-MGPS varies by at most 0.03, and GILC is stable with UDM (0.93–0.94) but improves with MDM, from 1.21 at 50 steps to 1.04 at 400 steps.

6 Conclusion

We introduced discrete diffusion priors for CML-based rain-field reconstruction. An explicit dry class preserves rainfall occurrence, while a codebook for positive intensities enables evaluation in physical units. On OpenMRG, masked in-filling achieves the highest wet/dry accuracy and an RMSE comparable to that of continuous diffusion at one to two orders of magnitude lower runtime. At the 95% level, VRG in-filling and G2D2 with UDM attain near-nominal coverage over all locations with narrow intervals, but cover only 0.42 and 0.69 of wet locations, respectively. At lower nominal levels, the discrete methods overcover across all locations and undercover on wet locations. Likelihood-guided discrete samplers remain less accurate than in-filling and, at 100 steps, underestimate total rainfall more.

Limitations and future work. Quantization distorts positive intensities and caps them at the largest codebook level. VRG constraints do not ensure agreement with link attenuation, and all conditioning methods sample the posterior only approximately. All methods with reported coverage undercover on wet locations at the 95% level, and all-location coverage can mask this deficit. Our study uses a single dataset; future work should improve likelihood guidance and evaluate other link densities, error levels, locations, and seasons.

AI USE STATEMENT

We used generative AI tools (large language models, including an AI coding assistant) in preparing this paper. They were used to edit and reorganize the manuscript: language editing, restructuring of sections, and drafting text, including the related-work section and the ethics and reproducibility statements, from the authors' notes, derivations, and results. They were also used to check the consistency of notation, cross-references, and bibliography entries; to cross-check the numbers reported in the text against the authors' experimental outputs; to review the manuscript for unsupported or inconsistent claims; and to write scripts that numerically verify analytic formulas in the appendices. Generative AI tools were not used to run the experiments or to generate or modify experimental results. The authors reviewed, verified, and edited all AI-assisted content, and take full responsibility for the final content of this paper, including text, claims, and artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work uses the publicly available OpenMRG dataset of CML, weather-radar, and rain-gauge measurements, which contains no personal data. Rain-field reconstruction supports hydrological applications such as flood warning and water-resource management. The reconstructions and uncertainty estimates reported here have not been validated for operational use; in particular, the methods can underestimate total rainfall and intensity extremes, and relying on them without such validation could mislead decisions.

REPRODUCIBILITY STATEMENT

All experiments use the public OpenMRG dataset. Appendix A describes data preprocessing, model architectures, and training, with hyperparameters in Table 3. Appendix B and Table 4 describe all reconstruction methods and their settings, and Appendix C defines the evaluation metrics. The quantizer, the diffusion bridges, and the conditioning methods are specified in Appendices D–I, including Algorithms 7 and 8.

REFERENCES

Pierre Ailliot, Craig Thompson, and Peter Thomson. Space–time modelling of precipitation by using a hidden Markov model and censored Gaussian distributions. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 58(3):405–426, 2009.

Jafet C. M. Andersson, Jonas Olsson, Remco C. Z. van de Beek, and Jonas Hansryd. OpenMRG: Open data from microwave links, radar, and gauges for rainfall quantification in Gothenburg, Sweden. *Earth System Science Data*, 14(12):5411–5426, 2022. doi: 10.5194/essd-14-5411-2022.

Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

András Bárdossy and Erich J Plate. Space-time model for daily rainfall using atmospheric circulation patterns. *Water Resources Research*, 28(5):1247–1259, 1992.

Nico Blettner, Christian Chwala, Barbara Haese, Sebastian Hörning, and Harald Kunstmann. Combining commercial microwave link and rain gauge observations to estimate countrywide precipitation: A stochastic reconstruction and pattern analysis approach. *Water Resources Research*, 58 (10):e2022WR032563, 2022. doi: 10.1029/2022WR032563.

Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 28266–28279, 2022. doi: 10.52202/068431-2049. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/b5b528767aa35f5b1a60fe0aaeca0563-Abstract-Conference.html.

Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T. Freeman. MaskGIT: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11315–11325, 2022. doi: 10.1109/CVPR52688.2022.01103.

Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023.

Christian Chwala and Harald Kunstmann. Commercial microwave link networks for rainfall observation: Assessment of the current status and future challenges. *WIREs Water*, 6(2):e1337, 2019. doi: 10.1002/wat2.1337.

Hongkun Dou, Zike Chen, Fengji Li, Hongjue Li, and Yue Deng. Plug-and-play guidance for discrete diffusion models via gradient-informed logit correction. *arXiv preprint arXiv:2606.06303*, 2026. doi: 10.48550/arXiv.2606.06303.

Adam Eshel, Hagit Messer, Harald Kunstmann, Pinhas Alpert, and Christian Chwala. Quantitative analysis of the performance of spatial interpolation methods for rainfall estimation using commercial microwave links. *Journal of Hydrometeorology*, 22(4):831–843, 2021. doi: 10.1175/JHM-D-20-0164.1.

Martin Fencl, Roberto Nebuloni, Jafet CM Andersson, Vojtech Bares, Nico Blettner, Greta Cazzaniga, Christian Chwala, Matteo Colli, Lotte de Vos, Abbas El Hachem, et al. Data formats and standards for opportunistic rainfall sensors. *Open Research Europe*, 3:169, 2024.

Oren Goldshtein, Hagit Messer, and Artem Zinevich. Rain rate estimation using measurements from commercial telecommunications links. *IEEE Transactions on Signal Processing*, 57:1616–1625, 2009. doi: 10.1109/TSP.2009.2012554.

Samson Gourevitch, Yazid Janati, Dario Shariatian, Umut Simsekli, Eric Moulines, Eric P Xing, and Alain Durmus. Uniform diffusion models revisited: Leave-one-out denoiser and absorbing state reformulation. *arXiv preprint arXiv:2605.22765*, 2026.

Hai Victor Habi. PyNNcml: Rain estimation and classification from commercial microwave links data, 2026. URL `https://github.com/haihabi/PyNNcml`.

Hai Victor Habi and Hagit Messer. Recurrent neural network for rain estimation using commercial microwave links. *IEEE Transactions on Geoscience and Remote Sensing*, 59(5):3672–3681, 2021.

Barbara Haese, Sebastian Hörning, Christian Chwala, András Bárdossy, Bernd Schalge, and Harald Kunstmann. Stochastic reconstruction and interpolation of precipitation fields using combined information of commercial microwave links and rain gauges. *Water Resources Research*, 53(12): 10740–10756, 2017. doi: 10.1002/2017WR021015.

Ye He, Kevin Rojas, and Molei Tao. What exactly does guidance do in masked discrete diffusion models. In *International Conference on Learning Representations*, 2026. URL `https://proceedings.iclr.cc/paper_files/paper/2026/hash/a14193e9d9fb0b03af0b717de1cac8ac-Abstract-Conference.html`.

Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. In *Advances in Neural Information Processing Systems*, volume 34, pp. 12454–12465, 2021. URL `https://proceedings.neurips.cc/paper/2021/hash/67d96d458abdef21792e6d8e590244e7-Abstract.html`.

International Telecommunication Union. Recommendation ITU-R P.838-3: Specific attenuation model for rain for use in prediction methods. Technical Report P.838-3, ITU-R, March 2005. URL `https://www.itu.int/rec/R-REC-P.838-3-200503-I/en`. Approved in March 2005.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations*, 2017. URL `https://openreview.net/forum?id=rkE3y85ee`.

Xinwu Ji, Yicheng Cai, Xin Zheng, JiaQuan Wan, Yufang Shen, Tao Yang, and Ji Chen. Rainfall field reconstruction from commercial microwave links via a conditional diffusion model. *Water Resources Research*, 62(7):e2025WR042876, 2026. doi: 10.1029/2025WR042876. URL https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2025WR042876.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

Hidde Leijnse, Remko Uijlenhoet, and J. N. M. Stricker. Rainfall measurement using radio links from cellular communication networks. *Water Resources Research*, 43(3):W03201, 2007. doi: 10.1029/2006WR005631.

Stuart P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28 (2):129–137, 1982. doi: 10.1109/TIT.1982.1056489.

David Lopez-Paz and Maxime Oquab. Revisiting classifier two-sample tests. In *International Conference on Learning Representations*, 2017.

Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 32819–32848. PMLR, 2024. URL https://proceedings.mlr.press/v235/lou24a.html.

James B. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pp. 281–297. University of California Press, 1967.

Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The Concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=S1jE5L5gl.

Hagit Messer, Artem Zinevich, and Pinhas Alpert. Environmental monitoring by wireless communication networks. *Science*, 312(5774):713–713, 2006.

D. B. Michelson and J. Koistinen. Gauge-radar network adjustment for the Baltic Sea experiment. *Physics and Chemistry of the Earth, Part B: Hydrology, Oceans and Atmosphere*, 25(10):915–920, 2000. ISSN 1464-1909. doi: 10.1016/S1464-1909(00)00125-8. URL https://www.sciencedirect.com/science/article/pii/S1464190900001258. First European Conference on Radar Meteorology.

Badr Moufad, Yazid Janati, Lisa Bedin, Alain Durmus, Randal Douc, Eric Moulines, and Jimmy Olsson. Variational diffusion posterior sampling with midpoint guidance. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025.

Badr Moufad, Albina Ilina, Hai Victor Habi, Salem Lahlou, Yazid Janati, Hagit Messer, and Eric Moulines. Bayesian rain field reconstruction using commercial microwave links and diffusion model priors. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026.

Naoki Murata, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Bac Nguyen, Stefano Ermon, and Yuki Mitsufuji. G2D2: Gradient-guided discrete diffusion for inverse problem solving. *Transactions on Machine Learning Research*, 2025.

Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking guidance for discrete state-space diffusion and flow models. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/597254dc45be8c166d3ccf0ba2d56325-Abstract-Conference.html.

Aart Overeem, Hidde Leijnse, and Remko Uijlenhoet. Country-wide rainfall maps from cellular communication networks. *Proceedings of the National Academy of Sciences*, 110(8):2741–2745, 2013. doi: 10.1073/pnas.1217961110.

Aart Overeem, Hidde Leijnse, and Remko Uijlenhoet. Two and a half years of country-wide rainfall maps using radio links from commercial cellular telecommunication networks. *Water Resources Research*, 52(10):8039–8065, 2016. doi: 10.1002/2016WR019412.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pp. 234–241. Springer, 2015.

Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T. Chiu, Alexander M. Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language models. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Yair Schiff, Subham Sekhar Sahoo, Hao Phung, Guanghan Wang, Sam Boshar, Hugo Dalla-Torre, Bernardo P. de Almeida, Alexander M. Rush, Thomas Pierrot, and Volodymyr Kuleshov. Simple guidance mechanisms for discrete diffusion models. In *International Conference on Learning Representations*, 2025.

Marc Schleiss, Jörg Rieckermann, and Alexis Berne. Quantification and modeling of wet-antenna attenuation for commercial microwave links. *IEEE Geoscience and Remote Sensing Letters*, 10 (5):1195–1199, 2013. doi: 10.1109/LGRS.2012.2236074.

Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis K. Titsias. Simplified and generalized masked diffusion for discrete data. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Robert L Siddon. Fast calculation of the exact radiological path for a three-dimensional CT array. *Medical Physics*, 12(2):252–255, 1985.

Zhengyan Wan, Yidong Ouyang, Liyan Xie, Fang Fang, Hongyuan Zha, and Guang Cheng. Discrete guidance matching: Exact guidance for discrete flow matching. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/f0470962fc538827f3130757fb707b71-Abstract-Conference.html.

Luhuan Wu, Brian L. Trippe, Christian A. Naesseth, David M. Blei, and John P. Cunningham. Practical and asymptotically exact conditional sampling in diffusion models. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Peng Zhang, Xichuan Liu, and Mingzhong Zou. Reconstructing and nowcasting the rainfall field by a CML network. *Earth and Space Science*, 10(9):e2023EA002909, 2023. doi: 10.1029/2023EA002909.

Lingxiao Zhao, Xueying Ding, Lijun Yu, and Leman Akoglu. Unified discrete diffusion for categorical data. *Journal of Machine Learning Research*, 26(215):1–49, 2025. URL https://www.jmlr.org/papers/v26/25-0171.html.

APPENDIX OUTLINE

A Implementation details **15**

A.1 Dataset preprocessing . . . 15

A.2 Model architectures and training . . . 15

B Implementation of the reconstruction methods **16**

C Evaluation metrics and additional results **17**

D Quantization of rain fields **18**

E Discrete diffusion bridges and training **19**

E.1 State spaces, schedules, and forward corruption . . . 19

E.2 Bridges conditional on the clean field . . . 20

E.3 Exact reverse laws and the learned parameterization . . . 21

E.4 Denoising loss and its variational interpretation . . . 22

E.5 Training and sampling algorithms . . . 23

F Shared CML calculations and gradient estimators **24**

F.1 Exact expectation over clean classes . . . 24

F.2 One construction for Gumbel–Softmax and categorical gradients . . . 26

G G2D2 for the CML likelihood **31**

G.1 The current state and the clean-class probabilities . . . 31

G.2 The local variational objective . . . 31

G.3 Exact derivatives and sequential probability updates . . . 32

G.4 Generating the state at the next time . . . 33

G.5 Relation to the other samplers and scope of the guarantees . . . 34

H Analytic CML guidance with GILC **36**

H.1 The prediction used at the current time . . . 36

H.2 Applying the shared CML derivative . . . 36

H.3 Correcting the probabilities and generating the next state . . . 36

H.4 Using the correction with uniform diffusion . . . 38

H.5 Using a sampled reward gradient . . . 38

I D-MGPS: discrete midpoint guidance **38**

I.1 An exact midpoint identity . . . 38

I.2 Constructing the reference distribution and the observation reward . . . 39

I.3 Optimizing categorical probabilities on the allowed states . . . 41

I.4 Using the shared fit over clean fields . . . 42

I.5 Complete sampling algorithms . . . 43

I.6 Midpoint choice, initialization, and practical limits . . . 44

J Extended related work **45**

J.1 Rainfall retrieval and reconstruction from microwave links . . . 45

J.2 Diffusion and iterative generation on discrete spaces . . . 46

J.3 Observation guidance for discrete diffusion . . . 47

A Implementation details

A.1 Dataset preprocessing

The OpenMRG dataset (Andersson et al., 2022) provides $26{,}495$ rain-rate fields. We split them chronologically into $21{,}196$ training fields and $5{,}299$ test fields (an 80%–20% split). We crop each field from its original width of 37 locations to 36, which gives $H \times W = 48 \times 36$.

Data for model training. We quantize rain intensities with the learned codebook described in Appendix D. Each resulting field contains 48×36 class indices in $\{0, \dots, K-1\}$, where K is the codebook size. During training, we flip each field horizontally with probability 0.1 and, independently, vertically with probability 0.1.

Data for CML-based reconstruction. We evaluate reconstruction on $1{,}000$ randomly sampled reference fields from the test split. We use the CML network topology provided by OpenMRG to simulate observations with noise level $\sigma = 0.05$. Following Moufad et al. (2026), we process the CML data with PyNNcml (Habi, 2026) to obtain the power-law coefficients for each link.

Data for the classifier two-sample test. We use equal numbers of quantized reference fields and fields generated by the discrete diffusion model under evaluation. We assign label 1 to reference fields and label 0 to generated fields. We shuffle the combined dataset and divide it into five folds. For each fold, we train the classifier on the other four folds and evaluate it on the held-out fold.

A.2 Model architectures and training

We describe the architectures and training procedures of the discrete diffusion models and of the C2ST classifier. Table 3 summarizes the settings.

Discrete diffusion architecture. We use a time-conditioned U-Net (Ronneberger et al., 2015) for both uniform and masked diffusion. The categorical value at each location is first mapped to a 32-dimensional embedding. The encoder contains residual blocks with 32 and 64 channels, separated by 2×2 max pooling. The bottleneck contains two 128-channel residual blocks and a four-head self-attention layer. The decoder applies nearest-neighbor upsampling, encoder skip connections, and residual blocks with 64 and 32 channels. Each residual block uses two 3×3 convolutions, group normalization with eight groups, and SiLU activations. A 32-dimensional sinusoidal timestep embedding is projected into every residual block. The final layer concatenates the decoder features with a 1×1 projection of the network input and maps them to one logit per intensity class.

For uniform diffusion, the network output is not used directly as the denoiser logits. We instead use the leave-one-out (LOO) parameterization of Gourevitch et al. (2026):

$$\operatorname{logit}\big(P^\theta(z_t, t)\big) = \mathrm{Network}_\theta(z_t, t) + \log\left(1 + \frac{K\alpha_t}{1-\alpha_t}\right) z_t,$$

where z_t denotes the one-hot representation of the corrupted field. The added term is the log-ratio of the forward probabilities $q_{t|0}(z_t^\ell \mid z_0^\ell)$ for $z_0^\ell = z_t^\ell$ and for $z_0^\ell \neq z_t^\ell$. At the cross-entropy optimum, the network output therefore models the clean class at each location given the other locations of z_t.

We train the masked and uniform diffusion models for 500 epochs with Adam (Kingma & Ba, 2015) and an initial learning rate of 3×10^{-3}. Training one model takes approximately 30 minutes on a single NVIDIA L40S GPU.

Classifier two-sample test. The C2ST classifier applies a 3×3 convolution from one to 32 channels, a ReLU activation, and 2×2 max pooling. A second 3×3 convolution maps the features to 64 channels and is followed by a ReLU activation and global average pooling. A linear layer maps the pooled features to a 64-dimensional representation, and a final linear layer outputs the prediction. For each of the five folds described in Appendix A.1, we train the classifier on the other four folds for 40 epochs with the cross-entropy loss, Adam (Kingma & Ba, 2015) with a learning rate of 10^{-3}, and a batch size of 64. We report the mean classification accuracy over the five held-out folds.

Table 3: Architecture and training hyperparameters of the discrete diffusion models (MDM and UDM) and of the classifier used in the classifier two-sample test (C2ST).

	Hyperparameter	Value
	Discrete diffusion models	
Architecture	Embedding tokens	K (UDM), $K+1$ (MDM)
	Categorical embedding dimension	32
	Encoder channels	$32, 64$
	Bottleneck channels	128
	Decoder channels	$64, 32$
	Residual blocks	2 encoder, 2 bottleneck, 2 decoder
	Convolution kernel / group normalization	3×3 / 8 groups
	Activation	SiLU
	Pooling / upsampling	2×2 max / nearest-neighbor $\times 2$
	Bottleneck self-attention heads	4
	Timestep embedding	32-dimensional sinusoidal
Training	Diffusion steps / schedule	1000 / linear
	Diffusion loss	Weighted cross-entropy
	Number of epochs	500
	Batch size	512
	Optimizer / learning rate	Adam / 3×10^{-3}
	Learning-rate schedule	5-epoch warmup; plateau decay 0.99
	C2ST classifier	
Architecture	Input channels	1
	Convolutional channels	$32, 64$
	Convolution kernel / padding	3×3 / 1
	Activation	ReLU
	Pooling	2×2 max; global average
	Feature dimension	64
Training	Number of folds	5
	Loss function	Cross-entropy
	Number of epochs per fold	40
	Batch size	64
	Optimizer	Adam
	Learning rate	10^{-3}

Table 4: Hyperparameters of the reconstruction methods in the CML experiment. n_{steps}: number of reverse steps; lr: learning rate; $n_{\text{particles}}$: number of particles. For VRG and E-VRG in-filling, the only hyperparameter is $n_{\text{steps}} = 100$.

Discrete diffusion			Continuous diffusion		Meteorological		
G2D2	GILC	D-MGPS	MGPS	TDS	IDW	GMZ	OK
n_{steps}=100 lr=2.5 n_{grad}=30 τ=1	n_{steps}=100 guidance=10 τ=1 $n_{\text{particles}}$=10	n_{steps}=100 lr=2.5 n_{grad}=20 τ=1 $n_{\text{particles}}$=1 η=0.5	ρ=7 n_{steps}=64 lr=3×10^{-2} n_{grad}=10	ρ=7 n_{steps}=320 γ=4 $n_{\text{particles}}$=10	p=2 ROI=6 n_{VRG}=1	p=2 ROI=6 n_{VRG}=5 iterations=20	variogram= `exponential`

B Implementation of the reconstruction methods

This appendix describes the implementation of the discrete-diffusion reconstruction methods and of the continuous-diffusion and meteorological baselines. For the continuous-diffusion and meteorological baselines, we adopt the configurations and released implementation of Moufad et al. (2026). Table 4 summarizes the hyperparameters. All discrete-diffusion methods in Table 1 use 100 reverse steps. For the step-sensitivity study, VRG and E-VRG in-filling, G2D2, GILC, and D-MGPS are also run with 50 and 400 reverse steps.

MDM in-filling. VRG and E-VRG in-filling use the same procedure (Section 3); they differ only in the observed locations and their prescribed tokens. We initialize each observed location with its prescribed token and mask all other locations. When multiple observations fall in the same location, we average their values before quantizing. Let r be the fraction of masked locations. We choose the initialization time t^* such that $1-\alpha_{t^*}$ is closest to r, and set the number of reverse steps according to the interval $[0, t^*]$. The reverse process then reveals the masked locations while keeping the observed tokens fixed.

GILC. We implement Algorithms 1 and 2 of Dou et al. (2026). For UDM, we convert the trained denoiser to its leave-one-out form (Gourevitch et al., 2026) (Appendix A.2) before substituting its probabilities into the bridge $q_{s|0,t}$. Cross-entropy training targets the clean-class probabilities used by the marginalization (bridge-mixture) parameterization of Appendix E.3. Substituting probabilities directly into the bridge (the plug-in parameterization, Equation (14)) recovers the same reverse transition when it is applied to the leave-one-out form. For MDM, the two forms are equivalent because the bridge is linear in z_0 (Gourevitch et al., 2026).

G2D2. We implement Algorithm 1 of Murata et al. (2025), using the released code as a reference[1]. We replace the original star-shaped transition with the Markov variant, which samples z_s from the bridge $q_{s|0,t}$ (Appendix G.4). This variant performed better on rain fields in our experiments. At each reverse step, we initialize the variational distribution with the denoiser prediction, namely the prior logits. For UDM, we sample a clean field from the fitted variational distribution, then sample z_s from the uniform bridge conditioned on that field and z_t (Equation (10)). With the marginalization parameterization, G2D2 does not require a leave-one-out denoiser for UDM.

D-MGPS. We run D-MGPS (Algorithm 8) with both MDM and UDM priors, using the same reverse-step grids as G2D2 and GILC (50, 100, and 400 steps). For each transition from t to $s < t$, we set the midpoint time to $u = \eta s$ (the ratio denoted γ in Appendix I.6), where $\eta \in (0, 1)$ is a hyperparameter.

MGPS and TDS. We use the variance-exploding versions of MGPS (Moufad et al., 2025) and TDS (Wu et al., 2023), with the released implementation and hyperparameters of Moufad et al. (2026).

C Evaluation metrics and additional results

Evaluation metrics. All metrics compare reconstructions (dequantized for the discrete methods) with the reference rain-rate fields on the 48×36 grid, in mm/h. RMSE is the root mean squared difference between reconstructed and reference rain rates over the grid cells. Wet/dry accuracy is the fraction of grid cells whose wet/dry status agrees with that of the reference. PCC is the Pearson correlation coefficient between reconstructed and reference rain rates over the grid cells. The cumulative rain-rate difference is the sum over grid cells of the reconstructed minus the reference rain rate. Its target is zero, and negative values indicate that the total rain rate over the domain is underestimated; it is a spatial sum of rates, not an accumulation over time. The reported values are means over test problems. For uncertainty quantification (Table 2), we draw 500 reconstructions for each of 1,000 test problems and form location-wise intervals at nominal levels of 50%, 75%, and 95%. The empirical marginal coverage is the fraction of locations at which the interval contains the reference rain rate, computed for each problem and averaged over problems; we report it both over all locations and over locations with positive reference rain rate. The interval width is the mean length of these intervals, in mm/h.

Uncertainty quantification. For each reference field $x = (x^1, \ldots, x^L)$, we generate $M = 500$ reconstructions $\{\hat{x}_j\}_{j=1}^M$. At each nominal level $q \in \{0.5, 0.75, 0.95\}$, the prediction interval $I^\ell(q)$ at location ℓ spans the empirical $(1-q)/2$ and $(1+q)/2$ quantiles of all M reconstructions. We compute the all-location coverage as

$$C(x; q) = \frac{1}{L} \sum_{\ell=1}^{L} \mathbf{1}\{x^\ell \in I^\ell(q)\},$$

[1] https://github.com/sony/g2d2

Table 5: Sensitivity of the discrete reconstruction methods to the number of reverse steps (codebook size $K = 20$). Values are means over test problems. The 100-step values are those of Table 1; confidence intervals and wet/dry accuracy are not available for the other runs.

			RMSE ↓			PCC ↑			Cum. rain-rate diff. → 0		
Prior	Method	Steps	50	100	400	50	100	400	50	100	400
MDM	VRG in-filling		0.84	0.84	0.84	0.25	0.26	0.26	−138.10	−135.96	−148.25
	E-VRG in-filling		0.86	0.85	0.85	0.23	0.24	0.23	−134.35	−136.16	−134.79
	G2D2		0.95	0.94	0.92	0.14	0.15	0.16	−488.29	−490.96	−487.31
	GILC		1.21	1.14	1.04	0.10	0.10	0.12	−132.06	−177.26	−169.37
	D-MGPS		1.01	1.01	1.00	0.07	0.07	0.07	−601.02	−606.89	−615.02
UDM	G2D2		0.95	0.95	0.96	0.15	0.14	0.14	−442.07	−423.81	−349.36
	GILC		0.93	0.93	0.94	0.16	0.15	0.15	−363.78	−364.01	−359.12
	D-MGPS		0.98	0.99	1.00	0.06	0.06	0.07	−558.50	−531.26	−492.24

and the wet-location coverage as

$$C_{\mathrm{wet}}(x;q) = \frac{\sum_{\ell=1}^{L} \mathbf{1}\{x^\ell > 0\}\mathbf{1}\{x^\ell \in I^\ell(q)\}}{\sum_{\ell=1}^{L} \mathbf{1}\{x^\ell > 0\}},$$

for fields with at least one wet location. The mean interval width is $W(x;q) = \frac{1}{L}\sum_{\ell=1}^{L} |I^\ell(q)|$ over all locations and $W_{\mathrm{wet}}(x;q) = \sum_{\ell=1}^{L} \mathbf{1}\{x^\ell > 0\}\,|I^\ell(q)| / \sum_{\ell=1}^{L} \mathbf{1}\{x^\ell > 0\}$ over wet locations, where $|I^\ell(q)|$ denotes the length of the interval. We report each metric averaged over the $1{,}000$ selected test fields.

Sensitivity to the number of reverse steps. Table 5 reports RMSE, PCC, and the cumulative rain-rate difference of the discrete reconstruction methods for 50, 100, and 400 reverse steps, with $K = 20$. The 100-step values also appear in Table 1. VRG and E-VRG in-filling remain the most accurate discrete methods at every step count, with RMSE 0.84–0.86 and PCC 0.23–0.26, compared with RMSE of at least 0.92 and PCC of at most 0.16 for the likelihood-guided samplers. With MDM, the RMSE of G2D2 decreases slightly from 0.95 at 50 steps to 0.92 at 400 steps and its PCC increases from 0.14 to 0.16, while its cumulative rain-rate difference stays between -491 and -487. With MDM, GILC has the largest RMSE change across step counts: its RMSE decreases from 1.21 to 1.04 and its PCC increases from 0.10 to 0.12, whereas its cumulative rain-rate difference does not vary monotonically (-132, -177, and -169 at 50, 100, and 400 steps). With UDM, the underestimation of total rain decreases as the number of steps increases for G2D2 (from -442 to -349) and for D-MGPS (from -559 to -492), while their RMSE changes by at most 0.02; GILC with UDM is stable in all three metrics (RMSE 0.93–0.94, PCC 0.15–0.16, cumulative rain-rate difference between -364 and -359).

D-MGPS, in the configuration used here, is less accurate than G2D2 and GILC with both priors and at every step count, with two exceptions relative to GILC with MDM: its RMSE is lower at every step count, and its wet/dry accuracy at 100 steps is higher (0.81 versus 0.79; Table 1). Its RMSE is 0.98–1.01, compared with 0.92–0.96 for G2D2 with both priors and for GILC with UDM, and its PCC (0.06–0.07) is the lowest in the table, compared with 0.10–0.16 for G2D2 and GILC. It also underestimates total rain the most, with cumulative rain-rate differences between -615 and -492; with MDM, its underestimation increases slightly with the number of steps (from -601 to -615). Like G2D2 and GILC, D-MGPS is currently less accurate than in-filling and underestimates total rain; improving these likelihood-guided samplers is left for future work.

D Quantization of rain fields

Separating occurrence from positive intensity. We reserve class 0 for zero rainfall and learn the remaining classes only from strictly positive intensities pooled over training locations. Throughout, K denotes the total number of classes: one dry class and $K - 1$ positive levels. Write $\mathcal{C}_+ = \{a_1, \ldots, a_{K-1}\}$ with $0 < a_1 < \cdots < a_{K-1}$, and $\mathcal{C} = \{0\} \cup \mathcal{C}_+$. The scalar encoder is $\mathcal{Q}(0) = 0$ and $\mathcal{Q}(r) = \arg\min_{1 \le k < K} |r - a_k|^2$ for $r > 0$, with ties assigned to the smaller positive index; decoding returns $\mathcal{D}(k) = a_k$, with $a_0 = 0$. These maps act independently at each grid location. In

particular, any r with $0 < r < a_1/2$ is encoded as class 1 and decoded as $a_1 > 0$: the zero value is never a candidate in the nearest-level search for a positive input.

Squared-error objective on wet cells. Let R denote a scalar rain intensity drawn from the marginal distribution over training locations. Assume $\Pr(R > 0) > 0$ and $\mathbb{E}[R^2 \mid R > 0] < \infty$. The optimal positive codebook solves

$$\mathcal{C}_+^\star \in \operatorname*{arg\,min}_{0<a_1<\cdots<a_{K-1}} \mathbb{E}\left[\min_{1\le k<K} |R - a_k|^2 \,\middle|\, R > 0\right]. \tag{4}$$

For the occurrence-preserving encoder, the unconditional reconstruction error satisfies $\mathbb{E}[|R - \mathcal{D}(\mathcal{Q}(R))|^2] = \Pr(R > 0)\,\mathbb{E}[\min_{1\le k<K} |R - a_k|^2 \mid R > 0]$. Dry cells contribute zero error; the wet-cell probability $\Pr(R > 0)$ scales the objective without changing its optimal positive levels.

Given training intensities $r_1, \ldots, r_M$, let $I_+ = \{i : r_i > 0\}$ and $M_+ = |I_+| > 0$. The empirical objective is

$$\widehat{J}_+(\mathcal{C}_+) = \frac{1}{M_+} \sum_{i\in I_+} \min_{1\le k<K} |r_i - a_k|^2, \qquad 0 < a_1 < \cdots < a_{K-1}. \tag{5}$$

Only wet training cells enter this fit; the codebook is then fixed and applied unchanged to the test fields.

Lloyd iterations. We approximate the empirical minimum with the classical scalar Lloyd algorithm (Lloyd, 1982), that is, one-dimensional $(K-1)$-means clustering of the positive training intensities. The algorithm initializes $K-1$ distinct positive centers and alternates nearest-center assignments with centroid updates. If C_k is the set of positive training indices assigned to class k, the update is $a_k \leftarrow |C_k|^{-1} \sum_{i\in C_k} r_i$ for every nonempty cluster. Every such mean is strictly positive. An empty cluster may retain its previous positive center until the next assignment; it must not be replaced by the dry value. Fitting $K-1$ nonempty classes requires at least $K-1$ distinct positive training values. We sort the fitted positive centers, retaining $a_0 = 0$ as a separate dry class. Assignment and centroid steps do not increase the empirical squared error, but Lloyd's algorithm need not reach its global minimum. Empirical fitting also introduces a sampling error relative to the population objective in Equation (4). For $K = 2$, there is a single positive level, whose empirical optimum is the mean of the positive training intensities.

What quantization preserves. For every field x and cell ℓ, $\mathbf{1}\{\mathcal{D}(\mathcal{Q}(x))^\ell > 0\} = \mathbf{1}\{x^\ell > 0\}$. Consequently, encoding preserves the entire wet–dry mask and its joint distribution across the grid; in particular, $\Pr(\mathcal{Q}(R) = 0) = \Pr(R = 0)$. The wet–dry distinction is between zero and positive values in the preprocessed data; no additional rain-detection threshold is introduced. Quantization can change positive amplitudes, their contrasts, and rain totals, and a finite codebook cannot reproduce arbitrary positive intensities or preserve all extremes. These identities concern the representation only; they do not imply that the learned diffusion prior exactly reproduces the occurrence probabilities or the spatial dependence of the data.

E Discrete diffusion bridges and training

This appendix specifies the forward process, derives its conditional bridges, and distinguishes the exact reverse law from the factorized model used in Section 2. It then gives the training objective and algorithms for learning and unconditional generation. The forward constructions and variational framework follow discrete diffusion models (Austin et al., 2021; Sahoo et al., 2024; Shi et al., 2024). We make the choice of reverse parameterization explicit because it matters for uniform corruption (Gourevitch et al., 2026).

E.1 State spaces, schedules, and forward corruption

A clean field contains L class indices in $\{0, \ldots, K-1\}$, with $K \ge 2$; class 0 is the dry class and classes $1, \ldots, K-1$ are the positive intensity levels (Section 2). We identify class k with its one-hot vector e_k and write $z_0^\ell = e_{k_\ell}$. Uniform diffusion (UDM) uses these K states throughout.

Masked diffusion (MDM) adds the absorbing state $e_{\mathtt{M}}$; clean vectors have a zero mask coordinate. The terminal probability vector is $\bar{q} = \mathbf{1}/K$ for uniform diffusion and $\bar{q} = e_{\mathtt{M}}$ for masked diffusion.

We use discrete times $0, \dots, T$ and a nonincreasing schedule satisfying $\alpha_0 = 1$, $0 < \alpha_t < 1$ for $0 < t < T$, and $\alpha_T = 0$. For $0 \le s < t \le T$, define $\alpha_{t|s} = \alpha_t/\alpha_s$. The denominator is positive because $s < T$. At one location, the forward kernel retains the current class with probability $\alpha_{t|s}$ and otherwise draws from $\bar{q}$:

$$q_{t|s}(z_t^\ell = e_h \mid z_s^\ell = e_j) = \alpha_{t|s}\mathbf{1}\{h = j\} + (1 - \alpha_{t|s})\bar{q}_h. \tag{6}$$

The field transition is the product of these kernels over locations. For uniform corruption, a replacement can draw the original class again; the replacement probability is therefore not the probability that the observed token changes. For masked corruption, a clean token can only stay unchanged or become masked, and a masked token stays masked.

These transition kernels satisfy the Chapman–Kolmogorov equations. Indeed, $R = \mathbf{1}\bar{q}^\top$ satisfies $R^2 = R$, so the matrices $Q_{t|s} = \alpha_{t|s}I + (1 - \alpha_{t|s})R$ obey $Q_{u|s}Q_{t|u} = Q_{t|s}$ for $s < u < t$. In particular, the marginal corruption is

$$q_{t|0}(z_t \mid z_0) = \prod_{\ell=1}^{L} \mathrm{Cat}\big(z_t^\ell; \alpha_t z_0^\ell + (1 - \alpha_t)\bar{q}\big)\,. \tag{7}$$

Thus a training pair (z_0, z_t) can be sampled directly at any selected time without simulating intermediate states. At T, the conditional law is exactly $\bar{q}^{\otimes L}$ and does not depend on z_0. For $0 < t < T$, the tokens of z_t are independent conditional on the clean field z_0. Their marginal law p_t need not factorize, because it averages over spatially dependent clean fields drawn from p_0.

E.2 Bridges conditional on the clean field

Fix times $0 \le s < t \le T$ and a pair (z_0, z_t) with $q_{t|0}(z_t \mid z_0) > 0$. Bayes' rule combined with the Markov property gives

$$q_{s|0,t}(z_s \mid z_0, z_t) = \frac{q_{t|s}(z_t \mid z_s)q_{s|0}(z_s \mid z_0)}{q_{t|0}(z_t \mid z_0)}. \tag{8}$$

Every factor in this ratio is a product over locations, so the bridge factorizes conditional on the *complete clean field* and the corrupted field.

For one location, let k, h, and j denote the clean class, the class at t, and a candidate class at s, respectively. The corresponding bridge probability is

$$B_{s,t}(j \mid k, h) = \frac{[\alpha_s\mathbf{1}\{j = k\} + (1 - \alpha_s)\bar{q}_j]\,[\alpha_{t|s}\mathbf{1}\{h = j\} + (1 - \alpha_{t|s})\bar{q}_h]}{\alpha_t\mathbf{1}\{h = k\} + (1 - \alpha_t)\bar{q}_h}. \tag{9}$$

The numerator is the joint probability of the intermediate and final classes conditional on the clean class. Its sum over j equals the denominator by the forward composition identity, which proves normalization. A zero denominator means that the conditioned pair is impossible; the formula must not be evaluated on that pair.

Uniform bridge. Here all K classes are allowed, and the denominator $d_{t,k}(h) = \alpha_t\mathbf{1}\{k = h\} + (1 - \alpha_t)/K$ is positive for every $t > 0$. The bridge probabilities simplify to

$$B^{\mathrm{unif}}_{s,t}(j \mid k, h) = \frac{[\alpha_s\mathbf{1}\{j = k\} + (1 - \alpha_s)/K]\,[\alpha_{t|s}\mathbf{1}\{j = h\} + (1 - \alpha_{t|s})/K]}{d_{t,k}(h)}. \tag{10}$$

The full-field bridge is the product of these categorical probabilities. Equivalently, its probability vector is the normalized element-wise product of $\alpha_s e_k + (1-\alpha_s)\mathbf{1}/K$ and $\alpha_{t|s}e_h + (1-\alpha_{t|s})\mathbf{1}/K$. This form is convenient for sampling because normalization can be performed directly on the nonnegative numerator.

Masked bridge. If $z_t^\ell = e_{\mathtt{M}}$, define $r_{s,t} = (\alpha_s - \alpha_t)/(1 - \alpha_t)$. The only possible states at s are the clean class and the mask, with

$$q^{\text{mask}}_{s|0,t}(z_s^\ell \mid z_0^\ell, z_t^\ell = e_{\mathtt{M}}) = \mathrm{Cat}\big(z_s^\ell; r_{s,t} z_0^\ell + (1 - r_{s,t}) e_{\mathtt{M}}\big)\,. \tag{11}$$

The numerator of $r_{s,t}$ is the probability that a token is still clean at s and becomes masked by t; the denominator is its probability of being masked at t. For an unmasked endpoint $z_t^\ell = e_h$, compatibility requires $z_0^\ell = e_h$, and the bridge is the point mass at e_h. In particular, an unmasked endpoint at T has zero probability and is never a valid conditioning event.

Endpoint and schedule checks. At $s = 0$, both bridges return the known clean token with probability one. If $\alpha_s = \alpha_t$, no corruption occurs between these times, so the bridge returns the current token. At $t = T$, the bridge equals the forward law at s conditional on z_0, because the terminal draw is independent of all earlier states. For masked diffusion, $r_{0,t} = 1$, so the last reverse step removes every remaining mask. These identities provide direct checks on an implementation.

E.3 Exact reverse laws and the learned parameterization

The data distribution and forward kernels define $p_{0:T}(z_{0:T}) = p_0(z_0) \prod_{t=1}^T q_{t|t-1}(z_t \mid z_{t-1})$ and $p_t(z_t) = \sum_{z_0} p_0(z_0) q_{t|0}(z_t \mid z_0)$. For $p_t(z_t) > 0$, the exact clean-field conditional is $p_{0|t}(z_0 \mid z_t) = p_0(z_0) q_{t|0}(z_t \mid z_0)/p_t(z_t)$. Consequently, the exact reverse transition satisfies

$$p_{s|t}(z_s \mid z_t) = \sum_{\substack{z_0: \\ q_{t|0}(z_t|z_0)>0}} q_{s|0,t}(z_s \mid z_0, z_t)\, p_{0|t}(z_0 \mid z_t). \tag{12}$$

The clean-field conditional is generally unavailable. It also need not factorize: observing a corrupted field does not remove dependence between its unknown clean tokens. Thus the factorized hard-clean bridge does not imply a factorized exact reverse transition after marginalizing z_0.

Clean-token prediction. The denoiser returns K logits per location; for uniform diffusion, they are obtained from the network output by the leave-one-out parameterization described in Appendix A.2. Let $P^\theta_{\ell k}(z_t, t)$ denote their softmax probabilities. For masked diffusion, these probabilities are used only at masked locations; at a revealed location $z_t^\ell = e_h$, set $P^\theta_{\ell k} = \mathbf{1}\{k = h\}$. With this carry-over rule, $P^\theta_\ell(z_t, t)$ is the denoiser prediction of Section 2, and the factorized proxy $p^\theta_{0|t}(z_0 \mid z_t) = \prod_\ell \mathrm{Cat}(z_0^\ell; P^\theta_\ell)$ has the support required by the masked bridge. No clean prediction assigns probability to the mask.

We use the normalized-bridge mixture of Section 2. Its per-location reverse probabilities are

$$\Pi^\theta_{s,t,\ell}(j \mid z_t) = \sum_{k=0}^{K-1} P^\theta_{\ell k}(z_t, t)\, B_{s,t}(j \mid k, h_\ell), \qquad z_t^\ell = e_{h_\ell}, \tag{13}$$

where zero-weight incompatible terms are omitted rather than evaluated. The field kernel is the product of these categorical distributions. Every prediction can depend on the complete z_t, so the learned generative process can still produce spatial dependence. However, even exact per-location clean marginals would not generally recover the joint reverse law in Equation (12). Moreover, $p^\theta_{0|t}$ is a denoiser proxy, not generally the clean-field conditional obtained by running the learned reverse chain to completion. For example, suppose a two-cell binary prior assigns probability $1/2$ to each of $(0, 0)$ and $(1, 1)$. At complete masking, both exact cell marginals are $(1/2, 1/2)$, but a single factorized jump from T to 0 assigns probability $1/4$ to each of the four fields, including the two absent from the prior.

Masked reverse step. At revealed locations, Equation (13) copies the token. At masked locations, linearity of the bridge in the clean one-hot vector gives the probability vector $r_{s,t} P^\theta_\ell + (1 - r_{s,t}) e_{\mathtt{M}}$, with a zero mask coordinate appended to P^θ_ℓ. A draw therefore consists of a Bernoulli reveal decision with probability $r_{s,t}$, followed, if revealed, by a categorical clean-class draw from P^θ_ℓ. Revealed tokens are never changed or remasked.

Uniform mixture versus soft substitution. For uniform corruption, the hard-clean normalizer $d_{t,k}(h)$ depends on k. It must remain inside the mixture in Equation (13). Replacing e_k by P_ℓ^θ directly in the bridge's numerator and denominator instead gives

$$\widehat{\Pi}_j = \frac{[\alpha_s P_{\ell j}^\theta + (1-\alpha_s)/K]\,[\alpha_{t|s}\mathbf{1}\{j=h\} + (1-\alpha_{t|s})/K]}{\alpha_t P_{\ell h}^\theta + (1-\alpha_t)/K}. \tag{14}$$

This is a different normalized distribution: it mixes the hard-clean bridges with weights proportional to $P_{\ell k}^\theta d_{t,k}(h)$, rather than $P_{\ell k}^\theta$. For example, with $K=2$, $s=0$, $\alpha_t = 1/2$, $h=0$, and $P_\ell^\theta = (1/2, 1/2)$, the bridge mixture returns $(1/2, 1/2)$, whereas soft substitution returns $(3/4, 1/4)$. This distinction concerns the reverse parameterization, not the validity of Bayes' rule. In particular, the joint-kernel normalization in Austin et al. (2021, Equation (4)) should not be identified with our mixture of individually normalized bridges; Gourevitch et al. (2026) examine this distinction for uniform models.

The mixture can nevertheless be evaluated in $O(K)$ operations per location. Set $\psi_k = P_{\ell k}^\theta / d_{t,k}(h)$ and $\Psi = \sum_k \psi_k$. Then its exact probability vector is

$$\Pi_{s,t,\ell}^\theta(j \mid z_t) = \left[\alpha_{t|s}\mathbf{1}\{j=h\} + \frac{1-\alpha_{t|s}}{K}\right]\left[\alpha_s \psi_j + \frac{1-\alpha_s}{K}\Psi\right]. \tag{15}$$

Equivalently, substituting the normalized vector ψ/Ψ for P_ℓ^θ in Equation (14) returns exactly these mixture probabilities. Since $d_{t,h}(h)/d_{t,k}(h) = 1 + K\alpha_t/(1-\alpha_t)$ for $k \neq h$, the logits of ψ/Ψ are those of P_ℓ^θ with $\log(1 + K\alpha_t/(1-\alpha_t))$ subtracted at the current class h; this is the leave-one-out form used in our implementation (Gourevitch et al., 2026). Alternatively, first draw a clean class from P_ℓ^θ and then draw its hard-clean bridge. This two-stage procedure samples exactly the same mixture without explicitly evaluating every clean-class component.

E.4 Denoising loss and its variational interpretation

Let ϖ be the distribution on $\{1, \dots, T\}$ from which training times are drawn, and let $w_t^\ell(z_t) \geq 0$ be specified loss weights. For a clean token $z_0^\ell = e_{k_\ell}$, its cross-entropy is $-\log P_{\ell k_\ell}^\theta(z_t, t)$. The weighted denoising objective is

$$\mathcal{L}_{\mathrm{CE}}(\theta) = \mathbb{E}_{\substack{z_0 \sim p_0,\ t\sim\varpi \\ z_t \sim q_{t|0}(\cdot|z_0)}}\left[-\sum_{\ell=1}^{L} w_t^\ell(z_t)\log P_{\ell k_\ell}^\theta(z_t, t)\right]. \tag{16}$$

For masked diffusion, revealed locations have target probability one under carry-over and contribute zero; only masked locations need to be evaluated. For uniform diffusion, all locations are included. An unchanged token under uniform corruption is not known to be clean and must not be copied deterministically.

What cross-entropy learns. At a fixed (z_t, t) with $\varpi(t)p_t(z_t) > 0$ and a location with positive weight, the conditional expected cross-entropy is minimized by the true marginal probabilities $p_{0|t}(z_0^\ell = e_k \mid z_t)$, if the model can represent them. This statement uses weights that do not depend on the unknown clean label after conditioning on (z_t, t). Clean-class-dependent weights generally change that optimum. Recovering these marginals is distinct from recovering the full clean-field conditional, which may contain dependencies that the product proxy cannot represent.

Negative evidence lower bound. Consider the learned model on the full consecutive grid, $p_{0:T}^\theta(z_{0:T}) = \bar{q}^{\otimes L}(z_T)\prod_{t=1}^T p_{t-1|t}^\theta(z_{t-1} \mid z_t)$. For a fixed clean field, the usual variational decomposition gives

$$\begin{aligned} -\log p_0^\theta(z_0) \;\leq\; \mathcal{B}_\theta(z_0) = {} & \mathrm{KL}\big(q_{T|0}(\cdot \mid z_0) \,\big\|\, \bar{q}^{\otimes L}\big) \\ & + \sum_{t=2}^{T} \mathbb{E}_{q_{t|0}(\cdot|z_0)}\left[\mathrm{KL}\Big(q_{t-1|0,t}(\cdot \mid z_0, z_t) \,\Big\|\, p_{t-1|t}^\theta(\cdot \mid z_t)\Big)\right] \\ & - \mathbb{E}_{q_{1|0}(\cdot|z_0)}[\log p_{0|1}^\theta(z_0 \mid z_1)]. \end{aligned} \tag{17}$$

The first term is zero because $\alpha_T = 0$ makes the forward terminal law exactly equal to the model's terminal prior. If $\alpha_T > 0$, it would not generally vanish; for a point-mass mask prior, the remaining unmasked forward states would even give an infinite terminal KL. The last term is the reconstruction term, not another bridge KL. For our normalized-bridge parameterization, the final-step identity $B_{0,t}(j \mid k, h) = \mathbf{1}\{j = k\}$ makes this decoder equal to the factorized denoiser prediction.

Exact finite-step reduction for masked diffusion. Write $r_t = (\alpha_{t-1} - \alpha_t)/(1 - \alpha_t)$. At a masked location, the true bridge assigns mass r_t to the known clean class and $1 - r_t$ to the mask. The learned bridge has the same mask mass and assigns clean-class mass $r_t P^\theta_\ell$. Thus, for $t \geq 2$, the location-wise KL is exactly $r_t[-\log P^\theta_{\ell k_\ell}]$. At a revealed location, both distributions are the same point mass, giving zero KL. The reconstruction term at $t = 1$ has the same form because $r_1 = 1$. Consequently,

$$\mathbb{E}_{z_0 \sim p_0}[\mathcal{B}_\theta(z_0)] = \sum_{t=1}^{T} r_t \, \mathbb{E}_{\substack{z_0 \sim p_0 \\ z_t \sim q_{t|0}(\cdot|z_0)}} \left[- \sum_{\ell : z_t^\ell = e_{\mathtt{M}}} \log P^\theta_{\ell k_\ell}(z_t, t) \right]. \tag{18}$$

This is the finite-grid weighted cross-entropy reduction underlying masked-diffusion objectives (Sahoo et al., 2024; Shi et al., 2024). Sampling $t \sim \varpi$ gives an unbiased estimator of this sum if the weight on each masked location is $r_t/\varpi(t)$, with $\varpi(t) > 0$ whenever $r_t > 0$. Changing ϖ therefore requires the corresponding importance correction when the intended objective is this bound. The identity also relies on deterministic carry-over and on retaining the prescribed mask probabilities in the model.

Uniform diffusion and alternative weightings. For uniform diffusion with our bridge-mixture parameterization, the intermediate KL terms in Equation (17) involve the full categorical reverse probabilities. They do not generally reduce to a scalar multiple of clean-token cross-entropy. Weighted cross-entropy remains a valid denoising objective, but it is not automatically this model's negative ELBO. Results for alternative reverse parameterizations or objectives must be interpreted with their own assumptions (Austin et al., 2021; Gourevitch et al., 2026). In particular, sharing a population minimizer does not imply equality of the objective values.

E.5 Training and sampling algorithms

Algorithm 1 samples the forward marginal directly and trains either model. Its weights are those of Equation (16); for the exact masked bound, use $r_t/\varpi(t)$ on active locations. A field with no masked location contributes zero to the masked loss. The displayed objective sums over locations: dividing by a fixed L only rescales it, whereas dividing by the random number of masked locations generally changes its weighting.

Algorithm 2 implements the learned transition on a selected grid of times $T = \tau_N > \cdots > \tau_0 = 0$ in $\{0, \ldots, T\}$. It evaluates the denoiser once per step and samples different locations independently conditional on the current field. For uniform diffusion, the auxiliary clean-class draw implements the normalized-bridge mixture exactly. For masked diffusion, the simpler reveal-then-class construction avoids evaluating incompatible bridges. Changing the sampling grid can change the generated distribution: learned finite-step kernels do not generally satisfy the composition identity of the exact reverse process.

Numerical and reproducibility details. Evaluate training losses with log-softmax, and exclude revealed masked-diffusion locations rather than taking logarithms of their hard carry-over probabilities. In uniform bridge sampling, normalize the nonnegative numerator; log-space evaluation can be used when the noise level is very small. For the direct mixture formula in Equation (15), small $d_{t,k}(h)$ require particular care with floating-point range. Validate probability sums and the final $s = 0$ transition, and never silently replace an impossible masked conditioning pair by an arbitrary distribution. The algorithms specify the mathematical procedures. Table 3 lists the number of diffusion steps, the schedule, the optimizer, and the training budget used in our experiments.

Algorithm 1 Weighted denoising training for masked or uniform diffusion

Input: Quantized training fields, mode, schedule α, time law ϖ, and weights w
Input: Denoiser parameters θ, batch size B, and an optimizer
Output: Trained clean-token denoiser

for each optimization step **do**
  for $b = 1, \dots, B$ **do**
    Draw a clean field z_0 and time $t \sim \varpi$
    Sample the forward marginal
    for each location ℓ **do**
      Draw $C_\ell \sim \mathrm{Bernoulli}(\alpha_t)$
      if $C_\ell = 1$ **then**
        $z_t^\ell \leftarrow z_0^\ell$
      else if mode is masked **then**
        $z_t^\ell \leftarrow e_{\mathtt{M}}$
      else
        Draw J_ℓ uniformly from $\{0, \dots, K-1\}$; $z_t^\ell \leftarrow e_{J_\ell}$
      end if
    end for
    Evaluate the supervised clean-token loss
    $\eta \leftarrow \mathrm{Logits}_\theta(z_t, t)$
    $A \leftarrow \{\ell : z_t^\ell = e_{\mathtt{M}}\}$ if masked, otherwise $A \leftarrow \{1, \dots, L\}$
    Read clean indices k_ℓ from $z_0^\ell = e_{k_\ell}$
    $\widehat{\mathcal{L}}_b \leftarrow -\sum_{\ell \in A} w_t^\ell(z_t)[\mathrm{logsoftmax}(\eta_\ell)]_{k_\ell}$
  end for
  $\widehat{\mathcal{L}} \leftarrow B^{-1} \sum_{b=1}^{B} \widehat{\mathcal{L}}_b$
  Update θ with the chosen optimizer using $\nabla_\theta \widehat{\mathcal{L}}$
end for

F SHARED CML CALCULATIONS AND GRADIENT ESTIMATORS

This appendix derives the expected log-likelihood and its derivatives for the Gaussian model in Equation (1); they are used by our CML adaptations of G2D2 and GILC and by D-MGPS. The calculation takes an arbitrary product distribution over clean fields as its input. We first derive the exact expectation over clean classes, then give one common construction of Gumbel–Softmax, straight-through, and score-function estimators for the three methods. The following appendices specify their reference distributions and sampling transitions.

F.1 EXACT EXPECTATION OVER CLEAN CLASSES

Attenuation and log-likelihood on the codebook. At cell ℓ, let $k_\ell \in \{0, \dots, K-1\}$ be the clean rain class. We represent this class by the one-hot vector $z_0^\ell = e_{k_\ell}$, which corresponds to rain intensity a_{k_ℓ} in the codebook of Section 2. For link i, precompute $c_{i\ell k} = \kappa_i \Delta_\ell^i a_k^{b_i}$. Here $a_k \geq 0$ is the intensity assigned to class k, Δ_ℓ^i is the length of link i inside cell ℓ, and $\kappa_i > 0$ and $b_i > 0$ are the fixed power-law coefficients of that link. Thus $c_{i\ell k}$ is the attenuation contributed by that cell when its class is k. The dry class has $a_0 = 0$ and contributes $c_{i\ell 0} = 0$; a cell outside the link has $c_{i\ell k} = 0$ for every class. The total attenuation is $A_i(z_0) = \sum_\ell c_{i\ell k_\ell} = [\mathcal{M}_{\mathrm{CML}}(\mathcal{D}(z_0))]_i$. With independent Gaussian observation errors of standard deviations $\sigma_i > 0$, the log-likelihood reward is $r_y(z_0) = -\frac{1}{2} \sum_i (y_i - A_i(z_0))^2 / \sigma_i^2$, after omitting the constant independent of z_0. The observations, noise scales, codebook, and attenuation coefficients are fixed throughout sampling.

Product distribution and attenuation moments. Let $\Xi = (\Xi_{\ell k})$ be an arbitrary matrix of clean-class probabilities, with $\Xi_{\ell k} \geq 0$ and $\sum_k \Xi_{\ell k} = 1$ at each cell. It defines a random clean field Z_0 by $\Pr_\Xi\{Z_0 = (e_{k_1}, \dots, e_{k_L})\} = \prod_{\ell=1}^{L} \Xi_{\ell k_\ell}$. Write $\mathbb{E}_\Xi$ and Var_Ξ for expectation and variance under this distribution. The mean contribution of cell ℓ to link i is $m_{i\ell}(\Xi) = \sum_k \Xi_{\ell k} c_{i\ell k}$, and the mean total attenuation is $\mu_i(\Xi) = \sum_\ell m_{i\ell}(\Xi)$. Independence across cells gives the attenuation variance

Algorithm 2 Unconditional generation with the normalized-bridge mixture

Input: Trained denoiser θ, mode, codebook $\mathcal{C}$, schedule α
Input: Reverse grid $T = \tau_N > \cdots > \tau_0 = 0$
Output: One generated rain field $\widetilde{x}$

Draw $z_T^\ell \sim \mathrm{Cat}(\bar{q})$ independently at every location
for $n = N, \ldots, 1$ **do**
  $(t, s) \leftarrow (\tau_n, \tau_{n-1})$
  $\eta \leftarrow \mathrm{Logits}_\theta(z_t, t)$; $P_\ell \leftarrow \mathrm{softmax}(\eta_\ell)$
  $z_s \leftarrow z_t$
  if mode is masked **then**
    $r \leftarrow (\alpha_s - \alpha_t)/(1 - \alpha_t)$
    for each location ℓ with $z_t^\ell = e_{\mathtt{M}}$ **do**
      Draw $C_\ell \sim \mathrm{Bernoulli}(r)$
      if $C_\ell = 1$ **then**
        Draw k with probabilities P_ℓ; set $z_s^\ell \leftarrow e_k$
      end if
    end for
    All previously revealed tokens remain unchanged
  else
    $\alpha_{t|s} \leftarrow \alpha_t/\alpha_s$
    for each location ℓ **do**
      Read h from $z_t^\ell = e_h$; draw a clean class k with probabilities P_ℓ
      $b_j \leftarrow [\alpha_s \mathbf{1}\{j = k\} + (1 - \alpha_s)/K][\alpha_{t|s} \mathbf{1}\{j = h\} + (1 - \alpha_{t|s})/K]$ for all j
      Draw j with probabilities $b_j / \sum_{j'} b_{j'}$; set $z_s^\ell \leftarrow e_j$
    end for
  end if
end for
return $\widetilde{x} = \mathcal{D}(z_0)$

$v_i(\Xi) = \sum_{\ell,k} \Xi_{\ell k}(c_{i\ell k} - m_{i\ell}(\Xi))^2$. If a cell is fixed to one class, its probability row is a point mass; it contributes that class's attenuation to $\mu_i(\Xi)$ and zero to $v_i(\Xi)$.

Expected log-likelihood. For each link, the mean squared residual satisfies $\mathbb{E}_\Xi[(y_i - A_i(Z_0))^2] = (y_i - \mu_i(\Xi))^2 + v_i(\Xi)$. Substituting this identity into the log-likelihood gives

$$\mathcal{R}(\Xi) = \mathbb{E}_\Xi[r_y(Z_0)] = -\frac{1}{2} \sum_{i=1}^{m} \frac{(y_i - \mu_i(\Xi))^2 + v_i(\Xi)}{\sigma_i^2}. \tag{19}$$

Although two links can share cells and have correlated attenuations, no covariance between links enters this expression: the assumed Gaussian log-likelihood is a sum of separate squared residuals. Evaluating the log-likelihood only at the mean attenuation omits $v_i(\Xi)$ and can change the guidance direction. All powers $a_k^{b_i}$ are evaluated on fixed codebook values, so this calculation requires neither continuous replacements of categorical samples nor differentiation of the power law at zero intensity. Equation (19) is exact under the chosen product distribution.

Changing the class at one cell. Fix a cell ℓ and a candidate class k. Set that cell to e_k and draw all other cells independently using their rows of Ξ. Write $\overline{r}_{\ell k}(\Xi)$ for the resulting expected reward. This definition also applies when $\Xi_{\ell k} = 0$: we fix the candidate class rather than condition on an event of probability zero. Writing k_h for the class drawn at another cell h, the attenuation from the other cells is $A_i^{(-\ell)} = \sum_{h \neq \ell} c_{ihk_h}$, with mean $\mu_{i,-\ell}(\Xi) = \mu_i(\Xi) - m_{i\ell}(\Xi)$ and variance $v_{i,-\ell}(\Xi) = \sum_{h \neq \ell} \sum_j \Xi_{hj}(c_{ihj} - m_{ih}(\Xi))^2$. For link i, the mean squared residual after fixing class k is therefore $(y_i - \mu_{i,-\ell}(\Xi) - c_{i\ell k})^2 + v_{i,-\ell}(\Xi)$. Expanding this square gives $\overline{r}_{\ell k}(\Xi) = C_\ell(\Xi) + H_{\ell k}(\Xi)$, where $C_\ell(\Xi) = -\frac{1}{2} \sum_i [(y_i - \mu_{i,-\ell}(\Xi))^2 + v_{i,-\ell}(\Xi)]/\sigma_i^2$ and

$$H_{\ell k}(\Xi) = \sum_i \frac{(y_i - \mu_i(\Xi) + m_{i\ell}(\Xi))c_{i\ell k} - \frac{1}{2}c_{i\ell k}^2}{\sigma_i^2}. \tag{20}$$

The term $C_\ell(\Xi)$ is the same for every candidate class. Both $C_\ell(\Xi)$ and $H_{\ell k}(\Xi)$ depend only on probabilities at the other cells, since $\mu_i(\Xi) - m_{i\ell}(\Xi)$ excludes cell ℓ. Averaging over the class of cell ℓ with probabilities Ξ_ℓ gives $\mathcal{R}(\Xi) = C_\ell(\Xi) + \sum_k \Xi_{\ell k} H_{\ell k}(\Xi)$. Thus the expected reward is linear in one probability row when all other rows are held fixed. It is generally not linear in all rows together.

Derivative with respect to class scores. For a row that can change, let S_ℓ be its allowed classes and write $\Xi_{\ell k} = \exp(\phi_{\ell k}) / \sum_{j \in S_\ell} \exp(\phi_{\ell j})$ for $k \in S_\ell$, with zero probability elsewhere. The scores $\phi_{\ell k}$ are real numbers; adding the same number to all of them leaves the row unchanged. Since $\partial \Xi_{\ell j} / \partial \phi_{\ell k} = \Xi_{\ell j}(\mathbf{1}\{j = k\} - \Xi_{\ell k})$, differentiating the preceding linear expression yields

$$\frac{\partial \mathcal{R}(\Xi)}{\partial \phi_{\ell k}} = \Xi_{\ell k} \left(H_{\ell k}(\Xi) - \sum_{j \in S_\ell} \Xi_{\ell j} H_{\ell j}(\Xi) \right). \tag{21}$$

The term $C_\ell(\Xi)$ cancels because the probabilities in the row sum to one. No derivative is taken through the class lookup or the coefficients $c_{i\ell k}$. A row fixed to one class has no free score and no update. The functions in this appendix have no time index because their argument Ξ specifies the entire distribution being averaged. The following appendices supply that argument with its time dependence: the denoiser probabilities P^t, the fitted clean-class probabilities V^t, or the prediction $P^u(v)$ at an intermediate state.

F.2 One construction for Gumbel–Softmax and categorical gradients

Gumbel–Softmax replaces a sampled class by a differentiable vector during optimization (Jang et al., 2017; Maddison et al., 2017). Its straight-through variant evaluates the reward at a sampled class but uses the derivative of that vector. We describe both variants here, together with the score-function estimator, often called REINFORCE. The same construction applies to the clean-field fits in G2D2, the single correction in GILC, and the midpoint fit in D-MGPS. It does not change the forward diffusion schedule or the final categorical sampling step.

F.2.1 Times, probabilities, and the exact categorical objective

Fix a reverse step $t \to s$, its current state z_t, and a time $0 \le u \le s < t$ at which to fit a distribution. G2D2 and GILC use $u = 0$; D-MGPS can also use $u > 0$. The state z_t and the times (u, t) stay fixed throughout this local calculation. For each cell, let $S^\ell_{u,t}$ be its allowed clean classes, let $A^{u,t}_{\ell k} > 0$ be reference probabilities on this set, and let $\chi^{u,t}_\ell \in [0, 1]$ be the fixed probability of being unmasked at time u. Each reference row $A^{u,t}_\ell$ sums to one. At an adjustable cell, real scores $\phi^{u,t}_{\ell k}$, called logits, define $W^{u,t}_{\ell k} = \mathrm{softmax}_{S^\ell_{u,t}}(\phi^{u,t}_\ell)_k$. The two superscripts identify the time of the fitted state and the current reverse time; the scores change during optimization at these fixed times. Scores for forbidden classes are absent, rather than finite approximations to $-\infty$.

For uniform diffusion, set $\chi^{u,t}_\ell = 1$ and use $V^{u,t}_{\ell k} = W^{u,t}_{\ell k}$. For masked diffusion, set

$$V^{u,t}_{\ell k} = \chi^{u,t}_\ell W^{u,t}_{\ell k}, \qquad V^{u,t}_{\ell \mathtt{M}} = 1 - \chi^{u,t}_\ell, \qquad q_{\phi^{u,t}}(v) = \prod_\ell V^{u,t}_\ell(v^\ell). \tag{22}$$

Thus $W^{u,t}_\ell$ specifies the class conditional on being unmasked, not the mask probability. For a cell masked in z_t, the schedule gives $\chi^{u,t}_\ell = (\alpha_u - \alpha_t)/(1 - \alpha_t)$. A cell already revealed in z_t has its class fixed, with $\chi^{u,t}_\ell = 1$ and a singleton support. At $u = 0$, every $\chi^{0,t}_\ell = 1$. Cells with $\chi^{u,t}_\ell = 0$ stay masked; their unused $W^{u,t}_\ell$ may be set equal to $A^{u,t}_\ell$. Only cells in $I_{u,t} = \{\ell : \chi^{u,t}_\ell > 0,\ |S^\ell_{u,t}| > 1\}$ have adjustable logits.

The reference distribution has the same factors as Equation (22), with $A^{u,t}$ in place of $W^{u,t}$. For G2D2 and GILC, $A^{0,t}$ is P^t restricted to its positive support. For D-MGPS, the reference is $\rho_{u|t}$ from Equation (45); its clean-class factors $A^{u,t}$ and mask probabilities are specified in Appendix I.3. Both the reference and its support are computed once from (z_t, t) and held fixed.

Write $\mathcal{R}_0(v) = r_y(v)$ for a clean field and $\mathcal{R}_u(v) = \mathcal{R}(P^u(v))$ for $u > 0$, where $P^u(v)$ is the clean-class prediction at the categorical state (v, u). These are the rewards in Equation (50). The

categorical KL term is $\mathcal{K}_{u,t}(\phi^{u,t}) = \sum_{\ell \in I_{u,t}} \chi_\ell^{u,t} \sum_{k \in S_{u,t}^\ell} W_{\ell k}^{u,t} \log(W_{\ell k}^{u,t}/A_{\ell k}^{u,t})$. The common variational objective is $\mathcal{F}_{u,t} = \mathcal{K}_{u,t} - \lambda_t \mathbb{E}_{q_{\phi^{u,t}}}[\mathcal{R}_u(Z_u)]$, with observation strength $\lambda_t \geq 0$. Its specializations are Equations (34) and (51). GILC uses the reward derivative at the reference without running this iterative fit.

Let $d_\ell^{u,t} = \sum_k W_{\ell k}^{u,t} \log(W_{\ell k}^{u,t}/A_{\ell k}^{u,t})$. The softmax derivative gives the exact KL derivative

$$\mathsf{K}_{\ell k}^{u,t} := \frac{\partial \mathcal{K}_{u,t}}{\partial \phi_{\ell k}^{u,t}} = \chi_\ell^{u,t} W_{\ell k}^{u,t} \left(\log \frac{W_{\ell k}^{u,t}}{A_{\ell k}^{u,t}} - d_\ell^{u,t} \right). \tag{23}$$

Indeed, differentiating $w \log(w/a)$ gives $\log(w/a) + 1$, and the constant one cancels because a softmax row sums to one. This derivative array, written $\mathsf{K}^{u,t}$ to distinguish it from the codebook size K, is never estimated from samples. In particular, we neither replace $W^{u,t}$ by a relaxed draw in the KL nor use a KL between continuous Concrete distributions.

F.2.2 Exact class draws and their continuous replacements

At one optimization iteration, draw independent uniforms $\omega_{\ell k} \in (0,1)$ and set $g_{\ell k}^{u,t} = -\log(-\log \omega_{\ell k})$ on each adjustable support. The variables $g_{\ell k}^{u,t}$ have distribution function $\exp(-e^{-x})$. For a temperature $\vartheta^{u,t} > 0$, form

$$Y_{\ell k}^{u,t} = \frac{\exp((\phi_{\ell k}^{u,t} + g_{\ell k}^{u,t})/\vartheta^{u,t})}{\sum_{a \in S_{u,t}^\ell} \exp((\phi_{\ell a}^{u,t} + g_{\ell a}^{u,t})/\vartheta^{u,t})}, \qquad U_\ell^{u,t} = e_{\arg\max_{k \in S_{u,t}^\ell} (\phi_{\ell k}^{u,t} + g_{\ell k}^{u,t})}. \tag{24}$$

The Gumbel-max identity gives $\Pr(U_\ell^{u,t} = e_k) = W_{\ell k}^{u,t}$ (Jang et al., 2017). To see this, $\exp(-g_{\ell k}^{u,t}) = -\log \omega_{\ell k}$ is exponential with rate one. Maximizing $\phi_{\ell k}^{u,t} + g_{\ell k}^{u,t}$ is equivalent to minimizing independent exponential variables with rates $\exp(\phi_{\ell k}^{u,t})$. The probability that class k has the smallest value is its rate divided by the sum of all rates, namely $W_{\ell k}^{u,t}$. The temperature does not enter this categorical law: dividing every score by the same positive number does not change its maximizer. Ties have probability zero; a deterministic index rule handles finite-precision ties. By contrast, $Y_\ell^{u,t}$ is a continuous vector, and its mean is not generally $W_\ell^{u,t}$. On a singleton support, set $Y_\ell^{u,t} = U_\ell^{u,t}$ to its prescribed one-hot vector without drawing Gumbel noise.

Independently of all class draws, sample $B_\ell^{u,t} \sim \mathrm{Bernoulli}(\chi_\ell^{u,t})$. For uniform diffusion and for $u = 0$, this indicator equals one. Define a relaxed input $X^{\mathrm{R},u,t}$ and a categorical input $X^{\mathrm{H},u,t}$ by

$$\begin{aligned} X_{\ell k}^{\mathrm{R},u,t} &= B_\ell^{u,t} Y_{\ell k}^{u,t}, \quad X_{\ell k}^{\mathrm{H},u,t} = B_\ell^{u,t} U_{\ell k}^{u,t}, \qquad k < K, \\ X_{\ell \mathtt{M}}^{\mathrm{R},u,t} &= 1 - B_\ell^{u,t}, \quad X_{\ell \mathtt{M}}^{\mathrm{H},u,t} = 1 - B_\ell^{u,t}. \end{aligned} \tag{25}$$

Omit the mask coordinate for uniform diffusion. When $\chi_\ell^{u,t} = 0$, directly set both input rows to $e_{\mathtt{M}}$; no class scores or class noise are needed. The hard field $X^{\mathrm{H},u,t}$ has exactly the distribution $q_{\phi^{u,t}}$. Even a singleton clean support requires a Bernoulli draw when $0 < \chi_\ell^{u,t} < 1$: its class is known if revealed, but revelation is still random. We do not replace the Bernoulli indicator by its mean, optimize its probability, or apply a Gumbel relaxation to the mask decision. This preserves the schedule and makes the mask pattern independent of the fitted class logits.

F.2.3 A specified differentiable CML reward

A derivative through Y requires a reward defined at fractional inputs, not only at one-hot fields. The values of the likelihood on categorical fields alone do not specify this extension. At $u = 0$, we use

$$f_0(X) = -\frac{1}{2} \sum_i \frac{\left(y_i - \sum_{\ell,k} c_{i\ell k} X_{\ell k} \right)^2}{\sigma_i^2}, \qquad \frac{\partial f_0}{\partial X_{\ell k}} = \sum_i \frac{c_{i\ell k} \left(y_i - \sum_{h,a} c_{iha} X_{ha} \right)}{\sigma_i^2}. \tag{26}$$

At every clean one-hot field, $f_0(X) = r_y(X)$. The extension interpolates the precomputed attenuation coefficients $c_{i\ell k}$. It does not first interpolate the rain intensities and then apply the power law: generally $\sum_k X_{\ell k} a_k^{b_i} \neq (\sum_k X_{\ell k} a_k)^{b_i}$. It therefore introduces no derivative of the power law at

zero rainfall. For a fractional probability matrix X, $f_0(X)$ also differs from $\mathcal{R}(X)$ in Equation (19), which includes the attenuation variance. They agree at one-hot fields but can give different gradients there as continuous extensions. Here f_0 is the chosen extension of a sampled reward; it is not a substitute for the exact categorical expectation $\mathcal{R}(W^{0,t})$.

For $u > 0$, the reward first uses a new denoiser prediction at time u. Let E_j be the network embedding of class j, including the mask when present. Replace an embedding lookup at cell ℓ by the linear combination $\sum_j X_{\ell j} E_j$, and denote the resulting raw network probabilities by $D^u_{\theta,\ell k}(X)$. If the network directly consumes one-hot arrays, feed X to that input instead. At a categorical input, this is its usual evaluation before imposing the masked copy rule. At a fractional input, this is an explicitly chosen differentiable extension of the network. Define its clean-class output by

$$\overline{P}^u_{\ell k}(X) = \begin{cases} D^u_{\theta,\ell k}(X), & \text{uniform diffusion,} \\ X_{\ell k} + X_{\ell \mathtt{M}} D^u_{\theta,\ell k}(X), & \text{masked diffusion,} \end{cases} \qquad f_u(X) = \mathcal{R}(\overline{P}^u(X)). \tag{27}$$

For our samples, $X_{\ell \mathtt{M}} = 1 - B^{u,t}_\ell$ is fixed during differentiation. The clean output rows are nonnegative and sum to one. At hard inputs, $\overline{P}^u(X) = P^u(X)$ and $f_u(X) = \mathcal{R}_u(X)$. At a relaxed revealed cell, the clean vector is copied as a probability vector; at a masked cell, the network supplies its clean probabilities. The variance term in $\mathcal{R}$ must be retained: $f_0(\overline{P}^u(X))$ would generally give a different reward even at a hard midpoint input.

The per-cell rewards $H_{\ell k}$ in Equation (20) supply the differential of $\mathcal{R}$ along normalized probability rows. For an input perturbation δX with fixed mask entries and $\sum_k \delta X_{\ell k} = 0$ at every cell, define

$$\Gamma^u_{\ell k}(X) := \sum_{h,a} H_{ha}(\overline{P}^u(X)) \frac{\partial \overline{P}^u_{ha}(X)}{\partial X_{\ell k}}, \qquad \mathrm{d}f_u(X)[\delta X] = \sum_{\ell,k} \Gamma^u_{\ell k}(X)\, \delta X_{\ell k}. \tag{28}$$

This follows from $\mathrm{d}\mathcal{R} = \sum_{h,a} H_{ha}\, \mathrm{d}\overline{P}_{ha}$ along normalized output rows. The array $\Gamma^u(X)$ represents the input derivative along the allowed directions; adding a constant to one row has no effect and cancels under the softmax Jacobian. For masked diffusion with the mask entries fixed, the Jacobian $\partial \overline{P}^u_{ha} / \partial X_{\ell k}$ has a direct copy term $\mathbf{1}\{h = \ell, a = k\}$ and a network term $X_{h\mathtt{M}}\, \partial D^u_{\theta,ha} / \partial X_{\ell k}$. Thus changing the class of a newly revealed cell affects both its own copied prediction and potentially the predictions at other, still masked cells. An argmax or a stopped derivative in the copy term would remove the first contribution. The parameters θ are fixed, but derivatives with respect to the denoiser input must remain enabled. The current reference, computed from (z_t, t), is held constant and is not differentiated.

F.2.4 Soft and straight-through gradients

For the soft variant, evaluate f_u at $X^{\mathrm{R},u,t}$. For the straight-through variant, define

$$X^{\mathrm{ST},u,t} = X^{\mathrm{R},u,t} + \mathrm{stopgrad}(X^{\mathrm{H},u,t} - X^{\mathrm{R},u,t}), \tag{29}$$

where stopgrad keeps its argument's value but sets its derivative to zero. The forward value is $X^{\mathrm{H},u,t}$, while the derivative with respect to the logits is that of $X^{\mathrm{R},u,t}$. The chosen extension f_u, including the masked copy rule, must be used in both variants.

Holding the sampled Gumbel noise and Bernoulli indicators fixed gives $\partial Y^{u,t}_{\ell a} / \partial \phi^{u,t}_{\ell k} = Y^{u,t}_{\ell a}(\mathbf{1}\{a = k\} - Y^{u,t}_{\ell k})/\vartheta^{u,t}$. For sample b, let $\Gamma^{u,t}_{b,\ell k}$ be the input derivative of f_u from Equation (26) ($u = 0$) or Equation (28) ($u > 0$), evaluated at $X^{\mathrm{R},u,t}_b$ for the soft variant and at $X^{\mathrm{H},u,t}_b$ for straight-through. With B_t independent samples, the reward-gradient estimate is

$$\widehat{\mathcal{G}}^{u,t}_{\ell k} = \frac{1}{B_t} \sum_{b=1}^{B_t} \frac{B^{u,t}_{b,\ell}}{\vartheta^{u,t}} Y^{u,t}_{b,\ell k} \left(\Gamma^{u,t}_{b,\ell k} - \sum_{a \in S^\ell_{u,t}} Y^{u,t}_{b,\ell a} \Gamma^{u,t}_{b,\ell a} \right). \tag{30}$$

The batch size B_t and the Bernoulli indicators $B^{u,t}_{b,\ell}$ are different objects. A cell with $B^{u,t}_{b,\ell} = 0$ receives no class-logit derivative from that sample. However, the reward at that masked cell can contribute to derivatives at other cells through the denoiser. Fixed current classes have no trainable score and no update.

In the soft variant, $\mathsf{K}^{u,t} - \lambda_t \widehat{\mathcal{G}}^{u,t}$ is a Monte Carlo derivative of $\mathcal{F}^{\vartheta}_{u,t} = \mathcal{K}_{u,t} - \lambda_t \mathbb{E}_{g,B}[f_u(X^{\mathrm{R},u,t})]$. For fixed positive temperature, it is unbiased for the gradient of this relaxed objective when differentiation and expectation may be interchanged, for example with bounded input derivatives on the product of simplices. It is generally not unbiased for the original categorical objective. In the straight-through variant, the forward reward is an unbiased sample of the categorical reward expectation, but the backward expression replaces a discrete derivative by a continuous one and is generally biased. It need not be the gradient of any scalar objective. More samples reduce Monte Carlo variation but do not remove either relaxation bias. Although Y approaches U almost surely as the temperature tends to zero, this fact alone does not justify exchanging that limit and differentiation, or prove that the straight-through gradient becomes unbiased.

F.2.5 The score-function alternative

The same exact hard samples give an estimator that does not require a differentiable extension. Since the reveal probabilities are fixed, differentiation of the categorical log probability gives

$$\mathsf{S}^{u,t}_{\ell k}(X^{\mathrm{H},u,t}) = X^{\mathrm{H},u,t}_{\ell k} - B^{u,t}_{\ell} W^{u,t}_{\ell k}, \qquad \frac{\partial \mathcal{F}_{u,t}}{\partial \phi^{u,t}_{\ell k}} = \mathsf{K}^{u,t}_{\ell k} - \lambda_t \mathbb{E}_q[\mathcal{R}_u(X^{\mathrm{H},u,t}) \mathsf{S}^{u,t}_{\ell k}]. \tag{31}$$

This follows by differentiating the finite sum over categorical fields and using $\partial q = q\, \partial \log q$. The mean score is zero. For sample b, write $R^{u,t}_b = \mathcal{R}_u(X^{\mathrm{H},u,t}_b)$ and choose a baseline $\beta^{u,t}_b$ independent of that sample, conditional on the current parameters and earlier information. Then an unbiased reward-gradient estimate is $\widehat{\mathcal{G}}^{u,t}_{\ell k} = B_t^{-1} \sum_b (R^{u,t}_b - \beta^{u,t}_b) \mathsf{S}^{u,t}_{\ell k}(X^{\mathrm{H},u,t}_b)$. For $B_t > 1$, use $\beta^{u,t}_b = (\sum_{a \neq b} R^{u,t}_a)/(B_t - 1)$; for one sample, use zero or a baseline from previous iterations. Using the mean of the entire batch instead multiplies the expected reward gradient by $(B_t - 1)/B_t$. Treat the sampled field, reward, and baseline as constants when multiplying them by the score; no derivative through the denoiser is needed. This estimator targets the categorical objective itself, at the cost of sampling variation that can be large.

F.2.6 Shared algorithms and their use by the three methods

Algorithm 3 returns both the exact KL derivative and a chosen reward-gradient estimate. Algorithm 4 uses them for a local variational fit. Temperatures $\vartheta^{u,t}_j > 0$ and step sizes $\varepsilon^{u,t}_j > 0$ are indexed by optimization iteration j at fixed (u, t); they are distinct from the forward schedule α_t and observation strength λ_t. Changing temperature during optimization changes the relaxed objective in the soft variant. Subtract the maximum before each softmax and sample uniforms strictly inside $(0, 1)$. Do not add mass to a forbidden class to make its logarithm finite. Record the variant, temperature sequence, batch size, optimizer steps, initialization, and reward extension when comparing results.

G2D2. Set $u = 0$, $\chi^{0,t} = 1$, $A^{0,t} = P^t$, and $\phi^{0,t} = \phi^t$ from Appendix G.3. Algorithm 4 replaces the call to the exact coordinate fit (Algorithm 5) in Algorithm 6; the outer denoiser prediction and sampling transition stay as specified there. The soft variant follows the relaxation principle of Murata et al. (2025, Section 3.4), here with our specified Gaussian CML reward extension. The straight-through and REINFORCE variants are alternatives, not claims about that paper's implementation. For this CML likelihood, the analytic expectation and coordinate updates remain available without sampling or a temperature.

GILC. Set $u = 0$, $\chi^{0,t} = 1$, $A^{0,t} = W^{0,t} = P^t$, and $\phi^{0,t} = \eta^t$ on the eligible cells of Appendix H. Call Algorithm 3 once and update $\eta^{t,y} = \eta^t + \lambda_t \widehat{\mathcal{G}}^{0,t}$. Here λ_t is the score-correction step size, and there is no iterative variational fit. The KL derivative is zero at this starting point. The straight-through choice follows GILC-DB (Dou et al., 2026); the other choices specify alternative estimators of the same reward derivative. The REINFORCE option here uses an independent baseline; it is not a reproduction of the separate group-normalized GILC-PG implementation. For CML, Equation (40) supplies the exact derivative and thus a direct reference for assessing estimator error.

D-MGPS. Use the reference factors at the chosen u from Appendix I.3. Algorithm 4 can replace the local fit at either $u = 0$ or $u > 0$ in Algorithm 8. At $u > 0$, only the soft and straight-through variants require differentiation through the midpoint denoiser.

Algorithm 3 One shared reward-gradient estimate and the exact categorical KL derivative

Input: Fixed $(z_t, u, t, S^{u,t}, A^{u,t}, \chi^{u,t})$, current logits $\phi^{u,t}$, batch size B_t
Input: Variant: soft, straight-through, or REINFORCE; temperature $\vartheta^{u,t} > 0$ for the first two
Input: Reward f_u from Equations (26)–(27); fixed model parameters
Output: $\mathsf{K}^{u,t}$ from Equation (23) and estimated reward derivative $\widehat{\mathcal{G}}^{u,t}$

Compute $W^{u,t} = \mathrm{softmax}(\phi^{u,t})$ on each allowed support; retain fixed rows
Compute $\mathsf{K}^{u,t}$ exactly; set its entries outside $I_{u,t}$ to zero
for $b = 1, \ldots, B_t$ **do**
    Draw each $B_{b,\ell}^{u,t} \sim \mathrm{Bernoulli}(\chi_\ell^{u,t})$, including singleton supports
    Construct $X_b^{\mathrm{H},u,t}$ by Gumbel-max; use the fixed rows and zero-reveal rules above
    if variant is REINFORCE **then**
        Evaluate $R_b^{u,t} = \mathcal{R}_u(X_b^{\mathrm{H},u,t})$ without input derivatives
        Compute the categorical scores $\mathsf{S}_{\ell k}^{u,t}$ in Equation (31)
    else
        Construct $Y_b^{u,t}$ and $X_b^{\mathrm{R},u,t}$ using the same class noise and mask indicators
        Use $X_b^{\mathrm{R},u,t}$ for soft, or $X_b^{\mathrm{ST},u,t}$ for straight-through
        Evaluate f_u and its input derivatives, keeping the sampled mask indicators fixed
        Compute the summand in Equation (30), without its $1/B_t$ factor
    end if
end for
Average the sample derivatives; for REINFORCE use the independent baselines in Appendix F.2.5
Set derivatives outside $I_{u,t}$ to zero and return $(\mathsf{K}^{u,t}, \widehat{\mathcal{G}}^{u,t})$

Algorithm 4 Local variational fit using the common gradient estimator

Input: Fixed reference and supports from Algorithm 3; observation strength $\lambda_t \geq 0$
Input: Positive initial $W^{u,t}$ on allowed supports (default $A^{u,t}$), M_t iterations
Input: Estimator variant, batch size B_t, steps $\varepsilon_j^{u,t}$, temperatures $\vartheta_j^{u,t}$ when needed
Input: CML reward inputs and fixed denoiser from Algorithm 3
Output: Fitted categorical probabilities $V^{u,t}$ with the prescribed mask probabilities

if $\lambda_t = 0$ **then**
    Set $W^{u,t} \leftarrow A^{u,t}$ and return the reference factors
end if
Set $\phi_{\ell k}^{u,t} \leftarrow \log W_{\ell k}^{u,t}$ for adjustable cells and allowed classes
for $j = 1, \ldots, M_t$ **do**
    Obtain $(\mathsf{K}^{u,t}, \widehat{\mathcal{G}}^{u,t})$ from Algorithm 3, using fresh draws
    $\phi^{u,t} \leftarrow \phi^{u,t} - \varepsilon_j^{u,t}(\mathsf{K}^{u,t} - \lambda_t \widehat{\mathcal{G}}^{u,t})$ on adjustable entries
    Subtract each adjustable row's mean score; keep the reference and mask probabilities fixed
end for
Compute $W^{u,t} = \mathrm{softmax}(\phi^{u,t})$ on adjustable rows; retain every other row
Form $V^{u,t}$ by Equation (22)
return $V^{u,t}$ ▷ Final sampling uses these categorical probabilities, not a relaxed draw

At $u > 0$, only the outer expectation over midpoint fields is sampled; the expectation over clean predictions remains the exact CML function $\mathcal{R}$. At $u = 0$, the shared analytic coordinate fit is also available. After fitting, sample a new categorical field with the fitted probabilities, independently of all optimization draws, and apply the exact categorical bridge $q_{s|u,t}$ to time s. No relaxed field is passed to that bridge. The same rule applies to the final clean draw at $s = 0$.

Finite gradient updates need not decrease even their chosen local objective at every iteration. These estimator choices do not change the posterior approximation issues discussed in the method-specific appendices.

G G2D2 FOR THE CML LIKELIHOOD

G2D2 (Murata et al., 2025) solves inverse problems with a discrete diffusion prior; it is not limited to observations of classes at selected cells. Here we apply its variational objective to the path-integrated Gaussian likelihood in Equation (1). The observations y_i are link attenuations, not cell labels. We use the common CML functions $\mathcal{R}(\Xi)$ and $H_{\ell k}(\Xi)$ from Appendix F rather than repeating their derivation. They allow us to evaluate the categorical objective and its derivatives exactly and to minimize it over one cell at a time. We give both an analytic CML fit and sampled fits based on the common estimators of Appendix F.2. That appendix specifies, with our CML reward, the soft Gumbel–Softmax construction used in the original paper, its straight-through alternative, and the categorical score-function alternative. We specify both the original G2D2 sampling transition and its Markov variant, because they give different sampling laws.

G.1 THE CURRENT STATE AND THE CLEAN-CLASS PROBABILITIES

Use the reverse grid $T = t_N > \cdots > t_0 = 0$ and the schedule of Appendix E.1: $\alpha_0 = 1$, $\alpha_T = 0$, and $0 < \alpha_t < 1$ at interior times. At step n, the current time is $t = t_n$ and the next time is $s = t_{n-1}$. The current field z_t contains one categorical state at each of the L cells. As in GILC and D-MGPS, evaluate the fixed denoiser once at (z_t, t) and write $P^t_{\ell k} = P^\theta_{\ell k}(z_t, t)$ for its probability of clean class k at cell ℓ. These probabilities describe classes at time zero; t records when the prediction was made. The resulting clean-field prediction is $p^\theta_{0|t}(z_0 \mid z_t) = \prod_\ell P^t_{\ell k_\ell}$ for $z_0 = (e_{k_1}, \ldots, e_{k_L})$.

For pure masked corruption, an unmasked current cell $z^\ell_t = e_h$ forces its clean class to be h. We therefore set $P^t_{\ell k} = \mathbf{1}\{k = h\}$ at that cell. This rule follows from the corruption process, not from a point observation in y. For uniform corruption, the current class does not determine the clean class, and this rule does not apply. In both cases, let $S^\ell_t = \{k : P^t_{\ell k} > 0\}$ be the allowed clean classes. Let $F_t = \{\ell : |S^\ell_t| > 1\}$ contain the cells whose probabilities can change during the fit.

Write $V^t = (V^t_{\ell k})$ for the clean-class probabilities that we will fit using y. They define $q_{V^t}(z_0 \mid z_t, y) = \prod_\ell V^t_{\ell k_\ell}$, with each row nonnegative and summing to one. The superscript t identifies the reverse step; V^t changes during its local optimization, whereas P^t stays fixed. We require $V^t_{\ell k} = 0$ outside S^ℓ_t, since otherwise the KL divergence to the denoiser prediction is infinite. If only one class is allowed, its probability stays equal to one. By default, we initialize $V^t = P^t$ at every reverse step. All probability rows depend on the complete current field z_t, even though their product assumes independence across cells at this fixed step.

G.2 THE LOCAL VARIATIONAL OBJECTIVE

Let $r_y(z_0)$ be the Gaussian log-likelihood with its constant independent of z_0 removed, as defined in Appendix F. Choose an observation strength $\lambda_t \geq 0$ to set the weight of the log-likelihood at reverse time t. Define the local probability distribution

$$\widetilde{\pi}^t(z_0 \mid z_t, y) = \frac{p^\theta_{0|t}(z_0 \mid z_t)\exp\{\lambda_t r_y(z_0)\}}{\mathcal{N}_t}, \qquad \mathcal{N}_t = \sum_z p^\theta_{0|t}(z \mid z_t)\exp\{\lambda_t r_y(z)\}. \tag{32}$$

The Gaussian likelihood is positive and finite because every $\sigma_i > 0$ and the codebook is finite, so $\mathcal{N}_t > 0$. At $\lambda_t = 1$, this distribution reweights the denoiser prediction by the stated likelihood. It is a local approximation: the denoiser prediction need not be the exact clean-field conditional of the prior p^θ_0.

We fit q_{V^t} by minimizing its KL divergence to $\widetilde{\pi}^t$. To obtain an objective that can be evaluated, insert Equation (32) into the definition of the KL divergence:

$$\begin{aligned}\mathrm{KL}(q_{V^t} \| \widetilde{\pi}^t) &= \sum_{z_0} q_{V^t}(z_0 \mid z_t, y)\left[\log \frac{q_{V^t}(z_0 \mid z_t, y)}{p^\theta_{0|t}(z_0 \mid z_t)} - \lambda_t r_y(z_0) + \log \mathcal{N}_t\right] \\ &= \sum_\ell \sum_{k \in S^\ell_t} V^t_{\ell k} \log \frac{V^t_{\ell k}}{P^t_{\ell k}} - \lambda_t \mathcal{R}(V^t) + \log \mathcal{N}_t.\end{aligned} \tag{33}$$

The sum over z_0 includes only fields with $q_{V^t}(z_0 \mid z_t, y) > 0$; zero-weight terms are omitted, so every displayed log ratio is defined. The first term separates by cell because both q_{V^t} and the denoiser prediction are products. For example, in the expectation of $\log(V^t_{\ell k_\ell}/P^t_{\ell k_\ell})$, the probabilities of all other cells sum to one, leaving only the sum over k_ℓ. The reward term is $\mathcal{R}(V^t)$ by definition of the common expectation. The last term is independent of V^t and can be omitted when minimizing, but is needed for the KL identity itself. We use the convention $0 \log(0/p) = 0$ for $p > 0$.

Substituting Equation (19) and omitting $\log \mathcal{N}_t$ gives the CML objective

$$\mathcal{F}_{0,t}(V^t) = \sum_\ell \sum_{k \in S_t^\ell} V^t_{\ell k} \log \frac{V^t_{\ell k}}{P^t_{\ell k}} + \frac{\lambda_t}{2} \sum_i \frac{(y_i - \mu_i(V^t))^2 + v_i(V^t)}{\sigma_i^2}. \tag{34}$$

The zero in $\mathcal{F}_{0,t}$ indicates that the fitted distribution is over clean fields, consistent with the D-MGPS notation. The functions $m_{i\ell}(V^t)$, $\mu_i(V^t)$, and $v_i(V^t)$ are those of Appendix F, evaluated at the current fitted probabilities, not at P^t. Fixed cells still contribute their assigned attenuation to the link means. At $\lambda_t = 0$, the initial value $V^t = P^t$ is already a global minimum.

Although the likelihood is a product over links, it is generally not a product over cells. For a link, expanding $(y_i - \sum_\ell c_{i\ell k_\ell})^2$ includes terms $2c_{i\ell k_\ell} c_{ihk_h}$ for distinct cells ℓ and h. Their product-distribution averages are $2m_{i\ell}(V^t)m_{ih}(V^t)$. Thus changing one cell's probabilities changes the objective terms involving other cells on the same link. The fit cannot in general be solved by independent reweighting at all cells. By contrast, an independent categorical observation at each measured cell gives a sum of separate cell losses; this in-filling case has a direct solution for each row. It is a special case of G2D2, not a restriction of the method.

G.3 Exact derivatives and sequential probability updates

For each cell $\ell \in F_t$, let $\phi^t_\ell = (\phi^t_{\ell k})_{k \in S^\ell_t}$ be real class scores, called logits, that parameterize the fitted probabilities through $V^t_\ell = \mathrm{softmax}_{S^\ell_t}(\phi^t_\ell)$. The superscript t identifies the reverse step; these scores change during the local fit, while the denoiser probabilities P^t remain fixed. Let $d^t_\ell = \sum_{j \in S^\ell_t} V^t_{\ell j} \log(V^t_{\ell j}/P^t_{\ell j})$. Differentiating the row KL gives $V^t_{\ell k}[\log(V^t_{\ell k}/P^t_{\ell k}) - d^t_\ell]$: the constant one in the derivative of $v \log(v/p)$ cancels under the softmax derivative. Equation (21) gives the reward derivative, so the complete score derivative is

$$\frac{\partial \mathcal{F}_{0,t}}{\partial \phi^t_{\ell k}} = V^t_{\ell k} \left[\log \frac{V^t_{\ell k}}{P^t_{\ell k}} - d^t_\ell - \lambda_t \left(H_{\ell k}(V^t) - \sum_{j \in S^\ell_t} V^t_{\ell j} H_{\ell j}(V^t) \right) \right]. \tag{35}$$

Every term can be evaluated by finite sums. During gradient optimization, the link moments and $H(V^t)$ must be recomputed after changing V^t; the reference P^t must not be changed. Neither the denoiser parameters nor the categorical field z_t is differentiated.

Minimizing over one cell. A second option requires no gradient step size. Hold all rows except ℓ fixed. The common calculation in Equation (20) writes the expected reward as $C_\ell(V^t) + \sum_k V^t_{\ell k} H_{\ell k}(V^t)$. Both C_ℓ and the candidate values $H_{\ell k}$ are independent of the row being optimized. After dropping terms independent of that row, its objective is

$$f^t_\ell(w) = \sum_{k \in S^\ell_t} w_k \log \frac{w_k}{P^t_{\ell k}} - \lambda_t \sum_{k \in S^\ell_t} w_k H_{\ell k}(V^t), \qquad w_k \geq 0, \quad \sum_k w_k = 1. \tag{36}$$

At a minimum with positive probabilities on the allowed classes, a multiplier ν for the sum constraint gives $\log(w_k/P^t_{\ell k}) + 1 - \lambda_t H_{\ell k}(V^t) + \nu = 0$. Exponentiating and choosing ν to make the probabilities sum to one gives

$$V^{t,\mathrm{new}}_{\ell k} = \frac{P^t_{\ell k} \exp\{\lambda_t H_{\ell k}(V^t)\}}{\sum_{j \in S^\ell_t} P^t_{\ell j} \exp\{\lambda_t H_{\ell j}(V^t)\}}, \qquad k \in S^\ell_t. \tag{37}$$

Every allowed reference probability is positive, and all candidate rewards are finite, so this row has positive entries on its support. It is the unique global minimum over that row, including its boundary: substituting the formula back into f^t_ℓ gives $f^t_\ell(w) - f^t_\ell(V^{t,\mathrm{new}}_\ell) = \mathrm{KL}(w \| V^{t,\mathrm{new}}_\ell) \geq 0$. This identity proves that replacing a single row cannot increase the full objective in exact arithmetic.

Algorithm 5 Shared CML fit over clean-class probabilities at reverse time t

Input: Fixed denoiser probabilities P^t, observations y, noise scales $\sigma_i > 0$, coefficients c
Input: Observation strength $\lambda_t \geq 0$, number of sweeps $M_t \geq 0$, specified cell order
Input: Initial probabilities $V^{t,\mathrm{init}}$ with the same support as P^t (default: P^t)
Output: Fitted probabilities $V^{t,\mathrm{fit}}$ for use by G2D2 or D-MGPS at $u = 0$

Set $S_t^\ell \leftarrow \{k : P_{\ell k}^t > 0\}$ and $F_t \leftarrow \{\ell : |S_t^\ell| > 1\}$
if $\lambda_t = 0$ **then**
  return P^t ▷ The exact minimum without the observation term
end if
Initialize $V^t \leftarrow V^{t,\mathrm{init}}$ ▷ Keep the reference P^t fixed throughout
if $M_t = 0$ or $F_t = \varnothing$ **then**
  return $V^{t,\mathrm{fit}} \leftarrow V^t$
end if
$m_{i\ell}^t \leftarrow \sum_k V_{\ell k}^t c_{i\ell k}$; $\mu_i^t \leftarrow \sum_\ell m_{i\ell}^t$
for $j = 1, \ldots, M_t$ **do**
  for each $\ell \in F_t$ in the specified order **do**
    Set $\mu_{i,-\ell}^t \leftarrow \mu_i^t - m_{i\ell}^t$ for links crossing cell ℓ
    for $k \in S_t^\ell$ **do**
      $H_{\ell k}^t \leftarrow \sum_{i:\Delta_\ell^i > 0} [(y_i - \mu_{i,-\ell}^t) c_{i\ell k} - c_{i\ell k}^2/2]/\sigma_i^2$
      $\phi_{\ell k}^t \leftarrow \log P_{\ell k}^t + \lambda_t H_{\ell k}^t$
    end for
    $V_\ell^{t,\mathrm{new}} \leftarrow \mathrm{softmax}_{S_t^\ell}(\phi_\ell^t)$
    for each link i crossing cell ℓ **do**
      $m_{i\ell}^{t,\mathrm{new}} \leftarrow \sum_k V_{\ell k}^{t,\mathrm{new}} c_{i\ell k}$
      $\mu_i^t \leftarrow \mu_i^t - m_{i\ell}^t + m_{i\ell}^{t,\mathrm{new}}$; $m_{i\ell}^t \leftarrow m_{i\ell}^{t,\mathrm{new}}$
    end for
    Replace $V_\ell^t \leftarrow V_\ell^{t,\mathrm{new}}$ before visiting the next cell
  end for
end for
return $V^{t,\mathrm{fit}} \leftarrow V^t$ ▷ Sampling is performed after the fit

After replacing row ℓ, update $m_{i\ell}$ and μ_i for every link crossing that cell before processing the next row. This ensures that the candidate rewards of the next row use the latest probabilities at all other cells. Updating every row simultaneously from the old means is a different procedure and does not have this descent guarantee. For a cell crossed by no link, all $H_{\ell k}$ are zero and the update returns P_ℓ^t. For a cell with only one allowed class, no update is required. The full objective is generally not convex in all rows together. Its values do not increase under sequential updates and are bounded below, but this establishes neither a global minimum nor exact posterior sampling.

The quantities m^t, μ^t, and H^t in Algorithm 5 are stored values of the common functions evaluated at the changing V^t. The time t stays fixed during all sweeps. The default initialization is P^t; the optional initial matrix allows the reuse described for D-MGPS in Appendix I.6 without changing the reference or the objective. Every initial row must sum to one and be positive exactly on S_t^ℓ; rows with only one allowed class remain fixed. One sweep visits every eligible cell once. With sparse link–cell intersections, a sweep costs $O(K\,\mathrm{nnz}(\Delta) + LK)$ operations, where $\mathrm{nnz}(\Delta)$ counts the intersections with $\Delta_\ell^i > 0$; no denoiser evaluation is needed inside it. Evaluate each softmax by subtracting its largest allowed score before exponentiating, and retain zero probability outside the support. Report the sweep count, cell order, strengths, and any stopping rule when evaluating the method.

G.4 Generating the state at the next time

After fitting, draw a new clean field $\widehat{z}_0$ with independent cell probabilities $V^{t,\mathrm{fit}}$. This field is an auxiliary draw for the current step; it is not kept as the final reconstruction unless $s = 0$. G2D2 and its Markov variant differ in how they use this draw to generate z_s.

The original G2D2 transition. The original method uses a corruption construction in which the states at different positive times are independent conditional on the clean field (Murata et al., 2025, Section 3.1). It is called star-shaped corruption because each noisy state is drawn directly from the same clean field. Its conditional law at each single time is still $q_{t|0}$ from Equation (7). Thus, given $\widehat{z}_0$, generate $z_s \sim q_{s|0}(\cdot \mid \widehat{z}_0)$ without conditioning this second draw on z_t. After summing over $\widehat{z}_0$, the cell transition is

$$K^{t,y}_{s,\star}(z^\ell_s = e_j \mid z_t) = \sum_{k \in S^\ell_t} V^{t,\mathrm{fit}}_{\ell k}\big[\alpha_s \mathbf{1}\{j = k\} + (1-\alpha_s)\bar{q}_j\big] = \alpha_s V^{t,\mathrm{fit}}_{\ell j} + (1-\alpha_s)\bar{q}_j. \tag{38}$$

For masked corruption, append a zero mask entry to $V^{t,\mathrm{fit}}_\ell$ in this formula. Every cell is masked at the next time with probability $1-\alpha_s$, even if it is currently unmasked. An unmasked cell retains its current clean class whenever it is not remasked; once remasked, a later fit can assign it another class. For uniform corruption, the replacement draws uniformly over the clean classes. The full transition is the product of the displayed cell transitions, conditional on the fitted probabilities and z_t.

The Markov variant. The alternative in Murata et al. (2025, Appendix E.3) keeps the forward Markov process and draws $z_s \sim q_{s|0,t}(\cdot \mid \widehat{z}_0, z_t)$. Its cell transition is

$$K^{t,y}_{s,\mathrm{Markov}}(z^\ell_s \mid z_t) = \sum_{k \in S^\ell_t} V^{t,\mathrm{fit}}_{\ell k} q_{s|0,t}(z^\ell_s \mid e_k, z^\ell_t). \tag{39}$$

Each bridge is normalized for its specified clean class before it is averaged; Appendix E.2 gives its probabilities. The support rules above ensure that the conditioned endpoints are compatible. For masked corruption, a currently masked cell is revealed with probability $(\alpha_s - \alpha_t)/(1-\alpha_t)$, with its clean class drawn from $V^{t,\mathrm{fit}}_\ell$. A currently unmasked cell remains unchanged. For uniform corruption, sample the bridge using Equation (10); the current class may change. These are the same bridge operations used by GILC, with fitted probabilities in place of the probabilities obtained by a single score correction.

At $s = 0$, both choices return $\widehat{z}_0$. At earlier times they generally differ, even if their fitted clean probabilities are identical. In particular, the original transition uses α_s, whereas the Markov bridge uses the conditional schedule involving both s and t. The same numerical time grid does not make their reveal or remasking probabilities equal.

G.5 Relation to the other samplers and scope of the guarantees

What is shared and what differs. At $V^t = P^t$, every d^t_ℓ and log probability ratio in Equation (35) is zero. One gradient descent step on the scores, with step size $\delta > 0$, therefore adds $\delta\lambda_t$ times the reward score gradient to $\log P^t$. With the Markov transition, this is the analytic GILC step of Appendix H, with correction strength $\delta\lambda_t$. Later gradient steps include a generally nonzero KL derivative and therefore differ from repeatedly applying a reward correction alone. The exact row update in Equation (37) is also different from one such gradient step.

For D-MGPS at $u = 0$, the reference is $\rho^\ell_{0|t}(k \mid z_t) = P^t_{\ell k}$ and the midpoint reward is the clean-field log-likelihood. Its objective is exactly Equation (34). We therefore use Algorithm 5 for both methods. If the support, strength, initialization, sweep count, and cell order agree, this fit followed by the Markov bridge gives the same reverse transition for both methods. The original star-shaped G2D2 transition does not satisfy this equivalence. At $u > 0$, D-MGPS optimizes a distribution over states at the intermediate time u and evaluates new denoiser predictions there; this additional construction is given in Appendix I.

Difference from the published implementation. The original G2D2 implementation estimates the likelihood term through Gumbel–Softmax samples and optimizes the variational scores with a gradient optimizer (Murata et al., 2025, Sections 3.3–3.4). It also allows initialization from probabilities fitted at the preceding reverse step. As written, Algorithm 6 uses the exact categorical CML expectation, sequential row minimization, and initialization from the current prediction alone. These choices remove the need for a relaxation temperature, likelihood-gradient sampling, and differentiation through the rain-intensity power law. They define an analytic adaptation, which must be distinguished from a reproduction of the published implementation. Our G2D2 experiments follow

Algorithm 6 G2D2 adaptation to CML data with the shared analytic fit

Input: Fixed denoiser θ, codebook $\mathcal{C}$, uniform or masked corruption, schedule α
Input: CML observations y, coefficients (κ, b, Δ), noise scales $\sigma_i > 0$
Input: Times $T = t_N > \cdots > t_0 = 0$, strengths λ_t, sweep counts M_t, cell order
Input: Transition choice: original star-shaped or Markov variant
Output: One approximate conditional rain field $\widetilde{x}$

Precompute $c_{i\ell k} = \kappa_i \Delta_\ell^i a_k^{b_i}$ at all link–cell intersections
Draw $z_T^\ell \sim \mathrm{Cat}(\bar{q})$ independently at each cell
for $n = N, \ldots, 1$ **do**
    $(t, s) \leftarrow (t_n, t_{n-1})$
    Evaluate $P^t \leftarrow P^\theta(z_t, t)$ once
    if corruption is masked **then**
        Set $P^t_{\ell k} \leftarrow \mathbf{1}\{k = h\}$ whenever $z_t^\ell = e_h$ is unmasked
    end if
    Fit $V^{t,\mathrm{fit}}$ by Algorithm 5 using $(P^t, y, \sigma, c, \lambda_t, M_t)$ and the given cell order
    Draw $\widehat{z}_0^\ell \sim \mathrm{Cat}(V_\ell^{t,\mathrm{fit}})$ independently over cells
    if the transition choice is star-shaped **then**
        Draw $z_s \sim q_{s|0}(\cdot \mid \widehat{z}_0)$ ▷ Equation (7)
    else
        Draw $z_s \sim q_{s|0,t}(\cdot \mid \widehat{z}_0, z_t)$ ▷ Appendix E.2
    end if
end for
return $\widetilde{x} = \mathcal{D}(z_0)$

Algorithm 1 of Murata et al. (2025), with the Markov transition and initialization from the current prediction (Appendix B and Table 4). For a sampled fit, replace the call to Algorithm 5 by Algorithm 4, with $u = 0$, $A^{0,t} = P^t$, and $\chi^{0,t} = 1$. Appendix F.2 gives, once for all methods, the reward extension, the exact KL derivative, the temperature and step-size indices, and the complete soft, straight-through, and REINFORCE procedures. The selected star-shaped or Markov transition is then applied to a fresh draw from the fitted categorical probabilities.

Exact identities do not make the implemented sampler exact. To define the ideal comparison, start with the prior p_0^θ from the main text and generate each noisy state directly from Z_0 using $q_{t|0}$. The states at different positive times are then independent conditional on Z_0, and the observations depend only on Z_0 through Equation (1). The ideal clean-field conditional given (z_t, y) follows by Bayes' rule in this model. If this conditional could be used in place of q_{V^t} at every reverse step, the star-shaped transition would map the posterior marginal at t to that at s. Indeed, averaging that conditional over the posterior law of Z_t returns $\pi_0^\theta(\cdot \mid y)$, and subsequently applying $q_{s|0}$ returns its corrupted marginal at s. Since $\alpha_T = 0$, the terminal posterior is $\bar{q}^{\otimes L}$, which is also the initialization in Algorithm 6. Induction over the reverse grid then gives the target clean posterior at time zero. This is the ideal construction underlying Murata et al. (2025, Lemma D.1). Their Theorem 3.1 bounds the endpoint KL error by a sum of local KL errors relative to the ideal clean-field conditionals, averaged over the states generated by the approximate sampler. It does not establish that fitting to the denoiser-based distribution in Equation (32) makes those errors vanish.

In our adaptation, the denoiser may differ from the prior's clean-field conditional, the product family generally cannot represent the dependencies induced by link measurements, and finitely many sweeps need not find a global minimum. Even a global minimum within that family would not generally equal the local distribution $\widetilde{\pi}^t$. The exact expectation, score gradient, and row minimizer remove errors in those calculations only. Furthermore, a strength λ_t that changes across steps need not correspond to one fixed power of the likelihood in the final posterior.

With $\lambda_t = 0$, the Markov variant recovers the learned unguided reverse step because $V^t = P^t$. The star-shaped choice instead recovers its own unguided step, based on $q_{s|0}$. For an approximate denoiser, the clean output law of this star-shaped unguided sampler need not equal the prior p_0^θ defined by the learned Markov sampler in the main text. This distinction must be retained when comparing posterior-sampling claims for the two transitions. Finally, product distributions at individual steps

do not imply a product law for the final field: each new denoiser prediction depends on the complete sampled grid.

H Analytic CML guidance with GILC

GILC (Dou et al., 2026) conditions a pretrained discrete diffusion model on observations by correcting its predicted class scores once at each sampling step, with the model parameters fixed. We apply the common CML expectation and score derivative from Appendix F to the denoiser prediction at each time, and then specify the resulting sampling step. Our GILC experiments follow the published algorithms; their implementation is described in Appendix B.

H.1 The prediction used at the current time

Use the reverse grid, current state z_t, and clean-class probabilities P^t defined in Appendix G.1. At step n, $t = t_n$ is the current diffusion time and $s = t_{n-1} < t$ is the next time; neither is the physical observation time. For each cell not fixed by masked corruption, the denoiser returns K finite real scores $\eta^t_{\ell k}$, called logits, with $\eta^t = \mathrm{Logits}_\theta(z_t, t)$ and $P^t_\ell = \mathrm{softmax}(\eta^t_\ell)$. Apply the same probability-one constraints at revealed cells as in Appendix G.1. With finite scores at the other cells, the set F_t defined there contains exactly the currently masked cells for masked diffusion, or all cells for uniform diffusion. GILC corrects the scores once at these cells, converts them to probabilities, and samples the next state. The current (z_t, t) and the denoiser parameters remain fixed during this correction.

H.2 Applying the shared CML derivative

Evaluate the common functions of Appendix F at $\Xi = P^t$. Their time-indexed values are $m^t_{i\ell} = m_{i\ell}(P^t)$, $\mu^t_i = \mu_i(P^t)$, $v^t_i = v_i(P^t)$, $C^t_\ell = C_\ell(P^t)$, and $H^t_{\ell k} = H_{\ell k}(P^t)$. Write $J_t = (J^1_t, \ldots, J^L_t)$ for hypothetical clean class indices drawn independently with probabilities $\Pr(J^\ell_t = k) = P^t_{\ell k}$. The index t identifies the denoiser prediction used for this draw; J_t is a clean field, not the corrupted state z_t. The analytic calculation averages over these classes without sampling them.

Let $\mathcal{R}_t(\eta^t) = \mathcal{R}(P^t)$, with the free probability rows obtained from the scores and all fixed rows held constant. Equation (21), evaluated at $\Xi = P^t$ and $\phi = \eta^t$, gives

$$G^t_{\ell k} = \frac{\partial \mathcal{R}_t(\eta^t)}{\partial \eta^t_{\ell k}} = P^t_{\ell k}\big(H^t_{\ell k} - \sum_j P^t_{\ell j} H^t_{\ell j}\big). \tag{40}$$

For fixed cells, set $G^t_\ell = 0$. A class receives a positive correction when its expected reward, with that class fixed, exceeds the current average at the cell. This derivative is evaluated once at P^t; GILC does not repeatedly fit a probability matrix V^t as in Appendix G.3. Appendix G.5 establishes the relation between this correction and the first gradient step of that fit.

The attenuation variance must be retained: omitting it can reverse the sign of this correction. For one cell with codebook $\{0, 1\}$, unit attenuation coefficient, noise standard deviation one, and observation $y = 0.4$, let $p_t = P^t_{1,1}$ be the clean wet-class probability predicted at time t. The exact expected reward is $-\frac{1}{2}[(1 - p_t)0.4^2 + p_t 0.6^2]$, with derivative -0.1 with respect to p_t. Using only the mean attenuation gives $-\frac{1}{2}(0.4 - p_t)^2$, whose derivative is $0.4 - p_t$. At $p_t = 0.1$, this replacement favors increasing the wet probability, although the dry class has the larger likelihood.

H.3 Correcting the probabilities and generating the next state

For each eligible cell, apply the single update $\eta^{t,y}_\ell = \eta^t_\ell + \lambda_t G^t_\ell$, with guidance strength $\lambda_t \geq 0$, and set $P^{t,y}_\ell = \mathrm{softmax}(\eta^{t,y}_\ell)$. The superscripts (t, y) indicate that the correction uses the observations y and the denoiser prediction made at time t. For a fixed cell $\ell \notin F_t$, set $P^{t,y}_\ell = P^t_\ell$. The positive sign gives the direction of increasing expected reward; a finite update need not increase it, and a zero gradient leaves the scores unchanged. All gradients are computed from the original P^t before any sampling or correction is performed. The observation errors σ_i define the likelihood, whereas λ_t controls the size of this algorithmic update.

Algorithm 7 CML reconstruction with one GILC correction per masked reverse step

Input: Pretrained denoiser θ, fixed codebook $\mathcal{C}$, link parameters (κ, b, Δ)
Input: Observations y, standard deviations $\sigma_i > 0$, strengths $\lambda_t \geq 0$
Input: Times $T = t_N > \cdots > t_0 = 0$ and nonincreasing schedule α as defined above
Output: One approximate conditional rain field $\widetilde{x}$

Precompute $c_{i\ell k} = \kappa_i \Delta_\ell^i a_k^{b_i}$ for cells crossed by links
Treat omitted coefficients as zero; initialize every cell of z_T to the mask
for $n = N, \ldots, 1$ **do**
    Set $t \leftarrow t_n$ and $s \leftarrow t_{n-1}$ ▷ Current and next diffusion times
    Predict clean classes from the current grid
    Evaluate the denoiser once: $\eta^t \leftarrow \mathrm{Logits}_\theta(z_t, t)$
    $F_t \leftarrow \{\ell : z_t^\ell = e_{\mathtt{M}}\}$ ▷ Cells that are still masked
    Set $P_\ell^t \leftarrow \mathrm{softmax}(\eta_\ell^t)$ for $\ell \in F_t$
    For $\ell \notin F_t$ with $z_t^\ell = e_h$, set $P_{\ell k}^t \leftarrow \mathbf{1}\{k = h\}$
    Compute the reward gradient from these probabilities
    $m_{i\ell}^t \leftarrow \sum_k P_{\ell k}^t c_{i\ell k}$ and $\mu_i^t \leftarrow \sum_\ell m_{i\ell}^t$
    For $\ell \in F_t$ and every k, compute
    $H_{\ell k}^t \leftarrow \sum_{i:\Delta_\ell^i > 0} \big((y_i - \mu_i^t + m_{i\ell}^t) c_{i\ell k} - \frac{1}{2} c_{i\ell k}^2\big)/\sigma_i^2$
    $G_{\ell k}^t \leftarrow P_{\ell k}^t (H_{\ell k}^t - \sum_j P_{\ell j}^t H_{\ell j}^t)$ for $\ell \in F_t$
    Apply one correction, then generate the grid at time s
    $z_s \leftarrow z_t$ ▷ Already assigned clean classes remain unchanged
    for $\ell \in F_t$ **do**
        $\eta_\ell^{t,y} \leftarrow \eta_\ell^t + \lambda_t G_\ell^t$; $P_\ell^{t,y} \leftarrow \mathrm{softmax}(\eta_\ell^{t,y})$
        Draw $B_{s|t}^\ell \sim \mathrm{Bernoulli}((\alpha_s - \alpha_t)/(1 - \alpha_t))$
        if $B_{s|t}^\ell = 1$ **then**
            Draw clean class k with probabilities $P_\ell^{t,y}$; set $z_s^\ell \leftarrow e_k$
        else
            Keep $z_s^\ell \leftarrow e_{\mathtt{M}}$
        end if
    end for
end for
return $\widetilde{x} = \mathcal{D}(z_0)$ ▷ Replace each clean class k by intensity a_k

Use the Markov transition in Equation (39), replacing the fitted probabilities $V^{t,\mathrm{fit}}$ by the corrected probabilities $P^{t,y}$. For a currently masked cell, this gives the following time-indexed transition, with a zero mask coordinate appended to $P_\ell^{t,y}$:

$$p_{s|t}^y(z_s^\ell \mid z_t) = \mathrm{Cat}\left(z_s^\ell; \frac{(\alpha_s - \alpha_t) P_\ell^{t,y} + (1 - \alpha_s) e_{\mathtt{M}}}{1 - \alpha_t}\right), \qquad z_t^\ell = e_{\mathtt{M}}. \tag{41}$$

A cell already assigned a clean class is copied without change. The observations therefore affect which clean class is drawn, while the probability of removing a mask remains determined by the diffusion schedule. Given z_t and the computed probabilities, the draws at different cells are independent. At the final step $s = 0$, every remaining mask is removed because $\alpha_0 = 1$.

Algorithm 7 keeps the denoiser parameters fixed throughout. The corrected scores $\eta^{t,y}$ and probabilities $P^{t,y}$ are used only for the transition from t to s. If $s > 0$, the next reverse step evaluates $\eta^s = \mathrm{Logits}_\theta(z_s, s)$ and computes P^s and the score gradient G^s from this new prediction; it does not reuse $\eta^{t,y}$ as the denoiser scores at time s. The attenuation moments for that step are $m_{i\ell}^s$, μ_i^s, and v_i^s, defined under P^s. Setting $\lambda_t = 0$ recovers the unguided sampler, and independent runs give an ensemble of approximate conditional reconstructions. In the notation of Dou et al. (2026), choosing $\lambda_t = 1/\beta$ for a fixed $\beta > 0$ matches their score-correction strength. This correspondence does not imply that the complete sampler generates exactly from a prior multiplied by $p(y \mid z_0)^{1/\beta}$.

One correction uses the sparse link–cell sums of Appendix G.3 and costs $O(K\,\mathrm{nnz}(\Delta))$ operations, in addition to the denoiser evaluation and the probability normalization over the grid. Only m^t and

μ^t are needed to compute H^t and G^t; computing the variance v^t is necessary only when evaluating the reward itself.

H.4 Using the correction with uniform diffusion

Initialize each cell of z_T independently and uniformly over the K clean classes. At each reverse step, compute $P^{t,y}$ by the same score correction, now at every cell. Then use Equation (39) with $V^{t,\mathrm{fit}} = P^{t,y}$: draw a clean class k with probabilities $P_\ell^{t,y}$, followed by the uniform bridge in Equation (10), conditional on e_k and z_t^ℓ. The normalization of each bridge and the order of these two draws are specified in Appendix G.4. In our GILC experiments with uniform diffusion, the implementation computes this transition in the equivalent plug-in form: the corrected probabilities, converted to their leave-one-out form, are substituted directly into the bridge (Equation (14)), which gives the same transition as the two-stage draw above (Appendices E.3 and B).

H.5 Using a sampled reward gradient

The GILC-DB implementation of Dou et al. (2026) uses a straight-through Gumbel–Softmax estimate of the reward derivative. Appendix F.2 derives that estimate once for all three methods, specifies a differentiable extension of the CML reward between one-hot fields, and gives soft and categorical score-function alternatives. To use it here, set $u = 0$, $\phi^{0,t} = \eta^t$, $A^{0,t} = W^{0,t} = P^t$, and $\chi^{0,t} = 1$ on the supports defined above. Call Algorithm 3 once, and replace G^t in Algorithm 7 by the returned $\widehat{\mathcal{G}}^{0,t}$. The corrected probabilities and reverse bridge are then computed exactly as above. There is no repeated fit: the KL derivative is zero at the reference and is not used in this correction. At $u = 0$, differentiating the CML reward requires no derivative through the denoiser. For straight-through, the sampled forward reward is categorical, but its backward gradient is generally biased; the soft variant instead differentiates a temperature-dependent relaxed reward. The exact derivative in Equation (40) remains available for this likelihood and requires neither sampled fields nor a temperature.

I D-MGPS: discrete midpoint guidance

D-MGPS (discrete midpoint guidance posterior sampling), introduced in Section 3, adapts the midpoint decomposition and reverse-KL variational approximation of MGPS (Moufad et al., 2025, Section 3) to the discrete models in Appendix E. At each reverse step, we approximate the distribution of a less corrupted field conditional on the current field and the observations. We draw an intermediate field from this approximation, then sample the next reverse state conditional on both this draw and the current state. We give the construction for uniform and masked corruption, using the CML expectation from Appendix F and the clean-field fit from Appendix G.3. D-MGPS is one of the conditioning methods evaluated in our CML benchmark, with both priors. In its current configuration, it has a lower PCC than G2D2 and GILC, a higher RMSE than G2D2 and than GILC with UDM, and the most negative cumulative rain-rate difference among the discrete samplers (Table 5).

I.1 An exact midpoint identity

Consider a reverse step from t to s, and select a midpoint time u satisfying $0 \le u \le s < t \le T$. Here "midpoint" refers to an auxiliary time $u \le s$, closer to the clean endpoint than the current time t. We use the schedule of Appendix E.1, including $\alpha_T = 0$ and $\alpha_{b|a} = \alpha_b/\alpha_a$ for $a < b$. When the two times coincide, the state does not change and $\alpha_{a|a} = 1$.

The target remains $\pi_0^\theta(z_0 \mid y) \propto p(y \mid z_0)p_0^\theta(z_0)$ from Section 3. Here p_0^θ is the distribution of the final clean field produced by the unguided learned sampler on the chosen time grid. To derive the identity below, consider an auxiliary process: first draw $Z_0 \sim p_0^\theta$, then apply the forward corruption kernels. Its joint distribution is $\mathsf{P}^\theta(z_{0:T}) = p_0^\theta(z_0)\prod_{r=1}^T q_{r|r-1}(z_r \mid z_{r-1})$, and the corresponding posterior is $\pi_{0:T}^\theta(z_{0:T} \mid y) \propto p(y \mid z_0)\mathsf{P}^\theta(z_{0:T})$. The conditional distributions of P^θ are generally intractable. The denoiser predictions and learned reverse transitions need not equal these conditionals. By construction, the auxiliary process and the learned sampler share the same clean marginal p_0^θ. Write $\pi_{u|t}^\theta(\cdot \mid z_t, y)$ for the distribution of Z_u conditional on $Z_t = z_t$ under

$\pi^\theta_{0:T}(\cdot \mid y)$. We write $v = (v^1, \dots, v^L)$ for a possible value of Z_u, the grid state at time u. Each component v^ℓ specifies one class at cell ℓ, or the mask in masked diffusion when $u > 0$. Thus summing over v means summing over complete categorical grid configurations.

Under this forward construction, Z_s is conditionally independent of (Z_0, Y) given (Z_u, Z_t). Consequently, whenever $\pi^\theta_t(z_t \mid y) > 0$,

$$\pi^\theta_{s|t}(z_s \mid z_t, y) = \sum_v q_{s|u,t}(z_s \mid v, z_t)\pi^\theta_{u|t}(v \mid z_t, y). \tag{42}$$

Terms with $\pi^\theta_{u|t}(v \mid z_t, y) = 0$ are omitted. This is the finite-state counterpart of the midpoint identity in Moufad et al. (2025, Lemma 3.1). It follows by conditioning on Z_u and applying the Markov property; it requires neither a Gaussian state space nor differentiability of the likelihood. An exact sampler would first draw $v \sim \pi^\theta_{u|t}(\cdot \mid z_t, y)$, then draw $Z_s \sim q_{s|u,t}(\cdot \mid v, z_t)$. The second distribution does not depend on y once v and z_t are given. Since $q_{T|0}(\cdot \mid z_0) = \bar{q}^{\otimes L}$ for every z_0, the terminal posterior is exactly $\pi^\theta_T(\cdot \mid y) = \bar{q}^{\otimes L}$. We can therefore sample the terminal state and evaluate the bridge exactly; the conditional distribution of Z_u is the part that requires approximation.

Categorical bridge between two noisy endpoints. For one location, let k, j, and h index the states at u, s, and t; for masked diffusion, any of them may be the mask. Bayes' rule gives

$$B_{s|u,t}(j \mid k, h) = \frac{[\alpha_{s|u}\mathbf{1}\{j = k\} + (1 - \alpha_{s|u})\bar{q}_j][\alpha_{t|s}\mathbf{1}\{h = j\} + (1 - \alpha_{t|s})\bar{q}_h]}{\alpha_{t|u}\mathbf{1}\{h = k\} + (1 - \alpha_{t|u})\bar{q}_h}. \tag{43}$$

The denominator is the probability of state h at t given state k at u, and must be positive. Given the two endpoint fields, the bridge samples each location independently with these probabilities. For uniform corruption, k, j, h range over the K clean classes and $\bar{q}_j = 1/K$. The probabilities can be evaluated by normalizing the numerator over candidate classes j. For masked corruption, the midpoint may itself contain masks. An unmasked endpoint at t forces the same class at u and s. If both endpoints are masked, the state at s is masked. If the midpoint is a clean class k and the endpoint at t is masked, the bridge returns k with probability $(\alpha_s - \alpha_t)/(\alpha_u - \alpha_t)$ and the mask otherwise. This last case can occur only when $\alpha_u > \alpha_t$. If $\alpha_u = \alpha_t$, there is no corruption between u and t: the endpoint classes must agree, and the state at s equals that class. The ratio above is then unnecessary and must not be evaluated as $0/0$. For $u = s$, the bridge simply returns the midpoint.

I.2 Constructing the reference distribution and the observation reward

Reference distribution from the current denoiser. Use the clean-class probabilities $P^t_{\ell k} = P^\theta_{\ell k}(z_t, t)$ and allowed classes S^ℓ_t defined in Appendix G.1, including its rule for unmasked cells. The denoiser is evaluated once at the current state (z_t, t) to obtain these probabilities. To see why the reference factorizes, keep the complete current field z_t fixed and write $z^\ell_t = e_{h_\ell}$, $v^\ell = e_{j_\ell}$, and $z^\ell_0 = e_{k_\ell}$. The denoiser approximation assigns probability $p^\theta_{0|t}(z_0 \mid z_t) = \prod_{\ell=1}^L P^t_{\ell k_\ell}$ to this clean field. Independent corruption across cells gives $q_{u|0,t}(v \mid z_0, z_t) = \prod_{\ell=1}^L B_{u,t}(j_\ell \mid k_\ell, h_\ell)$. Here $B_{u,t}$ is the single-cell bridge from Equation (9). Each bridge factor depends only on the clean class at the same cell. Every clean class in S^ℓ_t is compatible with the current class h_ℓ. Substituting the two products into Equation (2) gives

$$\begin{aligned}\rho_{u|t}(v \mid z_t) &= \sum_{k_1 \in S^1_t} \cdots \sum_{k_L \in S^L_t} \prod_{\ell=1}^L \left[P^t_{\ell k_\ell} B_{u,t}(j_\ell \mid k_\ell, h_\ell)\right] \\ &= \prod_{\ell=1}^L \left[\sum_{k \in S^\ell_t} P^t_{\ell k} B_{u,t}(j_\ell \mid k, h_\ell)\right].\end{aligned} \tag{44}$$

The second equality follows because the summed index k_ℓ appears only in the factor for cell ℓ. Thus $\rho_{u|t}(v \mid z_t) = \prod_{\ell=1}^L \rho^\ell_{u|t}(v^\ell \mid z_t)$, with

$$\rho^\ell_{u|t}(j \mid z_t) = \sum_{k=0}^{K-1} P^t_{\ell k}\, B_{u,t}(j \mid k, h_\ell). \tag{45}$$

The superscript ℓ labels a location, and the argument j denotes the one-hot state e_j. Each factor depends on the full z_t through the denoiser prediction; this does not contradict the factorization, because z_t is fixed throughout this calculation. A mixture of product distributions need not factorize in general. Here it does because the clean-class weights themselves form a product and each bridge factor uses only its own clean-class index. The exact clean-field conditional under P^θ generally has dependencies between cells; replacing it by the denoiser's product distribution is an approximation. Accordingly, this calculation establishes factorization of the chosen reference, without asserting factorization of the exact prior conditional or of the clean prior p_0^θ. Each bridge is normalized before averaging over k with weights $P^t_{\ell k}$; classes with zero weight are omitted. For uniform corruption, Equation (15), with u in place of s, computes this average in $O(K)$ operations. In general, inserting the probability vector P^t_ℓ in place of a clean class in the bridge formula does not give the same result. At $u = 0$, the reference equals the current denoiser distribution, so $\rho^\ell_{0|t}(k \mid z_t) = P^t_{\ell k}$. We compute this reference once per reverse step and keep it fixed while optimizing the distribution of Z_u.

For a currently masked location, set $\chi_{u,t} = (\alpha_u - \alpha_t)/(1 - \alpha_t)$. The reference probabilities are

$$\rho^\ell_{u|t}(k \mid z_t) = \chi_{u,t} P^t_{\ell k} \quad (0 \le k < K), \qquad \rho^\ell_{u|t}(\mathtt{M} \mid z_t) = 1 - \chi_{u,t}. \tag{46}$$

At a currently revealed location, $\rho^\ell_{u|t}(\cdot \mid z_t)$ is the point mass at that class. The reference $\rho_{u|t}$ approximates $\mathsf{P}^\theta(Z_u \in \cdot \mid Z_t = z_t)$ by a product over cells. Sampling from this reference and then from the bridge to s recovers the original unguided reverse transition:

$$\sum_v q_{s|u,t}(z_s \mid v, z_t)\rho_{u|t}(v \mid z_t) = \sum_{z_0} q_{s|0,t}(z_s \mid z_0, z_t) p^\theta_{0|t}(z_0 \mid z_t) = p^\theta_{s|t}(z_s \mid z_t). \tag{47}$$

To prove the identity, substitute Equation (45) into the left-hand side and exchange the finite sums. For each fixed z_0, summing over v combines the bridge to u and the bridge from u to s into $q_{s|0,t}$. The same current clean prediction P^t must be used throughout. This equality holds for every $u \le s$ when the midpoint is drawn from the unchanged reference. It does not assert a composition law for learned reverse steps that use separate denoiser predictions at different times.

Using the observations to weight midpoint fields. Bayes' rule factors the exact midpoint posterior as $\pi^\theta_{u|t}(v \mid z_t, y) \propto h^\theta_u(v)\mathsf{P}^\theta(Z_u = v \mid Z_t = z_t)$, where $h^\theta_u(v) = \mathbb{E}_{\mathsf{P}^\theta}[p(y \mid Z_0) \mid Z_u = v]$. The factor $h^\theta_u(v)$ averages the observation likelihood over clean fields conditional on $Z_u = v$. Both this expectation and the prior conditional require approximation. We assume that the likelihood is finite and strictly positive on the clean states under consideration, as holds for the Gaussian CML model with $\sigma_i > 0$. The continuous MGPS construction evaluates the likelihood at a denoised conditional-mean estimate. For D-MGPS, we average the log-likelihood reward r_y, which equals $\log p(y \mid \cdot)$ up to an additive constant, under the product distribution predicted by the denoiser:

$$\mathcal{R}_u(v) = \begin{cases} \displaystyle\sum_{z_0} p^\theta_{0|u}(z_0 \mid v)\, r_y(z_0), & u > 0, \\ r_y(v), & u = 0. \end{cases} \tag{48}$$

When $u > 0$, evaluating this reward requires a new denoiser prediction at the sampled categorical field v. For masked diffusion, the prediction assigns probability one to the class at each unmasked location of v. The prediction at z_t cannot be reused as the prediction at v. At $u = 0$, the field is already clean and no denoiser call is needed.

We use these approximations to define a distribution that favors midpoint fields with larger reward:

$$\nu_{u|t}(v \mid z_t, y) = \frac{\rho_{u|t}(v \mid z_t)\exp\{\lambda_t \mathcal{R}_u(v)\}}{\sum_w \rho_{u|t}(w \mid z_t)\exp\{\lambda_t \mathcal{R}_u(w)\}}. \tag{49}$$

The strength $\lambda_t \ge 0$ controls the weight of the observations; $\lambda_t = 1$ uses the reward without rescaling it. Even then, the weight is proportional to $\exp\{\mathbb{E}[\log p(y \mid Z_0)]\}$, which generally differs from the expected likelihood $\mathbb{E}[p(y \mid Z_0)]$ required by the exact posterior. Jensen's inequality bounds the first quantity by the second only when both expectations use the same distribution. It does not compare our approximate weight with $h^\theta_u(v)$, which uses the generally different exact clean-field conditional. This choice of observation weight differs from the likelihood evaluated at a denoised mean in continuous MGPS.

Using the shared CML expectation. For a sampled categorical field v at $u > 0$, evaluate the denoiser at (v, u) and write $P^u_{\ell k}(v) = P^\theta_{\ell k}(v, u)$, assigning probability one to revealed classes for masked diffusion. The common calculation in Appendix F gives

$$\mathcal{R}_u(v) = \mathcal{R}(P^u(v)). \tag{50}$$

Its link means and variances are $\mu_i(P^u(v))$ and $v_i(P^u(v))$; they must be recomputed for this prediction, rather than reused from P^t. At $u = 0$, no denoiser call is needed: give each class in v probability one and apply the same function, yielding $\mathcal{R}_0(v) = r_y(v)$. For each specified v, this calculation is exact under the chosen clean-field distribution. Optimizing a noisy midpoint distribution still requires averaging these reward values over the fields v drawn from it. Feeding mean token vectors to the denoiser in place of those categorical fields would change the calculation.

I.3 Optimizing categorical probabilities on the allowed states

We approximate $\nu_{u|t}$ by a product distribution $q_\phi(v) = \prod_\ell V^\phi_\ell(v^\ell)$ and minimize $\mathrm{KL}(q_\phi \| \nu_{u|t})$. Here ϕ collects the real class scores, called logits, that we adjust while keeping (z_t, u, t) fixed. For masked diffusion, these scores determine the class probabilities conditional on being unmasked; the mask probabilities remain fixed. As in MGPS, the approximation appears as the first argument of the KL divergence. Here, its factors are categorical distributions at each cell. The following parameterization specifies which classes can receive positive probability for each corruption mechanism.

Uniform corruption. Let $S_\ell = \{k : \rho^\ell_{u|t}(k \mid z_t) > 0\}$ and $A_{\ell k} = \rho^\ell_{u|t}(k \mid z_t)$ on this support. Set $\chi_\ell = 1$, $W_{\ell k} = \mathrm{softmax}_{S_\ell}(\phi_\ell)_k$, and $V^\phi_\ell(k) = W_{\ell k}$. We optimize logits only for classes in S_ℓ, so a class with zero reference probability keeps probability zero. If S_ℓ contains only one class, that class has probability one and no optimization is needed at this cell.

Masked corruption. Currently revealed locations remain fixed: for $z^\ell_t = e_{h_\ell}$, set $S_\ell = \{h_\ell\}$, $A_{\ell h_\ell} = 1$, $\chi_\ell = 1$, and $V^\phi_\ell = \delta_{h_\ell}$. At a currently masked location, set $\chi_\ell = \chi_{u,t}$ from Equation (46), $S_\ell = \{k : P^t_{\ell k} > 0\}$, and $A_{\ell k} = P^t_{\ell k}$. Set $W_\ell = \mathrm{softmax}_{S_\ell}(\phi_\ell)$ and define the variational factor by $V^\phi_\ell(k) = \chi_\ell W_{\ell k}$ for $k \in S_\ell$ and $V^\phi_\ell(\mathtt{M}) = 1 - \chi_\ell$. Thus χ_ℓ is the fixed probability of being unmasked at time u, and $W_{\ell k}$ is the probability of class k conditional on being unmasked. Only W_ℓ is optimized. If $\chi_\ell = 0$, the cell must remain masked and is excluded from optimization.

Fixing the mask probabilities preserves a property of the exact absorbing process. For each cell masked at t, consider the indicator that it was unmasked at u. Conditional on Z_t and on observations that depend only on Z_0, these indicators are independent Bernoulli variables with success probability $\chi_{u,t}$. The forward masking decisions are independent across cells and independent of the clean field and observations; conditioning on being masked at t gives Equation (46). The clean labels are independent of these indicators conditional on (Z_t, y), but may depend on one another across cells. Our product approximation does not represent those dependencies. The approximate target $\nu_{u|t}$ need not have the correct mask probabilities, so we impose them explicitly on q_ϕ. Optimizing the mask probabilities as well would define a different method.

Let F contain the locations with $\chi_\ell > 0$ and $|S_\ell| > 1$, whose class probabilities can be optimized; all other factors are fixed. Apply the KL expansion from Equation (33) with q_ϕ as the fitted distribution, $\rho_{u|t}$ as the reference, and $\mathcal{R}_u$ as the reward. Dropping the log normalizing constant, which is independent of ϕ, gives

$$\mathcal{F}_{u,t}(\phi) = \sum_{\ell \in F} \chi_\ell \sum_{k \in S_\ell} W_{\ell k} \log \frac{W_{\ell k}}{A_{\ell k}} - \lambda_t \mathbb{E}_{Z_u \sim q_\phi}[\mathcal{R}_u(Z_u)]. \tag{51}$$

At fixed cells, the variational and reference distributions agree, so their KL contribution is zero. At a masked cell in F, both distributions assign the same probability to the mask. Its KL contribution is therefore zero, leaving $\chi_\ell \mathrm{KL}(W_\ell \| A_\ell)$ for the clean classes. Initialize $W_\ell = A_\ell$, equivalently $\phi_{\ell k} = \log A_{\ell k}$. Then $q_\phi = \rho_{u|t}(\cdot \mid z_t)$. If $\lambda_t = 0$, this is already a global minimum and no optimization is needed. Equation (47) then shows that sampling from this reference followed by the bridge gives the unguided reverse step.

Choosing the gradient estimator. Appendix F.2 supplies the common construction for this fit, as well as for G2D2 and GILC. Its time-indexed variables $(S^{\ell}_{u,t}, A^{u,t}, \chi^{u,t}, W^{u,t}, \phi^{u,t}, I_{u,t})$ are the $(S_\ell, A, \chi, W, \phi, F)$ defined here at fixed (z_t, u, t). Equation (23) gives the exact KL derivative. For $u > 0$, soft or straight-through Gumbel–Softmax uses the midpoint reward extension in Equation (27) and the gradient in Equation (30). This requires input derivatives through the denoiser at u, with its weights fixed; the reference computed at t remains constant. The mask indicators are sampled independently of the class logits, and their probabilities are not optimized.

For the score-function alternative, omit the time indices and write $\mathsf{S}_{\ell k}(v)$ for the score $\mathsf{S}^{u,t}_{\ell k}(v)$ in Equation (31) and d_ℓ for $d^{u,t}_\ell$ in Equation (23). This gives an unbiased estimate of the categorical gradient, with independent baselines as detailed in Appendix F.2.5. Algorithm 9 retains this explicit REINFORCE version. Algorithm 4 gives the alternative soft or straight-through fits and can also use the same REINFORCE estimator. In every case, a finite stochastic fit need not decrease its objective at every iteration or reach a global minimum.

An alternative is to sum exactly over the class at one cell and sample only the other cells. Let $v^{-\ell}$ denote the sampled states at all cells other than ℓ, and evaluate $\mathcal{R}_u(v^{-\ell}, k)$ for each allowed class k at that cell. For this fixed choice of the other cells, the reward derivative is $\chi_\ell W_{\ell k}[\mathcal{R}_u(v^{-\ell}, k) - \sum_a W_{\ell a}\mathcal{R}_u(v^{-\ell}, a)]$. Averaging over independent draws of $v^{-\ell}$ gives a Rao–Blackwellized estimator, which replaces one random class by its exact conditional average. This generally requires a new denoiser evaluation for each candidate field: changing one input cell may change predictions at other cells.

What the variational bound controls. For fixed z_t, let $\mathcal{K}(v, z_s) = q_{s|u,t}(z_s \mid v, z_t)$. The distribution after drawing a midpoint from q_ϕ and applying this bridge is $(q_\phi\mathcal{K})(z_s) = \sum_v q_\phi(v)\mathcal{K}(v, z_s)$; define $\nu_{u|t}\mathcal{K}$ in the same way. The log-sum inequality, applied separately to each z_s, gives

$$\mathrm{KL}(q_\phi\mathcal{K} \| \nu_{u|t}\mathcal{K}) \le \mathrm{KL}(q_\phi \| \nu_{u|t}). \tag{52}$$

Applying the same bridge to both midpoint distributions cannot increase their KL divergence. This bounds the error relative to sampling from the approximate target $\nu_{u|t}$, as in MGPS. A bound relative to the true posterior transition $\pi^{\theta}_{s|t}$ would also need to control the errors in the reference distribution and observation weights.

After the last optimization update, draw a new midpoint from the fitted q_ϕ using new sampling randomness. Reusing a sample from an earlier optimization iteration would not in general give a draw from the final distribution. If the optimization is stochastic, its final parameters are random, as noted for continuous MGPS in Moufad et al. (2025, Remark 3.3). Conditional on those parameters the midpoint has distribution q_ϕ; conditional only on z_t, its distribution averages over the possible fitted parameters and need not factorize. Equation (52) holds for each fixed fitted parameter. By convexity of KL in its first argument, the divergence of the averaged transition is bounded by the expected fitted KL on the right-hand side.

I.4 Using the shared fit over clean fields

At $u = 0$, the midpoint field is clean, $\rho^{\ell}_{0|t}(k \mid z_t) = P^t_{\ell k}$, and $\mathcal{R}_0(v) = r_y(v)$. The reveal probability is one at every cell, so the variational clean-class probabilities are $V = W$. Substituting these identities into Equation (51) gives exactly the CML objective $\mathcal{F}_{0,t}$ in Equation (34). The matrix called V here is called V^t there; both contain the clean-class probabilities fitted while the current state (z_t, t) is held fixed. Appendix G.3 derives its score gradient, its exact minimizer over one cell, and the fact that sequential updates cannot increase the objective. For the analytic fit, use the shared Algorithm 5 with the current reference P^t, strength λ_t, and specified number of sweeps. The sampled alternatives use Algorithm 4 at $u = 0$ with the same reference.

After fitting, draw a clean field v from its final probabilities and apply the Markov bridge $q_{s|0,t}(\cdot \mid v, z_t)$. Appendix G.5 explains when this procedure agrees with Markov G2D2 and how the shared objective relates to a single GILC correction. For $u > 0$, D-MGPS instead fits the noisy midpoint distribution described above.

Algorithm 8 D-MGPS for CML reconstruction

Input: Pretrained denoiser with fixed parameters θ, corruption mode, codebook $\mathcal{C}$, schedule α
Input: CML data y, coefficients (κ, b, Δ), noise scales $\sigma_i > 0$
Input: Reverse grid $T = \tau_N > \cdots > \tau_0 = 0$ and midpoint times $0 \le u_n \le \tau_{n-1}$
Input: Strengths $\lambda_n \ge 0$, iteration counts $M_n \ge 0$, batch size $B \ge 1$, step sizes $\varepsilon_{n,j} > 0$
Input: Sampled estimator (soft, straight-through, or REINFORCE); temperatures $\vartheta_{n,j} > 0$ when needed
Input: Clean fit ($u = 0$): sampled or coordinate; cell order for the coordinate choice
Output: One approximate posterior rain field $\widetilde{x}$

Precompute sparse $c_{i\ell k} = \kappa_i \Delta_\ell^i a_k^{b_i}$
Draw $z_T^\ell \sim \mathrm{Cat}(\bar{q})$ independently at all locations
for $n = N, \ldots, 1$ **do**
  $(t, s, u) \leftarrow (\tau_n, \tau_{n-1}, u_n)$
  $P^t \leftarrow \mathrm{CleanPred}_\theta(z_t, t)$ ▷ Keep unmasked classes fixed
  Form the fixed factors $\rho_{u|t}^\ell(\cdot \mid z_t)$ using Equation (45)
  Construct allowed classes S_ℓ, reference probabilities A_ℓ, reveal probabilities χ_ℓ, and optimized cells F
  Use the factors $V(W)$ from Appendix I.3; initialize $W_\ell \leftarrow A_\ell$
  if $\lambda_n > 0$, $M_n > 0$, and $F \neq \varnothing$ **then**
    if $u = 0$ and the clean fit uses coordinate updates **then**
      $W \leftarrow$ Algorithm 5 with (P^t, λ_n, M_n), initial matrix W, and the given cell order
    else if $u > 0$ and the sampled estimator is REINFORCE **then**
      $W \leftarrow$ fitted factors from Algorithm 9, using $(M_n, B, \varepsilon_{n,\cdot}, \lambda_n)$
    else
      Fit V by Algorithm 4 with the chosen estimator, steps, and temperatures
      Recover $W_{\ell k} \leftarrow V_{\ell k}/\chi_\ell$ where $\chi_\ell > 0$; retain other rows
    end if
  end if
  Compute all midpoint probabilities V_ℓ from the final W and unchanged χ_ℓ
  Draw $v^\ell \sim \mathrm{Cat}(V_\ell)$ independently for all ℓ, using new sampling randomness
  if $u = s$ **then**
    $z_s \leftarrow v$
  else
    for each location ℓ, with $v^\ell = e_k$ and $z_t^\ell = e_h$ (either may be the mask) **do**
      Draw z_s^ℓ from $B_{s|u,t}(\cdot \mid k, h)$ in Equation (43)
    end for
  end if
end for
return $\widetilde{x} = \mathcal{D}(z_0)$

I.5 Complete sampling algorithms

Algorithm 8 repeats the guided reverse steps for the CML likelihood. At $u > 0$, it optimizes the midpoint probabilities using the chosen soft, straight-through, or REINFORCE estimator. At $u = 0$, the fit can use the exact coordinate minimizer or the common sampled estimator. Algorithm 9 specifies the REINFORCE fit at $u > 0$; Algorithm 4 gives all three sampled variants, and Algorithm 5 supplies the analytic clean-field fit. Soft Gumbel–Softmax passes fractional class vectors to the differentiable denoiser extension; straight-through and REINFORCE evaluate it at categorical fields. All three estimators use fresh categorical draws for the reverse sampling transition. The function $\mathrm{CleanPred}_\theta$ returns clean-class probabilities and, for masked diffusion, assigns probability one to the class at every unmasked input cell. The number of iterations, step sizes, midpoint times, and guidance strengths must be specified for each experiment. When Algorithm 4 is called at reverse step n, its inputs are $(u, t) = (u_n, \tau_n)$, $\lambda_t = \lambda_n$, $M_t = M_n$, $B_t = B$, $\varepsilon_j^{u,t} = \varepsilon_{n,j}$, and $\vartheta_j^{u,t} = \vartheta_{n,j}$. Here τ_n is a diffusion time, whereas $\vartheta_{n,j}$ is a relaxation temperature.

Algorithm 9 REINFORCE fit of the categorical midpoint distribution when $u > 0$

Input: Fixed $(z_t, u, t, A, \chi, S, F)$ and initial W from Algorithm 8
Input: Denoiser with fixed parameters, CML data, M iterations, batch size B, strength λ, steps ε_m
Output: Optimized class probabilities W, defining all midpoint factors $V(W)$

$\phi_{\ell k} \leftarrow \log W_{\ell k}$ for $\ell \in F$ and $k \in S_\ell$
for $m = 1, \dots, M$ **do**
  $W_\ell \leftarrow \mathrm{softmax}_{S_\ell}(\phi_\ell)$ for $\ell \in F$; keep other cells fixed
  Compute all V_ℓ from W_ℓ and unchanged mask probabilities
  for $b = 1, \dots, B$ **do**
    Draw a categorical field $v_b \sim \prod_\ell \mathrm{Cat}(V_\ell)$
    $P^{u,b} \leftarrow \mathrm{CleanPred}_\theta(v_b, u)$ ▷ New prediction at this field
    Evaluate $R_b \leftarrow \mathcal{R}_u(v_b)$ by the CML moments in Equation (50)
  end for
  $\beta_b \leftarrow (\sum_{a \neq b} R_a)/(B-1)$ if $B > 1$; otherwise $\beta_b \leftarrow 0$ ▷ Leave-one-out baselines
  for each optimized cell $\ell \in F$ **do**
    $d_\ell \leftarrow \sum_{a \in S_\ell} W_{\ell a} \log(W_{\ell a}/A_{\ell a})$
    for $k \in S_\ell$ **do**
      Evaluate $\mathsf{S}_{\ell k}(v_b)$ for all b using Equation (31)
      $\widehat{g}_{\ell k} \leftarrow \chi_\ell W_{\ell k}[\log(W_{\ell k}/A_{\ell k}) - d_\ell]$
      $\widehat{g}_{\ell k} \leftarrow \widehat{g}_{\ell k} - (\lambda/B) \sum_b (R_b - \beta_b)\mathsf{S}_{\ell k}(v_b)$
    end for
  end for
  $\phi_{\ell k} \leftarrow \phi_{\ell k} - \varepsilon_m \widehat{g}_{\ell k}$ simultaneously for $\ell \in F$, $k \in S_\ell$
  Subtract the mean from each optimized logit vector ▷ Probabilities are unchanged
end for
return $W_\ell = \mathrm{softmax}_{S_\ell}(\phi_\ell)$ for $\ell \in F$; keep other cells fixed

Final step and mask probabilities. At the final reverse step, $s = 0$ forces $u = 0$. The selected clean-field fit uses the likelihood directly, and the final clean field is drawn from the fitted probabilities. Because $u = s = 0$, no further bridge draw is needed. Thus our final output is a categorical draw, whereas the original continuous MGPS algorithm uses deterministic final denoising. For masked diffusion, no revealed token is ever changed or remasked. At a currently masked location, the probability of becoming revealed after the two-stage draw is $\chi_{u,t}(\alpha_s - \alpha_t)/(\alpha_u - \alpha_t) = (\alpha_s - \alpha_t)/(1 - \alpha_t)$ when $\alpha_u > \alpha_t$. If $\alpha_u = \alpha_t$, no reveal occurs and the ratio is not evaluated. The identity also covers $u = s$: if the sampled midpoint is unmasked, the bridge keeps that class with probability one. Thus optimization changes which class is revealed, while preserving the unguided probability of revealing a cell at each reverse step. At $s = 0$, every remaining mask is removed.

I.6 Midpoint choice, initialization, and practical limits

Choosing the midpoint. One choice is $u = \lfloor \gamma s \rfloor$ with $0 \leq \gamma \leq 1$. At $u = 0$, we optimize a distribution over clean fields; at $u = s$, we optimize the distribution at the next reverse time directly. When α is nonlinear in time, one can instead choose an intermediate corruption level: select the integer $u \in \{0, \dots, s\}$ whose α_u is closest to $1 - \gamma(1 - \alpha_s)$, with a specified tie rule. With either rule, choosing $u > 0$ requires denoiser predictions at sampled midpoint fields. Choosing $u = 0$ avoids those calls and permits exact likelihood evaluation at each midpoint field, but the product distribution must then approximate a conditional over fully clean fields. Errors in the denoiser and dependencies between cells can make that approximation difficult when z_t is strongly corrupted. There is no established guarantee that decreasing u improves accuracy, or that a midpoint choice effective for continuous MGPS transfers to categorical models.

Initialization and optional reuse across steps. By default, each reverse step initializes $W = A$, so the initial midpoint distribution equals the reference computed at the current z_t. No predictions from earlier reverse steps need to be retained. To reuse earlier predictions as in MGPS, one may store the clean-class probabilities $\overline{P}$ obtained by denoising the previous sampled midpoint. If that

midpoint was at time zero, use its one-hot classes instead. For masked diffusion, first replace $\overline{P}$ at every currently unmasked cell by probability one on its current class. Then apply Equation (45) at the new (u, t) with weights $\overline{P}$ to obtain starting midpoint probabilities. The reference $\rho_{u|t}(\cdot \mid z_t)$ in the KL penalty must still be computed using the current prediction P^t.

At each cell in F, restrict the starting class probabilities to S_ℓ and renormalize; if no mass remains, use A_ℓ. Mix the result with a specified positive fraction of A_ℓ to give positive probability, and hence finite logits, to every allowed class. For masked cells, apply these operations to the clean-class probabilities conditional on being unmasked, and retain the mask probability $1 - \chi_\ell$. Cells outside F remain fixed. Only the starting probabilities change; the reference and the objective remain the same. For $u = 0$, pass these starting probabilities as $V^{t,\mathrm{init}}$ to Algorithm 5; its default remains P^t. Because these starting probabilities use stored predictions, statements about the sampling law must also condition on that stored information, in addition to z_t. The additional two-stage WarmStart procedure in Moufad et al. (2025, Appendix C.2) uses a continuous prior score and is not included in this categorical construction.

Cost and numerical evaluation. When $u > 0$ and optimization is performed, each reverse step evaluates the denoiser once at z_t and $M_n B$ times at sampled midpoint fields; these calls may be batched. If optimization is skipped, only the prediction at z_t is needed. No backpropagation through the denoiser is required for the score-function estimator. Soft and straight-through Gumbel–Softmax additionally backpropagate each midpoint reward to the input class scores, while keeping the network parameters fixed. Each CML reward evaluation costs $O(K \operatorname{nnz}(\Delta))$ arithmetic operations, in addition to the network evaluation and the gridwise categorical operations. At $u = 0$, the denoiser is evaluated only at z_t; Appendix G.3 gives the cost of the analytic coordinate fit, while a sampled gradient costs $O(BK \operatorname{nnz}(\Delta))$ per iteration in addition to the gridwise operations. For the noisy-midpoint fit, apply the same rule of evaluating softmax only over allowed classes, keeping probability zero elsewhere. Clipping rewards or gradient estimates changes the estimator or optimization procedure and must be reported. Small observation variances can concentrate class probabilities and increase the variance of the gradient estimate, so the batch size, step sizes, and iteration counts must be chosen empirically.

Scope of the approximation. The midpoint identity and categorical bridge are exact. For the CML model, the expected log-likelihood under any specified product distribution over clean fields and the minimizer over one cell's probabilities at $u = 0$ are also exact. The full sampler remains approximate because it uses a denoiser-based reference, approximate observation weights when $u > 0$, a product family, and a finite number of optimization steps. The soft and straight-through variants introduce the additional gradient approximations described in Appendix F.2.4. Solving the variational problem exactly would not remove the other approximation errors. Varying λ_t across reverse steps does not generally correspond to raising the likelihood to one fixed power in the final posterior. If the likelihood vanishes on some fields, the variational distributions must respect those constraints, or a noisy observation model must be specified instead. Replacing $\log 0$ by an arbitrary finite number changes the model. Finite-state checks can verify bridge normalization, the midpoint identity, recovery of the unguided step, the constrained logit gradients, and the monotonicity of sequential coordinate updates. In the CML benchmark, D-MGPS currently has a lower PCC and a more negative cumulative rain-rate difference than G2D2 and GILC (Table 5). We have not determined which of the approximations above limits its accuracy; its runtime is reported in Table 1, but its interval coverage is not.

J Extended related work

J.1 Rainfall retrieval and reconstruction from microwave links

Commercial microwave links provide precipitation information by measuring attenuation along communication paths. Early studies established the potential of operational networks for environmental monitoring and rain-rate retrieval (Messer et al., 2006; Leijnse et al., 2007). Countrywide reconstructions subsequently demonstrated spatial rainfall mapping from large link networks (Overeem et al., 2013; 2016). The review by Chwala & Kunstmann (2019) covers these de-

velopments, as well as measurement processing, data access, and uncertainties. CMLs complement rain gauges and weather radar; the three sources differ in spatial coverage and error characteristics.

Before reconstruction, the reference signal level, the finite power resolution, the temporal sampling, and wet-antenna effects all affect the relation between recorded attenuation and rainfall (Leijnse et al., 2007). Schleiss et al. (2013) quantify and model the temporal evolution of wet-antenna attenuation. Our Gaussian observation model describes the residual uncertainty after preprocessing, not every physical error mechanism. OpenMRG provides colocated link, radar, and gauge observations for evaluation (Andersson et al., 2022), and shared data conventions support reproducible processing across opportunistic sensors (Fencl et al., 2024).

Spatial retrieval methods make different assumptions about rainfall within and between observed paths. Virtual-gauge methods assign representative point intensities to links before interpolation (Eshel et al., 2021; Zhang et al., 2023). Iterative link-based reconstruction instead adjusts local intensities to match aggregate measurements (Goldshtein et al., 2009). Neural methods learn rain-rate retrieval from attenuation time series (Habi & Messer, 2021). These procedures differ from conditional field generation, which targets a distribution over spatial configurations compatible with the measurements.

Uncertainty-aware reconstruction predates diffusion models. Transformed or censored latent Gaussian models represent rainfall intermittency and spatial dependence (Bárdossy & Plate, 1992; Ailliot et al., 2009). Random Mixing constructs conditional precipitation fields from gauge observations and link-averaged rain-rate constraints (Haese et al., 2017; Blettner et al., 2022). Recent approaches combine a pretrained continuous diffusion prior with a CML observation model (Moufad et al., 2026), or learn the observation conditioning when training the diffusion model (Ji et al., 2026). Our work differs primarily in the prior: a discrete diffusion prior on a categorical encoding with a dedicated dry class and a codebook learned on positive intensities, combined with conditioning at inference time. The dry class preserves the wet–dry pattern exactly under encoding, while positive intensities are quantized separately. This property alone does not guarantee calibrated reconstruction uncertainty, which we assess empirically through interval coverage (Table 2).

J.2 Diffusion and iterative generation on discrete spaces

Discrete diffusion defines corruption and generation on categorical states. Multinomial diffusion progressively corrupts categorical data toward a uniform distribution (Hoogeboom et al., 2021). D3PM (Austin et al., 2021) permits more general transition matrices, including absorbing corruption. These constructions underlie the uniform and masked processes considered here. Uniform noise can replace a token with another clean class, whereas absorbing corruption replaces it with a mask. The reverse processes therefore differ in their support and in their ability to revise sampled classes: uniform diffusion can change a class at later reverse steps, whereas masked diffusion keeps a revealed class fixed.

The time parameter need not be discrete. Campbell et al. (2022) formulate corruption and generation as continuous-time Markov chains and derive a continuous-time variational training objective. The states remain categorical; transition rates describe jumps between configurations. SEDD learns ratios of state probabilities through a score-entropy objective (Lou et al., 2024). These ratios determine the reverse jump rates and act as discrete scores, not as derivatives with respect to token indices. Zhao et al. (2025) relate discrete-time and continuous-time formulations through shared forward processes and backward parameterizations, with simplified variational objectives. Closed-form transitions simplify inference, but generation still depends on the learned model and the sampling scheme.

Masked models connect this framework with parallel token prediction. MaskGIT reconstructs masked image tokens with a bidirectional transformer and generates images through iterative refinement (Chang et al., 2022). Its confidence-based token selection differs from the transitions of an absorbing Markov process. Recent masked-diffusion formulations provide simpler denoising objectives with explicit probabilistic interpretations (Sahoo et al., 2024; Shi et al., 2024). In our setting, each clean token encodes a quantized rain intensity at a fixed grid location, and decoding is a codebook lookup. The denoiser uses the complete corrupted field, so its factorized predictions can depend on spatial context. However, guidance methods that fit or correct these factorized predictions apply the observation at each reverse step through a product distribution, which cannot represent the

dependence that a path-integrated observation induces between the cells of a link. A suitable prior representation therefore does not by itself solve the conditioning problem.

J.3 Observation guidance for discrete diffusion

For continuous-valued diffusion, posterior-sampling methods include approximate likelihood gradients (Chung et al., 2023), variational midpoint inference (Moufad et al., 2025), and sequential Monte Carlo (Wu et al., 2023). Their assumptions and guarantees differ; an asymptotic particle guarantee does not imply exact sampling with a finite budget. Adapting these ideas to categorical states requires specifying how the observations modify transition probabilities or rates.

Guidance for discrete diffusion either relies on trained conditional models or predictors, or applies a reward at inference time. Nisonoff et al. (2025) derive conditional jump rates from ratios of predictive probabilities at corrupted states. They also formulate predictor-free guidance from conditional and unconditional rates, and approximate predictor guidance to reduce its computational cost. Schiff et al. (2025) develop classifier-based and classifier-free guidance and study uniform corruption, which permits later changes to generated tokens. These approaches require a trained predictor or conditional model for each new measurement model.

Discrete Guidance Matching expresses guided clean-token probabilities through conditional expectations of a target-to-source density ratio (Wan et al., 2026). The conditioning identity is exact under its assumptions, but its implementation learns these expectations with a guidance network. This differs from conditioning a fixed rain-field prior on new measurements without additional training. He et al. (2026) analyze classifier-free guidance in low-dimensional masked models, deriving its effects on output probabilities, covariance, and sampling dynamics. Their results concern classifier-free guidance rather than CML inference, but they motivate examining the complete generated distribution.

For conditioning a fixed prior, G2D2 iteratively fits a categorical clean-state approximation, using continuous relaxations for optimization (Murata et al., 2025). Its original sampler uses star-shaped corruption; we also consider a variant based on Markov bridges. Appendix G derives both versions for our CML likelihood, including the analytic probability updates (Appendix G.3); Appendix F.2 gives their Gumbel–Softmax, straight-through, and score-function fits. GILC instead corrects the denoiser logits using gradients of an expected reward (Dou et al., 2026). For the Gaussian CML likelihood, this expectation and its logit gradient can be computed exactly under a categorical product distribution; Appendices F and H give the calculation and its GILC specialization, with the sampler in Algorithm 7. Exact evaluation of the local reward does not make posterior sampling exact.

We do not propose a new general guidance principle for discrete diffusion. Instead, we adapt G2D2 and GILC to the CML likelihood, introduce D-MGPS as a discrete variant of MGPS, and compare these samplers with masked in-filling on the same discrete priors. D-MGPS places the variational fit at an adjustable intermediate corruption time and passes a new categorical draw through a diffusion bridge. It retains approximation errors in the reference distribution, the likelihood surrogate, the product family, and the finite optimization. In our benchmark, it is currently less accurate than G2D2 and GILC, except that its RMSE and wet/dry accuracy are better than those of GILC with MDM (Tables 1 and 5).