Probabilistic Invertibility of Rectified Flows Beyond Global Monotonicity

Elias Firisa

Abstract

Rectified Flow (RF) has recently emerged as an efficient alternative to diffusion-based generative modeling, with a theoretical explanation for fast sampling that hinges on invertibility properties of a rectified transport map. Existing analyses establish global invertibility of an auxiliary map $H_t(\cdot)$ under a strong uniform monotonicity condition on the endpoint Jacobian, yielding a sufficient route to straight couplings and quantitative error control. Empirical evidence, however, indicates that this monotonicity condition is frequently violated while invertibility persists for structured targets such as Gaussian mixtures.

This preprint develops a complementary "generic-in-time" theory for the Gaussian-to-Gaussian-mixture regime. The central technical observation is a transversality identity: at any collision point $H_t(x) = H_t(y)$ with $x \neq y$ and $t \in (0, 1)$, the time derivative $\partial_t(H_t(x) - H_t(y))$ is nonzero and explicitly proportional to $-(x - y)/t$. This *transversality engine* implies that collision sets form a Lebesgue-null subset of $(0, 1) \times B_R \times B_R$ on any compact domain B_R, and consequently, for Lebesgue-almost every $t \in (0, 1)$, collision pairs in $B_R \times B_R$ have zero Lebesgue measure. Since Gaussian initialization is absolutely continuous, it follows that for almost every time t, the map H_t is injective almost surely with respect to $(Z_0, Z_0') \sim \rho_0 \times \rho_0$ on compacts, and hence in the limit over radii.

The argument couples (i) strip-wise flow regularity on $t \in [\varepsilon, 1 - \varepsilon]$ (supported by explicit mixture-score formulae) with (ii) the transversality identity and (iii) an implicit-function/Fubini slicing argument. This replaces uniform global monotonicity by an almost-sure injectivity statement aligned with the probabilistic manner in which RF is used in practice. Placeholders for numerical verification on synthetic mixtures and image-scale datasets are provided.

1 Introduction

Rectified Flow (RF) constructs a deterministic probability flow ODE that transports a simple base distribution ρ_0 (often a standard Gaussian) to a target ρ_1. In contrast to diffusion models, RF aims to learn (or approximate) velocity fields that generate near-straight trajectories between noise and data, enabling fast generation.

A recurring analytical object in RF theory is an auxiliary map $H_t(\cdot)$ that interpolates between the identity and the time-1 endpoint of the RF flow. In recent work [1], global invertibility of $H_t(\cdot)$ is obtained under a sufficient assumption (Assumption 4.6 therein) imposing uniform global monotonicity through the symmetric part of the endpoint Jacobian. That route supports strong structural consequences (e.g., straight couplings under 1-RF) and quantitative error bounds. At the same time, simulations reported in [1] suggest that Assumption 4.6 is not necessary, and an explicit open problem is raised for Gaussian mixtures.

Objective. This note develops a complementary analysis for the regime

$$\rho_0 = \mathcal{N}(0, I_d), \qquad \rho_1 = \sum_{k=1}^{K} \pi_k \, \mathcal{N}(\mu_k, \Sigma_k),$$

in which the RF drift admits a closed-form score representation. The principal aim is to formalize a probabilistic notion of invertibility: *for Lebesgue-almost every* $t \in (0,1)$, *the collision event* $H_t(Z_0) = H_t(Z_0')$ *with* $Z_0 \neq Z_0'$ *occurs with probability zero under* $(Z_0, Z_0') \sim \rho_0 \times \rho_0$.

Contributions. The contributions are as follows.

1. **Transversality engine.** A structural identity is established: at any collision point $H_t(x) = H_t(y)$ with $x \neq y$ and $t \in (0,1)$, the derivative $\partial_t(H_t(x) - H_t(y))$ is nonzero and equals $-(x-y)/t$. This identity is independent of monotonicity or convexity assumptions and isolates time-mixing as the mechanism that destroys persistent collisions.

2. **Generic injectivity on compacts.** Under standard flow regularity on strips $t \in [\varepsilon, 1-\varepsilon]$, the transversality engine implies that collision sets are Lebesgue-null in $(0,1) \times B_R \times B_R$ for every radius R. Consequently, for Lebesgue-almost every t, the set of collision pairs in $B_R \times B_R$ has $2d$-dimensional Lebesgue measure zero.

3. **Probabilistic invertibility under Gaussian initialization.** Since ρ_0 has a smooth density, the preceding Lebesgue-null collision property yields that, for almost every t, the collision event under $(Z_0, Z_0') \sim \rho_0 \times \rho_0$ has probability zero on each B_R and hence in the limit as $R \to \infty$.

4. **Working preprint structure.** A regularity pathway is documented using (i) existence/uniqueness and non-explosion via an Osgood criterion and (ii) analytic mixture-score formulae on time strips. A dedicated experiments section is included as a placeholder with expected outcomes aligned with theory.

Positioning. The results complement sufficient global monotonicity conditions by replacing uniform (worst-case) invertibility with a generic-in-time, almost-sure injectivity guarantee that matches the probabilistic use of RF (random initialization and continuous time sampling).

2 Mathematical setup

2.1 Independent coupling, interpolation, and RF drift

Let $(X_0, X_1) \sim \rho_0 \times \rho_1$ with $\rho_0 = \mathcal{N}(0, I_d)$. Define the linear interpolation

$$X_t := (1-t)X_0 + tX_1, \qquad t \in [0,1]. \tag{1}$$

Denote by p_t the density of X_t and by $s_t(x) := \nabla_x \log p_t(x)$ its score.

The RF velocity field (as in [1]) is

$$v_t(x) := \mathbb{E}[X_1 - X_0 \mid X_t = x], \qquad t \in (0,1). \tag{2}$$

The RF sampling ODE is

$$\frac{d}{dt} Z_t = v_t(Z_t), \qquad Z_0 = z_0. \tag{3}$$

The endpoint map is $Z_1 : \mathbb{R}^d \to \mathbb{R}^d$, $Z_1(z_0) := Z_{t=1}(z_0)$.

2.2 The rectified map H_t

Following [1], define for each $t \in [0,1]$ the auxiliary map

$$H_t(z_0) := (1-t)z_0 + t\, Z_1(z_0). \tag{4}$$

Invertibility of $H_t(\cdot)$ plays a central role in straightness arguments in [1].

2.3 Gaussian-mixture targets

Let the target be a (nondegenerate) Gaussian mixture

$$\rho_1(x) = \sum_{k=1}^{K} \pi_k \, \varphi(x; \mu_k, \Sigma_k), \qquad \pi_k > 0, \ \sum_{k=1}^{K} \pi_k = 1, \ \Sigma_k \succ 0. \tag{5}$$

In this regime, [1] derives a Tweedie-type identity (Appendix A.4) expressing v_t in terms of the score:

$$v_t(x) = \frac{x}{t} + \frac{1-t}{t} s_t(x), \qquad t \in (0,1), \tag{6}$$

and provides an explicit responsibility-weighted affine form for s_t (general covariance) as well as an explicit expression for $\nabla_x v_t$ on time strips. These explicit formulae imply real-analytic dependence on x for each fixed $t \in (0,1)$.

2.4 Existence, uniqueness, and non-explosion

The vector field in (6) exhibits an apparent singularity at $t = 0$ (and related behavior can occur near $t = 1$ depending on parametrization). The RF analysis in [1] establishes existence and uniqueness of solutions on $[0,1]$ under an Osgood non-explosion criterion (Assumption 4.4) and a mild moment condition (Theorem 4.5). In the Gaussian mixture regime, the Osgood condition is verified explicitly in Appendix A.3.2 for general mixtures, yielding global well-posedness of (3). The present work takes this well-posedness as the foundation for the endpoint map $Z_1(\cdot)$.

3 A general transversality theorem for rectified maps

Recall $H_t(z) = (1-t)z + tZ_1(z)$, where Z_1 is the time-1 endpoint map of the RF ODE. For $t \in (0,1)$ and $(x,y) \in \mathbb{R}^d \times \mathbb{R}^d$, define

$$F(t,x,y) := H_t(x) - H_t(y).$$

A collision at time t is a pair (x,y) with $x \neq y$ such that $F(t,x,y) = 0$.

Definition 1 (Collision sets on compacts)**.** For $R > 0$, let $B_R := \{x \in \mathbb{R}^d : \|x\| \leq R\}$ and $\Delta := \{(x,x) : x \in B_R\}$. Define the collision slice at time t by

$$\mathcal{C}_{t,R} := \{(x,y) \in (B_R \times B_R) \setminus \Delta : \ H_t(x) = H_t(y)\}.$$

Assumption 1 (Strip-wise regularity)**.** For every $\varepsilon \in (0, 1/2)$ and $R > 0$, the mapping $(t,x) \mapsto H_t(x)$ is C^1 on $[\varepsilon, 1-\varepsilon] \times B_R$.

Lemma 1 (Transversality engine)**.** *Let $t \in (0,1)$ and $x \neq y$. If $H_t(x) = H_t(y)$, then*

$$\partial_t\big(H_t(x) - H_t(y)\big) = -\frac{1}{t}(x-y) \neq 0.$$

Proof. Since $H_t(z) = (1-t)z + tZ_1(z)$, one has

$$H_t(x) - H_t(y) = (1-t)(x-y) + t(Z_1(x) - Z_1(y)).$$

At a collision, $(1-t)(x-y) + t(Z_1(x) - Z_1(y)) = 0$, hence $Z_1(x) - Z_1(y) = -(1-t)(x-y)/t$. Differentiating in t and substituting yields $-(x-y) - (1-t)(x-y)/t = -(x-y)/t \neq 0$. □

Theorem 1 (Generic injectivity on compact sets)**.** *Under Assumption 1, for every $R > 0$ and for Lebesgue-almost every $t \in (0, 1)$,*

$$\mathrm{Leb}_{2d}\big(\mathcal{C}_{t,R}\big) = 0.$$

Consequently, if ρ_0 admits a density with respect to Lebesgue measure, then for Lebesgue-almost every t,

$$(\rho_0 \times \rho_0)\big(\mathcal{C}_{t,R}\big) = 0.$$

Proof. Fix $\varepsilon \in (0, 1/2)$ and consider the domain $D_{\varepsilon,R} := [\varepsilon, 1-\varepsilon] \times ((B_R \times B_R) \setminus \Delta)$. By Assumption 1, F is C^1 on $D_{\varepsilon,R}$. Lemma 1 implies that on the zero set $\{F = 0\} \subset D_{\varepsilon,R}$, the partial derivative $\partial_t F$ is nonzero. The implicit function theorem therefore implies that near each point of $\{F = 0\}$, the solution set can be represented as a C^1 graph $t = \tau(x, y)$ over (x, y). Hence the full collision set in (t, x, y) has $(2d + 1)$-dimensional Lebesgue measure zero on $D_{\varepsilon,R}$. Fubini's theorem yields $\mathrm{Leb}_{2d}(\mathcal{C}_{t,R}) = 0$ for a.e. $t \in [\varepsilon, 1 - \varepsilon]$. Letting $\varepsilon \downarrow 0$ along a sequence and intersecting full-measure sets completes the proof. Absolute continuity of $\rho_0 \times \rho_0$ implies the probabilistic statement. □

Remark 1 (Beyond specific targets). Theorem 1 separates the transversality mechanism (Lemma 1) from distribution-specific verification of Assumption 1. In particular, for $\rho_0 = \mathcal{N}(0, I)$ and general ρ_1, the intermediate density p_t is a Gaussian convolution of a scaled version of ρ_1, suggesting strip-wise smoothness of score-based drifts under mild tail control. Verifying Assumption 1 for broader classes of targets is left as a direction for future work.

4 Verification for Gaussian mixture targets

Corollary 1 (Gaussian mixtures satisfy strip-wise regularity)**.** *Let $\rho_0 = \mathcal{N}(0, I_d)$ and let ρ_1 be a nondegenerate Gaussian mixture. Then Assumption 1 holds. Consequently, Theorem 1 implies that for every $R > 0$, $\mathrm{Leb}_{2d}(\mathcal{C}_{t,R}) = 0$ for Lebesgue-almost every $t \in (0, 1)$, and $(\rho_0 \times \rho_0)(\mathcal{C}_{t,R}) = 0$ for almost every t.*

Proof sketch. In the Gaussian-to-GMM regime, the RF drift admits an explicit score representation and its spatial derivatives are explicit and smooth on every time strip $[\varepsilon, 1 - \varepsilon]$ (Appendix A.4 in [1]). Moreover, global existence and uniqueness of the RF ODE are ensured by an Osgood non-explosion criterion (Assumption 4.4 and Theorem 4.5 in [1]), which is verified for general mixtures (Appendix A.3.2). Standard ODE flow regularity on compact time intervals then yields the C^1 property required by Assumption 1. □

5 Relation to monotonicity-based invertibility results

Assumption 4.6 in [1] imposes a uniform condition on the symmetric part of the endpoint Jacobian that yields local invertibility and properness of H_t, and thus global invertibility via classical global inverse function theorems. The present results provide an alternative explanation for the empirical robustness of invertibility in Gaussian-mixture settings: even when monotonicity fails, collisions are structurally transversal in time, which enforces a generic-in-time injectivity property on compacts and an almost-sure collision-free statement under Gaussian initialization.

6 Experiments

This section outlines numerical experiments designed to validate and illustrate the theoretical predictions. Results and plots are placeholders to be filled in.

6.1 Synthetic mixtures: collision frequency vs. time

Setup. Let $\rho_0 = \mathcal{N}(0, I_d)$ and ρ_1 be a K-component GMM with random means and covariances. Numerically approximate $Z_1(\cdot)$ by integrating the RF ODE using the analytic drift in (6). Sample pairs $(x, y) \sim \rho_0 \times \rho_0$ and evaluate whether $\|H_t(x) - H_t(y)\|$ is below a tolerance. Repeat over a grid of $t \in (0, 1)$ and multiple random seeds.

Expected outcome. For each fixed t, near-collisions should be extremely rare; empirically, the set of t exhibiting any detected collisions should shrink with improved numerical accuracy and should appear as isolated values rather than intervals, consistent with transversality.

6.2 Jacobian singularity diagnostics

Setup. Estimate $\sigma_{\min}(\nabla H_t(x))$ or $\det(\nabla H_t(x))$ for sampled $x \sim \rho_0$ and $t \sim \mathrm{Unif}(0, 1)$. Compare instances where Assumption 4.6 fails (e.g., negative eigenvalues of the symmetric part of $\nabla Z_1(x)$) to the observed singularity statistics.

Expected outcome. Violations of monotonicity should be observable, while actual near-singular ∇H_t events should remain rare, supporting the distinction between sufficient monotonicity and generic injectivity.

6.3 Image-scale mixtures (placeholder)

Setup. Construct a low-dimensional embedding of an image dataset and fit a GMM in the embedding space, or use a synthetic high-dimensional GMM. Repeat the collision diagnostics and sampling behavior.

Expected outcome. The transversality trend should persist qualitatively; collision probability should remain negligible for random draws.

7 Conclusion

This work develops a generic-in-time invertibility theory for the rectified map H_t in the Gaussian-to-Gaussian-mixture regime. The core lemma identifies a transversality mechanism intrinsic to the convex time-mixing structure: at any collision $H_t(x) = H_t(y)$ with $x \neq y$, the derivative in t is explicitly nonzero. Combined with strip-wise regularity of RF flows (supported by analytic mixture-score formulae and global well-posedness under an Osgood criterion), this mechanism yields that, for Lebesgue-almost every time t, collision sets on any compact domain have Lebesgue measure zero. As a result, for almost every t, the map H_t is injective almost surely under Gaussian initialization, providing a probabilistic notion of invertibility that aligns with the operational regime of rectified flow sampling.

Several directions remain natural. Establishing full global invertibility for almost every t may be approached by coercivity/properness estimates leveraging Gaussian-mixture tail structure. More broadly, generic-in-time arguments of the present type may inform sharper error bounds that depend on the geometric "straightness" of learned flows rather than uniform monotonicity.

A ODE flow regularity on time strips

This appendix records a standard statement from ODE flow theory suitable for citation and for verifying Assumption 1.

Proposition 1 (Flow regularity on compact time intervals)**.** *Let* $[\varepsilon, 1-\varepsilon] \subset (0,1)$ *and let* $U \subset \mathbb{R}^d$ *be open. Suppose* $v : [\varepsilon, 1-\varepsilon] \times U \to \mathbb{R}^d$ *is continuous in* t *and* C^1 *in* x*, and that for every compact* $K \subset U$ *there exists* $L_K < \infty$ *such that* $\|\nabla_x v_t(x)\| \le L_K$ *for all* $t \in [\varepsilon, 1-\varepsilon]$ *and* $x \in K$*. Then for each* $t_0 \in [\varepsilon, 1-\varepsilon]$ *and* $x_0 \in U$*, the ODE* $\dot{Z}_t = v_t(Z_t)$ *admits a unique solution on* $[\varepsilon, 1-\varepsilon]$ *(until exiting* U*), and the flow map* $x_0 \mapsto Z_t(x_0)$ *is* C^1 *on compact subsets that remain in* U*.*

Remark 2. In the Gaussian-to-GMM regime, Appendix A.4 of [1] provides explicit expressions for v_t and $\nabla_x v_t$, and these are smooth in x for every fixed $t \in (0,1)$, with uniform bounds on compacts for t restricted to $[\varepsilon, 1-\varepsilon]$. Together with global existence from the Osgood criterion (Assumption 4.4 and Theorem 4.5 in [1]), Proposition 1 yields the strip-wise regularity used in the main argument.

B Additional background: collision sets and slicing

For completeness, this appendix notes the measure-theoretic step used in Theorem 1.

Proposition 2 (Fubini slicing for null sets)**.** *Let* $A \subset (0,1) \times \mathbb{R}^{2d}$ *be measurable. If* $\mathrm{Leb}_{2d+1}(A) = 0$*, then for Lebesgue-almost every* $t \in (0,1)$*, the slice* $A_t := \{(x,y) \in \mathbb{R}^{2d} : (t,x,y) \in A\}$ *satisfies* $\mathrm{Leb}_{2d}(A_t) = 0$.

References

[1] Vansh Bansal, Saptarshi Roy, Purnamrita Sarkar, and Alessandro Rinaldo. *Straightness of Rectified Flow: A Theoretical Insight into Wasserstein Convergence.* arXiv preprint arXiv:2410.14949, 2024.

[2] Xingchao Liu, Chengyue Gong, and Qiang Liu. *Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow.* arXiv preprint arXiv:2209.03003, 2022.

[3] Qiang Liu. *Rectified Flow: A Marginal Preserving Approach to Optimal Transport.* arXiv preprint arXiv:2209.14577, 2022.

[4] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. *Flow Matching for Generative Modeling.* arXiv preprint arXiv:2210.02747, 2022.

[5] Johannes Hertrich, Antonin Chambolle, Julie Delon. *On the Relation between Rectified Flows and Optimal Transport.* arXiv preprint arXiv:2505.19712, 2025.

[6] John M. Lee. *Introduction to Smooth Manifolds.* Springer, 2nd edition, 2012.

[7] Victor Guillemin and Alan Pollack. *Differential Topology.* Prentice-Hall, 1974.