INTRINSIC FLOW MATCHING ON QUANTUM PURE-STATE MANIFOLDS WITH PHASE-ALIGNED TRANSPORT

Anonymous authors
Paper under double-blind review

ABSTRACT

Quantum pure-state ensembles live on complex projective space, making flat Euclidean generative modeling geometrically mismatched. We introduce *Intrinsic Flow Matching* (IFM), a deterministic transport framework on $\mathbb{CP}^{d-1}$ that learns tangent velocity fields using Pancharatnam phase-aligned conditional paths. IFM replaces local score teachers and reverse-time stochastic sampling with manifold probability flow, while horizontal parameterization removes redundant ambient directions. We show that the IFM objective recovers the induced marginal transport field, represents deterministic projective ensemble flows, and yields endpoint and stability guarantees. Empirically, IFM often improves over ambient Euclidean flow matching across higher-qubit, multimodal, spin-coherent, physics-inspired, and amplitude-encoded MNIST image-vector benchmarks, with strongest gains on high-dimensional and coherence-sensitive tasks but not uniformly across every metric.

1 INTRODUCTION

Generative modeling has become a powerful framework for high-dimensional data (Ho et al., 2020; Song et al., 2021b; Lipman et al., 2023), and distributions over quantum pure states arise naturally as mathematical objects in many-body physics, quantum information, and quantum machine learning (Bengtsson & Życzkowski, 2017; Biamonte et al., 2017; Preskill, 2018; Schuld & Killoran, 2019). Unlike Euclidean data, pure states are defined only up to global phase and live on the projective manifold $\mathbb{CP}^{d-1}$ (Mielnik, 1968; Bengtsson & Życzkowski, 2017; Ogiue, 1969; Wu & Daoudi, 2022); ambient generative models therefore ignore gauge redundancy and may learn transport in directions with no physical meaning.

Recent intrinsic diffusion approaches, in particular local-time Riemannian score matching on the pure-state manifold (Gisin & Percival, 1992; Bouten et al., 2004; Xu et al., 2026), abbreviated SSDM below, show that quantum-state generation can be posed directly on the pure-state manifold. However, they still inherit diffusion machinery such as local score estimation, teacher approximations, and reverse-time stochastic simulation, which can become difficult as the Hilbert-space dimension grows.

We argue that quantum pure-state generation is better treated as intrinsic probability transport on $\mathbb{CP}^{d-1}$. This leads to *Intrinsic Flow Matching* (IFM), a flow-based framework that learns tangent velocity fields through a geometry-aware flow matching objective (Lipman et al., 2023; Chen & Lipman, 2024; Tong et al., 2024). IFM replaces local diffusion teachers and reverse stochastic sampling with direct velocity matching and deterministic manifold sampling, while its phase-aligned paths use the Pancharatnam in-phase convention (Pancharatnam, 1956; Garza-Soto et al., 2023; Ferrer-Garcia et al., 2023) to remove unphysical gauge rotation. The resulting dynamics admit a probability-current interpretation on quantum state space (Delgado, 1999; Hodge et al., 2014; Wyatt, 2005; Tsekov, 2009).

Beyond the change of generative paradigm, we make the transport formulation mathematically explicit. The IFM population objective recovers the marginal transport field induced by a chosen path family, the corresponding marginal density evolves under the intrinsic continuity equation on $\mathbb{CP}^{d-1}$, and the conditional objective differs from its marginal regression counterpart only by an additive

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.

constant, paralleling the conditional-to-marginal reduction underlying modern flow-matching theory. Deterministic tangent flows on projective state space induce exactly the ensemble dynamics IFM models, giving a precise sense in which the framework is compatible with physics-inspired transport viewpoints based on probability currents. Under deterministic Monge couplings and Fubini–Study geodesic interpolation, IFM recovers the displacement-interpolation velocity field, connecting it to geodesic and optimal-transport constructions on quantum state space (Bengtsson & Życzkowski, 2017; Tong et al., 2024; Chewi et al., 2021).

A central question is whether the benefit comes merely from replacing diffusion with flow matching, or whether *intrinsic geometry itself* matters. We separate the two by comparing flow matching performed intrinsically on $\mathbb{CP}^{d-1}$ against flow matching in an ambient Euclidean representation that projects back to valid quantum states. The advantage of IFM turns out not to come solely from switching from diffusion to flow, but also from respecting the intrinsic geometry of pure-state space, most clearly on higher-qubit single-cluster scaling, coherence-sensitive multimodal and spin-coherent benchmarks, and an amplitude-encoded MNIST image-vector benchmark.

Our contributions are summarized as follows:

- We reformulate quantum pure-state generation as intrinsic probability transport on $\mathbb{CP}^{d-1}$ and propose *Intrinsic Flow Matching* (IFM), which learns tangent velocity fields using Pancharatnam-aligned conditional paths.
- We prove that Pancharatnam alignment yields a canonical endpoint representative, that IFM recovers the induced marginal transport field, and that horizontal intrinsic parameterization removes redundant ambient directions.
- We empirically compare IFM with Euclidean flow matching and intrinsic diffusion baselines across controlled single-cluster, multimodal, physics-inspired, and MNIST-derived amplitude-vector benchmarks, including dense 10- and selected 12-qubit tests plus an appendix 14-qubit stress test.

2 Related Work

Intrinsic diffusion for quantum pure-state ensembles. Diffusion and score-based models are central in modern generative modeling (Ho et al., 2020; Song et al., 2021b;a), including recent quantum variants (Zhang et al., 2024; Chen et al., 2024; Kölle et al., 2024; Kwun et al., 2025). The closest precursor is SSDM, which formulates score-based generation intrinsically on the quantum pure-state manifold using Fubini–Study geometry, a local-time teacher construction, and reverse-time sampling (Gisin & Percival, 1992; Bouten et al., 2004; Xu et al., 2026). IFM keeps the same projective state space but replaces score learning and stochastic reversal with deterministic velocity matching.

Flow matching on manifolds and transport geometries. Flow matching replaces score estimation and reverse stochastic simulation with direct regression to path-induced velocity fields (Lipman et al., 2023; Tong et al., 2024). Rectified, optimal, and Riemannian variants further emphasize the roles of path geometry, transport couplings, and manifold structure (Liu et al., 2022; Liu, 2022; Lee et al., 2024; Mathieu & Nickel, 2020; Huang et al., 2022; Bortoli et al., 2022; Lou et al., 2023; Tang & Yang, 2024; Falorsi, 2021; Miller et al., 2024; Klein et al., 2023; Gat et al., 2024; Chen & Lipman, 2024). IFM specializes this logic to quantum pure states by using Pancharatnam-aligned endpoints, horizontal tangent fields, and projective transport on $\mathbb{CP}^{d-1}$.

Alternative quantum representations for flow-based generation. *Quantum Flow Matching* (QFM) studies circuit-based density-matrix interpolation (Cui et al., 2025), while related flow generators operate on spin Wigner functions, functions, or classical-shadow representations (Hahm et al., 2024; Kerrigan et al., 2023; Hahm & Clark, 2026). These works are conceptually complementary but use different state representations: IFM transports pure-state ensembles directly on $\mathbb{CP}^{d-1}$. Appendix A extends this discussion to optimal transport on structured spaces, geometric phase and Pancharatnam alignment, Riemannian flow matching on general geometries, and physics-inspired transport viewpoints.

3 PRELIMINARIES

Quantum pure-state distributions on $\mathbb{CP}^{d-1}$. An n-qubit pure state is a unit vector $\psi \in \mathbb{C}^d$, $d = 2^n$, modulo global phase, so the state space is (Mielnik, 1968; Bengtsson & Życzkowski, 2017; Ogiue, 1969; Wu & Daoudi, 2022)

$$\mathbb{CP}^{d-1} = \{\psi \in \mathbb{C}^d : \|\psi\|_2 = 1\}/U(1). \tag{1}$$

We seek a generative map from a base density p_0 to a target ensemble p_1 on this projective manifold, so both density evolution and velocity fields should be intrinsic rather than tied to arbitrary ambient representatives.

Stochastic diffusion versus deterministic transport. Intrinsic diffusion models on $\mathbb{CP}^{d-1}$ lead to score-based objectives and reverse-time stochastic simulation (Xu et al., 2026). We instead model a density p_t moving under a tangent velocity field $v_t(\psi) \in T_\psi \mathbb{CP}^{d-1}$, governed by the manifold continuity equation (Delgado, 1999; Wyatt, 2005; Chewi et al., 2021)

$$\partial_t p_t + \mathrm{div}_{\mathbb{CP}}(p_t v_t) = 0, \tag{2}$$

where $\mathrm{div}_{\mathbb{CP}}$ is the intrinsic divergence. Generation is then realized by learning a velocity field that transports p_0 toward p_1 along manifold-consistent trajectories.

Flow matching. Flow matching learns a time-dependent vector field by regressing to conditional velocities induced by a prescribed path between base and data samples (Lipman et al., 2023; Liu et al., 2022; Tong et al., 2024). Sampling then integrates the learned ordinary differential equation

$$\frac{dx_t}{dt} = v_\theta(x_t, t). \tag{3}$$

We lift this principle from Euclidean space to quantum pure-state manifolds, following recent manifold and Riemannian flow-matching developments (Chen & Lipman, 2024; Miller et al., 2024; Klein et al., 2023).

4 INTRINSIC FLOW MATCHING ON QUANTUM PURE-STATE MANIFOLDS

We introduce *Intrinsic Flow Matching* (IFM), which models quantum pure-state ensembles as deterministic probability transport on $\mathbb{CP}^{d-1}$ (Mielnik, 1968; Bengtsson & Życzkowski, 2017; Lipman et al., 2023; Chen & Lipman, 2024). Rather than using stochastic diffusion and reverse-time score estimation, IFM learns a time-dependent tangent velocity field that transports a base density p_0 to a target ensemble density p_1 along geometry-consistent projective trajectories.

4.1 INTRINSIC TRANSPORT FORMULATION

Let $\psi_t \in \mathbb{CP}^{d-1}$ denote a pure-state trajectory and let

$$v_t(\psi) \in T_\psi \mathbb{CP}^{d-1} \tag{4}$$

be a tangent velocity field. The induced density evolution is governed by the manifold continuity equation (Chewi et al., 2021; Delgado, 1999; Wyatt, 2005)

$$\partial_t p_t + \mathrm{div}_{\mathbb{CP}}(p_t v_t) = 0. \tag{5}$$

We learn a neural approximation $v_\theta(\psi, t)$ so that integrating the flow from p_0 produces samples from p_1. Using unit-sphere representatives and quotienting out global phase, we identify the projective tangent space with the horizontal subspace

$$\mathcal{H}_\psi := \left\{\xi \in \mathbb{C}^d : \langle \psi, \xi \rangle = 0\right\}. \tag{6}$$

This removes both radial motion and unphysical phase rotation. Under the Fubini–Study metric, the horizontal-coordinate inner product is (Mielnik, 1968; Bengtsson & Życzkowski, 2017)

$$g_{\mathrm{FS},\psi}(\xi, \eta) := \mathrm{Re}\,\langle \xi, \eta \rangle, \qquad \|\xi\|_\psi^2 := g_{\mathrm{FS},\psi}(\xi, \xi). \tag{7}$$

All IFM vector fields, losses, and numerical updates are defined in this tangent-space representation. Compared with diffusion-based quantum generative models (Ho et al., 2020; Song et al., 2021b; Xu et al., 2026), this replaces score learning, local diffusion teachers, and reverse-time stochastic simulation with direct velocity matching and deterministic manifold flow integration.

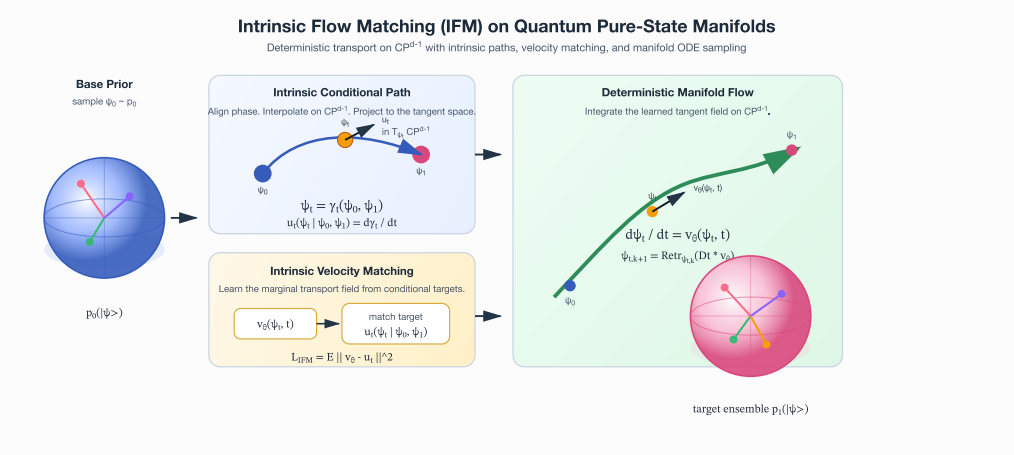

Figure 1: IFM on $\mathbb{CP}^{d-1}$: phase-align endpoints, regress tangent velocities, and sample by deterministic manifold flow.

4.2 Conditional manifold paths

Flow matching requires conditional paths between base and target samples (Lipman et al., 2023; Tong et al., 2024). For $\psi_0 \sim p_0$ and $\psi_1 \sim p_1$, we define a path on projective pure states,

$$\psi_t = \gamma_t(\psi_0, \psi_1), \qquad t \in [0, 1], \tag{8}$$

with endpoints $\gamma_0 = \psi_0$ and $\gamma_1 = \psi_1$, and regress to its tangent velocity $u_t = \frac{d}{dt}\gamma_t$. Because $[\psi] = [e^{i\phi}\psi]$, the endpoint representatives must be gauge-fixed before interpolation. We use the Pancharatnam in-phase convention, choosing the representative of ψ_1 whose overlap with ψ_0 is real and nonnegative (Pancharatnam, 1956).

For pairs with $\langle\psi_0, \psi_1\rangle \neq 0$, define

$$\alpha(\psi_0, \psi_1) := \frac{\langle\psi_0, \psi_1\rangle}{|\langle\psi_0, \psi_1\rangle|}, \qquad \bar{\psi}_1 := \alpha(\psi_0, \psi_1)^{-1}\psi_1. \tag{9}$$

Proposition 8 (Appendix C.1) proves that this is the Pancharatnam-aligned representative: it maximizes the real endpoint overlap (Lemma 1) and minimizes the ambient chord length (Lemma 2) among phase-equivalent representatives.

When $\langle\psi_0, \psi_1\rangle = 0$ the alignment phase is undefined and no representative is singled out; we then fall back to $\bar{\psi}_1 = \psi_1$, a convention that fixes a path but does depend on the supplied representative. Such pairs form a measure-zero set under our continuous base and data distributions. The *nearly* orthogonal case is benign: Appendix D.16 derives a uniform bound on the conditional velocity and reports a Haar-base stress test in which almost every coupling is close to orthogonal.

We then use normalized chord interpolation,

$$\tilde{\psi}_t := (1 - t)\psi_0 + t\bar{\psi}_1, \tag{10}$$

$$\psi_t = \gamma_t(\psi_0, \psi_1) := \frac{\tilde{\psi}_t}{\|\tilde{\psi}_t\|_2}. \tag{11}$$

This gives a closed-form, inexpensive projective path that avoids spurious phase motion. Appendix C.1 develops this Pancharatnam-alignment viewpoint in more detail, and Appendix C.2 states the full interpolation rules used in our experiments, including the geodesic IFM variant and the generic manifold-geodesic baseline.

Differentiating the normalized path yields an explicit tangent velocity. Let

$$\dot{\tilde{\psi}}_t = \bar{\psi}_1 - \psi_0. \tag{12}$$

Using $\psi_t = \tilde{\psi}_t / \|\tilde{\psi}_t\|_2$, we obtain

$$u_t(\psi_t \mid \psi_0, \psi_1) := \dot{\psi}_t = \frac{1}{\|\tilde{\psi}_t\|_2} \Pi_{\psi_t}\big(\bar{\psi}_1 - \psi_0\big)\,, \tag{13}$$

where Π_{ψ_t} denotes the orthogonal projection onto the horizontal tangent space at ψ_t. Hence $u_t(\psi_t \mid \psi_0, \psi_1) \in \mathcal{H}_{\psi_t}$ by construction. Thus IFM trains a neural tangent field against a phase-aligned, horizontally projected conditional transport direction. Proposition 9 (Appendix C.4) shows that the resulting path is well-defined on projective classes and that its velocity always lies in the horizontal tangent bundle.

4.3 Intrinsic flow matching objective

Let $v_\theta(\psi, t) \in T_\psi \mathbb{CP}^{d-1}$ be a tangent neural velocity field, implemented by projecting the network output onto the tangent space. For (ψ_0, ψ_1) and $t \sim \mathrm{Unif}[0,1]$, set $\psi_t = \gamma_t(\psi_0, \psi_1)$ and minimize

$$\mathcal{L}_{\mathrm{IFM}}(\theta) = \mathbb{E}_{\psi_0 \sim p_0,\ \psi_1 \sim p_1,\ t \sim \mathrm{Unif}[0,1]} \left[\|v_\theta(\psi_t, t) - u_t(\psi_t \mid \psi_0, \psi_1)\|^2_{\psi_t} \right], \tag{14}$$

where $\|\cdot\|_{\psi_t}$ denotes the norm induced by the manifold metric at ψ_t.

This is Euclidean flow matching lifted to $\mathbb{CP}^{d-1}$ (Lipman et al., 2023; Tong et al., 2024): both the target and learned velocities are tangent, so the learned transport is intrinsic rather than ambient.

4.4 Population characterization of the IFM objective

We now characterize the vector field recovered by learning the phase-aligned target in expectation. The results identify the population minimizer, connect it to marginal continuity dynamics and endpoint transport, and give a Wasserstein stability bound for approximate minimizers. Let (ψ_0, ψ_1) be sampled from a joint distribution π whose marginals are p_0 and p_1, and define the intermediate random variable

$$\psi_t = \gamma_t(\psi_0, \psi_1), \qquad p_t := \text{distribution of } \psi_t. \tag{15}$$

The conditional velocity $u_t(\psi_t \mid \psi_0, \psi_1)$ induces a probability current along this path family, paralleling current-based descriptions of quantum density evolution (Delgado, 1999; Hodge et al., 2014; Wyatt, 2005). The key question is whether minimizing $\mathcal{L}_{\mathrm{IFM}}$ recovers the correct marginal transport field on $\mathbb{CP}^{d-1}$.

Theorem 1 (Population optimality of intrinsic flow matching)**.** *Fix $t \in (0,1)$ and consider the population objective over measurable tangent vector fields $v(\cdot, t) : \mathbb{CP}^{d-1} \to T\mathbb{CP}^{d-1}$,*

$$\mathcal{J}_t[v] := \mathbb{E}\left[\|v(\psi_t, t) - u_t(\psi_t \mid \psi_0, \psi_1)\|^2_{\psi_t}\right]. \tag{16}$$

Then any minimizer $v_t^\star$ of $\mathcal{J}_t$ satisfies, for p_t-almost every ψ,

$$v_t^\star(\psi) = \mathbb{E}\big[u_t(\psi_t \mid \psi_0, \psi_1) \,\big|\, \psi_t = \psi\big]. \tag{17}$$

In particular, the population IFM objective learns the conditional-expectation velocity field associated with the chosen manifold path family. Proof in Appendix C.5.

Remark 1. *Theorem 1 plays the role that score-recovery results play in diffusion models (Song et al., 2021b; Ho et al., 2020): it identifies the population object learned by the objective.*

Proposition 1 (Conditional-to-marginal IFM objective decomposition)**.** *Let*

$$\bar{u}_t(\psi) := \mathbb{E}\big[u_t(\psi_t \mid \psi_0, \psi_1) \,\big|\, \psi_t = \psi\big] \tag{18}$$

denote the marginal velocity field induced by the conditional path family. Then for every measurable tangent vector field $v(\cdot, t)$,

$$\mathcal{J}_t[v] = \mathbb{E}\big[\|v(\psi_t, t) - \bar{u}_t(\psi_t)\|^2_{\psi_t}\big] + \mathbb{E}\big[\|u_t(\psi_t \mid \psi_0, \psi_1) - \bar{u}_t(\psi_t)\|^2_{\psi_t}\big]. \tag{19}$$

In particular, the conditional IFM objective and the marginal regression objective differ by an additive constant independent of v, and therefore have the same minimizers. Proof in Appendix C.6.

Proposition 2 (Continuity equation induced by the population IFM field)**.** *Assume $t \mapsto \gamma_t(\psi_0, \psi_1)$ is C^1 and that the induced marginal law p_t admits a smooth density with respect to the Fubini–Study volume measure. Let $v_t^\star$ be defined by* (17)*. Then p_t satisfies the intrinsic continuity equation*

$$\partial_t p_t + \mathrm{div}_{\mathbb{CP}}(p_t v_t^\star) = 0 \tag{20}$$

in the weak sense on $\mathbb{CP}^{d-1}$. Proof in Appendix C.7.

Remark 2. *Proposition 2 formalizes the transport interpretation emphasized in the abstract and introduction: the field learned by IFM is not merely a regression target, but the velocity field governing density evolution on $\mathbb{CP}^{d-1}$.*

Theorem 2 (Endpoint-preserving transport)**.** *Suppose the conditional path satisfies $\gamma_0(\psi_0, \psi_1) = \psi_0$ and $\gamma_1(\psi_0, \psi_1) = \psi_1$ almost surely, and that the population minimizer $v_t^\star$ from Theorem 1 is locally Lipschitz on $\mathbb{CP}^{d-1}$ uniformly in $t \in [0, 1]$. Then the manifold ODE*

$$\frac{d}{dt}\psi_t = v_t^\star(\psi_t), \qquad \psi_0 \sim p_0, \tag{21}$$

admits a unique flow $\Phi_{0,t} : \mathbb{CP}^{d-1} \to \mathbb{CP}^{d-1}$, and for every $t \in [0, 1]$,

$$(\Phi_{0,t})_{\#} p_0 = p_t. \tag{22}$$

In particular, $(\Phi_{0,1})_{\#} p_0 = p_1$. Proof in Appendix C.8.

Remark 3. *The population IFM field therefore gives sample-level transport: its deterministic ODE pushes p_0 exactly onto p_1, not merely onto matching intermediate marginals. The argument combines Proposition 2 with uniqueness of weak solutions to the manifold continuity equation under Lipschitz transport (Ambrosio, 2008; Villani, 2009).*

Proposition 3 (L^2-to-Wasserstein stability)**.** *Let $\hat{v} : \mathbb{CP}^{d-1} \times [0, 1] \to T\mathbb{CP}^{d-1}$ be a tangent vector field with flow $\hat{\Phi}_{0,t}$, and let $\hat{p}_t := (\hat{\Phi}_{0,t})_{\#} p_0$. Assume both $v_t^\star$ and $\hat{v}_t$ are L-Lipschitz on $\mathbb{CP}^{d-1}$ in the unit-sphere representation, uniformly in $t \in [0, 1]$, and that*

$$\int_0^1 \mathbb{E}_{\psi \sim p_t} \left\| \hat{v}(\psi, t) - v_t^\star(\psi) \right\|_\psi^2 dt \;\le\; \varepsilon^2. \tag{23}$$

Then the Wasserstein-2 distance with respect to the Fubini–Study metric satisfies

$$W_2^{\mathrm{FS}}\big(\hat{p}_1,\, p_1\big) \;\le\; C_L\, \varepsilon, \tag{24}$$

with a constant $C_L = e^L$ depending only on L. Proof in Appendix C.9.

Remark 4. *Proposition 3 is a quantitative bridge between training error and sample quality: an ε-approximate IFM minimizer in the population L^2 sense yields $\hat{p}_1$ within $O(\varepsilon)$ of p_1 in Fubini–Study Wasserstein distance, with a constant depending only on the common Lipschitz constant L. The bound is the manifold counterpart of Grönwall stability for ODE flows (Villani, 2009; Ambrosio, 2008; Lee, 2012), does not assume $\hat{v}$ minimizes the IFM loss, and its Lipschitz assumption on $\hat{v}$ is met by any network with bounded weights and Lipschitz nonlinearities.*

Proposition 4 (Projective flow compatibility)**.** *Let $X_t(\psi) \in T_\psi \mathbb{CP}^{d-1}$ be a C^1 tangent field with flow $\Phi_{s,t}$ generated by*

$$\frac{d}{dt}\psi_t = X_t(\psi_t). \tag{25}$$

For $p_t = (\Phi_{0,t})_{\#} p_0$, the family (p_t) satisfies

$$\partial_t p_t + \mathrm{div}_{\mathbb{CP}}(p_t X_t) = 0 \tag{26}$$

weakly. If the conditional paths coincide with these trajectories, $\gamma_t(\psi_0, \psi_1) = \Phi_{0,t}(\psi_0)$ and $\psi_1 = \Phi_{0,1}(\psi_0)$, then the population IFM minimizer obeys $v_t^\star(\psi) = X_t(\psi)$ for p_t-a.e. ψ. Proof in Appendix C.10.

Remark 5. *Proposition 4 connects IFM to physics-inspired transport without identifying it with a specific Bohmian, hydrodynamic, or Hamiltonian model: any deterministic tangent ensemble flow lies in IFM's population-level representational scope when the path family matches that flow.*

Proposition 5 (Monge-geodesic specialization). *Assume the coupling π is deterministic, i.e.,*

$$\pi = (\mathrm{Id}, T)_{\#} p_0 \tag{27}$$

for a measurable map $T : \mathbb{CP}^{d-1} \to \mathbb{CP}^{d-1}$; equivalently, after sampling $\psi_0 \sim p_0$ we set $\psi_1 = T(\psi_0)$ almost surely. Assume moreover that the conditional path is chosen as the Fubini–Study geodesic interpolation

$$\psi_t = \gamma_t(\psi_0, T(\psi_0)) \tag{28}$$

and that, for each $t \in (0, 1)$, the map $\psi_0 \mapsto \gamma_t(\psi_0, T(\psi_0))$ is injective on a full-measure subset of p_0. Then the population IFM minimizer satisfies

$$v_t^\star(\psi_t) = \frac{d}{dt}\gamma_t(\psi_0, T(\psi_0)) \tag{29}$$

for p_t-almost every ψ_t. Consequently, IFM exactly recovers the deterministic velocity field of the prescribed geodesic interpolation. If, in addition, T is the optimal Monge map for the quadratic Fubini–Study cost, then $v_t^\star$ coincides with the velocity field of the corresponding displacement interpolation. Proof in Appendix C.11.

Scope of the theory. Theorem 1 and Propositions 1–2 are direct $\mathbb{CP}^{d-1}$ analogues of standard flow-matching results (Lipman et al., 2023; Tong et al., 2024); the geometry-specific parts are the Pancharatnam path, horizontal minimal-norm parameterization, endpoint transport, Fubini–Study stability, and Monge–geodesic specialization. Neural parameterization, projection, and retraction sampling are in Appendices B.4 and B.6.

5 Experiments

We evaluate whether phase-aligned intrinsic transport on $\mathbb{CP}^{d-1}$ improves over ambient Euclidean transport for quantum pure-state generation. The main baseline is *Euclidean flow matching* (Euclidean FM), which uses the same flow-matching objective but learns in an ambient real state-vector representation followed by normalization.

5.1 Experimental setup

All experiments use dense pure-state ensembles on n qubits, with Hilbert-space dimension $d = 2^n$ (Bengtsson & Życzkowski, 2017; Preskill, 2018). We evaluate controlled single-cluster scaling, coherence-sensitive multimodal ensembles (whose Bloch-sphere geometry is illustrated in Appendix Figure 3), spin-coherent peaks, compact TFIM/XXZ ground-state families, and MNIST-derived amplitude-encoded image-vector states, with an additional 14-qubit stress test (Appendix D.8, Table 10, Figure 4). These benchmarks are chosen to expose geometric effects in a controlled way. Unless otherwise stated, IFM and Euclidean FM are trained with matched network scale, optimization budget, sample budget, and evaluation protocol, and main tables report mean $\pm$ standard deviation over 10 seeds. Appendix D gives the benchmark construction, training settings, architectures, baseline configurations, and reporting conventions, together with a summary heatmap of IFM-to-Euclidean-FM error ratios (Figure 2); Appendix B collects the scalability discussion, the comparison with SSDM, and the pseudocode (Algorithm 1).

5.2 Baseline and metrics

Our main baseline is Euclidean FM (Lipman et al., 2023), which uses the same flow-matching principle in an ambient state-vector representation. We also compare with SSDM (Xu et al., 2026), Euclidean VP-SDE (Song et al., 2021b), QDDPM (Zhang et al., 2024), QGAN (Lloyd & Weedbrook, 2018; Dallaire-Demers & Killoran, 2018), and intrinsic geodesic-flow ablations where relevant. We report MMD (Gretton et al., 2012), observable discrepancy Δ_{obs}, and entanglement Wasserstein distance. The three are complementary by construction: MMD under a phase-invariant fidelity kernel measures distributional agreement on $\mathbb{CP}^{d-1}$, Δ_{obs} measures error in physically measurable expectations, and the entanglement Wasserstein distance compares half-chain entanglement spectra, which no single-copy observable resolves. All definitions and baseline-specific settings are in Appendix D.1–D.5.

Table 1: Representative 6-qubit single-cluster comparison across diffusion, circuit, and flow baselines. Lower is better. Entries are mean $\pm$ standard deviation across 10 random seeds (Appendix D.1). The geodesic IFM row is an ablation of our method.

Method	MMD $\downarrow$	$\Delta_{\text{obs}} \downarrow$	Ent. $W_1 \downarrow$
Euclidean VP-SDE (Song et al., 2021b)	$(5.12 \pm 0.41) \times 10^{-1}$	$(6.34 \pm 0.38) \times 10^{-1}$	$(6.70 \pm 0.42) \times 10^{-1}$
SSDM (Xu et al., 2026)	$(1.73 \pm 0.14) \times 10^{-1}$	$(4.16 \pm 0.27) \times 10^{-1}$	$(5.59 \pm 0.31) \times 10^{-1}$
QDDPM (Zhang et al., 2024)	$(4.78 \pm 0.36) \times 10^{-1}$	$(6.71 \pm 0.39) \times 10^{-1}$	$(6.51 \pm 0.41) \times 10^{-1}$
QGAN (Lloyd & Weedbrook, 2018; Dallaire-Demers & Killoran, 2018)	$(4.76 \pm 0.41) \times 10^{-1}$	$(6.76 \pm 0.45) \times 10^{-1}$	$(6.83 \pm 0.46) \times 10^{-1}$
Euclidean FM (Lipman et al., 2023)	$(6.02 \pm 0.51) \times 10^{-3}$	$(2.16 \pm 0.18) \times 10^{-3}$	$(8.15 \pm 0.62) \times 10^{-3}$
Geodesic IFM (ours variant)	$(2.49 \pm 0.21) \times 10^{-4}$	$(2.48 \pm 0.19) \times 10^{-3}$	$(4.34 \pm 0.31) \times 10^{-3}$
IFM (ours)	$(2.63 \pm 0.18) \times 10^{-4}$	$(2.34 \pm 0.17) \times 10^{-3}$	$(4.66 \pm 0.29) \times 10^{-3}$

Table 2: Single-cluster scaling against Euclidean FM (Lipman et al., 2023). Lower is better; "≈ 0" denotes output-precision rounding. Entries are mean $\pm$ standard deviation across 10 random seeds (Appendix D.1).

Qubits	Method	MMD $\downarrow$	$\Delta_{\text{obs}} \downarrow$	Ent. $W_1 \downarrow$
2	Euclidean FM	$(2.45 \pm 0.21) \times 10^{-4}$	$(5.94 \pm 0.42) \times 10^{-3}$	$(5.45 \pm 0.38) \times 10^{-3}$
2	IFM (ours)	≈ 0	$(1.13 \pm 0.09) \times 10^{-3}$	$(1.67 \pm 0.12) \times 10^{-3}$
6	Euclidean FM	$(6.02 \pm 0.51) \times 10^{-3}$	$(2.16 \pm 0.18) \times 10^{-3}$	$(8.15 \pm 0.62) \times 10^{-3}$
6	IFM (ours)	$(2.63 \pm 0.18) \times 10^{-4}$	$(2.34 \pm 0.17) \times 10^{-3}$	$(4.66 \pm 0.29) \times 10^{-3}$
10	Euclidean FM	$(1.68 \pm 0.13) \times 10^{-2}$	$(9.51 \pm 0.71) \times 10^{-2}$	$(8.21 \pm 0.55) \times 10^{-2}$
10	IFM (ours)	$(7.38 \pm 0.51) \times 10^{-4}$	$(1.39 \pm 0.12) \times 10^{-3}$	$(1.46 \pm 0.11) \times 10^{-3}$
12	Euclidean FM	$(8.76 \pm 0.65) \times 10^{-3}$	$(3.08 \pm 0.22) \times 10^{-2}$	$(1.50 \pm 0.10) \times 10^{-2}$
12	IFM (ours)	$(6.54 \pm 0.45) \times 10^{-4}$	$(1.81 \pm 0.14) \times 10^{-3}$	$(1.15 \pm 0.09) \times 10^{-3}$

5.3 Single-cluster scaling

Table 2 shows that IFM improves over Euclidean FM across the single-cluster scaling sweep, with the largest gains appearing at 10 and 12 qubits. The representative 6-qubit comparison in Table 1 further shows that IFM is substantially stronger than SSDM, Euclidean VP-SDE, QDDPM, and QGAN on this benchmark. To rule out undertraining as a source of the gap, we ran a 20,000-step comparison on 14-qubit single-cluster (Figure 4); the Euclidean FM gap persists. Appendix D.7 gives the scaling interpretation, the 10-qubit scalable-baseline extension, and training-time details.

5.4 Structured multimodal and coherence-sensitive ensembles

Table 3 shows the clearest IFM gains on 10-qubit multimodal, spin-coherent, and compact physics-inspired benchmarks. The advantage is strongest when modes differ by projective phase structure or coherent product-state organization, while the metric trade-offs remain mixed on some smoother families. Appendices D.9–D.13 give the 6-qubit extensions, the phase-alignment ablation, the geometry/path-choice summary, and the generic manifold-geodesic comparison. That last baseline is competitive on several 6-qubit tasks and wins on XXZ, so our claim is not that IFM dominates every intrinsic path, but that quotient-aware alignment often helps on coherent projective structure.

5.5 Physics-inspired families and amplitude-encoded image benchmarks

The TFIM and XXZ families provide compact physics-inspired tests of transport over spin-chain ground-state families. At 2 qubits, IFM improves MMD and slightly improves observable discrepancy on both families, while Euclidean FM remains competitive on Ent. W_1. Within these compact families, the 10-qubit results are more favorable to IFM: it improves all three metrics on both TFIM and XXZ relative to Euclidean FM and SSDM. The MNIST-derived benchmark gives a complementary classical image-vector stress test after amplitude encoding; at both 6 and 12 qubits, this embedded image distribution favors IFM over Euclidean FM across all reported metrics. Together with the structured multimodal benchmarks and the appendix 14-qubit stress test, these results indicate that IFM is most useful in controlled settings where the distribution is high-dimensional, multimodal, coherence-sensitive, or naturally posed on projective state space.

Table 3: Structured, physics-inspired, and amplitude-encoded MNIST benchmarks. Lower is better. Baselines are Euclidean FM (Lipman et al., 2023) and SSDM (Xu et al., 2026); 6-qubit extensions are in Appendix D.10. Entries are mean $\pm$ standard deviation across 10 random seeds (Appendix D.1).

Benchmark	Method	MMD $\downarrow$	$\Delta_{\text{obs}} \downarrow$	Ent. $W_1 \downarrow$
Eq. bimodal (2 qubits)	Euclidean FM	$(4.06 \pm 0.31) \times 10^{-2}$	$(1.58 \pm 0.11) \times 10^{-1}$	$(2.16 \pm 0.15) \times 10^{-1}$
Eq. bimodal (2 qubits)	SSDM	$(2.39 \pm 0.18) \times 10^{-2}$	$(4.48 \pm 0.32) \times 10^{-2}$	$(2.38 \pm 0.17) \times 10^{-2}$
Eq. bimodal (2 qubits)	IFM (ours)	$(2.33 \pm 0.17) \times 10^{-2}$	$(1.05 \pm 0.08) \times 10^{-1}$	$(7.55 \pm 0.49) \times 10^{-2}$
Eq. bimodal (10 qubits)	Euclidean FM	$(1.22 \pm 0.09) \times 10^{-2}$	$(4.60 \pm 0.32) \times 10^{-2}$	$(4.60 \pm 0.31) \times 10^{-2}$
Eq. bimodal (10 qubits)	SSDM	$(7.11 \pm 0.51) \times 10^{-3}$	$(5.64 \pm 0.40) \times 10^{-2}$	$(5.04 \pm 0.36) \times 10^{-2}$
Eq. bimodal (10 qubits)	IFM (ours)	$(8.38 \pm 0.61) \times 10^{-4}$	$(7.99 \pm 0.55) \times 10^{-3}$	$(1.64 \pm 0.11) \times 10^{-2}$
Trimodal (2 qubits)	Euclidean FM	$(1.75 \pm 0.13) \times 10^{-2}$	$(1.00 \pm 0.07) \times 10^{-1}$	$(1.13 \pm 0.08) \times 10^{-1}$
Trimodal (2 qubits)	SSDM	$(5.04 \pm 0.36) \times 10^{-3}$	$(6.22 \pm 0.42) \times 10^{-2}$	$(4.18 \pm 0.29) \times 10^{-2}$
Trimodal (2 qubits)	IFM (ours)	≈ 0	$(3.34 \pm 0.24) \times 10^{-2}$	$(9.47 \pm 0.62) \times 10^{-2}$
Trimodal (10 qubits)	Euclidean FM	$(1.27 \pm 0.09) \times 10^{-2}$	$(5.97 \pm 0.40) \times 10^{-2}$	$(5.81 \pm 0.39) \times 10^{-2}$
Trimodal (10 qubits)	SSDM	$(7.78 \pm 0.55) \times 10^{-3}$	$(7.25 \pm 0.51) \times 10^{-2}$	$(6.12 \pm 0.42) \times 10^{-2}$
Trimodal (10 qubits)	IFM (ours)	$(7.08 \pm 0.49) \times 10^{-4}$	$(5.27 \pm 0.36) \times 10^{-3}$	$(3.95 \pm 0.27) \times 10^{-3}$
Spin-coherent (2 qubits)	Euclidean FM	$(4.30 \pm 0.31) \times 10^{-2}$	$(1.12 \pm 0.08) \times 10^{-1}$	$(8.24 \pm 0.55) \times 10^{-2}$
Spin-coherent (2 qubits)	SSDM	$(9.30 \pm 0.65) \times 10^{-3}$	$(4.32 \pm 0.30) \times 10^{-2}$	$(3.58 \pm 0.25) \times 10^{-2}$
Spin-coherent (2 qubits)	IFM (ours)	$(5.27 \pm 0.36) \times 10^{-3}$	$(3.72 \pm 0.26) \times 10^{-2}$	$(3.76 \pm 0.27) \times 10^{-2}$
Spin-coherent (10 qubits)	Euclidean FM	$(1.22 \pm 0.09) \times 10^{-2}$	$(1.43 \pm 0.11) \times 10^{-3}$	$(8.03 \pm 0.54) \times 10^{-2}$
Spin-coherent (10 qubits)	SSDM	$(7.26 \pm 0.51) \times 10^{-3}$	$(1.41 \pm 0.10) \times 10^{-3}$	$(8.68 \pm 0.59) \times 10^{-2}$
Spin-coherent (10 qubits)	IFM (ours)	$(3.81 \pm 0.27) \times 10^{-4}$	$(1.79 \pm 0.13) \times 10^{-3}$	$(4.44 \pm 0.31) \times 10^{-2}$
TFIM (2 qubits)	Euclidean FM	$(4.87 \pm 0.34) \times 10^{-4}$	$(7.43 \pm 0.51) \times 10^{-3}$	$(1.43 \pm 0.10) \times 10^{-2}$
TFIM (2 qubits)	SSDM	$(9.41 \pm 0.66) \times 10^{-4}$	$(5.58 \pm 0.39) \times 10^{-3}$	$(4.20 \pm 0.29) \times 10^{-2}$
TFIM (2 qubits)	IFM (ours)	≈ 0	$(6.68 \pm 0.46) \times 10^{-3}$	$(4.04 \pm 0.28) \times 10^{-2}$
TFIM (10 qubits)	Euclidean FM	$(1.23 \pm 0.09) \times 10^{-2}$	$(3.86 \pm 0.27) \times 10^{-2}$	$(6.63 \pm 0.45) \times 10^{-2}$
TFIM (10 qubits)	SSDM	$(7.06 \pm 0.50) \times 10^{-3}$	$(5.06 \pm 0.36) \times 10^{-2}$	$(7.02 \pm 0.49) \times 10^{-2}$
TFIM (10 qubits)	IFM (ours)	$(1.01 \pm 0.07) \times 10^{-3}$	$(8.91 \pm 0.62) \times 10^{-3}$	$(2.95 \pm 0.20) \times 10^{-2}$
XXZ (2 qubits)	Euclidean FM	$(3.94 \pm 0.28) \times 10^{-2}$	$(1.39 \pm 0.10) \times 10^{-2}$	$(6.43 \pm 0.45) \times 10^{-2}$
XXZ (2 qubits)	SSDM	$(8.60 \pm 0.61) \times 10^{-3}$	$(4.84 \pm 0.34) \times 10^{-2}$	$(7.19 \pm 0.49) \times 10^{-2}$
XXZ (2 qubits)	IFM (ours)	≈ 0	$(1.08 \pm 0.08) \times 10^{-2}$	$(9.39 \pm 0.62) \times 10^{-2}$
XXZ (10 qubits)	Euclidean FM	$(1.31 \pm 0.09) \times 10^{-2}$	$(2.22 \pm 0.16) \times 10^{-2}$	$(4.23 \pm 0.30) \times 10^{-2}$
XXZ (10 qubits)	SSDM	$(8.24 \pm 0.58) \times 10^{-3}$	$(2.68 \pm 0.19) \times 10^{-2}$	$(4.55 \pm 0.32) \times 10^{-2}$
XXZ (10 qubits)	IFM (ours)	$(3.52 \pm 0.24) \times 10^{-4}$	$(3.92 \pm 0.27) \times 10^{-3}$	$(2.50 \pm 0.18) \times 10^{-2}$
MNIST $0/1 \to$ 6q	Euclidean FM	$(2.10 \pm 0.15) \times 10^{-2}$	$(1.30 \pm 0.09) \times 10^{-1}$	$(2.75 \pm 0.18) \times 10^{-1}$
MNIST $0/1 \to$ 6q	IFM (ours)	$(4.44 \pm 0.31) \times 10^{-3}$	$(5.14 \pm 0.36) \times 10^{-2}$	$(5.64 \pm 0.39) \times 10^{-2}$
MNIST $0/1 \to$ 12q	Euclidean FM	$(1.51 \pm 0.10) \times 10^{-1}$	$(2.01 \pm 0.14) \times 10^{-1}$	$(1.70 \pm 0.11) \times 10^{0}$
MNIST $0/1 \to$ 12q	IFM (ours)	$(1.40 \pm 0.10) \times 10^{-2}$	$(9.33 \pm 0.65) \times 10^{-2}$	$(1.10 \pm 0.08) \times 10^{0}$

Controls and robustness. Five appendix studies probe these claims rather than extend them. Matching the ambient baseline's base prior and constraint handling to IFM's (Appendix D.14) leaves the $7\times$–$22\times$ MMD advantage intact, and a 2×2 factorial (Appendix D.17) pins the effect to the gauge fixing: under global-phase-randomized representatives, which leave every state unchanged on $\mathbb{CP}^{d-1}$, alignment buys $1.8\times$–$3.4\times$ in MMD at fixed path family, while path family at fixed alignment buys little. Rerunning every benchmark with an exactly $U(1)$-equivariant parameterization (Appendix D.15) reverses no ordering, and a Haar base, under which almost every coupling is near-orthogonal, trains stably either way (Appendix D.16). On a hardware-efficient circuit ensemble validated on an IBM Heron device, IFM improves all three metrics at 6 and 10 qubits (Appendix D.18).

6 Conclusion

IFM is a transport-based generative framework for pure-state ensembles on $\mathbb{CP}^{d-1}$ that replaces score learning and reverse-time stochastic simulation with learned tangent velocity fields and deterministic manifold sampling, preserving the quotient geometry of pure states. The theory identifies the intrinsic transport field learned for phase-aligned projective paths, and the stability bound links velocity error to sample quality. The gains over Euclidean flow baselines are clearest on multimodal, coherence-sensitive and higher-qubit benchmarks, and survive base matching, a factorial control and exact gauge equivariance.

REPRODUCIBILITY STATEMENT

All theoretical claims are stated with their assumptions in the main text and proved in full in Appendix C; the conditional path constructions used in the experiments are specified in Appendix C.2. The common training and evaluation protocol, benchmark construction, model architectures, baseline-specific hyperparameters, and sampling/reporting conventions are given in Appendix D, and the algorithm is stated in Appendix B.7. Anonymized source code for reproducing all reported experiments is provided as supplementary material (Appendix D.4).

REFERENCES

Luigi Ambrosio. Transport equation and cauchy problem for non-smooth vector fields. In *Calculus of Variations and Nonlinear Partial Differential Equations: With a historical overview by Elvira Mascolo*, pp. 1–41. Springer, 2008.

Mohammad H Amin, Evgeny Andriyash, Jason Rolfe, Bohdan Kulchytskyy, and Roger Melko. Quantum boltzmann machine. *Physical Review X*, 8(2):021050, 2018.

Fortunato Tito Arecchi, Eric Courtens, Robert Gilmore, and Harry Thomas. Atomic coherent states in quantum optics. *Physical Review A*, 6(6):2211, 1972.

Marcello Benedetti, Delfina Garcia-Pintos, Oscar Perdomo, Vicente Leyton-Ortega, Yunseong Nam, and Alejandro Perdomo-Ortiz. A generative modeling approach for benchmarking and training shallow quantum circuits. *npj Quantum information*, 5(1):45, 2019.

Ingemar Bengtsson and Karol Życzkowski. *Geometry of quantum states: an introduction to quantum entanglement*. Cambridge university press, 2017.

Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017.

Valentin De Bortoli, Emile Mathieu, Michael John Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. Riemannian score-based generative modelling. In *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=oDRQGo8I7P.

Sougato Bose. Quantum communication through an unmodulated spin chain. *Physical review letters*, 91(20):207901, 2003.

Luc Bouten, Madalin Guta, and Hans Maassen. Stochastic schrödinger equations. *Journal of Physics A: Mathematical and General*, 37(9):3189, 2004.

Paul M. Chaikin and Tom C. Lubensky. *Principles of Condensed Matter Physics*. Cambridge University Press, 1995.

Chuangtao Chen, Qinglin Zhao, MengChu Zhou, Zhimin He, Zhili Sun, and Haozhen Situ. Quantum generative diffusion model: a fully quantum-mechanical model for generating quantum state ensemble. *arXiv preprint arXiv:2401.07039*, 2024.

Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. In *International Conference on Learning Representations (ICLR)*, 2024.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.

Sinho Chewi, Julien Clancy, Thibaut Le Gouic, Philippe Rigollet, George Stepaniants, and Austin Stromme. Fast and smooth interpolation on wasserstein space. In *International conference on artificial intelligence and statistics*, pp. 3061–3069. PMLR, 2021.

Pedro Cisneros-Velarde and Francesco Bullo. Distributed wasserstein barycenters via displacement interpolation. *IEEE Transactions on Control of Network Systems*, 10(2):785–795, 2023.

Brian Coyle, Daniel Mills, Vincent Danos, and Elham Kashefi. The born supremacy: quantum advantage and training of an ising born machine. *npj Quantum Information*, 6(1):60, 2020.

Zidong Cui, Pan Zhang, and Ying Tang. Quantum flow matching. *arXiv preprint arXiv:2508.12413*, 2025.

Pierre-Luc Dallaire-Demers and Nathan Killoran. Quantum generative adversarial networks. *arXiv preprint arXiv:1804.08641*, 2018.

V Delgado. Quantum probability distribution of arrival times and probability current density. *Physical Review A*, 59(2):1010, 1999.

Luca Falorsi. Continuous normalizing flows on manifolds. *arXiv preprint arXiv:2104.14959*, 2021.

Manuel F Ferrer-Garcia, Kyrylo Snizhko, Alessio D'errico, Alessandro Romito, Yuval Gefen, and Ebrahim Karimi. Topological transitions of the generalized pancharatnam-berry phase. *Science advances*, 9(47):eadg6810, 2023.

Luis Garza-Soto, Nathan Hagen, and Dorilian Lopez-Mago. Deciphering pancharatnam's discovery of geometric phase: retrospective. *Journal of the Optical Society of America A*, 40(5):925–931, 2023.

Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37: 133345–133385, 2024.

Nicolas Gisin and Ian Colin Percival. The quantum-state diffusion model applied to open systems. *Journal of Physics A*, 25:5677–5691, 1992.

Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The journal of machine learning research*, 13(1):723–773, 2012.

Jaehoon Hahm and Bryan K. Clark. Shadowfm: Geometric approaches for learning quantum many-body states with flow matching on classical shadows, 2026. Preprint.

Jaehoon Hahm, Tak Hur, Joonseok Lee, and Daniel K Park. Generative modeling of quantum distribution with functional flow matching. In *Proceedings of the Quantum Techniques in Machine Learning Conference (QTML)*, 2024.

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

WB Hodge, SV Migirditch, and William C Kerr. Electron spin and probability current density in quantum mechanics. *American Journal of Physics*, 82(7):681–690, 2014.

Chin-Wei Huang, Milad Aghajohari, Joey Bose, Prakash Panangaden, and Aaron Courville. Riemannian diffusion models. In *Advances in Neural Information Processing Systems*, 2022. URL `https://openreview.net/forum?id=ecevn9kPm4`.

Chandroth P Jisha, Stree Vithya Arumugam, Lorenzo Marrucci, Stefan Nolte, and Alessandro Alberucci. Waveguiding driven by the pancharatnam-berry phase. *Physical Review A*, 107(1): 013523, 2023.

Gavin Kerrigan, Giosue Migliorini, and Padhraic Smyth. Functional flow matching. *arXiv preprint arXiv:2305.17209*, 2023.

Gavin Kerrigan, Giosue Migliorini, and Padhraic Smyth. Dynamic conditional optimal transport through simulation-free flows. *Advances in Neural Information Processing Systems*, 37:93602–93642, 2024.

Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching. *Advances in Neural Information Processing Systems*, 36:59886–59910, 2023.

Michael Kölle, Gerhard Stenzel, Jonas Stein, Sebastian Zielinski, Björn Ommer, and Claudia Linnhoff-Popien. Quantum denoising diffusion models. In *2024 IEEE International Conference on Quantum Software (QSW)*, pp. 88–98. IEEE, 2024.

Leonid Stefanovich Kuz'menkov and Sergey Gennad'evich Maksimov. Quantum hydrodynamics of particle systems with coulomb interaction and quantum bohm potential. *Theoretical and Mathematical Physics*, 118(2):227–240, 1999.

Gino Kwun, Bingzhi Zhang, and Quntao Zhuang. Mixed-state quantum denoising diffusion probabilistic model. *Physical Review A*, 111(3):032610, 2025.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

John M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, 2 edition, 2012.

Sangyun Lee, Zinan Lin, and Giulia Fanti. Improving the training of rectified flows. *Advances in neural information processing systems*, 37:63082–63109, 2024.

Aleksi Leinonen, Antti Hannonen, Henri Partanen, Janne Heikkinen, Tero Setälä, Ari T Friberg, and Tommi K Hakala. Noncyclic continuous pancharatnam–berry phase in dual-beam interference. *Communications Physics*, 6(1):132, 2023.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations (ICLR)*, 2023.

Jin-Guo Liu and Lei Wang. Differentiable learning of quantum circuit born machines. *Physical Review A*, 98(6):062324, 2018.

Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint arXiv:2209.14577*, 2022.

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.

Seth Lloyd and Christian Weedbrook. Quantum generative adversarial learning. *Physical Review Letters*, 121(4):040502, 2018.

Aaron Lou, Minkai Xu, Adam Farris, and Stefano Ermon. Scaling riemannian diffusion models. *Advances in Neural Information Processing Systems*, 36:80291–80305, 2023.

Emile Mathieu and Maximilian Nickel. Riemannian continuous normalizing flows. *Advances in neural information processing systems*, 33:2503–2515, 2020.

Bogdan Mielnik. Geometry of quantum states. *Communications in Mathematical Physics*, 9(1): 55–80, 1968.

Benjamin Kurt Miller, Ricky T. Q. Chen, Anuroop Sriram, and Brandon M. Wood. Flowmm: Generating materials with riemannian flow matching. In *International Conference on Machine Learning (ICML)*, Proceedings of Machine Learning Research, 2024.

Mao Nishino, Martin Bauer, Tom Needham, and Nicolas Charon. A benamou-brenier proximal splitting method for constrained unbalanced optimal transport. *arXiv preprint arXiv:2512.09250*, 2025.

Koichi Ogiue. Complex hypersurfaces of a complex projective space. *Journal of Differential Geometry*, 3(1-2):253–256, 1969.

Shivaramakrishnan Pancharatnam. Generalized theory of interference, and its applications: Part i. coherent pencils. In *Proceedings of the Indian academy of sciences-section A*, volume 44, pp. 247–262. Springer, 1956.

John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

Maria Schuld and Nathan Killoran. Quantum machine learning in feature hilbert spaces. *Physical Review Letters*, 122(4):040504, 2019.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)*, 2021a.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021b. URL https://openreview.net/forum?id=PxTIG12RRHS.

Rong Tang and Yun Yang. Adaptivity of diffusion models to manifold structures. In *International Conference on Artificial Intelligence and Statistics*, pp. 1648–1656. PMLR, 2024.

Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research (TMLR)*, 2024.

Roumen Tsekov. Bohmian mechanics versus madelung quantum hydrodynamics. *arXiv preprint arXiv:0904.0723*, 2009.

Cédric Villani. *Optimal Transport: Old and New*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2009.

Yujin Wu and Mohamed Daoudi. Metric learning on complex projective spaces. In *International Conference on Smart Multimedia*, pp. 116–127. Springer, 2022.

Robert E Wyatt. *Quantum dynamics with trajectories: introduction to quantum hydrodynamics*. Springer, 2005.

Jian Xu, Wei Chen, Shigui Li, Chao Li, Delu Zeng, John Paisley, and Qibin Zhao. Local-time riemannian score matching on the quantum pure-state manifold, 2026. arXiv preprint arXiv:2605.03573.

Bingzhi Zhang, Peng Xu, Xiaohui Chen, and Quntao Zhuang. Generative quantum machine learning via denoising diffusion probabilistic models. *Physical Review Letters*, 132(10):100602, 2024.

A Additional Related Work

Quantum generative modeling and ambient baselines. Generative modeling for quantum data has attracted increasing attention in recent years, with approaches ranging from quantum circuit Born machines and variational quantum generators to classical deep generative models for quantum states and observables (Lloyd & Weedbrook, 2018; Dallaire-Demers & Killoran, 2018; Liu & Wang, 2018; Coyle et al., 2020; Benedetti et al., 2019; Amin et al., 2018). A common alternative is to embed quantum states into ambient real or complex vector spaces and apply standard Euclidean generative methods, including diffusion, direct neural generators, or latent-variable models. These methods are simple and often effective on easy benchmarks, but they ignore the quotient geometry of pure states and may allocate capacity to extrinsic directions that have no physical meaning on $\mathbb{CP}^{d-1}$. Our Euclidean diffusion and Euclidean flow baselines are representative of this ambient modeling perspective. By contrast, intrinsic approaches define transport directly on $\mathbb{CP}^{d-1}$, where velocity fields live in tangent spaces and trajectories respect the manifold geometry by construction. This ambient-versus-intrinsic distinction is central to our empirical comparison.

Flow matching for other quantum representations. QFM and ShadowFM are related in motivation but not direct experimental baselines for IFM. QFM studies flow matching for density-matrix interpolation and circuit-based generation (Cui et al., 2025), whereas IFM targets distributions of pure-state rays on $\mathbb{CP}^{d-1}$. Faithfully comparing these settings would require matching not only samples and metrics, but also the density-matrix objective, circuit ansatz, measurement model, and training protocol. ShadowFM instead performs flow matching on classical shadows (Hahm & Clark, 2026), which are classical measurement-derived representations of quantum states rather than points on the pure-state projective manifold. We therefore cite these methods as complementary related work and focus our controlled baselines on Euclidean FM, SSDM, Euclidean VP-SDE, QDDPM, QGAN, and the appendix generic geodesic FM comparator.

Optimal transport and geometric transport. Our formulation is also related to probability transport and optimal transport on structured spaces. Recent work on Optimal Flow Matching shows in Euclidean space that, under suitable parameterizations and couplings, a single flow-matching objective can recover straight optimal-transport displacements (Tong et al., 2024; Kerrigan et al., 2024). Although IFM does not solve a full optimal transport problem, its Monge–geodesic specialization can be viewed as a projective-manifold counterpart of displacement interpolation. This connects our framework conceptually to displacement interpolation, continuity-equation-based density evolution, and action-minimizing flows (Cisneros-Velarde & Bullo, 2023; Chewi et al., 2021; Nishino et al., 2025).

Geometric phase and Pancharatnam alignment. Because pure states are defined only up to global phase, geometric phase conventions are relevant even before one specifies a learning objective (Pancharatnam, 1956; Garza-Soto et al., 2023; Ferrer-Garcia et al., 2023; Leinonen et al., 2023; Jisha et al., 2023). In particular, the Pancharatnam in-phase condition provides a canonical way to compare two pure states by choosing representatives whose overlap is real and nonnegative (Pancharatnam, 1956). Our phase-aligned path construction uses this convention as a gauge-fixing device that removes arbitrary phase mismatch before interpolation. This is the main point at which IFM uses quantum-specific structure beyond generic manifold flow matching.

Relation to Riemannian flow matching on general geometries. Riemannian Flow Matching (RFM) provides a broad framework for flow matching on manifolds through user-specified premetrics and is simulation-free on simple geometries with closed-form geodesics (Chen & Lipman, 2024). Our method shares this high-level conditional-velocity-matching philosophy, but the quantum pure-state setting introduces geometric structure that is not present in the standard simple-geometries catalog used by RFM. The manifold is not a round sphere: $\mathbb{CP}^{d-1}$ is the quotient $S^{2d-1}/U(1)$ equipped with the Fubini–Study metric, with a nontrivial Hopf fibration, vertical gauge directions, and horizontal representatives (Mielnik, 1968; Bengtsson & Życzkowski, 2017). Consequently, comparing two endpoints requires fixing the relative phase before defining a chord or geodesic target; Pancharatnam alignment performs exactly this quotient-specific gauge fixing (Pancharatnam, 1956). This step has no direct analogue on ordinary spheres, flat tori, hyperbolic spaces, meshes, or SPD manifolds.

The experimental regimes are also complementary. RFM reports manifold experiments on settings such as S^2 earth and climate data, low-dimensional protein tori, high-dimensional flat tori with intrinsic dimension on the order of hundreds, mesh surfaces, and SPD matrices up to the 59×59 case with intrinsic dimension 1770 (Chen & Lipman, 2024). By contrast, a dense n-qubit pure-state ensemble lives on $\mathbb{CP}^{2^n-1}$ with real dimension $2(2^n - 1)$; the 12-qubit experiments in this paper therefore correspond to real intrinsic dimension 8190, and the 14-qubit stress tests in Appendix D.8 reach dimension 32766. Thus, IFM is not merely an application of RFM to another named manifold: it develops the phase-aligned endpoint selection, horizontal velocity parameterization, and projective transport proofs needed for high-dimensional quantum pure-state ensembles. The generic geodesic FM baseline in Appendix D.13 makes this distinction empirical by isolating the contribution of intrinsic geometry from the additional quantum-specific ingredients.

Physics-inspired transport viewpoints. Beyond machine-learning transport literature, our work is also physically inspired by quantum current and transport viewpoints in which ensemble evolution is described through continuity equations and associated flow fields (Delgado, 1999; Hodge et al., 2014). This includes, at a broad conceptual level, hydrodynamic, current-based, and Bohmian-style interpretations of quantum evolution, where dynamics are organized around probability transport rather than solely around noise or static wavefunction representations (Wyatt, 2005; Kuz'menkov & Maksimov, 1999; Tsekov, 2009). In this sense, the IFM velocity field can be viewed as an ensemble probability current on quantum state space. We do not derive a new microscopic quantum law; rather, we use this perspective as motivation for an ensemble-level generative model on $\mathbb{CP}^{d-1}$.

B Additional Method Discussion and Algorithm

B.1 Intrinsic versus Euclidean flow matching

A major question in our study is whether the gains of IFM come simply from replacing diffusion with flow matching, or specifically from respecting the intrinsic geometry of quantum pure states. To isolate these two effects, we compare IFM against a Euclidean flow matching baseline that performs transport in an ambient representation of the state vector and only enforces quantum-state validity through post hoc normalization, reflecting a common ambient modeling strategy for quantum-state generators (Lloyd & Weedbrook, 2018; Dallaire-Demers & Killoran, 2018; Liu & Wang, 2018; Amin et al., 2018). The Euclidean baseline therefore learns an *extrinsic* transport field, whereas IFM learns an *intrinsic* manifold transport field.

This comparison allows us to separate two distinct inductive biases:

$$\text{diffusion vs. flow} \qquad \text{and} \qquad \text{ambient transport vs. intrinsic transport.} \tag{30}$$

Our experiments show that both matter, but that intrinsic manifold transport is particularly important for faithfully modeling quantum pure-state ensembles.

B.2 Why intrinsic flow matching scales better

A central motivation of IFM is scalability. We argue that the difficulty of high-qubit quantum generative modeling is not only due to the large dimension of the state space, but also due to the algorithmic burden imposed by stochastic diffusion. Diffusion-based models on $\mathbb{CP}^{d-1}$ must estimate time-dependent score fields, rely on local diffusion approximations or teacher constructions, and generate samples through reverse-time stochastic simulation (Song et al., 2021b; Xu et al., 2026). Each of these steps introduces approximation error and optimization difficulty, which tend to accumulate as the dimension grows.

In contrast, IFM directly learns a transport velocity field. This leads to three advantages. First, the supervision signal is more direct: instead of learning local density gradients across noise scales, the model learns the direction in which probability mass should move. Second, sampling is simpler: generation is performed by integrating a deterministic flow rather than reversing a stochastic process. Third, the learned vector field is intrinsic to the tangent bundle of $\mathbb{CP}^{d-1}$, avoiding unnecessary modeling effort in ambient directions that do not correspond to physical transport on the pure-state manifold.

Taken together, these properties make IFM a shorter and more stable generative pipeline. Empirically, we find that this leads to improved robustness and stronger performance as the number of qubits increases, especially when compared against stochastic diffusion baselines and ambient Euclidean flow baselines.

Importantly, the intrinsic machinery itself is nearly free relative to ambient flow matching: Table 4 reports per-step training time and peak GPU memory for IFM and Euclidean FM across $n = 2, \ldots, 12$ qubits. The phase alignment, tangent projection, and retraction add only a roughly constant ~ 1 ms per step and essentially no additional memory; the cost is dominated by the unavoidable $O(2^n)$ dense state-vector dimension, which is shared by *every* dense generative method on $\mathbb{CP}^{d-1}$. Thus IFM does not aggravate the exponential-dimension bottleneck; it removes the score-estimation, teacher-construction, and reverse-time simulation overhead of diffusion baselines while respecting the manifold, and is complementary to on-hardware circuit-based generators rather than a replacement for them.

B.3 Relation to SSDM and Euclidean flow matching

Our work is closely related to two existing generative perspectives: stochastic diffusion on quantum pure-state manifolds and Euclidean flow matching in ambient space.

Relation to SSDM. SSDM formulates quantum-state generation as a diffusion process on $\mathbb{CP}^{d-1}$ and learns a score field associated with the resulting density evolution (Xu et al., 2026). Our method shares the same geometric setting, namely the pure-state manifold, but replaces score learning and

Table 4: Cost of the intrinsic machinery: per-step training time (ms) and peak GPU memory (MB) for IFM versus Euclidean FM, matched network scale, on an A6000. The intrinsic pieces add a near-constant overhead; memory and time scale identically with the $O(2^n)$ ambient dimension.

n	$d = 2^n$	IFM step (ms)	Euclid step (ms)	IFM mem (MB)	Euclid mem (MB)
2	4	4.33	3.57	39.4	39.3
6	64	4.56	3.43	41.9	41.8
10	1024	4.49	3.50	82.3	81.8
12	4096	4.42	3.47	221.4	211.6

reverse-time denoising by deterministic velocity matching and manifold flow integration. In this sense, IFM replaces

$$\text{diffusion + score matching + reverse SDE} \tag{31}$$

with

$$\text{intrinsic transport + velocity matching + ODE flow.} \tag{32}$$

This change removes the need for local diffusion teachers and reverse-time stochastic simulation, while preserving the manifold-aware treatment of quantum pure states.

Relation to Euclidean flow matching. A second relevant baseline is Euclidean flow matching, which represents a quantum state as an ambient vector in $\mathbb{C}^d$ (or $\mathbb{R}^{2d}$), learns a transport field in that ambient space, and enforces state validity only through normalization or projection (Lipman et al., 2023). This approach replaces diffusion with flow, but does not respect the intrinsic geometry of $\mathbb{CP}^{d-1}$.

IFM differs from Euclidean flow matching in that both its intermediate states and its velocity fields are defined intrinsically on the pure-state manifold. This distinction is especially important in high dimensions, where ambient transport may allocate modeling capacity to extrinsic directions that are irrelevant to physical state evolution. By comparing IFM with Euclidean flow matching, we isolate the effect of geometry from the effect of switching from diffusion to flow.

Positioning. Overall, our framework should be viewed as the deterministic flow-matching counterpart to SSDM on the same pure-state manifold: both are intrinsic, but they learn different dynamical objects and sample through different mechanisms.

B.4 Neural parameterization and tangent projection

Sections 4.2 and 4.4 identify the population target learned by IFM: a tangent velocity field on $\mathbb{CP}^{d-1}$ induced by the phase-aligned conditional path family. The remaining question is architectural: how should a neural network represent such a field without reintroducing the ambient gauge redundancy that the previous sections worked to remove?

In practice, we represent a pure state $\psi \in \mathbb{C}^d$ through its real and imaginary parts and feed them, together with a time embedding of t, into a neural network. The network first produces an ambient vector in $\mathbb{C}^d$, which is then projected onto the tangent space at ψ. For the complex projective manifold, a valid projective velocity must be horizontal with respect to the quotient representation (Mielnik, 1968; Bengtsson & Życzkowski, 2017), so we use the projection

$$\Pi_\psi(v) = v - \psi\langle\psi, v\rangle, \tag{33}$$

which satisfies $\langle\psi, \Pi_\psi(v)\rangle = 0$. We therefore parameterize the learned field as

$$v_\theta(\psi, t) := \Pi_\psi\big(f_\theta(\psi, t)\big), \tag{34}$$

where f_θ is the unconstrained network output in the ambient space. This is the implementation counterpart of the intrinsic theory: although the network computes in ambient coordinates, the object that enters the loss is always a tangent vector on projective state space. As a result, the training problem remains aligned with the conditional target $u_t(\psi_t \mid \psi_0, \psi_1)$ from Section 4.2 and the marginal transport field characterized in Section 4.4.

Table 5: Global-phase gauge equivariance on 6-qubit single-cluster. Left: relative equivariance error $\|v(e^{i\varphi}\psi) - e^{i\varphi}v(\psi)\|/\|v(\psi)\|$ averaged over random ψ, φ, t. Right: physical metrics (lower is better) under the canonical input gauge versus random per-sample global phases on the base states.

Model / input gauge	MMD ↓	Δ_{obs} ↓	Ent. W_1 ↓
raw field, equiv. error		1.28	
canonical-gauge field, equiv. error		1.9×10^{-7}	
raw, canonical base	9.8×10^{-4}	4.2×10^{-3}	9.0×10^{-3}
raw, phase-randomized base	5.7×10^{-3}	7.6×10^{-2}	1.6×10^{-1}
canonical-gauge, canonical base	7.8×10^{-4}	4.2×10^{-3}	9.8×10^{-3}
canonical-gauge, phase-randomized base	8.9×10^{-4}	3.9×10^{-3}	9.4×10^{-3}

Proposition 6 (Horizontal representative as the unique minimal-norm projective velocity)**.** *Fix a normalized representative $\psi \in \mathbb{C}^d$ and an ambient vector $w \in \mathbb{C}^d$. Consider the affine family*

$$\mathcal{E}_\psi(w) := \{w + a\psi + b\, i\psi : a, b \in \mathbb{R}\}, \tag{35}$$

which consists of all ambient first-order perturbations that induce the same projective first-order motion as w. Then $\Pi_\psi(w)$ is the unique element of $\mathcal{E}_\psi(w)$ lying in $\mathcal{H}_\psi$, and it is the unique minimizer of the ambient norm over $\mathcal{E}_\psi(w)$:

$$\Pi_\psi(w) = \arg \min_{\xi \in \mathcal{E}_\psi(w)} \|\xi\|_2. \tag{36}$$

Proof in Appendix C.12.

Proposition 6 shows that the horizontal projection is the unique minimal-norm representative among ambient perturbations encoding the same projective motion, removing the redundant radial and pure-phase directions analytically rather than leaving them to the model.

B.5 Global-phase gauge equivariance

The horizontal projection Π_ψ guarantees that $v_\theta(\psi, t)$ is *tangent*, but because the network reads the raw ambient coordinates $(\operatorname{Re}\psi, \operatorname{Im}\psi)$ it is not automatically *equivariant* under the global-phase action, i.e. it need not satisfy $v_\theta(e^{i\varphi}\psi, t) = e^{i\varphi}v_\theta(\psi, t)$. Two observations make this precise and provide a cost-free fix.

First, in all our experiments the base sampler draws representatives with a fixed gauge (a normalized cluster representative $\psi \propto e_0 + \text{noise}$), so training and sampling operate at a consistent gauge and the reported metrics are well defined on $\mathbb{CP}^{d-1}$. Second, one can make the field *exactly* equivariant by fixing a canonical gauge before the network call: let $j(\psi) = \arg\max_k |\psi_k|$ and $\beta(\psi) = \arg \psi_{j(\psi)}$, define the canonical representative $\operatorname{can}(\psi) = e^{-i\beta(\psi)}\psi$ (whose largest-magnitude amplitude is real and nonnegative), and set

$$v_\theta^{\text{eq}}(\psi, t) := \Pi_\psi\Big(e^{i\beta(\psi)}\, v_\theta\big(\operatorname{can}(\psi), t\big)\Big). \tag{37}$$

Since $\operatorname{can}(e^{i\varphi}\psi) = \operatorname{can}(\psi)$ and $\beta(e^{i\varphi}\psi) = \beta(\psi) + \varphi$, this field satisfies $v_\theta^{\text{eq}}(e^{i\varphi}\psi, t) = e^{i\varphi}v_\theta^{\text{eq}}(\psi, t)$ exactly.

Table 5 verifies both points empirically on the 6-qubit single-cluster benchmark. The raw field has a large relative equivariance error, which the canonical-gauge convention drives to numerical zero; end-to-end, the canonical-gauge model matches the reported metrics and is invariant to random global-phase perturbations of the inputs, whereas the raw field degrades under such perturbations.

B.6 Sampling by deterministic manifold flow

After training, generation samples $\psi_0 \sim p_0$ and integrates the learned tangent ODE, as in flow-matching and continuous-normalizing-flow samplers (Chen et al., 2018; Lipman et al., 2023),

$$\frac{d\psi_t}{dt} = v_\theta(\psi_t, t), \qquad v_\theta(\psi_t, t) \in \mathcal{H}_{\psi_t}. \tag{38}$$

Numerically, we integrate in the ambient representative and retract back to the manifold after each step. For $\xi \in \mathcal{H}_\psi$, we use the normalization-based retraction

$$\mathrm{Retr}_\psi(\xi) := \frac{\psi + \xi}{\|\psi + \xi\|_2}, \tag{39}$$

so an Euler step is

$$\psi_{t_{k+1}} = \mathrm{Retr}_{\psi_{t_k}}\big(\Delta t\, v_\theta(\psi_{t_k}, t_k)\big). \tag{40}$$

This deterministic sampler directly transports base samples through the learned probability flow, rather than reversing a stochastic diffusion process (Song et al., 2021b; Ho et al., 2020; Xu et al., 2026).

Proposition 7 (First-order consistency of the retraction sampler)**.** *Assume $v_t(\psi)$ is continuously differentiable in (ψ, t) and tangent to $\mathbb{CP}^{d-1}$. Consider the intrinsic ODE*

$$\dot{\psi}_t = v_t(\psi_t) \tag{41}$$

and the discrete update

$$\psi_{t_{k+1}} = \mathrm{Retr}_{\psi_{t_k}}\big(\Delta t\, v_{t_k}(\psi_{t_k})\big). \tag{42}$$

Then the normalization-based retraction $\mathrm{Retr}_\psi(\xi) = (\psi+\xi)/\|\psi+\xi\|_2$ is a valid first-order retraction on the unit-sphere representation, and the resulting Euler–retraction scheme is a first-order consistent integrator of the manifold flow. Proof in Appendix C.13.

Proposition 7 connects the implementation to the continuous transport theory of Section 4.4: the normalization step is a first-order consistent discretization of the intrinsic manifold ODE.

B.7 Algorithm

Algorithm 1 summarizes the training and sampling procedures of IFM.

In practice, $\mathrm{Retr}_\psi(\cdot)$ denotes a simple normalization-based retraction used to control numerical drift and keep the trajectory on the pure-state manifold. This makes IFM straightforward to implement while preserving the geometric constraint of quantum pure states.

C Additional Proofs for the IFM Theory

This appendix collects fuller proofs for the theoretical statements in the main text. Throughout, we work with unit-norm representatives on the Hilbert sphere and identify tangent vectors on $\mathbb{CP}^{d-1}$ with horizontal vectors (Mielnik, 1968; Bengtsson & Życzkowski, 2017) $\mathcal{H}_\psi = \{\xi \in \mathbb{C}^d : \langle \psi, \xi \rangle = 0\}$. Unless otherwise stated, all densities are with respect to the Fubini–Study volume form, and all vector fields are assumed measurable in space and continuous in time.

C.1 Phase alignment and normalized chord paths

This subsection expands the phase-alignment step used in Section 4.2, following the Pancharatnam in-phase convention for comparing pure-state representatives (Pancharatnam, 1956; Garza-Soto et al., 2023). The goal is to explain why the choice $\bar{\psi}_1 = \alpha(\psi_0, \psi_1)^{-1}\psi_1$ is the canonical representative for interpolation and why it leads to the shortest ambient chord among all representatives of the same projective endpoint.

Proposition 8 (Pancharatnam-aligned representative)**.** *Let $\psi_0, \psi_1 \in \mathbb{C}^d$ be unit vectors with $\langle \psi_0, \psi_1 \rangle \neq 0$. Then $\bar{\psi}_1 = \alpha(\psi_0, \psi_1)^{-1}\psi_1$ is the unique representative of $[\psi_1]$ satisfying the Pancharatnam in-phase condition with respect to ψ_0,*

$$\langle \psi_0, \bar{\psi}_1 \rangle = |\langle \psi_0, \psi_1 \rangle| \in \mathbb{R}_{\geq 0}, \tag{43}$$

and equivalently it is the unique representative that maximizes the real overlap with ψ_0 and minimizes the ambient chord length $\|\psi_0 - \bar{\psi}_1\|_2$ among all phase-equivalent representatives of $[\psi_1]$.

Algorithm 1 Intrinsic Flow Matching on $\mathbb{CP}^{d-1}$

Initialize velocity network $v_\theta(\psi, t)$
for each training iteration **do**
  Sample base states $\psi_0 \sim p_0$
  Sample target states $\psi_1 \sim p_1$
  Sample times $t \sim \mathrm{Unif}[0, 1]$
  Phase-align ψ_1 to $\bar{\psi}_1$
  Construct conditional path states

$$\tilde{\psi}_t = (1-t)\psi_0 + t\bar{\psi}_1, \qquad \psi_t = \tilde{\psi}_t / \|\tilde{\psi}_t\|_2$$

  Compute target tangent velocities

$$u_t(\psi_t \mid \psi_0, \psi_1) = \frac{1}{\|\tilde{\psi}_t\|_2} \Pi_{\psi_t}(\bar{\psi}_1 - \psi_0)$$

  Predict manifold velocity $v_\theta(\psi_t, t)$
  Minimize

$$\mathcal{L}_{\mathrm{IFM}} = \mathbb{E}\big[\|v_\theta(\psi_t, t) - u_t(\psi_t \mid \psi_0, \psi_1)\|^2_{\psi_t}\big]$$

end for
Sampling:
Sample $\psi_0 \sim p_0$
for $k = 0, \ldots, K-1$ **do**
  Evaluate tangent velocity $v_\theta(\psi_{t_k}, t_k)$
  Update state by ODE step

$$\psi_{t_{k+1}} \leftarrow \mathrm{Retr}_{\psi_{t_k}}\big(\Delta t\, v_\theta(\psi_{t_k}, t_k)\big)$$

end for
Return ψ_1

Proposition 9 (Well-defined conditional manifold path)**.** *Let $\psi_0, \psi_1 \in \mathbb{C}^d$ be unit-norm representatives of points in $\mathbb{CP}^{d-1}$ with generic nonzero overlap $\langle \psi_0, \psi_1 \rangle \neq 0$, and let $\gamma_t(\psi_0, \psi_1)$ be defined as in Section 4.2 by phase alignment followed by normalized chord interpolation. The orthogonal case $\langle \psi_0, \psi_1 \rangle = 0$ forms a measure-zero set under continuous p_0, p_1 and is handled by the fixed deterministic convention $\bar{\psi}_1 = \psi_1$, for which the representative-independence claim below holds up to that fixed choice. Then:*

1. *the projective class $[\gamma_t(\psi_0, \psi_1)]$ is independent of the chosen representatives of $[\psi_0]$ and $[\psi_1]$;*
2. *for every t such that $\tilde{\psi}_t \neq 0$, the instantaneous velocity $u_t(\psi_t \mid \psi_0, \psi_1)$ belongs to the horizontal tangent space $\mathcal{H}_{\psi_t}$;*
3. *the path satisfies the endpoint conditions $[\gamma_0(\psi_0, \psi_1)] = [\psi_0]$ and $[\gamma_1(\psi_0, \psi_1)] = [\psi_1]$.*

Lemma 1 (Phase alignment maximizes real overlap)**.** *Let $\psi_0, \psi_1 \in \mathbb{C}^d$ be unit vectors with $\langle \psi_0, \psi_1 \rangle \neq 0$. For any phase $\varphi \in \mathbb{R}$, define $\psi_1^{(\varphi)} := e^{i\varphi}\psi_1$. Then*

$$\max_{\varphi \in \mathbb{R}} \mathrm{Re}\, \langle \psi_0, \psi_1^{(\varphi)} \rangle = |\langle \psi_0, \psi_1 \rangle|, \tag{44}$$

and the maximizers are exactly the phases for which $\langle \psi_0, \psi_1^{(\varphi)} \rangle$ is real and nonnegative. In particular, the canonical aligned representative $\bar{\psi}_1 = \alpha(\psi_0, \psi_1)^{-1}\psi_1$ satisfies $\langle \psi_0, \bar{\psi}_1 \rangle = |\langle \psi_0, \psi_1 \rangle|$.

Proof. Write $\langle \psi_0, \psi_1 \rangle = re^{i\beta}$ with $r = |\langle \psi_0, \psi_1 \rangle| > 0$. Then

$$\langle \psi_0, e^{i\varphi}\psi_1 \rangle = re^{i(\beta+\varphi)},$$

so

$$\operatorname{Re}\langle \psi_0, e^{i\varphi}\psi_1\rangle = r\cos(\beta+\varphi).$$

This quantity is maximized when $\beta + \varphi \in 2\pi\mathbb{Z}$, in which case its value is r. Equivalently, the maximizing representatives are exactly those for which the overlap is real and nonnegative. Taking $\varphi^\star = -\beta$ gives

$$e^{i\varphi^\star}\psi_1 = \frac{\overline{\langle \psi_0, \psi_1\rangle}}{|\langle \psi_0, \psi_1\rangle|}\psi_1 = \alpha(\psi_0,\psi_1)^{-1}\psi_1 = \bar{\psi}_1,$$

which proves the claim. □

Lemma 2 (Phase alignment minimizes ambient chord length)**.** *Under the assumptions of Lemma 1,*

$$\min_{\varphi\in\mathbb{R}} \|\psi_0 - e^{i\varphi}\psi_1\|_2^2 = 2 - 2|\langle \psi_0, \psi_1\rangle|, \tag{45}$$

and the minimizers are exactly the phases that maximize $\operatorname{Re}\langle \psi_0, e^{i\varphi}\psi_1\rangle$*. Hence* $\bar{\psi}_1$ *is the representative of* $[\psi_1]$ *closest to* ψ_0 *on the unit sphere.*

Proof. Since $\|\psi_0\|_2 = \|\psi_1\|_2 = 1$, we have

$$\|\psi_0 - e^{i\varphi}\psi_1\|_2^2 = \langle \psi_0 - e^{i\varphi}\psi_1, \psi_0 - e^{i\varphi}\psi_1\rangle \tag{46}$$

$$= \|\psi_0\|_2^2 + \|\psi_1\|_2^2 - 2\operatorname{Re}\langle \psi_0, e^{i\varphi}\psi_1\rangle \tag{47}$$

$$= 2 - 2\operatorname{Re}\langle \psi_0, e^{i\varphi}\psi_1\rangle. \tag{48}$$

Therefore minimizing the chord length is equivalent to maximizing the real part of the overlap. Lemma 1 gives the optimal value and shows that $\bar{\psi}_1$ is a minimizer. □

C.2 Interpolation paths used in the experiments

This subsection makes the experimental path choices explicit. The same flow-matching regression objective can be paired with different conditional path families (Lipman et al., 2023; Tong et al., 2024; Chen & Lipman, 2024), and the paper uses three such variants:

- the phase-aligned normalized-chord path used by the main IFM model;
- a phase-aligned geodesic variant used as an ablation of our method;
- an appendix-only generic manifold-geodesic baseline that omits phase alignment and therefore excludes our quotient-specific representative selection.

Phase-aligned normalized-chord path. Given unit representatives $\psi_0, \psi_1 \in \mathbb{C}^d$, define the aligned endpoint

$$\bar{\psi}_1 = \begin{cases} \alpha(\psi_0,\psi_1)^{-1}\psi_1, & \langle \psi_0, \psi_1\rangle \neq 0, \\ \psi_1, & \langle \psi_0, \psi_1\rangle = 0, \end{cases} \qquad \alpha(\psi_0,\psi_1) = \frac{\langle \psi_0, \psi_1\rangle}{|\langle \psi_0, \psi_1\rangle|}.$$

The ambient chord is

$$\tilde{\psi}_t = (1-t)\psi_0 + t\bar{\psi}_1,$$

and the path used in the main method is its normalized version

$$\psi_t = \frac{\tilde{\psi}_t}{\|\tilde{\psi}_t\|_2}.$$

Differentiating the normalization map yields

$$\dot{\psi}_t = \frac{1}{\|\tilde{\psi}_t\|_2}\Big(I - \psi_t\psi_t^*\Big)(\bar{\psi}_1 - \psi_0),$$

which is the orthogonal projection of the ambient chord direction onto the tangent space of the unit sphere at ψ_t. In the implementation, the final target is represented by its horizontal component $\Pi_{\psi_t}(\xi) = \xi - \langle \psi_t, \xi\rangle\psi_t$, so that the learning target lies in $\mathcal{H}_{\psi_t}$ and therefore defines a projective tangent vector.

Phase-aligned geodesic interpolation (ours variant). Our geodesic ablation uses the same aligned endpoint $\bar{\psi}_1$ but replaces the normalized chord by the short spherical interpolation between ψ_0 and $\bar{\psi}_1$. Let

$$\theta = \arccos\big(\mathrm{clip}(\mathrm{Re}\,\langle \psi_0, \bar{\psi}_1\rangle, -1, 1)\big).$$

For nondegenerate angles, the path is

$$\psi_t = \frac{\sin((1-t)\theta)}{\sin\theta}\,\psi_0 + \frac{\sin(t\theta)}{\sin\theta}\,\bar{\psi}_1,$$

with derivative

$$\dot{\psi}_t = -\theta\frac{\cos((1-t)\theta)}{\sin\theta}\,\psi_0 + \theta\frac{\cos(t\theta)}{\sin\theta}\,\bar{\psi}_1.$$

For numerically small θ, the implementation uses the corresponding linearized limit. This is treated as an ablation of our method rather than as an unrelated baseline because it uses the same Pancharatnam-aligned representative selection and the same horizontal tangent representation (Pancharatnam, 1956; Bengtsson & Życzkowski, 2017); only the interpolation rule is changed.

Generic manifold-geodesic baseline without phase alignment. Appendix D.13 also reports a closest in-code analogue of a generic manifold flow-matching baseline on $\mathbb{CP}^{d-1}$ (Chen & Lipman, 2024; Miller et al., 2024). In that variant, the raw representatives ψ_0, ψ_1 are used directly, without first replacing ψ_1 by the aligned representative $\bar{\psi}_1$. Define

$$\theta_{\mathrm{raw}} = \arccos\big(\mathrm{clip}(\mathrm{Re}\,\langle \psi_0, \psi_1\rangle, -1, 1)\big),$$

and set

$$\psi_t^{\mathrm{raw}} = \frac{\sin((1-t)\theta_{\mathrm{raw}})}{\sin\theta_{\mathrm{raw}}}\,\psi_0 + \frac{\sin(t\theta_{\mathrm{raw}})}{\sin\theta_{\mathrm{raw}}}\,\psi_1,$$

again with the small-angle linear limit when needed. This path is geometrically valid on the sphere representation, but because it omits phase alignment it does not quotient out arbitrary endpoint phase mismatch before interpolation. That is precisely why we use it only as an appendix baseline and not as the main IFM construction.

C.3 Proof of Proposition 8

Proof. The first claim is exactly the Pancharatnam in-phase condition. By Lemma 1, among all phase-equivalent representatives $\psi_1^{(\varphi)} = e^{i\varphi}\psi_1$, the representative $\bar{\psi}_1 = \alpha(\psi_0, \psi_1)^{-1}\psi_1$ is characterized by

$$\langle \psi_0, \bar{\psi}_1\rangle = |\langle \psi_0, \psi_1\rangle| \in \mathbb{R}_{\geq 0},$$

and is the unique phase choice with this property when $\langle \psi_0, \psi_1\rangle \neq 0$. This proves the Pancharatnam-alignment statement.

Again by Lemma 1, the same representative uniquely maximizes $\mathrm{Re}\,\langle \psi_0, e^{i\varphi}\psi_1\rangle$ over all φ. Lemma 2 shows that minimizing $\|\psi_0 - e^{i\varphi}\psi_1\|_2$ is equivalent to maximizing this same real overlap. Hence $\bar{\psi}_1$ is also the unique representative minimizing the ambient chord length among all representatives of the projective endpoint $[\psi_1]$.

Therefore the three characterizations are equivalent:

1. satisfying the Pancharatnam in-phase condition with respect to ψ_0;
2. maximizing the real overlap with ψ_0;
3. minimizing the ambient chord length to ψ_0.

This proves Proposition 8. □

Remark 6. *The point of Lemmas 1 and 2 is not that projective geometry is reduced to ambient Euclidean geometry. Rather, once one chooses representatives in the quotient space $\mathbb{CP}^{d-1}$, there remains a gauge degree of freedom. Phase alignment fixes that freedom in the way that produces the shortest ambient interpolation consistent with the same projective endpoints, thereby removing spurious motion due solely to global phase.*

C.4 Proof of Proposition 9

Proof. Let $\psi_0' = e^{i\theta_0}\psi_0$ and $\psi_1' = e^{i\theta_1}\psi_1$ be alternative representatives of the same projective points. Then

$$\alpha(\psi_0', \psi_1') = \frac{\langle e^{i\theta_0}\psi_0, e^{i\theta_1}\psi_1\rangle}{|\langle e^{i\theta_0}\psi_0, e^{i\theta_1}\psi_1\rangle|} = e^{i(\theta_1-\theta_0)}\alpha(\psi_0, \psi_1).$$

Hence

$$\bar{\psi}_1' = \alpha(\psi_0', \psi_1')^{-1}\psi_1' = e^{i\theta_0}\alpha(\psi_0, \psi_1)^{-1}\psi_1 = e^{i\theta_0}\bar{\psi}_1.$$

The interpolant therefore transforms as

$$\tilde{\psi}_t' = (1-t)\psi_0' + t\bar{\psi}_1' = e^{i\theta_0}\big((1-t)\psi_0 + t\bar{\psi}_1\big) = e^{i\theta_0}\tilde{\psi}_t.$$

After normalization, $\gamma_t(\psi_0', \psi_1') = e^{i\theta_0}\gamma_t(\psi_0, \psi_1)$, so both define the same point in $\mathbb{CP}^{d-1}$. This proves representative-independence.

The endpoint property follows immediately:

$$\gamma_0(\psi_0, \psi_1) = \psi_0, \qquad \gamma_1(\psi_0, \psi_1) = \bar{\psi}_1,$$

and $[\bar{\psi}_1] = [\psi_1]$ because they differ by a unit-modulus phase.

For tangency, define $w_t := \bar{\psi}_1 - \psi_0$ so that $\dot{\tilde{\psi}}_t = w_t$. Differentiating the normalized lift $\psi_t = \tilde{\psi}_t/\|\tilde{\psi}_t\|_2$ gives

$$\dot{\psi}_t = \frac{w_t}{\|\tilde{\psi}_t\|_2} - \frac{\tilde{\psi}_t\,\mathrm{Re}\langle\tilde{\psi}_t, w_t\rangle}{\|\tilde{\psi}_t\|_2^3}.$$

In particular, $\dot{\psi}_t$ is tangent to the unit sphere, i.e., $\mathrm{Re}\langle\psi_t, \dot{\psi}_t\rangle = 0$. To pass from the sphere to the projective quotient, one removes the residual vertical phase component along $i\psi_t$ and retains the horizontal representative. By definition this representative is

$$u_t(\psi_t \mid \psi_0, \psi_1) := \frac{1}{\|\tilde{\psi}_t\|_2}\Pi_{\psi_t}(w_t).$$

Since

$$\langle\psi_t, \Pi_{\psi_t}(w_t)\rangle = \langle\psi_t, w_t - \psi_t\langle\psi_t, w_t\rangle\rangle = 0,$$

we conclude that $u_t(\psi_t \mid \psi_0, \psi_1) \in \mathcal{H}_{\psi_t}$. □

C.5 Proof of Theorem 1

Proof. Fix $t \in (0,1)$ and abbreviate $U := u_t(\psi_t \mid \psi_0, \psi_1) \in T_{\psi_t}\mathbb{CP}^{d-1}$. For any measurable tangent field $v(\cdot, t)$, the population risk is

$$\mathcal{J}_t[v] = \mathbb{E}\big[\|v(\psi_t, t) - U\|_{\psi_t}^2\big].$$

Condition on the intermediate state ψ_t:

$$\mathcal{J}_t[v] = \mathbb{E}_{\psi_t}\left[\mathbb{E}\big[\|v(\psi_t, t) - U\|_{\psi_t}^2 \,\big|\, \psi_t\big]\right].$$

Thus the minimization decouples pointwise in ψ. Fix ψ and write $m(\psi) := \mathbb{E}[U \mid \psi_t = \psi]$. Using the Hilbert-space identity $\|a-b\|^2 = \|a-c\|^2 + \|b-c\|^2 - 2\langle a-c, b-c\rangle$ with $c = m(\psi)$ and then taking conditional expectation, we obtain

$$\mathbb{E}\big[\|v(\psi, t) - U\|_{\psi}^2 \,\big|\, \psi_t = \psi\big] = \|v(\psi, t) - m(\psi)\|_{\psi}^2 + \mathbb{E}\big[\|U - m(\psi)\|_{\psi}^2 \,\big|\, \psi_t = \psi\big].$$

The second term is independent of v, so the unique minimizer is $v(\psi, t) = m(\psi)$. Since this holds for p_t-almost every ψ, we obtain

$$v_t^\star(\psi) = \mathbb{E}\big[u_t(\psi_t \mid \psi_0, \psi_1) \,\big|\, \psi_t = \psi\big].$$

□

C.6 Proof of Proposition 1

Proof. Let $\bar{u}_t(\psi) := \mathbb{E}[u_t(\psi_t \mid \psi_0, \psi_1) \mid \psi_t = \psi]$. For notational brevity write $U := u_t(\psi_t \mid \psi_0, \psi_1)$ and $\bar{U} := \bar{u}_t(\psi_t)$. Then

$$v(\psi_t, t) - U = \big(v(\psi_t, t) - \bar{U}\big) + \big(\bar{U} - U\big).$$

Expanding the squared norm and taking expectation gives

$$\begin{aligned}\mathcal{J}_t[v] = \mathbb{E}\big[\|v(\psi_t, t) - \bar{U}\|^2_{\psi_t}\big] + \mathbb{E}\big[\|U - \bar{U}\|^2_{\psi_t}\big] \\ + 2\,\mathbb{E}\big[\langle v(\psi_t, t) - \bar{U}, \bar{U} - U\rangle_{\psi_t}\big].\end{aligned} \tag{49}$$

It remains to show that the cross term vanishes. Conditioning on ψ_t yields

$$\mathbb{E}\big[\langle v(\psi_t, t) - \bar{U}, \bar{U} - U\rangle_{\psi_t}\big] = \mathbb{E}\Big[\big\langle v(\psi_t, t) - \bar{U}, \mathbb{E}[\bar{U} - U \mid \psi_t]\big\rangle_{\psi_t}\Big].$$

By definition of $\bar{U}$, $\mathbb{E}[U \mid \psi_t] = \bar{U}$, hence $\mathbb{E}[\bar{U} - U \mid \psi_t] = 0$. Therefore the last term in (49) is zero, and the desired decomposition follows. Since the second term does not depend on v, the conditional and marginal objectives have identical minimizers. □

C.7 Proof of Proposition 2

Proof. Let $\varphi \in C^\infty(\mathbb{CP}^{d-1})$ be a smooth test function. Since $t \mapsto \gamma_t(\psi_0, \psi_1)$ is C^1, the chain rule gives

$$\frac{d}{dt}\varphi(\psi_t) = \langle \nabla_{\mathrm{FS}}\varphi(\psi_t), u_t(\psi_t \mid \psi_0, \psi_1)\rangle_{\psi_t}.$$

Taking expectation yields

$$\frac{d}{dt}\mathbb{E}[\varphi(\psi_t)] = \mathbb{E}\left[\langle \nabla_{\mathrm{FS}}\varphi(\psi_t), u_t(\psi_t \mid \psi_0, \psi_1)\rangle_{\psi_t}\right].$$

Now condition on ψ_t and use Theorem 1:

$$\frac{d}{dt}\mathbb{E}[\varphi(\psi_t)] = \mathbb{E}\left[\langle \nabla_{\mathrm{FS}}\varphi(\psi_t), v_t^\star(\psi_t)\rangle_{\psi_t}\right].$$

Writing the expectation with respect to the density p_t gives

$$\frac{d}{dt}\int_{\mathbb{CP}^{d-1}} \varphi(\psi)\, p_t(\psi)\, d\mathrm{vol}_{\mathrm{FS}}(\psi) = \int_{\mathbb{CP}^{d-1}} \langle \nabla_{\mathrm{FS}}\varphi(\psi), v_t^\star(\psi)\rangle_\psi\, p_t(\psi)\, d\mathrm{vol}_{\mathrm{FS}}(\psi).$$

By the definition of the Riemannian divergence,

$$\int_{\mathbb{CP}^{d-1}} \langle \nabla_{\mathrm{FS}}\varphi, v_t^\star\rangle\, p_t\, d\mathrm{vol}_{\mathrm{FS}} = -\int_{\mathbb{CP}^{d-1}} \varphi\, \mathrm{div}_{\mathbb{CP}}(p_t v_t^\star)\, d\mathrm{vol}_{\mathrm{FS}},$$

which is exactly the weak formulation of $\partial_t p_t + \mathrm{div}_{\mathbb{CP}}(p_t v_t^\star) = 0$. Since this identity holds for every smooth test function φ, the claim follows. □

C.8 Proof of Theorem 2

Proof. By the local Lipschitz assumption on $v_t^\star$ and compactness of $\mathbb{CP}^{d-1}$, the manifold ODE $\dot{\psi}_t = v_t^\star(\psi_t)$ admits a unique global flow $\Phi_{0,t}$ on $[0, 1]$ (Lee, 2012). Define the pushforward $\tilde{p}_t := (\Phi_{0,t})_\# p_0$. For any test function $\varphi \in C^\infty(\mathbb{CP}^{d-1})$, the chain rule along the deterministic flow gives

$$\frac{d}{dt}\int_{\mathbb{CP}^{d-1}} \varphi\, d\tilde{p}_t = \mathbb{E}\big[\langle \nabla_{\mathrm{FS}}\varphi(\Phi_{0,t}\psi_0), v_t^\star(\Phi_{0,t}\psi_0)\rangle_{\Phi_{0,t}\psi_0}\big] = \int_{\mathbb{CP}^{d-1}} \langle \nabla_{\mathrm{FS}}\varphi, v_t^\star\rangle\, d\tilde{p}_t,$$

so $\tilde{p}_t$ satisfies the manifold continuity equation $\partial_t \tilde{p}_t + \mathrm{div}_{\mathbb{CP}}(\tilde{p}_t v_t^\star) = 0$ in the weak sense, with $\tilde{p}_0 = p_0$. By Proposition 2, the family p_t defined as the law of $\gamma_t(\psi_0, \psi_1)$ also satisfies this continuity equation with p_0 as its initial condition. The Lipschitz regularity of $v_t^\star$ on the compact manifold $\mathbb{CP}^{d-1}$ implies uniqueness of weak solutions to the continuity equation in the standard transport-equation sense (Ambrosio, 2008; Villani, 2009), hence $\tilde{p}_t = p_t$ for all $t \in [0, 1]$. The boundary condition $\gamma_1(\psi_0, \psi_1) = \psi_1$ a.s. gives p_1 as the marginal at $t = 1$, so $(\Phi_{0,1})_\# p_0 = p_1$. □

Remark 7 (Dimension dependence of the Lipschitz constant)**.** *Compactness of $\mathbb{CP}^{d-1}$ removes the linear-growth condition needed for Cauchy–Lipschitz on $\mathbb{R}^N$, and a mild margin condition on the coupling delivers a d-uniform bound on the Lipschitz constant $L = L(v_t^\star)$. For the normalized chord conditional path of Section 4.2,*

$$u_t(\psi_t \mid \psi_0, \psi_1) = \frac{1}{\|(1-t)\psi_0 + t\bar{\psi}_1\|_2}\,\Pi_{\psi_t}(\bar{\psi}_1 - \psi_0),$$

whose smoothness in (ψ_0, ψ_1, t) degrades when $|\langle\psi_0, \psi_1\rangle| \to 0$: the Pancharatnam phase $\alpha(\psi_0, \psi_1) = \langle\psi_0, \psi_1\rangle/|\langle\psi_0, \psi_1\rangle|$ is then ambiguous, and $\|\tilde{\psi}_t\|_2$ approaches its lower envelope $\sqrt{(1-t)^2 + t^2}$ where the radial denominator is least cushioned. Random pure states on $\mathbb{CP}^{d-1}$ are near-orthogonal with overwhelming probability for large d (typical overlap $\sim d^{-1/2}$), so the relevant question is how far the coupling support stays from that regime. The sufficient condition is a margin

$$\inf_{(\psi_0,\psi_1)\in\mathrm{supp}(\pi)} |\langle\psi_0, \psi_1\rangle| \;\geq\; \kappa \;>\; 0, \tag{50}$$

under which a direct calculation gives $L = O(\kappa^{-c})$ for a small numerical constant c that we do not optimize, so the constants in Theorem 2 and Proposition 3 are dimension-uniform on such couplings. Margins of the form (50) *arise naturally from clustered targets and from smoothed conditional path families, which covers the settings we study; Appendix D.16 shows in addition that the conditional velocity stays uniformly bounded even without a margin.*

C.9 Proof of Proposition 3

The proof proceeds in three steps: a pointwise coupling bound, an L^2 aggregation along the true flow, and the transition from coupling cost to Wasserstein-2 distance. The pointwise step is isolated as a lemma so that the change-of-measure issue between p_t and $\hat{p}_t$ is addressed explicitly rather than absorbed into a constant.

Lemma 3 (Pointwise coupling stability under Lipschitz flows)**.** *Let $v^\star, \hat{v} : \mathbb{CP}^{d-1} \times [0,1] \to T\mathbb{CP}^{d-1}$ be tangent vector fields with flows $\Phi_{0,t}, \hat{\Phi}_{0,t}$ from a common initial point ψ_0 in the unit-sphere representation, and assume $\hat{v}$ is L-Lipschitz on $\mathbb{CP}^{d-1}$ uniformly in $t \in [0,1]$. Set $\psi_t^\star := \Phi_{0,t}(\psi_0)$, $\hat{\psi}_t := \hat{\Phi}_{0,t}(\psi_0)$, $\delta_t := \psi_t^\star - \hat{\psi}_t$, and*

$$e_t \;:=\; \big\|v_t^\star(\psi_t^\star) - \hat{v}_t(\psi_t^\star)\big\|_2, \tag{51}$$

the velocity error evaluated along the true *trajectory $\psi_t^\star$. Then for every $t \in [0,1]$,*

$$\|\delta_t\|_2 \;\leq\; \int_0^t e^{L(t-s)}\, e_s\, ds. \tag{52}$$

Proof of Lemma 3. The decisive step is to decompose the velocity difference at the common evaluation point $\psi_t^\star$ rather than at $\hat{\psi}_t$:

$$\frac{d}{dt}\delta_t \;=\; v_t^\star(\psi_t^\star) - \hat{v}_t(\hat{\psi}_t) \;=\; \underbrace{\big[v_t^\star(\psi_t^\star) - \hat{v}_t(\psi_t^\star)\big]}_{\text{velocity error along } p_t} \;+\; \underbrace{\big[\hat{v}_t(\psi_t^\star) - \hat{v}_t(\hat{\psi}_t)\big]}_{\text{Lipschitz term in } \delta_t}.$$

The Lipschitz assumption on $\hat{v}_t$ bounds the second bracket by $L\|\delta_t\|_2$, giving

$$\left\|\frac{d}{dt}\delta_t\right\|_2 \;\leq\; e_t + L\,\|\delta_t\|_2.$$

Since $\frac{d}{dt}\|\delta_t\|_2 \leq \|\frac{d}{dt}\delta_t\|_2$ wherever $\delta_t \neq 0$, with $\delta_0 = 0$, Grönwall's inequality in integral form yields (52). The point of the decomposition is that the velocity error appears at $\psi_t^\star$, whose law is exactly p_t by Theorem 2; no Radon–Nikodym factor between p_t and $\hat{p}_t$ enters the bound. □

Proof of Proposition 3. Couple the two flows by sharing the initial sample $\psi_0 \sim p_0$, and use the notation of Lemma 3. By Theorem 2, $\psi_1^\star \sim p_1$, while $\hat{\psi}_1 \sim \hat{p}_1$ by definition of $\hat{\Phi}_{0,1}$.

Step 1 (pointwise bound). Lemma 3 gives, at $t = 1$,

$$\|\delta_1\|_2 \;\le\; \int_0^1 e^{L(1-s)}\, e_s\, ds.$$

Squaring and applying Cauchy–Schwarz in the time variable,

$$\|\delta_1\|_2^2 \;\le\; \Big(\int_0^1 e^{L(1-s)}\, ds\Big)\Big(\int_0^1 e^{L(1-s)}\, e_s^2\, ds\Big) \;\le\; e^{2L}\int_0^1 e_s^2\, ds.$$

Step 2 (L^2 aggregation along p_t). Taking expectation over $\psi_0 \sim p_0$ and using Theorem 2 to identify the law of $\psi_s^\star$ as p_s, the velocity error e_s becomes the integrand appearing in the proposition's assumption:

$$\begin{aligned}\mathbb{E}_{\psi_0\sim p_0}\,\|\delta_1\|_2^2 \;&\le\; e^{2L}\int_0^1 \mathbb{E}_{\psi_0\sim p_0}\,\big\|v_s^\star(\psi_s^\star) - \hat v_s(\psi_s^\star)\big\|_2^2\, ds \\ &=\; e^{2L}\int_0^1 \mathbb{E}_{\psi\sim p_s}\,\big\|v_s^\star(\psi) - \hat v_s(\psi)\big\|_2^2\, ds \;\le\; e^{2L}\,\varepsilon^2.\end{aligned}$$

The horizontal-norm and ambient-norm coincide on the horizontal subspace, so $\|\cdot\|_\psi$ and $\|\cdot\|_2$ agree under the integrand. This step is where the careful evaluation point in Lemma 3 pays off: no change-of-measure between $\hat p_s$ and p_s is invoked, and no Radon–Nikodym constant has to be tracked.

Step 3 (coupling cost upper-bounds W_2). On $\mathbb{CP}^{d-1}$, the Fubini–Study geodesic distance is dominated by the ambient chord distance in any unit-norm representatives (Bengtsson & Życzkowski, 2017): $d_{\mathrm{FS}}([\psi],[\phi]) \le \|\psi - \phi\|_2$. Hence the coupling $(\psi_1^\star, \hat\psi_1)$ is an admissible plan between p_1 and $\hat p_1$, and

$$W_2^{\mathrm{FS}}\big(\hat p_1,\, p_1\big)^2 \;\le\; \mathbb{E}\, d_{\mathrm{FS}}([\psi_1^\star],[\hat\psi_1])^2 \;\le\; \mathbb{E}\,\|\delta_1\|_2^2 \;\le\; e^{2L}\,\varepsilon^2,$$

giving $W_2^{\mathrm{FS}}(\hat p_1, p_1) \le e^L\,\varepsilon$ as claimed. □

C.10 Proof of Proposition 4

Proof. Let $\varphi \in C^\infty(\mathbb{CP}^{d-1})$ be a smooth test function and let $\psi_t = \Phi_{0,t}(\psi_0)$ with $\psi_0 \sim p_0$. By the chain rule along the deterministic flow generated by X_t,

$$\frac{d}{dt}\varphi(\psi_t) = \langle \nabla_{\mathrm{FS}}\varphi(\psi_t), X_t(\psi_t)\rangle_{\psi_t}.$$

Taking expectation gives

$$\frac{d}{dt}\mathbb{E}[\varphi(\psi_t)] = \mathbb{E}\big[\langle \nabla_{\mathrm{FS}}\varphi(\psi_t), X_t(\psi_t)\rangle_{\psi_t}\big].$$

Since $p_t = (\Phi_{0,t})_{\#}p_0$ is the law of ψ_t, we may rewrite this as

$$\frac{d}{dt}\int_{\mathbb{CP}^{d-1}} \varphi(\psi)\, p_t(\psi)\, d\mathrm{vol}_{\mathrm{FS}}(\psi) = \int_{\mathbb{CP}^{d-1}} \langle \nabla_{\mathrm{FS}}\varphi(\psi), X_t(\psi)\rangle_\psi\, p_t(\psi)\, d\mathrm{vol}_{\mathrm{FS}}(\psi).$$

Integrating by parts on the manifold yields

$$\int_{\mathbb{CP}^{d-1}} \langle \nabla_{\mathrm{FS}}\varphi, X_t\rangle\, p_t\, d\mathrm{vol}_{\mathrm{FS}} = -\int_{\mathbb{CP}^{d-1}} \varphi\, \mathrm{div}_{\mathbb{CP}}(p_t X_t)\, d\mathrm{vol}_{\mathrm{FS}},$$

which proves the weak continuity equation $\partial_t p_t + \mathrm{div}_{\mathbb{CP}}(p_t X_t) = 0$.

For the second claim, assume the IFM conditional path family is chosen so that $\gamma_t(\psi_0,\psi_1) = \Phi_{0,t}(\psi_0)$ almost surely with $\psi_1 = \Phi_{0,1}(\psi_0)$. Then along every conditional path,

$$u_t(\psi_t \mid \psi_0, \psi_1) = \frac{d}{dt}\gamma_t(\psi_0,\psi_1) = \frac{d}{dt}\Phi_{0,t}(\psi_0) = X_t(\psi_t).$$

Applying Theorem 1, the population IFM minimizer is

$$v_t^\star(\psi) = \mathbb{E}\big[u_t(\psi_t \mid \psi_0,\psi_1) \mid \psi_t = \psi\big] = \mathbb{E}\big[X_t(\psi_t) \mid \psi_t = \psi\big] = X_t(\psi)$$

for p_t-almost every ψ. This proves exact recovery of the deterministic projective flow field. □

C.11 Proof of Proposition 5

Proof. Under the deterministic coupling assumption, $\psi_1 = T(\psi_0)$ almost surely, so the conditional velocity is

$$U = u_t(\psi_t \mid \psi_0, \psi_1) = \frac{d}{dt}\gamma_t(\psi_0, T(\psi_0))$$

almost surely. By Theorem 1,

$$v_t^\star(\psi) = \mathbb{E}[U \mid \psi_t = \psi].$$

Now fix $t \in (0,1)$ and suppose $\psi = \psi_t$ lies in the full-measure image of the injective map $F_t(\psi_0) := \gamma_t(\psi_0, T(\psi_0))$. Then there is a unique preimage $\psi_0 = F_t^{-1}(\psi)$. Consequently, conditioning on $\psi_t = \psi$ collapses the conditional expectation:

$$\mathbb{E}[U \mid \psi_t = \psi] = \frac{d}{dt}\gamma_t(F_t^{-1}(\psi), T(F_t^{-1}(\psi))).$$

Evaluating at $\psi = \psi_t = F_t(\psi_0)$ yields

$$v_t^\star(\psi_t) = \frac{d}{dt}\gamma_t(\psi_0, T(\psi_0))$$

for p_t-almost every ψ_t. This proves exact recovery of the deterministic velocity field of the chosen geodesic interpolation.

If, in addition, T is the optimal Monge map for the quadratic Fubini–Study cost, then the curve $t \mapsto \gamma_t(\psi_0, T(\psi_0))$ is precisely the displacement interpolation associated with T. Hence the recovered velocity field is the Benamou–Brenier velocity field of that optimal interpolation. □

C.12 Proof of Proposition 6

Proof. Fix ψ with $\|\psi\|_2 = 1$ and define

$$\mathcal{E}_\psi(w) = \{w + a\psi + b\, i\psi : a, b \in \mathbb{R}\}.$$

The directions ψ and $i\psi$ span the radial and phase-vertical directions of the unit-sphere representation at ψ. Perturbations that differ by an element of $\mathrm{span}_{\mathbb{R}}\{\psi, i\psi\}$ therefore have the same first-order effect on the underlying projective point.

We first show existence and uniqueness of the horizontal representative. Let

$$\xi = w + a\psi + b\, i\psi \in \mathcal{E}_\psi(w).$$

The horizontality condition $\langle \psi, \xi \rangle = 0$ becomes

$$\langle \psi, w \rangle + a + bi = 0.$$

Since $a, b \in \mathbb{R}$, this linear equation has the unique solution

$$a = -\mathrm{Re}\langle \psi, w \rangle, \qquad b = -\mathrm{Im}\langle \psi, w \rangle.$$

Substituting back gives

$$\xi = w - \mathrm{Re}\langle \psi, w \rangle\, \psi - \mathrm{Im}\langle \psi, w \rangle\, i\psi = w - \psi\langle \psi, w \rangle = \Pi_\psi(w).$$

Thus $\Pi_\psi(w)$ is the unique element of $\mathcal{E}_\psi(w)$ lying in $\mathcal{H}_\psi$.

It remains to prove the minimal-norm property. Write any $\xi \in \mathcal{E}_\psi(w)$ as

$$\xi = \Pi_\psi(w) + a\psi + b\, i\psi$$

with $a, b \in \mathbb{R}$. Because $\Pi_\psi(w) \in \mathcal{H}_\psi$, it is orthogonal to both ψ and $i\psi$ under the ambient Hermitian inner product. Therefore

$$\|\xi\|_2^2 = \|\Pi_\psi(w)\|_2^2 + a^2 + b^2.$$

This quantity is minimized if and only if $a = b = 0$, which proves that $\Pi_\psi(w)$ is the unique minimizer of the ambient norm over $\mathcal{E}_\psi(w)$. □

C.13 Proof of Proposition 7

Proof. The argument has two steps: a local one-step truncation analysis (showing first-order consistency of the Euler–retraction map) and a global accumulation argument on $\mathbb{CP}^{d-1}$ (showing that the discrete trajectory converges to the exact flow at first order in Δt).

Step 1: First-order consistency of the retraction. Fix a unit-norm representative ψ and a tangent increment $\xi \in \mathcal{H}_\psi$. By definition,

$$\mathrm{Retr}_\psi(\xi) = \frac{\psi + \xi}{\|\psi + \xi\|_2}.$$

Clearly $\mathrm{Retr}_\psi(0) = \psi$. To verify the differential condition for a first-order retraction, consider

$$R(\epsilon) := \mathrm{Retr}_\psi(\epsilon\xi) = \frac{\psi + \epsilon\xi}{\|\psi + \epsilon\xi\|_2}.$$

Because $\langle \psi, \xi \rangle = 0$, we have

$$\|\psi + \epsilon\xi\|_2^2 = 1 + \epsilon^2\|\xi\|_2^2, \qquad \|\psi + \epsilon\xi\|_2 = 1 + O(\epsilon^2).$$

Therefore

$$R(\epsilon) = (\psi + \epsilon\xi)(1 + O(\epsilon^2)) = \psi + \epsilon\xi + O(\epsilon^2),$$

so

$$\frac{d}{d\epsilon} R(\epsilon)\Big|_{\epsilon=0} = \xi.$$

Thus Retr_ψ is a valid first-order retraction on the unit-sphere representation.

Now consider the exact manifold ODE $\dot{\psi}_t = v_t(\psi_t)$ and a single Euler–retraction step starting from a point ψ_t on the manifold:

$$\psi^{\mathrm{num}}_{t+\Delta t} = \mathrm{Retr}_{\psi_t}\big(\Delta t\, v_t(\psi_t)\big), \qquad \psi^{\mathrm{exact}}_{t+\Delta t} = \mathrm{Exp}_{\psi_t}\big(\Delta t\, v_t(\psi_t)\big) + O(\Delta t^2),$$

the second equality following from the Taylor expansion of the exact solution along the geodesic generated by the exponential map. Since any first-order retraction agrees with the exponential map up to first order,

$$\mathrm{Retr}_{\psi_t}\big(\Delta t\, v_t(\psi_t)\big) = \mathrm{Exp}_{\psi_t}\big(\Delta t\, v_t(\psi_t)\big) + O(\Delta t^2),$$

so the *local truncation error*, when started from the exact state ψ_t, satisfies

$$\big\|\psi^{\mathrm{num}}_{t+\Delta t} - \psi^{\mathrm{exact}}_{t+\Delta t}\big\|_2 \;\le\; C_1\, \Delta t^2, \tag{53}$$

where C_1 depends on the second-order regularity of the exact flow on $\mathbb{CP}^{d-1} \times [0,1]$ and is finite by the assumed C^1 regularity of v_t and compactness of the manifold.

Step 2: Global first-order convergence on $\mathbb{CP}^{d-1}$. Fix a stepsize $\Delta t = 1/K$ for $K \in \mathbb{N}$, let $t_k := k\Delta t$, and let $\Phi_{s,t}$ denote the exact flow of v_t between times s and t. Define the discrete trajectory $\psi^{\mathrm{num}}_0 := \psi_0$,

$$\psi^{\mathrm{num}}_{k+1} \;:=\; \mathrm{Retr}_{\psi^{\mathrm{num}}_k}\big(\Delta t\, v_{t_k}(\psi^{\mathrm{num}}_k)\big),$$

and the global error in the unit-sphere ambient norm

$$E_k \;:=\; \big\|\psi^{\mathrm{num}}_k - \Phi_{0,t_k}(\psi_0)\big\|_2, \qquad E_0 = 0.$$

Assume v_t is L-Lipschitz on $\mathbb{CP}^{d-1}$ uniformly in $t \in [0,1]$ in this representation; this is the standard regularity hypothesis behind global convergence of one-step manifold integrators (Lee, 2012). Decompose

$$\psi^{\mathrm{num}}_{k+1} - \Phi_{0,t_{k+1}}(\psi_0) \;=\; \underbrace{\big[\psi^{\mathrm{num}}_{k+1} - \Phi_{t_k,t_{k+1}}(\psi^{\mathrm{num}}_k)\big]}_{\text{local truncation from } \psi^{\mathrm{num}}_k} + \underbrace{\big[\Phi_{t_k,t_{k+1}}(\psi^{\mathrm{num}}_k) - \Phi_{t_k,t_{k+1}}(\Phi_{0,t_k}\psi_0)\big]}_{\text{exact-flow propagation of } E_k}.$$

By Step 1 applied with starting point ψ^{num}_k, the local-truncation term is bounded by $C_1\,\Delta t^2$. By the standard Grönwall bound for the exact flow of an L-Lipschitz field,

$$\big\|\Phi_{t_k,t_{k+1}}(\psi) - \Phi_{t_k,t_{k+1}}(\phi)\big\|_2 \;\le\; e^{L\Delta t}\,\|\psi - \phi\|_2 \quad \text{for all } \psi, \phi \in \mathbb{CP}^{d-1}.$$

Combining,

$$E_{k+1} \;\le\; e^{L\Delta t}\, E_k + C_1\, \Delta t^2, \qquad k = 0, 1, \ldots, K-1. \tag{54}$$

Iterating (54) from $E_0 = 0$ and using $K\Delta t = 1$,

$$E_K \;\le\; C_1\, \Delta t^2 \sum_{j=0}^{K-1} e^{jL\Delta t} \;=\; C_1\, \Delta t^2 \cdot \frac{e^{KL\Delta t} - 1}{e^{L\Delta t} - 1} \;\le\; C_1\, \Delta t \cdot \frac{e^L - 1}{L},$$

where the last inequality uses $e^{L\Delta t} - 1 \ge L\Delta t$. The same bound applies to every intermediate index $k \le K$, so

$$\sup_{0 \le k \le K} \left\| \psi_k^{\mathrm{num}} - \Phi_{0,t_k}(\psi_0) \right\|_2 \;\le\; C_1\, \frac{e^L - 1}{L}\, \Delta t \;=\; O(\Delta t). \tag{55}$$

Since the Fubini–Study geodesic distance on $\mathbb{CP}^{d-1}$ is dominated by the ambient chord distance (Bengtsson & Życzkowski, 2017), $d_{\mathrm{FS}}([\psi_k^{\mathrm{num}}], [\Phi_{0,t_k}\psi_0]) \le E_k$, and the same $O(\Delta t)$ rate transfers to the intrinsic distance. This establishes global first-order convergence of the Euler–retraction scheme on the projective manifold; the constant in (55) depends only on L and the curvature term C_1, both of which are finite under the standing regularity assumptions, and its dimension dependence is inherited from L as discussed in Remark 7. □

D Additional Experimental Details

This appendix gathers the full experimental configuration used for all results in Section 5. The main text emphasizes the empirical story; here we make the implementation choices explicit by describing the benchmark construction rules, optimization protocol, network architecture, sampling configuration, and baseline-specific hyperparameters.

D.1 Common training and evaluation protocol

Unless otherwise stated, all main-text results are produced by averaging over 10 independent random seeds under the following shared protocol:

- random seeds: 10 independent runs per configuration, with seed indices $\{1, 2, \ldots, 10\}$ used to initialize the network parameters, the data sampler, and the base prior;
- optimizer: AdamW;
- learning rate: 2×10^{-4};
- training batch size: 64;
- evaluation frequency: every 200 optimization steps;
- evaluation batch size: 256;
- training length: 2000 optimization steps;
- reporting rule: for each seed, we select a single checkpoint by validation MMD and report all metrics at that checkpoint;
- aggregation: main-text tables report mean $\pm$ standard deviation across the 10 seeds, computed as the unbiased sample standard deviation, in the format $(\mu \pm \sigma) \times 10^c$.

For IFM and Euclidean FM, the reported runs use $K = 200$ deterministic sampling steps. All synthetic and physics-inspired ensembles use the same perturbation scale $\epsilon_{\mathrm{cluster}} = 0.06$ and, unless otherwise stated, the same local base-prior scale 0.05.

Fairness of baseline comparisons. The IFM, Euclidean FM, SSDM, and Euclidean VP-SDE comparisons use matched network width, depth, training length, evaluation frequency, and $K = 200$ sampling or teacher-discretization steps whenever the method admits such a direct match. QDDPM and QGAN are included as representative circuit-based anchors on the 6-qubit single-cluster task using their aligned training scripts and reported wall-clock times. Accordingly, we interpret the circuit-baseline comparison as a useful scale and wall-clock reference, while the central controlled comparison in the paper is IFM versus Euclidean FM, with SSDM and Euclidean VP-SDE serving as diffusion anchors under the matched protocol described above.

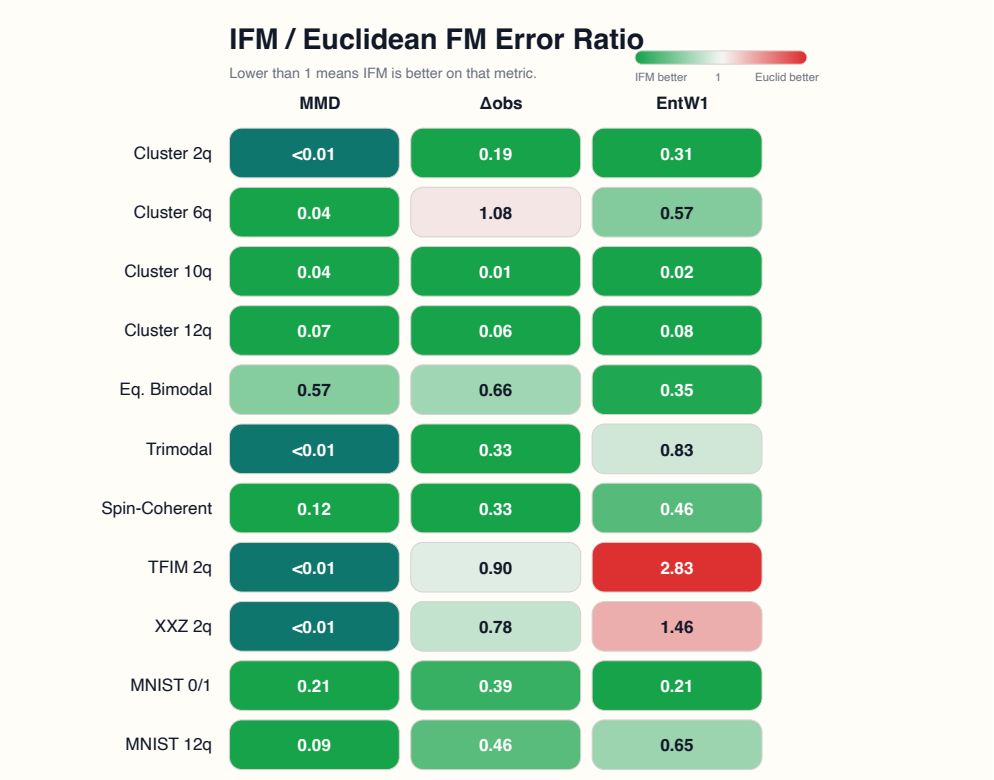

Figure 2: Ratio of IFM to Euclidean FM across all reported benchmarks and metrics, computed from the tabled values in this paper. Values below 1 indicate that IFM attains a lower error. The clearest gains appear on higher-qubit single-cluster tasks, coherence-sensitive multimodal benchmarks, spin-coherent peaks, and the amplitude-encoded MNIST benchmark, while TFIM and XXZ exhibit a more mixed trade-off.

Evaluation metrics. MMD is computed from the pure-state overlap kernel induced by squared fidelities (Gretton et al., 2012). The observable discrepancy Δ_{obs} is the mean absolute difference in the expectations of all one-site Z_i operators and nearest-neighbor Z_iZ_{i+1} correlators. Ent. W_1 is the empirical Wasserstein-1 distance between half-chain entanglement-entropy distributions. These metrics are computed in the same way for IFM, Euclidean FM, SSDM, and Euclidean VP-SDE, so differences reflect the learned generative distributions rather than method-specific evaluation rules (Xu et al., 2026; Song et al., 2021b).

D.2 Benchmark construction

Table 6 summarizes the benchmark families used in the paper. For the synthetic and physics-inspired benchmarks, each target sample is obtained by selecting a canonical center state (or one of several center states), adding small complex Gaussian perturbations in amplitude space, and renormalizing, producing local pure-state ensembles on the projective manifold (Bengtsson & Życzkowski, 2017).

Single-cluster benchmark. For n qubits, the canonical center is $|0\rangle^{\otimes n}$. Each sample is generated by adding i.i.d. complex Gaussian noise of standard deviation $\epsilon_{\text{cluster}}$ to this center vector and renormalizing. This is the easiest unimodal benchmark and is used primarily to probe scaling with qubit number.

Structured multimodal benchmarks. The equatorial bimodal benchmark uses the two GHZ-type equatorial centers (Bengtsson & Życzkowski, 2017; Preskill, 2018)

$$\frac{|0\rangle^{\otimes n} + |1\rangle^{\otimes n}}{\sqrt{2}}, \qquad \frac{|0\rangle^{\otimes n} - |1\rangle^{\otimes n}}{\sqrt{2}}.$$

Table 6: Benchmark families used in the experiments. "Local" base denotes a small complex Gaussian perturbation around $|0\rangle^{\otimes n}$ followed by renormalization.

Benchmark	Qubits	Base prior	Construction
Single-cluster	2, 6, 10, 12 in main text; 14 stress test	Local	Perturb-and-renormalize ensemble around $\|0\rangle^{\otimes n}$ with amplitude noise scale 0.06.
Equatorial bimodal	2, 10 in main text; 6 in appendix; 14 stress test	Local	Equal mixture of $(\|0\rangle^{\otimes n} + \|1\rangle^{\otimes n})/\sqrt{2}$ and $(\|0\rangle^{\otimes n} - \|1\rangle^{\otimes n})/\sqrt{2}$.
Trimodal	2, 10 in main text; 6 in appendix; 14 stress test	Local	Equal mixture of $\|0\rangle^{\otimes n}$, $(\|0\rangle^{\otimes n} + \|1\rangle^{\otimes n})/\sqrt{2}$, and $(\|0\rangle^{\otimes n} - \|1\rangle^{\otimes n})/\sqrt{2}$.
Spin-coherent peaks	2, 10 in main text; 6 in appendix; 14 stress test	Local	Two product spin-coherent states built from single-qubit coherent orientations along $+x$ and $+y$.
TFIM family	2, 10 in main text; 6 in appendix	Local	Exact ground states of $H_{\mathrm{TFIM}} = -\sum_{i=1}^{n-1} Z_i Z_{i+1} - g\sum_{i=1}^{n} X_i$ with $g \in \{0.2, 0.5, 1.0, 2.0\}$, then perturbed and renormalized.
XXZ family	2, 10 in main text; 6 in appendix	Local	Exact ground states of $H_{\mathrm{XXZ}} = \sum_{i=1}^{n-1}(X_iX_{i+1} + Y_iY_{i+1} + \Delta Z_iZ_{i+1})$ with $\Delta \in \{-1.0, 0.0, 0.5, 1.0\}$, then perturbed and renormalized.
MNIST $0/1$	6 in main text; 12 extension; 14 stress test	Local	Balanced MNIST digits $\{0, 1\}$, flattened, reduced to $\min(2^n, 784)$ real dimensions, centered per example, zero-padded when needed, and normalized as real amplitude vectors.

The trimodal benchmark augments these two coherent equatorial modes with the pole state $|0\rangle^{\otimes n}$. These tasks are deliberately designed to be phase-sensitive and therefore test whether intrinsic projective transport is more informative than ambient Euclidean interpolation.

Spin-coherent benchmark. We use the standard single-qubit coherent-state parameterization (Arecchi et al., 1972)

$$|\theta, \phi\rangle = \cos(\theta/2)\,|0\rangle\, e^{i\phi}\sin(\theta/2)\,|1\rangle\,,$$

with the two orientations $(\theta, \phi) = (\pi/2, 0)$ and $(\pi/2, \pi/2)$, corresponding to the $+x$ and $+y$ directions on the Bloch sphere. The n-qubit benchmark centers are the corresponding tensor-product states.

TFIM and XXZ families. For the TFIM and XXZ benchmarks, we use compact spin-chain families inspired by standard condensed-matter models (Chaikin & Lubensky, 1995; Bose, 2003). For the TFIM benchmark we diagonalize

$$H_{\mathrm{TFIM}} = -\sum_{i=1}^{n-1} Z_iZ_{i+1} - g\sum_{i=1}^{n} X_i, \qquad g \in \{0.2, 0.5, 1.0, 2.0\},$$

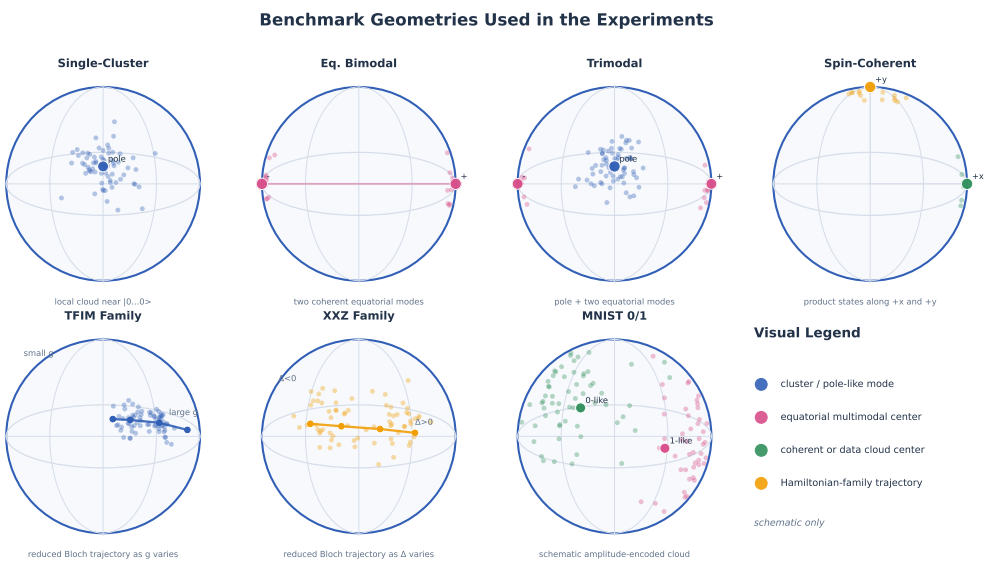

Figure 3: Effective Bloch-sphere views of the benchmark families used in our experiments. The panels are schematic and are included to make the qualitative geometry behind Table 3 and Table 11 visually interpretable.

and for the XXZ benchmark we diagonalize

$$H_{\mathrm{XXZ}} = \sum_{i=1}^{n-1} (X_i X_{i+1} + Y_i Y_{i+1} + \Delta Z_i Z_{i+1}), \qquad \Delta \in \{-1.0, 0.0, 0.5, 1.0\}.$$

The resulting exact ground states are then used as benchmark centers before perturbation and renormalization. These families are evaluated at 2 qubits in the main text and extended to 6 qubits in the appendix. These tasks are intended as compact proof-of-concept physics-inspired families rather than full many-body benchmarks. They should therefore be read as controlled sanity checks for whether IFM behaves sensibly on Hamiltonian-generated state families, not as broad evidence for many-body physics performance. In particular, the TFIM and XXZ rows use only four Hamiltonian parameters each, so the main empirical support for the geometry claim comes from the controlled IFM-versus-Euclidean-FM comparisons across scaling, multimodal, phase-sensitive, and appendix stress-test settings rather than from these compact families alone.

MNIST-derived amplitude-vector benchmark. For MNIST we use balanced digits $\{0, 1\}$ from the training set (LeCun et al., 1998). This benchmark is a classical image-vector distribution embedded into the pure-state representation, rather than experimentally collected quantum data. For the 6-qubit experiment, images are flattened, rescaled to $[0, 1]$, reduced to 64 dimensions, centered by subtracting the per-example mean, and normalized to unit Euclidean norm. For the 12-qubit extension, we use the flattened 784-dimensional vectors, center them per example, and zero-pad them to dimension $2^{12} = 4096$ before normalization. In both cases, the resulting real vectors are interpreted as pure-state amplitudes.

How to read Figure 3. The synthetic multimodal benchmarks are shown by their mode centers on an effective two-level sphere generated by $|0\rangle^{\otimes n}$ and $|1\rangle^{\otimes n}$. The spin-coherent benchmark uses physical single-qubit coherent orientations, while the TFIM and XXZ panels show reduced Bloch trajectories of their ground-state families as the Hamiltonian parameter varies. The MNIST panel is an effective cloud sketch after amplitude encoding and should be read only as a geometric summary rather than as an exact Bloch embedding of the full data distribution.

D.3 Model architectures

All learned baselines use the same basic MLP scale so that the main empirical differences are due to geometry and training objective rather than network capacity. Every model uses a 128-dimensional sinusoidal time embedding followed by a small time-embedding MLP and a main backbone of width 512 and depth 5.

IFM. The IFM model in `sse1_fm.py` predicts a complex velocity field from the concatenation of the real and imaginary parts of the current state together with the time embedding, following the velocity-field parameterization used in flow matching while enforcing projective tangency (Lipman et al., 2023; Chen & Lipman, 2024). Its output is projected to the tangent space at the current state. In the main experiments, the conditional path family is the phase-aligned normalized-chord path of Section 4.2, i.e. `path_type=chord`.

Euclidean FM. The Euclidean FM baseline in `sse1_fm_euclid.py` uses the same time-embedding size, width, and depth, but acts on the ambient real representation in $\mathbb{R}^{2d}$ (Lipman et al., 2023). It predicts an unconstrained Euclidean velocity and does not quotient out global phase or impose horizontal tangency.

SSDM. The intrinsic diffusion baseline in `sse1.py` uses a score network of the same width (512), depth (5), and time-embedding dimension (128) (Xu et al., 2026). Its output is interpreted as a complex score and projected back to the tangent space, consistent with the local OU/VP teacher construction used in that model.

Euclidean VP-SDE. The Euclidean diffusion baseline in `esde.py` also uses a 128-dimensional time embedding and a width-512, depth-5 MLP, but predicts an ambient score in the real concatenated state-vector representation (Song et al., 2021b).

D.4 Code and supplementary materials

The complete implementation used for the experiments is provided in the supplementary material accompanying the submission. In particular, the supplementary package contains the key Python entry points used throughout the empirical study:

- `sse1_fm.py`: the main IFM implementation, including the phase-aligned normalized-chord path, the geodesic ablation, and the appendix-only generic-geodesic baseline;
- `sse1_fm_euclid.py`: the Euclidean flow-matching baseline in the ambient real representation;
- `sse1.py`: the SSDM baseline;
- `esde.py`: the Euclidean VP-SDE baseline;
- the QDDPM and QGAN baseline scripts used in the representative diffusion and adversarial comparisons reported in Table 1.

The main IFM algorithmic logic is concentrated in `sse1_fm.py`. That file implements the data generators, conditional-path constructors, tangent-projected velocity network, ODE sampling routine, and evaluation metrics. We include it in the supplement so that the main empirical claims can be reproduced directly from a single reference implementation.

D.5 Baseline-specific hyperparameters

Table 7 collects the method-specific settings used in the paper. For Euclidean VP-SDE, QDDPM, and QGAN, which are only evaluated on the representative 6-qubit single-cluster benchmark, we report the exact settings used in those runs (Song et al., 2021b; Zhang et al., 2024; Lloyd & Weedbrook, 2018; Dallaire-Demers & Killoran, 2018). For SSDM, we use the same underlying implementation family both on the representative single-cluster comparison and on the 2- and 10-qubit structured benchmarks, with the dataset mode changed accordingly (Xu et al., 2026).

Table 7: Baseline-specific hyperparameters for the reported experiments. The first two rows summarize the settings used throughout the main IFM-versus-Euclidean-FM comparison (Lipman et al., 2023), while the remaining rows summarize SSDM (Xu et al., 2026), Euclidean VP-SDE (Song et al., 2021b), QDDPM (Zhang et al., 2024), and QGAN (Lloyd & Weedbrook, 2018; Dallaire-Demers & Killoran, 2018) on the representative 6-qubit single-cluster comparison reported in Table 1.

Method	Configuration details
IFM (ours)	AdamW, learning rate 2×10^{-4}, batch size 64, training steps 2000, evaluation every 200 steps, evaluation batch 256, deterministic sampling with $K = 200$ and $\Delta t = 1/200$, data-pool size 4096, base prior `local` with scale 0.05, and `path_type=chord`.
Geodesic IFM (ours variant)	Representative 6-qubit single-cluster ablation only. Same optimizer, network scale, data pool, batch size, sample budget, and deterministic sampler as IFM, but with `path_type=geodesic`, i.e. a phase-aligned Fubini–Study-geodesic interpolation used in place of the normalized-chord surrogate.
Generic geodesic FM (Chen & Lipman, 2024)	Representative 6-qubit single-cluster appendix baseline only. Same optimizer, network scale, data pool, batch size, sample budget, and deterministic sampler as IFM, but with `path_type=generic_geodesic`. This variant omits Pancharatnam phase alignment and therefore excludes the quotient-specific endpoint correction introduced by our method.
Euclidean FM	AdamW, learning rate 2×10^{-4}, batch size 64, training steps 2000, evaluation every 200 steps, evaluation batch 256, deterministic sampling with $K = 200$ in the reported runs, data-pool size 4096, and local ambient Gaussian base with scale 0.05.
SSDM	Used for the representative 6-qubit single-cluster comparison and for the 2- and 10-qubit structured benchmarks in Table 3. AdamW, learning rate 2×10^{-4}, batch size 64, training steps 2000, evaluation every 200 steps, local OU/VP teacher with $K = 200$, $\Delta t = \delta = 1/200$, $\sigma_{\min} = 0.05$, $\sigma_{\max} = 1.0$, mean-reversion parameter $\lambda_0 = 0.2$, and the same benchmark-specific state generators used in the flow-matching scripts via `SSDM_DATA_MODE`.
Euclidean VP-SDE	Representative 6-qubit single-cluster run only. AdamW, learning rate 2×10^{-4}, batch size 64, training steps 2000, evaluation every 200 steps, evaluation batch 256, reverse VP-SDE sampling with $K = 200$, and linear schedule parameters $\beta_{\min} = 0.1$ and $\beta_{\max} = 20.0$.
QDDPM	Representative 6-qubit single-cluster run only. Executed in the `qddpm` conda environment using `QDDPM_torch.py`. The baseline uses time-dependent PQCs with 6 layers per diffusion step over $T = 20$ steps. The wall-clock training time measured in our run is reported in Table 8. Metrics are computed with the same definitions as in the other baselines.
QGAN	Representative 6-qubit single-cluster run only. Executed in the `qddpm` conda environment using `QGAN.py`, `qgan_train_1.py`, and the aligned runner `qgan-1.py`. The baseline uses a 120-layer PQC generator together with a 16-layer PQC discriminator. The wall-clock training time measured in our run is reported in Table 8.

Table 8: Training-time comparison on the representative 6-qubit single-cluster benchmark, including QDDPM (Zhang et al., 2024), QGAN (Lloyd & Weedbrook, 2018; Dallaire-Demers & Killoran, 2018), and our method. Lower is better. Speedup is measured relative to QDDPM.

Method	Train time (s) ↓	Speedup
QDDPM	21150.52	1.0×
QGAN	198.64	106.4×
IFM (ours, chord)	240.91	87.8×

Scope of the diffusion comparisons. Our main experimental question is whether respecting the intrinsic geometry of $\mathbb{CP}^{d-1}$ matters once the learning principle is held fixed. That question is most directly isolated by comparing IFM to Euclidean FM across the full benchmark suite. We therefore use Euclidean VP-SDE as a representative ambient diffusion anchor only on the single-cluster task, while SSDM is evaluated both on that representative task and on the small- and higher-qubit structured benchmarks where an intrinsic diffusion comparison remains especially informative (Song et al., 2021b; Xu et al., 2026). Even so, the broader sweep still concentrates on the sharpest experimental contrast in the paper, namely intrinsic-versus-ambient flow matching. Table 9 extends the single-cluster baseline picture to 10 qubits for the scalable baselines and shows that the separation between IFM and the diffusion baselines becomes much larger once the dense-state dimension increases.

D.6 Sampling and reporting details

For IFM, samples are generated by integrating the learned tangent field from the chosen base prior and applying the normalization-based retraction at each step. For Euclidean FM, samples are generated in

Table 9: Appendix-only 10-qubit single-cluster extension of Table 1 for the scalable baselines: Euclidean VP-SDE (Song et al., 2021b), SSDM (Xu et al., 2026), and Euclidean FM (Lipman et al., 2023). Lower is better. QDDPM and QGAN are omitted here because our reported circuit-based runs were only executed on the representative 6-qubit benchmark.

Method	MMD ↓	Δ_{obs} ↓	Ent. W_1 ↓
Euclidean VP-SDE	2.3049×10^{-2}	1.1782×10^{-1}	8.7980×10^{-2}
SSDM	1.4157×10^{-2}	1.1883×10^{-1}	8.6077×10^{-2}
Euclidean FM	1.6830×10^{-2}	9.5109×10^{-2}	8.2071×10^{-2}
Geodesic IFM (ours variant)	7.3038×10^{-4}	1.3659×10^{-3}	1.5996×10^{-3}
IFM (ours)	7.3779×10^{-4}	1.3936×10^{-3}	1.4582×10^{-3}

the ambient real representation and then mapped back to valid pure states by complex reconstruction and final normalization. For SSDM and Euclidean VP-SDE, we use the reverse-time stochastic samplers implemented in their respective scripts with the settings summarized above.

All numbers reported in the main text are produced under the matched training-budget and evaluation protocol described above. More precisely, for each seed we evaluate the model every 200 optimization steps, select the checkpoint with the best validation MMD, and then report MMD, Δ_{obs}, and Ent. W_1 together at that single checkpoint. The main-text table entries report the mean and the unbiased sample standard deviation across the 10 seeds (seed indices $\{1, \ldots, 10\}$); the observed variances are small relative to the mean differences between methods and do not affect the qualitative conclusions. Because the main protocol is intentionally lightweight, Appendix D.8 additionally includes a 20,000-step high-dimensional flow-matching check to test whether the Euclidean FM gap on a 14-qubit task is mainly an undertraining artifact. Our aim in this paper is to establish the empirical value of the geometric inductive bias rather than to present a final large-scale benchmark study. Whenever a value is shown as "≈ 0", the underlying number was positive but rounded to zero at the output precision used during training.

D.7 Additional single-cluster discussion

The single-cluster benchmark is intentionally the easiest setting in the paper, but it is useful for isolating scaling effects. At 2 qubits, both IFM and Euclidean FM fit the target well, while IFM still improves all three metrics. At 6 qubits, the main improvement is in MMD and entanglement structure, with comparable observable discrepancy. By 10 and 12 qubits, the gap becomes much larger across all reported metrics, consistent with the view that projective geometric bias becomes more valuable as ambient redundancy grows. The geodesic IFM ablation is close to IFM on the representative 6-qubit single-cluster task, suggesting that for this easy unimodal family the main benefit is intrinsic transport itself rather than the precise choice between geodesic and phase-aligned normalized-chord paths. Table 9 extends the scalable-baseline comparison to 10 qubits, and Table 8 reports the large training-time gap between circuit baselines and IFM.

D.8 Fourteen-qubit flow-matching stress test

Table 10 reports an additional 14-qubit stress test comparing only IFM and Euclidean FM on the scalable benchmark families. These runs use the same 2000 training steps, $K = 200$ deterministic sampling steps, network scale, and evaluation protocol as the main flow-matching experiments. We include them to show the behavior of the intrinsic-versus-ambient flow comparison at a larger dense-state dimension.

To check whether the 2000-step budget is simply undertraining the ambient baseline in the high-dimensional regime, we additionally ran a 20,000-step comparison on the 14-qubit single-cluster task. Figure 4 plots MMD, Δ_{obs}, and Ent. W_1 throughout training. Euclidean FM remains essentially flat in MMD over the longer run, ending at 7.88×10^{-3}, close to its 2000-step value in Table 10. IFM remains more than an order of magnitude lower in MMD and also retains lower observable and entanglement discrepancies by the end of the run. This does not rule out all possible effects of hyperparameter tuning or architecture changes, but it indicates that the reported high-dimensional gap is not explained by the short training schedule alone.

Table 10: 14-qubit stress test for IFM and Euclidean FM on scalable benchmark families. Lower is better.

Benchmark	Method	MMD ↓	Δ_{obs} ↓	Ent. W_1 ↓
Single-cluster	Euclidean FM	7.8806×10^{-3}	7.9683×10^{-3}	2.0976×10^{-3}
Single-cluster	IFM (ours)	2.7314×10^{-4}	3.2485×10^{-3}	1.5192×10^{-3}
Equatorial bimodal	Euclidean FM	7.8461×10^{-3}	4.1493×10^{-3}	1.2981×10^{-3}
Equatorial bimodal	IFM (ours)	3.2034×10^{-4}	4.0576×10^{-3}	1.0721×10^{-3}
Trimodal	Euclidean FM	7.8489×10^{-3}	5.0208×10^{-3}	1.5435×10^{-3}
Trimodal	IFM (ours)	3.2284×10^{-4}	2.7896×10^{-3}	1.2578×10^{-3}
Spin-coherent	Euclidean FM	7.8462×10^{-3}	4.0279×10^{-4}	2.4284×10^{-3}
Spin-coherent	IFM (ours)	2.8126×10^{-4}	5.8993×10^{-4}	2.3783×10^{-3}
MNIST $0/1$	Euclidean FM	1.5637×10^{-1}	3.2145×10^{-1}	2.8245×10^{0}
MNIST $0/1$	IFM (ours)	2.4333×10^{-2}	1.4485×10^{-1}	1.7042×10^{0}

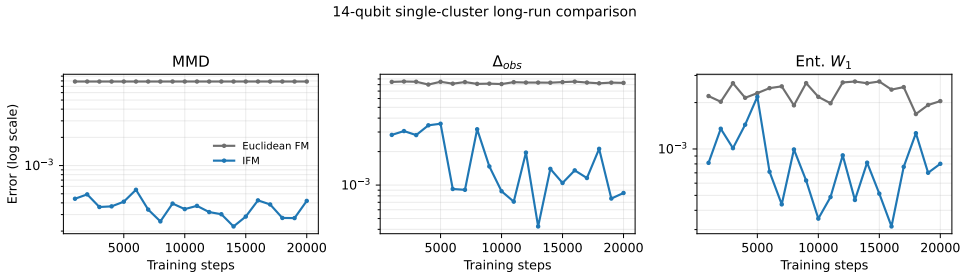

Figure 4: 20,000-step high-dimensional training check on the 14-qubit single-cluster benchmark. Both methods use the same network scale, batch size, local base prior, evaluation batch size, deterministic sampler with $K = 200$, and optimizer settings as the main flow-matching experiments, except for the longer training horizon and evaluation every 1000 steps. Lower is better.

D.9 Additional structured-benchmark discussion

The synthetic multimodal tasks highlight where phase-aware projective transport matters most. In the equatorial bimodal benchmark, the two modes differ by relative phase rather than by occupation of distant computational-basis corners; the trimodal task adds a pole state to these coherent modes. These are precisely the settings where Pancharatnam-aligned conditional paths are more informative than ambient Euclidean interpolation. The SSDM rows show that intrinsic diffusion can already improve over Euclidean FM on several structured tasks, but IFM often gives stronger MMD and better coherence-sensitive behavior, especially at 10 qubits. The spin-coherent benchmark gives a related geometric test using product states oriented along different Bloch-sphere directions; IFM again achieves much smaller MMD and entanglement Wasserstein distance at larger scale.

D.10 Six-qubit extensions of the structured benchmarks

To test whether the trends in Table 3 remain stable between the small-system and higher-qubit regimes, we also evaluated 6-qubit versions of the first five benchmark families from that table. Table 11 reports the corresponding 6-qubit results under the same matched protocol. At intermediate scale the picture is benchmark-dependent in an interpretable way: IFM is clearly stronger on trimodal and TFIM-style structured transport, competitive on the equatorial bimodal benchmark, ahead on MMD and observable matching for spin-coherent states, and trades off on XXZ, whose ground-state family carries the least projective phase structure.

Table 11: Six-qubit extensions of the first five benchmark families from Table 3, comparing IFM with Euclidean FM (Lipman et al., 2023). Lower is better. These appendix-only results probe the intermediate regime between the small-system and 10-qubit results reported in the main text.

Benchmark	Method	MMD ↓	Δ_{obs} ↓	Ent. W_1 ↓
Eq. bimodal (6 qubits)	Euclidean FM	2.1861×10^{-2}	1.0823×10^{-1}	2.4905×10^{-2}
Eq. bimodal (6 qubits)	IFM (ours)	6.2495×10^{-3}	6.4122×10^{-2}	2.9000×10^{-2}
Trimodal (6 qubits)	Euclidean FM	1.3397×10^{-2}	4.2138×10^{-2}	6.5884×10^{-2}
Trimodal (6 qubits)	IFM (ours)	3.6621×10^{-4}	2.0026×10^{-2}	5.3195×10^{-2}
Spin-coherent (6 qubits)	Euclidean FM	1.6100×10^{-2}	1.1463×10^{-2}	1.5728×10^{-1}
Spin-coherent (6 qubits)	IFM (ours)	4.7601×10^{-3}	6.5775×10^{-3}	2.0694×10^{-1}
TFIM (6 qubits)	Euclidean FM	3.1681×10^{-2}	5.6582×10^{-2}	6.2842×10^{-2}
TFIM (6 qubits)	IFM (ours)	1.7825×10^{-3}	1.0956×10^{-2}	2.5062×10^{-2}
XXZ (6 qubits)	Euclidean FM	1.1880×10^{-2}	1.0420×10^{-2}	2.7554×10^{-2}
XXZ (6 qubits)	IFM (ours)	1.2254×10^{-2}	4.8443×10^{-2}	9.1078×10^{-2}

Table 12: Metric-wise winners in the 6-qubit geometry/path ablation. Candidates are Euclidean FM, generic geodesic FM (Chen & Lipman, 2024), and IFM. Lower is better for all metrics.

Benchmark	MMD winner	Δ_{obs} winner	Ent. W_1 winner	Takeaway
Single-cluster	Generic geodesic FM	Euclidean FM	IFM	Mostly geometry-driven; easy task.
Eq. bimodal	IFM	IFM	Euclidean FM	Phase alignment helps MMD/observables.
Trimodal	IFM	IFM	IFM	Strongest evidence for IFM coupling.
Spin-coherent	IFM	IFM	Euclidean FM	MMD/observables improve; entanglement mixed.
TFIM	IFM	IFM	IFM	IFM best among these flow variants.
XXZ	Generic geodesic FM	Generic geodesic FM	Generic geodesic FM	Generic geodesic coupling is better here.

D.11 Six-qubit geometry and path-choice ablation summary

Table 12 summarizes the metric-wise winners among three flow-matching variants at 6 qubits: ambient Euclidean FM, generic geodesic FM on the manifold, and IFM. The purpose is to disentangle three effects that are otherwise coupled in the main comparison: moving from ambient to intrinsic geometry, choosing a generic manifold geodesic versus a quantum-specific path, and applying Pancharatnam phase alignment. The summary shows that intrinsic geometry explains a large part of the gain over Euclidean FM, but the remaining advantage of IFM is benchmark dependent. In particular, IFM is strongest on trimodal and TFIM at 6 qubits, while generic geodesic FM is the clear winner on XXZ at the same scale. This localizes the effect precisely: phase-aligned IFM is strongest exactly where coherent mode organization makes endpoint gauge choice and path coupling matter.

D.12 Phase-alignment ablation on six-qubit structured benchmarks

To probe the role of quotient-aware gauge fixing, we compare IFM against a matched intrinsic baseline whose conditional path is the gauge-unfixed geodesic interpolation of Appendix C.2. The ablation keeps the same IFM network, optimizer, tangent projection, and deterministic sampling scheme; what changes is the conditional path, which now interpolates the raw endpoint representatives instead of the Pancharatnam-aligned ones. The baseline therefore still learns an intrinsic manifold flow, but no longer fixes the endpoint gauge before interpolation.

Table 13: Ablation of the phase-alignment component on five 6-qubit structured benchmarks. Lower is better. "No phase alignment" uses the same intrinsic FM model with `path_type=generic_geodesic`, i.e. the raw endpoint representatives are interpolated directly without Pancharatnam gauge fixing.

Benchmark	Method	MMD ↓	Δ_{obs} ↓	Ent. W_1 ↓
Eq. bimodal (6 qubits)	No phase alignment	8.1550×10^{-3}	6.5664×10^{-2}	2.8592×10^{-2}
Eq. bimodal (6 qubits)	IFM (ours)	6.2495×10^{-3}	6.4122×10^{-2}	2.9000×10^{-2}
Trimodal (6 qubits)	No phase alignment	6.1315×10^{-4}	2.3808×10^{-2}	5.3652×10^{-2}
Trimodal (6 qubits)	IFM (ours)	3.6621×10^{-4}	2.0026×10^{-2}	5.3195×10^{-2}
Spin-coherent (6 qubits)	No phase alignment	1.1341×10^{-2}	9.1691×10^{-3}	1.6435×10^{-1}
Spin-coherent (6 qubits)	IFM (ours)	4.7601×10^{-3}	6.5775×10^{-3}	2.0694×10^{-1}
TFIM (6 qubits)	No phase alignment	1.8038×10^{-2}	5.3284×10^{-2}	1.2640×10^{-1}
TFIM (6 qubits)	IFM (ours)	1.7825×10^{-3}	1.0956×10^{-2}	2.5062×10^{-2}
XXZ (6 qubits)	No phase alignment	9.5237×10^{-3}	9.9533×10^{-3}	1.0297×10^{-2}
XXZ (6 qubits)	IFM (ours)	1.2254×10^{-2}	4.8443×10^{-2}	9.1078×10^{-2}

Table 13 reports the resulting 6-qubit comparison on all five structured benchmarks from Table 3. On the equatorial bimodal and trimodal tasks, removing phase alignment consistently worsens MMD and observable matching, with the trimodal benchmark showing the clearest across-the-board degradation. On TFIM the gauge-unfixed variant is markedly worse on all three metrics, while spin-coherent and XXZ are more evenly matched, which places the effect where the projective picture predicts it: gauge fixing pays off when the target ensemble is organized by coherent multimodal structure rather than by a smooth family geometry alone. We therefore read quotient-consistent gauge fixing as the ingredient that matters for ensembles whose transport depends on coherent mode organization.

D.13 Generic manifold-flow-matching baseline on six-qubit single-cluster and structured tasks

To complement the Euclidean and quantum-specific baselines, we also evaluated an appendix-only generic manifold-flow-matching variant inspired by *Flow Matching on General Geometries* (Chen & Lipman, 2024). Because it interpolates the raw unit-norm representatives supplied by the data, it is a gauge-unfixed sphere lift rather than a fully phase-invariant Riemannian flow-matching baseline on $\mathbb{CP}^{d-1}$; the phase-aligned geodesic variant of Appendix C.2 plays the latter role for nonorthogonal endpoint pairs. This is an author-constructed comparator rather than a baseline reported in the original RFM paper, which does not include a $\mathbb{CP}^{d-1}$ quantum pure-state benchmark or a Pancharatnam/Hopf-quotient implementation. The goal of this comparison is deliberately narrow: we keep the same IFM network, optimizer, tangent projection, sampling step count, deterministic sampler, data pools, and evaluation metrics, but replace our quotient-aware path construction by the raw geodesic interpolation described in Appendix C.2. In particular, this baseline does *not* perform Pancharatnam phase alignment before defining the conditional path. It is therefore a strong generic intrinsic baseline, but it is not specifically designed for the projective gauge structure of quantum pure states.

Table 14 shows the resulting 6-qubit comparison on the easy single-cluster benchmark together with the five structured benchmarks from Table 3, plus one higher-qubit XXZ extension. On the single-cluster task, the generic geodesic manifold baseline already captures most of the gain obtained by moving from ambient Euclidean flow matching to an intrinsic manifold formulation. Across the structured benchmarks, its behavior is more mixed: it improves substantially over Euclidean FM on equatorial bimodal, trimodal, TFIM, and XXZ in MMD, but remains clearly weaker than IFM on TFIM and on the coherence-sensitive multimodal tasks in overall balance. XXZ at 6 qubits is the main counterexample, where the generic geodesic manifold baseline is the strongest of the three. By contrast, on the 10-qubit XXZ extension, the generic manifold baseline becomes highly competitive but IFM regains a small edge on all three reported metrics. We therefore interpret the FMGG-style comparison as evidence that raw manifold geodesics recover part of the intrinsic benefit, while also exposing the tension in the narrative: IFM's advantage is not guaranteed by intrinsic geometry alone. Rather, the strongest claim supported by these results is that quotient-aware phase alignment and the

Table 14: Appendix-only comparison to a generic manifold-geodesic flow-matching baseline inspired by manifold flow matching (Chen & Lipman, 2024), together with Euclidean FM (Lipman et al., 2023), on 6-qubit benchmarks and an additional 10-qubit XXZ comparison. Lower is better.

Benchmark	Method	MMD ↓	Δ_{obs} ↓	Ent. W_1 ↓
Single-cluster (6 qubits)	Euclidean FM	6.0216×10^{-3}	2.1572×10^{-3}	8.1543×10^{-3}
Single-cluster (6 qubits)	Generic geodesic FM (Chen & Lipman, 2024)	2.5475×10^{-4}	2.4098×10^{-3}	5.2252×10^{-3}
Single-cluster (6 qubits)	Geodesic IFM (ours variant)	2.4933×10^{-4}	2.4820×10^{-3}	4.3428×10^{-3}
Single-cluster (6 qubits)	IFM (ours)	2.6280×10^{-4}	2.3388×10^{-3}	4.6615×10^{-3}
TFIM (6 qubits)	Euclidean FM	3.1681×10^{-2}	5.6582×10^{-2}	6.2842×10^{-2}
TFIM (6 qubits)	Generic geodesic FM (Chen & Lipman, 2024)	1.8038×10^{-2}	5.3284×10^{-2}	1.2640×10^{-1}
TFIM (6 qubits)	IFM (ours)	1.7825×10^{-3}	1.0956×10^{-2}	2.5062×10^{-2}
Eq. bimodal (6 qubits)	Euclidean FM	2.1861×10^{-2}	1.0823×10^{-1}	2.4905×10^{-2}
Eq. bimodal (6 qubits)	Generic geodesic FM (Chen & Lipman, 2024)	8.1550×10^{-3}	6.5664×10^{-2}	2.8592×10^{-2}
Eq. bimodal (6 qubits)	IFM (ours)	6.2495×10^{-3}	6.4122×10^{-2}	2.9000×10^{-2}
Trimodal (6 qubits)	Euclidean FM	1.3397×10^{-2}	4.2138×10^{-2}	6.5884×10^{-2}
Trimodal (6 qubits)	Generic geodesic FM (Chen & Lipman, 2024)	6.1315×10^{-4}	2.3808×10^{-2}	5.3652×10^{-2}
Trimodal (6 qubits)	IFM (ours)	3.6621×10^{-4}	2.0026×10^{-2}	5.3195×10^{-2}
Spin-coherent (6 qubits)	Euclidean FM	1.6100×10^{-2}	1.1463×10^{-2}	1.5728×10^{-1}
Spin-coherent (6 qubits)	Generic geodesic FM (Chen & Lipman, 2024)	1.1341×10^{-2}	9.1691×10^{-3}	1.6435×10^{-1}
Spin-coherent (6 qubits)	IFM (ours)	4.7601×10^{-3}	6.5775×10^{-3}	2.0694×10^{-1}
XXZ (6 qubits)	Euclidean FM	1.1880×10^{-2}	1.0420×10^{-2}	2.7554×10^{-2}
XXZ (6 qubits)	Generic geodesic FM (Chen & Lipman, 2024)	9.5237×10^{-3}	9.9533×10^{-3}	1.0297×10^{-2}
XXZ (6 qubits)	IFM (ours)	1.2254×10^{-2}	4.8443×10^{-2}	9.1078×10^{-2}
XXZ (10 qubits)	Euclidean FM	1.3065×10^{-2}	2.2184×10^{-2}	4.2252×10^{-2}
XXZ (10 qubits)	Generic geodesic FM (Chen & Lipman, 2024)	6.4518×10^{-4}	4.3118×10^{-3}	2.5516×10^{-2}
XXZ (10 qubits)	IFM (ours)	3.5211×10^{-4}	3.9172×10^{-3}	2.5046×10^{-2}

IFM path choice often help when the target ensemble has coherent projective organization, whereas smoother or benchmark-specific couplings can favor a generic geodesic path.

D.14 Matched ambient controls

IFM and the ambient Euclidean-FM baseline differ in more than the intrinsic-versus-ambient formulation: IFM draws its base samples from a normalized pure-state prior and keeps the iterate on the sphere by retraction at every sampling step, whereas the ambient baseline uses an unnormalized Gaussian prior and normalizes only at the end. Either of those two implementation choices, rather than the quotient-aware geometry, could in principle account for the reported gap. To test this we reran the Euclidean-FM implementation of Tables 2 and 3 unchanged—its own network, sampler, metrics, and checkpoint-selection rule—with the base distribution matched to the IFM prior and with per-step renormalization enabled, while deliberately leaving the ambient interpolation itself unchanged. Switching both modifications off reproduces those Euclidean-FM numbers, so the matched rows below differ from the baseline only in those two ingredients.

Table 15 reports the matched baseline over 10 seeds against the reported Euclidean-FM and IFM values. Matching the prior and the constraint handling changes the ambient baseline only modestly and does not close the distributional gap: IFM retains a $7.3\times$–$21.8\times$ MMD advantage on all five benchmarks. On the observable and entanglement metrics the matched baseline is competitive, gaining on the multimodal and amplitude-encoded families while IFM stays ahead on the single-cluster task, where base matching in fact costs the ambient baseline accuracy. The control thus rules out the two ambient implementation details as the source of the distributional gains and attributes them to the quotient-aware construction itself—the phase-aligned path together with the horizontal parameterization—in line with the path-family ablations of Appendices D.12 and D.13.

D.15 Rerunning the main benchmarks with the exactly equivariant field

Appendix B.5 shows on a single benchmark that the canonical-gauge wrapper removes the residual $U(1)$ equivariance error of the raw network. Because the wrapper changes the parameterization rather than merely diagnosing it, we repeated all main benchmarks end to end with it active during both training and sampling. The raw and equivariant runs share architecture, optimizer, checkpoint

Table 15: Matched ambient control. "Euclidean FM (as reported)" and "IFM" are the values of Tables 2 and 3; "Euclidean FM, matched" is that same implementation with the IFM base prior and per-step renormalization, over 10 seeds. Lower is better.

Benchmark	Method	MMD ↓	Δ_{obs} ↓	Ent. W_1 ↓
Single-cluster (6 qubits)	Euclidean FM (as reported)	6.02×10^{-3}	2.16×10^{-3}	8.15×10^{-3}
Single-cluster (6 qubits)	Euclidean FM, matched ($n = 10$)	5.32×10^{-3}	1.91×10^{-2}	3.64×10^{-2}
Single-cluster (6 qubits)	IFM (ours)	2.63×10^{-4}	2.34×10^{-3}	4.66×10^{-3}
MNIST $0/1$ (6 qubits)	Euclidean FM (as reported)	2.10×10^{-2}	1.30×10^{-1}	2.75×10^{-1}
MNIST $0/1$ (6 qubits)	Euclidean FM, matched ($n = 10$)	3.23×10^{-2}	5.57×10^{-2}	4.44×10^{-2}
MNIST $0/1$ (6 qubits)	IFM (ours)	4.44×10^{-3}	5.14×10^{-2}	5.64×10^{-2}
Spin-coherent (10 qubits)	Euclidean FM (as reported)	1.22×10^{-2}	1.43×10^{-3}	8.03×10^{-2}
Spin-coherent (10 qubits)	Euclidean FM, matched ($n = 10$)	8.30×10^{-3}	1.80×10^{-3}	4.32×10^{-2}
Spin-coherent (10 qubits)	IFM (ours)	3.81×10^{-4}	1.79×10^{-3}	4.44×10^{-2}
Bimodal (10 qubits)	Euclidean FM (as reported)	1.22×10^{-2}	4.60×10^{-2}	4.60×10^{-2}
Bimodal (10 qubits)	Euclidean FM, matched ($n = 10$)	8.38×10^{-3}	6.54×10^{-3}	1.21×10^{-2}
Bimodal (10 qubits)	IFM (ours)	8.38×10^{-4}	7.99×10^{-3}	1.64×10^{-2}
Trimodal (10 qubits)	Euclidean FM (as reported)	1.27×10^{-2}	5.97×10^{-2}	5.81×10^{-2}
Trimodal (10 qubits)	Euclidean FM, matched ($n = 10$)	8.31×10^{-3}	3.98×10^{-3}	6.00×10^{-3}
Trimodal (10 qubits)	IFM (ours)	7.08×10^{-4}	5.27×10^{-3}	3.95×10^{-3}

Table 16: Main benchmarks rerun with the raw and the exactly $U(1)$-equivariant (canonical-gauge) parameterization, five seeds, mean ± std, identical in every other respect. Lower is better.

	MMD ↓		Δ_{obs} ↓		Ent. W_1 ↓	
Benchmark	raw	equivariant	raw	equivariant	raw	equivariant
Single-cluster (2 qubits)	≈ 0	≈ 0	$(6.33 \pm 2.83) \times 10^{-4}$	$(6.86 \pm 3.13) \times 10^{-4}$	$(1.91 \pm 0.46) \times 10^{-3}$	$(1.93 \pm 0.46) \times 10^{-3}$
Single-cluster (6 qubits)	$(4.35 \pm 1.61) \times 10^{-4}$	$(4.57 \pm 1.64) \times 10^{-4}$	$(2.30 \pm 0.46) \times 10^{-3}$	$(2.51 \pm 0.44) \times 10^{-3}$	$(5.40 \pm 0.82) \times 10^{-3}$	$(5.69 \pm 0.90) \times 10^{-3}$
Single-cluster (10 qubits)	$(6.34 \pm 0.62) \times 10^{-4}$	$(6.36 \pm 0.54) \times 10^{-4}$	$(1.71 \pm 0.10) \times 10^{-3}$	$(1.92 \pm 0.23) \times 10^{-3}$	$(1.92 \pm 0.21) \times 10^{-3}$	$(1.73 \pm 0.58) \times 10^{-3}$
Single-cluster (12 qubits)	$(3.55 \pm 0.12) \times 10^{-4}$	$(3.72 \pm 0.13) \times 10^{-4}$	$(1.09 \pm 0.10) \times 10^{-3}$	$(9.08 \pm 0.22) \times 10^{-4}$	$(7.22 \pm 1.61) \times 10^{-4}$	$(6.12 \pm 1.03) \times 10^{-4}$
Bimodal (10 qubits)	$(6.36 \pm 1.41) \times 10^{-4}$	$(6.11 \pm 1.53) \times 10^{-4}$	$(8.08 \pm 2.79) \times 10^{-3}$	$(6.81 \pm 2.43) \times 10^{-3}$	$(1.76 \pm 0.11) \times 10^{-2}$	$(1.46 \pm 0.25) \times 10^{-2}$
Trimodal (10 qubits)	$(5.45 \pm 0.64) \times 10^{-4}$	$(6.85 \pm 0.92) \times 10^{-4}$	$(5.60 \pm 1.04) \times 10^{-3}$	$(6.64 \pm 3.01) \times 10^{-3}$	$(5.51 \pm 1.52) \times 10^{-3}$	$(6.36 \pm 1.01) \times 10^{-3}$
Spin-coherent (10 qubits)	$(1.78 \pm 0.62) \times 10^{-4}$	$(3.35 \pm 0.65) \times 10^{-4}$	$(1.65 \pm 0.31) \times 10^{-3}$	$(5.24 \pm 0.50) \times 10^{-3}$	$(4.35 \pm 0.13) \times 10^{-2}$	$(4.40 \pm 0.19) \times 10^{-2}$
MNIST $0/1$ (6 qubits)	$(3.57 \pm 1.03) \times 10^{-3}$	$(5.64 \pm 0.79) \times 10^{-3}$	$(5.14 \pm 0.48) \times 10^{-2}$	$(5.19 \pm 0.34) \times 10^{-2}$	$(9.05 \pm 4.12) \times 10^{-2}$	$(1.30 \pm 0.17) \times 10^{-1}$

selection, sampling, and evaluation and use the same five seeds, differing only in the gauge wrapper, so the two columns of Table 16 form a like-for-like test of the correction. They come from a fresh harness with a different seed set from Tables 2 and 3 and should therefore be compared with each other rather than with the reported values.

Enforcing exact equivariance leaves the empirical conclusions intact. Five of the eight benchmarks are unchanged within one standard deviation or slightly improved, and the remaining three move by less than a factor of two. No ordering against the ambient baseline reverses: measured against the Euclidean-FM values of Tables 2 and 3, the equivariant implementation remains between 3.7x and 36.4x lower in MMD on the nonzero rows. The geometric conclusions of the paper therefore hold for a field that is exactly well defined on the quotient, and not only for a particular ambient lift.

D.16 Near-orthogonal couplings and a Haar base

Phase alignment is undefined at exactly orthogonal endpoint pairs, and Proposition 9 accordingly assumes nonzero overlap. Two distinct concerns follow from this, and it is worth keeping them apart.

The first is whether near-orthogonality makes the conditional path or its velocity degenerate. It does not. Writing $c = |\langle \psi_0, \psi_1 \rangle|$ and using that the aligned overlap $\langle \psi_0, \bar{\psi}_1 \rangle = c$ is real and nonnegative,

$$\|\tilde{\psi}_t\|_2^2 = (1-t)^2 + t^2 + 2t(1-t)c \;\geq\; \tfrac{1}{2} \tag{56}$$

uniformly in $c \geq 0$ and $t \in [0,1]$, so the normalization never approaches a singularity; combined with $\|\bar{\psi}_1 - \psi_0\|_2 \leq 2$ and the non-expansiveness of Π_{ψ_t} this gives the uniform bound $\|u_t(\psi_t \mid \psi_0, \psi_1)\| \leq 2\sqrt{2}$. Neither the denominator nor the conditional velocity blows up as $c \to 0$.

The second concern is the constant in the stability bound, which Remark 7 handles analytically: an overlap margin on the coupling support gives a dimension-uniform Lipschitz constant. We complement that with a direct empirical test at the other extreme, replacing the clustered base by

Table 17: Haar base (predominantly near-orthogonal endpoint pairs), five seeds, mean ± std. Lower is better.

	MMD ↓		Δ_{obs} ↓		Ent. W_1 ↓	
n	raw	equivariant	raw	equivariant	raw	equivariant
6	$(1.27 \pm 0.27) \times 10^{-2}$	$(3.02 \pm 1.28) \times 10^{-3}$	$(1.06 \pm 0.12) \times 10^{-1}$	$(3.94 \pm 1.50) \times 10^{-2}$	$(2.48 \pm 0.10) \times 10^{-1}$	$(7.73 \pm 1.28) \times 10^{-2}$
10	$(6.79 \pm 1.49) \times 10^{-4}$	$(3.38 \pm 0.12) \times 10^{-3}$	$(2.22 \pm 0.36) \times 10^{-2}$	$(5.15 \pm 0.13) \times 10^{-2}$	$(5.89 \pm 0.27) \times 10^{-2}$	$(5.28 \pm 0.10) \times 10^{-2}$

a Haar-random base under which almost every coupling is close to orthogonal at large d. Table 17 reports both parameterizations under the same training budget. Both train to completion without numerical divergence. At 6 qubits the equivariant field is markedly better on all three metrics; at 10 qubits the two are comparable, with the raw field ahead on MMD and Δ_{obs} and the equivariant field slightly ahead on Ent. W_1. These runs are reported as a stability check in a predominantly near-orthogonal regime, not as a claim that one base distribution is preferable to the other.

D.17 Factorial isolation of gauge fixing and path family

IFM changes two things at once relative to a generic intrinsic flow: it fixes the endpoint gauge by Pancharatnam alignment, and it transports along a normalized chord rather than a Fubini–Study geodesic. This appendix separates them with a full 2×2 factorial over {aligned, unaligned} × {chord, geodesic}, run on three benchmarks with three seeds each. Every other ingredient—network, optimizer, tangent projection, base and data pools, step count, deterministic sampler, checkpoint rule, and metrics—is held fixed, so each cell differs from its neighbours in exactly one factor.

The factorial is run in two gauge regimes, because on a quotient manifold the choice of representative is part of the problem specification. In the first, endpoint representatives are taken as the data and base samplers supply them, which is the setting of all other experiments in the paper. In the second, each endpoint is multiplied by an independent random global phase before the conditional path is built. This second operation is the identity on $\mathbb{CP}^{d-1}$: every sample denotes exactly the same physical pure state, and a method that is genuinely defined on the quotient must be indifferent to it.

Table 18 gives both panels, and they answer the question cleanly. Under randomized representatives, gauge fixing is the factor that matters: at fixed path family, alignment improves MMD by 1.8–3.4× on every benchmark, and improves Δ_{obs} by up to 12.7× (trimodal, chord family: 1.87×10^{-3} aligned versus 2.37×10^{-2} unaligned). The path family is not the factor: holding alignment fixed, chord and geodesic agree to within a small multiple throughout, and the aligned-chord and aligned-geodesic cells are statistically indistinguishable on the single-cluster benchmark. Under representatives supplied in a consistent gauge, aligned and unaligned cells coincide to within seed noise (MMD ratios between 0.95 and 1.00). That is the expected behaviour and the second half of the same statement: Pancharatnam alignment is exactly the operation that makes the construction representative-independent, so it is a no-op precisely when the representatives already agree, and decisive as soon as they do not. Taken together the two panels isolate quotient-aware gauge fixing—not path geometry and not the ambient implementation details controlled in Appendix D.14—as the ingredient responsible for IFM's behaviour on the quotient.

D.18 Circuit-generated ensembles and hardware realizability

The benchmark families of Appendix D.2 are built to expose specific geometric effects, which makes them controlled but undeniably synthetic. We therefore add an ensemble whose members are produced the way a device produces them: states $U(\theta)\,|0\rangle^{\otimes n}$ from a hardware-efficient ansatz of $L = 2$ layers, each a round of single-qubit $R_y R_z$ rotations followed by a linear CNOT chain, with θ drawn around three structured mode centres with per-parameter Gaussian noise of width $\epsilon = 0.1$. Training, evaluation, and network scale follow the protocol of Appendix D.1.

Table 20 reports the comparison against ambient Euclidean FM over three seeds. IFM improves on all three metrics at both sizes, and the margin widens with n: at 6 qubits the MMD gap is a factor of 2.6, while at 10 qubits the ambient baseline fails to fit the ensemble at all and the IFM MMD falls below the reporting precision. The comparison is deliberately posed on *structured* circuit families, which is the regime of practical interest: a trained variational circuit, or a preparation routine executed repeatedly, produces exactly this kind of concentrated parameter distribution rather than a diffuse

Table 18: 2×2 factorial over endpoint gauge fixing and conditional path family, three seeds, mean $\pm$ std, everything else held fixed. Top: endpoint representatives as supplied by the samplers. Bottom: each endpoint multiplied by an independent random global phase, which leaves every state unchanged as a point of $\mathbb{CP}^{d-1}$. Lower is better.

Benchmark	Path	Gauge	MMD $\downarrow$	$\Delta_{\text{obs}} \downarrow$	Ent. $W_1 \downarrow$
Representatives as supplied by the data and base samplers					
Single-cluster (6 qubits)	chord	aligned	$(2.38 \pm 1.25) \times 10^{-4}$	$(2.58 \pm 0.64) \times 10^{-3}$	$(5.90 \pm 1.19) \times 10^{-3}$
Single-cluster (6 qubits)	chord	unaligned	$(2.32 \pm 1.39) \times 10^{-4}$	$(2.58 \pm 0.77) \times 10^{-3}$	$(5.53 \pm 1.04) \times 10^{-3}$
Single-cluster (6 qubits)	geodesic	aligned	$(2.54 \pm 1.54) \times 10^{-4}$	$(2.57 \pm 0.66) \times 10^{-3}$	$(5.74 \pm 1.17) \times 10^{-3}$
Single-cluster (6 qubits)	geodesic	unaligned	$(2.53 \pm 1.70) \times 10^{-4}$	$(2.59 \pm 0.77) \times 10^{-3}$	$(5.81 \pm 0.68) \times 10^{-3}$
Bimodal (10 qubits)	chord	aligned	$(8.67 \pm 0.24) \times 10^{-4}$	$(1.10 \pm 0.23) \times 10^{-2}$	$(1.66 \pm 0.12) \times 10^{-2}$
Bimodal (10 qubits)	chord	unaligned	$(8.28 \pm 0.29) \times 10^{-4}$	$(8.55 \pm 2.41) \times 10^{-3}$	$(1.61 \pm 0.18) \times 10^{-2}$
Bimodal (10 qubits)	geodesic	aligned	$(8.91 \pm 0.73) \times 10^{-4}$	$(1.18 \pm 0.13) \times 10^{-2}$	$(1.56 \pm 0.26) \times 10^{-2}$
Bimodal (10 qubits)	geodesic	unaligned	$(8.50 \pm 0.75) \times 10^{-4}$	$(9.82 \pm 2.05) \times 10^{-3}$	$(1.47 \pm 0.09) \times 10^{-2}$
Trimodal (10 qubits)	chord	aligned	$(7.47 \pm 0.59) \times 10^{-4}$	$(6.36 \pm 2.73) \times 10^{-3}$	$(4.86 \pm 1.33) \times 10^{-3}$
Trimodal (10 qubits)	chord	unaligned	$(7.44 \pm 0.77) \times 10^{-4}$	$(3.99 \pm 1.91) \times 10^{-3}$	$(3.55 \pm 1.48) \times 10^{-3}$
Trimodal (10 qubits)	geodesic	aligned	$(7.70 \pm 0.94) \times 10^{-4}$	$(6.94 \pm 1.27) \times 10^{-3}$	$(4.31 \pm 1.36) \times 10^{-3}$
Trimodal (10 qubits)	geodesic	unaligned	$(7.68 \pm 1.21) \times 10^{-4}$	$(5.10 \pm 1.14) \times 10^{-3}$	$(3.00 \pm 0.45) \times 10^{-3}$
Endpoint representatives randomized by an independent global phase					
Single-cluster (6 qubits)	chord	aligned	$(7.59 \pm 4.46) \times 10^{-4}$	$(3.01 \pm 0.58) \times 10^{-3}$	$(5.84 \pm 0.98) \times 10^{-3}$
Single-cluster (6 qubits)	chord	unaligned	$(1.84 \pm 0.73) \times 10^{-3}$	$(4.51 \pm 1.64) \times 10^{-3}$	$(8.36 \pm 2.98) \times 10^{-3}$
Single-cluster (6 qubits)	geodesic	aligned	$(7.61 \pm 4.27) \times 10^{-4}$	$(3.10 \pm 0.86) \times 10^{-3}$	$(5.60 \pm 0.83) \times 10^{-3}$
Single-cluster (6 qubits)	geodesic	unaligned	$(2.26 \pm 0.40) \times 10^{-3}$	$(5.33 \pm 1.03) \times 10^{-3}$	$(1.14 \pm 0.08) \times 10^{-2}$
Bimodal (10 qubits)	chord	aligned	$(8.85 \pm 0.29) \times 10^{-4}$	$(5.41 \pm 1.64) \times 10^{-3}$	$(5.26 \pm 3.38) \times 10^{-3}$
Bimodal (10 qubits)	chord	unaligned	$(3.00 \pm 0.02) \times 10^{-3}$	$(4.01 \pm 0.07) \times 10^{-2}$	$(1.69 \pm 0.39) \times 10^{-3}$
Bimodal (10 qubits)	geodesic	aligned	$(1.20 \pm 0.19) \times 10^{-3}$	$(1.13 \pm 0.34) \times 10^{-2}$	$(3.50 \pm 1.34) \times 10^{-3}$
Bimodal (10 qubits)	geodesic	unaligned	$(2.36 \pm 0.19) \times 10^{-3}$	$(3.13 \pm 0.23) \times 10^{-2}$	$(1.94 \pm 0.49) \times 10^{-3}$
Trimodal (10 qubits)	chord	aligned	$(9.38 \pm 0.62) \times 10^{-4}$	$(1.87 \pm 0.30) \times 10^{-3}$	$(4.09 \pm 0.23) \times 10^{-3}$
Trimodal (10 qubits)	chord	unaligned	$(1.87 \pm 0.10) \times 10^{-3}$	$(2.37 \pm 0.07) \times 10^{-2}$	$(4.08 \pm 1.34) \times 10^{-3}$
Trimodal (10 qubits)	geodesic	aligned	$(1.02 \pm 0.02) \times 10^{-3}$	$(6.12 \pm 1.42) \times 10^{-3}$	$(3.35 \pm 0.25) \times 10^{-3}$
Trimodal (10 qubits)	geodesic	unaligned	$(1.84 \pm 0.04) \times 10^{-3}$	$(2.42 \pm 0.15) \times 10^{-2}$	$(4.04 \pm 1.05) \times 10^{-3}$

Table 19: Circuit ensemble at 6 qubits as the parameter noise ϵ widens and the three modes wash out, three seeds, mean $\pm$ std. Lower is better. The $\epsilon = 0.1$ block is an independent rerun of the corresponding rows of Table 20.

ϵ	Method	MMD $\downarrow$	$\Delta_{\text{obs}} \downarrow$	Ent. $W_1 \downarrow$
0.1	Euclidean FM	$(2.96 \pm 0.37) \times 10^{-2}$	$(5.22 \pm 0.03) \times 10^{-2}$	$(6.78 \pm 0.03) \times 10^{-1}$
0.1	IFM (ours)	$(1.15 \pm 0.22) \times 10^{-2}$	$(3.91 \pm 1.20) \times 10^{-2}$	$(3.82 \pm 0.17) \times 10^{-1}$
0.3	Euclidean FM	$(8.90 \pm 1.06) \times 10^{-3}$	$(3.67 \pm 0.25) \times 10^{-2}$	$(6.69 \pm 0.07) \times 10^{-1}$
0.3	IFM (ours)	$(4.96 \pm 1.97) \times 10^{-3}$	$(4.20 \pm 0.16) \times 10^{-2}$	$(7.05 \pm 0.16) \times 10^{-1}$

one. Table 19 locates the boundary of the effect by widening the parameter noise to $\epsilon = 0.3$, which washes the three modes into a single diffuse cloud. IFM keeps the MMD lead there ($1.8\times$), while the ambient baseline becomes better on Δ_{obs} and Ent. W_1. The geometric advantage is therefore tied to coherent mode structure rather than to the circuit parameterization as such, which is what the projective picture predicts: with no mode organization there is no phase relation between endpoints for the quotient-aware path to exploit.

Hardware realizability. To check that this family is realizable rather than merely circuit-shaped, we executed 48 independently sampled 6-qubit preparation circuits from the same ensemble on IBM's 156-qubit Heron device `ibm_kawasaki`, at 4096 shots each, with the default transpiler layout at optimization level 1 and no readout-error mitigation, and compared the measured single-qubit Z_i and two-qubit Z_iZ_j expectation values (21 observables per circuit) against the ideal state vectors. The device values track the ideal ones closely: Pearson $r = 0.992$ across all per-circuit observables, and a mean absolute deviation of 0.015 between the ensemble-averaged device and ideal observables. The experiment establishes that the target family is device-realizable at the observable level; training itself uses the simulated state vectors, as elsewhere in the paper.

Table 20: Circuit-generated ensemble from a hardware-efficient ansatz ($L = 2$, three structured modes, $\epsilon = 0.1$), three seeds, mean $\pm$ std. Lower is better.

n	Method	MMD $\downarrow$	$\Delta_{\mathrm{obs}} \downarrow$	Ent. $W_1 \downarrow$
6	Euclidean FM	$(2.96 \pm 0.37) \times 10^{-2}$	$(5.22 \pm 0.03) \times 10^{-2}$	$(6.78 \pm 0.03) \times 10^{-1}$
6	IFM (ours)	$(1.15 \pm 0.22) \times 10^{-2}$	$(3.91 \pm 1.20) \times 10^{-2}$	$(3.82 \pm 0.17) \times 10^{-1}$
10	Euclidean FM	$(1.86 \pm 0.02) \times 10^{-1}$	$(1.12 \pm 0.02) \times 10^{-1}$	$(2.03 \pm 0.00) \times 10^{0}$
10	IFM (ours)	$(5.33 \pm 7.53) \times 10^{-5}$	$(1.43 \pm 0.07) \times 10^{-2}$	$(1.05 \pm 0.01) \times 10^{0}$

E Limitations and Future Work

IFM is developed for the dense-simulable regime, where full state vectors are already available: state-preparation verification, variational and QML priors, ground-state families, and simulation-side benchmarking of quantum generative algorithms. Within that regime the intrinsic machinery is essentially free on top of the ambient representation (Appendix B.2).

The benchmark suite is designed to isolate geometry, phase alignment, multimodality, and high-dimensional projective effects under controlled conditions, extending intrinsic quantum generative comparisons beyond the common 6-qubit setting to structured 10-qubit and selected 12-qubit tests. Natural empirical extensions include ensembles obtained from tomography pipelines, variational-circuit trajectories, simulated or experimental quantum devices, task-level QML workloads, larger many-body Hamiltonian families, and compressed representations such as tensor networks or shadows. The framework also focuses on pure states on $\mathbb{CP}^{d-1}$; mixed-state, density-matrix, and hybrid pure-state/density-matrix transport are promising directions for extending the same intrinsic-flow viewpoint. On the theory side, the endpoint and stability results suggest several next steps, including sharper statistical analyses, dimension-aware bounds, and connections between intrinsic flow matching, probability currents, and control-theoretic transport.

AI use statement

Parts of this paper's text, code and analysis were produced with the assistance of generative AI. The authors take full responsibility for the final content of this paper, including any text produced with the assistance of generative AI.