STOCHASTIC INTERPOLANTS VIII

ACE: Path Existence, Marginal Path Collapse, Adaptive Exponents, and Generator-Level Composition

Ziseok Lee

April 30, 2026

Document VII showed how FKC samples a composed path:

$$p_t^\star(x) = \frac{h_t(x)}{Z_t}, \qquad h_t(x) = \prod_i q_t^{(i)}(x)^{\gamma_i}.$$

But it assumed:

$$Z_t = \int h_t(x)\,\mathrm{d}x < \infty \qquad \forall t.$$

This document studies when that assumption fails.

Main topics:

Marginal Path Collapse	:	$Z_t = \infty$ at intermediate times,
Path Existence Criterion	:	a checkable integrability condition,
ACE	:	adaptive exponents guaranteeing valid paths,
Bridge to Semester II	:	Markov generators and path composition.

Let expert i have state space:

$$\mathbb{R}^{d_i}.$$

A one-sided stochastic-interpolant path has:

$$X_t^{(i)} = \alpha_t^{(i)} X_0^{(i)} + \beta_t^{(i)} X_1^{(i)},$$

where:

$$X_0^{(i)} \sim \mathcal{N}(0, I_{d_i}), \qquad X_1^{(i)} \sim q_1^{(i)}.$$

The time-t law is:

$$q_t^{(i)}.$$

Heterogeneity means:

$$\alpha_t^{(i)} \neq \alpha_t^{(j)}$$

and possibly:

$$d_i \neq d_j.$$

Thus each expert may have its own geometry, dimension, and noise schedule.

EMBEDDING INTO A COMMON AMBIENT SPACE

Let:

$$d = \max_i d_i.$$

For each expert i, choose:

$$I_i \subseteq \{1, \dots, d\}, \qquad |I_i| = d_i.$$

Projection:

$$\pi_i : \mathbb{R}^d \to \mathbb{R}^{d_i}.$$

Embedding:

$$\iota_i : \mathbb{R}^{d_i} \to \mathbb{R}^d.$$

Lifted density:

$$\tilde{q}_t^{(i)}(x) = q_t^{(i)}(\pi_i x).$$

Lifted score and velocity:

$$\tilde{s}_t^{(i)}(x) = \iota_i s_t^{(i)}(\pi_i x), \qquad \tilde{v}_t^{(i)}(x) = \iota_i v_t^{(i)}(\pi_i x).$$

RATIO-OF-DENSITIES PATH

Let $\gamma_i(t) \in \mathbb{R}$ be possibly time-dependent exponents.

Define:

$$\boxed{h_t(x) = \prod_{i=1}^{n} \left(\tilde{q}_t^{(i)}(x)\right)^{\gamma_i(t)}.}$$

If:

$$Z_t = \int_{\mathbb{R}^d} h_t(x)\,\mathrm{d}x < \infty,$$

then:

$$\boxed{p_t^\star(x) = \frac{h_t(x)}{Z_t}.}$$

Score:

$$\boxed{s_t^\star(x) = \nabla \log p_t^\star(x) = \sum_i \gamma_i(t)\tilde{s}_t^{(i)}(x).}$$

Negative exponents encode denominator experts.

PATH EXISTENCE PROPERTY

The family $\{h_t\}$ has the path existence property on $[0, t_{\text{end}}]$ if:

$$0 < Z_t = \int_{\mathbb{R}^d} h_t(x)\,\mathrm{d}x < \infty \qquad \forall t \in [0, t_{\text{end}}].$$

Then:

$$p_t^\star = h_t / Z_t$$

is a valid probability path.

This is not cosmetic.

If $Z_t = \infty$, then:

$$p_t^\star$$

does not exist, and:

$$s_t^\star = \nabla \log h_t$$

is not the score of any normalized distribution.

FKC and heuristic mixed-score samplers lose their intended target.

SUPPORT INCLUSION

When exponents may be negative, boundary singularities are possible.

Let:

$$I_+ = \{i : \gamma_i(t) > 0\}, \qquad I_- = \{i : \gamma_i(t) < 0\}.$$

A minimal support requirement is:

$$\operatorname{supp}\left(\prod_{i \in I_+} \tilde{q}_t^{(i)}\right) \subseteq \operatorname{supp}\left(\prod_{i \in I_-} \tilde{q}_t^{(i)}\right).$$

Otherwise the denominator can vanish where the numerator is positive.

In Gaussian-based interpolants with full-support Gaussian priors, the dominant issue is often not boundary support but tail integrability.

Marginal Path Collapse occurs when:

$$Z_t = \infty$$

for some intermediate t, even though:

$$Z_0 < \infty, \qquad Z_1 < \infty.$$

Consequences:

$p_t^\star$ is undefined,

$s_t^\star$ is not a probability score,

FKC weights no longer correct toward the intended path.

The sampler may still produce finite numerical updates.

But it follows another path:

$$p'_t \neq p_t^\star.$$

Consider:

$$h_t(x) = \frac{q_t^{(1)}(x) q_t^{(2)}(x)}{q_t^{(3)}(x) q_t^{(4)}(x)}.$$

Let each $q_t^{(i)}$ be Gaussian:

$$q_t^{(i)} = \mathcal{N}(0, \sigma_i^2(t) I).$$

Then:

$$h_t(x) = C_t \exp\left[-\frac{1}{2} \left(\frac{1}{\sigma_1^2(t)} + \frac{1}{\sigma_2^2(t)} - \frac{1}{\sigma_3^2(t)} - \frac{1}{\sigma_4^2(t)} \right) \|x\|^2 \right].$$

Define effective precision:

$$P_{\text{eff}}(t) = \frac{1}{\sigma_1^2(t)} + \frac{1}{\sigma_2^2(t)} - \frac{1}{\sigma_3^2(t)} - \frac{1}{\sigma_4^2(t)}.$$

If:

$$P_{\text{eff}}(t) < 0,$$

then h_t grows at infinity and is not integrable.

ENDPOINTS DO NOT GUARANTEE THE PATH

Even if:

$$P_{\text{eff}}(0) > 0, \qquad P_{\text{eff}}(1) > 0,$$

there may exist:

$$t^\star \in (0,1)$$

such that:

$$P_{\text{eff}}(t^\star) < 0.$$

Then:

$$Z_{t^\star} = \infty.$$

Interpretation:

denominator experts contract faster than numerator experts,

so the ratio becomes explosive in the tails.

This is the simplest instance of marginal path collapse.

COMPACT-TARGET SETUP

ACE studies the practically important case:

$$X_t^{(i)} = \alpha_t^{(i)} X_0^{(i)} + \beta_t^{(i)} X_1^{(i)},$$

where:

$$X_0^{(i)} \sim \mathcal{N}(0, I),$$

and:

$$X_1^{(i)}$$

has compact support.

Assume:

$$\mathrm{supp}(q_1^{(i)}) \subseteq B(0, R_i).$$

Then:

$$q_t^{(i)}(x) = \int \frac{\exp\left(-\frac{\|x-\beta_t^{(i)} y\|^2}{2(\alpha_t^{(i)})^2}\right)}{(2\pi(\alpha_t^{(i)})^2)^{d_i/2}} q_1^{(i)}(y)\,\mathrm{d}y.$$

The tails are Gaussian with scale $\alpha_t^{(i)}$.

Because:

$$\|y\| \le R_i,$$

we have:

$$\|x - \beta_t y\|^2 \ge \|x\|^2 - 2|\beta_t| R_i \|x\|.$$

Thus:

$$q_t^{(i)}(x) \le C_i(t) \exp\left[-\frac{\|x\|^2}{2(\alpha_t^{(i)})^2} + \frac{|\beta_t^{(i)}| R_i}{(\alpha_t^{(i)})^2} \|x\| \right].$$

Similarly, lower envelopes hold along appropriate rays.

Hence the quadratic tail coefficient of h_t is determined by:

$$\sum_i \frac{\gamma_i(t)}{(\alpha_t^{(i)})^2}.$$

For each coordinate $k \in \{1, \dots, d\}$, define:

$$\boxed{C_k(t) = \sum_{i:\, k \in I_i} \frac{\gamma_i(t)}{(\alpha_t^{(i)})^2}.}$$

Define:

$$C(t) = \min_k C_k(t).$$

Path Existence Criterion:
If:

$$h_1 \in L^1, \qquad C_k(t) > 0 \quad \forall k, \quad \forall t \in [0, t_{\text{end}}],$$

then:

$$h_t \in L^1 \qquad \forall t \in [0, t_{\text{end}}].$$

Conversely, if:

$$C_{k^\star}(t^\star) < 0$$

for some coordinate and time, then:

$$h_{t^\star} \notin L^1.$$

For each lifted expert:

$$\tilde{q}_t^{(i)}(x) = q_t^{(i)}(\pi_i x).$$

Using compact support:

$$\tilde{q}_t^{(i)}(x) \leq C_i \exp\left[-\frac{\|\pi_i x\|^2}{2(\alpha_t^{(i)})^2} + B_i(t)\|\pi_i x\|\right].$$

Raise to exponent $\gamma_i(t)$ and multiply over i.

The aggregate quadratic term is:

$$-\frac{1}{2}\sum_i \frac{\gamma_i(t)}{(\alpha_t^{(i)})^2}\|\pi_i x\|^2.$$

Since:

$$\|\pi_i x\|^2 = \sum_{k\in I_i} x_k^2,$$

this equals:

$$-\frac{1}{2}\sum_{k=1}^{d} C_k(t) x_k^2.$$

PROOF II: SUFFICIENCY

If:

$$C_k(t) > 0 \quad \forall k,$$

then:

$$\Lambda_t = \mathrm{diag}(C_1(t), \ldots, C_d(t))$$

is positive definite.

The product h_t is bounded by:

$$h_t(x) \leq A_t \exp\left[-\frac{1}{2} x^\top \Lambda_t x + B_t \|x\|\right].$$

Complete the square:

$$-\frac{1}{2}\lambda_{\min}\|x\|^2 + B_t\|x\| \leq -\frac{\lambda_{\min}}{4}\|x\|^2 + C_t.$$

Thus:

$$h_t(x) \leq A'_t e^{-c_t\|x\|^2}.$$

The Gaussian upper bound is integrable, hence:

$$h_t \in L^1.$$

PROOF III: COLLAPSE WHEN $C_k < 0$

Suppose:

$$C_{k^\star}(t^\star) < 0.$$

Consider the ray:

$$x = re_{k^\star}, \qquad r \to \infty.$$

Using lower tail bounds for numerator factors and upper tail bounds for denominator factors, one obtains:

$$h_{t^\star}(re_{k^\star}) \geq A \exp\left[+\frac{|C_{k^\star}(t^\star)|}{4} r^2 - Br\right]$$

for all sufficiently large r.

Then:

$$\int_R^\infty h_{t^\star}(re_{k^\star})\,\mathrm{d}r = \infty.$$

Thus:

$$h_{t^\star} \notin L^1.$$

WHAT ABOUT $C_k(t) = 0$?

The clean criterion states:

$$C_k(t) > 0 \Rightarrow \text{integrability},$$
$$C_k(t) < 0 \Rightarrow \text{non-integrability}.$$

The boundary case:

$$C_k(t) = 0$$

is delicate.

Then quadratic tails cancel and integrability depends on lower-order terms:

$$O(\|x\|), \qquad O(1), \qquad \text{support geometry}.$$

For sampling, ACE keeps:

$$C(t) \geq \delta > 0$$

on the discretization grid.

This avoids near-collapse as well as exact collapse.

If:

$$C(t) = \min_k C_k(t) > 0,$$

the same proof gives a sub-Gaussian tail:

$$p_t^\star(\|X\| > R) \lesssim \exp\left[-cC(t)R^2 + B_t R\right].$$

Therefore the $(1-\epsilon)$-quantile radius satisfies:

$$\boxed{R_t(\epsilon) = O\left(\frac{1}{\sqrt{C(t)}}\right)}$$

up to logarithmic and linear-envelope constants.

Interpretation:

$$C(t) \downarrow 0 \quad \Rightarrow \quad \text{intermediate law becomes diffuse.}$$

Even if $C(t) > 0$, small $C(t)$ destabilizes particle weights.

From the PEC proof:

$$p_t^\star(x) \leq A_t \exp\left[-\frac{1}{2}C(t)\|x\|^2 + B_t\|x\|\right].$$

For:

$$R \geq 4B_t/C(t),$$

we have:

$$-\frac{1}{2}C(t)R^2 + B_t R \leq -\frac{1}{4}C(t)R^2.$$

Then:

$$\mathbb{P}(\|X_t\| > R) \leq A_t \int_{\|x\|>R} e^{-\frac{1}{4}C(t)\|x\|^2}\,\mathrm{d}x.$$

Gaussian tail bounds give:

$$\mathbb{P}(\|X_t\| > R) \leq A_t' e^{-cC(t)R^2}.$$

Solving:

$$A_t' e^{-cC(t)R^2} \leq \epsilon$$

gives:

$$R = O(C(t)^{-1/2}).$$

ACE OBJECTIVE

We want to preserve endpoint composition:

$$p_0^\star, \qquad p_1^\star.$$

But we are free to change intermediate exponents:

$$\gamma_i(t) \quad t \in (0,1).$$

ACE constructs:

$$\tilde{\gamma}_i(t)$$

such that:

$$\tilde{\gamma}_i(0) = \gamma_i(0), \qquad \tilde{\gamma}_i(1) = \gamma_i(1),$$

while:

$$C(t) > 0 \quad \forall t \le t_{\text{end}}.$$

Thus ACE changes the path, not the endpoint task.

Choose one exponent index j.

Set:

$$\tilde{\gamma}_j(t) = \gamma_j(t) + b(t),$$

and:

$$\tilde{\gamma}_i(t) = \gamma_i(t) \qquad i \neq j.$$

Use a bump:

$$b(t) = B_1 Q(t) + B_2 L_\tau(t),$$

with:

$$Q(t) = t(1-t),$$

$$L_\tau(t) = \min(t, \tau(1-t)).$$

Endpoint preservation:

$$b(0) = b(1) = 0.$$

For differentiability, replace L_τ by a smoothed piecewise-linear bump if needed.

Assume:

$$C(0) > 0.$$

Then there exist differentiable functions:

$$\tilde{\gamma}_i(t)$$

such that:

$$\tilde{\gamma}_i(0) = \gamma_i(0), \qquad \tilde{\gamma}_i(1) = \gamma_i(1),$$

and:

$$\boxed{\tilde{C}(t) > 0 \qquad \forall t \in [0, t_{\text{end}}].}$$

A constructive choice is:

$$\tilde{\gamma}_j(t) = \gamma_j(t) + B_1 t(1-t) + B_2 L_\tau(t)$$

for sufficiently large:

$$B_1, B_2, \tau > 0.$$

PROOF I: CONTINUITY NEAR $t = 0$

Because:

$$C(0) > 0,$$

and $C(t)$ is continuous, there exists:

$$\delta_0 > 0$$

such that:

$$C(t) > 0 \qquad t \in [0, \delta_0].$$

Thus no correction is needed near 0.

The bump still satisfies:

$$b(0) = 0,$$

so the initial distribution is unchanged.

The main task is to ensure positivity on:

$$[\delta_0, t_{\text{end}}].$$

PROOF II: DOMINATE THE NEGATIVE PART

For $t \in [\delta_0, t_{\text{end}}]$, assume the chosen expert j covers the relevant coordinates. The corrected criterion is:

$$\tilde{C}_k(t) = C_k(t) + \frac{b(t)}{(\alpha_t^{(j)})^2} \qquad k \in I_j.$$

On the compact interval:

$$[\delta_0, t_{\text{end}}],$$

we have:

$$\alpha_t^{(j)} > 0,$$

and:

$$b(t) > 0$$

for $t \in (0, t_{\text{end}}]$. Thus by choosing B_1, B_2 sufficiently large:

$$\frac{b(t)}{(\alpha_t^{(j)})^2} > |C_k(t)| + \delta$$

for all relevant t, k. Hence:

$$\tilde{C}_k(t) \geq \delta > 0.$$

The bump satisfies:

$$b(0) = 0, \qquad b(1) = 0.$$

Therefore:

$$\tilde{\gamma}_j(0) = \gamma_j(0), \qquad \tilde{\gamma}_j(1) = \gamma_j(1).$$

For all other i:

$$\tilde{\gamma}_i = \gamma_i.$$

Thus:

$$p_0^\star \quad \text{and} \quad p_1^\star$$

are unchanged.

The only modification is:

$$p_t^\star \qquad t \in (0, 1).$$

This proves existence of a valid corrected path.

After correction:

$$p_t^{\mathrm{ACE}}(x) = \frac{\prod_i \tilde{q}_t^{(i)}(x)^{\tilde{\gamma}_i(t)}}{Z_t^{\mathrm{ACE}}}.$$

The score is:

$$\boxed{s_t^\star(x) = \sum_i \tilde{\gamma}_i(t)\tilde{s}_t^{(i)}(x).}$$

The target exists because:

$$\tilde{C}(t) > 0.$$

Next question:

How do we sample when $\dot{\tilde{\gamma}}_i(t) \neq 0$?

This requires a new term beyond standard FKC.

ACE SAMPLING DYNAMICS

Choose a base velocity:

$$v_t^\star.$$

Simulate:

$$\boxed{\mathrm{d}X_t = \left[v_t^\star(X_t) + \frac{\sigma_t^2}{2} s_t^\star(X_t)\right] \mathrm{d}t + \sigma_t \, \mathrm{d}W_t.}$$

where:

$$s_t^\star = \sum_i \gamma_i(t) \tilde{s}_t^{(i)}.$$

Weights satisfy:

$$\boxed{\mathrm{d}\log w_t = \left[F_t(X_t) + \sum_i \dot{\gamma}_i(t) \log \tilde{q}_t^{(i)}(X_t)\right] \mathrm{d}t.}$$

The extra ACE term is:

$$\sum_i \dot{\gamma}_i(t) \log \tilde{q}_t^{(i)}(X_t).$$

Let:

$$\ell_i(t,x) = \log \tilde{q}_t^{(i)}(x), \qquad \tilde{s}_i = \nabla \ell_i.$$

Each expert path satisfies:

$$\partial_t \tilde{q}_i = -\nabla \cdot (\tilde{v}_i \tilde{q}_i).$$

Equivalently:

$$\partial_t \ell_i = -\nabla \cdot \tilde{v}_i - \tilde{v}_i \cdot \tilde{s}_i.$$

Define:

$$\boxed{D_i(t,x) = -\nabla \cdot \tilde{v}_i(t,x) + (v_t^\star(x) - \tilde{v}_i(t,x)) \cdot \tilde{s}_i(t,x).}$$

The unnormalized target is:

$$h_t(x) = \prod_i \tilde{q}_i(t, x)^{\gamma_i(t)}.$$

Thus:

$$\log h_t = \sum_i \gamma_i(t) \ell_i(t, x).$$

Differentiate:

$$\partial_t \log h_t = \sum_i \gamma_i \partial_t \ell_i + \sum_i \dot{\gamma}_i \ell_i.$$

The continuity equation generated by $v_t^\star$ has log-density contribution:

$$-\nabla \cdot v_t^\star - v_t^\star \cdot s_t^\star.$$

The reweighting term is:

$$g_t = \partial_t \log h_t + \nabla \cdot v_t^\star + v_t^\star \cdot s_t^\star.$$

SIMPLIFYING g_t

Use:

$$\partial_t \ell_i = -\nabla \cdot \tilde{v}_i - \tilde{v}_i \cdot \tilde{s}_i,$$

and:

$$s_t^\star = \sum_i \gamma_i \tilde{s}_i.$$

Then:

$$g_t = \nabla \cdot v_t^\star + \sum_i \gamma_i \left[-\nabla \cdot \tilde{v}_i + (v_t^\star - \tilde{v}_i) \cdot \tilde{s}_i\right] + \sum_i \dot{\gamma}_i \ell_i.$$

Using D_i:

$$\boxed{g_t(x) = \nabla \cdot v_t^\star(x) + \sum_i \gamma_i(t) D_i(t,x) + \sum_i \dot{\gamma}_i(t) \log \tilde{q}_t^{(i)}(x).}$$

If $\dot{\gamma}_i = 0$, this reduces to FKC.

The normalized PDE uses:

$$\bar{g}_t = g_t - \mathbb{E}_{p_t^\star} g_t.$$

But particle weights may use uncentered:

$$\mathrm{d}A_t = g_t(X_t)\,\mathrm{d}t.$$

Reason:

$$g_t \quad \text{and} \quad g_t + c(t)$$

give the same normalized weights:

$$\frac{e^{A_T + \int c}}{\sum_j e^{A_T^{(j)} + \int c}} = \frac{e^{A_T}}{\sum_j e^{A_T^{(j)}}}.$$

Thus the unknown normalizing correction is unnecessary for self-normalized estimation.

TRACKING $\log \tilde{q}_i(t, X_t)$

ACE needs:

$$\ell_i(t, X_t) = \log \tilde{q}_t^{(i)}(X_t)$$

inside:

$$\sum_i \dot{\gamma}_i(t)\ell_i(t, X_t).$$

If exact log densities are unavailable, ACE tracks them by Itô.

Under:

$$\mathrm{d}X_t = \left(v_t^\star + \frac{\sigma_t^2}{2} s_t^\star\right) \mathrm{d}t + \sigma_t \, \mathrm{d}W_t,$$

Itô gives:

$$\mathrm{d}\ell_i(t, X_t) = \left[D_i + \frac{\sigma_t^2}{2} \left(s_t^\star \cdot \tilde{s}_i + \nabla \cdot \tilde{s}_i\right)\right] \mathrm{d}t + \sigma_t \tilde{s}_i \cdot \mathrm{d}W_t.$$

Itô:

$$\mathrm{d}\ell_i = \partial_t \ell_i \,\mathrm{d}t + \nabla \ell_i \cdot \mathrm{d}X_t + \frac{\sigma_t^2}{2} \Delta \ell_i \,\mathrm{d}t.$$

Use:

$$\nabla \ell_i = \tilde{s}_i, \qquad \Delta \ell_i = \nabla \cdot \tilde{s}_i.$$

Insert:

$$\mathrm{d}X_t = \left(v_t^\star + \frac{\sigma_t^2}{2} s_t^\star \right) \mathrm{d}t + \sigma_t \,\mathrm{d}W_t.$$

Then:

$$\mathrm{d}\ell_i = \left[\partial_t \ell_i + v_t^\star \cdot \tilde{s}_i + \frac{\sigma_t^2}{2} s_t^\star \cdot \tilde{s}_i + \frac{\sigma_t^2}{2} \nabla \cdot \tilde{s}_i \right] \mathrm{d}t + \sigma_t \tilde{s}_i \cdot \mathrm{d}W_t.$$

Since:

$$\partial_t \ell_i + v_t^\star \cdot \tilde{s}_i = D_i,$$

the formula follows.

ACE ALGORITHM

Initialize:

$$X_0^{(j)} \sim p_0^\star, \qquad w_0^{(j)} = 1/N.$$

For each time step:

$$s_t^\star(x) = \sum_i \gamma_i(t)\tilde{s}_t^{(i)}(x),$$

$$\mu_t(x) = v_t^\star(x) + \frac{\sigma_t^2}{2} s_t^\star(x).$$

Propagate:

$$X_{t+\Delta t}^{(j)} = X_t^{(j)} + \mu_t(X_t^{(j)})\Delta t + \sigma_t\sqrt{\Delta t}\xi_t^{(j)}.$$

Track:

$$\ell_i^{(j)}(t+\Delta t)$$

using the log-component SDE.

Update:

$$\log w^{(j)} \leftarrow \log w^{(j)} + g_t(X_t^{(j)})\Delta t.$$

Resample if:

$$\mathrm{ESS} < \tau N.$$

FKC AS A SPECIAL CASE OF ACE

If:

$$\gamma_i(t) = \gamma_i \quad \text{constant},$$

then:

$$\dot{\gamma}_i(t) = 0.$$

The ACE correction term:

$$\sum_i \dot{\gamma}_i(t)\ell_i(t, X_t)$$

vanishes.

Thus:

$$g_t = \nabla \cdot v_t^\star + \sum_i \gamma_i D_i.$$

This is the constant-exponent FKC weight.

Therefore:

$$\boxed{\mathrm{FKC} = \mathrm{ACE}\ \text{with constant exponents.}}$$

VALIDITY VS COMPUTABILITY

A mixed-score sampler may be numerically well-defined even when:

$$Z_t = \infty.$$

But if:

$$p_t^\star$$

does not exist, then:

$$s_t^\star$$

is not a score of any probability law. Therefore the sampler solves some well-posed SDE/FPE:

$$\partial_t p_t' = \text{FPE coefficients},$$

but:

$$p_t' \neq p_t^\star.$$

FKC correctors are justified only when:

$$p_t^\star = h_t / Z_t$$

exists at all sampling times. ACE repairs this by guaranteeing:

$$Z_t < \infty.$$

ACE uses the same particle logic as FKC:

propagate, weight, resample.

The difference is:

$$\gamma_i = \gamma_i(t)$$

and the log-density tracking term:

$$\dot{\gamma}_i(t) \log \tilde{q}_t^{(i)}(X_t).$$

Resampling duplicates high-weight particles and removes low-weight particles.

Interpretation:

ACE uses SMC to enforce the corrected valid path.

Without resampling, the method becomes a no-resampling weighted heuristic with larger variance.

ACE-LITE

ACE-lite keeps the adaptive exponent schedule:

$$\tilde{\gamma}_i(t)$$

but skips expensive resampling.

Thus it repairs:

path existence

but may not fully correct particle populations.

It is cheaper because:

no duplication/removal of particles.

But it can have:

higher variance and more bias.

Pedagogical distinction:

PEC repair is mathematical validity; SMC is statistical correction.

SYNTHETIC CHECKERBOARD COMPOSITION

A toy heterogeneous composition:

$$p^\star(X, Y \mid A, B) \propto \frac{p(X \mid A)\, p(X, Y \mid B)}{p(X)}.$$

This matches:

$$p_t^\star \propto q_t^{(1)} q_t^{(2)} / q_t^{(3)}.$$

The denominator exponent is negative:

$$\gamma_3 = -1.$$

With heterogeneous schedules:

$$\alpha_t^{(1)} \neq \alpha_t^{(2)} \neq \alpha_t^{(3)},$$

constant exponents can violate:

$$C(t) > 0.$$

ACE modifies one exponent with a bump so that:

$$C(t) \geq \delta > 0.$$

FLEXIBLE-POSE SCAFFOLD DECORATION

A scientific composition target:

$$p(M \mid T^{sc}, P) \propto \frac{p(M^{sc} \mid T(M^{sc}) = T^{sc})\, p(M \mid M \leftrightarrow P)}{p(M^{sc})}.$$

With guidance:

$$p_\omega(M) \propto p(M) \left(\frac{p(M \mid T^{sc}, P)}{p(M)} \right)^{\omega}.$$

This decomposes into multiple experts:

de-novo molecule model, conformer model, pocket-conditioned model.

The exponents include positive and negative terms.

Different experts often have different schedules, so:

$$C(t) < 0$$

can occur unless ACE corrects the path.

WHY SCHEDULE ALIGNMENT ALONE IS NOT ENOUGH

One possible response to heterogeneity is:

$$q_t^{(i)} \mapsto q_{\tau_i(t)}^{(i)}$$

to align schedules.

But reparameterization can distort sampling.

A nonuniform:

$$\tau_i(t)$$

can skip important regions of an expert's trajectory.

ACE instead preserves each expert's schedule and modifies exponents:

$$\gamma_i(t).$$

Thus ACE targets:

validity through exponent control, not forced schedule homogenization.

Even when:

$$C(t) > 0$$

for constant exponents, ACE can improve concentration by increasing:

$$C(t).$$

Since:

$$R_t(\epsilon) = O(C(t)^{-1/2}),$$

larger $C(t)$ gives tighter intermediate laws.

This may stabilize:

particle weights, sample quality, compositional constraints.

Thus ACE has two roles:

$$\boxed{\text{avoid nonexistence} \quad + \quad \text{control concentration.}}$$

Generator matching studies Markov processes through generators:

$$\mathcal{L}_t f(x) = \lim_{h \downarrow 0} \frac{\mathbb{E}[f(X_{t+h}) \mid X_t = x] - f(x)}{h}.$$

The marginal law satisfies the KFE:

$$\partial_t \langle p_t, f \rangle = \langle p_t, \mathcal{L}_t f \rangle.$$

In Euclidean space:

$$\mathcal{L}_t f = \nabla f \cdot u_t + \frac{1}{2} \nabla^2 f : \Sigma_t + \int [f(y) - f(x)] Q_t(\,\mathrm{d}y; x).$$

This includes:

flows, diffusions, jumps.

BRIDGE TO GENERATOR MATCHING II

Stochastic interpolants studied:

$$\partial_t \rho + \nabla \cdot (b\rho) = 0,$$

and:

$$\partial_t \rho + \nabla \cdot ((b + \varepsilon s)\rho) = \varepsilon \Delta \rho.$$

FKC/ACE add:

$$\partial_t p = \mathcal{L}_t^* p + \bar{g}_t p.$$

The new term:

$$\bar{g}_t p$$

is selection / branching / reweighting.

SMC and jump-process interpretations show that path composition naturally leads beyond pure ODE/SDE flows.

This motivates Semester II:

use arbitrary Markov generators as the primitive object.

Generator matching relies on linearity:

$$\partial_t \langle p_t, f \rangle = \langle p_t, \mathcal{L}_t f \rangle.$$

FKC uses a linear unnormalized equation:

$$\partial_t \tilde{p}_t = \mathcal{L}_t^* \tilde{p}_t + g_t \tilde{p}_t.$$

The normalized version is nonlinear:

$$\partial_t p_t = \mathcal{L}_t^* p_t + (g_t - \mathbb{E}_{p_t} g_t) p_t.$$

Particle systems restore linearity statistically:

unnormalized weights $\longrightarrow$ normalized empirical measure.

Thus FKC/ACE are a practical bridge from density-ratio composition to generator-level modeling.

Single stochastic interpolant $\Rightarrow$ q_t

Multiple trained paths $\Rightarrow$ $h_t = \prod_i q_t^{(i)\gamma_i(t)}$

Existence check $\Rightarrow$ $C_k(t) = \sum_{i:k\in I_i} \frac{\gamma_i(t)}{(\alpha_t^{(i)})^2}$

If valid $\Rightarrow$ FKC weighted SDE

If invalid $\Rightarrow$ ACE adaptive exponents

If general state space $\Rightarrow$ Markov generator viewpoint

PROOF EXERCISE 1: PEC

Prove the Path Existence Criterion in the compact-target setting.

Required steps:

1. Show:
$$q_t^{(i)}(x) \le C_i e^{-\|x\|^2/(2\alpha_i^2)+B_i\|x\|}.$$
2. Lift the bound to:
$$\tilde{q}_t^{(i)}(x) = q_t^{(i)}(\pi_i x).$$
3. Multiply powers:
$$h_t = \prod_i \tilde{q}_i^{\gamma_i}.$$
4. Identify:
$$-\frac{1}{2}\sum_k C_k(t)x_k^2.$$
5. Use Gaussian integrability when $C_k(t) > 0$.
6. Use a ray lower bound when some $C_k(t) < 0$.

Derive the ACE weight:

$$g_t = \nabla \cdot v_t^\star + \sum_i \gamma_i D_i + \sum_i \dot{\gamma}_i \ell_i .$$

Start with:

$$h_t = \prod_i \tilde{q}_i^{\gamma_i(t)}, \qquad \ell_i = \log \tilde{q}_i .$$

Use:

$$\partial_t \ell_i = -\nabla \cdot \tilde{v}_i - \tilde{v}_i \cdot \tilde{s}_i .$$

Compare the desired log-evolution:

$$\partial_t \log h_t$$

to the continuity evolution generated by $v_t^\star$:

$$-\nabla \cdot v_t^\star - v_t^\star \cdot s_t^\star .$$

The difference is the Feynman–Kac potential.

Show:

$$\mathrm{d}\ell_i(t, X_t) = \left[D_i + \frac{\sigma_t^2}{2}(s_t^\star \cdot \tilde{s}_i + \nabla \cdot \tilde{s}_i)\right] \mathrm{d}t + \sigma_t \tilde{s}_i \cdot \mathrm{d}W_t.$$

Use:

$$\mathrm{d}X_t = \left(v_t^\star + \frac{\sigma_t^2}{2} s_t^\star\right) \mathrm{d}t + \sigma_t \, \mathrm{d}W_t.$$

Apply Itô:

$$\mathrm{d}\ell_i = \partial_t \ell_i \, \mathrm{d}t + \tilde{s}_i \cdot \mathrm{d}X_t + \frac{\sigma_t^2}{2} \Delta \ell_i \, \mathrm{d}t.$$

Use:

$$\Delta \ell_i = \nabla \cdot \tilde{s}_i.$$

ACE adds a missing mathematical layer to FKC. FKC asks:

$$\text{given } p_t^\star, \text{ how do we sample it?}$$

ACE first asks:

$$\text{does } p_t^\star \text{ exist?}$$

The answer is controlled by:

$$C_k(t) = \sum_{i:k \in I_i} \frac{\gamma_i(t)}{(\alpha_t^{(i)})^2}.$$

If the path collapses:

$$C(t) < 0,$$

ACE constructs:

$$\tilde{\gamma}_i(t)$$

with the same endpoints but:

$$\tilde{C}(t) > 0.$$

Then ACE samples by Feynman–Kac weights with an extra term:

$$\sum_i \dot{\gamma}_i(t) \log \tilde{q}_t^{(i)}(X_t).$$

BRIDGE TO SEMESTER II

Semester I began with:

one stochastic interpolant path in $\mathbb{R}^d$.

It ends with:

composition of multiple paths using weighted Markov dynamics.

The new mathematical objects are:

Feynman–Kac semigroups, selection potentials, SMC particle systems, jump-process correctors.

Semester II abstracts all of this into:

generative modeling by arbitrary Markov generators.

Thus FKC and ACE are the natural bridge:

stochastic interpolants $\longrightarrow$ generator matching and general state spaces.