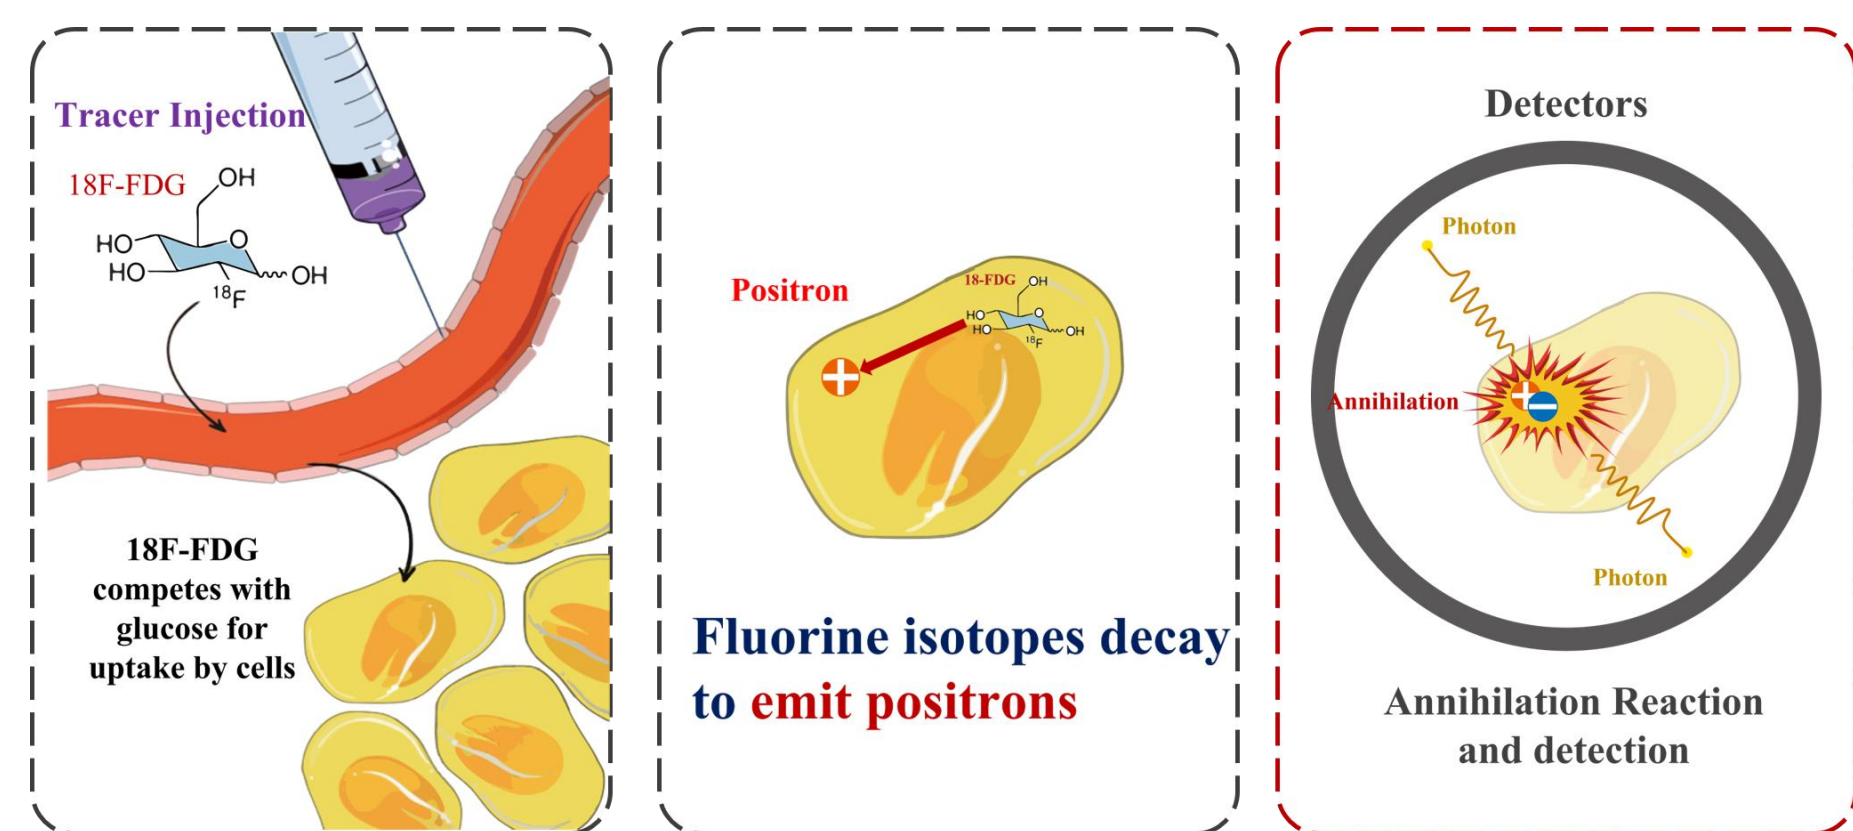


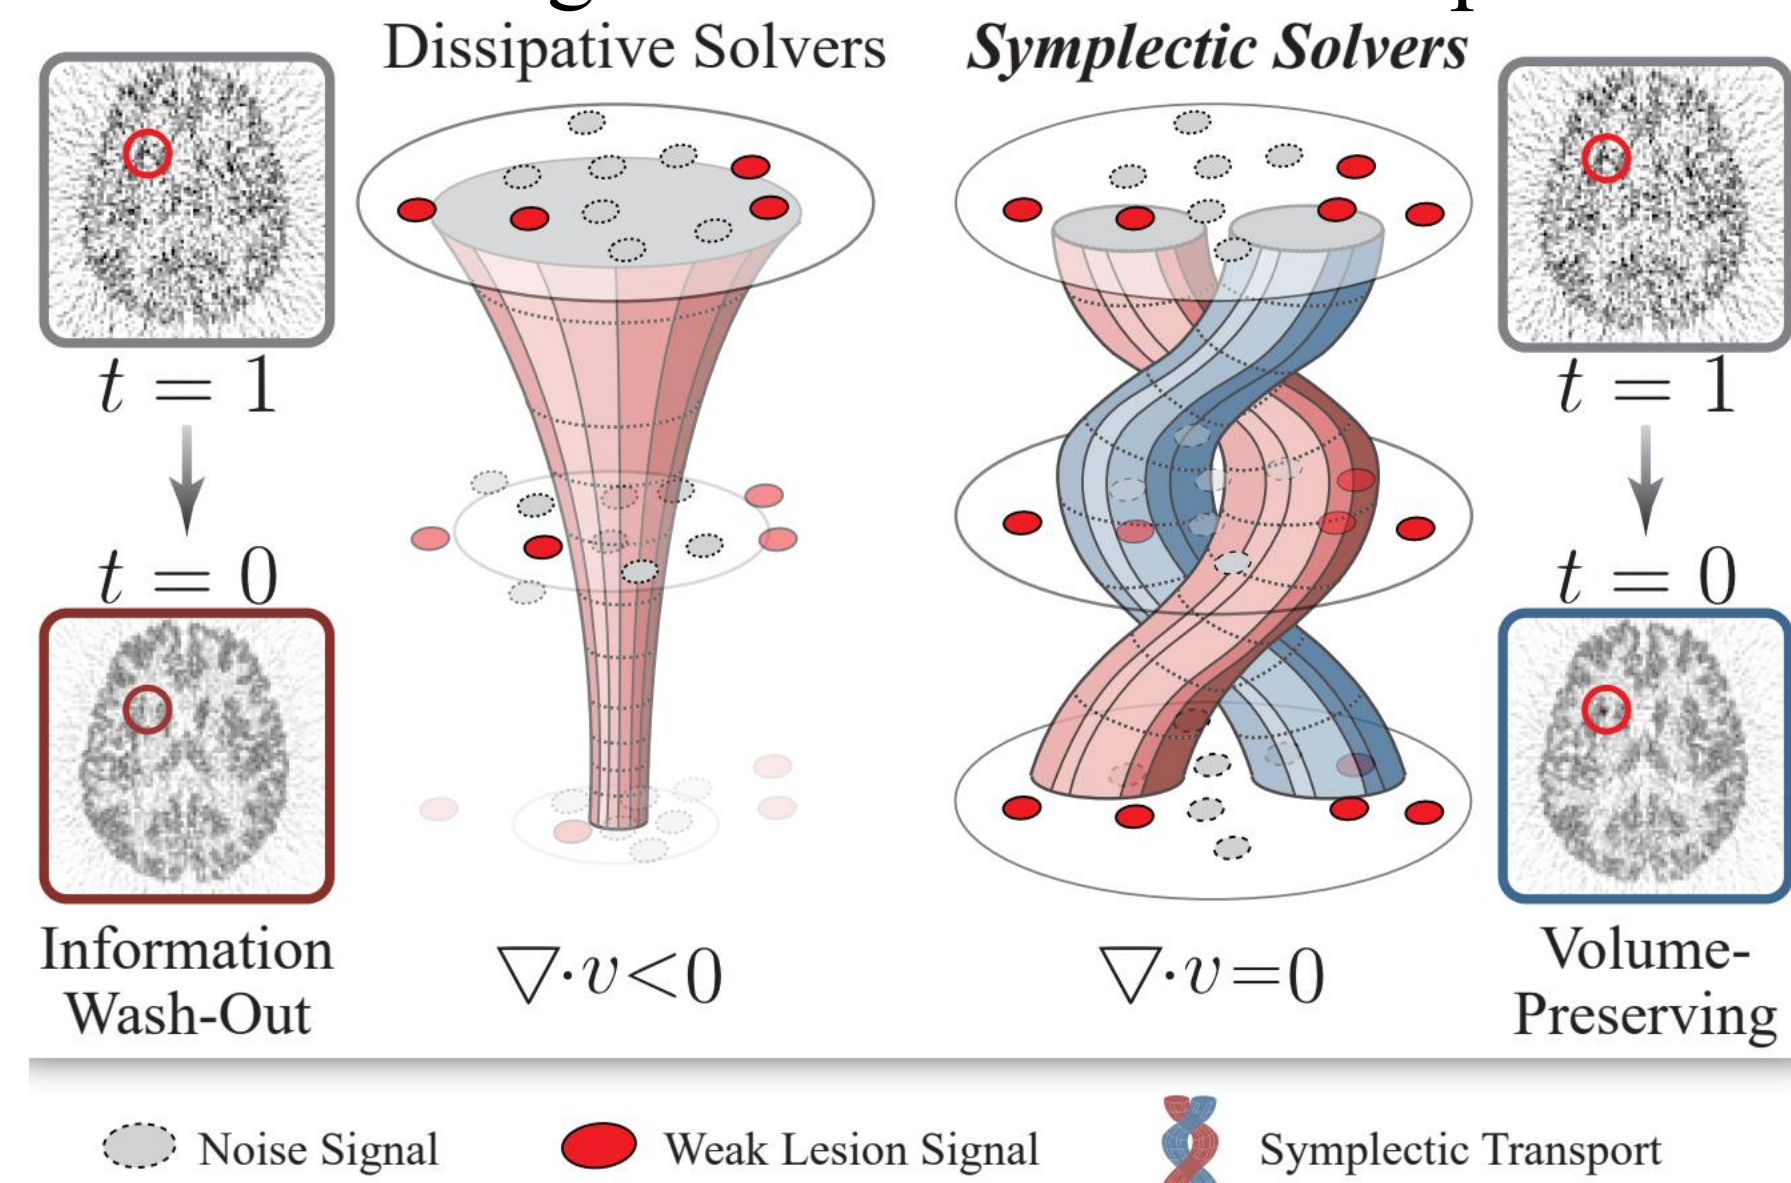


1. Background: PET imaging



Reducing tracer dose lowers radiation but degrades SNR due to Poisson statistics, making reconstruction severely ill-posed: $y \sim \text{Poisson}(Ax)$.

2. Motivation: Signal Wash-Out in Dissipative Flow



Dissipative solvers ($\nabla \cdot v < 0$) contract phase-space volume indiscriminately, washing out weak lesion signals alongside noise. FlowPET enforces symplectic transport ($\nabla \cdot v \equiv 0$), preserving signal mass by construction.

3. Hamiltonian Posterior Transport

$$H_\theta(t, x, p; y) = U_\psi(t, x; y) + K_\phi(t, p; y).$$

Separability ensures the Jacobian diagonal blocks vanish identically — divergence-free transport is guaranteed by architecture, not regularization.

$$v_\theta(t, z; y) = \begin{bmatrix} \dot{x} \\ \dot{p} \end{bmatrix} = \begin{bmatrix} \nabla_p K_\phi(t, p; y) \\ -\nabla_x U_\psi(t, x; y) \end{bmatrix}.$$

$$\nabla_z \cdot v_\theta = \text{tr} \left(\frac{\partial \dot{x}}{\partial x} \right) + \text{tr} \left(\frac{\partial \dot{p}}{\partial p} \right) = \text{tr}(\mathbf{0}) + \text{tr}(\mathbf{0}) \equiv 0.$$

4. Conjugate Phase-Space Boundaries

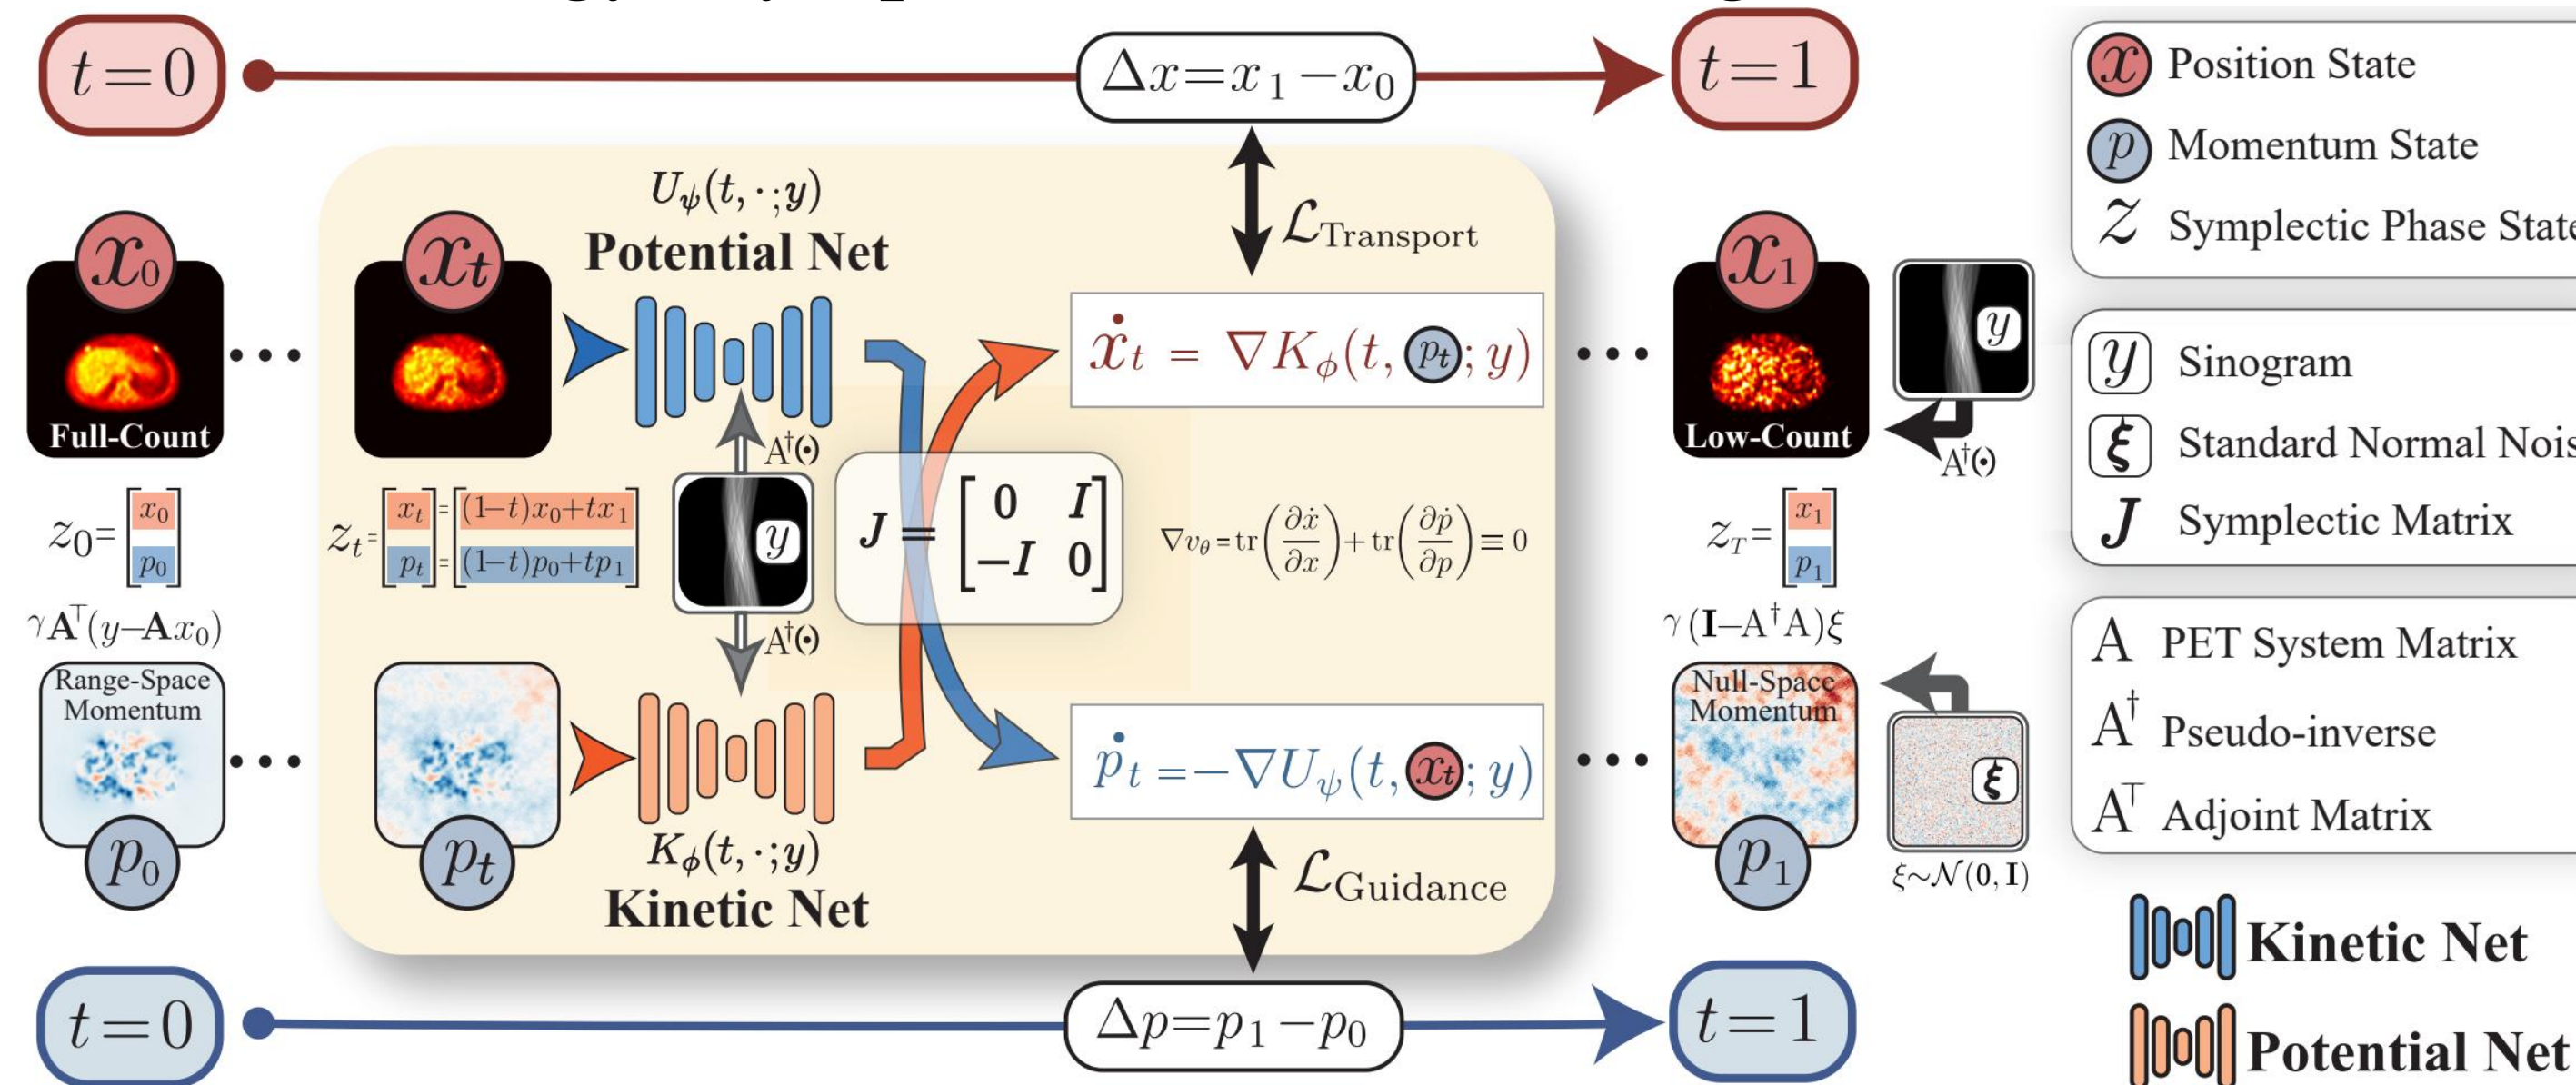
$$p_0 := \gamma \mathbf{A}^\top (y - \mathbf{A}x_0).$$

p_0 embeds data fidelity as a restoring force (range space);

$$p_1 = \gamma (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A}) \xi, \quad \xi \sim \mathcal{N}(0, \mathbf{I}).$$

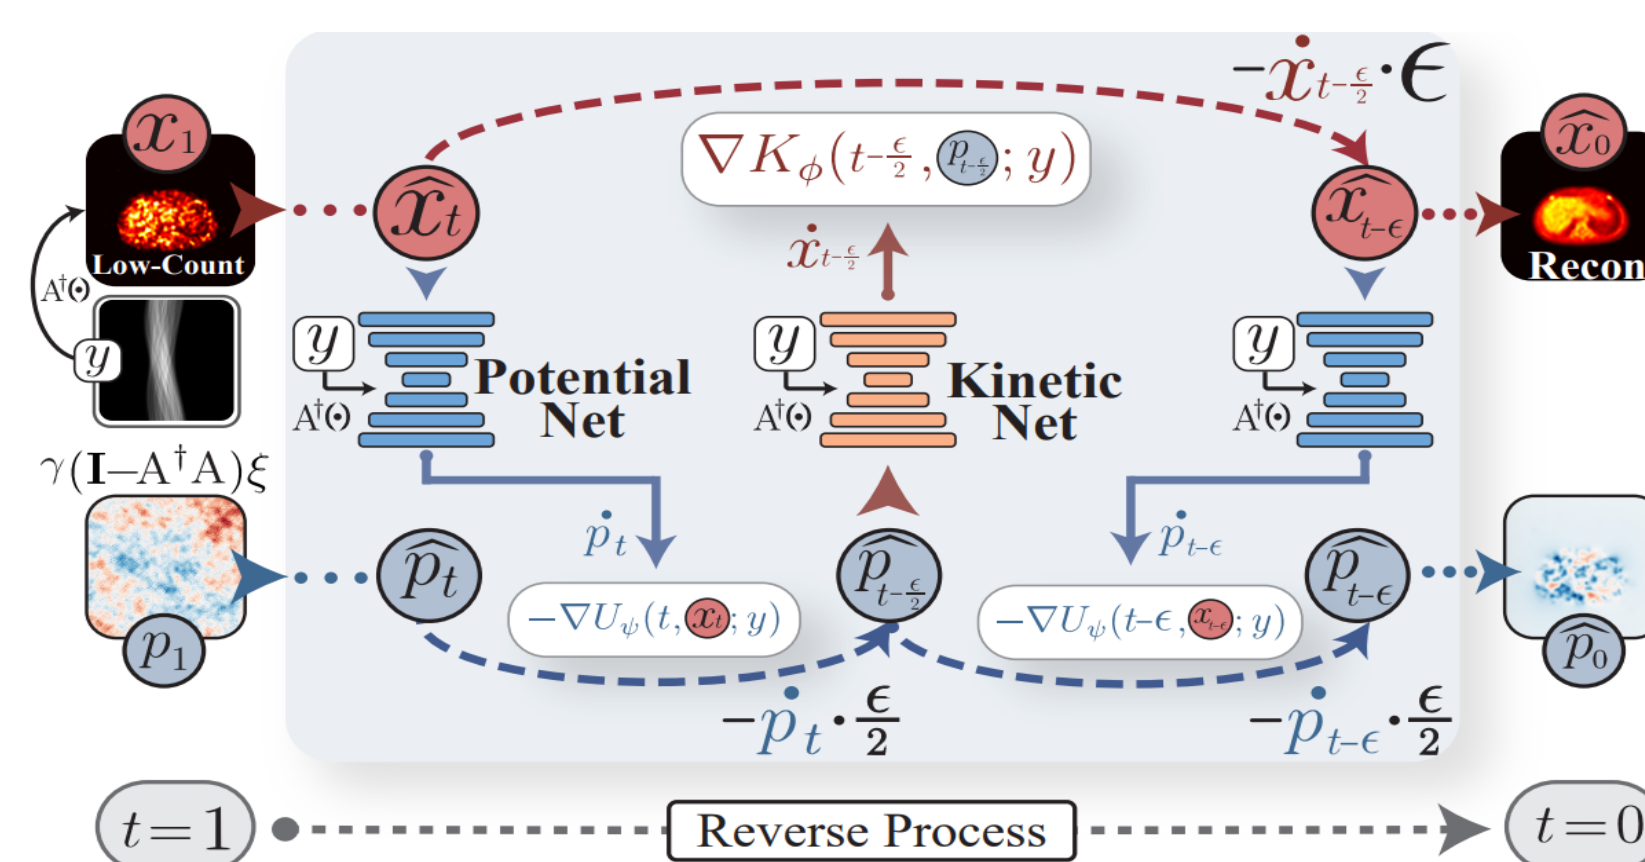
p_1 injects stochastic diversity only where physics is blind (null space).

5. Methodology: Symplectic Flow Matching



FlowPET formulates reconstruction as volume-preserving transport in symplectic phase space $\mathcal{Z} = \mathcal{X} \times \mathcal{P}$, driven by a Separable Hamiltonian System (U_ψ, K_ϕ). The trajectory bridges the full-count reference ($t = 0$) and the measurement-informed prior ($t = 1$), with physics-informed momentum initialization decoupling data consistency from stochastic exploration.

6. Methodology: Leapfrog Inference



Separable H enables explicit time-staggered updates. Unlike RK4, Leapfrog preserves symplectic structure exactly — no numerical dissipation.

7. Experiments & Wash-Out Analysis

Method	BrainWeb (20% Count)			In-House (1% Count)			UDPET Brain (1% Count)		
	SSIM↑	PSNR↑	RMSE↓	SSIM↑	PSNR↑	RMSE↓	SSIM↑	PSNR↑	RMSE↓
OSEM (Hudson & Larkin, 1994)	0.9078	28.35	0.0447	0.7456	23.59	0.0745	0.7607	19.87	0.1108
AutoContextCNN (Xiang et al., 2017)	0.9816	33.64	0.0233	0.9339	33.66	0.0226	0.8794	26.29	0.0541
DeepPET (Hägström et al., 2019)	0.9746	30.08	0.0331	0.8820	32.24	0.0263	0.8218	25.28	0.0581
CNNBPnet (Zhang et al., 2020)	0.9560	30.62	0.0329	0.9240	34.62	0.0200	0.7750	25.06	0.0621
FBPnet (Wang & Liu, 2020)	0.9327	33.62	0.0231	0.9592	34.19	0.0210	0.8907	27.36	0.0463
LCPR-Net (Xue et al., 2021)	0.9769	33.75	0.0224	0.9222	34.95	0.0206	0.8919	27.77	0.0446
Sino-cGAN (Liu et al., 2022)	0.9641	30.76	0.0306	0.9704	33.58	0.0223	0.8646	25.54	0.0569
IR-SDE (Luo et al., 2023)	0.9589	32.55	0.0266	0.9728	33.81	0.0222	0.9103	25.52	0.0549
DGLM_u (Zhang et al., 2024)	0.9785	33.58	0.0230	0.9551	32.93	0.0245	0.8905	25.95	0.0552
RED (Ai et al., 2025)	0.9664	34.45	0.0210	0.9472	34.15	0.0192	0.8890	26.51	0.0474
DREAM (Huang et al., 2025)	0.9777	33.07	0.0246	0.9661	35.77	0.0173	0.9081	28.49	0.0411
FourierPET (Zhang et al., 2026)	0.9859	35.36	0.0198	0.9740	35.19	0.0188	0.9083	27.98	0.0437
<i>FlowPET (Ours)</i>	0.9838	36.34	0.0188	0.9811	36.35	0.0160	0.9139	28.88	0.0387

SOTA on all three datasets. On In-House 1% count: 36.35 dB PSNR, 0.9811 SSIM (+0.58 dB over DREAM). Wash-out analysis confirms sustained lesion recovery vs. premature saturation in dissipative solvers.

