Rex: A Family of Reversible Exponential (Stochastic) Runge-Kutta Solvers

Zander W. Blasingame[1,2] Chen Liu[2]
[1]AITHYRA [2]Clarkson University

TL;DR

- ✗ Existing reversible diffusion solvers: **ODE-only**, **low-order**, **fixed-step** sizes.
- ✓ **Rex** turns *any* explicit (S)RK scheme into a reversible diffusion solver **ODEs *and* SDEs**.
- ✓ Inherits **convergence order**, **stability**, and enables **adaptive step sizing**.
- ✓ **Subsumes** DDIM, DPM-Solver, SEEDS-1, gDDIM, and more.
- ✓ **Exact** inversion ⇒ better editing & Boltzmann sampling.

Rex: (S)RK schemes ⇒ reversible solvers for diffusion

For any explicit (S)RK scheme Φ, we convert into a reversible exponential (**Rex**) solver via the following 3-step recipe.

1. Φ | **pick a scheme.** *e.g.* Euler, RK4, Dopri5, Euler-Maruyama, or ShARK.
2. Ψ | **build *Princeps*.** Lift to an exponentially weighted change-of-variables transforming the ODE/SDE into a non-stiff form, apply Φ in this space, then pull back to original variable, see Fig. 1.
3. Υ | **build Rex.** Make Ψ *algebraic reversible* via the McCallum-Foster coupling, see Fig. 2.

For simplicity assume the following general additive-noise Itô SDE:

$$\mathrm{d}\boldsymbol{X}_t = a(t)\boldsymbol{X}_t + b(t)\boldsymbol{f}_\theta(t, \boldsymbol{X}_t)\;\mathrm{d}t + g(t)\;\mathrm{d}\boldsymbol{W}_t, \tag{1}$$

where $\boldsymbol{f}_\theta$ denotes a noise or data prediction model. Then we can lift Φ into its exponentially integrated form via Fig. 1 with $\Xi(t) := \exp\left(\int_0^t a(\tau)\;\mathrm{d}\tau\right)$ and $\varsigma_t = \int \Xi^{-1}(t)b(t)\;\mathrm{d}t$.

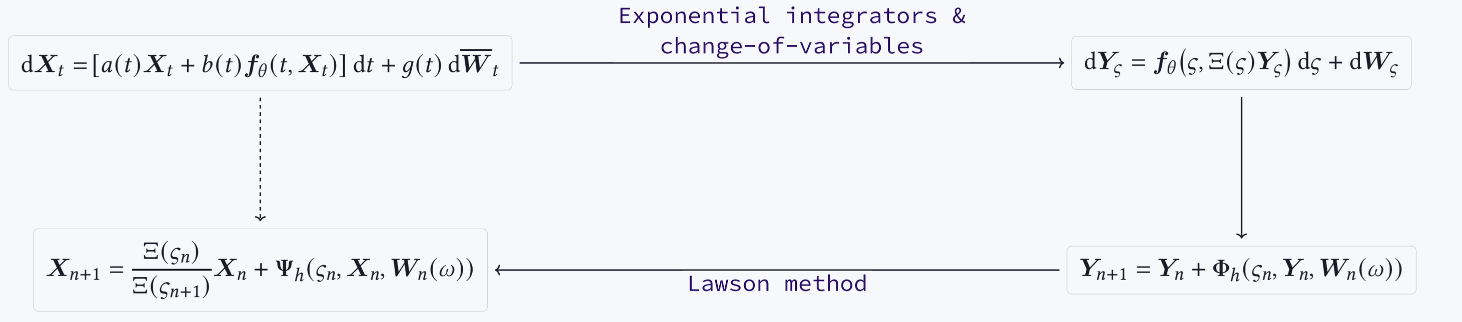

Figure 1. Construction of Ψ from an explicit (S)RK scheme Φ. Top path: Lawson change-of-variables; bottom path: Φ pulled back into the original state $\boldsymbol{X}_t$.

Now we use the reversible McCallum-Foster coupling to construct a reversible scheme in the lifted space, before bringing back to the original variable. See Fig. 2.

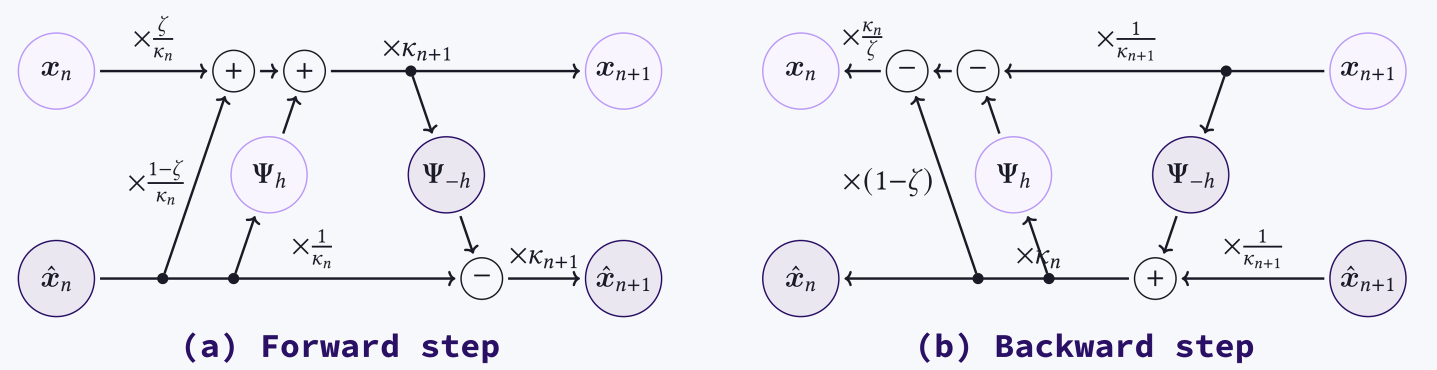

Figure 2. The Rex computation graph. Ψ_h is the exponentially weighted (S)RK scheme, $\zeta \in (0,1)$ is the coupling parameter, and κ_n is weighting term derived from the ODE/SDE formulation and if the model is parametrized via noise or data prediction.

The Rex update (Υ).
With step $h := \varsigma_{n+1} - \varsigma_n$ and weights $\kappa_n := \Xi(\varsigma_n)$, fix a single Brownian realization $\boldsymbol{W}_n(\omega)$.
Forward.

$$\boldsymbol{X}_{n+1} = \tfrac{\kappa_{n+1}}{\kappa_n}\left(\zeta\,\boldsymbol{X}_n + (1-\zeta)\,\hat{\boldsymbol{X}}_n\right) + \kappa_{n+1}\,\Psi_h(\varsigma_n, \hat{\boldsymbol{X}}_n, \boldsymbol{W}_n(\omega))$$
$$\hat{\boldsymbol{X}}_{n+1} = \tfrac{\kappa_{n+1}}{\kappa_n}\,\hat{\boldsymbol{X}}_n - \kappa_{n+1}\,\Psi_{-h}(\varsigma_{n+1}, \boldsymbol{X}_{n+1}, \boldsymbol{W}_n(\omega))$$

Backward (exact algebraic inverse, same path $\boldsymbol{W}_n(\omega)$).

$$\hat{\boldsymbol{X}}_n = \tfrac{\kappa_n}{\kappa_{n+1}}\,\hat{\boldsymbol{X}}_{n+1} + \kappa_n\,\Psi_{-h}(\varsigma_{n+1}, \boldsymbol{X}_{n+1}, \boldsymbol{W}_n(\omega))$$
$$\boldsymbol{X}_n = \tfrac{\kappa_n}{\kappa_{n+1}}\,\zeta^{-1}\boldsymbol{X}_{n+1} + (1-\zeta^{-1})\,\hat{\boldsymbol{X}}_n - \kappa_n\,\zeta^{-1}\,\Psi_h(\varsigma_n, \hat{\boldsymbol{X}}_n, \boldsymbol{W}_n(\omega))$$

***Princeps* subsumes prior solvers.** *Princeps* recovers, as special cases, DDIM; DPM-Solver; DPM-Solver++, SDE-DPM-Solver-1; SEEDS-1; and gDDIM. **Corollary.** **Rex** is the **reversible version of each**, *e.g.*, **Rex** (Euler) is Reversible DDIM.

Convergence & stability.
Order (ODE). If Φ is a k-th order explicit RK scheme for the reparametrized PF-ODE, **Rex** inherits **order** k: $\|\boldsymbol{x}_n - \boldsymbol{x}_{t_n}\| \leq C\,h^k$.
Order (SDE). If Φ is an SRK of strong order $\xi > 0$ for the reparametrized reverse-time SDE, then *Princeps* has **strong order** ξ.
Stability. **Rex** is inherits the **non-zero linear-stability region** from McCallum-Foster; BDIA and O-BELM are nowhere linearly stable.

Turn *any* explicit (S)RK scheme into a reversible diffusion solver

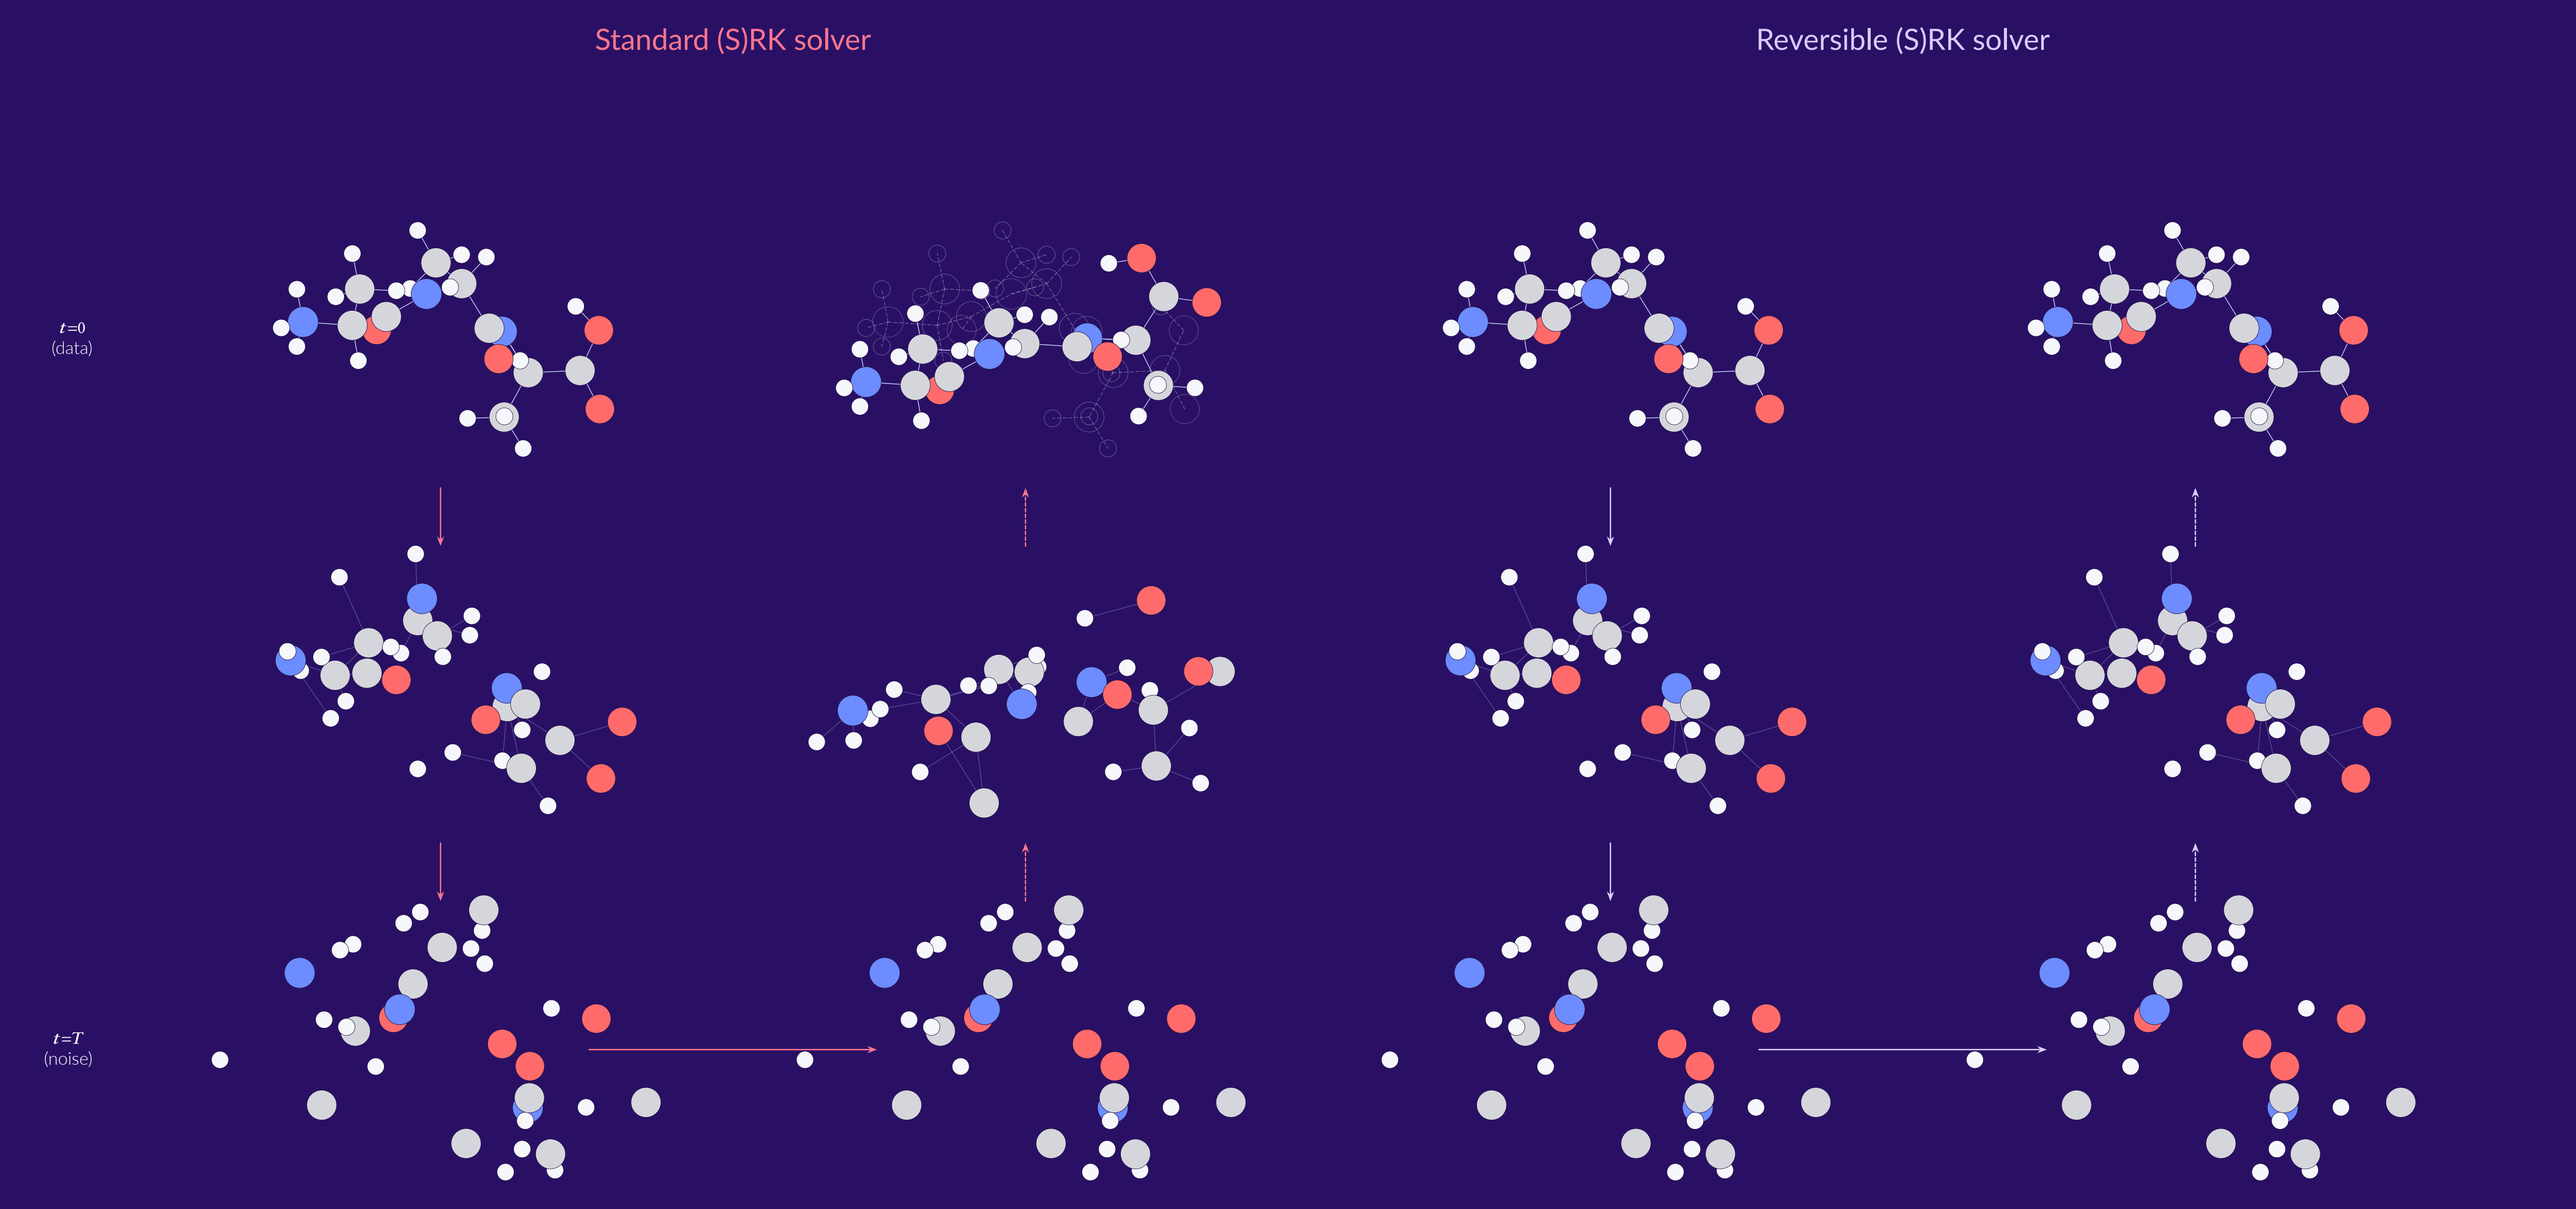

Rex is more robust to floating point errors

Round-trip latent MSE on SD v1.5 (512×512, fp32): forward then reverse solve, measuring $\|z_0 - \hat{z}_0\|_2^2$. The inversion error of **Rex** (Euler) orders of magnitude below every baseline (DDIM, EDICT, BDIA, O-BELM).

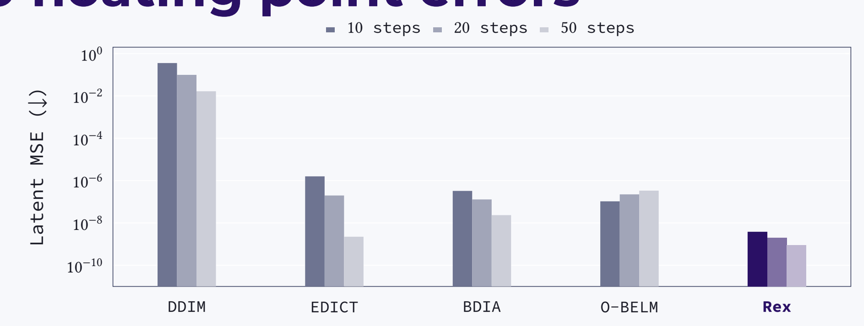

Figure 3. Latent reconstruction MSE (fp32, log scale); lower is better.

Rex improves image baselines

Unconditional sampling. Reversible solvers vs. a pre-trained DDPM on CelebA-HQ (256×256), 10^4 samples, six metrics (FD, FD$_\infty$, Precision, Recall, Density, Coverage).

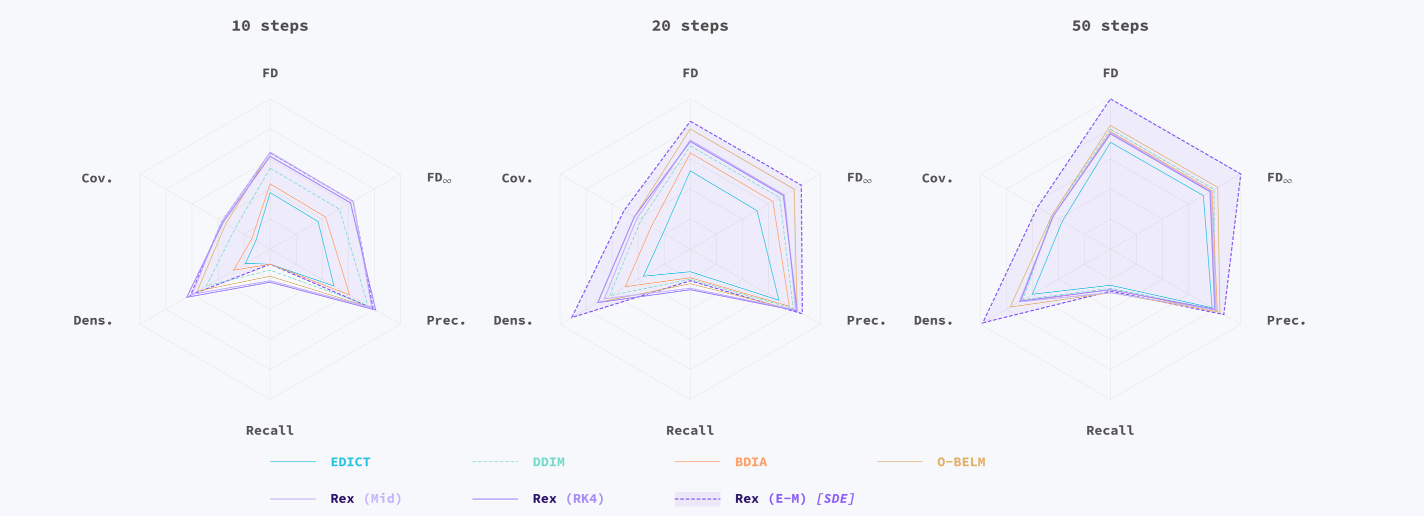

Figure 4. Six-axis radar; outer ring is best. Filled mauve = **Rex** (Euler-Maruyama, SDE).

Conditional generation & editing. **Text-to-image** (SD v1.5, COCO captions) and **editing** (pix2pix, 50+50 steps).

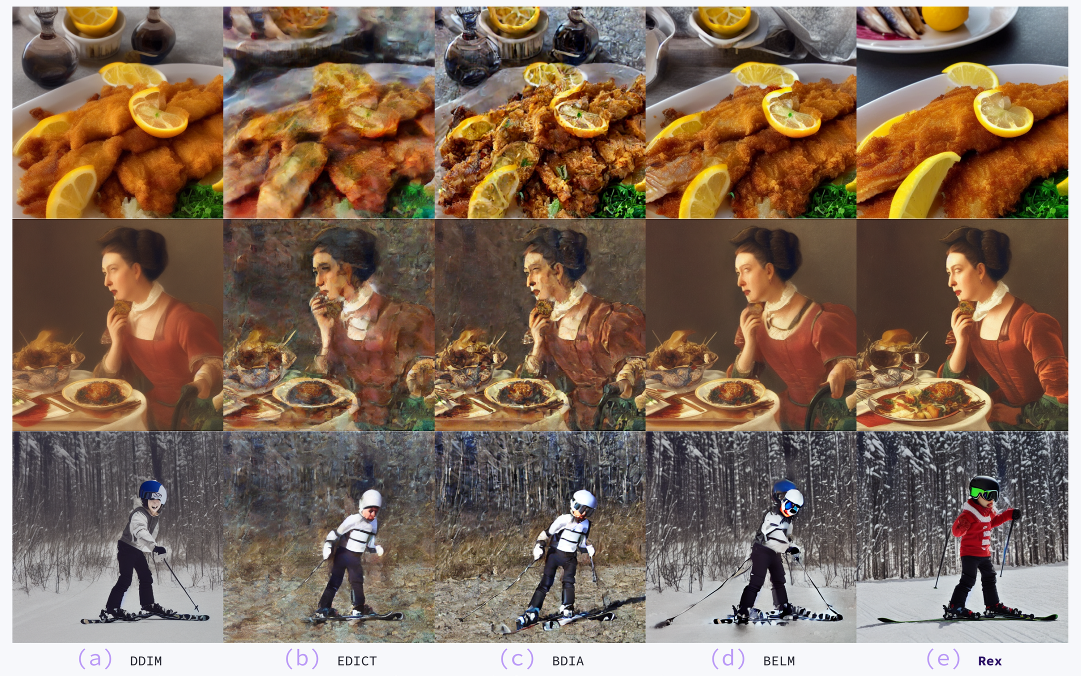

(a) DDIM (b) EDICT (c) BDIA (d) BELM (e) Rex

Figure 5. Text-to-image with SD v1.5 (512x512) at 10 steps.

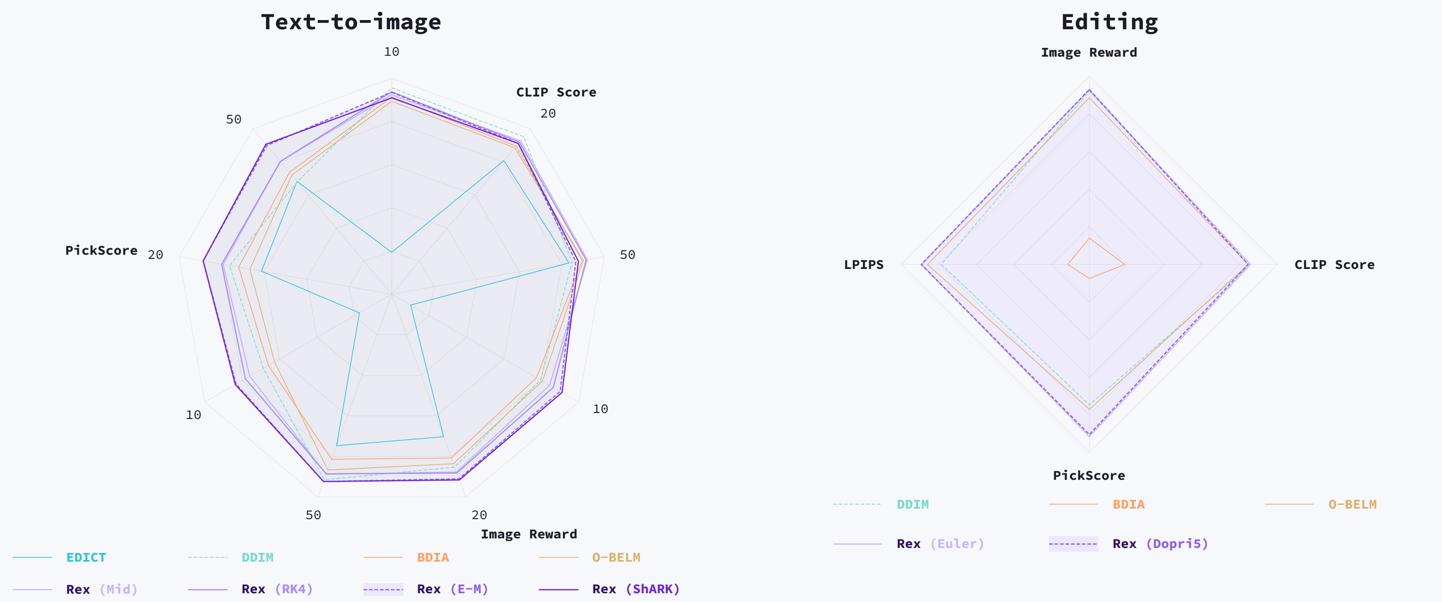

Improved Boltzmann sampling with bijective solvers

Swapping Dopri5 for **Rex** (Dopri5) in a pre-trained DiT makes the change-of-variables exact, improving importance sampling and beating the SBG and DiT baselines on energy Wasserstein.

Model	Scheme	ESS (↑)	$\mathcal{E}$-$\mathcal{W}_2$(↓)	$\mathbb{T}$-$\mathcal{W}_2$(↓)
RegFlow	---	0.029	1.051	1.612
SBG (IS)	---	0.052	0.758	0.502
SBG (SMC)	---	---	0.598	0.503
ECNF++	Dopri5	0.003	2.206	0.962
DiT	Dopri5	**0.140**	0.737	**0.468**
DiT	**Rex** (Dopri5)	0.104	**0.495**	0.497

Table 1. Boltzmann resampled metrics.