

Appendix

This appendix is organized as follows. In Section A, we present proofs for all of our propositions. Section A.1 presents a rigorous definition of our forward process using a more specific notation. This is then used in Section A.2.1 to prove the time reversal for our jump diffusions. We also present an intuitive proof of the time reversal using notation from the main text in Section A.2.2. In Section A.3 we prove Proposition 2 using the notation from the main text. We prove Proposition 3 in Section A.4 and we analyse the optimum of our objective directly without using stochastic process theory in Section A.5. In Section B we give more details on our objective and in Section C we detail how we apply diffusion guidance to our model. We give the full details for our experiments in Section D and finally, in Section E, we discuss the broader impacts of our work.

A Proofs

A.1 Notation and Setup

We here introduce a more rigorous notation for defining our trans-dimensional notation that will be used in a rigorous proof for the time-reversal of our jump diffusion. First, while it makes sense from a methodological and experimental point of view to present our setting as a *transdimensional* one, we slightly change the point of view in order to derive our theoretical results. We extend the space \mathbb{R}^d to $\hat{\mathbb{R}}^d = \mathbb{R}^d \cup \{\infty\}$ using the *one-point compactification* of the space. We refer to [43] for details on this space. The point ∞ will be understood as a mask. For instance, let $x_1, x_2, x_3 \in \mathbb{R}^d$. Then $X = (x_1, x_2, x_3) \in (\hat{\mathbb{R}}^d)^N$ with $N = 3$ corresponds to a vector for which all components are *observed* whereas $X' = (x_1, \infty, x_3) \in (\hat{\mathbb{R}}^d)^N$ corresponds to a vector for which only the components on the first and third dimension are observed. The second dimension is *masked* in that case. Doing so, we will consider diffusion models on the space $\mathbb{X} = (\hat{\mathbb{R}}^d)^N$ with $d, N \in \mathbb{N}$. In the case of a video diffusion model, N can be seen as the max number of frames. We will always consider that this space is equipped with its Borelian sigma-field \mathcal{X} and all probability measures will be defined on \mathcal{X} .

We denote $\dim : \mathbb{X} \rightarrow \{0, 1\}^N$ which is given for any $X = \{x_i\}_{i=1}^N \in \mathbb{X}$ by

$$\dim(X) = \{\delta_{\mathbb{R}^d}(x_i)\}_{i=1}^N.$$

In other words, $\dim(X)$ is a binary vector identifying the “dimension” of the vector X , i.e. which frames are observed. Going back to our example $X = (x_1, x_2, x_3) \in (\hat{\mathbb{R}}^d)^N$ and $X' = (x_1, \infty, x_3) \in (\hat{\mathbb{R}}^d)^N$, we have that $\dim(X) = \{1, 1, 1\}$ and $\dim(X') = \{1, 0, 1\}$. For any vector $u \in \{0, 1\}^N$ we denote $|u| = \sum_{i=1}^N u_i$, i.e. the *active dimensions* of u (or equivalently the non-masked frames). For any $X \in \mathbb{X}$ and $D \in \{0, 1\}^N$, we denote $X_D = \{X'_i\}_{i=1}^N$ with $X'_i = X_i$ if $D_i = 1$ and $X'_i = \infty$ if $D_i = 0$.

We denote $C_b^k(\mathbb{R}^d, \mathbb{R})$ the set of functions which are k differentiable and bounded. Similarly, we denote $C_b^k(\mathbb{R}^d, \mathbb{R})$ the set of functions which are k differentiable and compactly supported. The set $C_0^k(\mathbb{R}^d, \mathbb{R})$ denotes the functions which are k differentiable and vanish when $\|x\| \rightarrow +\infty$. We note that $f \in C(\hat{\mathbb{R}}^d)$, if $f \in C(\mathbb{R}^d)$ and $f - f(\infty) \in C_0(\mathbb{R}^d)$ and that $f \in C^k(\hat{\mathbb{R}}^d)$ for any $k \in \mathbb{N}$ if the restriction of f to \mathbb{R}^d is in $C^k(\mathbb{R}^d)$ and $f \in C(\hat{\mathbb{R}}^d)$.

A.1.1 Transdimensional infinitesimal generator

To introduce rigorously the *transdimensional* diffusion model defined in Section 3.1, we will introduce its *infinitesimal generator*. The infinitesimal generator of a stochastic process can be roughly defined as its “probabilistic derivative”. More precisely, assume that a stochastic process $(\mathbf{X}_t)_{t \geq 0}$ admits a transition semigroup $(P_t)_{t \geq 0}$, i.e. for any $t \geq 0$, $A \in \mathcal{X}$ and $X \in \mathbb{X}$ we have $\mathbb{P}(\mathbf{X}_t \in A \mid \mathbf{X}_0 = x) = P_t(x, A)$, then the infinitesimal generator is defined as $\mathcal{A}(f) = \lim_{t \rightarrow 0} (P_t(f) - f)/t$, for every f for which this quantity is well-defined.

Here, we start by introducing the infinitesimal generator of interest and give some intuition about its form. Then, we prove a time-reversal formula for this infinitesimal generator.

We consider $b : \mathbb{R}^d \rightarrow \mathbb{R}^d$, $\alpha : \{0, 1\}^{NM} \rightarrow \mathbb{R}_+$. For any $f \in C^2(X)$ and $X \in X$ we define

$$\begin{aligned} \mathcal{A}(f)(X) &= \sum_{i=1}^N \{ \langle b(X_i), \nabla_{x_i} f(X) \rangle + \frac{1}{2} \Delta_{x_i} f(X) \} \delta_{\mathbb{R}^d}(X_i) \\ &\quad - \sum_{D_1 \subset D_0^{\Delta_0}} \cdots \sum_{D_M \subset D_{M-1}^{\Delta_{M-1}}} \alpha(D_0, \dots, D_M) \sum_{i=0}^{M-1} (f(X) - f(X_{D_{i+1}})) \delta_{D_i}(\dim(X)), \end{aligned} \quad (5)$$

where $M \in \mathbb{N}$, $D_0 = \{1\}^N$, $\{\Delta_j\}_{j=0}^{M-1} \in \mathbb{N}^M$ such that $\sum_{j=0}^{M-1} \Delta_j < N$ and for any $j \in \{0, \dots, M-1\}$, $D_j^{\Delta_j}$ is the subset of $\{0, 1\}^{\{1, \dots, N\}}$ such that $D_{j+1} \in D_j^{\Delta_j}$ if and only if $D_j \cdot D_{j+1} = D_{j+1}$, where \cdot is the pointwise multiplication operator, and $|D_j| = |D_{j+1}| + \Delta_j$. The condition $D_j \cdot D_{j+1} = D_{j+1}$ means that the non-masked dimensions in D_{j+1} are also non-masked dimensions in D_j . The condition $|D_j| = |D_{j+1}| + \Delta_j$ means that in order to go from D_j to D_{j+1} , one needs to mask exactly Δ_j dimensions.

Therefore, a sequence $\{\Delta_j\}_{j=0}^{M-1} \in \mathbb{N}^M$ such that $\sum_{j=0}^{M-1} \Delta_j < N$ can be interpreted as a sequence of *drops* in dimension. At the core level, we have that $|D_M| = N - \sum_{j=0}^{M-1} \Delta_j$. For instance if $|D_M| = 1$, we have that at the end of the process, only one dimension is considered.

We choose α such that $\sum_{D_1 \subset D_0^{\Delta_0}} \cdots \sum_{D_M \subset D_{M-1}^{\Delta_{M-1}}} \alpha(D_0, \dots, D_M) = 1$. Therefore, $\alpha(D_0, \dots, D_M)$ corresponds to the probability to choose the *dimension path* $D_0 \rightarrow \cdots \rightarrow D_M$.

The part $X \mapsto \langle b(X_i), \nabla_{x_i} f(X) \rangle + \frac{1}{2} \Delta_{x_i} f(X)$ is more classical and corresponds to the *continuous part* of the diffusion process. We refer to [44] for a thorough introduction on infinitesimal generators. For simplicity, we omit the schedule coefficients in (5).

A.1.2 Justification of the form of the infinitesimal generator

For any *dimension path* $P = D_0 \rightarrow \cdots \rightarrow D_M$ (recall that $D_0 = \{1\}^N$), we define the *jump kernel* \mathbb{J}^P as follows. For any $x \in X$, we have $\mathbb{J}^P(X, dY) = \sum_{i=0}^{M-1} \delta_{D_i}(\dim(X)) \delta_{X_{D_{i+1}}}(dY)$. This operator corresponds to the *deletion* operator introduced in Section 3.1. Hence, for any *dimension path* $P = D_0 \rightarrow \cdots \rightarrow D_M$, we can define the associated infinitesimal generator: for any $f \in C^2(X)$ and $X \in X$ we define

$$\mathcal{A}^P(f)(X) = \sum_{i=1}^N \{ \langle b(x_i), \nabla_{x_i} f(X) \rangle + \frac{1}{2} \Delta_{x_i} f(X) \} \delta_{\mathbb{R}^d}(X_i) + \int_X (f(Y) - f(X)) \mathbb{J}^P(X, dY).$$

We can define the following *jump kernel*

$$\mathbb{J} = \sum_{D_1 \subset D_0^{\Delta_0}} \cdots \sum_{D_M \subset D_{M-1}^{\Delta_{M-1}}} \alpha(D_0, \dots, D_M) \mathbb{J}^P.$$

This corresponds to averaging the jump kernel over the different possible dimension paths. We have that for any $f \in C^2(X)$ and $X \in X$

$$\mathcal{A}(f)(X) = \sum_{i=1}^N \{ \langle b(x_i), \nabla_{x_i} f(X) \rangle + \frac{1}{2} \Delta_{x_i} f(X) \} \delta_{\mathbb{R}^d}(X_i) + \int_X (f(Y) - f(X)) \mathbb{J}(X, dY). \quad (6)$$

In other words, $\mathcal{A} = \sum_{D_1 \subset D_0^{\Delta_0}} \cdots \sum_{D_M \subset D_{M-1}^{\Delta_{M-1}}} \alpha(D_0, \dots, D_M) \mathcal{A}^P$.

In what follows, we assume that there exists a Markov process $(\mathbf{X}_t)_{t \geq 0}$ with infinitesimal generator \mathcal{A} . In order to sample from $(\mathbf{X}_t)_{t \geq 0}$, one choice is to first sample the dimension path P according to the probability α . Second sample from the Markov process associated with the infinitesimal generator \mathcal{A}^P . We can approximately sample from this process using the Lie-Trotter-Kato formula [44, Corollary 6.7, p.33].

Denote $(P_t)_{t \geq 0}$ the semigroup associated with \mathcal{A}^P , $(Q_t)_{t \geq 0}$ the semigroup associated with the continuous part of \mathcal{A}^P and $(J_t)_{t \geq 0}$ the semigroup associated with the jump part of \mathcal{A}^P . More precisely, we have that, $(Q_t)_{t \geq 0}$ is associated with $\mathcal{A}_{\text{cont}}$ such that for any $f \in C^2(X)$ and $X \in X$

$$\mathcal{A}_{\text{cont}}(f)(X) = \sum_{i=1}^N \{ \langle b(X_i), \nabla_{x_i} f(X) \rangle + \frac{1}{2} \Delta_{x_i} f(X) \}.$$

In addition, we have that, $(Q_t)_{t \geq 0}$ is associated with $\mathcal{A}_{\text{jump}}^P$ such that for any $f \in C^2(X)$ and $X \in X$

$$\mathcal{A}_{\text{jump}}^P(f)(X) = \int_X (f(Y) - f(X)) \mathbb{J}^P(X, dY).$$

First, note that $\mathcal{A}_{\text{cont}}$ corresponds to the infinitesimal generator of a classical diffusion on the components which are not set to ∞ . Hence, we can approximately sample from $(Q_t)_{t \geq 0}$ by sampling according to the Euler-Maruyama discretization of the associated diffusion, i.e. by setting

$$\mathbf{X}_t \approx \mathbf{X}_0 + tb(\mathbf{X}_0) + \sqrt{t}Z, \quad (7)$$

where Z is a Gaussian random variable.

Similarly, in order to sample from $(J_t)_{t \geq 0}$, one should sample from the jump process defined as follows. On the interval $[0, \tau)$, we have $\bar{\mathbf{X}}_t = \mathbf{X}_0$. At time τ , we define $\bar{\mathbf{X}}_1 \sim \mathbb{J}(\mathbf{X}_0, \cdot)$ and repeat the procedure. In this case τ is defined as an exponential random variable with parameter 1. For $t > 0$ small enough the probability that $t > \tau$ is of order t . Therefore, we sample from J , i.e. the deletion kernel, with probability t . Combining this approximation and (7), we get approximate samplers for $(Q_t)_{t \geq 0}$ and $(J_t)_{t \geq 0}$. Under mild assumptions, the Lie-Trotter-Kato formula ensures that for any $t \geq 0$

$$P_t = \lim_{n \rightarrow +\infty} (Q_{t/n} J_{t/n})^n.$$

This justifies sampling according to Algorithm 1 (in the case of the forward process).

A.2 Proof of Proposition 1

For the proof of Proposition 1, we first provide a rigorous proof using the notation introduced in A.1. We then follow this with a second proof that aims to be more intuitive using the notation used in the main paper.

A.2.1 Time-reversal for the transdimensional infinitesimal generator and Proof of Proposition 1

We are now going to derive the formula for the time-reversal of the transdimensional infinitesimal generator \mathcal{A} , see (5). This corresponds to a rigorous proof of Proposition 1. We refer to Section A.2.2 for a more intuitive, albeit less-rigorous, proof. We start by introducing the kernel \mathbb{K}^P given for any dimension path $D_0 \rightarrow \dots \rightarrow D_M$, for any $i \in \{0, \dots, M-1\}$, $Y \in D_{i+1}$ and $A \in \mathcal{X}$ by

$$\mathbb{K}^P(Y, A) = \sum_{i=0}^{M-1} \delta_{D_{i+1}}(\dim(Y)) \int_{A \cap D_i} \frac{p_t((X_{D_i \setminus D_{i+1}}, Y_{D_{i+1}}) | \dim(\mathbf{X}_t) = D_i) \mathbb{P}(\dim(\mathbf{X}_t) = D_i)}{p_t(Y_{D_{i+1}} | \dim(\mathbf{X}_t) = D_{i+1}) \mathbb{P}(\dim(\mathbf{X}_t) = D_{i+1})} dX_{D_i \setminus D_{i+1}}.$$

Note that this kernel is the same as the one considered in Proposition 1. It is well-defined under the following assumption.

Assumption 1. For any $t > 0$ and $D \subset \{0, 1\}^N$, we have that \mathbf{X}_t conditioned to $\dim(\mathbf{X}_t) = D$ admits a density w.r.t. the $|D|$ -dimensional Lebesgue measure, denoted $p_t(\cdot | \dim(\mathbf{X}_t) = D)$.

The following result will be key to establish the time-reversal formula.

Lemma 1. Assume A1. Let $A, B \in \mathcal{X}$. Let P be a dimension path $D_0 \rightarrow \dots \rightarrow D_M$ with $M \in \mathbb{N}$. Then, we have

$$\mathbb{E}[\mathbf{1}_A(\mathbf{X}_t) \mathbb{J}^P(\mathbf{X}_t, B)] = \mathbb{E}[\mathbf{1}_B(\mathbf{X}_t) \mathbb{K}^P(\mathbf{X}_t, A)].$$

Proof. Let $A, B \in \mathcal{X}$. We have

$$\begin{aligned} \mathbb{E}[\mathbf{1}_A(\mathbf{X}_t) \mathbb{J}^P(\mathbf{X}_t, B)] &= \sum_{i=0}^{M-1} \mathbb{E}[\mathbf{1}_A(\mathbf{X}_t) \delta_{D_i}(\dim(\mathbf{X}_t)) \mathbf{1}_B((\mathbf{X}_t)_{D_{i+1}})] \\ &= \sum_{i=0}^{M-1} \int_{A \cap D_i} p_t(X_{D_i} | \dim(\mathbf{X}_t) = D_i) \mathbb{P}(\dim(\mathbf{X}_t) = D_i) \mathbf{1}_B(X_{D_{i+1}}) dX_{D_i} \\ &= \sum_{i=0}^{M-1} \int_{A \cap D_i} p_t(X_{D_i} | \dim(\mathbf{X}_t) = D_i) \mathbb{P}(\dim(\mathbf{X}_t) = D_i) \mathbf{1}_B(X_{D_{i+1}}) dX_{D_{i+1}} dX_{D_i \setminus D_{i+1}} \\ &= \sum_{i=0}^{M-1} \int_{B \cap D_{i+1}} \mathbf{1}_B(X_{D_{i+1}}) \\ &\quad \times \left(\int_{A \cap D_i} \mathbf{1}_A(X_{D_i}) p_t(X_{D_i} | \dim(\mathbf{X}_t) = D_i) \mathbb{P}(\dim(\mathbf{X}_t) = D_i) dX_{D_i \setminus D_{i+1}} \right) dX_{D_{i+1}} \\ &= \sum_{i=0}^{M-1} \int_{B \cap D_{i+1}} \mathbf{1}_B(X_{D_{i+1}}) \\ &\quad \times \mathbb{K}^P(X_{D_{i+1}}, A) p_t(X_{D_{i+1}} | \dim(\mathbf{X}_t) = D_{i+1}) \mathbb{P}(\dim(\mathbf{X}_t) = D_{i+1}) dX_{D_{i+1}} \\ &= \sum_{i=0}^{M-1} \mathbb{E}[\delta_{D_{i+1}}(\dim(\mathbf{X}_t)) \mathbb{K}^P(\mathbf{X}_t, A) \mathbf{1}_B(\mathbf{X}_t)], \end{aligned}$$

which concludes the proof. \square

Lemma 1 shows that \mathbb{K}^P verifies the *flux equation* associated with \mathbb{J}^P . The flux equation is the discrete state-space equivalent of the classical time-reversal formula for continuous state-space. We refer to [45] for a rigorous treatment of time-reversal with jumps under entropic conditions.

We are also going to consider the following assumption which ensures that the integration by part formula is valid.

Assumption 2. For any $t > 0$ and $i \in \{1, \dots, N\}$, \mathbf{X}_t admits a smooth density w.r.t. the Nd -dimensional Lebesgue measure denoted p_t and we have that for any $f, h \in C_b^2((\mathbb{R}^d)^N)$ for any $u \in [0, t]$ and $i \in \{1, \dots, N\}$

$$\begin{aligned} & \mathbb{E}[\delta_{\mathbb{R}^d}((\mathbf{X}_u)_i) \langle \nabla_{x_i} f(\mathbf{X}_u), \nabla_{x_i} h(\mathbf{X}_u) \rangle] \\ &= -\mathbb{E}[\delta_{\mathbb{R}^d}((\mathbf{X}_u)_i) h(\mathbf{X}_u) (\Delta_{x_i} f(\mathbf{X}_u) + \langle \nabla_{x_i} \log p_u(\mathbf{X}_u), \nabla_{x_i} f(\mathbf{X}_u) \rangle)]. \end{aligned}$$

The second assumption ensures that we can apply the backward Kolmogorov evolution equation.

Assumption 3. For any $g \in C^2(\mathbb{X})$ and $t > 0$, we have that for any $u \in [0, t]$ and $X \in \mathbb{X}$, $\partial_u g(u, X) + \mathcal{A}(g)(u, X) = 0$, where for any $u \in [0, t]$ and $X \in \mathbb{X}$, $g(u, X) = \mathbb{E}[g(\mathbf{X}_t) \mid \mathbf{X}_u = X]$.

We refer to [19] for conditions under A2 and A3 are valid in the setting of diffusion processes.

Proposition 4. Assume A1, A2 and A3. Assume that there exists a Markov process $(\mathbf{X}_t)_{t \geq 0}$ solution of the martingale problem associated with (6). Let $T > 0$ and consider $(\mathbf{Y}_t)_{t \in [0, T]} = (\mathbf{X}_{T-t})_{t \in [0, T]}$. Then $(\mathbf{Y}_t)_{t \in [0, T]}$ is solution to the martingale problem associated with \mathcal{R} , where for any $f \in C^2(\mathbb{X})$, $t \in (0, T)$ and $x \in \mathbb{X}$ we have

$$\begin{aligned} \mathcal{R}(f)(t, X) &= \sum_{i=1}^N \{ -\langle b(X_i) + \nabla_{x_i} \log p_t(X), \nabla_{x_i} f(X) \rangle + \frac{1}{2} \Delta_{x_i} f(X) \} \delta_{\mathbb{R}^d}(X_i) \\ &+ \int_{\mathbb{X}} (f(Y) - f(X)) \mathbb{K}(X, dY). \end{aligned}$$

Proof. Let $f, g \in C^2(\mathbb{X})$. In what follows, we show that for any $s, t \in [0, T]$ with $t \geq s$

$$\mathbb{E}[(f(\mathbf{Y}_t) - f(\mathbf{Y}_s))g(\mathbf{Y}_s)] = \mathbb{E}[g(\mathbf{Y}_s) \int_s^t \mathcal{R}(f)(u, \mathbf{Y}_u) du].$$

More precisely, we show that for any $s, t \in [0, T]$ with $t \geq s$

$$\mathbb{E}[(f(\mathbf{X}_t) - f(\mathbf{X}_s))g(\mathbf{X}_t)] = \mathbb{E}[-g(\mathbf{X}_t) \int_s^t \mathcal{R}(f)(u, \mathbf{X}_u) du].$$

Let $s, t \in [0, T]$, with $t \geq s$. Next, we denote for any $u \in [0, t]$ and $X \in \mathbb{X}$, $g(u, X) = \mathbb{E}[g(\mathbf{X}_t) \mid \mathbf{X}_u = X]$. Using A3, we have that for any $u \in [0, t]$ and $X \in \mathbb{X}$, $\partial_u g(u, X) + \mathcal{A}(g)(u, X) = 0$, i.e. g satisfies the backward Kolmogorov equation. For any $u \in [0, t]$ and $X \in \mathbb{X}$, we have

$$\begin{aligned} \mathcal{A}(fg)(u, X) &= \partial_u g(u, X) f(X) + \sum_{i=1}^N (\langle b(X_i), \nabla_{x_i} g(u, X) \rangle + \frac{1}{2} \Delta_{x_i} g(u, X_i)) f(X) \delta_{\mathbb{R}^d}(X_i) \\ &+ \sum_{i=1}^N (\langle b(X_i), \nabla_{x_i} f(X) \rangle + \frac{1}{2} \Delta_{x_i} f(X)) g(u, X) \delta_{\mathbb{R}^d}(X_i) \\ &+ \sum_{i=1}^N \delta_{\mathbb{R}^d}(X_i) \langle \nabla_{x_i} f(X), \nabla_{x_i} g(u, X) \rangle + \mathbb{J}(X, fg) \\ &= \partial_u g(u, X) f(X) + \mathcal{A}(g)(u, X) f(X) + \mathbb{J}(X, fg) - \mathbb{J}(X, g) f(X) \\ &+ \sum_{i=1}^N (\langle b(X_i), \nabla_{x_i} f(X) \rangle + \frac{1}{2} \Delta_{x_i} f(X)) g(u, X) \delta_{\mathbb{R}^d}(X_i) \\ &+ \sum_{i=1}^N \delta_{\mathbb{R}^d}(X_i) \langle \nabla_{x_i} f(X), \nabla_{x_i} g(u, X) \rangle \\ &= \sum_{i=1}^N (\langle b(X_i), \nabla_{x_i} f(X) \rangle + \frac{1}{2} \Delta_{x_i} f(X)) g(u, X) \delta_{\mathbb{R}^d}(X_i) \\ &+ \sum_{i=1}^N \delta_{\mathbb{R}^d}(X_i) \langle \nabla_{x_i} f(X), \nabla_{x_i} g(u, X) \rangle + \mathbb{J}(X, fg) - \mathbb{J}(X, g) f(X). \quad (8) \end{aligned}$$

Using A2, we have that for any $u \in [0, t]$ and $i \in \{1, \dots, N\}$

$$\begin{aligned} & \mathbb{E}[\delta_{\mathbb{R}^d}((\mathbf{X}_u)_i) \langle \nabla_{x_i} f(\mathbf{X}_u), \nabla_{x_i} g(u, \mathbf{X}_u) \rangle] \\ &= -\mathbb{E}[\delta_{\mathbb{R}^d}((\mathbf{X}_u)_i) g(u, \mathbf{X}_u) (\Delta_{x_i} f(\mathbf{X}_u) + \langle \nabla_{x_i} \log p_u(\mathbf{X}_u), \nabla_{x_i} f(\mathbf{X}_u) \rangle)]. \quad (9) \end{aligned}$$

In addition, we have that for any $X \in \mathbb{X}$ and $u \in [0, t]$, $\mathbb{J}(X, fg) - \mathbb{J}(X, g) f(X) = \int_{\mathbb{X}} g(u, Y) (f(Y) - f(X)) \mathbb{J}(X, dY)$. Using Lemma 1, we get

$$\mathbb{E}[\mathbb{J}(\mathbf{X}_u, fg) - \mathbb{J}(\mathbf{X}_u, f) g(u, \mathbf{X}_u)] = -\mathbb{E}[g(u, \mathbf{X}_u) \mathbb{K}(\mathbf{X}_u, f)]. \quad (10)$$

Therefore, using (8), (9) and (10), we have

$$\mathbb{E}[\mathcal{A}(fg)(u, \mathbf{X}_u)] = \mathbb{E}[-\mathcal{R}(f)(u, \mathbf{X}_u)g(u, \mathbf{X}_u)].$$

Finally, we have

$$\begin{aligned} \mathbb{E}[(f(\mathbf{X}_t) - f(\mathbf{X}_s))g(\mathbf{X}_t)] &= \mathbb{E}[g(t, \mathbf{X}_t)f(\mathbf{X}_t) - f(\mathbf{X}_s)g(s, \mathbf{X}_s)] \\ &= \mathbb{E}[\int_s^t \mathcal{A}(fg)(u, \mathbf{X}_u)du] \\ &= -\mathbb{E}[\int_s^t \mathcal{R}(f)(u, \mathbf{X}_u)g(u, \mathbf{X}_u)du] = -\mathbb{E}[g(\mathbf{X}_t) \int_s^t \mathcal{R}(f)(u, \mathbf{X}_u)du], \end{aligned}$$

which concludes the proof. \square

A.2.2 Intuitive Proof of Proposition 1

We recall Proposition 1.

Proposition 1. *The time reversal of a forward jump diffusion process given by drift $\vec{\mathbf{b}}_t$, diffusion coefficient \vec{g}_t , rate $\vec{\lambda}_t(n)$ and transition kernel $\sum_{i=1}^n K^{\text{del}}(i|n)\delta_{\text{del}(\mathbf{X},i)}(\mathbf{Y})$ is given by a jump diffusion process with drift $\overleftarrow{\mathbf{b}}_t^*(\mathbf{X})$, diffusion coefficient \overleftarrow{g}_t^* , rate $\overleftarrow{\lambda}_t^*(\mathbf{X})$ and transition kernel $\int_{\mathbf{y}^{\text{add}}} \sum_{i=1}^{n+1} A_t^*(\mathbf{y}^{\text{add}}, i|\mathbf{X})\delta_{\text{ins}(\mathbf{X},\mathbf{y}^{\text{add}},i)}(\mathbf{Y})d\mathbf{y}^{\text{add}}$ as defined below*

$$\begin{aligned} \overleftarrow{\mathbf{b}}_t^*(\mathbf{X}) &= \vec{\mathbf{b}}_t(\mathbf{X}) - \vec{g}_t^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{X}), \quad \overleftarrow{g}_t^* = \vec{g}_t, \\ \overleftarrow{\lambda}_t^*(\mathbf{X}) &= \vec{\lambda}_t(n+1) \frac{\sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} p_t(\text{ins}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)) d\mathbf{y}^{\text{add}}}{p_t(\mathbf{X})}, \\ A_t^*(\mathbf{y}^{\text{add}}, i|\mathbf{X}) &\propto p_t(\text{ins}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)) K^{\text{del}}(i|n+1). \end{aligned}$$

Diffusion part. Using standard diffusion models arguments such as [18] [45], we get

$$\overleftarrow{\mathbf{b}}_t^*(\mathbf{X}) = \vec{\mathbf{b}}_t(\mathbf{X}) - \vec{g}_t^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{X}|n).$$

Jump part. We use the flux equation from [45] which intuitively relates the probability flow going in the forward direction with the probability flow going the backward direction with equality being achieved at the time reversal.

$$\begin{aligned} p_t(\mathbf{X}) \overleftarrow{\lambda}_t^*(\mathbf{X}) \overleftarrow{K}_t^*(\mathbf{Y}|\mathbf{X}) &= p_t(\mathbf{Y}) \vec{\lambda}_t(\mathbf{Y}) \vec{K}_t(\mathbf{X}|\mathbf{Y}) \\ p_t(\mathbf{X}) \overleftarrow{\lambda}_t^*(\mathbf{X}) \int_{\mathbf{y}^{\text{add}}} \sum_{i=1}^{n+1} A_t^*(\mathbf{y}^{\text{add}}, i|\mathbf{X}) \delta_{\text{ins}(\mathbf{X},\mathbf{y}^{\text{add}},i)}(\mathbf{Y}) d\mathbf{y}^{\text{add}} \\ &= p_t(\mathbf{Y}) \vec{\lambda}_t(\mathbf{Y}) \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \delta_{\text{del}(\mathbf{Y},i)}(\mathbf{X}). \end{aligned} \quad (11)$$

To find $\overleftarrow{\lambda}_t^*(\mathbf{X})$, we sum and integrate both sides over m and \mathbf{y} , with $\mathbf{Y} = (m, \mathbf{y})$,

$$\begin{aligned} \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} p_t(\mathbf{X}) \overleftarrow{\lambda}_t^*(\mathbf{X}) \int_{\mathbf{y}^{\text{add}}} \sum_{i=1}^{n+1} A_t^*(\mathbf{y}^{\text{add}}, i|\mathbf{X}) \delta_{\text{ins}(\mathbf{X},\mathbf{y}^{\text{add}},i)}(\mathbf{Y}) d\mathbf{y}^{\text{add}} d\mathbf{y} \\ = \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} p_t(\mathbf{Y}) \vec{\lambda}_t(\mathbf{Y}) \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \delta_{\text{del}(\mathbf{Y},i)}(\mathbf{X}) d\mathbf{y}. \end{aligned}$$

Now we use the fact that $\delta_{\text{del}(\mathbf{Y},i)}(\mathbf{X})$ is 0 for any $m \neq n+1$,

$$\begin{aligned} p_t(\mathbf{X}) \overleftarrow{\lambda}_t^*(\mathbf{X}) &= \vec{\lambda}_t(n+1) \int_{\mathbf{y} \in \mathbb{R}^{(n+1)d}} p_t(\mathbf{Y}) \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \delta_{\text{del}(\mathbf{Y},i)}(\mathbf{X}) d\mathbf{y} \\ &= \vec{\lambda}_t(n+1) \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y} \in \mathbb{R}^{(n+1)d}} p_t(\mathbf{Y}) \delta_{\text{del}(\mathbf{Y},i)}(\mathbf{X}) d\mathbf{y}. \end{aligned}$$

Now letting $\mathbf{Y} = \text{ins}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)$,

$$\begin{aligned} p_t(\mathbf{X}) \overleftarrow{\lambda}_t^*(\mathbf{X}) &= \vec{\lambda}_t(n+1) \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} p_t(\text{ins}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)) d\mathbf{y}^{\text{add}} \\ \overleftarrow{\lambda}_t^*(\mathbf{X}) x &= \vec{\lambda}_t(n+1) \frac{\sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} p_t(\text{ins}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)) d\mathbf{y}^{\text{add}}}{p_t(\mathbf{X})}. \end{aligned}$$

To find $A_t^*(\mathbf{y}^{\text{add}}, i|\mathbf{X})$, we start from (11) and set $\mathbf{Y} = \text{ins}(\mathbf{X}, \mathbf{z}^{\text{add}}, j)$ to get

$$p_t(\mathbf{X}) \overleftarrow{\lambda}_t^*(\mathbf{X}) A_t^*(\mathbf{z}^{\text{add}}, j|\mathbf{X}) = p_t(\mathbf{Y}) \vec{\lambda}_t(n+1) K^{\text{del}}(j|n+1).$$

By inspection, we see immediately that

$$A_t^*(\mathbf{z}^{\text{add}}, j|\mathbf{X}) \propto p_t(\text{ins}(\mathbf{X}, \mathbf{z}^{\text{add}}, j)) K^{\text{del}}(j|n+1).$$

With a re-labeling of \mathbf{z}^{add} and j we achieve the desired form

$$A_t^*(\mathbf{y}^{\text{add}}, i|\mathbf{X}) \propto p_t(\text{ins}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)) K^{\text{del}}(i|n+1).$$

A.3 Proof of Proposition 2

In this section we prove Proposition 2 using the notation from the main paper by following the framework of [17]. We operate on a state space $\mathcal{X} = \bigcup_{n=1}^N \{n\} \times \mathbb{R}^{nd}$. On this space the gradient operator $\nabla : \mathcal{C}(\mathcal{X}, \mathbb{R}) \rightarrow \mathcal{C}(\mathcal{X}, \mathcal{X})$ is defined as $\nabla f(\mathbf{X}) = \nabla_{\mathbf{x}}^{(nd)} f(\mathbf{X})$ where $\nabla_{\mathbf{x}}^{(nd)}$ is the standard gradient operator defined as $\mathcal{C}(\mathbb{R}^{nd}, \mathbb{R}) \rightarrow \mathcal{C}(\mathbb{R}^{nd}, \mathbb{R}^{nd})$ with respect to $\mathbf{x} \in \mathbb{R}^{nd}$. We will write integration with respect to a probability measure defined on \mathcal{X} as an explicit sum over the number of components and integral over \mathbb{R}^{nd} with respect to a probability density defined on \mathbb{R}^{nd} i.e. $\int_{\mathbf{X}} f(\mathbf{X}) \mu(d\mathbf{X}) = \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} f(\mathbf{X}) p(n) p(\mathbf{x}|n) d\mathbf{x}$ where, for $A \subset \mathbb{R}^{nd}$, $\int_{(n,A)} \mu(d\mathbf{X}) = \int_{\mathbf{x} \in A} p(n) p(\mathbf{x}|n) d\mathbf{x}$. We will write $p(\mathbf{X})$ as shorthand for $p(n)p(\mathbf{x}|n)$.

Following, [17], we start by augmenting our space with a time variable so that operators become time inhomogeneous on the extended space. We write this as $\bar{\mathbf{X}} = (\mathbf{X}, t)$ where $\bar{\mathbf{X}}$ lives in the extended space $\mathcal{S} = \mathcal{X} \times \mathbb{R}_{\geq 0}$. In the proof, we use the infinitesimal generators for the forward and backward processes. An infinitesimal generator is defined as

$$\mathcal{A}(f)(\bar{\mathbf{X}}) = \lim_{t \rightarrow 0} \frac{\mathbb{E}_{p_{t|0}(\bar{\mathbf{Y}}|\bar{\mathbf{X}})}[f(\bar{\mathbf{Y}})] - f(\bar{\mathbf{X}})}{t}$$

and can be understood as a probabilistic version of a derivative. For our process on the augmented space \mathcal{S} , our generators decompose as $\mathcal{A} = \partial_t + \hat{\mathcal{A}}_t$ where $\hat{\mathcal{A}}_t$ operates only on the spatial components of $\bar{\mathbf{X}}$ i.e. \mathbf{X} [17].

We now define the spatial infinitesimal generators for our forward and backward process. We will change our treatment of the time variable compared to the main text. Both our forward and backward processes will run from $t = 0$ to $t = T$, with the true time reversal of \mathbf{X} following the forward process satisfying $(\mathbf{Y}_t)_{t \in [0, T]} = (\mathbf{X}_{T-t})_{t \in [0, T]}$. Further, we will write \overrightarrow{g}_t as g_t and $\overleftarrow{g}_t = g_{T-t}$ as we do not learn g and this is the optimal relation from the time reversal. We define

$$\hat{\mathcal{L}}_t(f)(\mathbf{X}) = \overrightarrow{\mathbf{b}}_t(\mathbf{X}) \cdot \nabla f(\mathbf{X}) + \frac{1}{2} g_t^2 \Delta f(\mathbf{X}) + \overrightarrow{\lambda}_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} f(\mathbf{Y}) (\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y})) d\mathbf{y},$$

as well as

$$\hat{\mathcal{K}}_t(f)(\mathbf{X}) = \overleftarrow{\mathbf{b}}_t^\theta(\mathbf{X}) \cdot \nabla f(\mathbf{X}) + \frac{1}{2} g_{T-t}^2 \Delta f(\mathbf{X}) + \overleftarrow{\lambda}_t^\theta(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} f(\mathbf{Y}) (\overleftarrow{K}_t^\theta(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y})) d\mathbf{y}$$

where $\Delta = (\nabla \cdot \nabla)$ is the Laplace operator and δ is a dirac delta on \mathcal{X} i.e. $\sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} \delta_{\mathbf{X}}(\mathbf{Y}) d\mathbf{y} = 1$ and $\sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} f(\mathbf{Y}) \delta_{\mathbf{X}}(\mathbf{Y}) d\mathbf{y} = f(\mathbf{X})$.

Verifying Assumption 1. The first step in the proof is to verify Assumption 1 in [17]. Letting $\nu_t(\mathbf{X}) = p_{T-t}(\mathbf{X})$, we assume we can write $\partial_t p_t(\mathbf{X}) = \hat{\mathcal{K}}_t^* p_t(\mathbf{X})$ in the form $\mathcal{M}\nu + c\nu = 0$ for some function $c : \mathcal{S} \rightarrow \mathbb{R}$, where \mathcal{M} is the generator of another auxiliary process on \mathcal{S} and $\hat{\mathcal{K}}_t^*$ is the adjoint operator which satisfies $\langle \hat{\mathcal{K}}_t^* f, h \rangle = \langle f, \hat{\mathcal{K}}_t h \rangle$ i.e.

$$\sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} h(\mathbf{X}) \hat{\mathcal{K}}_t^*(f)(\mathbf{X}) d\mathbf{x} = \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} f(\mathbf{X}) \hat{\mathcal{K}}_t(h)(\mathbf{X}) d\mathbf{x}$$

We now find $\hat{\mathcal{K}}_t^*$. We start by substituting in the form for $\hat{\mathcal{K}}_t$,

$$\sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} f(\mathbf{X}) \hat{\mathcal{K}}_t(h)(\mathbf{X}) d\mathbf{x} = \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} f(\mathbf{X}) \{ (\overleftarrow{\mathbf{b}}_t^\theta(\mathbf{x}) \cdot \nabla h)(\mathbf{X}) + \frac{1}{2} g_{T-t}^2 \Delta h(\mathbf{X}) + \overleftarrow{\lambda}_t^\theta(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} h(\mathbf{Y}) (\overleftarrow{K}_t^\theta(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y})) d\mathbf{y} \} d\mathbf{x}$$

We first focus on the RHS terms corresponding to the diffusion part of the process

$$\begin{aligned} & \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} f(\mathbf{X}) \{ (\overleftarrow{\mathbf{b}}_t^\theta \cdot \nabla h)(\mathbf{X}) + \frac{1}{2} g_{T-t}^2 \Delta h(\mathbf{X}) \} d\mathbf{x} \\ &= \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} f(\mathbf{X}) (\overleftarrow{\mathbf{b}}_t^\theta \cdot \nabla h)(\mathbf{X}) + \frac{1}{2} g_{T-t}^2 f(\mathbf{X}) \nabla \cdot \nabla h(\mathbf{X}) d\mathbf{x} \\ &= \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} f(\mathbf{X}) (\overleftarrow{\mathbf{b}}_t^\theta \cdot \nabla h)(\mathbf{X}) + \frac{1}{2} g_{T-t}^2 h(\mathbf{X}) \nabla \cdot \nabla f(\mathbf{X}) d\mathbf{x} \\ &= \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} -h(\mathbf{X}) \nabla \cdot (f \overleftarrow{\mathbf{b}}_t^\theta)(\mathbf{X}) + \frac{1}{2} g_{T-t}^2 h(\mathbf{X}) \nabla \cdot \nabla f(\mathbf{X}) d\mathbf{x} \\ &= \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} h(\mathbf{X}) \{ -\nabla \cdot (f \overleftarrow{\mathbf{b}}_t^\theta)(\mathbf{X}) + \frac{1}{2} g_{T-t}^2 \nabla \cdot \nabla f(\mathbf{X}) \} d\mathbf{x} \\ &= \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} h(\mathbf{X}) \{ -f(\mathbf{X}) \nabla \cdot \overleftarrow{\mathbf{b}}_t^\theta(\mathbf{X}) - \nabla f(\mathbf{X}) \cdot \overleftarrow{\mathbf{b}}_t^\theta(\mathbf{X}) + \frac{1}{2} g_{T-t}^2 \nabla \cdot \nabla f(\mathbf{X}) \} d\mathbf{x}. \end{aligned}$$

where we apply integration by parts twice to arrive at the third line and once to arrive at the fourth line. We now focus on the RHS term corresponding to the jump part of the process

$$\begin{aligned} & \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} f(\mathbf{X}) \{ \overleftarrow{\lambda}_t^\theta(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} h(\mathbf{Y}) (\overleftarrow{K}_t^\theta(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y})) d\mathbf{y} \} d\mathbf{x} \\ &= \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} h(\mathbf{Y}) \{ \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} f(\mathbf{X}) \overleftarrow{\lambda}_t^\theta(\mathbf{X}) (\overleftarrow{K}_t^\theta(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y})) d\mathbf{x} \} d\mathbf{y} \\ &= \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} h(\mathbf{X}) \{ \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} f(\mathbf{Y}) \overleftarrow{\lambda}_t^\theta(\mathbf{Y}) (\overleftarrow{K}_t^\theta(\mathbf{X}|\mathbf{Y}) - \delta_{\mathbf{Y}}(\mathbf{X})) d\mathbf{y} \} d\mathbf{x}, \end{aligned}$$

where on the last line we have relabelled \mathbf{X} to \mathbf{Y} and \mathbf{Y} to \mathbf{X} . Putting both re-arranged forms for the RHS together, we obtain

$$\begin{aligned} & \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} h(\mathbf{X}) \hat{\mathcal{K}}_t^*(f)(\mathbf{X}) d\mathbf{x} = \\ & \quad \sum_{n=1}^N \int_{\mathbf{x} \in \mathbb{R}^{nd}} h(\mathbf{X}) \{ -f(\mathbf{X}) \nabla \cdot \overleftarrow{\mathbf{b}}_t^\theta(\mathbf{X}) - \nabla f(\mathbf{X}) \cdot \overleftarrow{\mathbf{b}}_t^\theta(\mathbf{X}) + \frac{1}{2} g_{T-t}^2 \nabla \cdot \nabla f(\mathbf{X}) + \\ & \quad \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} f(\mathbf{Y}) \overleftarrow{\lambda}_t^\theta(\mathbf{Y}) (\overleftarrow{K}_t^\theta(\mathbf{X}|\mathbf{Y}) - \delta_{\mathbf{Y}}(\mathbf{X})) d\mathbf{y} \} d\mathbf{x}. \end{aligned}$$

We therefore have

$$\begin{aligned} \hat{\mathcal{K}}_t^*(f)(\mathbf{X}) &= -f(\mathbf{X}) \nabla \cdot \overleftarrow{\mathbf{b}}_t^\theta(\mathbf{X}) - \nabla f(\mathbf{X}) \cdot \overleftarrow{\mathbf{b}}_t^\theta(\mathbf{X}) + \frac{1}{2} g_{T-t}^2 \Delta f(\mathbf{X}) + \\ & \quad \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} f(\mathbf{Y}) \overleftarrow{\lambda}_t^\theta(\mathbf{Y}) (\overleftarrow{K}_t^\theta(\mathbf{X}|\mathbf{Y}) - \delta_{\mathbf{Y}}(\mathbf{X})) d\mathbf{y}. \end{aligned}$$

Now we re-write $\partial_t p_t(\mathbf{X}) = \hat{\mathcal{K}}_t^* p_t(\mathbf{x})$ in the form $\mathcal{M}\nu + c\nu = 0$. We start by re-arranging

$$\partial_t p_t(\mathbf{X}) = \hat{\mathcal{K}}_t^* p_t(\mathbf{X}) \implies 0 = \partial_t \nu_t(\mathbf{X}) + \hat{\mathcal{K}}_{T-t}^* \nu_t(\mathbf{X}).$$

Substituting in our form for $\hat{\mathcal{K}}_t^*$ we obtain

$$\begin{aligned} 0 &= \partial_t \nu_t(\mathbf{X}) - \nu_t(\mathbf{X}) \nabla \cdot \overleftarrow{\mathbf{b}}_{T-t}^\theta(\mathbf{X}) - \overleftarrow{\mathbf{b}}_{T-t}^\theta(\mathbf{X}) \cdot \nabla \nu_t(\mathbf{X}) + \frac{1}{2} g_t^2 \Delta \nu_t(\mathbf{X}) \\ & \quad + \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} \nu_t(\mathbf{Y}) \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) (\overleftarrow{K}_{T-t}^\theta(\mathbf{X}|\mathbf{Y}) - \delta_{\mathbf{X}}(\mathbf{Y})) d\mathbf{y} \end{aligned} \quad (12)$$

We define our auxiliary process to have generator $\mathcal{M} = \partial_t + \hat{\mathcal{M}}_t$ with

$$\hat{\mathcal{M}}_t(f)(\mathbf{X}) = \mathbf{b}_t^M(\mathbf{X}) \cdot \nabla f(\mathbf{X}) + \frac{1}{2} g_t^2 \Delta f(\mathbf{X}) + \lambda_t^M(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} f(\mathbf{Y}) (K_t^M(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y})) d\mathbf{y}$$

which is a jump diffusion process with drift \mathbf{b}_t^M , diffusion coefficient g_t , rate λ_t^M and transition kernel K_t^M . Then if we have $\mathbf{b}_t^M = -\overleftarrow{\mathbf{b}}_{T-t}^\theta$,

$$\lambda_t^M(\mathbf{X}) = \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) \overleftarrow{K}_{T-t}^\theta(\mathbf{X}|\mathbf{Y}) d\mathbf{y}, \quad (13)$$

and

$$K_t^M(\mathbf{Y}|\mathbf{X}) \propto \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) \overleftarrow{K}_{T-t}^\theta(\mathbf{X}|\mathbf{Y}). \quad (14)$$

Then we have (12) can be rewritten as

$$\begin{aligned} 0 &= \partial_t \nu_t(\mathbf{X}) - \nu_t(\mathbf{X}) \nabla \cdot \overleftarrow{\mathbf{b}}_{T-t}^\theta(\mathbf{X}) + \mathbf{b}_t^M(\mathbf{X}) \cdot \nabla \nu_t(\mathbf{X}) + \frac{1}{2} g_t^2 \Delta \nu_t(\mathbf{X}) \\ & \quad - \nu_t(\mathbf{X}) \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{X}) + \nu_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) \overleftarrow{K}_{T-t}^\theta(\mathbf{X}|\mathbf{Y}) d\mathbf{y} + \\ & \quad \lambda_t^M(\mathbf{x}) \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} \nu_t(\mathbf{Y}) (K_t^M(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y})) d\mathbf{y} \end{aligned}$$

which is in the form $\mathcal{M}(\nu)(\bar{\mathbf{X}}) + c(\bar{\mathbf{X}})\nu(\bar{\mathbf{X}}) = 0$ if we let

$$c(\bar{\mathbf{X}}) = -\nabla \cdot \overleftarrow{\mathbf{b}}_{T-t}^\theta(\bar{\mathbf{X}}) - \overleftarrow{\lambda}_{T-t}^\theta(\bar{\mathbf{X}}) + \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) \overleftarrow{K}_{T-t}^\theta(\bar{\mathbf{X}}|\mathbf{Y}) d\mathbf{y}.$$

Verifying Assumption 2. Now that we have verified Assumption 1, the second step in the proof is Assumption 2 from [17]. We assume there is a bounded measurable function $\alpha : \mathcal{S} \rightarrow (0, \infty)$ such

that $\alpha \mathcal{M}f = \mathcal{L}(f\alpha) - f\mathcal{L}\alpha$ for all functions $f : \mathcal{X} \rightarrow \mathbb{R}$ such that $f \in \mathcal{D}(\mathcal{M})$ and $f\alpha \in \mathcal{D}(\mathcal{L})$. Substituting in \mathcal{M} and \mathcal{L} we get

$$\begin{aligned}
& \alpha_t(\mathbf{X})[\partial_t f(\mathbf{X}) - \overleftarrow{\mathbf{b}}_{T-t}^\theta(\mathbf{X}) \cdot \nabla f(\mathbf{X}) \\
& \quad + \frac{1}{2}g_t^2 \Delta f(\mathbf{X}) + \lambda_t^M(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} f(\mathbf{Y})(K_t^M(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y}))d\mathbf{y}] \\
& = \partial_t(f\alpha_t)(\mathbf{X}) + \overrightarrow{\mathbf{b}}_t(\mathbf{X}) \cdot \nabla(f\alpha_t)(\mathbf{X}) + \frac{1}{2}g_t^2 \Delta(f\alpha_t)(\mathbf{X}) \\
& \quad + \overrightarrow{\lambda}_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} f(\mathbf{Y})\alpha_t(\mathbf{Y})(\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y}))d\mathbf{y} \\
& \quad - f(\mathbf{X})[\partial_t \alpha_t(\mathbf{X}) + \overrightarrow{\mathbf{b}}_t(\mathbf{X}) \cdot \nabla \alpha_t(\mathbf{X}) + \frac{1}{2}g_t^2 \Delta \alpha_t(\mathbf{X}) \\
& \quad + \overrightarrow{\lambda}_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} \alpha_t(\mathbf{Y})(\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y}))d\mathbf{y}] \tag{15}
\end{aligned}$$

Since f does not depend on time, $\partial_t f(\mathbf{X}) = 0$ and $\partial_t(f\alpha_t)(\mathbf{X}) = f(\mathbf{X})\partial_t \alpha_t(\mathbf{X})$ thus the ∂_t terms on the RHS also cancel out. Comparing terms on the LHS and RHS relating to the diffusion part of the process we obtain

$$\begin{aligned}
& -\alpha_t(\mathbf{X})(\overleftarrow{\mathbf{b}}_{T-t}^\theta(\mathbf{X}) \cdot \nabla f(\mathbf{X})) + \frac{1}{2}\alpha_t(\mathbf{X})g_t^2 \Delta f(\mathbf{X}) = \\
& \overrightarrow{\mathbf{b}}_t(\mathbf{X}) \cdot \nabla(f\alpha_t)(\mathbf{X}) + \frac{1}{2}g_t^2 \Delta(f\alpha_t)(\mathbf{X}) - f(\mathbf{X})\overrightarrow{\mathbf{b}}_t(\mathbf{X}) \cdot \nabla \alpha_t(\mathbf{X}) - \frac{1}{2}f(\mathbf{X})g_t^2 \Delta \alpha_t(\mathbf{X}).
\end{aligned}$$

Therefore, we get

$$\begin{aligned}
& -\alpha_t(\mathbf{X})(\overleftarrow{\mathbf{b}}_{T-t}^\theta(\mathbf{X}) \cdot \nabla f(\mathbf{X})) + \frac{1}{2}\alpha_t(\mathbf{X})g_t^2 \Delta f(\mathbf{X}) = \\
& \overrightarrow{\mathbf{b}}_t(\mathbf{X}) \cdot (f(\mathbf{X})\nabla \alpha_t(\mathbf{X}) + \alpha_t(\mathbf{X})\nabla f(\mathbf{X})) \\
& \quad + \frac{1}{2}g_t^2 (2\nabla f(\mathbf{X}) \cdot \nabla \alpha_t(\mathbf{X}) + f(\mathbf{X})\Delta \alpha_t(\mathbf{X}) + \alpha_t(\mathbf{X})\Delta f(\mathbf{X})) \\
& \quad - f(\mathbf{X})\overrightarrow{\mathbf{b}}_t(\mathbf{X}) \cdot \nabla \alpha_t(\mathbf{X}) - \frac{1}{2}f(\mathbf{X})g_t^2 \Delta \alpha_t(\mathbf{X}).
\end{aligned}$$

Simplifying the above expression, we get

$$\begin{aligned}
& -\alpha_t(\mathbf{X})(\overleftarrow{\mathbf{b}}_{T-t}^\theta(\mathbf{X}) \cdot \nabla f(\mathbf{X})) = \alpha_t(\mathbf{X})\overrightarrow{\mathbf{b}}_t(\mathbf{X}) \cdot \nabla f(\mathbf{X}) + g_t^2 \nabla f(\mathbf{X}) \cdot \nabla \alpha_t(\mathbf{X}) \\
& (-\alpha_t(\mathbf{X})\overleftarrow{\mathbf{b}}_{T-t}^\theta(\mathbf{X})) \cdot \nabla f(\mathbf{X}) = (\alpha_t(\mathbf{X})\overrightarrow{\mathbf{b}}_t(\mathbf{X}) + g_t^2 \nabla \alpha_t(\mathbf{X})) \cdot \nabla f(\mathbf{X}).
\end{aligned}$$

This is true for any f implying

$$-\alpha_t(\mathbf{X})\overleftarrow{\mathbf{b}}_{T-t}^\theta(\mathbf{X}) = \alpha_t(\mathbf{X})\overrightarrow{\mathbf{b}}_t(\mathbf{X}) + g_t^2 \nabla \alpha_t(\mathbf{X}).$$

This implies that $\alpha_t(\mathbf{X})$ satisfies

$$\nabla \log \alpha_t(\mathbf{X}) = -\frac{1}{g_t^2}(\overrightarrow{\mathbf{b}}_t(\mathbf{X}) + \overleftarrow{\mathbf{b}}_{T-t}^\theta(\mathbf{X})) \tag{16}$$

Comparing terms from the LHS and RHS of (15) relating to the jump part of the process we obtain

$$\begin{aligned}
& \alpha_t(\mathbf{X})\lambda_t^M(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} f(\mathbf{Y})(K_t^M(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y}))d\mathbf{y} = \\
& \overrightarrow{\lambda}_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} f(\mathbf{Y})\alpha_t(\mathbf{Y})(\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y}))d\mathbf{y} \\
& - f(\mathbf{X})\overrightarrow{\lambda}_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} \alpha_t(\mathbf{Y})(\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y}))d\mathbf{y}.
\end{aligned}$$

Hence, we have

$$\begin{aligned}
& \alpha_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} f(\mathbf{Y})\lambda_t^M(\mathbf{X})K_t^M(\mathbf{Y}|\mathbf{X})d\mathbf{y} - \alpha_t(\mathbf{X})\lambda_t^M(\mathbf{X})f(\mathbf{X}) = \\
& \overrightarrow{\lambda}_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} f(\mathbf{Y})\alpha_t(\mathbf{Y})\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X})d\mathbf{y} \\
& - f(\mathbf{X})\overrightarrow{\lambda}_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} \alpha_t(\mathbf{Y})\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X})d\mathbf{y}.
\end{aligned}$$

Recalling the definitions of $\lambda_t^M(\mathbf{X})$ and $K_t^M(\mathbf{Y}|\mathbf{X})$, (13) and (14), we get

$$\begin{aligned}
& \alpha_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} f(\mathbf{Y})\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y})\overleftarrow{K}_{T-t}^\theta(\mathbf{X}|\mathbf{Y})d\mathbf{y} \\
& \quad - \alpha_t(\mathbf{X})f(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y})\overleftarrow{K}_{T-t}^\theta(\mathbf{X}|\mathbf{Y})d\mathbf{y} = \\
& \overrightarrow{\lambda}_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} f(\mathbf{Y})\alpha_t(\mathbf{Y})\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X})d\mathbf{y} \\
& - f(\mathbf{X})\overrightarrow{\lambda}_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} \alpha_t(\mathbf{Y})\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X})d\mathbf{y}.
\end{aligned}$$

This equality is satisfied if $\alpha_t(\mathbf{X})$ follows the following relation

$$\alpha_t(\mathbf{Y}) = \alpha_t(\mathbf{X}) \frac{\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) \overleftarrow{K}_{T-t}^\theta(\mathbf{X}|\mathbf{Y})}{\overrightarrow{\lambda}_t(\mathbf{X}) \overrightarrow{K}_t(\mathbf{Y}|\mathbf{X})} \quad \text{for } n \neq m \quad (17)$$

We only require this relation to be satisfied for $n \neq m$ because both $\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X})$ and $\overleftarrow{K}_{T-t}^\theta(\mathbf{X}|\mathbf{Y})$ are 0 for $n = m$. We note at this point, as in [17], that if we have $\alpha_t(\mathbf{X}) = 1/p_t(\mathbf{X})$ and $\overleftarrow{\lambda}_{T-t}^\theta$ and $\overleftarrow{K}_{T-t}^\theta(\mathbf{X}|\mathbf{Y})$ equal to the true time-reversals, then both (16), and (17) are satisfied. However, $\alpha_t(\mathbf{X}) = 1/p_t(\mathbf{X})$ is not the only α_t to satisfy these equations. (16) and (17) can be thought of as enforcing a certain parameterization of the generative process in terms of α_t [17].

Concluding the proof. Now for the final part of the proof, we substitute our value for α into the \mathcal{I}_{ISM} loss from [17] which is equal to the negative of the evidence lower bound on $\mathbb{E}_{p_{\text{data}}(\mathbf{x}_0)}[\log p_0^\theta(\mathbf{X}_0)]$ up to a constant independent of θ . Defining $\beta_t(\mathbf{X}_t) = 1/\alpha_t(\mathbf{X}_t)$, we have

$$\mathcal{I}_{\text{ISM}}(\beta) = \int_0^T \mathbb{E}_{p_t(\mathbf{x}_t)} \left[\frac{\hat{\mathcal{L}}_t^* \beta_t(\mathbf{X}_t)}{\beta_t(\mathbf{X}_t)} + \hat{\mathcal{L}}_t \log \beta_t(\mathbf{X}_t) \right] dt.$$

We split the spatial infinitesimal generator of the forward process into the generator corresponding to the diffusion and the generator corresponding to the jump part, $\hat{\mathcal{L}} = \hat{\mathcal{L}}_t^{\text{diff}} + \hat{\mathcal{L}}_t^{\text{J}}$ with

$$\hat{\mathcal{L}}_t^{\text{diff}}(f)(\mathbf{X}) = \overrightarrow{\mathbf{b}}_t(\mathbf{X}) \cdot \nabla f(\mathbf{X}) + \frac{1}{2} g_t^2 \Delta f(\mathbf{X}).$$

and

$$\hat{\mathcal{L}}_t^{\text{J}}(f)(\mathbf{X}) = \overrightarrow{\lambda}_t(\mathbf{X}) \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} f(\mathbf{Y}) (\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X}) - \delta_{\mathbf{X}}(\mathbf{Y})) d\mathbf{y}.$$

By comparison with the approach to find the adjoint $\hat{\mathcal{K}}_t^*$, we also have $\hat{\mathcal{L}}_t^* = \hat{\mathcal{L}}_t^{\text{diff}*} + \hat{\mathcal{L}}_t^{\text{J}*}$ with

$$\hat{\mathcal{L}}_t^{\text{diff}*}(f)(\mathbf{X}) = -f(\mathbf{X}) \nabla \cdot \overrightarrow{\mathbf{b}}_t(\mathbf{X}) - \nabla f(\mathbf{X}) \cdot \overrightarrow{\mathbf{b}}_t(\mathbf{X}) + \frac{1}{2} g_t^2 \Delta f(\mathbf{X}).$$

In addition, we get

$$\hat{\mathcal{L}}_t^{\text{J}*}(f)(\mathbf{X}) = \sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} f(\mathbf{Y}) \overrightarrow{\lambda}_t(\mathbf{Y}) (\overrightarrow{K}_t(\mathbf{X}|\mathbf{Y}) - \delta_{\mathbf{X}}(\mathbf{Y})) d\mathbf{y}.$$

Finally, \mathcal{I}_{ISM} becomes

$$\begin{aligned} \mathcal{I}_{\text{ISM}}(\beta) &= \int_0^T \mathbb{E}_{p_t(\mathbf{x}_t)} \left[\frac{\hat{\mathcal{L}}_t^{\text{diff}*} \beta_t(\mathbf{X}_t)}{\beta_t(\mathbf{X}_t)} + \hat{\mathcal{L}}_t^{\text{diff}} \log \beta_t(\mathbf{X}_t) \right] dt + \int_0^T \mathbb{E}_{p_t(\mathbf{x}_t)} \left[\frac{\hat{\mathcal{L}}_t^{\text{J}*} \beta_t(\mathbf{X}_t)}{\beta_t(\mathbf{X}_t)} + \hat{\mathcal{L}}_t^{\text{J}} \log \beta_t(\mathbf{X}_t) \right] dt \\ &= \mathcal{I}_{\text{ISM}}^{\text{diff}}(\beta) + \mathcal{I}_{\text{ISM}}^{\text{J}}(\beta), \end{aligned}$$

where we have named the two terms corresponding to the diffusion and jump part of the process as $\mathcal{I}_{\text{ISM}}^{\text{diff}}$, $\mathcal{I}_{\text{ISM}}^{\text{J}}$ respectively. For the diffusion part of the loss, we use the denoising form of the objective proven in Appendix E of [17] which is equivalent to $\mathcal{I}_{\text{ISM}}^{\text{diff}}$ up to a constant independent of θ

$$\mathcal{I}_{\text{ISM}}^{\text{diff}}(\beta) = \int_0^T \mathbb{E}_{p_{0,t}(\mathbf{x}_0, \mathbf{x}_t)} \left[\frac{\hat{\mathcal{L}}_t^{\text{diff}}(p_{t|0}(\cdot|\mathbf{x}_0)\alpha_t(\cdot))(\mathbf{x}_t)}{p_{t|0}(\mathbf{x}_t|\mathbf{x}_0)\alpha_t(\mathbf{x}_t)} - \hat{\mathcal{L}}_t^{\text{diff}} \log(p_{t|0}(\cdot|\mathbf{x}_0)\alpha_t(\cdot))(\mathbf{x}_t) \right] dt + \text{const.}$$

To simplify this expression, we first re-arrange $\hat{\mathcal{L}}_t^{\text{diff}}(h)$ for some general function $h : \mathcal{S} \rightarrow \mathbb{R}$.

$$\begin{aligned} \frac{\hat{\mathcal{L}}_t^{\text{diff}}(h)}{h} - \hat{\mathcal{L}}_t^{\text{diff}}(\log h) &= \frac{\overrightarrow{\mathbf{b}}_t \cdot \nabla h}{h} + \frac{1}{2} g_t^2 \frac{\Delta h}{h} - \overrightarrow{\mathbf{b}}_t \cdot \nabla \log h - \frac{1}{2} g_t^2 \Delta \log h \\ &= \frac{1}{2} g_t^2 \left(\frac{\nabla \cdot \nabla h}{h} - \nabla \cdot \nabla \log h \right) \\ &= \frac{1}{2} g_t^2 \|\nabla \log h\|^2. \end{aligned}$$

Setting $h = p_{t|0}(\cdot|\mathbf{x}_0)\alpha_t(\cdot)$, our diffusion part of the loss becomes

$$\mathcal{I}_{\text{ISM}}^{\text{diff}}(\beta) = \frac{1}{2} \int_0^T g_t^2 \mathbb{E}_{p_{0,t}(\mathbf{x}_0, \mathbf{x}_t)} [\|\nabla \log p_{t|0}(\mathbf{X}_t|\mathbf{X}_0) + \nabla \log \alpha_t(\mathbf{X}_t)\|^2] dt + \text{const}$$

We then directly parameterize $\nabla \log \alpha_t(\mathbf{X}_t)$ as $-s_t^\theta(\mathbf{X}_t)$

$$\mathcal{I}_{\text{ISM}}^{\text{diff}}(\beta) = \frac{1}{2} \int_0^T g_t^2 \mathbb{E}_{p_{0,t}(\mathbf{x}_0, \mathbf{x}_t)} [\|\nabla \log p_{t|0}(\mathbf{X}_t|\mathbf{X}_0) - s_t^\theta(\mathbf{X}_t)\|^2] dt + \text{const.}$$

We now focus on the expectation within the integral to re-write it in an easy to calculate form

$$\begin{aligned} & \mathbb{E}_{p_{0,t}(\mathbf{X}_0, \mathbf{X}_t)} [\|\nabla \log p_{t|0}(\mathbf{X}_t|\mathbf{X}_0) - s_t^\theta(\mathbf{X}_t)\|^2] \\ & = \mathbb{E}_{p_{0,t}(\mathbf{X}_0, \mathbf{X}_t)} [\|s_t^\theta(\mathbf{X}_t)\|^2 - 2s_t^\theta(\mathbf{X}_t)^T \nabla \log p_{0,t}(\mathbf{X}_0, \mathbf{X}_t)] + \text{const} \end{aligned}$$

Now we note that we can re-write $\nabla \log p_{0,t}(\mathbf{X}_0, \mathbf{X}_t)$ using M_t where M_t is a mask variable $M_t \in \{0, 1\}^{n_0}$ that is 0 for components of \mathbf{X}_0 that have been deleted to get to \mathbf{X}_t and 1 for components that remain in \mathbf{X}_t .

$$\begin{aligned} \nabla \log p_{0,t}(\mathbf{X}_0, \mathbf{X}_t) & = \frac{1}{p_{0,t}(\mathbf{X}_0, \mathbf{X}_t)} \nabla p_{0,t}(\mathbf{X}_0, \mathbf{X}_t) \\ & = \frac{1}{p_{0,t}(\mathbf{X}_0, \mathbf{X}_t)} \nabla \sum_{M_t} p_{0,t}(\mathbf{X}_0, \mathbf{X}_t, M_t) \\ & = \sum_{M_t} \frac{1}{p_{0,t}(\mathbf{X}_0, \mathbf{X}_t)} \nabla p_{0,t}(\mathbf{X}_0, \mathbf{X}_t, M_t) \\ & = \sum_{M_t} \frac{p(n_t, M_t, \mathbf{X}_0)}{p_{0,t}(\mathbf{X}_0, \mathbf{X}_t)} \nabla p_{t|0}(\mathbf{x}_t | n_t, \mathbf{X}_0, M_t) \\ & = \sum_{M_t} \frac{p(M_t | \mathbf{X}_0, \mathbf{X}_t)}{p(\mathbf{x}_t | n_t, \mathbf{X}_0, M_t)} \nabla p_{t|0}(\mathbf{x}_t | n_t, \mathbf{X}_0, M_t) \\ & = \mathbb{E}_{p(M_t | \mathbf{X}_0, \mathbf{X}_t)} [\nabla \log p_{t|0}(\mathbf{x}_t | n_t, \mathbf{X}_0, M_t)] \end{aligned}$$

Substituting this back in we get

$$\begin{aligned} & \mathbb{E}_{p_{0,t}(\mathbf{X}_0, \mathbf{X}_t)} [\|\nabla \log p_{t|0}(\mathbf{X}_t|\mathbf{X}_0) - s_t^\theta(\mathbf{X}_t)\|^2] \\ & = \mathbb{E}_{p_{0,t}(\mathbf{X}_0, \mathbf{X}_t)} [\|s_t^\theta(\mathbf{X}_t)\|^2 - 2s_t^\theta(\mathbf{X}_t)^T \mathbb{E}_{p(M_t | \mathbf{X}_0, \mathbf{X}_t)} [\nabla \log p_{t|0}(\mathbf{x}_t | n_t, \mathbf{X}_0, M_t)]] + \text{const} \\ & = \mathbb{E}_{p_{0,t}(\mathbf{X}_0, \mathbf{X}_t, M_t)} [\|\nabla \log p_{t|0}(\mathbf{x}_t | n_t, \mathbf{X}_0, M_t) - s_t^\theta(\mathbf{X}_t)\|^2] + \text{const}. \end{aligned}$$

Therefore, the diffusion part of \mathcal{I}_{ISM} can be written as

$$\mathcal{I}_{\text{ISM}}^{\text{diff}}(\beta) = \frac{T}{2} \mathbb{E}_{\mathcal{U}(t;0,T) p_{0,t}(\mathbf{X}_0, \mathbf{X}_t, M_t)} [g_t^2 \|\nabla \log p_{t|0}(\mathbf{x}_t | n_t, \mathbf{X}_0, M_t) - s_t^\theta(\mathbf{X}_t)\|^2] + \text{const}.$$

We now focus on the jump part of the loss $\mathcal{I}_{\text{ISM}}^J$. We first substitute in $\hat{\mathcal{L}}_t^J$ and $\hat{\mathcal{L}}_t^{J*}$

$$\begin{aligned} \mathcal{I}_{\text{ISM}}^J & = \int_0^T \mathbb{E}_{p_t(\mathbf{X}_t)} [\sum_m \int_{\mathbf{Y} \in \mathbb{R}^{md}} \vec{\lambda}_t(\mathbf{Y}) \frac{\beta_t(\mathbf{Y})}{\beta_t(\mathbf{X}_t)} (\vec{K}_t(\mathbf{X}_t|\mathbf{Y}) - \delta_{\mathbf{Y}}(\mathbf{X}_t)) d\mathbf{y} + \\ & \quad \vec{\lambda}_t(\mathbf{X}_t) \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} \vec{K}_t(\mathbf{Y}|\mathbf{X}_t) \log \beta_t(\mathbf{Y}) d\mathbf{y} - \vec{\lambda}_t(\mathbf{X}_t) \log \beta_t(\mathbf{X}_t)] dt. \end{aligned} \quad (18)$$

Noting that $\beta_t(\mathbf{X}_t) = 1/\alpha_t(\mathbf{X}_t)$, we get

$$\frac{\beta_t(\mathbf{X}_t)}{\beta_t(\mathbf{Y})} = \frac{\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) \overleftarrow{K}_{T-t}^\theta(\mathbf{X}_t|\mathbf{Y})}{\overleftarrow{\lambda}_t(\mathbf{X}_t) \overleftarrow{K}_t(\mathbf{Y}|\mathbf{X}_t)} \quad \text{for } n_t \neq m \quad (19)$$

or swapping labels for \mathbf{X}_t and \mathbf{Y} ,

$$\frac{\beta_t(\mathbf{Y})}{\beta_t(\mathbf{X}_t)} = \frac{\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{X}_t) \overleftarrow{K}_{T-t}^\theta(\mathbf{Y}|\mathbf{X}_t)}{\overleftarrow{\lambda}_t(\mathbf{Y}) \overleftarrow{K}_t(\mathbf{X}_t|\mathbf{Y})} \quad \text{for } n_t \neq m \quad (20)$$

Substituting (19) into the second line and (20) into the first line of (18) and using the fact that $\vec{K}_t(\mathbf{X}_t|\mathbf{Y}) = 0$ for $n_t = m$, we obtain

$$\begin{aligned} \mathcal{I}_{\text{ISM}}^J & = \int_0^T \mathbb{E}_{p_t(\mathbf{X}_t)} [\sum_{m=1 \setminus n_t}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} \vec{\lambda}_t(\mathbf{Y}) \frac{\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{X}_t) \overleftarrow{K}_{T-t}^\theta(\mathbf{Y}|\mathbf{X}_t)}{\overleftarrow{\lambda}_t(\mathbf{Y}) \overleftarrow{K}_t(\mathbf{X}_t|\mathbf{Y})} \vec{K}_t(\mathbf{X}_t|\mathbf{Y}) d\mathbf{y} \\ & \quad - \sum_{m=1}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} \vec{\lambda}_t(\mathbf{Y}) \frac{\beta_t(\mathbf{Y})}{\beta_t(\mathbf{X}_t)} \delta_{\mathbf{Y}}(\mathbf{X}_t) d\mathbf{y} \\ & \quad + \vec{\lambda}_t(\mathbf{X}_t) \sum_{m=1 \setminus n_t}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} \vec{K}_t(\mathbf{Y}|\mathbf{X}_t) \{\log \beta_t(\mathbf{X}_t) - \log \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) \\ & \quad - \log \overleftarrow{K}_{T-t}^\theta(\mathbf{X}_t|\mathbf{Y}) + \log \overleftarrow{\lambda}_t(\mathbf{X}_t) + \log \overleftarrow{K}_t(\mathbf{Y}|\mathbf{X}_t)\} d\mathbf{y} \\ & \quad - \vec{\lambda}_t(\mathbf{X}_t) \log \beta_t(\mathbf{X}_t)] dt. \end{aligned}$$

Hence, we have

$$\begin{aligned} \mathcal{I}_{\text{ISM}}^J & = \int_0^T \mathbb{E}_{p_t(\mathbf{X}_t)} [\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{X}_t) \sum_{m=1 \setminus n_t}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} \overleftarrow{K}_{T-t}^\theta(\mathbf{Y}|\mathbf{X}_t) d\mathbf{y} - \vec{\lambda}_t(\mathbf{X}_t) \frac{\beta_t(\mathbf{X}_t)}{\beta_t(\mathbf{X}_t)} + \\ & \quad \vec{\lambda}_t(\mathbf{X}_t) \sum_{m=1 \setminus n_t}^N \int_{\mathbf{Y} \in \mathbb{R}^{md}} \vec{K}_t(\mathbf{Y}|\mathbf{X}_t) \{-\log \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) - \log \overleftarrow{K}_{T-t}^\theta(\mathbf{X}_t|\mathbf{Y})\} d\mathbf{y}] dt + \text{const}. \end{aligned}$$

This can be rewritten as

$$\mathcal{I}_{\text{ISM}}^J = \int_0^T \mathbb{E}_{p_t(\mathbf{X}_t)} [\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{X}_t) + \overrightarrow{\lambda}_t(\mathbf{X}_t) \mathbb{E}_{\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X}_t)} [-\log \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) - \log \overleftarrow{K}_{T-t}^\theta(\mathbf{X}_t|\mathbf{Y})]] dt + \text{const.}$$

Therefore, we have

$$\mathcal{I}_{\text{ISM}}^J = T \mathbb{E}_{\mathcal{U}(t;0,T)p_t(\mathbf{X}_t)\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X}_t)} [\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{X}_t) - \overrightarrow{\lambda}_t(\mathbf{X}_t) \log \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) - \overrightarrow{\lambda}_t(\mathbf{X}_t) \log \overleftarrow{K}_{T-t}^\theta(\mathbf{X}_t|\mathbf{Y})] + \text{const.}$$

Finally, using the definition of the forward and backward kernels, i.e. $\overrightarrow{K}_t(\mathbf{Y}|\mathbf{X}_t) = \sum_{i=1}^n K^{\text{del}}(i|n) \delta_{\text{del}(\mathbf{x},i)}(\mathbf{Y})$ and $\overleftarrow{K}_{T-t}^\theta(\mathbf{X}_t|\mathbf{Y}) = \int_{\mathbf{x}^{\text{add}}} \sum_{i=1}^n A_t^\theta(\mathbf{x}^{\text{add}}, i|\mathbf{Y}) \delta_{\text{ins}(\mathbf{Y}, \mathbf{x}^{\text{add}}, i)}(\mathbf{X}_t) d\mathbf{x}^{\text{add}}$, we get

$$\begin{aligned} \mathcal{I}_{\text{ISM}}^J &= T \mathbb{E}_{\mathcal{U}(t;0,T)p_t(\mathbf{X}_t)} [\sum_{m=1}^N \int_{\mathbf{y} \in \mathbb{R}^{md}} \sum_{i=1}^{n_t} K^{\text{del}}(i|n_t) \delta_{\text{del}(\mathbf{x}_t,i)}(\mathbf{Y}) (\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{X}_t) - \overrightarrow{\lambda}_t(\mathbf{X}_t) \log \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) \\ &\quad - \overrightarrow{\lambda}_t(\mathbf{X}_t) \log \overleftarrow{K}_{T-t}^\theta(\mathbf{X}_t|\mathbf{Y})) d\mathbf{y}] + \text{const} \end{aligned}$$

We get

$$\begin{aligned} \mathcal{I}_{\text{ISM}}^J &= T \mathbb{E}_{\mathcal{U}(t;0,T)p_t(\mathbf{X}_t)K^{\text{del}}(i|n_t)\delta_{\text{del}(\mathbf{x}_t,i)}(\mathbf{Y})} \\ &\quad [\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{X}_t) - \overrightarrow{\lambda}_t(\mathbf{X}_t) \log \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) - \overrightarrow{\lambda}_t(\mathbf{X}_t) \log \overleftarrow{K}_{T-t}^\theta(\mathbf{X}_t|\mathbf{Y})] + \text{const.} \end{aligned}$$

Therefore, we have

$$\begin{aligned} \mathcal{I}_{\text{ISM}}^J &= T \mathbb{E}_{\mathcal{U}(t;0,T)p_t(\mathbf{X}_t)K^{\text{del}}(i|n_t)\delta_{\text{del}(\mathbf{x}_t,i)}(\mathbf{Y})} \\ &\quad [\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{X}_t) - \overrightarrow{\lambda}_t(\mathbf{X}_t) \log \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) - \overrightarrow{\lambda}_t(\mathbf{X}_t) \log A_{T-t}^\theta(\mathbf{x}^{\text{add}}, i|\mathbf{Y})] + \text{const.} \end{aligned}$$

Putting are expressions for $\mathcal{I}_{\text{ISM}}^{\text{diff}}$ and $\mathcal{I}_{\text{ISM}}^J$ together we obtain

$$\begin{aligned} \mathcal{I}_{\text{ISM}} &= \frac{T}{2} \mathbb{E}[g_t^2 \|\nabla \log p_{t|0}(\mathbf{x}_t|n_t, \mathbf{X}_0, M_t) - s_t^\theta(\mathbf{X}_t)\|^2] + \\ &\quad T \mathbb{E}[\overleftarrow{\lambda}_{T-t}^\theta(\mathbf{X}_t) - \overrightarrow{\lambda}_t(\mathbf{X}_t) \log \overleftarrow{\lambda}_{T-t}^\theta(\mathbf{Y}) - \overrightarrow{\lambda}_t(\mathbf{X}_t) \log A_{T-t}^\theta(\mathbf{x}^{\text{add}}, i|\mathbf{Y})] + \text{const.} \end{aligned}$$

We get that $-\mathcal{I}_{\text{ISM}}$ gives us our evidence lower bound on $\mathbb{E}_{p_{\text{data}}(\mathbf{X}_0)}[\log p_0^\theta(\mathbf{X}_0)]$ up to a constant that does not depend on θ . In the main text we have used a time notation such that the backward process runs backwards from $t = T$ to $t = 0$. To align with the notation of time used in the main text we change $T - t$ to t on subscripts for $\overleftarrow{\lambda}_{T-t}^\theta$ and A_{T-t}^θ . We also will use the fact that $\overrightarrow{\lambda}_t(\mathbf{X}_t)$ depends only on the number of components in \mathbf{X}_t , $\overrightarrow{\lambda}_t(\mathbf{X}_t) = \overrightarrow{\lambda}_t(n_t)$.

$$\begin{aligned} \mathcal{L}(\theta) &= -\frac{T}{2} \mathbb{E}[g_t^2 \|s_t^\theta(\mathbf{X}_t) - \nabla_{\mathbf{x}_t} \log p_{t|0}(\mathbf{x}_t|n_t, \mathbf{X}_0, M_t)\|^2] + \\ &\quad T \mathbb{E}[-\overleftarrow{\lambda}_t^\theta(\mathbf{X}_t) + \overrightarrow{\lambda}_t(n_t) \log \overleftarrow{\lambda}_t^\theta(\mathbf{Y}) + \overrightarrow{\lambda}_t(n_t) \log A_t^\theta(\mathbf{x}^{\text{add}}, i|\mathbf{Y})] + \text{const.} \end{aligned}$$

Tightness of the lower bound Now that we have derived the ELBO as in Proposition 2, we show that the maximizers of the ELBO are tight, i.e. that they close the variational gap. We do this by proving the general ELBO presented in [17] has this property and therefore ours, which is a special case of this general ELBO, also has that the optimum parameters close the variational gap.

To state our proposition, we recall the setting of [17]. The forward noising process is denoted $(\mathbf{Y}_t)_{t \geq 0}$ and associated with an infinitesimal generator $\hat{\mathcal{L}}$ its extension $(t, \mathbf{Y}_t)_{t \geq 0}$ is associated with the infinitesimal generator \mathcal{L} , i.e. $\mathcal{L} = \partial_t + \hat{\mathcal{L}}$. We also define the score-matching operator Φ given for any f for which it is defined by

$$\Phi(f) = \mathcal{L}(f)/f - \mathcal{L}(\log(f)).$$

We recall that according to [17, Equation (8)] and under [17, Assumption 1, Assumption2], we have

$$\log p_T(\mathbf{Y}_0) \geq \mathbb{E}[\log p_0(\mathbf{Y}_T) - \int_0^T \mathcal{L}(v/\beta)/(v/\beta) + \mathcal{L}(\log \beta) dt],$$

with $v_t = p_{T-t}$ for any $t \in [0, T]$. We define the *variational gap* Gap as follows

$$\text{Gap} = \mathbb{E}[\log p_T(\mathbf{Y}_0) - \log p_0(\mathbf{Y}_T) + \int_0^T \mathcal{L}(v/\beta)/(v/\beta) + \mathcal{L}(\log \beta) dt].$$

In addition, using Itô Formula, we have that $\log v_T(\mathbf{Y}_T) - \log v_0(\mathbf{Y}_0) = \int_0^T \mathcal{L}(v) dt$. Assuming that $\mathbb{E}[|\log v_T(\mathbf{Y}_T) - \log v_0(\mathbf{Y}_0)|] < +\infty$, we get

$$\text{Gap} = \mathbb{E}[\int_0^T -\mathcal{L}(\log v) + \mathcal{L}(v/\beta)/(v/\beta) + \mathcal{L}(\log \beta) dt] = \mathbb{E}[\int_0^T \Phi(v/\beta) dt].$$

In particular, using [17, Proposition 1], we get that $\text{Gap} \geq 0$ and $\text{Gap} = 0$ if and only if $\beta \propto v$. In addition, the ELBO is maximized if and only if $\beta \propto v$, see [17, Equation 10] and the remark that follows. Therefore, we have that: if we maximize the ELBO then the ELBO is tight. Combining this with the fact that the ELBO is maximized at the time reversal [17], then we have that when our jump diffusion parameters match the time reversal, our variational gap is 0.

Other approaches. Another way to derive the ELBO is to follow the steps of [15] directly, since [17] is a general framework extending this approach. The key formulae to derive the result and the ELBO is 1) a Feynman-Kac formula 2) a Girsanov formula. In the case of jump diffusions (with jump in \mathbb{R}^d) a Girsanov formula has been established by [22]. Extending this result to one-point compactification space would allow us to prove directly Proposition 2 without having to rely on the general framework of [17].

A.4 Proof of Proposition 3

We start by recalling the form for the time reversal given in Proposition 1

$$\overleftarrow{\lambda}_t^*(\mathbf{X}) = \overrightarrow{\lambda}_t(n+1) \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} p_t(\text{ins}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)) d\mathbf{y}^{\text{add}} / p_t(\mathbf{X}).$$

We then introduce a marginalization over \mathbf{X}_0

$$\begin{aligned} \overleftarrow{\lambda}_t^*(\mathbf{X}) &= \overrightarrow{\lambda}_t(n+1) \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} \sum_{n_0} \int_{\mathbf{x}_0} p_{0,t}(\mathbf{X}_0, \text{ins}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)) d\mathbf{x}_0 d\mathbf{y}^{\text{add}} / p_t(\mathbf{X}) \\ &= \overrightarrow{\lambda}_t(n+1) \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} \sum_{n_0} \int_{\mathbf{x}_0} \frac{p_{0,t}(\mathbf{X}_0)}{p_t(\mathbf{X})} p_{t|0}(\text{ins}(\mathbf{X}, \mathbf{y}^{\text{add}}, i) | \mathbf{X}_0) d\mathbf{x}_0 d\mathbf{y}^{\text{add}} \\ &= \overrightarrow{\lambda}_t(n+1) \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} \sum_{n_0} \int_{\mathbf{x}_0} \frac{p_{0|t}(\mathbf{X}_0 | \mathbf{X})}{p_{t|0}(\mathbf{X} | \mathbf{X}_0)} p_{t|0}(\text{ins}(\mathbf{X}, \mathbf{y}^{\text{add}}, i) | \mathbf{X}_0) d\mathbf{x}_0 d\mathbf{y}^{\text{add}} \\ &= \overrightarrow{\lambda}_t(n+1) \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} \sum_{n_0} \int_{\mathbf{x}_0} \frac{p_{0|t}(n_0 | \mathbf{X}) p_{0,t}(\mathbf{x}_0 | \mathbf{X}, n_0)}{p_{t|0}(n | \mathbf{X}_0) p_{t|0}(\mathbf{x} | \mathbf{X}_0, n)} \times \\ &\quad p_{t|0}(n+1 | \mathbf{X}_0) p_{t|0}(\mathbf{z}(\mathbf{X}, \mathbf{y}^{\text{add}}, i) | \mathbf{X}_0, n+1) d\mathbf{x}_0 d\mathbf{y}^{\text{add}} \end{aligned}$$

where $(n+1, \mathbf{z}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)) = \text{ins}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)$. Now using the fact the forward component deletion process does not depend on \mathbf{x}_0 , only n_0 , we have $p_{t|0}(n | \mathbf{X}_0) = p_{t|0}(n | n_0)$ and $p_{t|0}(n+1 | \mathbf{X}_0) = p_{t|0}(n+1 | n_0)$. Using this result, we get

$$\begin{aligned} \overleftarrow{\lambda}_t^*(\mathbf{X}) &= \overrightarrow{\lambda}_t(n+1) \sum_{n_0} \left\{ \frac{p_{t|0}(n+1 | n_0)}{p_{t|0}(n | n_0)} p_{0|t}(n_0 | \mathbf{X}) \times \right. \\ &\quad \left. \int_{\mathbf{x}_0} \frac{\sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} p_{t|0}(\mathbf{z}(\mathbf{X}, \mathbf{y}^{\text{add}}, i) | \mathbf{X}_0, n+1) d\mathbf{y}^{\text{add}}}{p_{t|0}(\mathbf{x} | \mathbf{X}_0, n)} p_{0|t}(\mathbf{x}_0 | \mathbf{X}, n_0) d\mathbf{x}_0 \right\}. \quad (21) \end{aligned}$$

We now focus on the probability ratio within the integral over \mathbf{x}_0 . We will show that this ratio is 1. We start with the numerator, introducing a marginalization over possible mask variables between \mathbf{X}_0 and $(n+1, \mathbf{z})$, denoted $M^{(n+1)}$ with $M^{(n+1)}$ having $n+1$ ones and $n_0 - (n+1)$ zeros.

$$\begin{aligned} &\sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} p_{t|0}(\mathbf{z}(\mathbf{X}, \mathbf{y}^{\text{add}}, i) | \mathbf{X}_0, n+1) d\mathbf{y}^{\text{add}} \\ &= \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \sum_{M^{(n+1)}} \int_{\mathbf{y}^{\text{add}}} p_{t|0}(M^{(n+1)}, \mathbf{z}(\mathbf{X}, \mathbf{y}^{\text{add}}, i) | \mathbf{X}_0, n+1) d\mathbf{y}^{\text{add}} \\ &= \sum_{M^{(n+1)}} \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) p_{t|0}(M^{(n+1)} | \mathbf{X}_0, n+1) \int_{\mathbf{y}^{\text{add}}} p_{t|0}(\mathbf{z}(\mathbf{X}, \mathbf{y}^{\text{add}}, i) | \mathbf{X}_0, n+1, M^{(n+1)}) d\mathbf{y}^{\text{add}} \end{aligned}$$

Now, for our forward process we have

$$p_{t|0}(\mathbf{z}(\mathbf{X}, \mathbf{y}^{\text{add}}, i) | \mathbf{X}_0, n+1, M^{(n+1)}) = \prod_{j=1}^{n+1} \mathcal{N}(\mathbf{z}^{(j)}; \sqrt{\alpha_t} M^{(n+1)}(\mathbf{X}_0)^j, (1 - \alpha_t) I_d)$$

where \mathbf{z} is shorthand for $\mathbf{z}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)$, $\mathbf{z}^{(j)}$ is the vector in \mathbb{R}^d for the j th component of \mathbf{z} and $M^{(n+1)}(\mathbf{X}_0)^j$ is the vector in \mathbb{R}^d corresponding to the component in \mathbf{X}_0 corresponding to the j th one in the $M^{(n+1)}$ mask. Integrating out \mathbf{y}^{add} we have

$$\int_{\mathbf{y}^{\text{add}}} p_{t|0}(\mathbf{z}(\mathbf{X}, \mathbf{y}^{\text{add}}, i) | \mathbf{X}_0, n+1, M^{(n+1)}) d\mathbf{y}^{\text{add}} = \prod_{j=1}^n \mathcal{N}(\mathbf{x}^{(j)}; \sqrt{\alpha_t} M^{(n+1) \setminus i}(\mathbf{X}_0)^j, (1 - \alpha_t) I_d),$$

where $M^{(n+1) \setminus i}$ denotes a mask variable obtained by setting the i th one of $M^{(n+1)}$ to zero. Hence, we have

$$\begin{aligned} &\sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} p_{t|0}(\mathbf{z}(\mathbf{X}, \mathbf{y}^{\text{add}}, i) | \mathbf{X}_0, n+1) d\mathbf{y}^{\text{add}} \\ &= \sum_{M^{(n+1)}} \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) p_{t|0}(M^{(n+1)} | \mathbf{X}_0, n+1) \\ &\quad \prod_{j=1}^n \mathcal{N}(\mathbf{x}^{(j)}; \sqrt{\alpha_t} M^{(n+1) \setminus i}(\mathbf{X}_0)^j, (1 - \alpha_t) I_d). \quad (22) \end{aligned}$$

We now re-write the denominator from (21) introducing a marginalization over mask variables, $M^{(n)}$

$$p_{t|0}(\mathbf{x} | \mathbf{X}_0, n) = \sum_{M^{(n)}} p_{t|0}(M^{(n)} | \mathbf{X}_0, n) p_{t|0}(\mathbf{x} | M^{(n)}, \mathbf{X}_0, n). \quad (23)$$

We use the following recursion for the probabilities assigned to mask variables

$$p_{t|0}(M^{(n)}|\mathbf{X}_0, n) = \sum_{M^{(n+1)}} \sum_{i=1}^{n+1} \mathbb{I}\{M^{(n+1)\setminus i} = M^{(n)}\} K^{\text{del}}(i|n+1) p_{t|0}(M^{(n+1)}|\mathbf{X}_0, n+1).$$

Substituting this into (23) gives

$$\begin{aligned} p_{t|0}(\mathbf{x}|\mathbf{X}_0, n) &= \sum_{M^{(n)}} \sum_{M^{(n+1)}} \sum_{i=1}^{n+1} \mathbb{I}\{M^{(n+1)\setminus i} = M^{(n)}\} K^{\text{del}}(i|n+1) \times \\ &\quad p_{t|0}(M^{(n+1)}|\mathbf{X}_0, n+1) p_{t|0}(\mathbf{x}|M^{(n)}, \mathbf{X}_0, n) \\ &= \sum_{M^{(n)}} \sum_{M^{(n+1)}} \sum_{i=1}^{n+1} \mathbb{I}\{M^{(n+1)\setminus i} = M^{(n)}\} K^{\text{del}}(i|n+1) \\ &\quad \times p_{t|0}(M^{(n+1)}|\mathbf{X}_0, n+1) \prod_{j=1}^n \mathcal{N}(\mathbf{x}^{(j)}; \sqrt{\alpha_t} M^{(n)}(\mathbf{X}_0)^j, (1-\alpha_t)I_d) \\ &= \sum_{M^{(n+1)}} \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) p_{t|0}(M^{(n+1)}|\mathbf{X}_0, n+1) \\ &\quad \times \prod_{j=1}^n \mathcal{N}(\mathbf{x}^{(j)}; \sqrt{\alpha_t} M^{(n+1)\setminus i}(\mathbf{X}_0)^j, (1-\alpha_t)I_d). \end{aligned}$$

By comparing with (22), we can see that

$$p_{t|0}(\mathbf{x}|\mathbf{X}_0, n) = \sum_{i=1}^{n+1} K^{\text{del}}(i|n+1) \int_{\mathbf{y}^{\text{add}}} p_{t|0}(\mathbf{z}(\mathbf{X}, \mathbf{y}^{\text{add}}, i)|\mathbf{X}_0, n+1) d\mathbf{y}^{\text{add}}.$$

This shows that the probability ratio in (21) is 1. Therefore, we have

$$\begin{aligned} \overleftarrow{\lambda}_t^*(\mathbf{X}) &= \overrightarrow{\lambda}_t(n+1) \sum_{n_0} \left\{ \frac{p_{t|0}(n+1|n_0)}{p_{t|0}(n|n_0)} p_{0|t}(n_0|\mathbf{X}) \int_{\mathbf{x}_0} p_{0|t}(\mathbf{x}_0|\mathbf{X}, n_0) d\mathbf{x}_0 \right\} \\ &= \overrightarrow{\lambda}_t(n+1) \sum_{n_0} \frac{p_{t|0}(n+1|n_0)}{p_{t|0}(n|n_0)} p_{0|t}(n_0|\mathbf{X}), \end{aligned}$$

which concludes the proof.

$p_{t|0}(n|n_0)$ can be analytically calculated when $\overrightarrow{\lambda}_t(n)$ is of a simple enough form. When $\overrightarrow{\lambda}_t(n)$ does not depend on n then the dimension deletion process simply becomes a time inhomogeneous Poisson process. Therefore, we would have

$$p_{t|0}(n|n_0) = \frac{(\int_0^t \overrightarrow{\lambda}_s ds)^{n_0-n}}{(n_0-n)!} \exp(-\int_0^t \overrightarrow{\lambda}_s ds).$$

In our experiments we set $\overrightarrow{\lambda}_t(n=1) = 0$ to stop the dimension deletion process when we reach a single component. If we have $\overrightarrow{\lambda}_t(n) = \overrightarrow{\lambda}_t(m)$ for all $n, m > 1$ then we can still use the time inhomogeneous Poisson process formula for $n > 1$ and find the probability for $n = 1$, $p_{t|0}(n=1|n_0)$ by requiring $p_{t|0}(n|n_0)$ to be a valid normalized distribution. Therefore, for the case that $\overrightarrow{\lambda}_t(n) = \overrightarrow{\lambda}_t(m)$ for all $n, m > 1$ and $\overrightarrow{\lambda}_t(n=1) = 0$, we have

$$p_{t|0}(n|n_0) = \begin{cases} \frac{(\int_0^t \overrightarrow{\lambda}_s ds)^{n_0-n}}{(n_0-n)!} \exp(-\int_0^t \overrightarrow{\lambda}_s ds) & 1 < n \leq n_0 \\ 1 - \sum_{m=2}^{n_0} \frac{(\int_0^t \overrightarrow{\lambda}_s ds)^{n_0-m}}{(n_0-m)!} \exp(-\int_0^t \overrightarrow{\lambda}_s ds) & n = 1 \end{cases}$$

In cases where $\overrightarrow{\lambda}_t(n)$ depends on n not just for $n = 1$, $p_{t|0}(n|n_0)$ can become more difficult to calculate analytically. However, since the probability distributions are all 1-dimensional over n , it is very cheap to simply simulate the forward dimension deletion process many times and empirically estimate $p_{t|0}(n|n_0)$ although we do not need to do this for our experiments.

A.5 The Objective is Maximized at the Time Reversal

In this section, we analyze the objective $\mathcal{L}(\theta)$ as a standalone object and determine the optimum values for s_t^θ , $\overleftarrow{\lambda}_t^\theta$ and A_t^θ directly. This is in order to gain intuition directly into the learning signal of $\mathcal{L}(\theta)$ without needing to refer to stochastic process theory.

The definition of $\mathcal{L}(\theta)$ as in the main text is

$$\begin{aligned} \mathcal{L}(\theta) &= -\frac{T}{2} \mathbb{E}[g_t^\theta \|s_t^\theta(\mathbf{X}_t) - \nabla_{\mathbf{x}_t} \log p_{t|0}(\mathbf{x}_t|\mathbf{X}_0, n_t, M_t)\|^2] + \\ &\quad T \mathbb{E}[-\overleftarrow{\lambda}_t^\theta(\mathbf{X}_t) + \overrightarrow{\lambda}_t(n_t) \log \overleftarrow{\lambda}_t^\theta(\mathbf{Y}) + \overrightarrow{\lambda}_t(n_t) \log A_t^\theta(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y})] + C. \end{aligned}$$

with the expectations taken over $\mathcal{U}(t; 0, T) p_{0,t}(\mathbf{X}_0, \mathbf{X}_t, M_t) K^{\text{del}}(i|n_t) \delta_{\text{del}(\mathbf{X}_t, i)}(\mathbf{Y})$.

Continuous optimum. We start by analysing the objective for s_t^θ . This part of $\mathcal{L}(\theta)$ can be written as

$$-\frac{1}{2} \int_0^T g_t^2 \mathbb{E}_{p_{0,t}(\mathbf{x}_0, \mathbf{x}_t, M_t)} [\|s_t^\theta(\mathbf{X}_t) - \nabla_{\mathbf{x}_t} \log p_{t|0}(\mathbf{x}_t | \mathbf{X}_0, n_t, M_t)\|^2] dt$$

We now use the fact that the function that minimizes an L_2 regression problem $\min_f \mathbb{E}_{p(x,y)} [\|f(x) - y\|^2]$ is the conditional expectation of the target $f^*(x) = \mathbb{E}_{p(y|x)}[y]$. Therefore the optimum value for $s_t^\theta(\mathbf{X}_t)$ is

$$\begin{aligned} s_t^*(\mathbf{X}_t) &= \mathbb{E}_{p(M_t, \mathbf{x}_0 | \mathbf{X}_t)} [\nabla_{\mathbf{x}_t} \log p_{t|0}(\mathbf{x}_t | \mathbf{X}_0, n_t, M_t)] \\ &= \sum_{M_t} \sum_{n_0=1}^N \int_{\mathbf{x}_0 \in \mathbb{R}^{n_0 d}} p(M_t, n_0, \mathbf{x}_0 | \mathbf{X}_t) \nabla_{\mathbf{x}_t} \log p_{t|0}(\mathbf{x}_t | \mathbf{x}_0, n_0, n_t, M_t) d\mathbf{x}_0 \\ &= \sum_{M_t} \sum_{n_0=1}^N \int_{\mathbf{x}_0 \in \mathbb{R}^{n_0 d}} \frac{p(M_t, n_0, \mathbf{x}_0 | \mathbf{X}_t)}{p_{t|0}(\mathbf{x}_0 | \mathbf{x}_0, n_0, n_t, M_t)} \nabla_{\mathbf{x}_t} p_{t|0}(\mathbf{x}_t | \mathbf{x}_0, n_0, n_t, M_t) d\mathbf{x}_0 \\ &= \sum_{M_t} \sum_{n_0=1}^N \int_{\mathbf{x}_0 \in \mathbb{R}^{n_0 d}} \frac{p(\mathbf{x}_0, n_0, n_t, M_t)}{p(n_t, \mathbf{x}_t)} \nabla_{\mathbf{x}_t} p_{t|0}(\mathbf{x}_t | \mathbf{x}_0, n_0, n_t, M_t) d\mathbf{x}_0 \\ &= \frac{1}{p(n_t, \mathbf{x}_t)} \sum_{M_t} \sum_{n_0=1}^N \int_{\mathbf{x}_0 \in \mathbb{R}^{n_0 d}} \nabla_{\mathbf{x}_t} p(\mathbf{x}_t, \mathbf{x}_0, n_0, n_t, M_t) d\mathbf{x}_0 \\ &= \frac{1}{p(n_t, \mathbf{x}_t)} \nabla_{\mathbf{x}_t} \sum_{M_t} \sum_{n_0=1}^N \int_{\mathbf{x}_0 \in \mathbb{R}^{n_0 d}} p(\mathbf{x}_t, \mathbf{x}_0, n_0, n_t, M_t) d\mathbf{x}_0 \\ &= \frac{1}{p(n_t, \mathbf{x}_t)} \nabla_{\mathbf{x}_t} p(\mathbf{x}_t, n_t) = \nabla_{\mathbf{x}_t} \log p(\mathbf{X}_t). \end{aligned}$$

Therefore, the optimum value for $s_t^\theta(\mathbf{X}_t)$ is $\nabla_{\mathbf{x}_t} \log p(\mathbf{X}_t)$ which is the value that gives $\overleftarrow{\mathbf{b}}_t$ to be the time reversal of $\overrightarrow{\mathbf{b}}_t$ as stated in Proposition 1.

Jump rate optimum. The learning signal for $\overleftarrow{\lambda}_t^\theta$ comes from these two terms in $\mathcal{L}(\theta)$

$$T \mathbb{E}[-\overleftarrow{\lambda}_t^\theta(\mathbf{Z}_t) + \overrightarrow{\lambda}_t(n_t) \log \overleftarrow{\lambda}_t^\theta(\mathbf{Y})] \quad (24)$$

This expectation is maximized when for each test input \mathbf{Z} and test time t , we have the following expression maximized

$$-p_t(\mathbf{Z}) \overleftarrow{\lambda}_t^\theta(\mathbf{Z}) + \sum_{i=1}^{n_z+1} \int_{\mathbf{y}^{\text{add}}} p_t(\text{ins}(\mathbf{Z}, \mathbf{y}^{\text{add}}, i)) K^{\text{del}}(i | n_z + 1) d\mathbf{y}^{\text{add}} \times \overrightarrow{\lambda}_t(n_z + 1) \log \overleftarrow{\lambda}_t^\theta(\mathbf{Z}),$$

because $p_t(\mathbf{Z})$ is the probability \mathbf{Z} gets drawn as a full sample from the forward process and $\sum_{i=1}^{n_z+1} \int_{\mathbf{y}^{\text{add}}} p_t(\text{ins}(\mathbf{Z}, \mathbf{y}^{\text{add}}, i)) K^{\text{del}}(i | n_z + 1) d\mathbf{y}^{\text{add}}$ is the probability that a sample one component bigger than \mathbf{Z} gets drawn from the forward process and then a component is deleted to get to \mathbf{Z} . Therefore the first probability is the probability that test input \mathbf{Z} and test time t appear as the first term in (24) whereas the second probability is the probability that test input \mathbf{Z} and test time t appear as the second term in (24).

We now use the fact that, for constants b and c ,

$$\text{argmax}_a \quad -ba + c \log a = \frac{c}{b}.$$

We therefore have the optimum $\overleftarrow{\lambda}_t^\theta(\mathbf{Z})$ as

$$\overleftarrow{\lambda}_t^*(\mathbf{Z}) = \overrightarrow{\lambda}_t(n_z + 1) \frac{\sum_{i=1}^{n_z+1} \int_{\mathbf{y}^{\text{add}}} p_t(\text{ins}(\mathbf{Z}, \mathbf{y}^{\text{add}}, i)) K^{\text{del}}(i | n_z + 1) d\mathbf{y}^{\text{add}}}{p_t(\mathbf{Z})}$$

which is the form for the time-reversal given in Proposition (1).

Jump kernel optimum. Finally, we analyse the part of $\mathcal{L}(\theta)$ for learning $A_t^\theta(\mathbf{x}_t^{\text{add}}, i | \mathbf{Y})$,

$$\begin{aligned} &T \mathbb{E}[\overrightarrow{\lambda}_t(n_t) \log A_t^\theta(\mathbf{x}_t^{\text{add}}, i | \mathbf{Y})] \\ &= \int_0^T \mathbb{E}_{p_t(\mathbf{x}_t)} K^{\text{del}}(i | n_t) \delta_{\text{del}(\mathbf{x}_t, i)}(\mathbf{Y}) [\overrightarrow{\lambda}_t(n_t) \log A_t^\theta(\mathbf{x}_t^{\text{add}}, i | \mathbf{Y})] dt \\ &= \int_0^T \mathbb{E}_{p_t(n_t)} [\overrightarrow{\lambda}_t(n_t) \mathbb{E}_{p_t(\mathbf{x}_t | n_t)} K^{\text{del}}(i | n_t) \delta_{\text{del}(\mathbf{x}_t, i)}(\mathbf{Y}) [\log A_t^\theta(\mathbf{x}_t^{\text{add}}, i | \mathbf{Y})]] dt. \end{aligned}$$

We now re-write the joint probability distribution that the inner expectation is taken with respect to,

$$p_t(\mathbf{x}_t | n_t) K^{\text{del}}(i | n_t) \delta_{\text{del}(\mathbf{x}_t, i)}(\mathbf{Y}) = \tilde{p}(\mathbf{Y} | n_t) p(\mathbf{x}_t^{\text{add}}, i | \mathbf{Y}) \delta_{\mathbf{y}}(\mathbf{x}_t^{\text{base}}).$$

with

$$\tilde{p}(\mathbf{Y}|n_t) = \sum_{i=1}^{n_t} \int_{\mathbf{x}_t} p_t(\mathbf{x}_t|n_t) K^{\text{del}}(i|n_t) \delta_{\text{del}(\mathbf{x}_t, i)}(\mathbf{Y}) d\mathbf{x}_t,$$

and

$$p(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y}) \propto p_t(\mathbf{x}_t|n_t) K^{\text{del}}(i|n_t),$$

and $\mathbf{x}_t^{\text{base}} \in \mathbb{R}^{(n_t-1)d}$ referring to the $n_t - 1$ components of \mathbf{x}_t , that are not $\mathbf{x}_t^{\text{add}}$ i.e. $\mathbf{X}_t = \text{ins}((\mathbf{x}_t^{\text{base}}, n_t - 1), \mathbf{x}_t^{\text{add}}, i)$. We then have

$$\begin{aligned} & T \mathbb{E}[\bar{\lambda}_t(n_t) \log A_t^\theta(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y})] \\ &= \int_0^T \mathbb{E}_{p_t(n_t)}[\bar{\lambda}_t(n_t) \mathbb{E}_{\tilde{p}(\mathbf{Y}|n_t)p(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y})\delta_{\mathbf{y}}(\mathbf{x}_t^{\text{base}})}[\log A_t^\theta(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y})]] dt \\ &= \int_0^T \mathbb{E}_{p_t(n_t)}[\bar{\lambda}_t(n_t) \mathbb{E}_{\tilde{p}(\mathbf{Y}|n_t)p(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y})\delta_{\mathbf{y}}(\mathbf{x}_t^{\text{base}})}[\log A_t^\theta(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y})]] dt \\ &\quad - \int_0^T \mathbb{E}_{p_t(n_t)}[\bar{\lambda}_t(n_t) \mathbb{E}_{\tilde{p}(\mathbf{Y}|n_t)p(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y})\delta_{\mathbf{y}}(\mathbf{x}_t^{\text{base}})}[\log p(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y})]] dt + \text{const} \\ &= \int_0^T \mathbb{E}_{p_t(n_t)}[\bar{\lambda}_t(n_t) \mathbb{E}_{\tilde{p}(\mathbf{Y}|n_t)\delta_{\mathbf{y}}(\mathbf{x}_t^{\text{base}})}[-\text{KL}(p(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y}) || A_t^\theta(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y}))]] dt + \text{const}. \end{aligned}$$

Therefore, the optimum $A_t^\theta(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y})$ which maximizes this part of $\mathcal{L}(\theta)$ is

$$A_t^*(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y}) = p(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y}) \propto p_t(\mathbf{X}_t) K^{\text{del}}(i|n_t).$$

which is the same form as given in Proposition 1.

B Training Objective

We estimate our objective $\mathcal{L}(\theta)$ by taking minibatches from the expectation $\mathcal{U}(t; 0, T) p_{0,t}(\mathbf{X}_0, \mathbf{X}_t, M_t) K^{\text{del}}(i|n_t) \delta_{\text{del}(\mathbf{X}_t, i)}(\mathbf{Y})$. We first sample $t \sim \mathcal{U}(t; 0, T)$ and then take samples from our dataset $\mathbf{X}_0 \sim p_{\text{data}}(\mathbf{X}_0)$. In order to sample $p_{t|0}(\mathbf{X}_t, M_t|\mathbf{X}_0)$ we need to both add noise, delete dimensions and sample a mask variable. Since the Gaussian noising process is isotropic, we can add a suitable amount of noise to all dimensions of \mathbf{X}_0 and then delete dimensions of that noised full dimensional value. More specifically, we first sample $\tilde{\mathbf{X}}_t = (n_0, \tilde{\mathbf{x}}_t)$ with $\tilde{\mathbf{x}}_t \sim \mathcal{N}(\tilde{\mathbf{x}}_t; \sqrt{\alpha_t} \mathbf{x}_0, (1 - \alpha_t) I_{n_0 d})$ for $\alpha_t = \exp(-\int_0^t \beta(s) ds)$ using the analytic forward equations for the VP-SDE derived in [3]. Then we sample the number of dimensions to delete. This is simple to do when our rate function is independent of n except for the case when $n = 1$ at which it is zero. We simply sample a Poisson random variable with mean parameter $\int_0^t \bar{\lambda}_s ds$ and then clamp its value such that the maximum number of possible components that are deleted is $n_0 - 1$. This gives the appropriate distribution over n , $p_{t|0}(n|n_0)$ as given in Section A.4. To sample which dimensions are deleted, we can sample $K^{\text{del}}(i_1|n_0) K^{\text{del}}(i_2|n_0 - 1) \dots K^{\text{del}}(i_{n_0 - n_t}|n_t + 1)$ from which we can create the mask M_t and apply it to $\tilde{\mathbf{X}}_t$ to obtain \mathbf{X}_t , $\mathbf{X}_t = M_t(\tilde{\mathbf{X}}_t)$. When $K^{\text{del}}(i|n) = 1/n$ this is especially simple to do by simply randomly permuting the components of $\tilde{\mathbf{X}}_t$, and then removing the final $n_0 - n_t$ components.

As is typically done in standard diffusion models, we parameterize s_t^θ in terms of a noise prediction network that predicts ϵ where $\mathbf{x}_t = \sqrt{\alpha_t} M_t(\mathbf{x}_0) + \sqrt{1 - \alpha_t} \epsilon$, $\epsilon \sim \mathcal{N}(0, I_{n_t d})$. We then re-weight the score loss in time such that we have a uniform weighting in time rather than the ‘likelihood weighting’ with g_t^2 [3, 21]. Our objective to learn s_t^θ then becomes

$$-\mathbb{E}_{\mathcal{U}(t; 0, T) p_{\text{data}}(\mathbf{X}_0) p(M_t, n_t|\mathbf{X}_0) \mathcal{N}(\epsilon; 0, I_{n_t d})} [\|\epsilon_t^\theta(\mathbf{X}_t) - \epsilon\|^2]$$

$$\text{with } \mathbf{x}_t = \sqrt{\alpha_t} M_t(\mathbf{x}_0) + \sqrt{1 - \alpha_t} \epsilon, s_t^\theta(\mathbf{X}_t) = \frac{-1}{\sqrt{1 - \alpha_t}} \epsilon_t^\theta(\mathbf{X}_t).$$

Further, by using the parameterization given in Proposition 3, we can directly supervise the value of $p_{0|t}^\theta(n_0|\mathbf{X}_t)$ by adding an extra term to our objective. We can treat the learning of $p_{0|t}^\theta(n_0|\mathbf{X}_t)$ as a standard prediction task where we aim to predict n_0 given access to \mathbf{X}_t . A standard objective for learning $p_{0|t}^\theta(n_0|\mathbf{X}_t)$ is then the cross entropy

$$\max_{\theta} \mathbb{E}_{p_{0,t}(\mathbf{X}_0, \mathbf{X}_t)} \left[\log p_{0|t}^\theta(n_0|\mathbf{X}_t) \right]$$

Our augmented objective then becomes

$$\tilde{\mathcal{L}}(\theta) = T\mathbb{E}\left[-\frac{1}{2}\|\epsilon_t^\theta(\mathbf{X}_t) - \epsilon\|^2 - \overleftarrow{\lambda}_t^\theta(\mathbf{X}_t) + \overrightarrow{\lambda}_t(n_t) \log \overleftarrow{\lambda}_t^\theta(\mathbf{Y}) + \overrightarrow{\lambda}_t(n_t) \log A_t^\theta(\mathbf{x}_t^{\text{add}}, i|\mathbf{Y}) + \gamma \log p_{0|t}^\theta(n_0|\mathbf{X}_t)\right] \quad (25)$$

where the expectation is taken with respect to

$$\mathcal{U}(t; 0, T) p_{\text{data}}(\mathbf{X}_0) p(M_t, n_t|\mathbf{X}_0) \mathcal{N}(\epsilon; 0, I_{n_t d}) K^{\text{del}}(i|n_t) \delta_{\text{del}(\mathbf{x}_t, i)}(\mathbf{Y})$$

where $\mathbf{x}_t = \sqrt{\alpha_t} M_t(\mathbf{x}_0) + \sqrt{1 - \alpha_t} \epsilon$ and γ is a loss weighting term for the cross entropy loss.

C Trans-Dimensional Diffusion Guidance

To guide an unconditionally trained model such that it generates datapoints consistent with conditioning information, we use the reconstruction guided sampling approach introduced in [9]. Our conditioning information will be the values for some of the components of \mathbf{X}_0 , and thus the guidance should guide the generative process such that the rest of the components of the generated datapoint are consistent with those observed components. Following the notation of [9], we denote the observed components as $\mathbf{x}^a \in \mathbb{R}^{n_a d}$ and the components to be generated as $\mathbf{x}^b \in \mathbb{R}^{n_b d}$. Our trained score function $s_t^\theta(\mathbf{X}_t)$ approximates $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{X}_t)$ whereas we would like the score to approximate $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{X}_t|\mathbf{x}_0^a)$. In order to do this, we will need to augment our unconditional score $s_t^\theta(\mathbf{X}_t)$ such that it incorporates the conditioning information.

We first focus on the dimensions of the score vector corresponding to \mathbf{x}^a . These can be calculated analytically from the forward process

$$\nabla_{\mathbf{x}_t^a} \log p(\mathbf{X}_t|\mathbf{x}_0^a) = \nabla_{\mathbf{x}_t^a} \log p_{t|0}(\mathbf{x}_t^a|\mathbf{x}_0^a, n_t)$$

with $p_{t|0}(\mathbf{x}_t^a|\mathbf{x}_0^a, n_t) = \mathcal{N}(\mathbf{x}_t^a; \sqrt{\alpha_t} \mathbf{x}_0^a, (1 - \alpha_t) I_{n_a d})$. Note that we assume a correspondence between \mathbf{x}_t^a and \mathbf{x}_0^a . For example, in video if we condition on the first and last frame, we assume that the first and last frame of the current noisy \mathbf{x}_t correspond to \mathbf{x}_0^a and guide them towards their observed values. For molecules, the point cloud is permutation invariant and so we can simply assume the first n_a components of \mathbf{x}_t correspond to \mathbf{x}_0^a and guide them to their observed values.

Now we analyse the dimensions of the score vector corresponding to \mathbf{x}^b . We split the score as

$$\nabla_{\mathbf{x}_t^b} \log p(\mathbf{X}_t|\mathbf{x}_0^a) = \nabla_{\mathbf{x}_t^b} \log p(\mathbf{x}_0^a|\mathbf{X}_t) + \nabla_{\mathbf{x}_t^b} \log p_t(\mathbf{X}_t)$$

$p(\mathbf{x}_0^a|\mathbf{X}_t)$ is intractable to calculate directly and so, following [9], we approximate it with $\mathcal{N}(\mathbf{x}_0^a; \hat{\mathbf{x}}_0^{\theta a}(\mathbf{X}_t), \frac{1 - \alpha_t}{\alpha_t} I_{n_a d})$ where $\hat{\mathbf{x}}_0^{\theta a}(\mathbf{X}_t)$ is a point estimate of \mathbf{x}_0^a given from $s_t^\theta(\mathbf{X}_t)$ calculated as

$$\hat{\mathbf{x}}_0^{\theta a}(\mathbf{X}_t) = \frac{\mathbf{x}_t^a + (1 - \alpha_t) s_t^\theta(\mathbf{X}_t)^a}{\sqrt{\alpha_t}}$$

where again we have assumed a correspondence between \mathbf{x}_t^a and \mathbf{x}_0^a . Our approximation for $\nabla_{\mathbf{x}_t^b} \log p(\mathbf{x}_0^a|\mathbf{X}_t)$ is then

$$\nabla_{\mathbf{x}_t^b} \log p(\mathbf{x}_0^a|\mathbf{X}_t) \approx -\nabla_{\mathbf{x}_t^b} \frac{\alpha_t}{2(1 - \alpha_t)} \|\mathbf{x}_0^a - \hat{\mathbf{x}}_0^{\theta a}(\mathbf{X}_t)\|^2$$

which can be calculated by differentiating through the score network s_t^θ .

We approximate $\overleftarrow{\lambda}_t^*(\mathbf{X}_t|\mathbf{x}_0^a)$ and $A_t^*(\mathbf{y}^{\text{add}}, i|\mathbf{X}_t, \mathbf{x}_0^a)$, with their unconditional forms $\overleftarrow{\lambda}_t^\theta(\mathbf{X}_t)$ and $A_t^\theta(\mathbf{y}^{\text{add}}, i|\mathbf{X}_t)$. We find this approximation still leads to valid generations because the guidance of the score network s_t^θ , results in \mathbf{X}_t containing the conditioning information which in turn leads to $\overleftarrow{\lambda}_t^\theta(\mathbf{X}_t)$ guiding the number of components in \mathbf{X}_t to be consistent with the conditioning information too as verified in our experiments. Further, any errors in the approximation for $A_t^\theta(\mathbf{y}^{\text{add}}, i|\mathbf{X}_t)$ are fixed by further applications of the guided score function, highlighting the benefits of our combined autoregressive and diffusion based approach.

D Experiment Details

Our code is available at <https://github.com/andrew-cr/jump-diffusion>

D.1 Molecules

D.1.1 Network Architecture

Backbone For our backbone network architecture, we used the EGNN used in [8]. This is a specially designed graph neural network applied to the point cloud treating it as a fully connected graph. A special equivariant update is used, operating only on distances between atoms. We refer to [8] for the specific details on the architecture. We used the same size network as used in [8]’s QM9 experiments, specifically there are 9 layers, with a hidden node feature size of 256. The output of the EGNN is fed into a final output projection layer to give the score network output $s_t^\theta(\mathbf{X}_t)$.

Component number prediction To obtain $p_{0|t}^\theta(n_0|\mathbf{X}_t)$, we take the embedding produced by the EGNN before the final output embedding layer and pass it through 8 transformer layers each consisting of a self-attention block and an MLP block applied channel wise. Our transformer model dimension is 128 and so we project the EGNN embedding output down to 128 before entering into the transformer layers. We then take the output of the transformer and take the average embedding over all nodes. This embedding is then passed through a final projection layer to give softmax logits over the $p_{0|t}^\theta(n_0|\mathbf{X}_t)$ distribution.

Autoregressive Distribution Our $A_t^\theta(\mathbf{y}^{\text{add}}, i|\mathbf{X}_t)$ network has to predict the position and features for a new atom when it is added to the molecule. Since the point cloud is permutation invariant, we do not need to predict i and so we just need to parameterize $A_t^\theta(\mathbf{y}^{\text{add}}|\mathbf{X}_t)$. We found the network to perform the best if the network first predicts the nearest atom and then a vector from that atom to the location of the new atom. To achieve this, we first predict softmax logits for a distribution over the nearest atom by applying a projection to the embedding output from the previously described transformer block. During training, the output of this distribution can be directly supervised by a cross entropy loss. Given the nearest atom, we then need to predict the position and features of the new atom to add. We do this by passing in the embedding generated by the EGNN and original point cloud features into a new transformer block of the same size as that used for $p_{0|t}^\theta(n_0|\mathbf{X}_t)$. We also input the distances from the nearest atom to all other atoms in the molecule currently as an additional feature. To obtain the position of the new atom, we will take a weighted sum of all the vectors between the nearest atom and other atoms in the molecule. This is to make it easy for the network to create new atoms ‘in plane’ with existing atoms which is useful for e.g. completing rings that have to remain in the same plane. To calculate the weights for the vectors, we apply an output projection to the output of the transformer block. The new atom features (atom type and charge) are generated by a separate output projection from the transformer block. For the position and features, $A_t^\theta(\mathbf{y}^{\text{add}}|\mathbf{X}_t)$ outputs both a mean and a standard deviation for a Gaussian distribution. For the position distribution, we set the standard deviation to be isotropic to remain equivariant to rotations. In total our model has around 7.3 million parameters.

D.1.2 Training

We train our model for 1.3 million iterations at a batch size of 64. We use the Adam optimizer with learning rate 0.00003. We also keep a running exponential moving average of the network weights that is used during sampling as is standard for training diffusion models [2, 3, 16] with a decay parameter of 0.9999. We train on the 100K molecules contained in the QM9 training split. We model hydrogens explicitly. Training a model requires approximately 7 days on a single GPU which was done on an Academic cluster.

In [8] the atom type is encoded as a one-hot vector and diffused as a continuous variable along with the positions and charge values for all atoms. They found that multiplying the one-hot vectors by 0.25 to boost performance by allowing the atom-type to be decided later on in the diffusion process. We instead multiply the one-hot vectors by 4 so that atom-type is decided early on in the diffusion process which improves our guided performance when conditioning on certain atom-types being

present. We found our model is robust to this change and achieves similar sample quality to [8] as shown in Table 2.

When deleting dimensions, we first shuffle the ordering of the nodes and then delete the final $n_0 - n_t$ nodes. The cross entropy loss weighting in (25) is set to 1.

Following [8] we train our model to operate within the center of mass (CoM) zero subspace of possible molecule positions. The means, throughout the forward and backward process, the average position of an atom is 0. In our transdimensional framework, this is achieved by first deleting any atoms required under the forward component deletion process. We then move the molecule such that its CoM is 0. We then add CoM free noise such that the noisy molecule also has CoM= 0. Our score model s_t^θ is parameterized through a noise prediction model ϵ_t^θ which is trained to predict the CoM free noise that was added. Therefore, our score network learns suitable directions to maintain the process on the CoM= 0 subspace. For the position prediction from $A_t^\theta(\mathbf{y}^{\text{add}}|\mathbf{X}_t)$ we train it to predict the new atom position from the current molecules reference frame. When the new atom is added, we then update all atom positions such that CoM= 0 is maintained.

D.1.3 Sampling

During sampling we found that adding corrector steps [3] improved sample quality. Intuitively, corrector steps form a process that has $p_t(\mathbf{X})$ as its stationary distribution rather than the process progressing toward $p_0(\mathbf{X})$. We use the same method to determine the corrector step size ζ as in [3]. For the conditional generation tasks, we also found it useful to include corrector steps for the component generation process. As shown in [29], corrector steps in discrete spaces can be achieved by simulating with a rate that is the addition of the forward and backward rates. We achieve this in the context of trans-dimensional modeling by first simulating a possible insertion using $\overleftarrow{\lambda}_t^\theta$ and then simulating a possible deletion using $\overrightarrow{\lambda}_t$. We describe our overall sampling algorithm in Algorithm 2.

Algorithm 2: Sampling the Generative Process with Corrector Steps

Input: Number of corrector steps C
 $t \leftarrow T$
 $\mathbf{X} \sim p_{\text{ref}}(\mathbf{X}) = \mathbb{I}\{n = 1\}\mathcal{N}(\mathbf{x}; 0, I_d)$
while $t > 0$ **do**
 if $u < \overleftarrow{\lambda}_t^\theta(\mathbf{X})\delta t$ with $u \sim \mathcal{U}(0, 1)$ **then**
 Sample $\mathbf{x}^{\text{add}}, i \sim A_t^\theta(\mathbf{x}^{\text{add}}, i|\mathbf{X})$
 $\mathbf{X} \leftarrow \text{ins}(\mathbf{X}, \mathbf{x}^{\text{add}}, i)$
 end
 $\mathbf{x} \leftarrow \mathbf{x} - \overleftarrow{\mathbf{b}}_t^\theta(\mathbf{X})\delta t + g_t\sqrt{\delta t}\epsilon$ with $\epsilon \sim \mathcal{N}(0, I_{nd})$
 for $c = [1, \dots, C]$ **do**
 $\mathbf{x} \leftarrow \mathbf{x} + \zeta s_{t-\delta t}^\theta(\mathbf{X}) + \sqrt{2\zeta}\epsilon$ with $\epsilon \sim \mathcal{N}(0, I_{nd})$
 if $u < \overleftarrow{\lambda}_{t-\delta t}^\theta(\mathbf{X})\delta t$ with $u \sim \mathcal{U}(0, 1)$ **then**
 Sample $\mathbf{x}^{\text{add}}, i \sim A_{t-\delta t}^\theta(\mathbf{x}^{\text{add}}, i|\mathbf{X})$
 $\mathbf{X} \leftarrow \text{ins}(\mathbf{X}, \mathbf{x}^{\text{add}}, i)$
 end
 if $u < \overrightarrow{\lambda}_{t-\delta t}(n)\delta t$ with $u \sim \mathcal{U}(0, 1)$ **then**
 $\mathbf{X} \leftarrow \text{del}(\mathbf{X}, i)$ with $i \sim K^{\text{del}}(i|n)$
 end
 end
 $\mathbf{X} \leftarrow (n, \mathbf{x}), t \leftarrow t - \delta t$
end

D.1.4 Evaluation

Unconditional For our unconditional sampling evaluation, we start adding corrector steps when $t < 0.1T$ in the backward process and use 5 corrector steps without the corrector steps on the number of components. We set $\delta = 0.05$ for $t > 0.5T$ and $\delta = 0.001$ for $t < 0.5T$ such that the total number

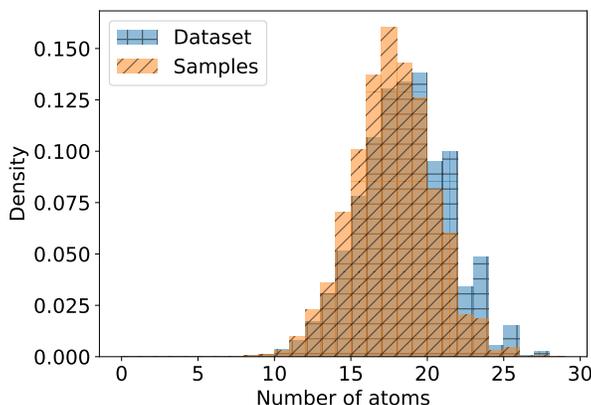


Figure 6: Distribution of the size of molecules in the QM9 dataset as measured through the number of atoms versus the distribution of the size of molecules generated by our unconditional model.

of network evaluations is 1000. We show the distribution of sizes of molecules generated by our model in Figure 6 and show more unconditional samples in Figure 7. We find our model consistently generates realistic molecules and achieves a size distribution similar to the training dataset even though this is not explicitly trained and arises from sampling our backward rate $\overleftarrow{\lambda}_t^\theta$. Since we are numerically integrating a continuous time process and approximating the true time reversal rate $\overleftarrow{\lambda}_t^*$, some approximation error is expected. For this experiment, sampling all of our models and ablations takes approximately 2 GPU days on Nvidia 1080Ti GPUs.

Conditional For evaluating applying conditional diffusion guidance to our model, we choose 10 conditioning tasks that each result in a different distribution of target dimensions. The task is to produce molecules that include at least a certain number of target atom types. We then guide the first set of atoms generated by the model to have these desired atom types. The tasks chosen are given in Table 5. Molecules in the training dataset that meet the conditions in each task have a different distribution of sizes. The tasks were chosen so that we have an approximately linearly increasing mean number of atoms for molecules that meet the condition. We also require that there are at least 100 examples of molecules that meet the condition within the training dataset.

For sampling when using conditional diffusion guidance, we use 3 corrector steps throughout the backward process with $\delta t = 0.001$. For these conditional tasks, we include the corrector steps on the number of components. We show the distribution of dimensions for each task from the training dataset and from our generated samples in Figure 8. Our metrics are calculated by first drawing 1000 samples for each conditioning task and then finding the Hellinger distance between the size distribution generated by our method (orange diagonal hashing in Figure 8) and the size distribution for molecules in the training dataset that match the conditions of the task (green no hashing in Figure 8). We find that indeed our model when guided by diffusion guidance can automatically produce a size distribution close to the ground truth size distribution found in the dataset for that conditioning value. We show samples generated by our conditionally guided model in Figure 9. We can see that our model can generate realistic molecules that include the required atom types and are of a suitable size. For this experiment, sampling all of our models and ablations takes approximately 13 GPU days on Nvidia 1080Ti GPUs.

Interpolations For our interpolations experiments, we follow the set up of [8] who train a new model conditioned on the polarizability of molecules in the dataset. We train a conditional version of our model which can be achieved by simply adding in the polarizability as an additional feature input to our backbone network and re-using all the same hyperparameters. We show more examples of interpolations in Figure 10.



Figure 7: Unconditional samples from our model.

D.1.5 Ablations

For our main model, we set $\vec{\lambda}_{t < 0.1T} = 0$ to ensure that all dimensions are added with enough generation time remaining for the diffusion process to finalize all state values. To verify this setting, we compare its performance with $\vec{\lambda}_{t < 0.03T} = 0$ and $\vec{\lambda}_{t < 0.3T} = 0$. We show our results in Table 6. We find that the $\vec{\lambda}_{t < 0.03T} = 0$ setting to generate reasonable sample quality but incur some extra dimension error due to the generative process sometimes observing a lack of dimensions near $t = 0$ and adding too many dimensions. We observed the same effect in the paper for when setting $\vec{\lambda}_t$ to be constant for all t in Table 3. Further, the setting $\vec{\lambda}_{t < 0.3T} = 0$ also results in increased dimension error due to there being less opportunity for the guidance model to supervise the number of dimensions. We find that $\vec{\lambda}_{t < 0.1T} = 0$ to be a reasonable trade-off between these effects.

D.1.6 Uniqueness and Novelty Metrics

We here investigate sample diversity and novelty of our unconditional generative models. We measure uniqueness by computing the chemical graph corresponding to each generated sample and measure what proportion of the 10000 produced samples have a unique chemical graph amongst this set of 10000 as is done in [8]. We show our results in Table 7 and find our TDDM method to have slightly lower levels of uniqueness when compared to the fixed dimension diffusion model baseline.

Table 5: The 10 conditioning tasks used for evaluation. The number of each atom type required for the task is given in columns 2 – 5 whilst the average number of atoms in molecules that meet this condition in the training dataset is given in the 6th column.

Task	Carbon	Nitrogen	Oxygen	Fluorine	Mean Number of Atoms
1	4	1	2	1	11.9
2	4	3	1	1	13.0
3	5	2	1	1	13.9
4	6	0	1	1	14.6
5	5	3	1	0	16.0
6	6	3	0	0	17.2
7	6	1	2	0	17.7
8	7	1	1	0	19.1
9	8	1	0	0	19.9
10	8	0	1	0	21.0

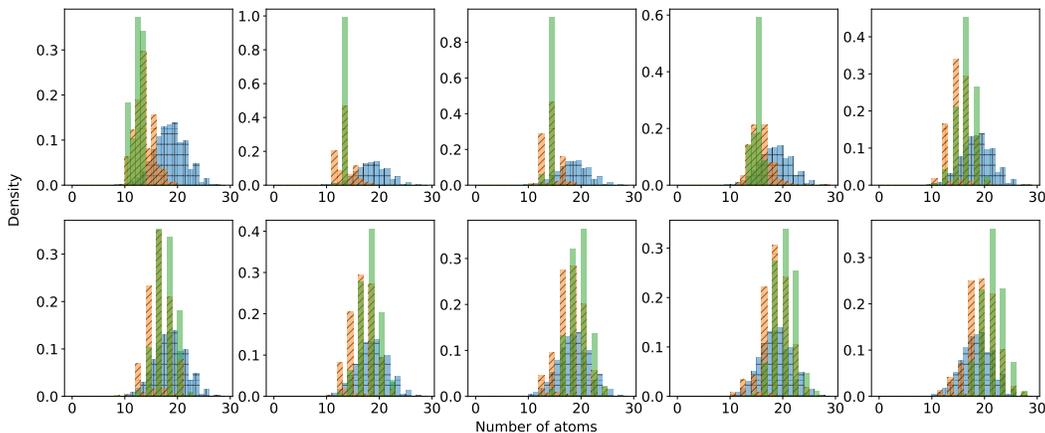


Figure 8: Distribution of molecule sizes for each conditioning task. Tasks 1 – 5 are shown left to right in the top row and tasks 6 – 10 are shown left to right in the bottom row. We show the unconditional size distribution from the dataset in blue vertical/horizontal hashing, the size distribution of our conditionally generated samples in orange diagonal hashing and finally the size distribution for molecules in the training dataset that match the conditions of each task (the ground truth size distribution) in green no hashing.

Measuring novelty on generative models trained on the QM9 dataset is challenging because the QM9 dataset contains an exhaustive enumeration of all molecules that satisfy certain predefined constraints [46], [8]. Therefore, if a novel molecule is produced it means the generative model has failed to capture some of the physical properties of the dataset and indeed it is found in [8] that during training, as the model improved, novelty decreased. Novelty is therefore not typically included in evaluating molecular diffusion models. For completeness, we include the novelty scores in Table 7 as a comparison to the results presented in [8] Appendix C. We find that our samples are closer to the statistics of the training dataset whilst still producing ‘novel’ samples at a consistent rate.

D.2 Video

D.2.1 Dataset

We used the VP² benchmark, which consists of 35 000 videos, each 35 frames long. The videos are evenly divided among seven tasks, namely: push {red, green, blue} button, open {slide, drawer}, push {upright block, flat block} off table. The 5000 videos for each task were collected using a scripted task-specific policy operating in the RoboDesk environment [40]. They sample an action vector at every step during data generation by adding i.i.d. Gaussian

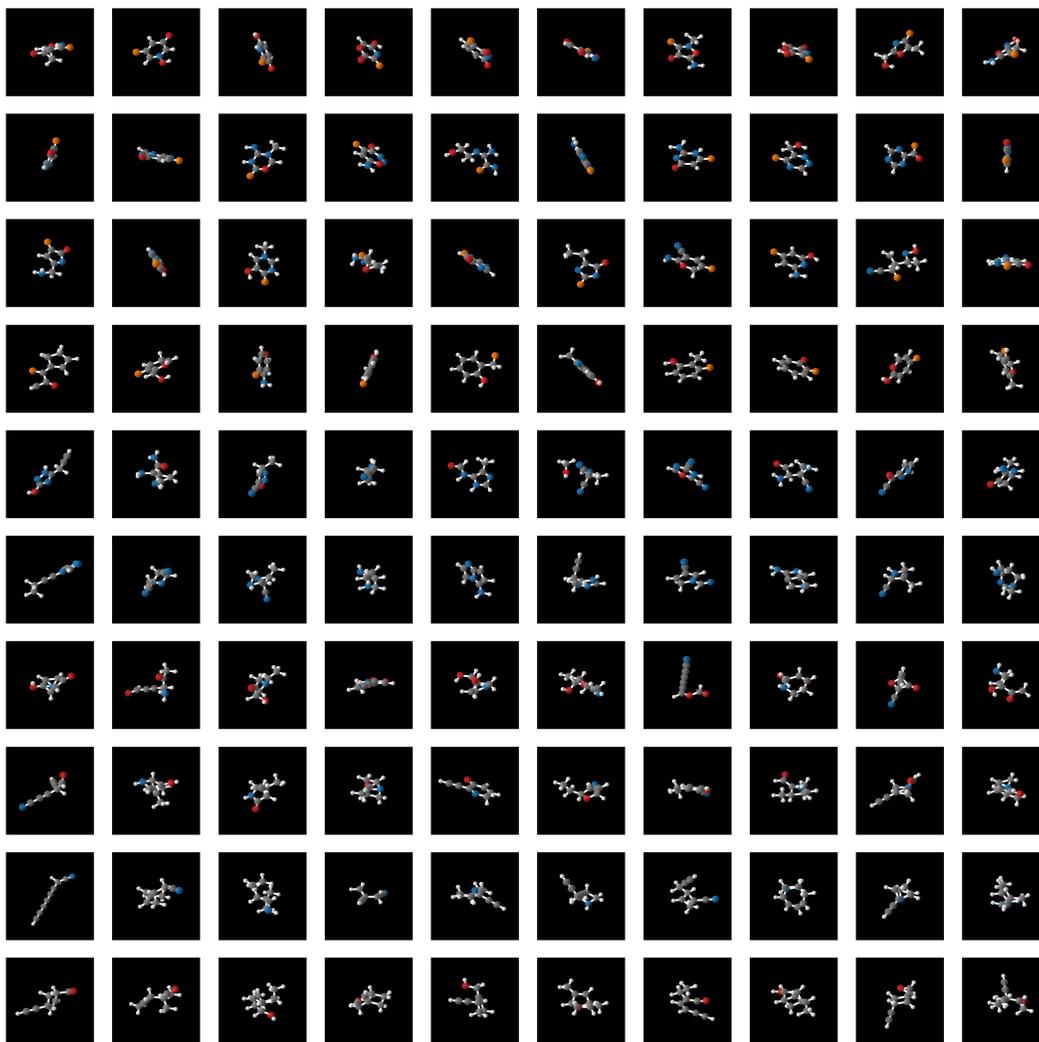


Figure 9: Samples generated by our model when conditional diffusion guidance is applied. Each row represents one task with task 1 at the top, down to task 10 at the bottom. For each task, 10 samples are shown in each row.

noise to each dimension of the action vector output by the scripted policy. For each task, they sample 2500 videos with noise standard deviation 0.1 and 2500 videos with standard deviation 0.2. We filter out the lower-quality trajectories sampled with noise standard deviation 0.2, and so use only the 17 500 videos (2500) per task with noise standard deviation 0.1. We convert these videos to 32×32 resolution and then, so that the data we train on has varying lengths, we create each training example by sampling a length l from a uniform distribution over $\{2, \dots, 35\}$ and then taking a random l -frame subset of the video.

D.2.2 Forward Process

The video domain differs from molecules in two important ways. The first is that videos cannot be reasonably treated as a permutation-invariant set. This is because the order of the frames matters. Secondly, generating a full new component for the molecules with a single pass autoregressive network is feasible, however, a component for the videos is a full frame which is challenging for a single pass autoregressive network to generate. We design our forward process to overcome these challenges.



Figure 10: Interpolations showing a sequence of generations for linearly increasing polarizability from 39 Bohr³ to 66 Bohr³ with fixed random noise. Each row shows an individual interpolation with Bohr³ increasing from left to right.

Table 6: Ablation of when to set the forward rate to 0 on the conditional molecule generation task. We report dimension error as the average Hellinger distance between the generated and ground truth conditional dimension distributions as well as average sample quality metrics. Metrics are reported after 620k training iterations.

Method	Dimension Error	% Atom stable	% Molecule Stable	% Valid
$\vec{\lambda}_{t < 0.03T} = 0$	0.227 \pm 0.16	91.5 \pm 3.7	56.5 \pm 9.8	72.0 \pm 11
$\vec{\lambda}_{t < 0.1T} = 0$	0.162 \pm 0.071	92.4 \pm 2.8	53.9 \pm 12	72.7 \pm 9.6
$\vec{\lambda}_{t < 0.3T} = 0$	0.266 \pm 0.11	92.0 \pm 3.2	53.5 \pm 13	66.6 \pm 12

Table 7: Uniqueness and novelty metrics on unconditional molecule generation. We produce 10000 samples for each method and measure validity using RDKit. Uniqueness is judged as whether the chemical graph is unique amongst the 10000 produced samples. Amongst the valid and unique molecules, we then find the percentage that have a chemical graph not present in the training dataset.

Method	Percentage of Valid and Unique		
	% Valid	% Valid and Unique	Molecules that are Novel
FDDM [8]	91.9	90.7	65.7
TDDM (ours) $\vec{\lambda}_t$	92.3	89.9	53.6
TDDM, const $\vec{\lambda}_t$	86.7	84.4	56.9
TDDM, $\vec{\lambda}_{t < 0.9T} = 0$	89.4	86.1	51.3
TDDM w/o Prop. 3	87.1	85.9	63.3

We define our forward process to delete frames in a random order. This means that during generation, frames can be generated in any order in the reverse process, enabling more conditioning tasks since we can always ensure that whichever frames we want to condition on are added first. Further, we use a non-isotropic noise schedule by adding noise just to the frame that is about to be deleted. Once it is deleted, we then start noising the next randomly chosen frame. This is so that, in the backward direction, when a new frame is added, it is simply Gaussian noise. Then the score network will fully denoise that new frame before the next new frame is added. We now specify exactly how our forward process is constructed.

We enable random-order deletion by applying an initial shuffling operation occurring at time $t = 0$. Before this operation, we represent the video \mathbf{x} as an ordered sequence of frames, $\mathbf{x}_0 = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n_0}]$. During shuffling, we sample a random permutation π of the integers $1, \dots, n_0$. Then the frames are kept in the same order, but annotated with an index variable so that we have $\mathbf{x}_{0+} = [(\mathbf{x}_{0+}^{(1)}, \pi(1)), (\mathbf{x}_{0+}^{(2)}, \pi(2)), \dots, (\mathbf{x}_{0+}^{(n_0)}, \pi(n_0))]$.

We will run the forward process from $t = 0$ to $t = 100N$. We will set the forward rate such we delete down from n_t to $n_t - 1$ at time $(N - n_t + 1)100$. This is achieved heuristically by setting

$$\vec{\lambda}_t(n_t) = \begin{cases} 0 & \text{for } t < (N - n_t + 1)100, \\ \infty & \text{for } t \geq (N - n_t + 1)100. \end{cases}$$

We can see that at time $t = (N - n_t + 1)100$ we will quickly delete down from n_t to $n_t - 1$ at which point $\vec{\lambda}_t(n_t)$ will become 0 thus stopping deletion until the process arrives at the next multiple of 100 in time. When we hit a deletion event, we delete the frame from \mathbf{X}_t that has the current highest index variable $\pi(n)$. In other words

$$K^{\text{del}}(i|\mathbf{X}_t) = \begin{cases} 1 & \text{for } n_t = \mathbf{x}_t^{(i)}[2], \\ 0 & \text{otherwise} \end{cases}$$

where we use $\mathbf{x}_t^{(i)}[2]$ to refer to the shuffle index variable for the i th current frame in \mathbf{x}_t .

We now provide an example progression of the forward deletion process. Assume we have $n_0 = 4$, $N = 5$ and sample a permutation such that $\pi(1) = 3, \pi(2) = 2, \pi(3) = 4$, and $\pi(4) = 1$. Initially

the state is augmented to include the shuffle index. Then the forward process progresses from $t = 0$ to $t = 500$ with components being deleted in descending order of the shuffle index

$$\begin{aligned}\mathbf{x}_{0+} &= [(\mathbf{x}_t^{(1)}, 3), (\mathbf{x}_t^{(2)}, 2), (\mathbf{x}_t^{(3)}, 4), (\mathbf{x}_t^{(4)}, 1)] \\ \mathbf{x}_{100+} &= [(\mathbf{x}_t^{(1)}, 3), (\mathbf{x}_t^{(2)}, 2), (\mathbf{x}_t^{(3)}, 4), (\mathbf{x}_t^{(4)}, 1)] \\ \mathbf{x}_{200+} &= [(\mathbf{x}_t^{(1)}, 3), (\mathbf{x}_t^{(2)}, 2), (\mathbf{x}_t^{(4)}, 1)] \\ \mathbf{x}_{300+} &= [(\mathbf{x}_t^{(2)}, 2), (\mathbf{x}_t^{(4)}, 1)] \\ \mathbf{x}_{400+} &= [(\mathbf{x}_t^{(4)}, 1)]\end{aligned}$$

In this example, due to the random permutation sampled, the final video frame remained after all others had been deleted. Note that the order of frames is preserved as we delete frames in the forward process although the spacing between them can change as we delete frames in the middle.

Between jumps, we use a noising process to add noise to frames. The noising process is non-isotropic in that it adds noise to different frames at different rates such that the a frame is noised only in the time window immediately preceding its deletion. For component $i \in [1, \dots, n_t]$, we set the forward noising process such that $p_{t|0}(\mathbf{x}_t^{(i)} | \mathbf{x}_0^{(i)}, M_t) = \mathcal{N}(\mathbf{x}_t^{(i)}; \mathbf{x}_0^{(i)}, \sigma_t(\mathbf{x}_t^{(i)})^2)$ where $\mathbf{x}_0^{(i)}$ is the clean frame corresponding to $\mathbf{x}_t^{(i)}$ as given by the mask M_t and $\sigma_t(\mathbf{x}_t^{(i)})$ follows

$$\sigma_t(\mathbf{x}_t^{(i)}) = \begin{cases} 0 & \text{for } t < (N - \mathbf{x}_t^{(i)}[2])100, \\ 100 & \text{for } t > (N - \mathbf{x}_t^{(i)}[2])100, \\ t - (N - \mathbf{x}_t^{(i)}[2])100 & \text{for } (N - \mathbf{x}_t^{(i)}[2])100 \leq t \leq (N - \mathbf{x}_t^{(i)}[2] + 1)100 \end{cases}$$

where we again use $\mathbf{x}_t^{(i)}[2]$ for the shuffle index of component i . This is the VE-SDE from [3] applied to each frame in turn. We note that we only add noise to the state values on not the shuffle index itself.

The SDE parameters that result in the VE-SDE are $\vec{\mathbf{b}}_t = 0$ and $\vec{\mathbf{g}}_t = \sqrt{2t - 2(N - \mathbf{x}_t^{(i)}[2])100}$.

D.2.3 Sampling the Backward Process

When t is not at a multiple of 100, the forward process is purely adding Gaussian noise, and so the reverse process is also purely operating on the continuous dimensions. We use the Heun sampler proposed by [16] to update the continuous dimensions in this case, and also a variation of their discretisation of t - specifically to update from e.g. $t = 600$ to $t = 500$, we use their discretization of t as if the maximum value was 100 and then offset all values by 500.

To invert the dimension deletion process, we can use Proposition 3 to derive our reverse dimension generation process. We re-write our parameterized $\overleftarrow{\lambda}_t^\theta$ using Proposition 3 as

$$\overleftarrow{\lambda}_t^\theta(\mathbf{X}_t) = \overrightarrow{\lambda}_t(n_t + 1) \mathbb{E}_{p_{0|t}^\theta(n_0 | \mathbf{X}_t)} \left[\frac{p_{t|0}(n_t + 1 | n_0)}{p_{t|0}(n_t | n_0)} \right]$$

At each time multiple of 100 in the backward process, we will have an opportunity to add a component. At this time point, we estimate the expectation with a single sample $n_0 \sim p_{0|t}^\theta(n_0 | \mathbf{X}_t)$. If $n_0 > n_t$ then $\overleftarrow{\lambda}_t^\theta(\mathbf{X}_t) = \infty$. The new component will then be added at which point $\overleftarrow{\lambda}_t^\theta(\mathbf{X}_t)$ becomes 0 for the remainder of this block of time due to n_t becoming $n_t + 1$. If $n_0 = n_t$ then $\overleftarrow{\lambda}_t^\theta(\mathbf{X}_t) = 0$ and no new component is added. $\overleftarrow{\lambda}_t^\theta(\mathbf{X}_t)$ will continue to be 0 for the remainder of the backward process once an opportunity to add a component is not used.

When a new frame is added, we use $A_t^\theta(\mathbf{y}^{\text{add}}, i | \mathbf{X}_t)$ to decide where the frame is added and its initial value. Since when we delete a frame it is fully noised, $A_t^\theta(\mathbf{y}^{\text{add}}, i | \mathbf{X}_t)$ can simply predict Gaussian noise for the new frame \mathbf{y}^{add} . However, $A_t^\theta(\mathbf{y}^{\text{add}}, i | \mathbf{X}_t)$ will still learn to predict a suitable location i to place the new frame such that backward process is the reversal of the forward.

We give an example simulation from the backward generative process in Figure 11.

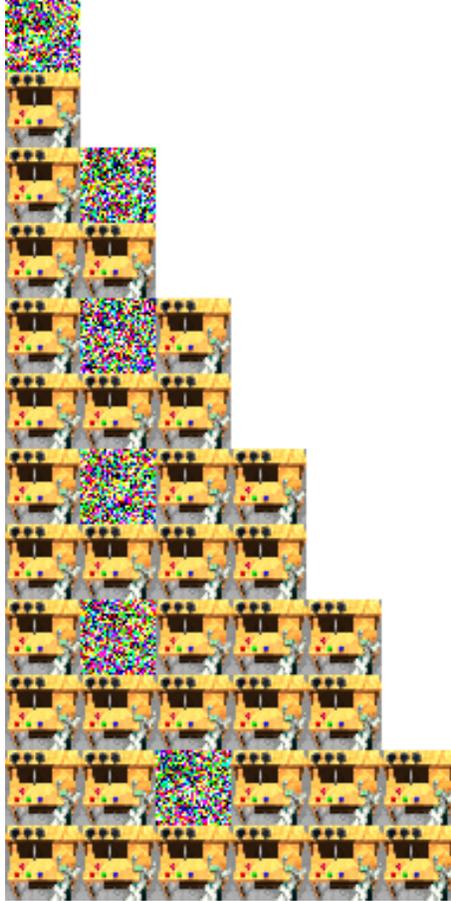


Figure 11: An example simulation of the backward generative process conditioned on the first and last frame. Note how the process first adds a new frame and then fully denoises it before adding the next frame. Since the first and last frame are very similar, the process produces a short video.

D.2.4 Network Architecture

Our video diffusion network architecture is based on the U-net used by [42], which takes as input the index of each frame within the video, and uses the differences between these indices to control the interactions between frames via an attention mechanism. Since, during generation, we do not know the final position of each frame within the \mathbf{x}_0 , we instead pass in its position within the ordered sequence \mathbf{x}_t .

One further difference is that, since we are perform non-isotropic diffusion, the standard deviation of the added noise will differ between frames. We adapt to this by performing preconditioning, and inputting the timestep embedding, separately for each frame $\mathbf{x}_t^{(i)}$ based on $\sigma_t(\mathbf{x}_t^{(i)})$ instead of basing them on the global diffusion timestep t . Our timestep embedding and pre- and post-conditioning of network inputs/outputs are as suggested by [16], other than being done on a per-frame basis. The architecture from [42] with these changes applied then gives us our score network s_t^θ .

While it would be possible to train a single network that estimates the score and all quantities needed for modelling jumps, we chose to train two separate networks in order to factorize our exploration of the design space. These were the score network s_t^θ , and the rate and index prediction network modeling $p_{0|t}^\theta(n_0|\mathbf{X}_t)$ and $A_t^\theta(i|\mathbf{X}_t)$. The rate and index prediction network is similar to the first half of the score network, in that it uses all U-net blocks up to and including the middle one. We then flatten the $512 \times 4 \times 4$ hidden state for each frame after this block such that, for an n_t frame input, we obtain a $n_t \times 8192$ hidden state. These are fed through a 1D convolution with kernel size 2 and zero-padding of size 1 on each end, reducing the hidden state to $(n_t + 1) \times 128$, which is in turn fed

through a ReLU activation function. This hidden state is then fed into three separate heads. One head maps it to the parameters of $A_t^\theta(i|\mathbf{X}_t)$ via a 1D convolution of kernel size 3. The output of size $(n_t + 1)$ is fed through a softmax to provide the categorical distribution $A_t^\theta(i|\mathbf{X}_t)$. The second head averages the hidden state over the “frame” dimension, producing a 128-dimensional vector. This is fed through a single linear layer and a softmax to parameterize $p_{0|t}^\theta(n_0|\mathbf{X}_t)$. Finally, the third head consists of a 1D convolution of kernel size 3 with 35 output channels. The $(n_t + 1) \times 35$ output is fed through a softmax to parameterize distributions over the number of frames that were deleted from \mathbf{X}_0 which came before the first in \mathbf{x}_t , the number of frames from \mathbf{X}_0 which were deleted between each pair of frames in \mathbf{x}_t , and the number deleted after the last frame in \mathbf{x}_t . We do not use this head at inference-time but found that including it improved the performance of the other heads by helping the network learn better representations.

For a final performance improvement, we note that under our forward process there is only ever one “noised” frame in \mathbf{x}_t , while there are sometimes many clean frames. Since the cost of running our architecture scales with the number of frames, running it on many clean frames may significantly increase the cost while providing little improvement to performance. We therefore only feed into the architecture the “noised” frame, the two closest “clean” frames before it, and the two closest “clean” frames after it. See our released source code for the full implementation of this architecture.

D.2.5 Training

To sample t during training, we adapt the log-normal distribution suggested by [16] in the context of isotropic diffusion over a single image. To apply it to our non-isotropic video diffusion, we first sample which frames have been deleted, which exist with no noise, and which have had noise added, by sampling the timestep from a uniform distribution and simulating our proposed forward process. We then simply change the noise standard deviation for the noisy frame, replacing it with a sample from the log-normal distribution. The normal distribution underlying our log-normal has mean -0.6 and standard deviation 1.8 . This can be interpreted as sampling the timestep from a mixture of log-normal distributions, $\frac{1}{N} \sum_{i=0}^{N-1} \mathcal{LN}(t - 100i; -0.6, 1.8^2)$. Here, the mixture index i can be interpreted as controlling the number of deleted frames.

We use the same loss weighting as [16] but, similarly to our use of preconditioning, compute the weighting separately for each frame $\mathbf{x}_t^{(i)}$ as a function of $\sigma_t(\mathbf{x}_t^{(i)})$ to account for the non-isotropic noise.

D.2.6 Perceptual Quality Metrics

We now verify that our reverse process does not have any degradation in quality during the generation as more dimensions are added. We generate 10000 videos and throw away the 278 that were sampled to have only two frames. We then compute the FID score for individual frames in each of the remaining 9722 videos. We group together the scores for all the first frames to be generated in the reverse process and then for the second frame to be generated and so on. We show our results in Table 8. We find that when a frame is inserted has no apparent effect on perceptual quality and conclude that there is no overall degradation in quality as our sampling process progresses. We note that the absolute value of these FID scores may not be meaningful due to the RoboDesk dataset being far out of distribution for the Inception network used to calculate FID scores. We can visually confirm good sample quality from Figure 5.

Table 8: FID for video frames grouped by when they were inserted during sampling.

1st	2nd	3rd	3rd last	2nd last	last
34.2	34.9	34.7	34.2	34.1	34.4

E Broader Impacts

In this work, we presented a general method for performing generative modeling on datasets of varying dimensionality. We have not focused on applications and instead present a generic method.

Along with other generic methods for generative modeling, we must consider the potential negative social impacts that these models can cause when inappropriately used. As generative modeling capabilities increase, it becomes simpler to generate fake content which can be used to spread misinformation. In addition to this, generative models are becoming embedded into larger systems that then have real effects on society. There will be biases present within the generations created by the model which in turn can reinforce these biases when the model's outputs are used within wider systems. In order to mitigate these harms, applications of generative models to real world problems must be accompanied with studies into their biases and potential ways they can be misused. Further, public releases of models must be accompanied with model cards [47] explaining the biases, limitations and intended uses of the model.