## A Details of the DDRM ELBO objective

DDRM is a Markov chain conditioned on $\mathbf{y}$, which would lead to the following ELBO objective [45]:

$$
\begin{align*}
& \mathbb{E}_{\mathbf{x}_{0} \sim q\left(\mathbf{x}_{0}\right), \mathbf{y} \sim q\left(\mathbf{y} \mid \mathbf{x}_{0}\right)}\left[\log p_{\theta}\left(\mathbf{x}_{0} \mid \mathbf{y}\right)\right]  \tag{9}\\
\geq & -\mathbb{E}\left[\sum_{t=1}^{T-1} D_{\mathrm{KL}}\left(q^{(t)}\left(\mathbf{x}_{t} \mid \mathbf{x}_{t+1}, \mathbf{x}_{0}, \mathbf{y}\right) \| p_{\theta}^{(t)}\left(\mathbf{x}_{t} \mid \mathbf{x}_{t+1}, \mathbf{y}\right)\right)\right]+\mathbb{E}\left[\log p_{\theta}^{(0)}\left(\mathbf{x}_{0} \mid \mathbf{x}_{1}, \mathbf{y}\right)\right] \\
& -\mathbb{E}\left[D_{\mathrm{KL}}\left(q^{(T)}\left(\mathbf{x}_{T} \mid \mathbf{x}_{0}, \mathbf{y}\right) \| p_{\theta}^{(T)}\left(\mathbf{x}_{T} \mid \mathbf{y}\right)\right)\right] \tag{10}
\end{align*}
$$

where $q\left(\mathbf{x}_{0}\right)$ is the data distribution, $q\left(\mathbf{y} \mid \mathbf{x}_{0}\right)$ follows Equation 1 in the main paper, the expectation on the right hand side is given by sampling $\mathbf{x}_{0} \sim q\left(\mathbf{x}_{0}\right), \mathbf{y} \sim q\left(\mathbf{y} \mid \mathbf{x}_{0}\right), \mathbf{x}_{T} \sim q^{(T)}\left(\mathbf{x}_{T} \mid \mathbf{x}_{0}, \mathbf{y}\right)$, and $\mathbf{x}_{t} \sim q^{(t)}\left(\mathbf{x}_{t} \mid \mathbf{x}_{t+1}, \mathbf{x}_{0}, \mathbf{y}\right)$ for $t \in[1, T-1]$.

## B Equivalence between "Variance Preserving" and "Variance Exploding" Diffusion Models

In our main paper, we describe our methods based on the "Variance Exploding" hyperparameters $\sigma_{t}$, where $\sigma_{t} \in[0, \infty)$ and

$$
\begin{equation*}
q\left(\mathbf{x}_{t} \mid \mathbf{x}_{0}\right)=\mathcal{N}\left(\mathbf{x}_{0}, \sigma_{t}^{2} \boldsymbol{I}\right) \tag{11}
\end{equation*}
$$

In DDIM [45], the hyperparameters are "Variance Preserving" ones $\alpha_{t}$, where $\alpha_{t} \in(0,1]$ and

$$
\begin{equation*}
q\left(\mathbf{x}_{t} \mid \mathbf{x}_{0}\right)=\mathcal{N}\left(\sqrt{\alpha_{t}} \mathbf{x}_{0},\left(1-\alpha_{t}\right) \boldsymbol{I}\right) \tag{12}
\end{equation*}
$$

We use the colored notation $\mathbf{x}_{t}$ to emphasize that this is different from $\mathbf{x}_{t}$ (an exception is $\mathbf{x}_{0}=\mathbf{x}_{0}$ ). Using the reparametrization trick, we have that:

$$
\begin{gather*}
\mathbf{x}_{t}=\mathbf{x}_{0}+\sigma_{t} \epsilon  \tag{13}\\
\mathbf{x}_{t}=\sqrt{\alpha_{t}} \mathbf{x}_{0}+\sqrt{1-\alpha_{t}} \epsilon \tag{14}
\end{gather*}
$$

where $\epsilon \sim \mathcal{N}(0, \boldsymbol{I})$. We can divide by $\sqrt{1+\sigma_{t}^{2}}$ in both sides of Equation 13:

$$
\begin{equation*}
\frac{\mathbf{x}_{t}}{\sqrt{1+\sigma_{t}^{2}}}=\frac{\mathbf{x}_{0}}{\sqrt{1+\sigma_{t}^{2}}}+\frac{\sigma_{t}}{\sqrt{1+\sigma_{t}^{2}}} \epsilon \tag{15}
\end{equation*}
$$

Let $\alpha_{t}=1 /\left(1+\sigma_{t}^{2}\right)$, and let $\mathbf{x}_{t}=\mathbf{x}_{t} / \sqrt{1+\sigma_{t}^{2}}$; then from Equation 15 we have that

$$
\begin{equation*}
\mathbf{x}_{t}=\sqrt{\alpha_{t}} \mathbf{x}_{0}+\sqrt{1-\alpha_{t}} \epsilon \tag{16}
\end{equation*}
$$

which is equivalent to the "Variance Preserving" case. Therefore, we can use "Variance Preserving" models, such as DDPM, directly in our DDRM updates, even though the latter uses the "Variance Exploding" parametrization:

1. From $\mathrm{x}_{t}$, obtain predictions $\epsilon$ and $\mathrm{x}_{t}=\mathrm{x}_{t} \sqrt{1+\sigma_{t}^{2}}$.
2. From $\mathbf{x}_{t}$ and $\epsilon$, apply DDRM updates to get $\mathbf{x}_{t-1}$.
3. From $\mathbf{x}_{t-1}$, get $\mathbf{x}_{t-1}=\mathbf{x}_{t-1} / \sqrt{1+\sigma_{t-1}^{2}}$.

Note that although the inference algorithms are shown to be equivalent, the choice between "Variance Preserving" and "Variance Exploding" may affect the training of diffusion networks.

## C Proofs

Proposition 3.1. The conditional distributions $q^{(t)}$ defined in Equations 4 and 5 satisfy the following:

$$
\begin{equation*}
q\left(\mathbf{x}_{t} \mid \mathbf{x}_{0}\right)=\mathcal{N}\left(\mathbf{x}_{0}, \sigma_{t}^{2} \boldsymbol{I}\right) \tag{6}
\end{equation*}
$$

defined by marginalizing over $\mathbf{x}_{t^{\prime}}\left(\right.$ for all $\left.t^{\prime}>t\right)$ and $\mathbf{y}$, where $q\left(\mathbf{y} \mid \mathbf{x}_{0}\right)$ is defined as in Equation (1) with $\mathbf{x}=\mathbf{x}_{0}$.

Proof. The proof uses a basic property of Gaussian marginals (see [4] for the complete version).

1. If $p\left(z_{1} \mid z_{0}\right)=\mathcal{N}\left(z_{0}, V_{1}\right), p\left(z_{2} \mid z_{1}\right)=\mathcal{N}\left(\alpha z_{1}, V_{2}\right)$, then $p\left(z_{2} \mid z_{0}\right)=\mathcal{N}\left(\alpha z_{0}, \alpha^{2} V_{1}+V_{2}\right)$.
2. If $p\left(z_{1}\right)=\mathcal{N}\left(\mu_{1}, V_{1}\right)$ and $p\left(z_{2}\right)=\mathcal{N}\left(\mu_{2}, V_{2}\right)$, then $p\left(z_{1}+z_{2}\right)=\mathcal{N}\left(\mu_{1}+\mu_{2}, V_{1}+V_{2}\right)$.

First, we note that $q\left(\mathbf{y} \mid \mathbf{x}_{0}\right)$ is defined from Equation 1 in the main paper, and thus for all $i$ :

$$
\begin{equation*}
q\left(\overline{\mathbf{y}}^{(i)} \mid \mathbf{x}_{0}\right)=\mathcal{N}\left(\overline{\mathbf{x}}_{0}^{(i)}, \sigma_{\mathbf{y}}^{2} / s_{i}^{2}\right) \tag{17}
\end{equation*}
$$

Case I For $\mathbf{x}_{T}$, it is obvious when $s_{i}=0$. When $s_{i}>0$, we have Equation 17 and that:

$$
\begin{equation*}
q^{(T)}\left(\overline{\mathbf{x}}_{T}^{(i)} \mid \mathbf{x}_{0}, \mathbf{y}\right)=\mathcal{N}\left(\overline{\mathbf{y}}^{(i)}, \sigma_{T}^{2}-\frac{\sigma_{\mathbf{y}}^{2}}{s_{i}^{2}}\right) \tag{18}
\end{equation*}
$$

and thus

$$
q^{(T)}\left(\overline{\mathbf{x}}_{T}^{(i)} \mid \mathbf{x}_{0}\right)=\mathcal{N}\left(\overline{\mathbf{x}}_{0}^{(i)}, \sigma_{\mathbf{y}}^{2} / s_{i}^{2}+\sigma_{T}^{2}-\frac{\sigma_{\mathbf{y}}^{2}}{s_{i}^{2}}\right)=\mathcal{N}\left(\overline{\mathbf{x}}_{0}^{(i)}, \sigma_{T}^{2}\right)
$$

Case II For any $t<T$ and $i$ such that $s_{i}>0$ and $\sigma_{t}>\sigma_{\mathbf{y}} / s_{i}$, we have Equation 17 and that:

$$
\begin{equation*}
q^{(t)}\left(\overline{\mathbf{x}}_{t}^{(i)} \mid \mathbf{x}_{t+1}, \mathbf{x}_{0}, \mathbf{y}\right)=\mathcal{N}\left(\left(1-\eta_{b}\right) \overline{\mathbf{x}}_{0}^{(i)}+\eta_{b} \overline{\mathbf{y}}^{(i)}, \sigma_{t}^{2}-\frac{\sigma_{\mathbf{y}}^{2}}{s_{i}^{2}} \eta_{b}^{2}\right) \tag{19}
\end{equation*}
$$

and thus we can safely remove the dependence on $\mathbf{x}_{t+1}$ via marginalization. $q^{(t)}\left(\overline{\mathbf{x}}_{t}^{(i)} \mid \mathbf{x}_{0}\right)$ is a Gaussian with the mean being $\left(1-\eta_{b}\right) \overline{\mathbf{x}}_{0}^{(i)}+\eta_{b} \overline{\mathbf{x}}_{0}^{(i)}=\overline{\mathbf{x}}_{0}^{(i)}$ and variance being

$$
\sigma_{t}^{2}-\frac{\sigma_{\mathbf{y}}^{2}}{s_{i}^{2}} \eta_{b}^{2}+\frac{\sigma_{\mathbf{y}}^{2}}{s_{i}^{2}} \eta_{b}^{2}=\sigma_{t}^{2}
$$

where we note that $\overline{\mathbf{y}}^{(i)}$ has a standard deviation of $\sigma_{\mathbf{y}} / s_{i}$.
Case III For any $t<T$ and $i$ such that $s_{i}>0$ and $\sigma_{t}<\sigma_{\mathbf{y}} / s_{i}$, we have Equation 17, so $\left(\overline{\mathbf{y}}^{(i)}-\overline{\mathbf{x}}_{0}^{(i)}\right) /\left(\sigma_{\mathbf{y}} / s_{i}\right)$ is distributed as a standard Gaussian. Moreover, similar to Case II, $q^{(t)}\left(\overline{\mathbf{x}}_{t}^{(i)} \mid \mathbf{x}_{0}\right)$ is a Gaussian with its mean being

$$
\overline{\mathbf{x}}_{0}^{(i)}+\sqrt{1-\eta^{2}} \sigma_{t} \frac{\overline{\mathbf{y}}^{(i)}-\overline{\mathbf{x}}_{0}^{(i)}}{\sigma_{\mathbf{y}} / s_{i}}
$$

and its variance being $\eta^{2} \sigma_{t}^{2}$, so $q^{(t)}\left(\overline{\mathbf{x}}_{t}^{(i)} \mid \mathbf{x}_{0}\right)$ is a Gaussian with a mean of $\overline{\mathbf{x}}_{0}^{(i)}$ and a variance of

$$
\left(1-\eta^{2}\right) \sigma_{t}^{2}+\eta^{2} \sigma_{t}^{2}=\sigma_{t}^{2}
$$

Case IV For any $t \leq T$ and $i$ such that $s_{i}=0$ (where there is no dependence on y), we apply mathematical induction. The base case $(t=T)$ is true, as we have shown earlier in Case I. In the step case $(t<T)$, we have that $q^{(t+1)}\left(\overline{\mathbf{x}}_{t+1}^{(i)} \mid \mathbf{x}_{0}\right)=\mathcal{N}\left(\overline{\mathbf{x}}_{0}^{(i)}, \sigma_{t+1}^{2}\right)$. Similar to Case II, $q^{(t)}\left(\overline{\mathbf{x}}_{t}^{(i)} \mid \mathbf{x}_{0}\right)$ is a Gaussian with its mean being

$$
\overline{\mathbf{x}}_{0}^{(i)}+\sqrt{1-\eta^{2}} \sigma_{t} \frac{\overline{\mathbf{x}}_{t+1}^{(i)}-\overline{\mathbf{x}}_{0}^{(i)}}{\sigma_{t+1}}
$$

and variance being $\eta^{2} \sigma_{t}^{2}$, which does not depend on $\mathbf{y}$. Therefore, $q^{(t)}\left(\overline{\mathbf{x}}_{t}^{(i)} \mid \mathbf{x}_{0}\right)$ is also Gaussian, with a mean of $\overline{\mathbf{x}}_{0}^{(i)}$ and a variance of

$$
\left(1-\eta^{2}\right) \sigma_{t}^{2}+\eta^{2} \sigma_{t}^{2}=\sigma_{t}^{2}
$$

Hence, the proof is completed via the four cases.
Theorem 3.2. Assume that the models $f_{\theta}^{(t)}$ and $f_{\theta}^{\left(t^{\prime}\right)}$ do not have weight sharing whenever $t \neq t^{\prime}$, then when $\eta=1$ and $\eta_{b}=\frac{2 \sigma_{t}^{2}}{\sigma_{t}^{2}+\sigma_{y}^{2} / s_{i}^{2}}$, the ELBO objective of DDRM (details in Appendix A) can be rewritten in the form of the DDPM / DDIM objective in Equation (2).

Proof. As there is no parameter sharing between models at different time steps $t$, let us focus on any particular time step $t$ and rewrite the corresponding objective as a denoising autoencoder objective.
Case I For $t>0$, the only term in Equation 10 that is related to $f_{\theta}^{(t)}$ (which is used to make the prediction $\mathbf{x}_{\theta, t}$ ) is:

$$
\begin{align*}
& D_{\mathrm{KL}}\left(q^{(t)}\left(\mathbf{x}_{t} \mid \mathbf{x}_{t+1}, \mathbf{x}_{0}, \mathbf{y}\right) \| p_{\theta}^{(t)}\left(\mathbf{x}_{t} \mid \mathbf{x}_{t+1}, \mathbf{y}\right)\right) \\
= & D_{\mathrm{KL}}\left(q^{(t)}\left(\overline{\mathbf{x}}_{t} \mid \mathbf{x}_{t+1}, \mathbf{x}_{0}, \mathbf{y}\right) \| p_{\theta}^{(t)}\left(\overline{\mathbf{x}}_{t} \mid \mathbf{x}_{t+1}, \mathbf{y}\right)\right) \\
= & \sum_{i=1}^{n} D_{\mathrm{KL}}\left(q^{(t)}\left(\overline{\mathbf{x}}_{t}^{(i)} \mid \mathbf{x}_{t+1}, \mathbf{x}_{0}, \mathbf{y}\right) \| p_{\theta}^{(t)}\left(\overline{\mathbf{x}}_{t}^{(i)} \mid \mathbf{x}_{t+1}, \mathbf{x}_{0}, \mathbf{y}\right)\right), \tag{20}
\end{align*}
$$

where the first equality is from the orthogonality of $\boldsymbol{V}^{\top}$ and the second equality is from the fact that both $q^{(t)}$ and $p_{\theta}^{(t)}$ over the spectral space are Gaussians with identical diagonal covariance matrices (so the KL divergence can factorize).
Here, we will use a simple property of the KL divergence between univariate Gaussians [27]:

$$
\begin{aligned}
& \text { If } p=\mathcal{N}\left(\mu_{1}, V_{1}\right), q=\mathcal{N}\left(\mu_{2}, V_{2}\right) \text {, then } \\
& \qquad D_{\mathrm{KL}}(p \| q)=\frac{1}{2} \log \frac{V_{2}}{V_{1}}+\frac{V_{1}+\left(\mu_{1}-\mu_{2}\right)^{2}}{2 V_{2}}-\frac{1}{2} .
\end{aligned}
$$

Since we constructed $p_{\theta}^{(t)}$ and $q^{(t)}$ to have the same variance, Equation 20 is a total squared error with weights for each dimension of $\overline{\mathbf{x}}_{t}$ (the spectral space), so the DDPM objective (which is a total squared error objective in the original space) is still a good approximation. In order to transform it into a denoising autoencoder objective (equivalent to DDPM), the weights have to be equal. Next, we will show that our construction of $\eta=1$ and $\eta_{b}=2 \sigma_{t}^{2} /\left(\sigma_{t}^{2}+\sigma_{\mathbf{y}}^{2} / s_{i}^{2}\right)$ satisfies this.
All the indices $i$ will fall into one of the three cases: $s_{i}=0, \sigma_{t}<\sigma_{\mathbf{y}} / s_{i}$, or $\sigma_{t}>\sigma_{\mathbf{y}} / s_{i}$.

- For $s_{i}=0$, the KL divergence is $\frac{\left(\overline{\mathbf{x}}_{\theta, t}^{(i)}-\overline{\mathbf{x}}_{0}^{(i)}\right)^{2}}{2 \sigma_{t}^{2}}$, where we recall $\overline{\mathbf{x}}_{\theta, t}=\boldsymbol{V}^{\top} f_{\theta}^{(t)}\left(\mathbf{x}_{t+1}\right)$.
- For $\sigma_{t}<\frac{\sigma_{\mathbf{y}}}{s_{i}}$, the KL divergence is also $\frac{\left(\overline{\mathbf{x}}_{\theta, t}^{(i)}-\overline{\mathbf{x}}_{0}^{(i)}\right)^{2}}{2 \sigma_{t}^{2}}$.
- For $\sigma_{t} \geq \frac{\sigma_{\mathrm{y}}}{s_{i}}$, we have defined $\eta_{b}$ as a solution to the following quadratic equation (the other solution is 0 , which is irrelevant to our case since it does not make use of information from y):

$$
\begin{equation*}
\left(\sigma_{t}^{2}+\frac{\sigma_{\mathbf{y}}^{2}}{s_{i}^{2}}\right) \eta_{b}^{2}-2 \sigma_{t}^{2} \eta_{b}=0 \tag{21}
\end{equation*}
$$

reorganizing terms, we have that:

$$
\begin{align*}
&\left(\sigma_{t}^{2}+\frac{\sigma_{\mathbf{y}}^{2}}{s_{i}^{2}}\right) \eta_{b}^{2}-2 \sigma_{t}^{2} \eta_{b}+\sigma_{t}^{2}=\sigma_{t}^{2} \\
& \sigma_{t}^{2}\left(1-\eta_{b}\right)^{2}= \sigma_{t}^{2} \eta_{b}^{2}-2 \sigma_{t}^{2} \eta_{b}+\sigma_{t}^{2}=\sigma_{t}^{2}-\frac{\sigma_{\mathbf{y}}^{2}}{s_{i}^{2}} \eta_{b}^{2} \\
& \frac{\left(1-\eta_{b}\right)^{2}}{\sigma_{t}^{2}-\frac{\sigma_{\mathbf{y}}^{2}}{s_{i}^{2}} \eta_{b}^{2}}=\frac{1}{\sigma_{t}^{2}} \tag{22}
\end{align*}
$$

So the KL divergence is

$$
\frac{\left(1-\eta_{b}\right)^{2}}{2\left(\sigma_{t}^{2}-\frac{\sigma_{\mathbf{v}}^{2}}{s_{i}^{2}} \eta_{b}^{2}\right)}\left(\overline{\mathbf{x}}_{\theta, t}^{(i)}-\overline{\mathbf{x}}_{0}^{(i)}\right)^{2}=\frac{\left(\overline{\mathbf{x}}_{\theta, t}^{(i)}-\overline{\mathbf{x}}_{0}^{(i)}\right)^{2}}{2 \sigma_{t}^{2}}
$$

Therefore, regardless of how the cases are distributed among indices, we will always have that:
$D_{\mathrm{KL}}\left(q^{(t)}\left(\overline{\mathbf{x}}_{t} \mid \mathbf{x}_{t+1}, \mathbf{x}_{0}, \mathbf{y}\right) \| p_{\theta}^{(t)}\left(\overline{\mathbf{x}}_{t} \mid \mathbf{x}_{t+1}, \mathbf{y}\right)\right)=\sum_{i=1}^{n^{2}} \frac{\left(\overline{\mathbf{x}}_{\theta, t}^{(i)}-\overline{\mathbf{x}}_{0}^{(i)}\right)^{2}}{2 \sigma_{t}^{2}}=\frac{\left\|\overline{\mathbf{x}}_{\theta, t}-\overline{\mathbf{x}}_{0}\right\|_{2}^{2}}{2 \sigma_{t}^{2}}=\frac{\left\|f_{\theta}^{(t)}\left(\mathbf{x}_{t+1}\right)-\mathbf{x}_{0}\right\|_{2}^{2}}{2 \sigma_{t}^{2}}$.
Case II For $t=0$, we will only have two cases ( $s_{i}=0$ or $\sigma_{t}<\frac{\sigma_{\mathrm{y}}}{s_{i}}$ ), and thus, similar to Case I,
$\log p_{\theta}^{(0)}\left(\overline{\mathbf{x}}_{0} \mid \mathbf{x}_{1}, \mathbf{y}\right)=\sum_{i=1}^{n^{2}} \log p_{\theta}^{(0)}\left(\overline{\mathbf{x}}_{0}^{(i)} \mid \mathbf{x}_{1}, \mathbf{y}\right) \propto \sum_{i=1}^{n^{2}}\left(\overline{\mathbf{x}}_{\theta, 0}^{(i)}-\overline{\mathbf{x}}_{0}^{(i)}\right)^{2}=\left\|\overline{\mathbf{x}}_{\theta, 0}-\overline{\mathbf{x}}_{0}\right\|_{2}^{2}=\left\|f_{\theta}^{(0)}\left(\mathbf{x}_{1}\right)-\mathbf{x}_{0}\right\|_{2}^{2}$,
as long as we have a constant variance for $p_{\theta}^{(0)}$. Thus, every individual term in Equation 10 can be written as a denoising autoencoder objective, completing the proof.

## D Memory Efficient SVD

Here we explain how we obtained the singular value decomposition (SVD) for different degradation models efficiently.

## D. 1 Denoising

In denoising, the corrupted image is the original image with additive white Gaussian noise. Therefore, $\boldsymbol{H}=\boldsymbol{I}$ and all the SVD elements of $\boldsymbol{H}$ are simply the identity matrix $\boldsymbol{I}$, which in turns makes their multiplication by different vectors trivial.

## D. 2 Inpainting

In inpainting, $\boldsymbol{H}$ retains a known subset of size $k$ of the image's pixels. This is equivalent to permuting the pixels such that the retained one are placed at the top, then keeping the first $k$ entries. Therefore,

$$
\begin{equation*}
H=I \Sigma P \tag{23}
\end{equation*}
$$

where $\boldsymbol{P}$ is the appropriate permutation matrix, $\boldsymbol{\Sigma}$ is a rectangular diagonal matrix of size $k \times n$ with ones in its main diagonal, and $\boldsymbol{I}$ is the identity matrix. Since permutation matrices are orthogonal, Equation 23 is the SVD of $\boldsymbol{H}$.
We can multiply a given vector by $\boldsymbol{P}$ and $\boldsymbol{P}^{T}$ by storing the permutation itself rather than the matrix. $\boldsymbol{\Sigma}$ can multiply a vector by simply slicing it. Therefore, by storing the appropriate permutation and the number $k$, we can apply each element of the SVD with $\Theta(n)$ space complexity.

## D. 3 Super Resolution

For super resolution, we assume that the original image of size $d \times d$ (i.e. $n=3 d^{2}$ ) is downscaled using a block averaging filter by $r$ in each dimension, such that $d$ is divisible by $r$. In this scenario, each pixel in the output image is the average of an $r \times r$ patch in the input image, and each such patch affects exactly one output pixel. Therefore, any output pixel is given by $(\boldsymbol{H} \mathbf{x})_{i}=\boldsymbol{k}^{T} \boldsymbol{p}_{i}$, where $\boldsymbol{k}$ is a vector of size $r^{2}$ with $\frac{1}{r^{2}}$ in each entry, and $\boldsymbol{p}_{i}$ is the vectorized $i$-th $r \times r$ patch. More formally, if $\boldsymbol{P}_{1}$ is a permutation matrix that reorders a vectorized image into patches, then

$$
\boldsymbol{H}=\left(\boldsymbol{I} \otimes \boldsymbol{k}^{T}\right) \boldsymbol{P}_{1}
$$

where $\otimes$ is the Kronecker product, and $\boldsymbol{I}$ is the identity matrix of size $\frac{d}{r} \times \frac{d}{r}$. In order to obtain the SVD of $\boldsymbol{H}$, we calculate the SVD of $\boldsymbol{k}^{T}$ :

$$
\boldsymbol{k}^{T}=\boldsymbol{U}_{\boldsymbol{k}} \boldsymbol{\Sigma}_{\boldsymbol{k}} \boldsymbol{V}_{\boldsymbol{k}}^{T}
$$

Using properties of the Kronecker product, we observe

$$
\begin{align*}
\boldsymbol{H} & =\left(\boldsymbol{I} \otimes \boldsymbol{k}^{T}\right) \boldsymbol{P}_{1}=\left((\boldsymbol{I I I}) \otimes\left(\boldsymbol{U}_{\boldsymbol{k}} \boldsymbol{\Sigma}_{\boldsymbol{k}} \boldsymbol{V}_{\boldsymbol{k}}^{T}\right)\right) \boldsymbol{P}_{1}  \tag{24}\\
& =\left(\boldsymbol{I} \otimes \boldsymbol{U}_{\boldsymbol{k}}\right)\left(\boldsymbol{I} \otimes \boldsymbol{\Sigma}_{\boldsymbol{k}}\right)\left(\boldsymbol{I} \otimes \boldsymbol{V}_{\boldsymbol{k}}^{T}\right) \boldsymbol{P}_{1}
\end{align*}
$$

The Kronecker product of two orthogonal matrices is an orthogonal matrix. Therefore, $\boldsymbol{I} \otimes \boldsymbol{U}_{\boldsymbol{k}}$ and $\boldsymbol{I} \otimes \boldsymbol{V}_{\boldsymbol{k}}^{T}$ are orthogonal. Observe that the matrix $\boldsymbol{I} \otimes \boldsymbol{\Sigma}_{\boldsymbol{k}}$ has one non-zero value $\left(\frac{1}{r^{2}}\right)$ in each row. By applying a simple permutation on its columns, these values can be reordered to be on the main diagonal. We denote the appropriate permutation matrix by $\boldsymbol{P}_{2}$, and obtain

$$
\begin{equation*}
\boldsymbol{H}=\boldsymbol{U} \boldsymbol{\Sigma} \boldsymbol{V}^{T} \tag{25}
\end{equation*}
$$

where $\boldsymbol{U}=\boldsymbol{I} \otimes \boldsymbol{U}_{\boldsymbol{k}}$ is orthogonal, $\boldsymbol{\Sigma}=\left(\boldsymbol{I} \otimes \boldsymbol{\Sigma}_{\boldsymbol{k}}\right) \boldsymbol{P}_{2}^{T}$ is a rectangular diagonal matrix with nonnegative entries, and $\boldsymbol{V}^{T}=\boldsymbol{P}_{2}\left(\boldsymbol{I} \otimes \boldsymbol{V}_{\boldsymbol{k}}^{T}\right) \boldsymbol{P}_{1}$ is orthogonal. As such, Equation 25 is the SVD of $\boldsymbol{H}$. By storing the permutations and the SVD elements of $\boldsymbol{k}^{T}$, we can simulate each element of the SVD of $\boldsymbol{H}$ with $\Theta(n)$ space complexity, without directly calculating the Kronecker products with $\boldsymbol{I}$.

## D. 4 Colorization

The grayscale image is obtained by averaging the red, green, and blue channels of each pixel. This means that every output pixel is given by $(\boldsymbol{H} \mathbf{x})_{i}=\boldsymbol{k}^{T} \boldsymbol{p}_{i}$, where $\boldsymbol{k}^{T}=\left(\begin{array}{lll}\frac{1}{3} & \frac{1}{3} & \frac{1}{3}\end{array}\right)$ and $\boldsymbol{p}_{i}$ is the 3 -valued $i$-th pixel of the original color image. The SVD of $\boldsymbol{H}$ is obtained exactly the same as in the super resolution case, with separate pixels replacing separate patches.

## D. 5 Deblurring

We focus on separable blurring, where the 2D blurring kernel is $\boldsymbol{K}=\boldsymbol{r} \boldsymbol{c}^{T}$, which means $\boldsymbol{c}$ is applied on the columns of the image, and $\boldsymbol{r}^{T}$ is applied on its rows. The blurred image can be obtained by $\boldsymbol{B}=\boldsymbol{A}_{c} \boldsymbol{X} \boldsymbol{A}_{r}^{T}$, where $\boldsymbol{A}_{c}$ and $\boldsymbol{A}_{r}$ apply a 1D convolution with kernels $\boldsymbol{c}$ and $\boldsymbol{r}$, respectively. Alternatively, $\boldsymbol{b}=\boldsymbol{H} \boldsymbol{x}$, where $\boldsymbol{x}$ is the vectorized image $\boldsymbol{X}, \boldsymbol{b}$ is the vectorized blurred image $\boldsymbol{B}$, and $\boldsymbol{H}$ is the matrix applying the 2 D convolution $\boldsymbol{K}$. It can be shown that $\boldsymbol{H}=\boldsymbol{A}_{r} \otimes \boldsymbol{A}_{c}$, where $\otimes$ is the Kronecker product. In order to calculate the SVD of $\boldsymbol{H}$, we calculate the SVD of $\boldsymbol{A}_{r}$ and $\boldsymbol{A}_{c}$ :

$$
\boldsymbol{A}_{r}=\boldsymbol{U}_{r} \boldsymbol{\Sigma}_{r} \boldsymbol{V}_{r}^{T}, \quad \boldsymbol{A}_{c}=\boldsymbol{U}_{c} \boldsymbol{\Sigma}_{c} \boldsymbol{V}_{c}^{T} .
$$

Using the properties of the Kronecker product, we observe

$$
\begin{align*}
\boldsymbol{H} & =\boldsymbol{A}_{r} \otimes \boldsymbol{A}_{c}=\left(\boldsymbol{U}_{r} \boldsymbol{\Sigma}_{r} \boldsymbol{V}_{r}^{T}\right) \otimes\left(\boldsymbol{U}_{c} \boldsymbol{\Sigma}_{c} \boldsymbol{V}_{c}^{T}\right) \\
& =\left(\boldsymbol{U}_{r} \otimes \boldsymbol{U}_{c}\right)\left(\boldsymbol{\Sigma}_{r} \otimes \boldsymbol{\Sigma}_{c}\right)\left(\boldsymbol{V}_{r} \otimes \boldsymbol{V}_{c}\right)^{T} \tag{26}
\end{align*}
$$

The Kronecker product preserves orthogonality. Therefore, Equation 26 is a valid SVD of $\boldsymbol{H}$, with the exception of the singular values not being on the main diagonal, and not being sorted descendingly. We reorder the columns so that the singular values are on the main diagonal and denote the corresponding permutation matrix by $\boldsymbol{P}_{1}$. We also sort the values descendingly and denote the sorting permutation matrix by $\boldsymbol{P}_{2}$, and obtain the following SVD:

$$
\begin{equation*}
\boldsymbol{H}=\boldsymbol{U} \boldsymbol{\Sigma} \boldsymbol{V}^{T} \tag{27}
\end{equation*}
$$

where $\boldsymbol{U}=\left(\boldsymbol{U}_{r} \otimes \boldsymbol{U}_{c}\right) \boldsymbol{P}_{2}^{T}, \boldsymbol{\Sigma}=\boldsymbol{P}_{2}\left(\boldsymbol{\Sigma}_{r} \otimes \boldsymbol{\Sigma}_{c}\right) \boldsymbol{P}_{1}^{T} \boldsymbol{P}_{2}^{T}$, and $\boldsymbol{V}^{T}=\boldsymbol{P}_{2} \boldsymbol{P}_{1}\left(\boldsymbol{V}_{r} \otimes \boldsymbol{V}_{c}\right)^{T}$.
For every matrix of the form $\boldsymbol{M}=\boldsymbol{N} \otimes \boldsymbol{L}$, it holds that $\boldsymbol{M} x$ is the vectorized version of $\boldsymbol{L} \boldsymbol{X} \boldsymbol{N}^{T}$. By using this property and applying the relevant permutation, we can simulate multiplying a vector by $\boldsymbol{U}, \boldsymbol{V}, \boldsymbol{U}^{T}$, or $\boldsymbol{V}^{T}$ without storing the full matrix. The space complexity of this approach is $\Theta(n)$, which is required for computing the SVD of $\boldsymbol{A}_{r}$ and $\boldsymbol{A}_{c}$, as well as storing the permutations.
The above calculations remain valid when the blurring is zero-padded, i.e., images are padded with zeroes so that the convolution is not circulant around the edges. We consider a zero-padded deblurring problem in our experiments. Note that the noiseless version of this problem has a simple solution applying the pseudo-inverse of the blurring matrix on the blurry image. This solution attains 32.41 dB in PSNR on ImageNet-1K, while DDRM improves upon it and achieves 35.64 dB . When noise is added to the blurry image, such a simple solution amplifies the noise and fails to provide a valid output. Therefore, we opt not to report its results.
Furthermore, the above calculations are also applicable to blurring with strided convolutions. We use this fact in our implementation of the bicubic super resolution SVD, which can be interpreted as a strided convolution with a fixed kernel.

Table 3: Ablation studies on $\eta$ and $\eta_{b}$.
(a) $\operatorname{PSNR}(\uparrow)$.
(b) KID $\times 10^{3}(\downarrow)$.

|  | $\eta_{b}$ | 0.7 | 0.8 | 0.9 | 1.0 |  | $\eta_{b}$ | 0.7 | 0.8 |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\eta$ |  | 25.16 | 25.19 | 25.20 | 25.20 | 0.9 | 1.0 |  |  |
| 0.7 | 25.17 | 25.23 | 25.27 | 25.29 | 0.8 | 16.27 | 14.30 | 12.76 | 11.65 |
| 0.8 | 25.07 | 19.07 | 17.37 | 15.98 |  |  |  |  |  |
| 0.9 | 25.07 | 25.18 | 25.26 | 25.32 | 0.9 |  | 27.85 | 25.64 | 23.81 |
| 1.0 | 24.54 | 25.75 | 24.91 | 25.04 | 1.0 |  | 45.10 | 42.50 | 40.10 |

## E Ablation Studies on Hyperparameters

$\eta$ and $\eta_{\mathbf{b}}$. Apart from the timestep schedules, DDRM has two hyperparameters $\eta$ and $\eta_{b}$, which control the level of noise injected at each timestep. To identify an ideal combination, we perform a hyperparameter search over $\eta, \eta_{b} \in\{0.7,0.8,0.9,1.0\}$ for the task of deblurring with $\sigma_{y}=0.05$ in 1000 ImageNet validation images, using the model trained in [13]. It is possible to also consider different $\eta$ values for $s_{i}=0$ and $\sigma_{i}<\sigma_{\mathbf{y}} / s_{i}$; we leave that as future work.
We report PSNR and KID results in Table 3. From the results, we observe that generally (i) as $\eta_{b}$ increases, PSNR increases while KID decreases, which is reasonable given that we wish to leverage the information from $\mathbf{y}$; (ii) as $\eta$ increases, PSNR increases (except for $\eta=1.0$ ) yet KID also increases, which presents a trade-off in reconstruction error and image quality (known as the perception-distortion trade-off [6]). Therefore, we choose $\eta_{b}=1$ and $\eta=0.85$ to balance performance on PSNR and KID when we report results.

Timestep schedules. The timestep schedule has a direct impact on NFEs, as the wall-clock time is roughly linear with respect to NFEs [45]. In Tables 5 and 6, we compare the PSNR, FID, and KID of DDRM with 20 or 100 timesteps (with or without conditioning) and default $\eta=0.85$ and $\eta_{b}=1$. We observe that DDRM with 20 or 100 timesteps have similar performance when other hyperparameters are identical, with DDRM (20) having a slight edge in FID and KID.

## F Experimental Setup of DGP, RED, and SNIPS

Recall that we evaluated DGP [38], RED [40], and SNIPS [25] on $256 \times 256$ ImageNet 1 K images, for the problems of $4 \times$ super resolution and deblurring without any noise in the measurements. Below we expand on the experimental setup of each one.
For DGP [38], we use the same hyperparameters introduced in the original paper for MSE-biased super resolution. We note that the downscaling applied in DGP is different from the block averaging filter that we used, and the numbers they reported are on the $128 \times 128$ resolution. Nevertheless, in our experiments, DGP achieved a PSNR of 23.06 on ImageNet $1 \mathrm{~K} 256 \times 256$ block averaging $4 \times$ super resolution, which is similar to the 23.30 reported in the original work. When applied on the deblurring problem, we retained the same DGP hyperparameters as well.

For RED [40], we apply the iterative algorithm only in the luminance channel of the image in the YCbCr space, as done in the original paper for deblurring and super resolution. As for the denoising engine enabling the algorithm, we use the same diffusion model used in DDRM to enable as fair a comparison as possible. We use the last step of the diffusion model (equivalent to denoising with $\sigma=0.005$ ), as we found it to work best empirically. We also chose the steepest-descent version (RED-SD), and $\lambda=500$ for best PSNR performance given the denoiser we used. We also set $\sigma_{0}=0.01$ when the measurements are noiseless, because $\sigma_{0}$ cannot be 0 as RED divides by it.

In super resolution, RED is initialized with the bicubic upsampled low-res image. In deblurring, it is initialized with the blurry image. We then run RED on the ImageNet 1 K for different numbers of steps (see Table 4), and choose the best PSNR for each problem. Namely, we show in our paper RED on super resolution with 100 steps, and on deblurring with 500 steps. Interestingly, RED achieves a PSNR close to its best for super resolution in just 20 steps. However, DDRM (with 20 steps) still

Table 4: RED results on ImageNet $1 \mathrm{~K}(256 \times 256)$ for $4 \times$ super resolution and deblurring for different numbers of steps.

|  | SUPER-RES |  | DEBLURRING |  |
| :--- | :---: | :---: | :---: | :---: |
| STEPS | PSNR $\uparrow$ | KID $\downarrow$ | PSNR $\uparrow$ | KID $\downarrow$ |
| 0 | 25.65 | 44.90 | 19.26 | 38.00 |
| 20 | 26.05 | 52.51 | 23.49 | 21.99 |
| 100 | 26.08 | 53.55 | 25.00 | 26.09 |
| 500 | 26.00 | 54.19 | 26.16 | 21.21 |

Table 5: ImageNet 50K validation set $(256 \times 256)$ results on $4 \times$ super resolution with additive noise of $\sigma_{\mathbf{y}}=0.05$.

| METHOD | PSNR $\uparrow$ | FID $\downarrow$ | KID $\downarrow$ | NFES $\downarrow$ |
| :--- | :---: | :---: | :---: | :---: |
| BICUBIC | 22.65 | 64.24 | 50.56 | $\mathbf{0}$ |
| DDRM | 24.70 | 20.16 | 15.25 | 100 |
| DDRM-CC | $\mathbf{2 4 . 7 1}$ | 18.22 | 13.57 | 100 |
| DDRM | 24.29 | 17.88 | 13.18 | 20 |
| DDRM-CC | 24.30 | $\mathbf{1 5 . 9 2}$ | $\mathbf{1 1 . 4 7}$ | 20 |

outperforms RED in PSNR, with substantially better perceptual quality (see Table 1 in the main paper).
Another interesting plug-and-play image restoration method is DPIR [57], which has recently achieved impressive results. It does so by applying the well-known Half Quadratic Splitting (HQS) plug-and-play algorithm using a newly proposed architecture. HQS requires an analytical solution of a minimization problem which is infeasible in general, due to the high memory requirements. DPIR provides efficient solutions for the specific degradation matrices $\mathbf{H}$ considered (circulant blurring, bicubic downsampling), which are different from the ones we consider (zero-padded blurring, block downsampling). In order to draw a fair comparison between the algorithms, one would have to use the same denosier architecture in both (as we have done for RED and SNIPS), and use the same degradation models. To apply DPIR on the same problems that we consider, we would need to substantially modify it and introduce efficient solutions. Therefore, we instead compare to RED, an alternative plug-and-play method.
SNIPS [25] did not originally work with ImageNet images. However, considering the method's similarity to DDRM (as both operate in the spectral space of $\boldsymbol{H}$ ), a comparison is necessary. We apply SNIPS with the same underlying diffusion model (with all 1000 timesteps) as DDRM for fairness. SNIPS evaluates the diffusion model $\tau$ times for each timestep. We set $\tau=1$ so that SNIPS' runtime remains reasonable in comparison to the rest of the considered methods, and do not explore higher values of $\tau$. It is worth mentioning that in the original work, $\tau$ was set to 3 for an LSUN bedrooms diffusion model with 1086 timesteps. We set $c=0.67$ as it achieved the best PSNR performance.
The original work in SNIPS calculates the SVD of $\boldsymbol{H}$ directly, which hinders its ability to handle $256 \times 256$ images on typical hardware. In order to draw comparisons, we replaced the direct calculation of the SVD with our efficient implementation detailed in Appendix D.

In Figure 4 and Table 2 in the main paper, we show that DGP, RED, and SNIPS all fail to produce viable results when significant noise is added to the measurements. For these results, we use the same hyperparameters used in the noiseless case for all algorithms (except $\sigma_{\mathbf{y}}$ where applicable). While tuning the hyperparameters may boost performance, we do not explore that option as we are only interested in algorithms where given $\boldsymbol{H}$ and $\sigma_{\mathbf{y}}$, the restoration process is automatic. To further demonstrate DDRM's capabilities and speed, we evaluate it on the entire 50, 000-image ImageNet validation set in Tables 5 and 6, reporting Fréchet Inception distance (FID) [18] as well as KID, as enough samples are available.

Table 6: ImageNet 50 K validation set $(256 \times 256)$ results on deblurring with additive noise of $\sigma_{\mathbf{y}}=0.05$.

| METHOD | PSNR $\uparrow$ | FID $\downarrow$ | KID $\downarrow$ | NFES $\downarrow$ |
| :--- | :---: | :---: | :---: | :---: |
| BLURRY | 18.05 | 93.36 | 74.13 | 0 |
| DDRM | 24.23 | 22.30 | 16.23 | 100 |
| DDRM-CC | 24.21 | 20.06 | 14.20 | 100 |
| DDRM | 24.60 | 21.60 | 15.65 | 20 |
| DDRM-CC | $\mathbf{2 4 . 6 1}$ | $\mathbf{1 9 . 6 6}$ | $\mathbf{1 3 . 9 4}$ | 20 |

## G Runtime of Algorithms

In the main paper, we show the number of neural function evaluations (NFEs) as a proxy for the runtime of algorithms. Here, we consider the case of noisy deblurring, and measure the runtime of DDRM, RED, SNIPS, and DGP on an Nvidia RTX 3080 GPU. For each image, DDRM, RED, and SNIPS all run at around 0.09 s/it (seconds per iteration), with negligible differences of $<0.01 \mathrm{~s} / i t$. We note that the denoiser model of DDRM, SNIPS, and RED is the same, so runtime is almost perfectly linearly correlated with NFEs. As for DGP, it uses a different model (a GAN), and it is slightly slower than our denoiser ( $0.11 \mathrm{~s} / \mathrm{it}$ ); this is partly because DGP requires additional gradient computations in order to perform an update. All in all, we observe that the runtime is indeed linear with NFEs, and since no algorithm has a significant runtime advantage over the rest, we prefer to use NFEs as a proxy for runtime, as it is a hardware-independent measure.

In this paper, we used pretrained generative models for image restoration. Since we didn't train any models, a single Nvidia RTX 3080 GPU was sufficient to run all experiments that were shown in the paper and the appendices.

## H ILVR as a special case of DDRM

Given a generative diffusion model (e.g. DDPM [19]) that can predict $\mathbf{x}$ given $\mathbf{x}_{t+1}$ and $t+1$ for $t \in[0, T-1]$, and a noiseless measurement $\mathbf{y}=\boldsymbol{H} \mathbf{x}$, where $\boldsymbol{H}$ is a downscaling matrix, the Iterative Latent Variable Refinement (ILVR) [9] algorithm can sample from the posterior distribution $p_{\theta}^{(t)}\left(\mathbf{x}_{t} \mid \mathbf{x}_{t+1}, \mathbf{y}\right)$ for $t \in[0, T-1]$.
We assume a variance exploding diffusion model, i.e. $\mathbf{x}_{t}=\mathbf{x}+\sigma_{t} \epsilon_{t}$ where $\epsilon_{t} \sim \mathcal{N}(0, \boldsymbol{I})$, without loss of generality (because it is equivalent to the variance preserving scheme, as we show in Appendix B). Under this setting, ILVR applies the following updates for $t=T-1, \ldots, 0$ :

$$
\begin{aligned}
\mathbf{x}_{t}^{\prime} & =\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t} \\
\mathbf{y}_{t} & =\boldsymbol{H}^{\dagger} \mathbf{y}+\sigma_{t} \epsilon_{t}^{\prime} \\
\mathbf{x}_{t} & =\mathbf{x}_{t}^{\prime}-\boldsymbol{H}^{\dagger} \boldsymbol{H} \mathbf{x}_{t}^{\prime}+\boldsymbol{H}^{\dagger} \boldsymbol{H} \mathbf{y}_{t}
\end{aligned}
$$

where $\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)$ is the prediction for $\mathbf{x}$ given by the diffusion model at timestep $t+1$, $\epsilon_{t} \sim \mathcal{N}(0, \boldsymbol{I})$, and $\epsilon_{t}^{\prime} \sim \mathcal{N}(0, \boldsymbol{I})$. Substituting $\mathbf{x}_{t}^{\prime}, \mathbf{y}_{t}$, and $\boldsymbol{H}=\boldsymbol{U} \boldsymbol{\Sigma} \boldsymbol{V}^{T}$, the last equation becomes

$$
\begin{aligned}
\mathbf{x}_{t} & =\mathbf{x}_{t}^{\prime}-\boldsymbol{H}^{\dagger} \boldsymbol{H} \mathbf{x}_{t}^{\prime}+\boldsymbol{H}^{\dagger} \boldsymbol{H} \mathbf{y}_{t} \\
& =\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}-\boldsymbol{H}^{\dagger} \boldsymbol{H}\left(\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}\right)+\boldsymbol{H}^{\dagger} \boldsymbol{H}\left(\boldsymbol{H}^{\dagger} \mathbf{y}+\sigma_{t} \epsilon_{t}^{\prime}\right) \\
& =\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}-\boldsymbol{V} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \boldsymbol{U} \boldsymbol{\Sigma} \boldsymbol{V}^{T}\left(\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}\right)+\boldsymbol{V} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \boldsymbol{U} \boldsymbol{\Sigma} \boldsymbol{V}^{T}\left(\boldsymbol{V} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \mathbf{y}+\sigma_{t} \epsilon_{t}^{\prime}\right) \\
& =\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}-\boldsymbol{V} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V}^{T}\left(\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}\right)+\boldsymbol{V} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V}^{T}\left(\boldsymbol{V} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \mathbf{y}+\sigma_{t} \epsilon_{t}^{\prime}\right) \\
& =\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}-\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V} \boldsymbol{V}^{T}\left(\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}\right)+\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V} \boldsymbol{V}^{T}\left(\boldsymbol{V} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \mathbf{y}+\sigma_{t} \epsilon_{t}^{\prime}\right) \\
& =\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}-\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma}\left(\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}\right)+\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma}\left(\boldsymbol{V} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \mathbf{y}+\sigma_{t} \epsilon_{t}^{\prime}\right)
\end{aligned}
$$

The second to last equality holds because $\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma}$ is a square diagonal matrix, and matrix multiplication with a square diagonal matrix is commutative. Recall that $\overline{\mathbf{x}}_{t}=\boldsymbol{V}^{T} \mathbf{x}_{t}, \overline{\mathbf{y}}=\boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \mathbf{y}$, and $\overline{\mathbf{x}}_{\theta, t}=$

$$
\begin{aligned}
& \boldsymbol{V}^{T} \mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right), \text { thus } \\
& \overline{\mathbf{x}}_{t}=\boldsymbol{V}^{T} \mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \boldsymbol{V}^{T} \epsilon_{t}-\boldsymbol{V}^{T} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma}\left(\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}\right)+\boldsymbol{V}^{T} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma}\left(\boldsymbol{V} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \mathbf{y}+\sigma_{t} \epsilon_{t}^{\prime}\right) \\
&=\boldsymbol{V}^{T} \mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \boldsymbol{V}^{T} \epsilon_{t}-\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V}^{T}\left(\mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \epsilon_{t}\right)+\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V}^{T}\left(\boldsymbol{V} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \mathbf{y}+\sigma_{t} \epsilon_{t}^{\prime}\right) \\
&=\boldsymbol{V}^{T} \mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)+\sigma_{t} \boldsymbol{V}^{T} \epsilon_{t}-\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V}^{T} \mathbf{x}_{\theta, t}\left(\mathbf{x}_{t+1}, t+1\right)-\sigma_{t} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V}^{T} \epsilon_{t}+\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V}^{T} \boldsymbol{V} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \mathbf{y}+\sigma_{t} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V}^{T} \epsilon_{t}^{\prime} \\
&=\overline{\mathbf{x}}_{\theta, t}+\sigma_{t} \boldsymbol{V}^{T} \epsilon_{t}-\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \overline{\mathbf{x}}_{\theta, t}-\sigma_{t} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V}^{T} \epsilon_{t}+\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \mathbf{y}+\sigma_{t} \boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \boldsymbol{V}^{T} \epsilon_{t}^{\prime} \\
&=\left(\boldsymbol{I}-\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma}\right) \overline{\mathbf{x}}_{\theta, t}+\left(\boldsymbol{I}-\boldsymbol{\Sigma}^{\dagger \boldsymbol{\Sigma}) \sigma_{t} \boldsymbol{V}^{T} \epsilon_{t}+\boldsymbol{\Sigma}^{\dagger} \boldsymbol{U}^{T} \mathbf{y}+\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \sigma_{t} \boldsymbol{V}^{T} \epsilon_{t}^{\prime}}\right. \\
&=\left(\boldsymbol{I}-\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma}\right) \overline{\mathbf{x}}_{\theta, t}+\left(\boldsymbol{I}-\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma}\right) \sigma_{t} \boldsymbol{V}^{T} \epsilon_{t}+\overline{\mathbf{y}}+\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma} \sigma_{t} \boldsymbol{V}^{T} \epsilon_{t}^{\prime} .
\end{aligned}
$$

The matrix $\boldsymbol{\Sigma}^{\dagger} \boldsymbol{\Sigma}$ is a square diagonal matrix with zeroes in its entries where the singular value is zero, and ones otherwise. In addition, $\boldsymbol{\Sigma}^{\dagger}$ has a row of zeroes when the singular value is zero. Therefore, it holds that

$$
\overline{\mathbf{x}}_{t}^{(i)}= \begin{cases}\overline{\mathbf{x}}_{\theta, t}^{(i)}+\left(\sigma_{t} \boldsymbol{V}^{T} \epsilon_{t}\right)^{(i)} & \text { if } s_{i}=0  \tag{28}\\ \overline{\mathbf{y}}^{(i)}+\left(\sigma_{t} \boldsymbol{V}^{T} \epsilon_{t}^{\prime}\right)^{(i)} & \text { if } s_{i} \neq 0\end{cases}
$$

which in turn implies

$$
p_{\theta}^{(t)}\left(\overline{\mathbf{x}}_{t}^{(i)} \mid \mathbf{x}_{t+1}, y\right)=\left\{\begin{array}{ll}
\mathcal{N}\left(\overline{\mathbf{x}}_{\theta, t}^{(i)}, \sigma_{t}^{2} \boldsymbol{I}\right) & \text { if } s_{i}=0  \tag{29}\\
\mathcal{N}\left(\overline{\mathbf{y}}^{(i)}, \sigma_{t}^{2} \boldsymbol{I}\right) & \text { if } s_{i} \neq 0
\end{array} .\right.
$$

This distribution is exactly the same as Equation 8 in the main paper when $\eta=\eta_{b}=1$ and $\sigma_{\mathbf{y}}=0$.
As for $\mathbf{x}_{T}$, ILVR initializes it by sampling from $\mathcal{N}\left(0, \sigma_{T}^{2} \boldsymbol{I}\right)$ (or $\mathcal{N}(0, \boldsymbol{I})$ in the variance preserving case) while DDRM samples according to Equation 7 in the main paper. The two initializations have the same variance but differ in the mean. This difference has a negligible effect on the end result since the variance is much larger than the difference in the means. Therefore, the above form of ILVR is a specific form of a DDRM (with $\eta=\eta_{b}=1$ ), posed as a solution for linear inverse problems without noise in the measurements.
In their experiments, ILVR only tested $\boldsymbol{H}$ which is the bicubic downscaling matrix with varying scale factors. In theory, ILVR can also work for any linear degradation $\boldsymbol{H}$, as long as $\mathbf{y}$ does not contain noise.

## I Additional Results

We provide additional figures below showing DDRM's versatility across different datasets, inverse problems, and noise levels (Figures 7, 8, 10, 11, and 12). We also showcase the sample diversity provided by DDRM in Figure 9; we present more uncurated samples from the ImageNet experiments in Figures 13 and 14. Moreover, we further illustrate DDRM's advantage over previous unsupervised methods by evaluating on two additional inverse problems: (i) $4 \times$ super-resolution with the popular bicubic downsampling kernel; and (ii) deblurring with an anisotropic Gaussian blur kernel (with $\sigma=20$ horizontally and $\sigma=1$ vertically), mimicing motion blur. We show both noiseless and noisy versions in Tables 7 and 8 , respectively. To maintain the unsupervised nature of the tested methods, we use the same hyperparameters as in block-averaging super-resolution and uniform deblurring.

Table 7: Noiseless $4 \times$ super-resolution (using a bicubic kernel) and anisotropic Gaussian deblurring results on ImageNet $1 \mathrm{~K}(256 \times 256)$.

| Method | $4 \times$ super-resolution (Bicubic) |  |  | Deblurring (Anisotropic) |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | PSNR $\uparrow$ | SSIM $\uparrow$ | KID $\downarrow$ | NFEs $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | KID $\downarrow$ | NFEs $\downarrow$ |
| Baseline | 26.06 | 0.73 | 72.41 | 0 | 19.96 | 0.58 | 25.23 | 0 |
| DGP | 20.82 | 0.50 | 29.62 | 1500 | 23.35 | 0.59 | 20.10 | 1500 |
| RED | 26.14 | 0.73 | 47.61 | 100 | 29.39 | 0.86 | 10.49 | 500 |
| SNIPS | 17.65 | 0.23 | 30.30 | 1000 | 33.34 | 0.86 | 0.58 | 1000 |
| DDRM | $\mathbf{2 7 . 0 9}$ | $\mathbf{0 . 7 6}$ | $\mathbf{1 2 . 7 8}$ | 20 | $\mathbf{3 6 . 0 2}$ | $\mathbf{0 . 9 3}$ | $\mathbf{0 . 4 1}$ | 20 |

Table 8: Noisy ( $\sigma_{\mathbf{y}}=0.05$ ) $4 \times$ super-resolution (using a bicubic kernel) and anisotropic Gaussian deblurring results on ImageNet $1 \mathrm{~K}(256 \times 256)$.

| Method | $4 \times$ super-resolution (Bicubic) |  |  | Deblurring (Anisotropic) |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | PSNR $\uparrow$ | SSIM $\uparrow$ | KID $\downarrow$ | NFEs $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | KID $\downarrow$ | NFEs $\downarrow$ |
| Baseline | 21.68 | 0.40 | 73.87 | 0 | 19.96 | 0.27 | 55.00 | 0 |
| DGP | 19.68 | 0.40 | 44.07 | 1500 | 22.64 | 0.53 | 25.38 | 1500 |
| RED | 22.65 | 0.46 | 54.90 | 100 | 11.97 | 0.10 | 130.30 | 500 |
| SNIPS | 16.16 | 0.14 | 69.69 | 1000 | 17.49 | 0.20 | 48.37 | 1000 |
| DDRM | $\mathbf{2 5 . 5 3}$ | $\mathbf{0 . 6 8}$ | $\mathbf{1 4 . 5 7}$ | 20 | $\mathbf{2 6 . 9 5}$ | $\mathbf{0 . 7 3}$ | $\mathbf{1 0 . 3 4}$ | 20 |



Figure 7: Pairs of low-res and recovered $256 \times 256$ face images with a 20-step DDRM. Noisy low-res images contain noise with a standard deviation of $\sigma_{\mathbf{y}}=0.1$.


Figure 8: Pairs of degraded and recovered $256 \times 256$ face images with a 20-step DDRM. Degraded images contain noise with a standard deviation of $\sigma_{\mathbf{y}}=0.1$.


Figure 9: Original, degraded, and 6 recovered $256 \times 256$ face images with a 20 -step DDRM. Degraded images contain noise with a standard deviation of $\sigma_{\mathbf{y}}=0.1$. The mean and standard deviation (scaled by 4 ) of the sampled solution is shown.


Figure 10: Pairs of degraded and recovered $256 \times 256$ bedroom images with a 20 -step DDRM.
Degraded images contain noise with a standard deviation of $\sigma_{\mathbf{y}}=0.05$.


Figure 11: Pairs of degraded and recovered $256 \times 256$ cat images with a 20 -step DDRM. Degraded images contain noise with a standard deviation of $\sigma_{\mathbf{y}}=0.05$.


Figure 12: Pairs of degraded and recovered $256 \times 256$ USC-SIPI images with a 20 -step DDRM using an ImageNet diffusion model. Degraded images contain noise with a standard deviation of $\sigma_{\mathbf{y}}=0.05$.


Figure 13: Uncurated samples from the noisy $4 \times$ super resolution ( $\sigma_{\mathbf{y}}=0.05$ ) task on $256 \times 256$ ImageNet 1 K . Each triplet contains (from left to right): the original image, the low-res image, and the restored image with DDRM-20.


Figure 14: Uncurated samples from the noisy deblurring ( $\sigma_{\mathbf{y}}=0.05$ ) task on $256 \times 256$ ImageNet 1 K . Each triplet contains (from left to right): the original image, the blurry image, and the restored image with DDRM-20.


Figure 15: Denoising ( $\sigma_{\mathbf{y}}=0.75$ ) face images. DDRM restores more fine details (e.g. hair) than an MMSE denoiser. The denoiser used here is the denoising diffusion function $f_{\theta}\left(\mathbf{x}_{t}, t\right)$ used by DDRM, where $t$ minimizes $\left|\sigma_{t}-\sigma_{\mathbf{y}}\right|$.

