Rebuttal PDF for the Submission
"Mirror Diffusion Models for Constrained and Watermarked Generation"

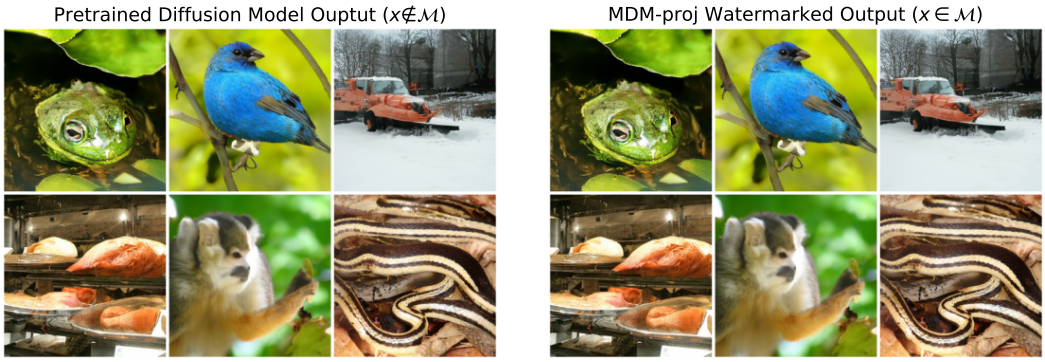

Figure 12: Unconditional watermarked generation on ImageNet 256×256. Similar to Section 5.2, we consider a polytope constraint set, $\mathcal{M} := \{x \in \mathbb{R}^d : c_i < a_i^\top x < b_i, \forall i \in [m]\}$, where $m = 100$, $b = 1.2$, $c = -1.2$, and a_i orthogonal Gaussian random vectors, except in a much higher-dimensional space, $d = 196608$. For comparison, we also include the results of pretrained diffusion model, which, despite being indistinguishable, actually violate the polytope constraint, *i.e.,* $x \notin \mathcal{M}$. It is clear that our MDM-proj is capable of embedding invisible watermarks to high-resolution images.

Figure 13: Conditional watermarked generation on ImageNet 256×256. Specifically, we consider the JPEG restoration task, where, given degraded, low-quality inputs y (*c.f.*, upper-left), we wish to generate their corresponding clean images x (*c.f.*, upper-right) by modeling $p(x|y)$. It is clear that both MDM-dual and MDM-proj are capable of solving this conditional generation task, generating clean images that additionally embed invisible watermarks, *i.e.,* $x \in \mathcal{M}$.