

384 **A Maintaining log-submodularity in the generative model**

385 **Theorem 1.** Let f be a strictly submodular function over subsets of a ground set \mathcal{Y} , and g be a
 386 function over the same space such that

$$\|f - g\|_\infty \leq \min_{\substack{S \neq T \\ S, T \notin \{\emptyset, \mathcal{Y}\}}} \frac{1}{4} [f(S) + f(T) - f(S \cup T) - f(S \cap T)]. \quad (4)$$

387 Then g is also submodular.

388 *Proof.* In all the following, we assume that S, T are subsets of a ground set \mathcal{Y} such that $S \neq T$ and
 389 $S, T \notin \{\emptyset, \mathcal{Y}\}$ (the inequalities being immediate in these corner cases). Let

$$\epsilon := \min_{S, T} f(S) + f(T) - f(S \cup T) - f(S \cap T)$$

390 By the strict submodularity hypothesis, we know $\epsilon > 0$.

391 Let $S, T \subseteq \mathcal{Y}$ such that $S \neq T$ and $S, T \neq \emptyset, \mathcal{Y}$. To show the log-submodularity of g , it suffices to
 392 show that

$$g(S) + g(T) \geq g(S \cup T) + g(S \cap T).$$

393 By definition of ϵ ,

$$f(S) + f(T) - f(S \cup T) - f(S \cap T) \geq \epsilon$$

394 From equation 4, we know that

$$\max_{S \subseteq \mathcal{Y}} |f(S) - g(S)| \leq \epsilon/4.$$

395 It follows that

$$\begin{aligned} g(S) + g(T) - g(S \cup T) + g(S \cap T) &\geq f(S) + f(T) - f(S \cup T) - f(S \cap T) - \epsilon \\ &\geq 0 \end{aligned}$$

396 which proves the submodularity of g . □

397 **B Encoder details**

398 For the MNIST encodings, the VAE encoder consists of a 2d-convolutional layer with 64 filters of
 399 height and width 4 and strides of 2, followed by a 2d convolution layer with 128 filters (same height,
 width and strides), then by a dense layer of 1024 neurons. The encodings are of length 32.

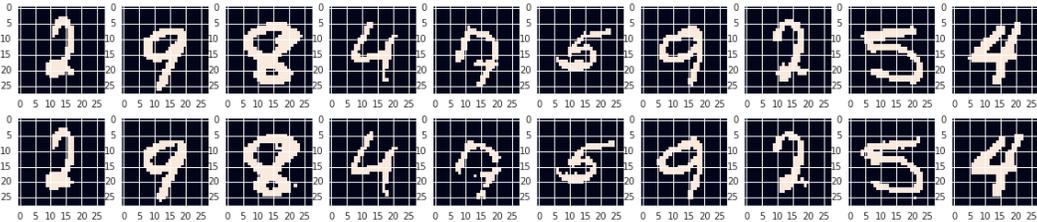


Figure 6: Digits and VAE reconstructions from the MNIST training set

400

401 CelebA encodings were generated by a VAE using a Wide Residual Network [47] encoder with 10
 402 layers and filter-multiplier $k = 4$, a latent space of 32 full-covariance Gaussians, and a deconvolutional
 403 decoder trained end-to-end using an ELBO loss. In detail, the decoder architecture consists of a 16K
 404 dense layer followed by a sequence of 4×4 convolutions with [512, 256, 128, 64] filters interleaved
 405 with $2 \times$ upsampling layers and a final 6×6 convolution with 3 output channels for each of 5
 406 components in a mixture of quantized logistic distributions representing the decoded image.