Appendix for *A Multi-Resolution Framework for U-Nets with Applications to Hierarchical VAEs*

A Framework Details and Technical Proofs

Here we provide proofs for the theorems in the main paper and additional theoretical results supporting these.

A.1 Definitions and Notations

The following provides an index of commonly used notation throughout this manuscript for reference.

The *function space* of interest in this work is $L^2(\mathbb{X})$, the space of square integrable functions, where $\mathbb{X}$ is a compact subset of $\mathbb{R}^m$ for some integer m, for instance, $\mathbb{X} = [0, 1]$. This set of functions is defined as

$$L^2(\mathbb{X}) = \{f : \mathbb{X} \to \mathbb{R} \,|\, \|f\|_2 < \infty,\ f \text{ Borel measurable}\}. \tag{A.1}$$

$L^2(\mathbb{X})$ forms a vector space with the standard operations.

We denote $V_{-j} \subset L^2(\mathbb{X})$ as a finite-dimensional approximation space. With the nesting property, $V_{-j+1} \subset V_{-j}$, the space U_{-j+1} is the orthogonal compliment of V_{-j+1} within V_{-j}, i.e. $V_{-j} = U_{-j+1} \oplus V_{-j+1}$.

The *integration* shorthand notations used are as follows. For an integrable function $t \mapsto f(t)$, we use

$$f(t)dt := \int_0^t f(s)ds. \tag{A.2}$$

The function f may be multi-dimensional in which case we mean the multi-dimensional integral in whichever basis is being used. For *stochastic* integrals, we only analyse dynamics within the truncation V_{-J} of $L^2(\mathbb{X})$. In this case, W_t refers to a Brownian motion on the same amount of dimensions as V_{-J} in the *standard*, or 'pixel', basis of V_{-J}. The shorthand

$$g(W_t)dW_t := \int_0^t g(W_s)dW_s, \tag{A.3}$$

is used for the standard Itô integral. Last, for a stochastic process X_t on V_{-J} we use

$$dX_t := X_t - X_0, \tag{A.4}$$

if not specified otherwise.

For *measures*, we use $\mathbb{D}$ to prefix a set for which we consider the space of probability measures over: for instance, $\mathbb{D}(\mathbb{X})$ denotes the space of probability measures over $\mathbb{X}$. We often refer to measures over functions (i.e. images): recall that V_{-J} is an L^2-function space and we take $\mathbb{D}(V_{-J})$ to be probability measures over this space.

When referenced in Definition 2, the distance metric between two measures ν_1 and ν_2 which yields the topology of *weak continuity* is the *Monge-Kantorovich* metric [50, 51]

$$d_{\mathbb{P}}(\nu_1, \nu_2) = \sup_{f \in \text{Lip}_1(\mathbb{X})} \int f d(\nu_1 - \nu_2), \tag{A.5}$$

where

$$\text{Lip}_1(\mathbb{X}) = \{f : \mathbb{X} \to \mathbb{R} \,\big|\, |f(x) - f(y)| \leq d(x, y), \forall x, y \in \mathbb{X}\}. \tag{A.6}$$

Further, we use the Wasserstein-2 metric which in comparison to the weak convergence above has additional moment assumptions. It is given by

$$\mathcal{W}_2(\nu_1, \nu_2) = \left(\inf_{\gamma \in \Gamma(\nu_1, \nu_2)} \mathbb{E}\|X_1 - X_2\|_2^2 \right)^{1/2}. \tag{A.7}$$

where $(X_1, X_2) \sim \gamma$ and $\Gamma(\nu_1, \nu_2)$ is the space of measures on $\mathbb{D}(V_{-J} \times V_{-J})$ with marginals ν_1 and ν_2.

A.2 Dimension Reduction Conjugacy

Assume momentarily the one dimensional case where $\mathbb{X} = [0, 1]$. Let V_{-j} be an *approximation space* contained in $L^2(\mathbb{X})$ (see Definition 1) pertaining to image pixel values

$$V_{-j} = \{f \in L^2([0,1]) \mid f_{[2^{-j}\cdot k, 2^{-j}\cdot(k+1))} = c_k,\ k \in \{0, \dots, 2^j - 1\},\ c_k \in \mathbb{R}\}. \tag{A.8}$$

For a function $f \in V_{-j}$, there are several ways to express f in different bases. Consider the *standard (or 'pixel') basis* for a fixed V_{-j} given via

$$e_{j,k} = \mathbb{1}_{[2^{-j}\cdot k, 2^{-j}\cdot(k+1)]}. \tag{A.9}$$

Clearly, the family $\mathbf{E}_j \coloneqq \{e_{j,k}\}_{k=0}^{2^j-1}$ is an orthogonal basis of V_{-j}, hence full rank with dimension 2^j. Functions in V_{-j} may be expressed as

$$f = \sum_{k=0}^{2^j-1} c_k \cdot e_{j,k}, \tag{A.10}$$

for $c_k \in \mathbb{R}$.

First, let us recall the average *pooling* operation in these bases $\mathbf{E}_j$ and $\mathbf{E}_{j-1}$ of V_{-j} and V_{-j+1}, where $\text{pool}_{-j,-j+1} : V_{-j} \to V_{-j+1}$. Its operation is given by

$$\text{pool}_{-j,-j+1}(f) = \text{pool}_{-j,-j+1}\left(\sum_{k=0}^{2^j-1} c_k \cdot e_{j,k}\right) = \sum_{i=0}^{2^{j-1}-1} \tilde{c}_i \cdot e_{j-1,i}, \tag{A.11}$$

where for $i \in \{0, \dots, 2^{j-1} - 1\}$ we have the coefficient relation

$$\tilde{c}_i = \frac{c_{2i} + c_{2i+1}}{2} = \frac{1}{2^{-j}} \int_{[2^{-j}\cdot(2i), 2^{-j}\cdot(2i+1))} f(x)dx. \tag{A.12}$$

Average pooling and its imposed basis representation are commonly used in U-Net architectures [1], for instance in state-of-the-art diffusion models [13] and HVAEs [9].

Note that across approximation spaces of two resolutions V_{-j} and V_{-j+1}, the standard bases $\mathbf{E}_j$ and $\mathbf{E}_{j-1}$ share no basis elements. As basis elements change at each resolution, it is difficult to analyse V_{-j} embedded in V_{-j+1}. What we seek is a basis for all V_{-j} such that any basis element in this set at resolution j is also a basis element in V_{-J}, the approximation space of highest resolution J we consider. This is where wavelets serve their purpose: We consider a *multi-resolution (or 'wavelet') basis* of V_{-J} [52]. For the purpose of our theoretical results below, we are here focusing on a *Haar wavelet* basis [32] which we introduce in the following, but note that our framework straight-forwardly generalises to other wavelet bases. Begin with $\phi_1 = \mathbb{1}_{[0,1)}$ as L^2-basis element for V_{-1}, the space of constant functions on $[0, 1)$. For V_{-2} we have the space of L^2 functions which are constant on $[0, 1/2)$ and $[1/2, 1)$, which we receive by adding the basis element $\psi = \sqrt{2}(\mathbb{1}_{[0,1/2)} - \mathbb{1}_{[1/2,1)})$. Here ϕ_1 is known as the father wavelet, and ψ as the mother wavelet. To make a basis for general V_{-j} we localise these two wavelets with scaling and translation, i.e

$$\psi_{i,k} = 2^{-i/2} \cdot \psi(2^i(\,\cdot - k)) \quad \text{where } i \in \{0, j\}, k \in \{0, 2^{-i+1}\}. \tag{A.13}$$

It is straight-forward to check that $\mathbf{\Psi}_j \coloneqq \{\psi_{i,k}\}_{i=0,k=0}^{j,2^{i-1}}$ is an orthonormal basis of V_{-j} on $[0, 1]$. Further, the truncated basis $\mathbf{\Psi}_{j-1}$, which is a basis for V_{-j+1}, is contained in the basis $\mathbf{\Psi}_j$. This is in contrast to $\mathbf{E}_{j-1}$ which has basis elements distinct from the elements in the basis $\mathbf{E}_j$ on a higher resolution.

The collections $\mathbf{E}_j$ and $\mathbf{\Psi}_j$ both constitute full-rank bases for V_{-j}. They further have the same dimension and so there is a linear isomorphism $\pi_j : V_{-j} \to V_{-j}$ for change of basis, i.e.

$$\pi_j(e_{j,i}) = \psi_{j,i}. \tag{A.14}$$

This can be normalised to be an isometry. We now analyse the pooling operation in our basis $\mathbf{\Psi}_j$, restating Theorem 2 from the main text and providing a proof.

Theorem 2. Given V_{-j} as in Definition 1, let $x \in V_{-j}$ be represented in the standard basis $\mathbf{E}_j$ and Haar basis $\mathbf{\Psi}_j$. Let $\pi_j : \mathbf{E}_j \mapsto \mathbf{\Psi}_j$ be the change of basis map illustrated in Fig. 3, then we have the conjugacy $\pi_{j-1} \circ \text{pool}_{-j,-j+1} = \text{proj}_{V_{-j+1}} \circ \pi_j$.

Proof. Define the conjugate pooling map in the wavelet basis, $\text{pool}^*_{j,j+1} : V_{-j} \to V_{-j+1}$ computed on the bases $\mathbf{\Psi}_j$ and $\mathbf{\Psi}_{j-1}$,

$$\text{pool}^*_{-j,-j+1} \coloneqq \pi_{j-1} \circ \text{pool}_{-j,-j+1} \circ \pi_j^{-1}. \tag{A.15}$$

$$\begin{array}{ccc} (V_{-j}, \mathbf{E}_j) & \xrightarrow{\text{pool}_{-j,-j+1}} & (V_{-j+1}, \mathbf{E}_{j-1}) \\ \uparrow{\scriptstyle \pi_j^{-1}} & & {\scriptstyle \pi_{j-1}}\downarrow \\ (V_{-j}, \mathbf{\Psi}_j) & \xrightarrow[\text{pool}^*_{-j,-j+1}]{} & (V_{-j+1}, \mathbf{\Psi}_{j-1}) \end{array}$$

Due to the scaling and translation construction in Eq. (A.13) and because the pooling operation is local, we need only consider the case for $\text{pool}_{-2,-1}$. This is because one can view pooling between the higher-resolution spaces as multiple localised pooling operations between V_{-2} and V_{-1}. Now note that $\text{pool}_{-2,-1}$ maps V_{-2} to V_{-1}. Further,

$$\int_{\mathbb{X}} \psi(x)dx = 0, \tag{A.16}$$

where $\psi = \sqrt{2}(\mathbb{1}_{[0,1/2)} - \mathbb{1}_{[1/2,1)})$ is the mother wavelet. For $v \in V_{-2}$ let v have the wavelet representation $v = \tilde{c}_2\psi + \tilde{c}_1\phi_1$, where $\phi_1 = \mathbb{1}_{[0,1)}$ is the father wavelet. To pool we compute the average of the two coefficients ('pixel values')

$$\text{pool}_{-2,-1}(v) = \int_{\mathbb{X}} v(x)dx = \int_{\mathbb{X}} \tilde{c}_2\psi(x) + \tilde{c}_1\phi_1(x)dx = \tilde{c}_1. \tag{A.17}$$

Thus average pooling here corresponds to truncation of the wavelet basis for V_{-2} to the wavelet basis for V_{-1}. As this basis is orthonormal over $L^2(\mathbb{X})$, truncation corresponds to L^2 projection, i.e. $\text{pool}^*_{-j,-j+1} = \text{proj}_{V_{-j+1}}$, as claimed. □

Theorem 2 shows that the pooling operation is conjugate to projection in the Haar wavelet approximation space, and computed by truncation in the Haar wavelet basis. The only quantity we needed for our basis over the V_{-j} was the vanishing moment quantity

$$\int_{\mathbb{X}} \psi(x)dx = 0. \tag{A.18}$$

To extend this property to higher dimensions, such as the two dimensions of gray-scale images, we use the tensor product of $[0, 1]$, and further, the tensor product of basis functions. This property is preserved, and hence the associated average pooling operation is preserved on the tensor product wavelet basis, too. To further extend it to color images, one may consider the cartesian product of several L2 spaces.

A.3 Average pooling Truncation Error

In this section we prove Theorem 3, which quantifies the regularisation imposed by an average pooling bottleneck trained by minimising the reconstruction error. The proof structure is as follows: First we give an intuition for autoencoders with an average pooling bottleneck, then derive the relevant assumptions for Theorem 3. We next prove our result under strong assumptions. Last, we weaken our assumptions so that our theorem is relevant to HVAE architectures.

Suppose we train an autoencoder on V_{-j} without dimension reduction, calling the parameterised forward (or encoder/bottom-up) and backward (or decoder/top-down) passes $F_{j,\phi}$, $B_{j,\theta} : V_{-j} \mapsto V_{-j}$ respectively. We can optimise $F_{j,\phi}$ and $B_{j,\theta}$ w.r.t. ϕ and θ to find a perfect reconstruction, i.e. $x = B_{j,\theta}F_{j,\phi}x$ for all x in our data as there is no bottleneck (no dimensionality reduction): $B_{j,\theta}$ need only be a left inverse of $F_{j,\phi}$, as in

$$B_{j,\theta}F_{j,\phi} = I. \tag{A.19}$$

Importantly, we can choose $F_{j,\phi}$ and $B_{j,\theta}$ satisfying A.19 *independent* of our data. For instance, they could both be the identity operator and achieve perfect reconstruction, but contain no information about the generative characteristics of our data. Compare this to a *bottleneck* with average pooling, i.e. an autoencoder with dimension reduction. Here, we consider the dimension reduction from V_{-j}

to V_{-j+1}, where we split $V_{-j} = V_{-j+1} \oplus U_{-j+1}$. As we have seen in Theorem 2, through average pooling, we keep information in V_{-j+1}, and discard the information in U_{-j+1}. For simplicity, let $\text{embd}_{V_{-j}}$ be the inclusion of the projection $\text{proj}_{V_{-j+1}}$. Now to achieve perfect reconstruction

$$x = (B_{j,\theta} \circ \text{embd}_{V_{-j}} \circ \text{proj}_{V_{-j+1}} \circ F_{j,\phi})x, \tag{A.20}$$

we require $(\text{proj}_{U_{-j+1}} F_{j,\phi})x = 0$. Simply put, the encoder $F_{j,\phi}$ should make sure that the discarded information in the bottleneck is nullified.

We may marry this observation with a simple U-Net structure (without skip connection) with L^2-reconstruction and average pooling dimension reduction. Let V_{-j} be one of our multi-resolution approximation spaces and $\mathbb{D}(V_{-j})$ be the space of probability measures over V_{-j}. Recall in a multi-resolution basis we have $V_{-j} = V_{-j+1} \oplus U_{-j+1}$ where U_{-j+1} is the $-j+1$ orthogonal compliment within V_{-j}. For any $v \in V_{-j}$ we may write $v = \text{proj}_{V_{-j+1}} v \oplus \text{proj}_{U_{-j+1}} v$ and analyse the truncation error in V_{-j+1}, i.e. the discarded information, via

$$\|v - \text{embd}_{V_{-j}} \circ \text{proj}_{V_{-j+1}} v\|_2^2 = \|\text{proj}_{U_{-j+1}} v\|_2^2. \tag{A.21}$$

If we normalise this value to

$$\frac{\|v - \text{embd}_{V_{-j}} \circ \text{proj}_{V_{-j+1}} v\|_2^2}{\|v\|_2^2} = \frac{\|\text{proj}_{U_{-j+1}} v\|_2^2}{\|v\|_2^2} \in [0, 1], \tag{A.22}$$

then this is zero when v is non-zero only within V_{-j+1} and zero everywhere within U_{-j+1}. Suppose now that we have a measure $\nu_j \in \mathbb{D}(V_{-j})$, we could quantify *how much* of the norm of a sample from ν_j comes from the U_{-j+1} components by computing

$$\mathbb{E}_{v \sim \nu_j} \frac{\|\text{proj}_{U_{-j+1}} v\|_2^2}{\|v\|_2^2} = \int \frac{\|\text{proj}_{U_{-j+1}} v\|_2^2}{\|v\|_2^2} d\nu_j(v) \in [0, 1]. \tag{A.23}$$

This value forms a convex sum with its complement projection to $\text{proj}_{V_{-j+1}}$, demonstrating the splitting of mass across V_{-j+1} and U_{-j+1}, as we show in Lemma 1.

Lemma 1. Let $\nu_j \in \mathbb{D}(V_{-j})$ be atom-less at 0, then

$$\mathbb{E}_{v \sim \nu_j} \frac{\|\text{proj}_{V_{-j+1}} v\|_2^2}{\|v\|_2^2} + \mathbb{E}_{v \sim \nu_j} \frac{\|\text{proj}_{U_{-j+1}} v\|_2^2}{\|v\|_2^2} = 1. \tag{A.24}$$

Proof. For any $v \in V_{-j}$ we have $\|v\|_2^2 = \|\text{proj}_{V_{-j+1}} v\|_2^2 + \|\text{proj}_{U_{-j+1}} v\|_2^2$ due to orthogonality of V_{-j+1} and U_{-j+1}. As both $\|\text{proj}_{V_{-j+1}} v\|_2^2$ and $\|\text{proj}_{U_{-j+1}} v\|_2^2$ are projections, they are bounded by $\|v\|_2^2$ giving that the integrands in Eq. (A.24) are bounded by one, and so for all $v \neq 0$ (no point mass at 0) the expectation is bounded. □

From the splitting behaviour of masses in the L^2-norm observed in Lemma 1 we see that

1. if $\mathbb{E}_{v \sim \nu_j} \|\text{proj}_{U_{-j+1}} v\|_2^2 / \|v\|_2^2$ is large, then, on average, samples from ν_j have most of their size in the U_{-j+1} subspace; or,
2. if $\mathbb{E}_{v \sim \nu_j} \|\text{proj}_{U_{-j+1}} v\|_2^2 / \|v\|_2^2$ is small, then, on average, samples from ν_j have most of their size in the V_{-j+1} subspace.

In the latter case, $\|\text{proj}_{U_{-j+1}} v\|_2^2 \approx 0$, i.e. $\text{embd}_{V_{-j}} \circ \text{proj}_{V_{-j+1}} v \approx v$. We get the heuristic $\text{embd}_{V_{-j}} \circ \text{proj}_{V_{-j+1}} \approx I$ on the measure ν_j, yielding a perfect reconstruction.

Let $\text{embd}_{V_{-j}} \circ \text{proj}_{V_{-j+1}}, I : V_{-j} \to V_{-j}$, then this heuristic performs the operator approximation

$$\mathbb{E}_{v \sim \nu_j} \|(\text{embd}_{V_{-j}} \circ \text{proj}_{V_{-j+1}} - I)v\|_2^2, \tag{A.25}$$

quantifying 'how close' these operators are on ν_j. For many measures, this (near) equivalence between operators will not hold. But what if instead, we had an operator $D : V_{-j} \to V_{-j}$ such that the push-forward of ν_j through this operator had this quality. Practically, this push-forward operator will be parameterised by neural networks, for instance later in the context of U-Nets. For simplicity, we will initially consider the case where D is linear on V_{-j}, then we consider when D is Lipschitz.

Lemma 2. Given V_{-j} with the L^2-orthogonal decomposition $V_{-j} = V_{-j+1} \oplus U_{-j+1}$, let $D_{-j} : V_{-j} \to V_{-j}$ be an invertible linear operator and define $F_j : V_{-j} \to V_{-j+1}$ and $B_j : V_{-j+1} \to V_{-j}$ through

$$F_j = \text{proj}_{V_{-j+1}} \circ D_j, \qquad B_j = D_j^{-1} \circ \text{embd}_{V_{-j}}. \tag{A.26}$$

Then $B_j F_j \equiv I$ on V_{-j}, or otherwise, we have the truncation bound

$$\frac{\left\|\text{proj}_{U_{-j+1}} F_j v\right\|_2^2}{\|D_j\|_2^2} \leq \frac{\left\|\text{proj}_{U_{-j+1}} F_j v\right\|_2^2}{\|F_j\|_2^2} \leq \|(I - B_j F_j)v\|_2^2 . \tag{A.27}$$

Proof. Consider the operator $D_j(I - B_j F_j) : V_{-j} \to V_{-j}$ which is linear and obeys the multiplicative bound $\|D_j(I - B_j F_j)\| \leq \|D_j\| \, \|I - B_j F_j\|$. This implies for any $v \in V_{-j}$,

$$\frac{\|D_j(I - B_j F_j)v\|_2^2}{\|D_j\|_2^2} \leq \|(I - B_j F_j)v\|_2^2 . \tag{A.28}$$

The numerator is equal to

$$\|D_j(I - B_j F_j)v\|_2^2 = \left\|(D_j - \text{embd}_{V_{-j}} \circ \text{proj}_{V_{-j+1}} \circ D_j)v\right\|_2^2 . \tag{A.29}$$

As we have the orthogonal decomposition $V_{-j} = V_{-j+1} \oplus U_{-j+1}$, we know

$$\begin{aligned} I &= \text{proj}_{V_{-j+1}} \oplus \text{proj}_{U_{-j+1}} && \text{(A.30)} \\ &= \text{embd}_{V_{-j}} \circ \text{proj}_{V_{-j+1}} + \text{embd}_{V_{-j}} \circ \text{proj}_{U_{-j+1}}, && \text{(A.31)} \end{aligned}$$

and as $\left\|\text{embd}_{V_{-j}}\right\|_2 = 1$, we get

$$\left\|(I - \text{embd}_{V_{-j}} \circ \text{proj}_{V_{-j+1}} \circ D_j)v\right\|_2^2 = \left\|\text{proj}_{U_{-j+1}} \circ D_j v\right\|_2^2 . \tag{A.32}$$

So now as $\|D_j(I - B_j F_j)v\|_2^2 = \left\|\text{proj}_{U_{-j+1}} \circ D_j v\right\|_2^2$, we may use $\|D_j(I - B_j F_j)v\|_2^2 \leq \|D_j\|^2 \, \|I - B_j F_j v\|_2^2$ to get the desired result. □

The quantity $\left\|\text{proj}_{U_{-j+1}} F_j v\right\|_2^2 / \|F_j\|_2^2$ is analogous to the in Lemma 1 discussed quantity $\|\text{proj}_{U_{-j+1}} v\|_2^2 / \|v\|_2^2$, but we now have a 'free parameter', the operator D_j.

Next, suppose D_j is trainable with parameters θ. We do so by minimising the reconstruction cost

$$\mathbb{E}_{v \sim \nu_j} \|(I - B_j F_j)v\|_2^2 , \tag{A.33}$$

which upper-bounds our 'closeness metric' in Lemma 2.

In the linear case (D_j is linear), to ensure that $D_{j,\theta}$ is invertible we may parameterise it by an (unnormalised) LU-decomposition of the identity

$$I = D_{j,\theta}^{-1} D_{j,\theta} = L_{j,\theta} U_{j,\theta}, \tag{A.34}$$

where the diagonal entries of $L_{j,\theta}$ and $U_{j,\theta}$ are necessarily inverses of one-another. This is a natural parameterisation when considering a U-Net with dimensionality reduction. Building from Lemma 2, we can now consider the stacked U-Net (without skip connections), i.e. a U-Net with multiple downsampling/upsampling and forward/backward operators stacked on top of each other, in the linear setting. In Proposition 1, we show that this LU-parameterisation forces the pivots of $U_{j,\theta}$ to tend toward zero.

Proposition 1. Let $\{V_{-j}\}_{j=0}^J$ be a multi-resolution hierarchy of V_{-J} with the orthogonal decompositions $V_{-j} = V_{-j+1} \oplus U_{-j+1}$ and $F_{j,\phi}$, $B_{j,\theta} : V_{-j} \to V_{-j}$ be bounded linear operators such that $B_{j,\theta} F_{j,\phi} = I$. Define $\boldsymbol{F}_{j,\phi} : V_{-j} \to V_{-j+1}$ and $\boldsymbol{B}_{j,\theta} : V_{-j+1} \to V_{-j}$ by

$$\boldsymbol{F}_{j,\phi} \coloneqq \text{proj}_{V_{-j+1}} \circ F_{j,\phi}, \qquad \boldsymbol{B}_{j,\theta} \coloneqq B_{j,\theta} \circ \text{embd}_{V_{-j}}, \tag{A.35}$$

with compositions

$$\boldsymbol{F}_{j_1|j_2,\phi} \coloneqq \boldsymbol{F}_{j_1,\phi} \circ \cdots \circ \boldsymbol{F}_{j_2,\phi}, \qquad \boldsymbol{B}_{j_1|j_2,\phi} \coloneqq \boldsymbol{B}_{j_1,\phi} \circ \cdots \circ \boldsymbol{B}_{j_2,\phi}. \tag{A.36}$$

Then

$$\sum_{j=1}^{J} \frac{\left\|\text{proj}_{U_{-j+1}} F_j v\right\|_2^2}{\|F_j\|_2^2} \leq \left\|(I - \boldsymbol{B}_{1|J,\theta}\boldsymbol{F}_{1|J,\phi})v\right\|_2^2. \tag{A.37}$$

Proof. The operator $\mathbf{F}_{1|J}$ is linear, and decomposes into a block operator form with pivots $\mathbf{F}_{j|J}$ for each $j \in \{1, \ldots, J\}$. Each $\mathbf{F}_{j|J}$ is L^2-operator norm bounded by $\|F_j\|_2$, so if

$$\lambda_{1|J} \coloneqq \text{diag}(\|F_1\|_2, \ldots, \|F_J\|_2), \tag{A.38}$$

then $\|\lambda_{1|J}^{-1}\mathbf{F}_{1|J}\|_2 \leq 1$. Last, as the spaces $\{U_{-j}\}_{j=0}^{J}$ are orthogonal and $\mathbf{F}_{1|J}$ has triangular form:

$$\|\lambda_{1|J}^{-1}(F_{1|J} - \mathbf{F}_{1|J})v\|_2^2 = \sum_{j=1}^{J} \frac{\left\|\text{proj}_{U_{-j+1}} F_j v\right\|_2^2}{\|F_j\|_2^2}, \tag{A.39}$$

and $\|\lambda_{1|J}^{-1}(F_{1|J} - \mathbf{F}_{1|J})v\|_2^2 \leq \|(I - \mathbf{B}_{1|J}\mathbf{F}_{1|J})v\|_2^2$. □

Here in the linear case, a U-Net's encoder is a triangular matrix where the basis vectors are the Haar wavelets. Proposition 1 states that the pivots of this matrix are minimised. Adversely, this diminishes the rank of the autoencoder and pushes our original underdetermined problem to a singular one. In other words, the U-Net is in this case demanding to approximate the identity (via an LU-like-decomposition), a linear operator, with an operator of diminishing rank.

Proposition 2. Let $\mathbb{D}(\mathbb{X})$ be the space of probability measures over $\mathbb{X}$, and assume for $\overline{F}_j, \overline{B}_j : \mathbb{D}(\mathbb{X}) \to \mathbb{D}(\mathbb{X})$ that these are inverses of one-another and $\overline{F}_j$ is Lipschitz, that is

$$\overline{F}_j\overline{B}_j = I, \qquad \mathcal{W}_2(\overline{F}_j\nu_1, \overline{F}_j\nu_2) \leq \|\overline{F}_j\|_2 W_2(\nu_1, \nu_2). \tag{A.40}$$

Then for any $\nu \in \mathbb{D}(\mathbb{X})$ with bounded second moment,

$$\mathbb{E}_{X_j \sim \overline{F}_j\nu} \frac{\|\text{proj}_{U_{-j}} X_j\|_2^2}{\|\overline{F}_j\|_2^2} \leq \mathcal{W}_2(\nu, \overline{B}_j \circ P_{V_{-j}} \circ \overline{F}_j\nu). \tag{A.41}$$

Proof. First as $\overline{F}_j\overline{B}_j = I$ we know that

$$\mathcal{W}_2(\overline{F}_j\nu, P_{V_{-j}}\overline{F}_j\nu) = \mathcal{W}_2(\overline{F}_j\overline{B}_j\overline{F}_j\nu, \overline{F}_j\overline{B}_j P_{V_{-j}}\overline{F}_j\nu). \tag{A.42}$$

But for any $X \in V_{-j}$ we have the orthogonal decomposition

$$X = \text{proj}_{V_{-j}} X \oplus \text{proj}_{U_{-j}} X, \tag{A.43}$$

which respects the L^2-norm by

$$\|X\|_2^2 = \|\text{proj}_{V_{-j}} X\|_2^2 + \|\text{proj}_{U_{-j}} X\|^2, \tag{A.44}$$

and in particular,

$$\|X - \text{proj}_{V_{-j}} X\|_2^2 = \|\text{proj}_{U_{-j}} X\|^2. \tag{A.45}$$

This grants

$$\begin{aligned}
(W_2(\overline{F}_j\nu, P_{V_{-j}}\overline{F}_j\nu))^2 &= \inf_{\gamma \in \Gamma(\overline{F}_j\nu, P_{V_{-j}}\overline{F}_j\nu)} \mathbb{E}\|X - Y\|_2^2 && \text{(A.46)}\\
&= \inf_{\gamma \in \Gamma(\overline{F}_j\nu, P_{V_{-j}}\overline{F}_j\nu)} \mathbb{E}\|\text{proj}_{V_{-j}} X - \text{proj}_{V_{-j}} Y\|_2^2 + \|\text{proj}_{U_{-j}} X\|_2^2 && \text{(A.47)}\\
&= \inf_{\gamma \in \Gamma(\overline{F}_j\nu, P_{V_{-j}}\overline{F}_j\nu)} \mathbb{E}\|\text{proj}_{U_{-j}} X\|_2^2 && \text{(A.48)}\\
&= (\mathcal{W}_2(\text{proj}_{U_{-j}}\overline{F}_j\nu, \delta_{\{0\}}))^2. && \text{(A.49)}
\end{aligned}$$

Now the Lipschitz of $\overline{F}_j$ yields

$$\frac{\mathcal{W}_2(\overline{F}_j\overline{B}_j\overline{F}_j\nu, \overline{F}_j\overline{B}_jP_{V_{-j}}\overline{F}_j\nu)}{\|\overline{F}_j\|_2} \leq \mathcal{W}_2(\overline{B}_j\overline{F}_j\nu, \overline{B}_jP_{V_{-j}}\overline{F}_j\nu) = \mathcal{W}_2(\nu, \overline{B}_jP_{V_{-j}}\overline{F}_j\nu). \quad \text{(A.50)}$$

Squaring and substituting grants

$$\frac{(\mathcal{W}_2(\text{proj}_{U_{-j}}\overline{F}_j\nu, \delta_{\{0\}}))^2}{\|\overline{F}_j\|_2^2} \leq \mathcal{W}_2(\nu, \overline{B}_jP_{V_{-j}}\overline{F}_j\nu). \quad \text{(A.51)}$$

□

To work on multiple resolution spaces, we need to define what the triangular operator over our space of measures is. For a cylinder set B on $V_{-J} = V_0 \oplus \bigoplus_{j=0}^{J} U_{-j}$ we can assume it has the form $\bigotimes_j B_j$ where B_j is a cylinder on U_j. Break ν_J into the multi-resolution sub-spaces by defining projection onto $\mathbb{D}(U_{-j})$ through

$$\text{proj}_{U_{-j}}(\nu_J)(B_j) \coloneqq \nu_J(B_j \otimes U_{-j}^{\perp}), \quad \text{(A.52)}$$

where B_j is a cylinder for U_{-j}. This projection of measures is respected by evaluation in that

$$\mathbb{E}_{X_j\sim\text{proj}_{U_{-j}}\nu_J}X_j = \int v_j d\text{proj}_{U_{-j}}\nu_J(v_j) = \int \text{proj}_{U_{-j}}v d\nu_J(v) = \mathbb{E}_{X_j\sim\nu_J}\text{proj}_{U_{-j}}X. \quad \text{(A.53)}$$

As $\|X\|_2^2 = \sum_j \text{proj}_{U_{-j}}\|\text{proj}_{U_{-j}}X\|_2^2$ due to the orthogonality of the spaces, then

$$\mathbb{E}_{X_j\sim\text{proj}_{U_{-j}}\nu_J}\|X_j\|_2^2 = \sum_j \mathbb{E}_{X_j\sim\text{proj}_{U_{-j}}\nu_J}\|X_j\|_2^2 = \sum_j \mathbb{E}_{X\sim\nu_J}\|\text{proj}_{U_{-j}}X\|_2^2. \quad \text{(A.54)}$$

Define the extension, with a convenient abuse of notation, of $\text{proj}_{V_{-j+1}}$ on $\mathbb{D}(V_{-J})$ to be

$$\text{proj}_{V_{-j+1}}(\nu_J) \coloneqq \text{proj}_{V_{-j+1}}(\nu_J) \otimes \text{proj}_{V_{-j+1}^{\perp}}(\nu_J). \quad \text{(A.55)}$$

If $F_{-j} : \mathbb{D}(V_{-j}) \to \mathbb{D}(V_{-j})$ are linear operators for $j \in \{0, \ldots, J\}$, extend each $F_j : \mathbb{D}(V_{-j}) \to \mathbb{D}(V_{-j})$ to $\mathbb{D}(V_{-j}) \times \mathbb{D}(V_{-j}^{\perp})$ through

$$\overline{F}_j \coloneqq F_j \oplus I. \quad \text{(A.56)}$$

For a measure $\nu_J \in \mathbb{D}(V_{-J})$ we can split it into $\mathbb{D}(V_{-j}) \times \mathbb{D}(V_{-j}^{\perp})$ via

$$\text{proj}_{V_{-j}}\nu_J \times \text{proj}_{V_{-j}^{\perp}}\nu_J, \quad \text{(A.57)}$$

which also remains a measure in $\mathbb{D}(V_{-J})$ as $\mathbb{D}(V_{-j}) \times \mathbb{D}(V_{-j}^{\perp}) \subset \mathbb{D}(V_{-J})$. Now the operator $\overline{F}_j$ acts on the product measure $\nu_j \otimes \nu_j^{\perp}$ by

$$\overline{F}_j(\nu_j \otimes \nu_j^{\perp}) = F_j\,\nu_j \otimes I\,\nu_j^{\perp}. \quad \text{(A.58)}$$

Now we may define the map $\boldsymbol{F}_j : \mathbb{D}(V_{-J}) \to \mathbb{D}(V_{-j}) \times \mathbb{D}(V_{-j}^{\perp})$ through

$$\boldsymbol{F}_j \coloneqq \overline{F}_j\text{proj}_{V_{-j}}, \quad \text{(A.59)}$$

and its compositions by

$$\boldsymbol{F}_{j_1|j_2} = \boldsymbol{F}_{j_1} \circ \cdots \circ \boldsymbol{F}_{j_2}, \quad \text{(A.60)}$$

which too is an operator on $\mathbb{D}(V_{-J})$.

Further if we have a measure ν_j on V_{-j} we can form the embedding map

$$\text{embd}_j\nu_j = \nu_j \otimes \bigotimes_{i=j}^{J} \delta_{\{0\}}, \quad \text{(A.61)}$$

which we extend to $\mathbb{D}(V_{-J})$ by a convenient abuse of notation

$$\text{proj}_j \nu_J = \text{proj}_j(\nu_J) \otimes \bigotimes_{i=j}^{J} \delta_{\{0\}}. \tag{A.62}$$

Let $B_{-j} : \mathbb{D}(V_{-j}) \to \mathbb{D}(V_{-j})$ be the linear operator which is the inverse of F_{-j}. Now if we extend $B_j : \mathbb{D}(V_{-j}) \to \mathbb{D}(V_{-j})$ to $\mathbb{D}(V_{-j}) \times \mathbb{D}(V_{-j}^{\perp})$ like before through

$$\overline{B}_j \coloneqq B_j \oplus I, \tag{A.63}$$

so the map $\overline{B}_j \text{embd}_j$ is well defined on $\mathbb{D}(V_{-J})$. Now analogously define $\boldsymbol{B}_j$ and its compositions by

$$\boldsymbol{B}_j \coloneqq \overline{B}_j \text{embd}_{V_{-j}} \qquad\qquad \boldsymbol{B}_{j_1|j_2} = \boldsymbol{B}_{j_2} \circ \cdots \circ \boldsymbol{B}_{j_1}. \tag{A.64}$$

In an analogous way, the operator $\boldsymbol{F}_{j_1|j_2}$ is 'upper triangular' and $\boldsymbol{B}_{j_1|j_2}$ is 'lower triangular'. In this way, we are again seeking a lower/upper (LU-) decomposition of the identity on $\mathbb{D}(V_{-J})$. Now we may prove Theorem 3.

Theorem 3. Let $\{V_{-j}\}_{j=0}^{J}$ be a multi-resolution hierarchy of V_{-J} where $V_{-j} = V_{-j+1} \oplus U_{-j+1}$, and further, let $F_{j,\phi}, B_{j,\theta} : \mathbb{D}(V_{-j}) \mapsto \mathbb{D}(V_{-j})$ be such that $B_{j,\theta} F_{j,\phi} = I$ with parameters ϕ and θ. Define $\boldsymbol{F}_{j_1|j_2,\phi} \coloneqq \boldsymbol{F}_{j_1,\phi} \circ \cdots \circ \boldsymbol{F}_{j_2,\phi}$ by $\boldsymbol{F}_{j,\phi} : \mathbb{D}(V_{-j}) \mapsto \mathbb{D}(V_{-j+1})$ where $\boldsymbol{F}_{j,\phi} \coloneqq \text{proj}_{V_{-j+1}} \circ F_{j,\phi}$, and analogously define $\boldsymbol{B}_{j_1|j_2,\theta}$ with $\boldsymbol{B}_{j,\theta} \coloneqq B_{j,\theta} \circ \text{embd}_{V_{-j}}$. Then, the sequence $\{\boldsymbol{B}_{1|j,\theta}(\boldsymbol{F}_{1|J,\phi}\nu_J)\}_{j=0}^{J}$ forms a discrete multi-resolution bridge between $\boldsymbol{F}_{1|J,\phi}\nu_J$ and $\boldsymbol{B}_{1|J,\theta}\boldsymbol{F}_{1|J,\phi}\nu_J$ at times $\{t_j\}_{j=1}^{J}$, and

$$\sum_{j=0}^{J} \mathbb{E}_{X_{t_J} \sim \nu_J} \left\| \textit{proj}_{U_{-j+1}} X_{t_j} \right\|_2^2 / \left\| \boldsymbol{F}_{j|J,\phi} \right\|_2^2 \leq (\mathcal{W}_2(\boldsymbol{B}_{1|J,\theta}\boldsymbol{F}_{1|J,\phi}\nu_J, \nu_J))^2, \tag{A.65}$$

where $\mathcal{W}_2$ is the Wasserstein-2 metric and $\left\| \boldsymbol{F}_{j|J,\phi} \right\|_2$ is the Lipschitz constant of $\boldsymbol{F}_{j|J,\phi}$.

Proof. All we must show is that successively chaining the projections from Proposition 2 decomposes like in Proposition 1. For $X_1, X_2 \sim \nu$, $\mathcal{W}_2(F_j F_{j+1}\nu, P_{-j+2}F_{j-1}P_{-j+1}F_j\nu)$ consider f_j, f_{j-1} as realised paths for our kernel and write $\|f_{j-1}f_j X_1 - \text{proj}_{V_{-j+2}} f_{j-1} \text{proj}_{V_{-j+1}} f_j X_2\|_2^2$

$$\begin{aligned} &= \|\text{proj}_{V_{-j+1}}(f_j f_{j+1} X_1 - \text{proj}_{V_{-j+2}} f_{j-1} \text{proj}_{V_{-j+1}} f_j X_2)\|_2^2 \\ &\quad + \|\text{proj}_{U_{-j+1}}(f_j f_{j+1} X_1 - \text{proj}_{V_{-j+2}} f_{j-1} \text{proj}_{V_{-j+1}} f_j X_2)\|_2^2 \end{aligned}$$

due to the triangular form and the orthogonality of the multi-resolution basis. Let $\nu_{-j+1} = \text{proj}_{V_{-j+1}} \nu_{-j}$, then as $\text{proj}_{V_{-j+1}}$ commutes with any term equivalent to the identity operator on V_{-j+1}, the first term becomes

$$\|f_{j-1} X_{1,t_{j+1}} - \text{proj}_{V_{-j+2}} f_{j-1} X_{2,t_{j+1}})\|_2^2, \tag{A.66}$$

where $X_{1,t_{j+1}}, X_{2,t_{j+1}} \sim \nu_{-j+1}$. When an optimal coupling is made, this term becomes $\|\text{proj}_{U_{-j+2}} X_{1,t_{j+1}}\|_2^2$. The second term has $\text{proj}_{U_{-j+1}} \text{proj}_{V_{-j+2}}$ nullified, and again commutes where appropriate making this

$$\|\text{proj}_{U_{-j+1}} X_{1,t_{j+1}}\|_2^2. \tag{A.67}$$

We may again use the triangular form to utilise the identify

$$\|\text{proj}_{U_{-j+1}} F_j\|_2^2 \leq \|F_j\|_2^2, \tag{A.68}$$

to define

$$\lambda_{j|J} \coloneqq \text{diag}(\|\text{proj}_{U_{-j+1}} F_{j|J}\|_2^2, \ldots, \|\text{proj}_{U_{-J+1}} F_{J|J}\|_2^2) \tag{A.69}$$

so that

$$\mathcal{W}_2(\lambda_{j|J}^{-1}(F_{j|J}\nu_1), \lambda_{j|J}^{-1}(F_{j|J}\nu_2)) \leq \mathcal{W}_2(\nu_1, \nu_2). \tag{A.70}$$

Piecing the decomposition and scaling together, we yield

$$\mathbb{E}_{\nu_{-j+2}}\|\text{proj}_{U_{-j+2}}X_{1,t_{j+1}}\|_2^2/\|F_{j-2|j}\|_2^2 + \mathbb{E}_{\nu_{-j+1}}\|\text{proj}_{U_{-j+1}}X_{1,t_{j+1}}\|_2^2/\|F_{j-2|j}\|_2^2 \tag{A.71}$$

$$\leq (\mathcal{W}_2(\nu, \mathbf{B}_{j-2|j}\mathbf{F}_{j-2|j}))^2. \tag{A.72}$$

Iterating over j in the fashion given yields the result. Last, measures within

$$\mathcal{U}_{\boldsymbol{BF}} \coloneqq \{\nu_J \,|\, \boldsymbol{F}_{j|J}\gamma_J = \overline{F}_{j|J} \otimes \bigotimes_{i=j}^{J} \delta_{\{0\}}\}, \tag{A.73}$$

are invariant under $\boldsymbol{B}_{J|1}\boldsymbol{F}_{1|J}$, further, $\boldsymbol{B}_{J|1}\boldsymbol{F}_{1|J}$ projects onto this set. To see this, take any measure $\nu_J \in \mathbb{D}(V_{-J})$ and apply $\boldsymbol{F}_{j|J}$. The information in $V_{-j}^{\perp}$ split by $\boldsymbol{P}_j$ is replaced by $\delta_{\{0\}}$ in the backward pass. Thus $\boldsymbol{B}_{J|1}\boldsymbol{F}_{1|J}\boldsymbol{B}_{J|1}\boldsymbol{F}_{1|J} = \boldsymbol{B}_{J|1}\boldsymbol{F}_{1|J}$. □

A.4 U-Nets in V_{-J}

Here we show how U-Nets can be seen as only computing the non-truncated components of a multi-resolution diffusion bridge on V_{-J} — the computations are performed in V_{-j} for $j < J$ at various layers. This amounts to showing the embedding presented in Theorem 1.

Theorem 1. Let $B_j : [t_j, t_{j+1}) \times \mathbb{D}(V_{-j}) \mapsto \mathbb{D}(V_{-j})$ be a linear operator (such as a diffusion transition kernel, see Appendix A) for $j < J$ with coefficients $\mu^{(j)}, \sigma^{(j)} : [t_j, t_{j+1}) \times V_{-j} \mapsto V_{-j}$, and define the natural extensions within V_{-J} in bold, i.e. $\boldsymbol{B}_j \coloneqq B_j \oplus \boldsymbol{I}_{V_{-j}^{\perp}}$. Then the operator $\boldsymbol{B} : [0, T] \times \mathbb{D}(V_{-J}) \mapsto \mathbb{D}(V_{-J})$ and the coefficients $\boldsymbol{\mu}, \boldsymbol{\sigma} : [0, T] \times V_{-J} \mapsto V_{-J}$ given by

$$\boldsymbol{B} \coloneqq \sum_{j=0}^{J} \mathbb{1}_{[t_j, t_{j+1})} \cdot \boldsymbol{B}_j, \quad \boldsymbol{\mu} \coloneqq \sum_{j=0}^{J} \mathbb{1}_{[t_j, t_{j+1})} \cdot \boldsymbol{\mu}^{(j)}, \quad \boldsymbol{\sigma} \coloneqq \sum_{j=0}^{J} \mathbb{1}_{[t_j, t_{j+1})} \cdot \boldsymbol{\sigma}^{(j)},$$

induce a multi-resolution bridge of measures from the dynamics for $t \in [0, T]$ and on the standard basis as $dX_t = \boldsymbol{\mu}_t(X_t)dt + \boldsymbol{\sigma}_t(X_t)dW_t$ (see Appendix A.4 for the details of this integration) for $X_t \in V_{-J}$, i.e. a (backward) multi-resolution diffusion process.

Proof. At time $t = 0$ we have that $\text{supp}\nu_0 \subset V_0 = \{0\}$, so $\mathbb{D}(V_0) = \delta_{\{0\}}$. For the s in the first time interval $[t_0, t_1)$ it must be the case $\nu_s = \delta_{\{0\}}$, so $\mu_s^{(j)}, \sigma_s^{(j)} = 0$ and $B_0(s) \equiv I$. The extension is thus $\boldsymbol{B}_0(s) \equiv I$ on V_{-J}. At $t = t_1$, the operator $\boldsymbol{B}_1 \equiv I$ on $V_1^{\perp}$ grants $\boldsymbol{\mu}_s, \boldsymbol{\sigma}_s = 0$ here, On V_1, $\boldsymbol{B}_1$, it is an operator with domain in V_1, granting $\text{supp}\nu_{t_1} \subset V_1$. For s within the interval (t_1, t_2) we maintain $\text{supp}\nu_s \subset V_1$, and by induction we can continue this for any $s \in [t_j, t_{j+1}) \subsetneq [0, 1]$. Let E_j be a basis of V_{-j}, then as $\boldsymbol{\mu}_s, \boldsymbol{\sigma}_s = 0$ on $V_{-j}^{\perp}$ the diffusion SDE on $[t_j, t_{j+1}) \times V_{-j}$ given in the basis E_j by

$$dX_t^{(j)} = \mu_t^{(j)}(X_t)dt + \sigma_t^{(j)}(X_t)dW_t \tag{A.74}$$

embeds into an SDE on V_{-J} with basis E_J by

$$dX_t = \boldsymbol{\mu}_t(X_t)dt + \boldsymbol{\sigma}_t(X_t)dW_t, \tag{A.75}$$

which maintains $X_t \in V_{-j}$ as $\boldsymbol{\sigma}_t \equiv 0$ on the complement. □

In practice, we will compute the sample paths made from Equation A.74, but we can in theory think of this as Equation A.75. The U-Net sequential truncation of spaces, then sequential inclusion of these spaces is what forms the multi-dimensional bridge with our sampling models.

A.5 Forward Euler Diffusion Approximations

Here we show that the backward cell structure of state-of-the-art HVAEs approximates an SDE evolution within each resolution.

Theorem 4. Let $t_J \coloneqq T \in (0, 1)$ and consider (the p_θ backward pass) $\boldsymbol{B}_{\theta,1|J} : \mathbb{D}(V_{-J}) \mapsto \mathbb{D}(V_0)$ given in multi-resolution Markov process in the standard basis:

$$dZ_t = (\overleftarrow{\mu}_{1,t}(Z_t) + \overleftarrow{\mu}_{2,t}(Z_t))dt + \overleftarrow{\sigma}_t(Z_t)dW_t, \tag{A.76}$$

where $\text{proj}_{U_{-j+1}}Z_{t_j} = 0$, $\|Z_t\|_2 > \|Z_s\|_2$ with $0 \leq s < t \leq T$ and for a measure $\nu_J \in \mathbb{D}(V_{-J})$ we have $Z_0 \sim \boldsymbol{F}_{\phi,J|1}\nu_J$. Then, VDVAEs approximates this process, and its residual cells are a type of two-step forward Euler discretisation of this Stochastic Differential Equation (SDE).

Proof. The evolution

$$dZ_t = (\overleftarrow{\mu}_{1,t}(Z_t) + \overleftarrow{\mu}_{2,t}(Z_t))dt + \overleftarrow{\sigma}_t(Z_t)dW_t, \tag{A.77}$$

subject to $Z_0 = 0, \|Z_t\|_2 > \|Z_s\|_2$ and X_T, $Z_0 \sim \boldsymbol{F}_{\phi,J|1}\nu_J$. By Theorem 3 we know $\boldsymbol{F}_{\phi,J|1}\nu_J$ enforces the form

$$\text{proj}_{V_1}\overline{F}_{\phi,J|1}\nu_J \otimes \bigotimes_{j=1}^{J-1}\delta_{\{0\}} \tag{A.78}$$

when ϕ is trained with a reconstruction loss. By Theorem 5, the full cost used imposes $\text{proj}_{V_1}\overline{F}_{\phi,J|1}\nu_J = \delta_{\{0\}}$, further, VDVAE initialises $Z_0 = \delta_{\{0\}}$. This enforces $Z_0 = 0$ as $Z_0 \sim \delta_{\{0\}}$. For the backward SDE, consider the splitting

$$dZ_t^{(1)} = \overleftarrow{\mu}_{1,t}(Z_t^{(1)})dt + \overleftarrow{\sigma}_t(Z_t^{(1)})dW_t, \qquad dZ_t^{(2)} = \overleftarrow{\mu}_{2,t}(Z_t^{(2)})dt, \tag{A.79}$$

where $dZ_t = dZ_t^{(1)} + dZ_t^{(2)}$ when $Z_t = Z_t^{(1)} = Z_t^{(2)}$. For the split SDE make the forward-Euler discretisation

$$Z_{i+1}^{(1)} = Z_i^{(1)} + \int_i^{i+1}\overleftarrow{\mu}_{1,t}(Z_t^{(1)})dt + \int_i^{i+1}\overleftarrow{\sigma}_t(Z_i^{(1)})dW_t \approx Z_i^{(1)} + \overleftarrow{\mu}_{1,i}(Z_i^{(1)}) + \overleftarrow{\sigma}_i(Z_i^{(1)})(W_1). \tag{A.80}$$

Now the second deterministic component can also be approximated with a forward-Euler discretisation

$$Z_{i+1}^{(2)} = Z_i^{(2)} + \int_i^{i+1}\overleftarrow{\mu}_{2,t}(Z_t^{(2)})dt \approx Z_i^{(2)} + \overleftarrow{\mu}_{2,i}(Z_i^{(2)}). \tag{A.81}$$

As $Z_0 = 0$, we need only show the update at a time i, so assume we have Z_i. First we update in the SDE step, so make the assignment and update

$$Z_i^{(1)} \leftarrow Z_i, \qquad Z_{i+1}^{(1)} = Z_i^{(1)} + \overleftarrow{\mu}_{i,1}(Z_i^{(1)}) + \overleftarrow{\sigma}_i(Z_i^{(1)})\Delta W_1. \tag{A.82}$$

Now assign $Z_i^{(2)} \leftarrow Z_{i+1}^{(1)}$ so we may update in the mean direction with

$$Z_{i+1}^{(2)} = Z_i^{(2)} + \overleftarrow{\mu}_{i,2}(Z_i^{(2)}), \tag{A.83}$$

with the total update $Z_{i+1} \leftarrow Z_{i+1}^{(2)}$. This gives the cell update for NVAE in Figure A.1. To help enforce the growth $\|Z_t\|_2 > \|Z_s\|_2$, VDVAE splits $Z_i^{(1)} = Z_i + Z_{i,+}$ where $Z_{i,+}$ increases the norm of the latent process Z_t. This connection and the associated update are illustrated in Figure A.1 [left]. Note here that if no residual connection through the cell was used (just the re-parameterisation trick in a VAE), then we degenerate to a standard Markovian diffusion process and yield the Euler-Maruyama VAE structure in Figure A.1 [right].

Remark 1. To simplify the stepping notation in the HVAE backward cells (Figures 4 and A.1), we use $Z_i^{(1)} = Z_i + \overleftarrow{\mu}_{1,i}(Z_i) + \overleftarrow{\sigma}_i(Z_i)(W_1)$ and $Z_i^{(2)} = Z_i^{(1)} + \overleftarrow{\mu}_{i,2}(Z_i^{(1)})$ so that the index i refers to all computations of the i^{th} backward cell.

□

A.6 Time-homogenuous model

Recall VDVAE has the continuous time analogue

$$dZ_t = (\overleftarrow{\mu}_{1,t}(Z_t) + \overleftarrow{\mu}_{2,t}(Z_t))dt + \overleftarrow{\sigma}_t(Z_t)dW_t, \tag{A.84}$$

where $Z_0 = 0$, $\|Z_t\|_2 > \|Z_s\|_2$ with $0 \leq s < t \leq T$ and for a measure $\nu_J \in \mathbb{D}(V_{-J})$. Due to Theorem 5, we know that the initial condition of VDVAE's U-Net is the point mass $\delta_{\{0\}}$. As the backwards pass flows from zero to positive valued functions, this direction is increasing and the equation is stiff with few layers. The distance progression from zero is our proxy for time, and we can use its 'position' to measure this. Thus, the coefficients $\overleftarrow{\mu}_{t,1}, \overleftarrow{\mu}_{t,2}, \overleftarrow{\sigma}_t$ need not have a time

dependence as this is already encoded in the norm of the Z_t processes. Thus, the time-homogeneous model postulated in the main text is:

$$dZ_t = (\overleftarrow{\mu}_1(Z_t) + \overleftarrow{\mu}_2(Z_t))dt + \overleftarrow{\sigma}(Z_t)dW_t, \tag{A.85}$$

$$Z_0 = 0, \ \|Z_t\|_2 > \|Z_s\|_2. \tag{A.86}$$

In practice, the loss of time dependence in the components corresponds to weight sharing the parameters across time, as explored in the experimental section. Weight sharing, or a time-homogeneous model, is common for score based diffusion models [40, 41], and due to our identification we are able to utilise this for HVAEs.

A.7 HVAE Sampling

Here we use our framework to comment on the sampling distribution imposed by the U-Net within VDVAE.

Theorem 5. Consider the SDE in Eq. (A.76), trained through the ELBO in Eq. B.101. Let $\tilde{\nu}_J$ denote the data measure and $\nu_0 = \delta_{\{0\}}$ be the initial multi-resolution bridge measure imposed by VDVAEs. If $q_{\phi,j}$ and $p_{\theta,j}$ are the densities of $B_{\phi,1|j}\boldsymbol{F}_{J|1}\tilde{\nu}_J$ and $B_{\theta,1|j}\nu_0$ respectively, then a VDVAE optimises the boundary condition $\min_{\theta,\phi} KL(q_{\phi,0,1}||q_{\phi,0}p_{\theta,1})$, where a double index indicates the joint distribution.

Proof. We need to only show two things. First, due to Theorems 3 and 4, we know that the architecture imposes

$$\text{proj}_{V_1}\overline{F}_{\phi,J|1}\nu_J \otimes \bigotimes_{j=1}^{J-1}\delta_{\{0\}}, \tag{A.87}$$

so we must analyse how $\text{proj}_{V_1}\overline{F}_{\phi,J|1}\nu_J$ is trained. Second, we use Theorem 4 to view the discretised version of the continuous problem, and identify the error in the two-step forward Euler splitting.

On the first point, VDVAE uses an ELBO reconstruction with a KL divergence between the backwards pass of the data $\overline{B}_{\phi,J|1}\overline{F}_{\phi,J|1}\tilde{\nu}_J$ (the 'q_ϕ-distribution'), and the backwards pass of the model imposed by the U-Net $\overline{B}_{\phi,J|1}\nu_0$ (the 'p_θ distribution'). As Z_0 is zero initialised, we know $\nu_0 = \delta_{\{0\}}$. We need to show the cost function used imposes this initialisation on $\overline{B}_{\phi,1|0}\overline{F}_{\phi,J|0}\tilde{\nu}_J$. Let $X_T \sim \overline{F}_{\phi,J|0}\tilde{\nu}_J$, call the distribution of this $q_{\phi,0}$. We also use $Z_1 \sim q_{\phi,1}$ for a sample from $\overline{B}_{\phi,1|0}\overline{F}_{\phi,J|0}\tilde{\nu}_J$ and $Z_1 \sim p_{\theta,1}$ for a sample from $\overline{B}_{\phi,1|0}\nu_0$. For a realisation x of X_T, VDVAE computes

$$KL(q_{\phi,1|0}(\,\cdot\,|X_T = x)||p_{\theta,1|0}(\,\cdot\,|Z_0 = 0)) = KL(q_{\phi,1|0}(\,\cdot\,|X_T = x)||p_{\theta,1}(\,\cdot\,)), \tag{A.88}$$

which in training is weighted by each datum, so the total cost in this term is

$$\int KL(q_{\phi,1|0}(\,\cdot\,|X_T = x)||p_{\theta,1}(\,\cdot\,))q_{\phi,0}(X_T = x)dx. \tag{A.89}$$

But this is equal to,

$$\int\int \log\left(\frac{q_{\phi,1|0}(Z_1 = z|X_T = x)}{p_{\theta,1}(Z_1 = z_1)}\right) q_{\phi,1|0}(Z_1 = z|X_T = x)q_{\phi,0}(X_T = x)dzdx \tag{A.90}$$

$$= \int\int \log\left(\frac{q_{\phi,0,1}(Z_1 = z, X_T = x)}{p_{\theta,1}(Z_1 = z_1)q_{\phi,0}(X_T = x)}\right) q_{\phi,0,1}(Z_1 = z, X_T = x)dzdx \tag{A.91}$$

$$= KL(q_{\phi,0,1}(Z_1, X_T)||p_{\theta,1}(Z_1)q_{\phi,0}(X_T)) = KL(q_{\phi,0,1}||p_{\theta,1}q_{\phi,0}). \tag{A.92}$$

The distribution of $p_{\theta,1}$ is Gaussian as a one time step diffusion evolution from the initial point mass $\nu_0 = \delta_{\{0\}}$.

□

Theorem 1 states that the choice of the initial latent variable in VDVAE imposes a boundary condition on the continuous SDE formulation. Further, this boundary condition is enforced into the final output X_T of the encoder within VDVAE.

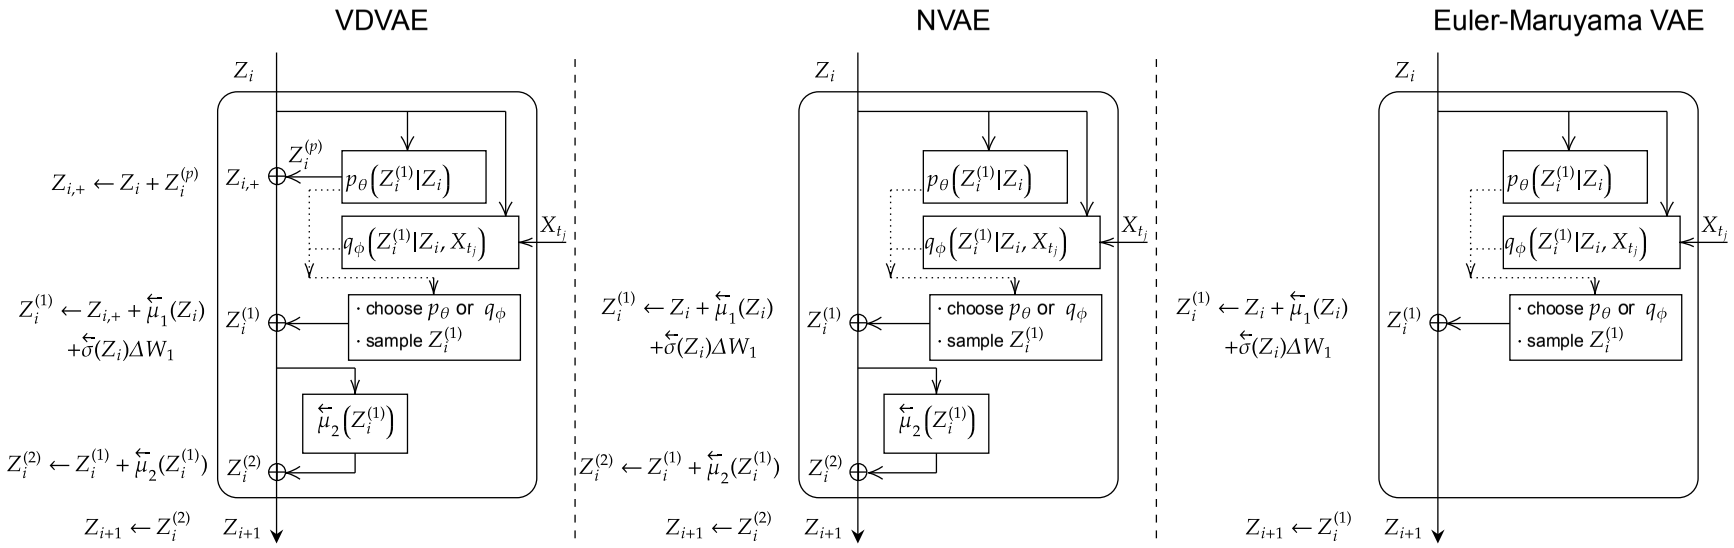

Figure A.1: HVAE top-down cells are resembling two-step forward Euler discretisations of a continuous-time diffusion process in Eq. A.76. We here provide the residual cell structures of [left] VDVAE [9], [middle] NVAE [10] and [right] a Euler-Maruyama VAE. Either q_ϕ (conditional) or p_θ (unconditional) are used in the sampling step (indicated by the dotted lines) during training and generation, respectively. X_{t_j} is an input from the in effect non-stochastic bottom-up pass, Y_i is the input from the previous, and Y_{i+1} the output to the next residual cell. $\oplus$ indicates element-wise addition.

B Background

B.1 Multi-Resolution Hierarchy and thought experiment

Let $\mathbb{X} \subset \mathbb{R}^m$ be compact and $L^2(\mathbb{X})$ be the space of square-integrable functions over this set. We are interested in decomposing $L^2(\mathbb{X})$ across multiple resolutions.

Definition 1 (abbreviated). A *multi-resolution hierarchy* is one of the form

$$\cdots \subset V_1 \subset V_0 \subset V_{-1} \subset \cdots \tag{B.93}$$

$$\overline{\bigcup_{j\in\mathbb{Z}} V_{-j}} = L^2(\mathbb{R}^m) \tag{B.94}$$

$$\bigcap_{j\in\mathbb{Z}} V_{-j} = \{0\} \tag{B.95}$$

$$f(\cdot) \in V_{-j} \iff f(2^j\,\cdot) \in V_0 \tag{B.96}$$

$$f(\cdot) \in V_0 \iff f(\,\cdot - n) \in V_0, \text{ for } n \in \mathbb{Z}. \tag{B.97}$$

Each V_{-j} is a finite truncation of $L^2(\mathbb{X})$. What we are interested in is to consider a function $f \in L^2(\mathbb{X})$ and finding a finite dimensional approximation in V_{-J}, say, for $J > 0$. Further, for gray-scale images, $\mathbb{X} = [0,1]^2$, the space of pixel-represented images. To simplify notation, we just consider $\mathbb{X} = [0,1]$ for the examples below, but we can extend this to gray-scale images, and to colour images with a Cartesian product.

The 'pixel' multi-resolution hierarchy is given by the collection of sub-spaces

$$V_{-j} = \{f \in L^2([0,1]) \mid f|_{[2^{-j}\cdot k, 2^{-j}\cdot(k+1))} = c_k,\ k \in \{0,\ldots,2^j-1\},\ c_k \in \mathbb{R}\}. \tag{B.98}$$

It can be readily checked that these sub-spaces obey the assumptions of Definition 1. An example image projected into such sub-spaces, obtained from a discrete Haar wavelet transform, is illustrated in Fig. B.2. We call it the pixel space of functions as elements of this set are piece-wise constant on dyadically split intervals of resolution 2^j, i.e a pixelated image. For each V_{-j} there is an obvious basis of size 2^j where we store the coefficients $(c_0, c_1, \ldots, c_{2^j-1}) \in \mathbb{R}^{2^j}$. The set of basis vectors for it is the *standard basis* $\{e_i\}_{i=0}^{2^j-1}$ which are 0 for all co-ordinates except for the $i^{th}+1$ entry which is 1. This basis is not natural to the multi-resolution structure of V_{-j}. This is because all the basis functions change when we project down to V_{-j+1}. We want to use the multi-resolution structure to create a basis which naturally relates V_{-j}, V_{-j+1}, and any other sub-space. To do this consider $V_{-j} \cap V_{-j+1}^{\perp} \subset V_{-j}$. Define this orthogonal compliment to be $U_{-j+1} \coloneqq V_{-j+1}^{\perp}$, then see $V_{-j} = V_{-j+1} \oplus U_{-j+1}$. Doing this recursively finds $V_{-j} = V_0 \oplus \bigoplus_{i=0}^{-j+1} U_i$, and taking the limit

$$L^2(\mathbb{X}) = \bigoplus_{i=0}^{-\infty} U_i \oplus V_0. \tag{B.99}$$

Each of the sub-spaces $\{U_{-j}\}_{j=0}^{\infty}$ are mutually orthogonal as each $V_{-j} \perp U_{-j}$. Now suppose we had a basis set $\mathbf{\Psi}_j$ for each U_{-j} and $\mathbf{\Phi}_0$ for V_0. As these spaces are orthogonal, so are the basis sets to each other, too. We can make a basis for V_{-j} with span$(\mathbf{\Phi}_0, \mathbf{\Psi}_0, \cdots, \mathbf{\Psi}_{-j+1})$. For the above examples, V_0 needs only a single basis function $\phi_{0,k} = \mathbb{1}_{[k,k+1)}$, further if $\psi = \sqrt{2}(\mathbb{1}_{[0,1/2)} - \mathbb{1}_{[1/2,1)})$, then given the functions $x \mapsto \psi_{j,k}(x) \coloneqq 2^{j/2}\cdot\psi(2^{-j}(x-k))$ we have $\{\psi_{j,k}\}$ is a basis for V_{-j}.

original (512×512)

gray-scaled original (512×512)

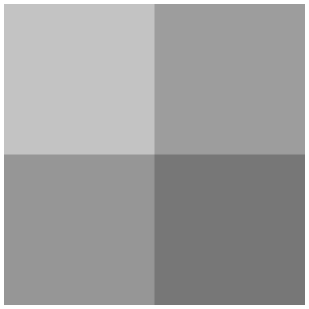

V_{-1} (2×2)

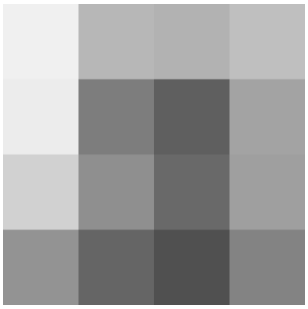

V_{-2} (4×4)

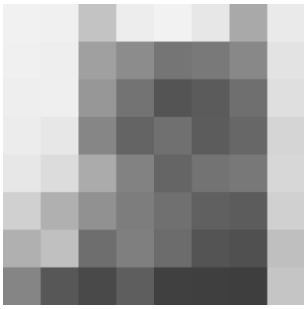

V_{-3} (8×8)

V_{-4} (16×16)

V_{-5} (32×32)

V_{-6} (64×64)

V_{-7} (128×128)

V_{-8} (256×256)

V_{-9} (512×512)

Figure B.2: The thought experiment discussed in §2. The original colour image [top-left], its gray-scale version [top-right], and its Haar wavelet projections to the approximation spaces V_{-j} for $j \in \{1, \ldots, 9\}$.

B.2 U-Net

In practice, a U-Net [1] is a neural network structure which consists of a forward pass (encoder) and backward pass (decoder), wherein layers in the forward pass have some form of dimension reduction, and layers in the backward pass have some form of dimension embedding. Furthermore, there are 'skip connection' between corresponding layers on the same level of the forward and backward pass.

We now formalise this notion, referring to an illustration of a U-Net in Fig. B.3. In black, we label the latent spaces to be V_{-j} for all j, where the original data is in V_{-J} and the U-Net 'bend' (bottleneck) occurs at V_0. We use $f_{j,\theta}$ to be the forward component, or encoder, of the U-Net, and similarly $b_{j,\theta}$ as the backward component or decoder, operating on the latent space V_{-j}. P_{-j+1} refers to the dimension reduction operation between latent space V_{-j} and V_{-j+1}, and E_{-j} refers to its corresponding dimension embedding operation between latent spaces V_{-j+1} and V_{-j}. A standard dimension reduction operation in practice is to take P_{-j+1} as average pooling, reducing the resolution of an image. Similarly, the embedding step may be some form of deterministic interpolation of the image to a higher resolution. We note that the skip connection in Fig. B.3 occur before the dimension reduction step, in this sense, lossless information is fed from the image of $f_{j,\theta}$ into the domain of $b_{j,\theta}$.

In blue, we show another backward process $b_{j,\phi}$ that is often present in U-Net architectures for generative models. This second backward process is used for unconditional sampling. In the context of HVAEs, we may refer to it as the (hierarchical) prior (and likelihood model). It is trained to match its counterpart in black, without any information from the forward process. In HVAEs, this is enforced by a KL-divergence between distributions on the latent spaces V_{-j}. The goal of either backward process is as follows:

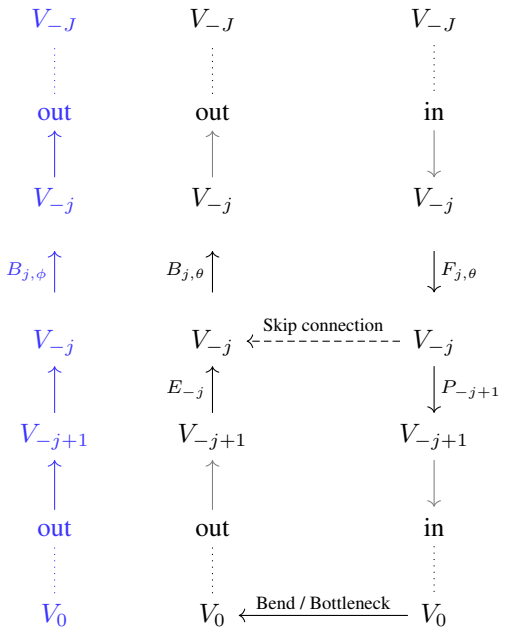

Figure B.3: The repeated structure in a U-Net, where V_{-j+1} is a lower dimensional latent space compared to V_{-j}. $f_{j,\theta}, b_{j,\theta}$ are in practice typically parameterised by neural networks (e.g. convolutional neural networks); P_{-j+1} is a dimension reduction operation (e.g. average pooling) to a lower-dimensional latent space; and, E_{-j} is a dimension embedding operation (e.g. deterministic interpolation) to a higher-dimensional latent space. This structure is repeated to achieve a desired dimension of the latent space at the U-Net bottleneck.

1. $b_{j,\theta}$ must be able to reconstruct the data from $f_{j,\theta}$, and in this sense it is reasonable to require $b_{j,\theta} f_{j,\theta} = I$;
2. $b_{j,\phi}$ must connect the data to a known sampling distribution.

The second backward process can be absent when the backward process $b_{j,\theta}$ is imposed to be the inverse of $f_{j,\theta}$, such as in Normalising Flow based models, or reversible score-based diffusion models. In this case the invertibility is assured, and the boundary condition that the encoder connects to a sampling distribution must be enforced. For the purposes of our study, we will assume that in the absence of dimension reduction, the decoder is constrained to be an inverse of the encoder. This is a reasonable assumption: for instance, in HVAEs near perfect data reconstructions are readily achieved.

For variational autoencoders, the encoder and decoder are not necessarily deterministic and involve resampling. To encapsulate this, we will work with the data as a measure and have $F_{\theta,j}$ and $B_{\theta,j}$ as the corresponding kernels imposed by $f_{j,\theta}$ and $b_{j,\theta}$, respectively.

With all of these considerations in mind, for the purposes of our framework we provide a definition of an idealised U-Net which is an approximate encapsulation of all models using a U-Net architecture.

Definition 3. (Idealised U-Net for generative modelling)
For each $j \in \{0, \dots, J\}$, let $F_{j,\theta}, B_{j,\theta} : \mathbb{D}(V_{-j}) \mapsto \mathbb{D}(V_{-j})$ such that $B_{j,\theta} F_{j,\theta} \equiv I_{V_{-j}}$. A U-Net

with (average pooling) dimension reduction P_{-j+1} and dimension embedding E_{-j} is the operator $\mathbf{U} : \mathbb{D}(V_{-J}) \mapsto \mathbb{D}(V_{-J})$ given by

$$\mathbf{U} \coloneqq B_{J,\theta}E_{-J} \circ \cdots \circ B_{1,\theta}E_{-1} \circ P_0 F_{1,\theta} \circ \cdots \circ P_{-J+1}F_{J,\theta}, \qquad B_j F_j \equiv I. \tag{B.100}$$

Remark 2. The condition $B_{j,\theta}F_{j,\theta} \equiv I_{V_{-j}}$ in our idealised U-Net (for unconditional sampling here) is either imposed directly (reversible flow based model), or approximated via skip connections. For instance, in our HVAE case, we have both a U-Net without skip connections (the p distribution) and a U-Net with skip connections (the q distribution). The U-Net related to the q distribution learns how to reconstruct our data from the reconstruction term in the ELBO cost function. The U-Net related to the p distribution learns to mimic the q distribution via the KL term in the ELBO of the HVAE, whose decoder is trained to invert its encoder — $B_{j,\phi}F_{j,\phi} \equiv I_{V_{-j}}$ — but the p U-Net lacks skip connections. Thus, in the HVAE context, we are analysing U-Nets which must simultaneously reconstruct our data and lose their reliance on their skip connections due to the condition that the q U-Net must be approximately equal to the p U-Net.

B.3 Hierarchical VAEs

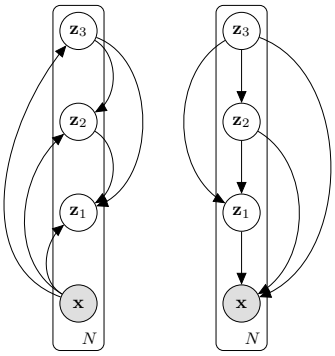

Figure B.4: Conditioning structure in state-of-the-art HVAE models (VDVAE[9] / NVAE[10]) with $L = 3$. [Left] Amortised variational posterior $q_\theta(\vec{\mathbf{z}} \mid \mathbf{x})$. [Right] Generative model $p_\phi(\mathbf{x}, \vec{\mathbf{z}})$.

A *hierarchical Variational Autoencoder (HVAE)* [4] is a VAE [48] where latent variables are separated into L groups $\vec{\mathbf{z}} = (\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_L)$ which conditionally depend on each other. L is often referred to as stochastic depth. For convenience, we refer to the observed variable $\mathbf{x}$ as $\mathbf{z}_0$, so $\mathbf{x} \equiv \mathbf{z}_0$. In HVAEs, latent variables typically follow a 'bow tie', U-Net [1] type architecture with an information bottleneck [53], so $\dim(\mathbf{z}_{l+1}) \leq \dim(\mathbf{z}_l)$ for all $l = 0, \ldots, L-1$. Latent variables live on multiple *resolutions*, either decreasing steadily [36, 54] or step-wise every few stochastic layers [10, 9] in dimension. We consider this multi-resolution property an important characteristic of HVAEs. It distinguishes HVAEs from other deep generative models, in particular vanilla diffusion models where latent and data variables are of equal dimension [13].

As in a plain VAE with only a single group of latent variables, an HVAE has a likelihood $p_\phi(\mathbf{x}|\vec{\mathbf{z}})$, a prior $p_\phi(\vec{\mathbf{z}})$ and an approximate posterior $q_\theta(\vec{\mathbf{z}}|\mathbf{x})$. To train the HVAE, one optimises the ELBO w.r.t. parameters ϕ and θ via stochastic gradient descent using the reparametrization trick

$$\log p(\mathcal{D}) \geq \mathcal{L}(\mathcal{D}; \theta, \phi) = \mathbb{E}_{\mathbf{x}\sim\mathcal{D}}\Big[\underbrace{\mathbb{E}_{\vec{\mathbf{z}}\sim q_\theta(\vec{\mathbf{z}}|\mathbf{x})}\left[\log p_\phi(\mathbf{x}|\vec{\mathbf{z}})\right]}_{\text{Reconstruction loss}} - \underbrace{\mathrm{KL}[q_\theta(\vec{\mathbf{z}}|\mathbf{x})||p_\phi(\vec{\mathbf{z}})]}_{\text{Prior loss}}\Big]. \tag{B.101}$$

Numerous conditioning structures of the latent variables in HVAEs exist, and we review them in §4. In this work, we follow [9, 10, 28]: the latent variables in the prior and approximate posterior are estimated in the same order, from $\mathbf{z}_L$ to $\mathbf{z}_1$, conditioning 'on all previous latent variables', i.e.

$$p_\phi(\vec{\mathbf{z}}) = p_\phi(\mathbf{z}_L) \prod_{l=1}^{L-1} p_\phi(\mathbf{z}_l|\mathbf{z}_{>l}) \tag{B.102}$$

$$q_\theta(\vec{\mathbf{z}}|\mathbf{x}) = q_\theta(\mathbf{z}_L|\mathbf{x}) \prod_{l=1}^{L-1} q_\theta(\mathbf{z}_l|\mathbf{z}_{>l}, \mathbf{x}) \tag{B.103}$$

We visualise the graphical model of this HVAE in Fig. B.4. Recent HVAEs [9, 10] capture this

[4] We closely follow the introduction of hierarchical VAEs in [9, §2.2].

dependence on all previous latent variables $\mathbf{z}_{>l}$ in their residual state as shown in §2.3, imposing this conditional structure. This implies a 1st-order Markov chain conditional on the previous residual state, not the previous $\mathbf{z}_l$. Such 1st-order Markov processes have shown great success empirically, such as in LSTMs [38]. Further, note that in all previous work on HVAEs, the neural networks estimating the inference and generative distributions of the l-th stochastic layer are *not* sharing parameters with those estimating other stochastic layers.

Intuitively, HVAEs' conditional structure together with a U-Net architecture imposes an inductive bias on the model to learn a *hierarchy* of latent variables where each level corresponds to a different degree of abstraction. In this work, we characterise this intuition via the regularisation property of U-Nets in §2.2.

The distributions over the latent variables in both the inference and generative model are Gaussian with mean $\boldsymbol{\mu}$ and a diagonal covariance matrix $\boldsymbol{\Sigma}$, i.e. for all $l = 1, \dots, L$,

$$q_\theta(\mathbf{z}_l|\mathbf{z}_{>l}, \mathbf{x}) \sim \mathcal{N}(\boldsymbol{\mu}_{l,\theta}, \boldsymbol{\Sigma}_{l,\theta}), \tag{B.104}$$

$$p_\phi(\mathbf{z}_l|\mathbf{z}_{>l}) \sim \mathcal{N}(\boldsymbol{\mu}_{l,\theta}, \boldsymbol{\Sigma}_{l,\theta}), \tag{B.105}$$

where mean and variances are estimated by neural networks with parameters ϕ and θ corresponding to stochastic layer l. Note that $p_\phi(\mathbf{z}_L|\mathbf{z}_{>l}) = p_\phi(\mathbf{z}_L)$, where the top-down block estimating $p_\phi(\mathbf{z}_L)$ receives the zero-vector as input, and $q_\theta(\mathbf{z}_L|\mathbf{z}_{>L}, \mathbf{x}) = q_\theta(\mathbf{z}_L|\mathbf{x})$, meaning that we infer without conditioning on other latent groups at the L-th step. Further, VDVAE chooses $p_\phi(\mathbf{x}|\vec{\mathbf{z}})$ to be a discretized Mixture-of-Logistics likelihood.

B.4 Sampling of Time Steps in HVAEs

Monte Carlo sampling of time steps in ELBO of HVAEs.

We here provide one additional theoretical result. We show that the ELBO of an HVAE can be written as an expected value over uniformly distributed time steps.

Previous work [19] [41] (Eq. (13), respectively) showed that the diffusion loss term $\mathcal{L}_T(\mathbf{x})$ in the ELBO of discrete-time diffusion models can be written as

$$\mathcal{L}_T(\mathbf{x}) = \frac{T}{2}\mathbb{E}_{\boldsymbol{\epsilon}\sim\mathcal{N}(0,\mathbf{I}),i\sim U\{1,T\}}\left[(\mathrm{SNR}(s) - \mathrm{SNR}(t))\left\|\mathbf{x} - \hat{\mathbf{x}}_\theta\left(\mathbf{z}_t; t\right)\right\|_2^2\right] \tag{B.106}$$

which allows maximizing the variational lower-bound via a Monte Carlo estimator of Eq. B.106, sampling time steps.

Inspired by this result for diffusion models, we provide a similar form of the ELBO for an HVAE with factorisation as in Eqs. (B.102)-(B.103) (and the graphical model in Fig. B.4). An HVAE's ELBO can be written as

$$\log p(\mathbf{x}) \geq \mathbb{E}_{\vec{\mathbf{z}}\sim q(\vec{\mathbf{z}}|\mathbf{x})}\left[\log p(\mathbf{x}|\vec{\mathbf{z}})\right] - L\,\mathbb{E}_{l\sim\mathrm{Unif}(1,L)}\left[\mathbb{E}_{\vec{\mathbf{z}}\sim q(\vec{\mathbf{z}}|\mathbf{x})}\log\frac{q(\mathbf{z}_l|\mathbf{z}_{>l},\mathbf{x})}{p(\mathbf{z}_l|\mathbf{z}_{>l})}\right].$$

Proof.

$$\begin{aligned}
\log p(\mathbf{x}) &\geq \mathbb{E}_{\vec{\mathbf{z}}\sim q(\vec{\mathbf{z}}|\mathbf{x})}\left[\log p(\mathbf{x}|\vec{\mathbf{z}})\right] - \mathrm{KL}\left[q(\vec{\mathbf{z}}|\mathbf{x})||p(\vec{\mathbf{z}})\right] \\
&= \mathbb{E}_{\vec{\mathbf{z}}\sim q(\vec{\mathbf{z}}|\mathbf{x})}\left[\log p(\mathbf{x}|\vec{\mathbf{z}})\right] - \int \mathrm{d}\vec{\mathbf{z}} q(\vec{\mathbf{z}}|\mathbf{x}) \log\left[\frac{\prod_{l=1}^L q(\mathbf{z}_l|\mathbf{z}_{>l},\mathbf{x})}{\prod_{l=1}^L p(\mathbf{z}_l|\mathbf{z}_{>l})}\right] \\
&= \mathbb{E}_{\vec{\mathbf{z}}\sim q(\vec{\mathbf{z}}|\mathbf{x})}\left[\log p(\mathbf{x}|\vec{\mathbf{z}})\right] - \sum_{l=1}^L \int \mathrm{d}\vec{\mathbf{z}} q(\vec{\mathbf{z}}|\mathbf{x}) \log\left[\frac{q(\mathbf{z}_l|\mathbf{z}_{>l},\mathbf{x})}{p(\mathbf{z}_l|\mathbf{z}_{>l})}\right] \\
&= \mathbb{E}_{\vec{\mathbf{z}}\sim q(\vec{\mathbf{z}}|\mathbf{x})}\left[\log p(\mathbf{x}|\vec{\mathbf{z}})\right] - L\,\mathbb{E}_{l\sim\mathrm{Unif}(1,L)}\left[\mathbb{E}_{\vec{\mathbf{z}}\sim q(\vec{\mathbf{z}}|\mathbf{x})}\log\frac{q(\mathbf{z}_l|\mathbf{z}_{>l},\mathbf{x})}{p(\mathbf{z}_l|\mathbf{z}_{>l})}\right].
\end{aligned}$$

□

This allows reducing the computational and memory costs of the KL-terms in the loss and depends on how many Monte Carlo samples are drawn. However, in contrast to diffusion models, all intermediate stochastic layers (up to the top-most and bottom-most layer chosen when sampling time steps in the recognition and generative model, respectively) still need to be computed as each latent variable's distribution depends on all previous ones.

C Code, computational resources, existing assets used

Code. We provide our PyTorch code base at https://github.com/FabianFalck/unet-vdvae. Our implementation is based on, modifies and extends the official implementation of VDVAE [9]. Below, we highlight key contributions:

- We implemented weight-sharing of individual ResNet blocks for a certain number of repetitions.
- We implemented the datasets and the preprocessing of MNIST and CelebA, which were previously not used with VDVAE.
- We implemented the option of synchronous and asynchronous processing in time (see Appendix G.4.5).
- We implemented Fourier features with hyperparameters choosing their frequencies following VDM [19]. One can concatenate them at three different locations as options.
- We simplified the multi-GPU implementation.
- We implemented an option to convert the VDVAE cell into a non-residual cell (see Appendix G.4.4).
- We implemented logging of various metrics and plots with weight&biases.
- We implemented gradient checkpointing [55] as an option in the decoder of VDVAE where the bulk of the computation occurs. We provide two implementations of gradient checkpointing, one based on the official PyTorch implementation which is unfortunately slow when using multiple GPUs, and a prototype for a custom implementation based on https://github.com/csrhddlam/pytorch-checkpoint.

The `README.md` contains instructions on installation, downloading the required datasets, the setup of weights&biases, and how to reproduce our main results.

Computational resources. For the majority of time during this project, we used two compute clusters: The first cluster is a Microsoft Azure server with two Nvidia Tesla K80 graphic cards with 11GB of GPU memory each, which we had exclusive access to. The second cluster is an internal cluster with 12 Nvidia GeForce GTX 1080 graphic cards and 10GB of GPU memory each, shared with a large number of users. In the late stages of the project, in particular to perform runs on ImageNet32, ImageNet64 and CelebA, we used a large-scale compute cluster with A100 graphic cards with 40GB of GPU memory each. We refer to the acknowledgements section for further details.

In the following, we provide a rough estimate of the total compute required to reproduce our main experiments. Compute time until convergence scales with the depth of the HVAEs. For the shallower HVAEs in our small-scale experiments in §G.1, training times range from several days to a week. For our larger-scale experiments on MNIST and CIFAR10, training times range between 1 to 3 weeks. For our deepest runs on ImageNet32 and CelebA, training times range between 2.5 to 4 weeks.

For orientation, in Table C.1, we provide an estimate of the training times of our large-scale runs in Table 1. We note that these runs have been computed on different hardware, i.e. the training times are only to some degree comparable, yet give an indication.

Existing assets used. In the experiments, our work directly builds on top of the official implementation of VDVAE [9] (MIT License). We use the datasets reported in Appendix D. In our implementation, we make use of the following existing assets and list them together with their licenses: PyTorch [56], highlighting the torchvision package for image benchmark datasets, and the gradient checkpointing implementation (custom license), Numpy [57] (BSD 3-Clause License) Weights&Biases [58] (MIT License), Apex [59] (BSD 3-Clause "New" or "Revised" License), Pickle [60] (license not available), Matplotlib [61] (PSF License), ImageIO [62] (BSD 2-Clause "Simplified" License), MPI4Py [63] (BSD 2-Clause "Simplified" License), Scikit-learn [64] (BSD 3-Clause License), and Pillow [65] (custom license).

D Datasets

In our experiments, we make use of the following datasets: MNIST [42], CIFAR10 [43], ImageNet32 [44, 45], ImageNet64 [44, 45], and CelebA [46]. We briefly discuss these datasets, focussing on

Table C.1: A large-scale study of parameter efficiency in HVAEs. For all our runs in Table 1, we report their stochastic depth and estimated training time.

	Method	Depth	Training time
MNIST (28×28)			
	`WS-VDVAE` (ours)	57	$\approx$ 5 days
	`VDVAE`* (ours)	43	$\approx$ 5 days
CIFAR10 (32×32)			
	`WS-VDVAE` (ours)	268	$\approx$ 18 days
	`WS-VDVAE` (ours)	105	$\approx$ 13 days
	`VDVAE`* (ours)	43	$\approx$ 9 days
ImageNet (32×32)			
	`WS-VDVAE` (ours)	169	$\approx$ 20 days
	`WS-VDVAE` (ours)	235	$\approx$ 24 days
	`VDVAE`* (ours)	78	$\approx$ 16 days
CelebA (64×64)			
	`WS-VDVAE` (ours)	125	$\approx$ 27 days
	`VDVAE`* (ours)	75	$\approx$ 21 days

their preprocessing, data splits, data consent and commenting on potential personally identifiable information or offensive content in the data. We refer to the training set as images used during optimisation, the validation set as images used to guide training (e.g. to compute evaluation metrics during training) but not used for optimisation directly, and the test set as images not looked at during training and only to compute performance of completed runs. For all datasets, we fix the training-validation-test split over different runs, and we scale images to be approximately centred and having a standard deviation of one based on statistics computed on the respective training set. If not stated otherwise, we use a modified version of the implementation of these datasets in [9].

MNIST. The MNIST dataset [42] contains gray-scale handwritten images of 10 digit classes ('0' to '9') with resolution 28×28. It contains $60,000$ training and $10,000$ test images, respectively. From the training images, we use $55,000$ images as the training set and 5000 images as the validation set. We use all $10,000$ test images as the testing set. We build on top of the implementation provided in NVAE [10] (`https://github.com/NVlabs/NVAE/blob/master/datasets.py`), which itself uses torchvision [56], and dynamically binarize the images, meaning that pixel values are binary, as drawn from a Bernoulli distribution with the probabilities given by the scaled gray-scale values in $[0, 1]$. Furthermore, we pad each image with zeros so to obtain the resolution 32×32.

The dataset is highly standardised and cropped to individual digits so that offensive content or personally identifiable information can be excluded. As the original NIST database from which MNIST was curated is no longer available, we cannot comment on whether consent was obtained from the subjects writing and providing these digits [27].

CIFAR10. The CIFAR10 dataset [43] contains coloured images from 10 classes ('airplane', 'automobile', 'bird', 'cat', 'deer', 'dog', 'frog', 'horse', 'ship', 'truck') with resolution 32×32. It contains $50,000$ training and $10,000$ test images, respectively. We split the training images into $45,000$ images in the training set and 5000 images in the validation set, and use all $10,000$ test images as the test set.

CIFAR10 was constructed from the so-called 80 million tiny images dataset by Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton [66]. On the official website of the 80 million tiny images dataset, the authors state that this larger dataset was officially withdrawn by the authors on June 29th, 2020 due to offensive images being identified in it [67]. The authors of the 80 million tiny images dataset do not comment on whether CIFAR10, which is a subset of this dataset, likewise contains these offensive images or is unaffected. [43] states that the images in the 80 million tiny images dataset were retrieved by searching the web for specific nouns. The authors provide no information to which degree consent was obtained from the people who own these images.

ImageNet32. The ImageNet32 dataset, a downsampled version of the ImageNet database [44, 45], contains $1,281,167$ training and $50,000$ test images from $10,000$ classes with resolution 32×32.

From the training images, $5,000$ images as the validation set and the remaining $1,276,167$ as the training set, and further use all $50,000$ test images as the test set.

ImageNet is a human curated collection of images downloaded from the web via search engines. While ImageNet used Amazon Mechanical Turk to lable the images, we were unable to find information on processes which ensured no personally identifiable or offensive content was contained in the images, which is somewhat likely given the "in-the-wild" nature of the dataset. The ImageNet website states that the copyright of the images does not belong to authors of ImageNet.

ImageNet64. The ImageNet64 dataset, a second downsampled version of the ImageNet database [44, 45], likewise contains $1,281,167$ training and $50,000$ validation images with resolution 64×64. We use the same data splits as for ImageNet32. Refer to the above paragraph on ImageNet32 for discussion of personally identifiable information, offensive content and consent.

CelebA. The CelebA dataset [46] contains $162,770$ training, $19,867$ validation and $19,962$ test images with resolution 64×64 which we directly use as our training, validation and test set, respectively. Our implementation is a modified version of the one provided in NVAE [10] (`https://github.com/NVlabs/NVAE/blob/master/datasets.py`).

CelebA images are "obtained from the Internet". The authors state that these images are not the property of the authors of associated institutions [68]. As this dataset shows the faces of humans, these images are personally identifiable. We were unable to identify a process by which consent for using these images was obtained, or how potential offensive content was prevented.

E Potential negative societal impacts

Our work provides mainly theoretical and methodological contributions to U-Nets and HVAEs, and we hence see no direct negative societal impacts. Since U-Nets are widely used in applications, our theoretical results and any future work derived from them may downstream improve such applications, and thus also enhance their performance in malicious uses. In particular, U-Nets are widely used in generative modelling, and here, our work may have an effect on the quality of 'deep fakes', fake datasets or other unethical uses of generative modelling. For HVAEs, our work may inspire novel models which may lead to improved performance and stability of these models, also when used in applications with negative societal impact.

F Model and training details

On the stability of training runs. VDVAE uses several techniques to improve the training stability of HVAEs: First, gradient clipping [69] is used, which reduces the effective learning rate of a mini-batch if the gradient norm surpasses a specific threshold. Second, gradient updates are skipped entirely when gradient norms surpass a second threshold, typically chosen higher than the one for gradient clipping. Third, gradient updates are also skipped if the gradient update would cause an overflow, resulting in NaN values.

In spite of the above techniques to avoid deterioration of training in very deep HVAEs, particularly when using a lot of weight-sharing, we experienced stability problems during late stages of training. These were particularly an issue on CIFAR10 in the late stages of training (on average roughly after 2 weeks of computation time), and often resulted in NaN values being introduced or posterior collapse. We did not extensively explore ways to prevent these in order to do minimal changes compared to vanilla VDVAE. We believe that an appropriate choice of the learning rate (e.g. with a decreasing schedule in later iterations) in combination with other changes to the hyperparameters may greatly help with these issues, but principled fixes of, for instance, the instabilities identified in Theorem 5 are likewise important.

Gradient checkpointing, and other alternatives to reduce GPU memory cost. A practical limitation of training deep HVAEs (with or without weight-shared layers) is their GPU memory cost: Training deeper HVAEs means storing more intermediate activations in GPU memory during the forward pass, when memory consumption reaches its peak at the start of the backward pass. This limits the depth of the networks that can be trained on given GPU resources. To address this issue, particularly when training models which may not even fit on used hardware, we provide a

prototype of a custom[5] *gradient checkpointing* implementation. Checkpointing occurs every few ResNet blocks which trades off compute for memory and can be used as an option. In gradient checkpointing, activations are stored only at specific nodes (checkpoints) in the computation graph, saving GPU memory, and are otherwise recomputed on-demand, requiring one additional forward pass per mini-batch [55]. Training dynamics remain unaltered. We note that other techniques exist specifically targeted at residual networks: For example, [70] propose to stochastically drop out entire residual blocks at training time [6]. This technique has two disadvantages: It changes training dynamics, and peak memory consumption varies between mini-batches, where particularly the latter is an inconvenient property for the practitioner as it may cause out-of-memory errors.

G Additional experimental details and results

In this section, we provide additional experimental details and results.

Hyperparameters and hyperparameter tuning. In the following, we describe the hyperparameters chosen in our experiments. As highlighted in the main text, we use the state-of-the-art hyperparameters of VDVAE [9] wherever possible. This was possible for CIFAR10, ImageNet32 and ImageNet64. On MNIST and CelebA, VDVAE [9] did not provide experimental results. For MNIST, we took the hyperparameters of CIFAR10 as the basis and performed minimal hyperparameter tuning, mostly increasing the batch size and tuning the number and repetitions of residual blocks. For CelebA, we used the hyperparameters of ImageNet64 with minimal hyperparameter tuning, focussing on the number and repetitions of the residual blocks. For all datasets, the main hyperparameter we tuned was the number and repetitions (through weight-sharing) of residual blocks ceteris paribus, i.e. without searching over the space of other important hyperparameters. As a consequence, it is likely that further hyperparameter tuning would improve performance as changing the number of repetitions changes (the architecture of) the model.

We provide three disjunct sets of hyperparameters: *global* hyperparameters (Table G.2), which are applicable to all runs, *data-specific* hyperparameters (Table G.3), which are applicable to specific datasets, and *run-specific* hyperparameters, which vary by run. The run-specific hyperparameters will be provided in the respective subsections of §G, where applicable.

In the below tables, 'factor of # channels in conv. blocks' refers to the multiplicative factor of the number of channels in the bottleneck of a (residual) block used throughout VDVAE. '# channels of z_l' refers to C in the shape [C, H, W] of the latent conditional distributions in the approximate posterior and prior, where height H and width W are determined by the resolution of latent $\boldsymbol{z}_l$. Likewise, '# channels in residual state' refers to C in the shape [C, H, W] of the residual state flowing through the decoder of VDVAE. 'Decay rate γ of evaluation model' refers to the multiplicative factor by which the latest model parameters are weighted during training to update the evaluation model.

Table G.2: Global hyperparameters.

factor of # channels in conv. blocks	0.25
Gradient skipping threshold	3000
Adam optimizer: Weight decay	0.01
Adam optimizer: β_1	0.9
Adam optimizer: β_2	0.9

G.1 "More from less": Parameter efficiency in HVAEs

In this experiment, we investigate the effect of repeating ResNet blocks in the bottom-up and top-down pass via weight-sharing. `rN` indicates that a ResNet block is repeated `N` times through weight-sharing where `r` is to be treated like an operator and `N` is a positive integer. In contrast, `xN`, already used in the official implementation of VDVAE, indicates `N` number of ResNet blocks without weight-sharing.

[5] Our implementation deviates from the official PyTorch implementation of gradient checkpointing which is slow when using multiple GPUs, and is based on https://github.com/csrhddlam/pytorch-checkpoint.

[6] This technique is called "stochastic depth" as the active depth of the network varies at random. In this work, however, we go with our earlier definition of this term which refers to the number of stochastic layers in our network, and thus avoid using its name to prevent ambiguities.

Table G.3: Data-specific hyperparameters.

Dataset	MNIST	CIFAR10	ImageNet32	ImageNet64	CelebA
Learning rate	0.0001	0.0002	0.00015	0.00015	0.00015
# iterations for learning rate warm-up	100	100	100	100	100
Batch size	200	16	8	4	4
Gradient clipping threshold	200	200	200	220	220
# channels of $\boldsymbol{z}_l$	8	16	16	16	16
# channels in residual state	32	384	512	512	512
Decay rate γ of evaluation model	0.9999	0.9999	0.999	0.999	0.999

In Tables G.4 and G.5, we provide the NLLs on the test set at convergence corresponding to the NLLs on the validation set during training which we reported in Fig. 5. In general, weight-sharing tends to improve NLL, and models with significantly less parameters reach or even surpass other models with more parameters. We refer to the main text for the intuition of this behavior. In Table G.5 (CIFAR10), we find that the not weight-sharing runs have test NLLs noticeably deviating from the results on the validation set, yet the overall trend of more weight-sharing improving NLL tends to be observed. This is in line with our general observation that our HVAE models are particularly unstable on CIFAR10.

Table G.4: A small-scale study on parameter efficiency of HVAEs on *MNIST*. We compare models with one, two, three and four parameterised blocks per resolution ($\{x1, x2, x3, x4\}$) against models with a single parameterised block per resolution weight-shared $\{2, 3, 5, 10, 20\}$ times ($\{r2, r3, r5, r10, r20\}$). We report NLL ($\downarrow$) measured on the test set, corresponding to the results on the validation set in Fig. 5. NLL performance increases with more weight-sharing repetitions and surpasses models without weight-sharing but with more parameters.

Neural architecture	# Params	NLL ($\downarrow$)
`r1/x1`	107k	≤ 86.87
`r2`	107k	≤ 85.25
`r3`	107k	≤ 84.92
`r5`	107k	≤ 83.92
`r10`	107k	≤ 82.67
`r20`	107k	≤ 81.84
`x2`	140k	≤ 84.44
`x3`	173k	≤ 82.64
`x4`	206k	≤ 82.46

In Table G.6, we provide key run-specific hyperparameters for the large-scale runs corresponding to Table 1 in the main text. Two points on the architecture of the encoder and the decoder are worth noting: First, note that the decoder typically features more parameters and a larger stochastic depth than the encoder. We here follow VDVAE which observed this distribution of the parameters to be beneficial. Second, note that while we experienced a benefit of weight-sharing, there is a diminishing return of the number of times a specific cell is repeated. Hence, we typically repeat a single block for no more than 10-20 times, beyond which performance does not improve while computational cost increases linearly with the number of repetitions. Exploring how to optimally exploit the benefit of weight-sharing in HVAEs would be an interesting aspect for future work.

Table G.5: A small-scale study on parameter efficiency of HVAEs on *CIFAR10*. We compare models with with one, two, three and four parameterised blocks per resolution ($\{x1, x2, x3, x4\}$) against models with a single parameterised block per resolution weight-shared $\{2, 3, 5, 10, 20\}$ times ($\{r2, r3, r5, r10, r20\}$). We report NLL ($\downarrow$) measured on the test set, corresponding to the results on the validation set in Fig. 5. NLL performance tends to increase with more weight-sharing repetitions. However, in contrast to the validation set (see Fig. 5) where this trend is evident, it is less so on the test set.

Neural architecture	# Params	NLL ($\downarrow$)
`r1/x1`	8.7m	≤ 4.17
`r2`	8.7m	≤ 4.93
`r3`	8.7m	≤ 4.78
`r5`	8.7m	–
`r10`	8.7m	≤ 4.32
`r20`	8.7m	≤ 3.54
`x2`	13.0m	≤ 5.77
`x3`	17.3m	≤ 3.07
`x4`	21.6m	≤ 3.01

Table G.6: A large-scale study of parameter efficiency in HVAEs. We here provide key run-specific hyperparameters corresponding to the results reported in Table 1 in the main text. Note that the row order of our runs directly corresponds with Table 1. δ refers to gradient clipping threshold. γ refers to the gradient skipping threshold. We use the same nomenclature for number of cells (`x`) and number of repetitions for one block (`r`) as before. In addition, as in VDVAE's official code base, we use `d` to indicate average pooling, where the integer before `d` indicates the resolution on which we pool, and the integer after indicates the down-scaling factor. Further, `m` indicates interpolating, where we up-scale from a source (integer after `m`) to a target resolution (integer before `m`).

Dataset	Method	Batch size	δ	γ	Encoder Architecture	Decoder Architecture

MNIST 28×28	`WS-VDVAE` (ours)	70	-	200	`32r3,32r3,32r3,32r3,32r3,32d2,`
					`16r3,16r3,16r3,16d2,`
					`8x6,8d2,4x3,4d4,1x3`
						`1x1,4m1,4x2,8m4,`
						`8x5,16m8,16r3,16r3,16r3,16r3,16r3,32m16,`
						`32r3,32r3,32r3,32r3,32r3,32r3,32r3,`
						`32r3,32r3,32r3`

	`VDVAE`* (ours)	70	-	200	`32x11,32d2,16x6,16d2,` `8x6,8d2,4x3,4d4,1x3`	`1x1,4m1,4x2,8m4,8x5,16m8,16x10,32m16,32x21`

CIFAR10 32×32	`WS-VDVAE` (ours)	16	400	4000	`32r12,32r12,32r12,32r12,32r12,32r12,`
					`32d2,16r12,16r12,16r12,16r12,16d2,`
					`8r12,8r12,8r12,8r12,8r12,8r12,8d2,`
					`4r12,4r12,4r12,4d4,1r12,1r12,1r12`
						`1r12,4m1,4r12,4r12,8m4,`
						`8r12,8r12,8r12,8r12,8r12,`
						`16m8,16r12,16r12,16r12,16r12,16r12,`
						`32m16,32r12,32r12,32r12,32r12,`
						`32r12,32r12,32r12,32r12,32r12`

	`WS-VDVAE` (ours)	16	200	2500	`32r3,32r3,32r3,32r3,32r3,32r3,`
					`32r3,32r3,32r3,32r3,32r3,32d2,`
					`16r3,16r3,16r3,16r3,16r3,16d2,`
					`8x6,8d2,4x3,4d4,1x3`
						`1x1,4m1,4x2,8m4,8x5,`
						`16m8,16r3,16r3,16r3,16r3,16r3,`
						`16r3,16r3,16r3,16r3,16r3,32m16,`
						`32r3,32r3,32r3,32r3,32r3,32r3,`
						`32r3,32r3,32r3,32r3,32r3,32r3,`
						`32r3,32r3,32r3,32r3,32r3,32r3,`
						`32r3,32r3,32r3`

	`VDVAE`* (ours)	16	200	400	`32x11,32d2,16x6,16d2,` `8x6,8d2,4x3,4d4,1x3`	`1x1,4m1,4x2,8m4,8x5,` `16m8,16x10,32m16,32x21`

ImageNet 32×32	`WS-VDVAE` (ours)	8	200	5000	`32r10,32r10,32r10,32r10,32d2,`
					`16r10,16r10,16r10,16d2,8x8,8d2,`
					`4x6,4d4,1x6`
						`1x2,4m1,4x4,8m4,8x9,16m8,`
						`16r10,16r10,16r10,16r10,16r10,32m16,`
						`32r10,32r10,32r10,32r10,32r10,32r10,`
						`32r10,32r10,32r10,32r10`

	`WS-VDVAE` (ours)	8	200	5000	`32r6,32r6,32r6,32r6,32r6,`
					`32r6,32r6,32r6,32r6,32d2,`
					`16r6,16r6,16r6,16r6,16r6,16d2,`
					`8x8,8d2,4x6,4d4,1x6`
						`1x2,4m1,4x4,8m4,8x9,`
						`16m8,16r6,16r6,16r6,`
						`16r6,16r6,16r6,16r6,16r6,16r6,`
						`16r6,16r6,32m16,32r6,32r6,`
						`32r6,32r6,32r6,32r6,32r6,32r6,`
						`32r6,32r6,32r6,32r6,32r6,32r6,32r6,32r6,`
						`32r6,32r6,32r6,32r6,32r6,32r6,32r6,32r6,32r6`

	`VDVAE`* (ours)	8	200	300	`32x15,32d2,16x9,16d2,8x8,8d2,` `4x6,4d4,1x6`	`1x2,4m1,4x4,8m4,8x9,16m8,16x19,32m16,32x40`

CelebA 64×64	`WS-VDVAE` (ours)	4	220	3000	`64r3,64r3,64r3,64r3,64r3,64r3,64r3,`
					`64d2,32r3,32r3,32r3,32r3,32r3,32r3,`
					`32r3,32r3,32r3,32r3,32r3,32d2,`
					`16r3,16r3,16r3,16r3,16r3,16r3,16d2,`
					`8r3,8r3,8r3,8d2,4r3,4r3,4r3,4d4,`
					`1r3,1r3,1r3`
						`1r3,1r3,4m1,4r3,4r3,4r3,`
						`8m4,8r3,8r3,8r3,8r3,`
						`16m8,16r3,16r3,16r3,16r3,16r3,16r3,16r3,`
						`32m16,32r3,32r3,32r3,32r3,32r3,32r3,`
						`32r3,32r3,32r3,32r3,32r3,32r3,32r3,`
						`32r3,32r3,32r3,64m32,64r3,64r3,64r3,64r3,`
						`64r3,64r3,64r3,64r3`

	`VDVAE`* (ours)	4	220	3000	`64x11,64d2,32x20,32d2,` `16x9,16d2,8x8,8d2,4x7,4d4,1x5`	`1x2,4m1,4x3,8m4,8x7,16m8,` `16x15,32m16,32x31,64m32,64x12`

G.2 HVAEs secretly represent time and make use of it

In this experiment, we measure the L_2 norm of the residual state at every ResNet block in both the forward (bottom-up/encoder) and backward (top-down/decoder) model. Let x_i be the output of ResNet block i in the bottom-up model, and y_i be the input of ResNet block i in the top-down model for one batch. In the following, augmenting Fig. 6 on MNIST in the main text, we measure $\|x_i\|_2$ or $\|y_i\|_2$, respectively, over 10 batches. We also use this data to compute appropriate statistics (mean and standard deviation) which we plot. We measure the state norm in the forward and backward pass for models trained on CIFAR10 and ImageNet32 in Figs. G.5 and G.6, respectively. We note that the forward pass of the ImageNet32 has a slightly unorthodox, yet striking pattern in terms of state norm magnitude, presumably caused by an overparameterisation of the model. In summary, these findings provide further evidence that the residual state norm of VDVAEs represents time.

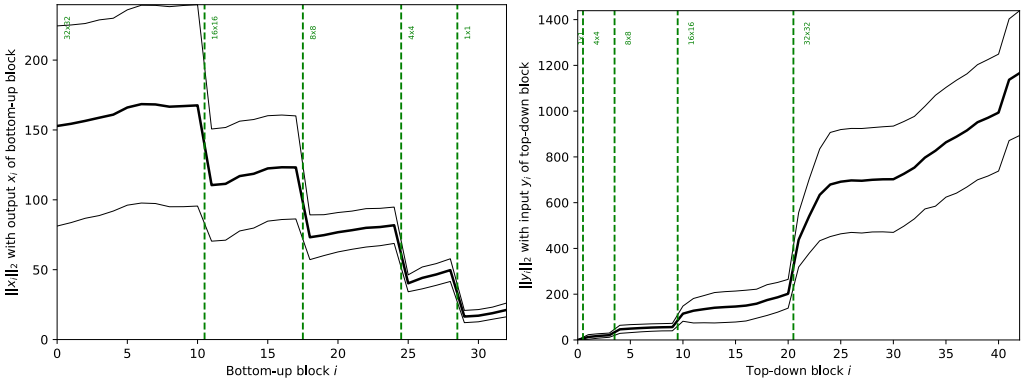

Figure G.5: HVAEs are secretly representing time *on CIFAR10*: We measure the L_2-norm of the residual state at every residual block i for the [Left] forward (bottom-up) pass, and [Right] the backward (top-down) pass, respectively, over 10 batches with 100 data points each. The thick line refers to the average and the thin, outer lines refer to ± 2 standard deviations.

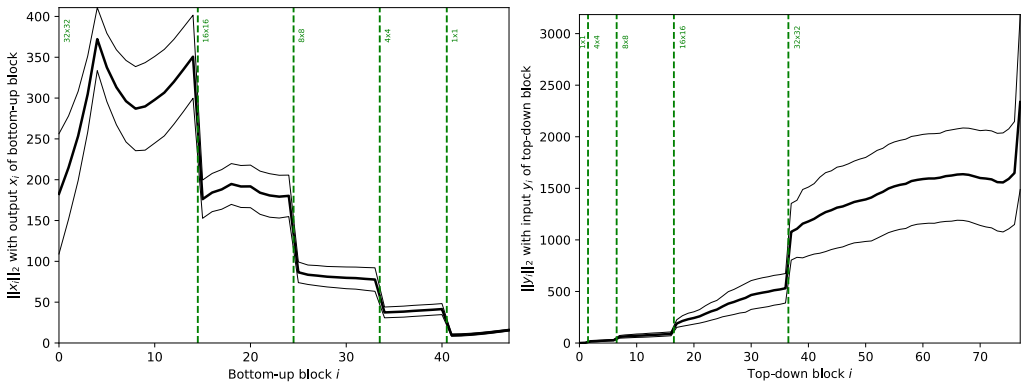

Figure G.6: HVAEs are secretly representing time *on ImageNet32*: We measure the L_2-norm of the residual state at every residual block i for the [Left] forward (bottom-up) pass, and [Right] the backward (top-down) pass, respectively, over 10 batches with 100 data points each. The thick line refers to the average and the thin, outer lines refer to ± 2 standard deviations.

When normalising the residual state in our experiments in Table 2 (case "normalised"), we do so at the same positions where we measure the state norm above. At the output of every forward ResNet block x_i and the input of every backward ResNet block y_i, we assign

$$x_i \leftarrow \frac{x_i}{\|x_i\|_2} \qquad\qquad y_i \leftarrow \frac{y_i}{\|y_i\|_2}$$

for every mini-batch during training. This results in a straight line in these plots for the "normalised" case. As the natural behavior of VDVAEs is—as we measured—to learn a non-constant norm, normalising the state norm has a deteriorating consequence, as we observe in Table 2. In contrast, the regular, unnormalised runs (case "non-normalised") show well-performing results.

We further analysed the normalised state norm experiments in Table 2. The normalised MNIST and CIFAR10 runs terminated early (indicated by ✗), more precisely after 18 hours and 4.5 days of training, respectively. From the very start of the optimisation, the normalised models have poor training behavior. To show this, in Fig. G.7, we illustrate the NLL on the validation set during training for the three normalised runs as compared to regular, non-normalised training. Validation ELBO only improves for a short time, after which the normalised runs deteriorate, showing no further improvement or even a worse NLL.

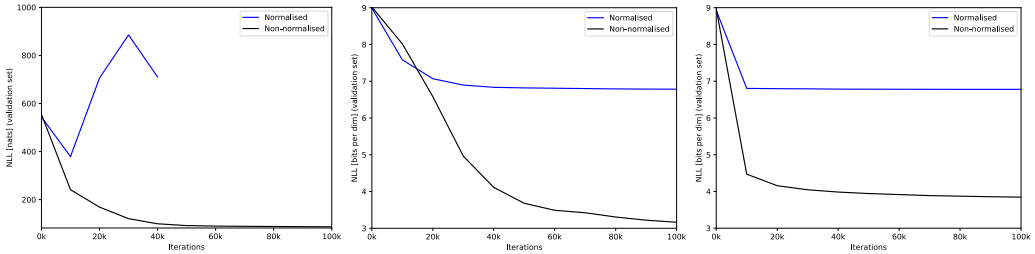

Figure G.7: On the training dynamics of VDVAE with and without a normalised residual state norm. NLL (↓) measured on the validation set of MNIST [left], CIFAR10 [middle] and ImageNet32 [right]. The normalised runs suffer from poor training dynamics from the very start of the optimisation and even terminate early on MNIST and CIFAR10, indicating that VDVAE makes use of the time representing state norm during training.

G.3 Sampling instabilities in HVAEs

When retrieving unconditional samples from our models, we scale the variances in the unconditional distributions with a temperature factor τ, as is common practice. We tune τ "by eye" to improve the fidelity of the retrieved images, yet do not cherry pick these samples. In Figs. G.8 to G.12, we provide additional, not cherry-picked unconditional samples for models trained on CIFAR10, ImageNet32, ImageNet64, MNIST and CelebA, extending those presented in Fig. 7. As shown earlier, the instabilities in VDVAE result in poor unconditional samples for CIFAR10, ImageNet32 and ImageNet64, but relatively good samples for MNIST and CelebA.

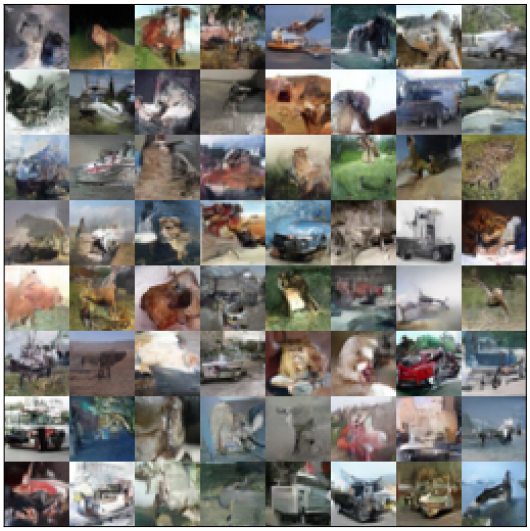

Figure G.8: Further unconditional samples (not cherry-picked) of VDVAE* on *CIFAR10*, augmenting those presented in Fig. 7. While samples on MNIST and CelebA demonstrate high fidelity and diversity, samples on CIFAR10, ImageNet32 and ImageNet64 are diverse, but unrecognisable, demonstrating the instabilities identified by Theorem 5. We chose the temperature as $\tau = 0.9$.

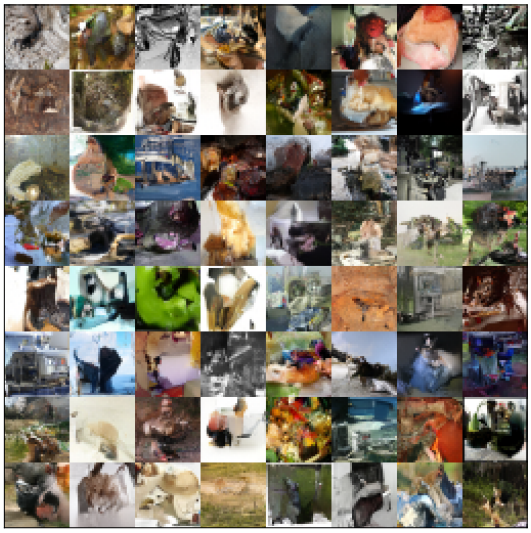

Figure G.9: Further unconditional samples (not cherry-picked) of VDVAE* on *ImageNet32*, augmenting those presented in Fig. 7. While samples on MNIST and CelebA demonstrate high fidelity and diversity, samples on CIFAR10, ImageNet32 and ImageNet64 are diverse, but unrecognisable, demonstrating the instabilities identified by Theorem 5. We chose the temperature as $\tau = 1.0$.

Figure G.10: Further unconditional samples (not cherry-picked) of VDVAE* on *ImageNet64*, augmenting those presented in Fig. 7. While samples on MNIST and CelebA demonstrate high fidelity and diversity, samples on CIFAR10, ImageNet32 and ImageNet64 are diverse, but unrecognisable, demonstrating the instabilities identified by Theorem 5. Temperatures τ are tuned for maximum fidelity. We chose the temperature as $\tau = 0.9$.

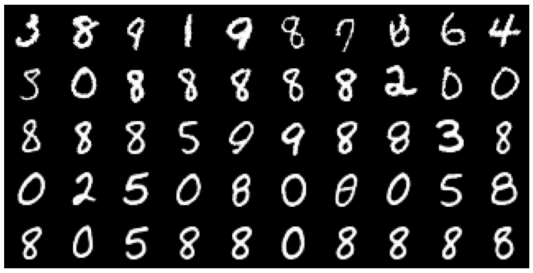

Figure G.11: Further unconditional samples (not cherry-picked) of VDVAE* on *MNIST*, augmenting those presented in Fig. 7. While samples on MNIST and CelebA demonstrate high fidelity and diversity, samples on CIFAR10, ImageNet32 and ImageNet64 are diverse, but unrecognisable, demonstrating the instabilities identified by Theorem 5. We chose the temperatures as $\tau \in \{1.0, 0.9, 0.8, 0.7, 0.5\}$ (corresponding to the rows).

Figure G.12: Further unconditional samples (not cherry-picked) of VDVAE* on *CelebA*, augmenting those presented in Fig. 7. While samples on MNIST and CelebA demonstrate high fidelity and diversity, samples on CIFAR10, ImageNet32 and ImageNet64 are diverse, but unrecognisable, demonstrating the instabilities identified by Theorem 5. We chose the temperature as $\tau = 0.5$.

In addition, we here also visualise the representational advantage of HVAEs. Fig. G.13 shows samples where we gradually increase the number of samples from the posterior vs. the prior distributions in each resolution across the columns. This means that in column 1, we sample the first latent $\boldsymbol{z}_l$ *in each resolution* from the (on encoder activations conditional) posterior q, and all other latents from the prior p. A similar figure, but gradually increasing the contribution of the posterior across the blocks of all resolutions (i.e. column 1 samples $\boldsymbol{z}_l$ from the posterior in the very first resolution only) is shown in VDVAE [9, Fig. 4]. Fig. G.13

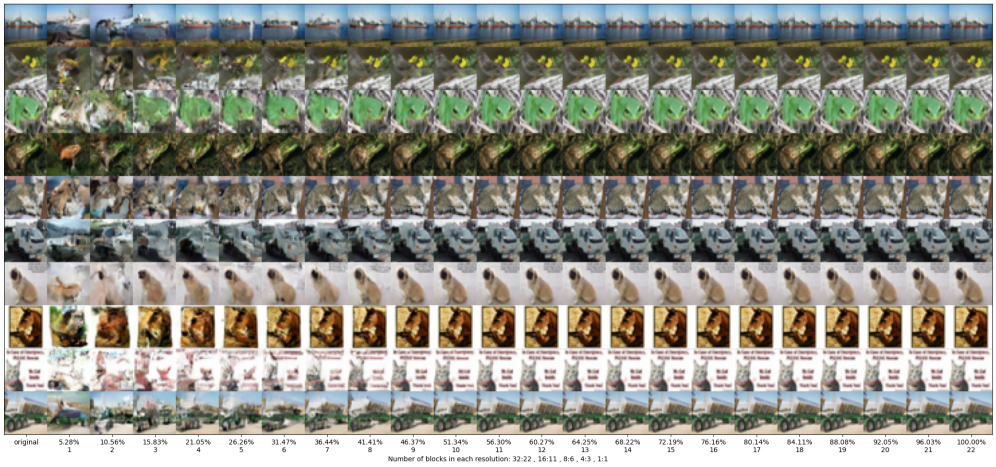

Figure G.13: Samples drawn from our model when gradually increasing the contribution of the approximate posterior. In each column with integer s, we sample the first s latent variables from the approximate posterior in each resolution, i.e. $\mathbf{z}_i \sim q(\mathbf{z}_i|\mathbf{z}_{>i})$ (up to the maximum number of latent variables in each resolution), and $\mathbf{z}_j \sim p(\mathbf{z}_j|\mathbf{z}_{>j})$ for all other latent variables. The percentage number indicates the fraction of the number of latent variables among all latent variables sampled from the approximate posterior. In the left-most column, we visualise corresponding input images.

G.4 Ablation studies

G.4.1 Number of latent variables

The number of latent variables increase when increasing the stochastic depth through weight-sharing. Thus, an important ablation study is the question whether simply increasing the number of latent variables improves HVAE performance, which may explain the weight-sharing effect. On CIFAR10, we find that this is not the case: In Table G.7, we analyse the effect of increasing the number of latent variables ceteris paribus. Furthermore, in Fig. G.14, we report validation NLL during training for the same runs. In this experiment, we realise the increase in number of latent variables by increasing the number of channels in each latent variable $\mathbf{z}_l$ exponentially while slightly decreasing the number of blocks so to keep the number of parameters roughly constant. Both results indicate that the number of latent variables, at least for this configuration on CIFAR10, do not add performance and hence cannot explain the weight-sharing performance.

Table G.7: On the effect of the number of latent variables on CIFAR10. We report the NLL on the test set at convergence.

# of latent variables	# Params	NLL (↓)
396k	39m	2.88
792k	39m	2.88
1.584m	39m	2.87
3.168m	39m	2.88

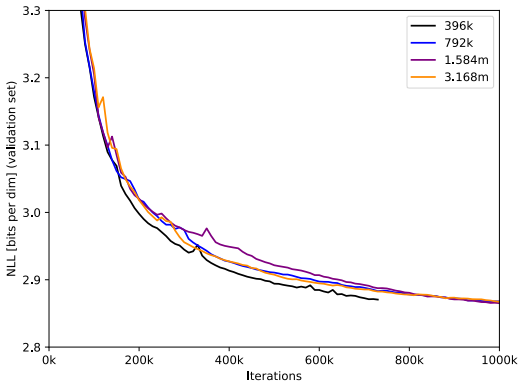

Figure G.14: On the effect of the number of latent variables. We report NLL on the validation set of CIFAR10 during training.

G.4.2 Fourier features

In this experiment, we are interested in the effect of Fourier features imposed onto a Haar wavelet basis due to the inductive bias of the U-Net. Intuitively, we would expect that Fourier features do not add performance as the U-Net already imposes a good basis for images. We now validate this hypothesis experimentally: We compute Fourier features in every ResNet block at three different locations as additional channels and varying frequencies. We implement Fourier features closely following VDM [19]: Let $h_{i,j,k}$ be an element of a hidden activation of the network, for example of a sampled latent $h = \mathbf{z}_l$, in spatial position (i, j) and channel k. Then, for each scalar $h_{i,j,k}$, we add two transformed scalars for each frequency governed by β as follows:

$$f^{\beta}_{i,j,k} = \sin\left(z_{i,j,k} 2^{\beta} \pi\right), \text{ and } g^{\beta}_{i,j,k} = \cos\left(z_{i,j,k} 2^{\beta} \pi\right).$$

In our experiments, we experiment with different choices for β, but typically select two values at a time (as in VDM), increasing the number of channels in the resulting activation by a factor of five. Fourier features are computed on and concatenated to activations at three different locations (in three separate experiments): At the input of the ResNet block, after sampling, and for the input of the two branches parameterising the posterior and prior distributions.

In Tables G.8 and G.9, we report performance when concatenating Fourier features at every ResNet block in these three locations. In all cases, Fourier features deteriorate performance in this multi-resolution wavelet basis, particularly for high-frequencies which often lead to early termination due to numeric overflows. However, if training only a single-resolution model where no basis is enforced, training does not deteriorate, not even for high-frequency Fourier features, yet performance can neither be improved. Furthermore, we experimented with computing and concatenating the Fourier features only to the input image of the model, hypothesising numerical instabilities caused by computing Fourier transforms at every ResNet block, and report results in Table G.10. Here, performance is significantly better as runs no longer deteriorate, but Fourier features still do not improve performance compared to not using Fourier features at all.

Table G.8: Fourier features introduced and concatenated in every ResNet block at three different locations on *MNIST*. VDVAE typically deteriorates or has poor performance.

Exponent β	NLL
Loc. 1	
$[1, 2]$	≤ 78.4
$[3, 4]$	≤ 80.55
$[5, 6]$ (✗)	–
Loc. 2	
$[1, 2]$	≤ 554.50
$[3, 4]$ (✗)	–
$[5, 6]$ (✗)	–
Loc. 3	
$[1, 2]$ (✗)	–
$[2, 3]$	≤ 306.67
$[3, 4]$	≤ 345.67
Loc. 1 & single-res.	
$[3, 4]$	≤ 87.55
$[5, 6]$	≤ 86.96
$[7, 8]$	≤ 91.67
No Fourier Features	≤ 79.81

Table G.9: Fourier features introduced and concatenated in every ResNet block at three different locations on *CIFAR10*. VDVAE typically deteriorates or achieves a poor performance.

Exponent β	NLL
Loc. 1	
$[3, 4]$ (✗)	–
$[5, 6]$ (✗)	–
$[7, 8]$	≤ 8.94
Loc. 2	
$[3, 4]$ (✗)	–
$[5, 6]$ (✗)	–
$[7, 8]$ (✗)	–
Loc. 3	
$[3, 4]$ (✗)	–
$[5, 6]$	≤ 8.94
$[7, 8]$	≤ 8.99
No Fourier Features	≤ 2.87

Table G.10: Fourier features introduced on the input image of the model only, with results on *CIFAR10*. While performing better than if introduced at every ResNet block, still Fourier features do not improve performance compared to using no Fourier features at all.

Exponent β	NLL
Fourier Features on input only	
$[3, 4]$	≤ 2.95
$[5, 6]$	≤ 2.96
$[7, 8]$	≤ 2.89
No Fourier Features	≤ 2.87

G.4.3 On the effect of a multi-resolution bridge.

State-of-the-art HVAEs have a U-Net architecture with pooling and, hence, are multi-resolution bridges (see Theorem 4). We investigate the effect of multiple resolutions in HVAEs (here with spatial dimensions $\{32^2, 16^2, 8^2, 4^2, 1^2\}$) against a single resolution (here with spatial dimension 32^2). We choose the number of blocks for the single resolution model such that they are distributed in the encoder and decoder proportionally to the multi-resolution model and the total number of parameters are equal in both, ensuring a fair comparison. As we show in Table G.11, the multi-resolution models perform slightly better than their single-resolution counterparts, yet we would have expected this difference to be more pronounced. We also note that it may be worth measuring other metrics for instance on fidelity, such as the FID score [71]. Additionally, multi-resolution models have a representational advantage due to their Haar wavelet basis representation (illustrated in Appendix G.4, Fig. G.13).

Table G.11: Single- vs. multi-resolution HVAEs.

# Resolutions	# Params	NLL
MNIST		
Single	328k	≤ 81.40
Multiple	339k	≤ 80.14
CIFAR10		
Single	39m	≤ 2.89
Multiple	39m	≤ 2.87
ImageNet32		
Single	119m	≤ 3.68
Multiple	119m	≤ 3.67

G.4.4 On the importance of a stochastic differential equation structure in HVAEs

A key component of recent HVAEs is a residual cell, as outlined in §4. The residual connection makes HVAEs discretise an underlying SDE, as we outlined in this work. Experimentally, it was previously noted as being crucial for stability of very deep HVAEs. Here, we are interested in ablating the importance of imposing an SDE structure into HVAEs: We compare models with a residual HVAE cell (as in VDVAE) with a non-residual HVAE cell which is as close to VDVAE as possible to ensure a fair comparison. The non-residual VDVAE cell does not possess a residual state which flows through the backbone architecture. We achieve this by removing the connection between the first and second element-wise addition in VDVAE's cell (see [9, Fig. 3]), which is equivalent to setting $Z_{i,+} = 0$. Hence, in the non-residual cell, during training and evaluation, the reparameterised sample is directly taken forward. Note that this is distinct from the Euler-Maruyama cell which features a residual connection. Our experiments confirm that a *residual* cell is key for training stability, as illustrated in Table G.12 and Fig. G.15: Without a residual state flowing through the decoder, models quickly experience posterior collapse of the majority of layers during training.

G.4.5 Synchronous vs. asynchronous processing in time

During the bottom-up pass, VDVAE takes forward the activation of the last time step in each resolution which is passed to every time step in the top-down pass on the same resolution (see Fig. 3 in VDVAE [9]). In this ablation study, we were interested in this slightly peculiar choice of an *asynchronous* forward and backward process and to what degree it is important for performance. We thus compare an asynchronous model, with skip connections as in VDVAE [9], with a *synchronous* model, where activations from the bottom-up pass are taken forward to the corresponding time step in the top-down pass. In other words, in the synchronous case, the skip connection mapping between time steps in the encoder and decoder is 'bijective', and it is not 'injective', but 'surjective' in the asynchronuous case. We realise the synchronous case by choosing the same number of blocks in the encoder as `VDVAE`* has in the decoder, i.e. constructing a 'symmetric' model. To ensure a fair comparison, both models (synchronous and asynchronous) are further constructed to have the same number of parameters. In Table G.13, we find that synchronous and asynchronous processing achieve comparable NLL, indicating that the asynchronous design is not an important contributor to performance in VDVAE. We

Table G.12: Residual vs. non-residual VDVAE cell. The residual HVAE strongly outperforms a non-residual VDVAE cell, where the latter's training deteriorates. This is also analysed in Fig. G.13. We report NLL on the test set at convergence, or at the last model checkpoint before deterioration of training.

Cell type	NLL
MNIST	
Residual VDVAE cell	≤ 80.05
Non-residual VDVAE cell	≤ 112.58
CIFAR10	
Residual VDVAE cell	≤ 2.87
Non-residual VDVAE cell	≤ 3.66
ImageNet	
Residual VDVAE cell	≤ 3.667
Non-residual VDVAE cell (✗)	≤ 4.608

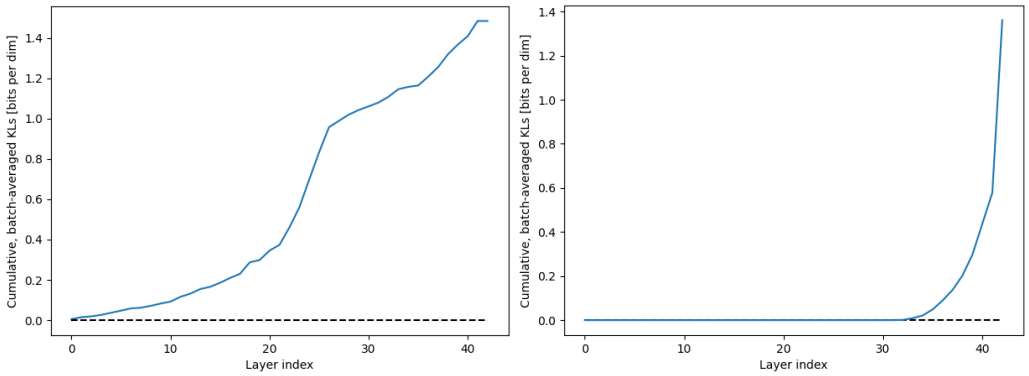

Figure G.15: Cumulative sum of KL-terms in the ELBO of a residual and non-residual VDVAE, averaged over a batch at convergence. We report the two CIFAR10 runs in Table G.12. The posterior collapses for the majority of the latent variables in the non-residual VDVAE cell case [right], but carries information for all latent variables in the regular, residual cell case [left].

note, however, that an advantage of the asynchronous design, which is exploited by VDVAE, is that the bottom-up and top-down architectures can have different capacities, i.e. have a different number of ResNet blocks. VDVAE found that a more powerful decoder was beneficial for performance [9].

Table G.13: Synchronous vs. asynchronous processing in time. We report NLL on the test set on CIFAR10 and ImageNet32, respectively.

Processing	NLL
CIFAR10	
Synchronous	≤ 2.85
Asynchronous	≤ 2.86
ImageNet32	
Synchronous	≤ 3.69
Asynchronous	≤ 3.69