arXiv:2510.02826v1 [cs.LG] 3 Oct 2025

Multi-scale Autoregressive Models are Laplacian, Discrete, and Latent Diffusion Models in Disguise

Steve Hong
University of Cambridge
mdh58@cam.ac.uk

Samuel Belkadi
University of Cambridge
bk2764@cam.ac.uk

Abstract

We revisit Visual Autoregressive (VAR) models through the lens of an iterative-refinement framework. Rather than viewing VAR solely as next-scale autoregression, we formalise it as a deterministic forward process that constructs a Laplacian-style latent pyramid, paired with a learned backward process that reconstructs it in a small number of coarse-to-fine steps. This view connects VAR to denoising diffusion and isolates three design choices that help explain its efficiency and fidelity: refining in a learned latent space, casting prediction as discrete classification over code indices, and partitioning the task by spatial frequency. We run controlled experiments to quantify each factor's contribution to fidelity and speed, and we outline how the same framework extends to permutation-invariant graph generation and to probabilistic, ensemble-style medium-range weather forecasting. The framework also suggests practical interfaces for VAR to leverage tools from the diffusion ecosystem while retaining few-step, scale-parallel generation.

1 Introduction

Generative modelling spans two complementary paradigms with distinct strengths and limitations [1]. Next-token autoregressive (AR) models fit discrete sequence generation (e.g., text) efficiently [2], but struggle with high-fidelity image generation because their imposed causal orders, such as raster-scan, disrupt natural spatial structures [3, 4]. Diffusion models (DDPM) [5], on the other hand, excel on continuous data, such as images and audio [6]. However, they typically require hundreds to thousands of refinement steps [5] and requires adjustments for discrete outputs [7].

Visual Autoregressive (VAR) models [4] were proposed as a coarse-to-fine alternative that narrows the fidelity gap between diffusion and AR. VAR generates images by predicting the next higher-resolution scale conditioned on all lower-resolution scales, and producing every token at a given scale in one pass rather than token-by-token. In practice, this scale-parallel strategy achieves competitive image quality with significantly fewer iterations, rendering it faster than diffusion.

VAR attributes its gains primarily to the next-scale prediction framework [4]. However, scale-wise refinement alone is not new: the classic Laplacian pyramid [8] follows the same coarse-to-fine idea, and several diffusion models already explore multi-scale structures [9–12], yet they arguably did not achieve the same fidelity and speed as VAR. It is therefore not immediately clear how scale refinement is the main reason for VAR's success. This observation motivates a further analysis to explain such attributions.

Recent work advances along two tracks: incremental VAR modifications and broader developments in diffusion. On the VAR side, most studies address engineering bottlenecks: reducing KV-cache and long-sequence overhead with collaborative decoding [13], localising scale refinement to first-order Markov to boost efficiency [14], and reducing codebook bottlenecks by introducing an infinite-vocabulary classifier with bitwise self-correction [15]. On the diffusion side, methods aim to reduce

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: SPIGM workshop @ NeurIPS 2025.

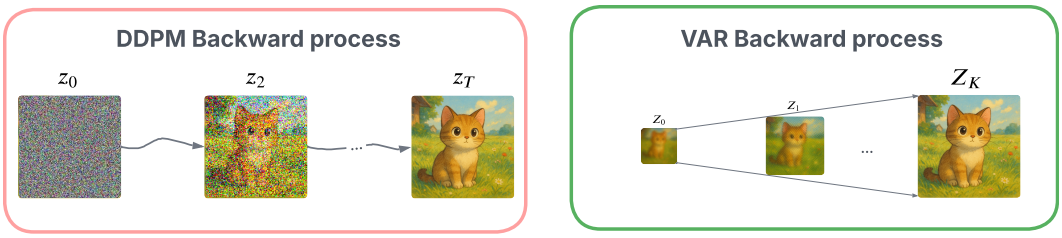

Figure 1: Iterative refinement in two paradigms. VAR refines by predicting next-scale discrete codes across a small number of scales, while diffusion refines by denoising continuous latents over multiple timesteps.

the step count or per-step cost by employing one- or few-step consistency objectives [16], progressive distillation [17], and fast ODE solvers [18], often in conjunction with latent-consistency adapters. Because these efforts proceed in parallel rather than jointly, a unified iterative-refinement perspective would not only allow for consolidation in research efforts but also open paths for VAR to leverage the diffusion ecosystem.

Core contributions. In this work, we derive VAR as an iterative refinement model rather than a purely multi-scale autoregressive model, enabling a direct connection to denoising diffusion.
We address two questions:

1. In the iterative-refinement formulation of VAR, how can we decompose and quantify which elements drive its success?
2. What implications does this insight have for generative tasks beyond images, including molecular graph generation and weather forecasting?

2 Background

2.1 Visual Autoregressive Models

VQ-VAE tokeniser. VAR uses a vector-quantized autoencoder that turns an input image $x \in \mathbb{R}^{H\times W\times 3}$ into a grid of discrete codes. An encoder E_ψ produces a continuous feature map $z_e = E_\psi(x)$. Each vector in z_e is quantized to its nearest codebook embedding e_i from a learnable codebook $\mathcal{E} = \{e_i\}_{i=1}^{V}$, yielding the quantized map z_q that is fed to a decoder D_ϕ to reconstruct $\hat{x} = D_\phi(z_q)$. Training uses the VQ-VAE objective with a reconstruction term and the codebook/commitment terms:

$$\mathcal{L}_{\mathrm{VQ}} \;=\; \underbrace{\|x-\hat{x}\|_2^2}_{\text{reconstruction}} \;+\; \underbrace{\|\,\mathrm{sg}[z_e]-e\,\|_2^2}_{\text{codebook}} \;+\; \beta\,\underbrace{\|\,z_e-\mathrm{sg}[e]\,\|_2^2}_{\text{commitment}}. \tag{1}$$

where $\mathrm{sg}[\cdot]$ denotes the stop-gradient operator. This loss causes the encoder to commit to discrete codes, while the codebook tracks encoder outputs via vector quantisation, providing discrete token grids suitable for autoregression. VAR replaces a single-scale grid with a stack of token maps $(r^{(1)},\ldots,r^{(S)})$ at increasing resolutions produced by a multi-scale VQ tokeniser.

Next-scale autoregression. Let $r^{(k)}$ denote the token map at scale k. VAR models the joint as

$$p(r^{(1)},\ldots,r^{(S)}) \;=\; \prod_{k=1}^{S} p\Big(r^{(k)} \mid r^{(<k)}\Big), \tag{2}$$

and predicts all tokens within a scale in parallel, using attention that is conditioned only on the prefix lower-resolution maps $r^{(\le k-1)}$ (block-causal across scales). This shift from next-token to next-scale prediction preserves 2D locality and reduces generation complexity relative to raster AR, as each refinement step operates on a 2D grid at a single scale.

It is important to note that the tokeniser and generator operate in *residual* scales: at level k the encoder forms a residual token map by subtracting the upsampled contribution of coarser scales ($x \leftarrow x - \phi_k(z_k)$), and at generation time the model *predicts residuals* and reconstructs by upsample-and-add across all scales ($\hat{x} \leftarrow \hat{x} + \phi_k(z_k)$). [4]

2.2 Denoising Diffusion Probabilistic Models

DDPMs [5] define a fixed forward (noising) Markov chain

$$q(x_t \mid x_{t-1}) = \mathcal{N}\big(\sqrt{1-\beta_t}\, x_{t-1},\, \beta_t I\big), \quad t = 1, \dots, T, \tag{3}$$

which admits a closed-form marginal

$$q(x_t \mid x_0) = \mathcal{N}\big(\sqrt{\bar{\alpha}_t}\, x_0,\, (1-\bar{\alpha}_t) I\big), \quad \bar{\alpha}_t = \textstyle\prod_{s=1}^{t} \alpha_s,\ \alpha_t = 1 - \beta_t. \tag{4}$$

The reverse chain is learned as Gaussian conditionals

$$p_\theta(x_{t-1} \mid x_t) = \mathcal{N}\big(\mu_\theta(x_t, t),\, \Sigma_\theta(t)\big), \tag{5}$$

by optimizing a variational bound that decomposes into tractable KL terms between Gaussians (thanks to the closed-form forward posterior $q(x_{t-1} \mid x_t, x_0)$). In practice, parameterizing μ_θ via ε_θ (predicting the added noise) yields the simplified training step

$$\min_\theta\ \mathbb{E}_{t, x_0, \varepsilon} \left\| \varepsilon - \varepsilon_\theta\big(\sqrt{\bar{\alpha}_t}\, x_0 + \sqrt{1-\bar{\alpha}_t}\, \varepsilon,\ t\big) \right\|_2^2, \tag{6}$$

and sampling iterates $x_T \sim \mathcal{N}(0, I) \to x_0$ using the learned reverse conditionals.

There are two important remarks about DDPM: **(i)** Closed-form marginals and posteriors make training low-variance and fully supervised at each noise level (the KL decomposition compares Gaussians in closed form). **(ii)** Many-step generation: quality improves with the number of reverse steps T, but naïve DDPM sampling is slow; learning (or interpolating) reverse variances enables strong quality with far fewer steps (tens instead of hundreds), materially improving practicality.

2.3 Laplacian pyramid

The Laplacian pyramid [8] represents an image as a low-pass residual plus a stack of band-pass detail layers arranged from coarse to fine. Let $G_0 = x$ be the original image and let $\mathrm{down}(\cdot)$ denote low-pass filtering followed by subsampling. The Gaussian pyramid is built by

$$G_{k+1} \ = \ \mathrm{down}(G_k), \qquad k = 0, \dots, S-1. \tag{7}$$

Define $\mathrm{up}(\cdot)$ as expand (upsample) followed by low-pass filtering that is matched to the analysis filter. The Laplacian (detail) levels are then

$$L_k \ = \ G_k \ - \ \mathrm{up}(G_{k+1}), \qquad k = 0, \dots, S-1, \tag{8}$$

together with the coarsest residual G_S. Perfect reconstruction follows by upsample-and-add:

$$\widehat{G}_S \ = \ G_S, \qquad \widehat{G}_k \ = \ L_k \ + \ \mathrm{up}(\widehat{G}_{k+1}) \ \text{ for } k = S-1, \dots, 0, \tag{9}$$

yielding $\widehat{x} = \widehat{G}_0 = x$ when the synthesis operator $\mathrm{up}(\cdot)$ is the adjoint (or appropriately matched) to the analysis operator in $\mathrm{down}(\cdot)$. Each L_k behaves as a spatially localised band-pass component, while G_S captures the low frequencies.

3 Multi-scale Autoregressive Models as Iterative Refinement Models

3.1 Iterative refinement

We view *iterative refinement* as generating by a short sequence of learned updates on a state x. From an easy initialiser $x_0 \sim \pi_0$, apply step-specific maps to approach the target:

$$x_{t+1} = \Phi_\theta^{(t)}(x_t, c), \qquad t = 0, \dots, T-1, \quad \hat{x} = x_T. \tag{10}$$

Each step is trained to make a local improvement, e.g. by decreasing a task energy or satisfying a consistency relation:

$$E(x_{t+1}; c) \le E(x_t; c) - \delta_t \quad \text{or} \quad \mathcal{C}(x_{t+1}, x_t, c) \approx 0. \tag{11}$$

A coarse-to-fine schedule arises by constraining $\Phi_\theta^{(t)}$, or the representation, such that early steps set global structure while later steps add detail; this applies to discrete or continuous x; in pixel, latent, or other spaces.

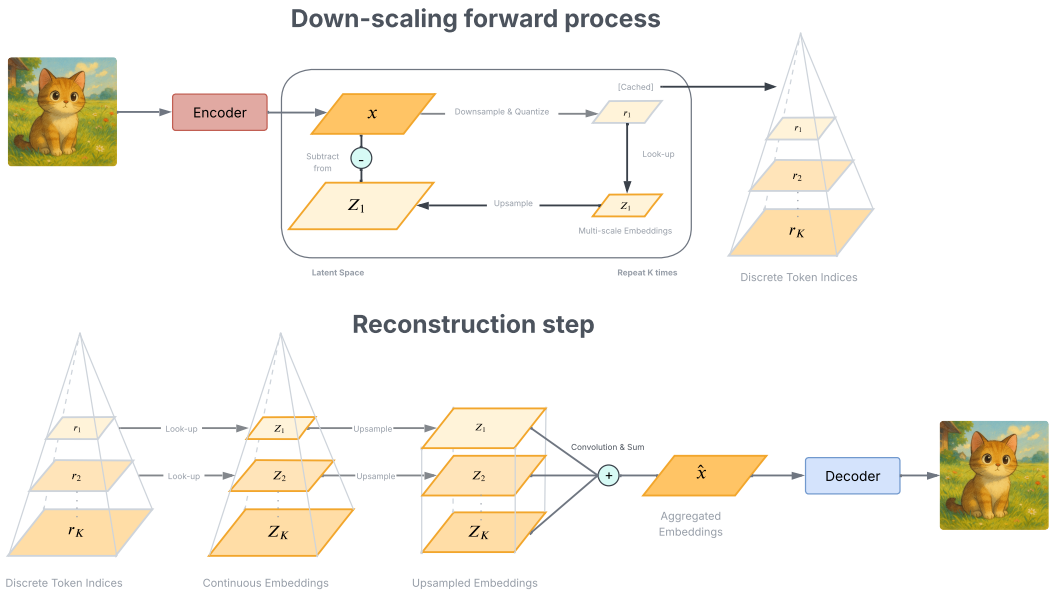

Figure 2: **Forward process of VAR.** The image is encoded into a latent grid, progressively downsampled to build a pyramid. At each level, we upsample the coarser level to compute a residual, then quantise that residual into discrete indices, and also quantise the coarsest level as base codes. This forward path is deterministic and provides the training targets for the generator.

3.2 Rewriting VAR with iterative refinement

We now define VAR as an iterative refinement framework. We work in the discrete latent space produced by a VQ-VAE tokeniser: the encoder maps x to a continuous feature z_e, which is vector-quantised to an index grid $z^{(0)}$, and the decoder maps quantised latents back to pixels. Let $L^{(0)} \in \mathbb{R}^{H_0 \times W_0 \times d}$ denote the corresponding d-channel latent obtained by embedding each index into $\mathbb{R}^d$. The multi-scale pyramid is then built by deterministic downscaling of latents, $L^{(k)} = \mathcal{S}_k\big(L^{(k-1)}\big)$, and, in the spirit of a Laplacian pyramid, we represent the hierarchy in residuals $R^{(k-1)} = L^{(k-1)} - \mathcal{U}_k\big(L^{(k)}\big)$. At each level, these residuals are quantised with per-scale codebooks whose entries lie in $\mathbb{R}^d$, and the coarsest latent $L^{(S)}$ is likewise quantised to provide the base of the refinement process.

The forward process. As in Figure 2 first encode the image into latent space,

$$L^{(0)} \;=\; F(x). \tag{12}$$

We then form a latent Gaussian pyramid by deterministic downscaling,

$$L^{(k)} \;=\; \mathcal{S}_k\big(L^{(k-1)}\big), \qquad k = 1, \ldots, S, \tag{13}$$

and define latent residuals in Laplacian form,

$$R^{(k-1)} \;=\; L^{(k-1)} \;-\; \mathcal{U}_k\big(L^{(k)}\big), \qquad k = 1, \ldots, S. \tag{14}$$

Each residual is pointwise quantised with a d-channel codebook, while the coarsest latent may be quantised as the base:

$$r^{(k-1)} \;=\; Q_k^{\text{res}}\big(R^{(k-1)}\big), \quad k = 1, \ldots, S, \qquad b^{(S)} \;=\; Q_S^{\text{base}}\big(L^{(S)}\big). \tag{15}$$

The forward distribution is therefore fully deterministic,

$$q\Big(b^{(S)}, r^{(0)}, \ldots, r^{(S-1)} \mid x\Big) \;=\; \delta\big(b^{(S)} - Q_S^{\text{base}}(L^{(S)})\big) \prod_{k=1}^{S} \delta\big(r^{(k-1)} - Q_k^{\text{res}}(R^{(k-1)})\big), \tag{16}$$

The backward process. Generation proceeds from coarse to fine in latent space by predicting residual indices in parallel at each scale and reconstructing latents via upsample-and-add. Let the d-channel embeddings associated with the residual and base codebooks be $\mathrm{emb}_k^{\mathrm{res}} : \{0, \ldots, V_k - 1\} \to \mathbb{R}^d$ and $\mathrm{emb}_S^{\mathrm{base}} : \{0, \ldots, V_S - 1\} \to \mathbb{R}^d$. We first sample a factorised prior over the coarsest base codes and embed them,

$$b^{(S)} \sim p_\theta\big(b^{(S)}\big), \qquad L^{(S)} \;=\; \mathrm{emb}_S^{\mathrm{base}}\big(b^{(S)}\big). \tag{17}$$

For $k = S, \ldots, 1$ we form the upscaled context and a Transformer ψ_θ predicts the entire residual index map at the next finer scale in one shot,

$$C^{(k-1)} \;=\; \mathcal{U}_k\big(L^{(k)}\big), \qquad p_\theta\big(r^{(k-1)} \mid L^{(k)}\big) \;=\; \prod_{i=1}^{H_{k-1}W_{k-1}} \mathrm{Cat}\Big(r_i^{(k-1)} \,;\, \psi_\theta\big(C^{(k-1)}\big)_{i,\cdot}\Big), \tag{18}$$

followed by residual embedding and latent reconstruction,

$$\widehat{R}^{(k-1)} \;=\; \mathrm{emb}_k^{\mathrm{res}}\big(r^{(k-1)}\big), \qquad L^{(k-1)} \;=\; C^{(k-1)} \;+\; \widehat{R}^{(k-1)}. \tag{19}$$

The finest latent is finally decoded to pixels,

$$\hat{x} \;=\; G\big(L^{(0)}\big). \tag{20}$$

Remark. In this formulation, **(i)** tokens within each scale are generated fully in parallel: conditioned on the upsampled context $C^{(k-1)}$, the categorical factors across spatial sites are independent, so the entire residual map $r^{(k-1)}$ is produced in one shot. **(ii)** The number of refinement steps is small: just S scales (typically S=8 in the original setup), versus the hundreds to a thousand time steps used by diffusion samplers. **(iii)** The (latent) likelihood is exact and tractable because the model defines a finite autoregressive factorization across scales, whereas DDPMs posit a first-order Markov chain $p_\theta(x_{0:T}) = p(x_T) \prod_{t=1}^{T} p_\theta(x_{t-1} \mid x_t)$ and are trained/evaluated via a variational bound (ELBO) built from Gaussian KL terms rather than exact token-level likelihood.

3.3 Discussion

Given this view, both DDPM and VAR are instances of iterative refinement; this perspective surfaces three design choices that may explain why VAR performs strongly relative to DDPM:

Latent space. Refining in a learned latent space may reduce spatial dimensionality and attenuate pixel-level redundancy, so updates operate on more aggregated features and reduce task complexity. This could improve the signal-to-noise ratio of supervision and shorten the effective sequence length, broadly in line with observations from latent diffusion [19].

Discrete domain. Framing prediction over categorical indices can rephrase a continuous regression task as classification, which may yield steadier gradients and mitigate regression-to-the-mean effects [20, 21]. A finite codebook also acts as a mild prior over reusable visual primitives and might enable simple stepwise error-correction, though DDPM's continuous noise targets [5] offer different benefits.

Frequency refinement. Replacing denoising with next-scale upscaling explicitly partitions the image into S bands of increasing spatial frequency, akin to a Laplacian pyramid [8]. Each refinement step can specialise: early steps set low-frequency layout and colour; later steps add higher-frequency texture and edges. This frequency factorisation may reduce interference present at intermediate DDPM noise levels, where a single update mixes multiple frequencies. It may also align supervision with perceptual importance per band and enable within-scale parallel token generation, thereby improving the fidelity-speed trade-off, even though S must still be tuned.

4 Experiments

We design three small MLP experiments on MNIST [22] to probe the three hypotheses from Section 3.3: **(i)** refining in a compact latent may ease optimisation, **(ii)** discrete prediction may stabilise training relative to continuous regression, and **(iii)** coarse-to-fine banded refinement may reduce interference between frequencies. We maintain models and budgets in alignment and report the mean $\pm$ standard deviation across three seeds.

Common setup. We use MNIST 28×28 images normalised to $[0,1]$. We apply additive Gaussian noise with $\sigma{=}0.20$ at train and test unless stated otherwise. All predictors are 2-layer MLPs (width 512, GELU, with LayerNorm applied before the output layer), trained with AdamW (learning rate 10^{-3}, weight decay 10^{-4}) for 30 epochs with a batch size of 256. For latent-space experiments, we pretrain a linear autoencoder with a $d{=}64$ bottleneck using MSE and then freeze it; we denote its encoder and decoder by E and D. For Laplacian experiments, we use a fixed separable $[1,4,6,4,1]/16$ blur for analysis.

Setup A: Pixel vs. latent denoising

Task. The task is to predict the clean target given noisy $x{+}\varepsilon$. We compare: *Pixel-MLP* $f_\theta : \tilde{x} \mapsto \hat{x}$ with MSE to x; and *Latent-MLP* $g_\phi : E(\tilde{x}) \mapsto \widehat{z}$ with MSE to $z{=}E(x)$, followed by reconstruction $\hat{x}{=}D(\widehat{z})$.

Metrics. We report test MSE ↓, PSNR ↑, epochs-to-target (MSE ≤ 0.018) ↓, and robustness under a noise shift when testing at $\sigma{=}0.30$.

Table 1: Setup A: Pixel vs. latent denoising on MNIST.

Method	MSE @ $\sigma{=}0.20$ ↓	PSNR (dB) ↑	Epochs to 0.018 ↓	MSE @ $\sigma{=}0.30$ ↓
Pixel-MLP	0.0195 ± 0.0003	17.6 ± 0.1	14 ± 1	0.0280 ± 0.0010
Latent-MLP	$\mathbf{0.0158}\pm0.0002$	$\mathbf{18.9}\pm0.1$	$\mathbf{9}\pm1$	$\mathbf{0.0221}\pm0.0010$

Observation. The latent model reaches a given error in fewer epochs and maintains lower error under a noise shift. This pattern is consistent with the idea that a compact representation may improve signal-to-noise and shorten the effective optimisation path.

Setup B: Discrete classification vs. continuous regression

Task. We first build a k-means codebook (here $k{=}64$) on clean *latent* features $E(x)$. We then train a shared MLP trunk with two alternative heads: a *Regression* head that predicts continuous z (MSE), and a *Softmax-64* head that predicts the nearest code index (cross-entropy), which we dequantise to the centroid for reconstruction $\hat{x}{=}D(\widehat{z})$. As a third variant, *Bitwise-8* uses eight independent Bernoulli heads (logistic) over a learned 8-bit code for the same codebook.

Metrics. We report dequantised test MSE ↓, mini-batch gradient-norm variance (arbitrary units, lower is steadier) ↓, and relative training time per epoch.

Table 2: Setup B: Discrete prediction stabilises training at fixed capacity.

Head	Dequantised MSE ↓	Grad var ($\times10^{-3}$) ↓	Time/epoch (rel.)
Regression (MSE)	0.0204 ± 0.0004	2.4 ± 0.2	1.00
Softmax-64 (CE)	$\mathbf{0.0181}\pm0.0003$	1.5 ± 0.1	1.02
Bitwise-8 (BCE)	0.0185 ± 0.0003	$\mathbf{1.1}\pm0.1$	**0.92**

Observation. Classification-based targets yield steadier gradients and lower reconstruction error at the same budget. The bitwise factorisation further reduces per-step cost and variance. These results suggest that reframing regression as classification can simplify optimisation.

Setup C: Coarse-to-fine residual learning vs. single-shot

Task. We construct a 2-level Laplacian pyramid with G_1 (14×14 low-pass) and R_0 (28×28 residual). We compare: a *Single-shot* MLP that predicts the full clean image from the noisy input; and a *Coarse→fine* two-stage approach in which (i) $\text{MLP}_{\text{coarse}}$ denoises G_1, and (ii) $\text{MLP}_{\text{resid}}$ predicts R_0 conditioned on the upsampled $\widehat{G}_1$ concatenated with the noisy input, after which we reconstruct $\hat{x}{=}\text{up}(\widehat{G}_1){+}\widehat{R}_0$. We match the total number of parameters to the single-shot baseline.

Metrics. We evaluate test MSE ↓, epochs-to-target (MSE ≤ 0.018) ↓, and high-frequency PSNR (HF-PSNR on R_0) ↑.

Table 3: Setup C: Coarse-to-fine specialisation reduces interference.

Method	MSE ↓	Epochs to 0.018 ↓	HF-PSNR (dB) ↑
Single-shot (one MLP)	0.0193±0.0003	12±1	16.8±0.2
Coarse→fine (2 MLPs)	**0.0164**±0.0002	**7**±1	**18.2**±0.2

Depth sweep. Increasing the number of bands beyond two yields diminishing returns:

Table 4: Effect of refinement depth S with matched total parameters.

S (bands)	MSE ↓	Train time/epoch (rel.)
1	0.0193	1.00
2	**0.0164**	1.15
3	0.0162	1.22

Observation. Two banded stages already capture most of the gain, while additional splits offer only small improvements at higher cost. This outcome aligns with the idea that a small number of coarse-to-fine refinements can be sufficient in practice.

5 Beyond Images: Applications of Iterative Refinement

5.1 Graph generation

Graph generation [23, 24] seeks to learn distributions over graphs for applications in chemistry, biology, and network science, where the target law should be invariant to node permutations.

A central challenge is ensuring permutation invariance of the induced distribution. Many autoregressive constructions impose an arbitrary node ordering, which can cause the likelihood or sampling trajectory to depend on that ordering and therefore break invariance [25]. Typical remedies include using exchangeable priors and permutation-equivariant architectures (e.g., GNNs where message passing and pooling commute with node permutations), but care is still required in training objectives and decoding to avoid reintroducing order dependence.

Recent work [26] has instantiated a VAR-style, coarse-to-fine generator for graphs and reported competitive sample quality, along with significant inference speedups, suggesting that next-scale prediction may transfer beyond images.

Diffusion-based graph models, on the other hand, establish that, if the score/denoising network is permutation equivariant, the resulting generative distribution is permutation invariant [27].

Why the iterative refinement perspective may help. Using the iterative-refinement perspective, we can reach an analogous conclusion for VAR. First, VAR generates by a short sequence of refinement steps, mirroring diffusion's iterative updates. Second, if the refinement map at each scale is a Transformer that omits order-dependent positional encodings, then each step is order-invariant and, thus, permutation-invariant. With an exchangeable base prior at the coarsest scale, the composition of permutation-invariant steps yields a permutation-equivariant distribution over graphs. While this is an informal argument, it provides a concrete design rule: use exchangeable priors and set-structured Transformers (or equivariant GNNs) for all refinement maps to ensure permutation invariance.

5.2 Medium-range weather forecasting

Medium-range (3–15 day) forecasts inform public safety and energy planning, where distribution-aware evaluation matters. GenCast [28] established a skilful, probabilistic ML forecaster by framing the task as generative modelling with diffusion, providing calibrated ensemble trajectories rather than a single path. In contrast, leading deterministic ML forecasters (e.g., GraphCast [29], Pangu [30]) are single-pass generators trained for point predictions and do not provide distributional forecasts.

Why the iterative refinement perspective may help. An iterative-refinement (VAR-style) forecaster is probabilistic by construction: each scale predicts categorical distributions over discrete latents, so diverse ensemble members arise with low incremental cost. Coarse-to-fine frequency updates also align with atmospheric structure: early scales capture planetary/synoptic patterns while later scales add mesoscale detail, and replace many diffusion timesteps with a small number of scale transitions, thereby improving efficiency without abandoning distributional modelling.

6 Limitations and Future Work

Scale of experiments. Our empirical study is intentionally small: we use controlled MLP baselines to isolate effects under limited compute. A more complete assessment would train a reference VAR model end to end on standard image benchmarks and compare to its diffusion analogue under similar settings. Specifically: (i) *latent vs. pixel*: run next-scale refinement in VQ-VAE latents and, as a "no-latent" control, in a Laplacian pixel pyramid; (ii) *discrete vs. continuous*: keep the pyramid but replace the categorical head with a Gaussian regression head to test quantisation; (iii) *multi-scale vs. single-scale*: collapse to a single latent grid to assess the value of coarse to fine. For each setting, attach a *diffusion head* that predicts ε with T steps in the same space (pixel diffusion, latent diffusion, and discrete diffusion, respectively). Report quality–cost frontiers (e.g., FID or PSNR versus wall-clock and versus network evaluations), calibration (e.g., bits/dim with standard dequantization), and sensitivity to S scales versus T steps.

Deterministic forward process. In vanilla VAR, the forward path that builds the latent/residual pyramid is fully deterministic. Stochasticity appears only in bitwise self-correction variants (e.g., [15]) that inject Bernoulli noise at the code level. Whether mild forward randomness helps remains an open question and requires further experiments.

Leverage the diffusion ecosystem. Several techniques from diffusion may transfer directly to VAR's iterative-refinement view. Examples include classifier-free guidance adapted to categorical heads, self-conditioning across scales, loss re-weighting akin to simplified diffusion objectives, few-step distillation of a large-S refiner into a small-S student, early-exit or schedule search analogous to noise-schedule tuning, and solver-style improvements such as consistency training to reduce the number of refinement transitions. One might also explore rectified-flow style objectives between consecutive scales, latent-consistency adapters for faster sampling, and temperature or top-k calibration per scale to trade fidelity for diversity. Given the fruitful research in the diffusion literature, it is now the time for VAR to leverage many techniques from this field.

7 Conclusion

We reframed Visual Autoregressive models as iterative refiners operating on a latent Laplacian pyramid. This view makes the generation path explicit: a short sequence of coarse-to-fine residual updates, and clarifies how VAR relates to denoising diffusion while retaining exact cross-scale factorisation and within-scale parallelism. Under this lens, three design choices appear central: working in a compact latent space, predicting discrete indices rather than continuous targets, and allocating supervision across frequency bands. Small toy on MNIST support these hypotheses, indicating that a modest number of banded steps can capture most of the benefit.

Beyond images, the same template suggests practical routes for permutation-equivariant graph generation and for efficient probabilistic weather forecasting, where calibrated ensembles are required. The formulation also exposes straightforward interfaces to the diffusion ecosystem, including guidance, consistency-style objectives, and few-step distillation adapted to categorical heads and scale schedules. A fuller assessment will require end-to-end studies with matched backbones on standard benchmarks, but the iterative-refinement framing already offers a common language for analysing, ablating, and improving both VAR and diffusion families.

References

[1] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM computing surveys*, 56(4):1–39, 2023.

[2] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[3] Aaron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, and Koray Kavukcuoglu. Conditional image generation with pixelcnn decoders, 2016.

[4] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. In *Proceedings of the 38th Annual Conference on Neural Information Processing Systems*, 2024.

[5] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

[6] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis, 2021.

[7] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in neural information processing systems*, 34:17981–17993, 2021.

[8] Peter J Burt and Edward H Adelson. A multiresolution spline with application to image mosaics. *ACM Transactions on Graphics (ToG)*, 2(4):217–236, 1983.

[9] Jonathan Ho, Chitwan Saharia, William Chan, David J. Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation, 2021.

[10] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement, 2021.

[11] Wan-Cyuan Fan, Yen-Chun Chen, DongDong Chen, Yu Cheng, Lu Yuan, and Yu-Chiang Frank Wang. Frido: Feature pyramid diffusion for complex scene image synthesis. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 579–587, 2023.

[12] Yuval Atzmon, Maciej Bala, Yogesh Balaji, Tiffany Cai, Yin Cui, Jiaojiao Fan, Yunhao Ge, Siddharth Gururani, Jacob Huffman, Ronald Isaac, et al. Edify image: High-quality image generation with pixel space laplacian diffusion models. *arXiv preprint arXiv:2411.07126*, 2024.

[13] Zigeng Chen, Xinyin Ma, Gongfan Fang, and Xinchao Wang. Collaborative decoding makes visual auto-regressive modeling efficient, 2024.

[14] Jinhua Zhang, Wei Long, Minghao Han, Weiyi You, and Shuhang Gu. Mvar: Visual autoregressive modeling with scale and spatial markovian conditioning, 2025.

[15] Jian Han, Jinlai Liu, Yi Jiang, Bin Yan, Yuqi Zhang, Zehuan Yuan, Bingyue Peng, and Xiaobing Liu. Infinity: Scaling bitwise autoregressive modeling for high-resolution image synthesis. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 15733–15744, 2025.

[16] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. 2023.

[17] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022.

[18] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in neural information processing systems*, 35:5775–5787, 2022.

[19] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.

[20] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.

[21] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019.

[22] Li Deng. The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE signal processing magazine*, 29(6):141–142, 2012.

[23] Jiaxuan You, Rex Ying, Xiang Ren, William Hamilton, and Jure Leskovec. Graphrnn: Generating realistic graphs with deep autoregressive models. In *Proceedings of the 35th International Conference on Machine Learning*, pages 5708–5717. PMLR, 2018.

[24] Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. In *Proceedings of the Eleventh International Conference on Learning Representations*, 2023.

[25] Lingxiao Zhao, Xueying Ding, and Leman Akoglu. Pard: Permutation-invariant autoregressive diffusion for graph generation. In *Proceedings of the 38th Annual Conference on Neural Information Processing Systems*, 2024.

[26] Samuel Belkadi, Steve Hong, Marian Chen, Miruna Cretu, Charles Harris, and Pietro Lio. Diffusion-free graph generation with next-scale prediction. *arXiv preprint arXiv:2503.23612*, 2025.

[27] Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon. Permutation invariant graph generation via score-based generative modeling. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.

[28] Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R Andersson, Andrew El-Kadi, Dominic Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, et al. Gencast: Diffusion-based ensemble forecasting for medium-range weather. *arXiv preprint arXiv:2312.15796*, 2023.

[29] Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, Alexander Merose, Stephan Hoyer, George Holland, Oriol Vinyals, Jacklynn Stott, Alexander Pritzel, Shakir Mohamed, and Peter Battaglia. Learning skillful medium-range global weather forecasting. *Science*, 382(6677):1361–1366, 2023. doi: 10.1126/science.adi2336. URL https://arxiv.org/abs/2212.12794.

[30] Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Accurate medium-range global weather forecasting with 3d neural networks. *Nature*, 619:533–538, 2023. doi: 10.1038/s41586-023-06185-3. See also arXiv:2211.02556.