

PATTERN-GUIDED DIFFUSION MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Diffusion models have shown promise in forecasting future data from multivariate time series. However, few existing methods account for recurring structures, or patterns, that appear within the data. We present Pattern-Guided Diffusion Models (PGDM), which leverage inherent patterns within temporal data for forecasting future time steps. PGDM first extracts patterns using archetypal analysis and estimates the most likely next pattern in the sequence. By guiding predictions with this pattern estimate, PGDM makes more realistic predictions that fit within the set of known patterns. We additionally introduce a novel uncertainty quantification technique based on archetypal analysis, and we dynamically scale the guidance level based on the pattern estimate uncertainty. We apply our method to two well-motivated forecasting applications, predicting visual field measurements and motion capture frames. On both, we show that pattern guidance improves PGDM’s performance (MAE / CRPS) by up to 40.67% / 56.26% and 14.12% / 14.10%, respectively. PGDM also outperforms baselines by up to 65.58% / 84.83% and 93.64% / 92.55%.

1 INTRODUCTION

Diffusion models are a class of generative models that perform generation by iteratively removing noise from a noisy sample. These models are easier to train and generate higher quality images compared to the previous state-of-the-art, generative adversarial networks (Dhariwal & Nichol, 2021). Recent work has found success in using diffusion models to forecast future steps of temporal data (Chang et al., 2024; Feng et al., 2024; Gu et al., 2022; Hu et al., 2024; Li et al., 2022; Lv et al., 2024; Rasul et al., 2021; Wen et al., 2023). Such methods, however, rarely leverage the recurring structures that often manifest in temporal data. These appear, for example, in medical modalities due to the physiology and anatomy of the human body. Basketball videos also contain repeated structures due to standardized courts, player positions, and strategies. The few diffusion-based forecasters that exploit these *patterns* often overlook changes over time and uncertainties in pattern representation (Wang et al., 2024; Westny et al., 2024; Zhao et al., 2024).

In this paper, we present Pattern-Guided Diffusion Models (PGDM) for forecasting temporal data with inherent patterns. Using archetypal analysis (Cutler & Breiman, 1994), we extract patterns from training data, then train a guidance function to predict future pattern contributions to the data. PGDM then forecasts future points guided by these predictions. To handle evolving patterns, we introduce a novel uncertainty metric that dynamically tunes the scale of pattern guidance.

We evaluate PGDM on two impactful applications. First, we consider the clinical application of visual field (VF) prediction. Visual field tests measure a patient’s functional vision, and the resulting measurements manifest common patterns across patients due to eye anatomy. Furthermore, forecasting future VF measurements can serve as a decision aid for clinicians. On a real-world VF dataset, we find that pattern guidance improves the performance (MAE / CRPS) of PGDM by up to 40.67% / 56.26% on average, surpassing baseline models by up to 65.58% / 84.83% on average. Next, we consider the application of forecasting future motion capture frames for human motion prediction. Pose patterns frequently appear in common human movements, such as walking and running. Predicting human motion may aid advancements in human robot collaboration and autonomous driving. On motion capture frames for a variety of dance genres, we show that PGDM is able to leverage even the diverse, highly dynamic patterns that present in dance motion. Pattern guidance improves the performance (MAE / CRPS) of PGDM by up to 14.12% / 14.10% on average, allowing PGDM to surpass baselines by up to 93.64% / 92.55% on average.

054 In summary, our contributions are as follows.

- 055
- 056 1. We present Pattern-Guided Diffusion Models (PGDM), which leverage inherent patterns
- 057 within temporal data.
- 058 2. We introduce a novel uncertainty quantification method based on archetypal analysis, and
- 059 we show that this uncertainty metric captures geometric distance from the training set.
- 060 3. We show that the proposed uncertainty quantification metric approximately lower bounds
- 061 the error of the pattern predictions that guide PGDM.
- 062 4. We propose a method to dynamically tune the level of pattern guidance based on the proposed
- 063 uncertainty metric.
- 064
- 065

066 2 RELATED WORKS

067

068 Diffusion models have been widely used for temporal forecasting. TimeGrad (Rasul et al., 2021) and

069 LDT (Feng et al., 2024) forecast multivariate time series conditioned on histories, while BVAE (Li

070 et al., 2022) uses a bi-directional VAE for the reverse diffusion process. Models like USTD (Hu et al.,

071 2024) and DiffSTG (Wen et al., 2023) capture spatial dependencies using spatio-temporal graphs.

072 Other applications include pedestrian trajectory (Lv et al., 2024; Gu et al., 2022) and medical sensor

073 signal prediction (Chang et al., 2024). For a comprehensive overview, see Yang et al. (2024).

074 Few such diffusion-based forecasters attempt to leverage patterns within the data. Hypothesizing that

075 past patterns likely reappear later, Diff-MGR (Zhao et al., 2024) conditions predictions on previous

076 patterns. Westny et al. (2024) also proposed to guide predictions of traffic trajectories using patterns,

077 formalized as environment maps, as agent behaviors are often dictated by the environment (e.g., cars

078 stay within the road lanes). Similarly to Diff-MGR, Westny et al. (2024) assumes that manifested

079 patterns will remain constant over time, as predictions are conditioned on a fixed map. Wang

080 et al. (2024) instead proposed to guide predictions with dynamically changing patterns, captured by

081 segmented real-time camera readings, as humans are most likely to walk towards specific destinations

082 within a scene, such as a door, stairs, or a hallway. However, these approaches do not account for

083 uncertainty in the guiding patterns. In contrast, our PGDM model adapts the level of pattern guidance

084 based on the estimated reliability of dynamically evolving patterns.

085 3 BACKGROUND

086

087 Let the data of interest be $x \in \mathbb{R}^d$ sampled from distribution $p(x)$, which arrives in a temporal

088 sequence $\{x_t\}$ with time index t . We are concerned with such data that contains *patterns*, or repeating

089 structures. Given an observed history of length T over time $t \in \{1, 2, \dots, T\}$, we aim to predict

090 a horizon of length H over time $t \in \{T + 1, T + 2, \dots, T'\}$ with $T' = T + H$. Denote a set of

091 n of history and horizon pairs by $\{x_{1:T}^i, x_{T:T'}^i\}_{i=1}^n$. We use archetypal analysis, which identifies

092 patterns resembling real data rather than an abstraction, to overcome the challenge of extracting

093 useful patterns. Here, we provide a brief overview of diffusion models and archetypal analysis.

094 3.1 DIFFUSION MODELS

095

096 Diffusion models (Ho et al., 2020; Sohl-Dickstein et al., 2015) aim to learn and generate data from a

097 distribution $p(x)$ through a forward and reverse process, in which noise is iteratively added to and

098 removed from the data, respectively. Given $x_{t,0} \sim p(x_t)$ at time t , the fixed S -step forward process

099 creates a sequence of increasingly noisy samples $x_{t,1}, x_{t,2}, \dots, x_{t,S}$. Note that here we use the

100 notation $x_{t,s}$, where t denotes the time index and s denotes the diffusion step. The noisy samples are

101 drawn from the distribution $q(x_{t,s}|x_{t,s-1}) := \mathcal{N}(\sqrt{1-\beta_s}x_{t,s-1}, \beta_s\mathbf{I})$, where $\beta_1, \beta_2, \dots, \beta_S$ is a

102 noise variance schedule. With appropriately chosen variance schedule, this distribution approaches a

103 standard normal as $S \rightarrow \infty$. Conveniently, the noising step $x_{t,s}$ can be sampled in closed form given

104 $x_{t,0}$, $q(x_{t,s}|x_{t,0}) = \mathcal{N}(\sqrt{\bar{\alpha}_s}x_{t,0}, (1-\bar{\alpha}_s)\mathbf{I})$, where $\alpha_s = 1 - \beta_s$ and $\bar{\alpha}_s = \prod_{i=1}^s \alpha_i$.

105

106 Conversely, the reverse process removes noise from $x_{t,S} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to ultimately recover the training

107 distribution. The goal is to learn the distribution $p_\theta(x_{t,s-1}|x_{t,s}) := \mathcal{N}(\mu_\theta(x_{t,s}, s), \sigma_s^2\mathbf{I})$. The

denoising parameters θ are learned by optimizing the evidence lower bound (ELBO) on negative

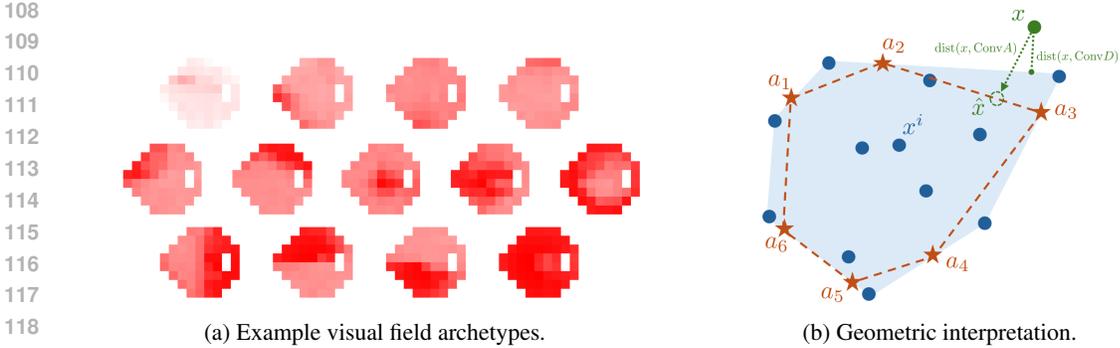


Figure 1: Overview of archetypal analysis (AA). (a) Example archetypes extracted from a visual field dataset. The archetypes capture visual loss patterns that consistently appear across glaucoma patients. Darker regions indicate greater vision loss. (b) Given a dataset D (blue dots), AA identifies a set of archetypes A (red stars). Within $\text{Conv}A \subseteq \text{Conv}D$ (within red dashed line), any point can be reconstructed from A without error. For any point $x \notin \text{Conv}A$, the reconstruction error is the distance between x and $\text{Conv}A$.

log likelihood, $-\log p_\theta(x_{t,0})$. At inference time, samples from the learned distribution $p_\theta(x_{t,0})$ are generated by applying the reverse process over S steps to noisy samples.

3.2 ARCHETYPAL ANALYSIS

Archetypal analysis (AA) (Cutler & Breiman, 1994) extracts extremal patterns, or archetypes, from a given dataset. These archetypes are themselves a combination of the data and therefore are realistic and interpretable representations of the data’s significant patterns. Figure 1a shows examples extracted from a visual field dataset. Furthermore, any point within the given dataset can be constructed as a combination of these archetypes, allowing for the contribution of each pattern to be quantified.

More formally, given dataset $D = \{x^i\}_{i=1}^n$, AA finds the p archetypes $a_1, a_2, \dots, a_p \in \mathbb{R}^d$ that minimize the residual sum of squares (RSS) error,

$$\min_{c_j^i, a_j^i} \forall i, j \sum_{i=1}^n \left\| x^i - \sum_{j=1}^p c_j^i a_j \right\|^2, \quad (1)$$

subject to $c_j^i \geq 0 \forall j$ and $\sum_{j=1}^p c_j^i = 1$ for each $i \in \{1, 2, \dots, n\}$, and where $a_j = \sum_{k=1}^n \beta_j^k x^k$ with $\beta_j^k \geq 0 \forall k$ and $\sum_{k=1}^n \beta_j^k = 1$ for each $j \in \{1, 2, \dots, p\}$. That is, if zero reconstruction error is achieved, then each $x^i \in D$ can be reconstructed as a convex combination of the archetypes, and the archetypes are themselves a convex combination of the data. Then, given a set of identified archetypes, any new data point can be reconstructed as a convex combination of the archetypes, with coefficients found by minimizing the objective in equation 1 with fixed a_1, a_2, \dots, a_p .

Figure 1b visualizes a geometric interpretation of AA. Intuitively, AA identifies a region in which any point can be perfectly reconstructed by the archetypes. This region is the convex hull of the archetypes, denoted $\text{Conv}A$ for the archetype set A . In the ideal case when $p = n$, the set of archetypes is exactly D , and the region of reconstructible points fully encompasses the dataset. In practice, $p < n$ is typically chosen, and the resulting archetypes instead lie on the boundary of $\text{Conv}D$ (Cutler & Breiman, 1994, Proposition 1). The archetype set therefore defines a reconstructible region that closely, but often not fully, captures the dataset D . Furthermore, for any $x \notin \text{Conv}A$, the reconstruction error can be thought of as the distance between x and $\text{Conv}A$. In Section 4.2, we show that this reconstruction error can be used to estimate the distance between x and the dataset D .

4 METHODS

Our goal is to learn the conditional distribution $p(x_{T:T'} | x_{1:T}, \hat{P}_{T:T'})$, where $\hat{P}_{T:T'}$ represents the patterns estimated to manifest in the future. To determine a pattern representation space, we first extract archetype set A from the training data. The representation space is the contribution of archetypes to each data point. We then train our proposed Pattern-Guided Diffusion Model (PGDM) to learn the desired conditional distribution.

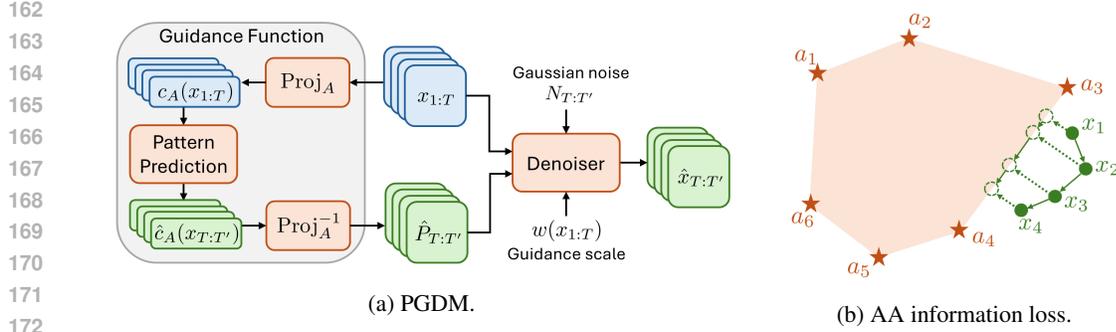


Figure 2: (a) Overview of Pattern-Guided Diffusion Models. A pattern guidance function estimates the future patterns over a horizon. This prediction is performed in the archetype space, where the projected representation captures the contribution of each pattern to the data. The predicted future patterns are supplied as conditioning context to the diffusion model. The scale of the pattern guidance is dynamically tuned based on the uncertainty of the input sequence. (b) Projection into the lower-dimensional archetype space loses information for any point outside the convex hull of the archetypes $\text{Conv}A$. In this example, the sequence $x_1, x_2, x_3, x_4 \notin \text{Conv}A$ is perturbed to the boundary of $\text{Conv}A$ after reconstruction, losing important dynamical information.

Figure 2a summarizes PGDM, which is constructed from three key components. 1) A pattern guidance function predicts the patterns \hat{P} that will appear over the future horizon. First, the input sequence $x_{1:T}$ is projected to the archetype space, where the projected pattern representation $c_A(x_{1:T})$ quantifies the contribution of each pattern to the data. The projected representation for the future steps is then predicted, and the resulting estimate $\hat{c}_A(x_{T:T'})$ is lifted back to the original data space as a predicted pattern $\hat{P}_{T:T'}$. 2) While prediction in the lower p -dimensional archetype space helps to evade the curse of dimensionality, the projection operation naturally incurs an information loss. By design, while any point $x_t \in \text{Conv}A$ can be projected with no information loss, this cannot be said for a point $x_t \notin \text{Conv}A$, which may appear when the data’s patterns change over time. Figure 2b visualizes this loss of information. We therefore introduce a novel Archetypal Analysis Uncertainty Quantification (AAUQ) technique, which we show captures geometric distance from the training set. This uncertainty metric also estimates a lower bound on the loss of the guidance function. 3) We use the predicted pattern from the guidance function as additional conditioning context for the diffusion model, following the general methodology of classifier-free diffusion guidance (Ho & Salimans, 2022). To account for uncertainties in the guidance function, we dynamically tune the guidance scale based on our uncertainty metric. We now describe each of the three components in further detail.

4.1 GUIDANCE FUNCTION

First, we employ archetypal analysis to extract significant patterns from the training data. For convenience, when the meaning is clear, we overload the notation A to indicate both the resulting set of $p < d$ archetypes $\{a_1, a_2, \dots, a_p\} \subset \mathbb{R}^d$ and the matrix constructed from these archetypes $A \in \mathbb{R}^{d \times p}$. These archetypes define an archetype, or pattern, space in p dimensions. Given archetypes A , let the projection function that determines the pattern representation (i.e., the coefficients minimizing objective equation 1) be $c_A : \mathbb{R}^d \rightarrow \mathbb{R}^p$. For any point x_t , the reconstruction is $\hat{x}_t = Ac_A(x_t)$. We denote the error of this reconstruction as $L_{c_A}(x_t) = \|x_t - Ac_A(x_t)\|$.

Next, we train a lightweight neural network $f_A : \mathbb{R}^{p \times T} \rightarrow \mathbb{R}^{p \times H}$ to predict H future pattern representations based on T past pattern representations. The error of this prediction function is $L_{f_A}(c_{1:T}) = \|Ac_{T:T'} - Af_A(c_{1:T})\|$.

The guidance function $f_G : \mathbb{R}^{d \times T} \rightarrow \mathbb{R}^{d \times H}$ is

$$f_G(x_{1:T}) = Af_A \circ c_A(x_{1:T}). \quad (2)$$

Note that $f_G(x_{1:T})$ is $\hat{P}_{T:T'}$ in Figure 2a. We now show that a lower bound on the guidance function error L_{f_G} is a function of the projection error L_{c_A} and pattern prediction error L_{f_A} .

Theorem 1 (Bound on Guidance Function Error). *For any sequence $x_{1:T}$ and horizon $x_{T:T'}$ pair, let the error of guidance function f_G defined in Equation equation 2 be $L_{f_G}(x_{1:T}) =$*

216 $\|x_{T:T'} - Af_A(c_A(x_{1:T}))\|$. Then

$$217 \quad L_{f_G}(x_{1:T}) \geq L_{c_A}(x_{T:T'}) - L_{f_A}(x_{1:T}). \quad (3)$$

218
219 *Proof.* Please see Appendix A for the proof. \square

220
221 It is clear from Theorem 3 that, if the prediction function f_A has a reasonably low error, then the error
222 of the projection function c_A can serve as an approximate lower bound for the error of the guidance
223 function f_G . Therefore, an estimate of L_{c_A} may quantify the degree to which PGDM should "trust"
224 the guidance function, allowing the level of pattern guidance to be dynamically tuned. Next, we
225 introduce a novel uncertainty quantification technique that can be used as a proxy for L_{c_A} .
226

227 4.2 UNCERTAINTY QUANTIFICATION FOR ARCHETYPAL ANALYSIS

228
229 Projection to the lower-dimensional archetype space may lose information. We therefore introduce a
230 novel uncertainty quantification metric based on archetypal analysis that captures this loss.

231 **Definition 1** (Archetypal Analysis Uncertainty Quantification). For any sequence $x_{1:T}$ and archetype
232 set A , the uncertainty u_A of the archetype projection is

$$233 \quad u_A(x_{1:T}) = \frac{1}{T} \sum_{t=1}^T \|x_t - Ac_A(x_t)\|. \quad (4)$$

234
235 Intuitively, this Archetypal Analysis Uncertainty Quantification (AAUQ) metric is simply the average
236 reconstruction loss of the history sequence. AAUQ can also be geometrically interpreted as estimating
237 the average distance of $x_{1:T}$ from the training dataset.

238 **Theorem 2** (AAUQ as Geometric Distance). Assume that a set of archetypes $A = \{a_j\}_{j=1}^p$ is
239 extracted from a dataset D . Define d as the closest point in $\text{Conv}D$ to x_t . For any x_t ,

$$240 \quad u_A(x_t) - \delta \leq \text{dist}(x_t, \text{Conv}D) \leq u_A(x_t) + \delta, \quad (5)$$

241 where $\delta = \|Ac_A(x_t) - d\|$ and $\delta = 0$ when $p = n$.

242
243 *Proof.* Please see Appendix B for the proof. \square

244
245 **Remark 1.** In Theorem 2, δ captures the ability of the archetypes to express the dataset D . This
246 can be seen in the proof of Theorem 2. This can also be understood from the example in Figure 1b,
247 in which $\text{dist}(x, \text{Conv}A)$ is exactly $u_A(x)$, and δ is the distance between \hat{x} and the point on the
248 boundary of $\text{Conv}D$. This has interesting implications for the geometric interpretation of AAUQ
249 in the case that the archetypes perfectly reconstruct the data, which occurs when the number of
250 archetypes is equal to the size of the dataset. It is clear from Theorem 2 that if $p = n$,

$$251 \quad u_A(x_t) = \text{dist}(x_t, \text{Conv}D). \quad (6)$$

252 4.3 PATTERN-GUIDED DIFFUSION MODELS

253
254 PGDM predicts future sequences $x_{T:T'}$ conditioned on the pattern prediction from the guidance func-
255 tion. We follow the methodology of classifier-free diffusion guidance, in which the model is trained
256 with conditioning dropout. This effectively learns two denoising models, $\epsilon_\theta(z_{T:T',s}, x_{1:T}, \hat{P}_{T:T'})$
257 and $\epsilon_\theta(z_{T:T',s}, x_{1:T}, \emptyset)$, where \emptyset is a null value and $z_{T:T',s}$ is the sample to be denoised at diffusion
258 step s . Algorithm 1 summarizes the training process of PGDM. For history and horizon sequences
259 sampled in Line 2, the guidance conditioning is set to the pattern prediction or a null value in Line 4.
260 With the loss function on Line 5, the denoising model learns to estimate the noise added to $x_{T:T'}$.
261

262 When generating samples with traditional classifier-free guidance, each denoising step uses a linear
263 combination of the conditional and unconditional predictions:

$$264 \quad \hat{\epsilon}_\theta(z_{T:T',s}, x_{1:T}) = w\epsilon_\theta(z_{T:T',s}, x_{1:T}, \hat{P}_{T:T'}) + (1-w)\epsilon_\theta(z_{T:T',s}, x_{1:T}, \emptyset),$$

265 where $w \geq 0$ is the *guidance scale*. When $w = 0$, generation is unguided and sampled data are
266 more diverse. The guidance level increases with w , leading to less diverse but higher quality samples.
267 While traditional classifier-free guidance is performed with constant guidance scale w , we instead
268 use a dynamic guidance scale $w(x_{1:T})$ that captures the trustworthiness of the guidance function:
269

$$269 \quad \hat{\epsilon}_\theta(z_{T:T',s}, x_{1:T}) = w(x_{1:T})\epsilon_\theta(z_{T:T',s}, x_{1:T}, \hat{P}_{T:T'}) + (1-w(x_{1:T}))\epsilon_\theta(z_{T:T',s}, x_{1:T}, \emptyset), \quad (7)$$

Algorithm 1 PGDM Training

```

1: for all epochs do
2:   Sample  $x_{1:T}, x_{T:T'}$  from the training set
3:    $s \sim \text{Uniform}(1, \dots, S), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
4:    $\hat{P}_{T:T'} = \emptyset$  with probability  $p_{\text{drop}}$ , else  $\hat{P}_{T:T'} = f_G(x_{1:T})$  from eqn. equation 2
5:   Take gradient descent step on  $\nabla_{\theta} \|\epsilon - \epsilon_{\theta}(\sqrt{\alpha_s}x_{T:T'} + \sqrt{1 - \alpha_s}\epsilon, x_{1:T}, \hat{P}_{T:T'})\|^2$ 
6: end for

```

Algorithm 2 PGDM Inference

```

1: given  $x_{1:T}$ 
2:  $x_{T:T',S} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
3:  $\hat{P}_{T:T'} = f_G(x_{1:T})$  from eqn. equation 2
4: for  $s = S, \dots, 1$  do
5:    $n \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  if  $s > 1$ , else  $n = 0$ 
6:   Compute  $\hat{\epsilon}_{\theta}(z_{T:T',s}, x_{1:T})$  from eqn. equation 7
7:   Sample  $x_{T:T',s-1} = \frac{1}{\sqrt{\alpha_s}} \left( x_{T:T',s} - \frac{1-\alpha_s}{\sqrt{1-\alpha_s}} \right) \hat{\epsilon}_{\theta}(z_{T:T',s}, x_{1:T}) + \sqrt{\beta_s}n$ 
8: end for
9: Compute  $w^* = w(x_{1:T}, \bar{w}^*, \gamma)$  from eqn. equation 8, for  $\bar{w}^* \in [0, 1]$ 
10: Mix  $\hat{x}_{T:T'} = w^* \hat{P}_{T:T'} + (1 - w^*)x_{T:T',0}$ 

```

Definition 2 (Dynamic Guidance Scale). *Let A be the set of archetypes extracted from the training dataset. Given a maximum guidance scale \bar{w} and maximum tolerable uncertainty γ , the dynamic guidance scale for sequence $x_{1:T}$ is*

$$w(x_{1:T}, \bar{w}, \gamma) = \text{ReLU}\left(-\frac{\bar{w}}{\gamma} u_A(x_{1:T}) + \bar{w}\right). \quad (8)$$

Here, we measure the trustworthiness of our guidance function by our AAUQ uncertainty metric. The uncertainty $u_A(x_{1:T})$ is easy to compute and in practice, we find that the $u_A(x_{1:T})$ is a good proxy for estimating $L_{c_A}(x_{T:T'})$ and therefore the lower bound in Theorem 1. Hence, we design the dynamic guidance scale so that, as uncertainty increases, the guidance scale $w \in (0, \bar{w})$ decreases. When the uncertainty exceeds γ , $w = 0$. In other words, PGDM follows the pattern guidance most strictly when the data is in-distribution. For out-of-distribution data with unseen patterns, PGDM relies less on pattern-guidance, reverting to a standard diffusion model in the extreme case.

Algorithm 2 summarizes the inference process of PGDM. In Line 2, noisy data is sampled from a standard normal distribution. In Line 3, the patterns of the future sequence are predicted. Then, over S reverse diffusion steps in Lines 4–7, the guided and unguided denoising models are combined with dynamic guidance scale to iteratively remove noise. The resulting sequence prediction is sampled from the learned distribution, which we expect to be tightly centered around the pattern prediction. Finally in Lines 9 and 10, we mix the raw pattern prediction with the pattern-guided sequence prediction using a dynamic mixing scale with maximum value \bar{w}^* . We include this mixing step to mitigate some potential practical challenges of PGDM (see Appendix C for a full discussion).

5 APPLICATIONS

We validate PGDM on two applications, visual field prediction and human motion prediction. To help demonstrate the benefits of pattern guidance, we show the performance of two PGDM models selected from our hyperparameter search. The first model, PGDM_{MAE}, achieved the lowest validation mean absolute error (MAE) with pattern guidance. The second, PGDM_{GDE}, achieved the highest capacity for pattern guidance, or the greatest achievable improvement in MAE by using guided predictions ($\bar{w} > 0$) over unguided predictions ($\bar{w} = 0$). That is, we selected PGDM_{GDE} as the model that achieved the largest gain in MAE by adding any level of guidance in the range of 1 to 5.

For these two models, we compare performance with guidance to performance without guidance and multiple baselines. While we would ideally compare PGDM to baselines that use some pattern

conditioning (Wang et al., 2024; Westny et al., 2024; Zhao et al., 2024), most existing methods formalize patterns in a manner that does not translate to our applications. For others, we were unable to obtain sufficient implementation details or code. Therefore, we instead select more general baseline techniques that do not use pattern conditioning. We compare PGDM to multiple diffusion-based baselines for forecasting, TimeGrad (Rasul et al., 2021), CSDI (Tashiro et al., 2021), and ARMD (Gao et al., 2025). While CSDI is a data imputation technique, it can easily be extended for forecasting. For the visual field application, we further compare PGDM to GenViT (Yang et al., 2022), a diffusion model with a vision transformer backbone that has been applied to VF prediction by Tian et al. (2023). Finally, we perform multiple ablations, one serving as a proxy for evaluating Westny et al. (2024).

We evaluate all models using mean absolute error (MAE). Additionally, to evaluate PGDM as a probabilistic forecaster, we measure the continuous ranked probability score (CRPS). Following Rasul et al. (2021) and Tashiro et al. (2021), we report CRPS_{SUM} , computed as the CRPS over the summed dimensions of the data.

Source code is supplied in the supplementary material. Complete implementation details for pattern extraction and training, including hyperparameter selection, model architectures, compute resources, parameter counts, and inference latency are provided in Appendix D.

Visual Field Prediction. Pattern-Guided Diffusion models are especially useful in medical settings, where data often reflects consistent patterns due to anatomy. For instance, 24-2 visual field (VF) tests measure light sensitivity in decibels (dB) at 52 central points of vision, with specific loss patterns linked to structural eye damage (e.g., nerve fiber bundle loss) (Keltner et al., 2003). Figure 1a illustrates archetypal patterns from VF data, where darker areas indicate reduced vision. Forecasting VF outcomes can support clinicians in diagnosis, progression identification, and treatment planning.

We evaluate PGDM on the public the University of Washington Humphrey Visual Field (UWHVF) dataset (Montesano et al., 2022). To the best of our knowledge, UWHVF is the only publicly available 24-2 VF dataset, containing 7,428 sequences from 3,871 patients. The UWHVF measurements capture the patient’s light sensitivity compared to normative data, ranging from -38 dB to 50 dB. That is, a negative (positive) dB indicates worse (better) vision than typical. Due to the few follow-up visits per patient, we predict $H = 1$ step into the future based on the past $T = 3$ steps. Additionally, as VF measurements are taken at non-constant time increments, we also condition predictions on the recorded age at each VF measurement and the desired time horizon for prediction. Thus, the one-step-ahead prediction can be made for an arbitrary length time period. We create multiple forecasting sequences from each patient in a sliding window fashion, resulting in 6,171 sequences.

Human Motion Prediction. We also apply PGDM to predict future frames of human motion capture data, a task relevant to domains including human robot interaction and autonomous driving (Lyu et al., 2022). Human motion often involves repeated body positions when executing common movements (e.g., walking, running, dancing). While our visual field application demonstrates PGDM’s utility in the clinical domain, the motion prediction application presents a more challenging task with more rapidly evolving signals and longer prediction horizons.

We evaluate PGDM on the AIST++ dataset (Li et al., 2021), which contains motion capture frames capturing 3D motion from 10 dance genres. Predicting dance motion is more challenging for PGDM compared to locomotion, as the variety of dance genres and styles leads to a rich set of patterns with less periodic progressions over time. The AIST++ dataset captures the skeleton with 3D pose data for 17 keypoints, which represent specific joints or locations in the body. For consistency across heights, we normalize all data to the scale $[0, 100]$. We predict $H = 5$ steps into the future based on a past sequence of $T = 3$ steps. We create multiple forecasting sequences from each motion capture video, resulting between 36,739 and 105,504 sequences across the 10 genres.

5.1 RESULTS

Pattern Extraction. We extract $p = 13$ archetypal patterns from UWHVF, shown in Figure 1a. We extract between $p = 12$ and $p = 22$ archetypes for each genre of AIST++. Figure 3 shows the archetypes extracted from the break dancing frames (see Appendix E for remaining genres). Appendix F reports the reconstruction error of the extracted patterns and the guidance function error.

AAUQ approximately lower bounds the guidance function error. Motivated by Theorem 1, PGDM uses AAUQ to determine the appropriate level of guidance. In Figure 4, we compare AAUQ

378
379
380
381
382
383
384
385
386

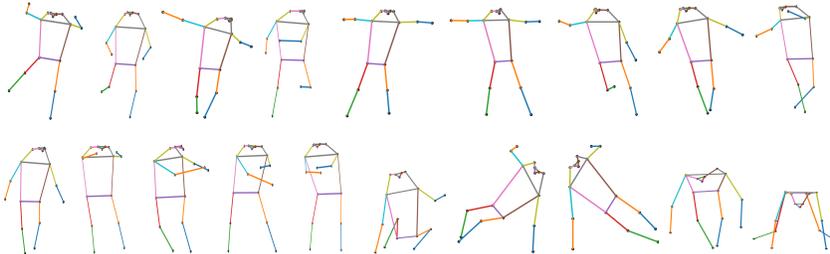
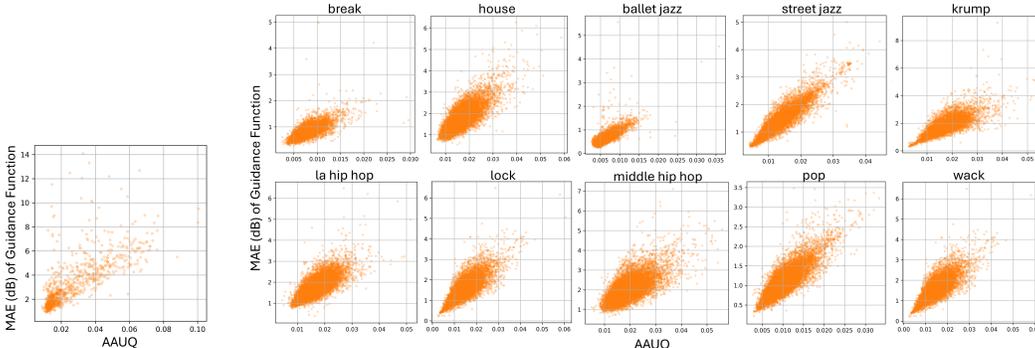


Figure 3: Nineteen archetypes extracted from AIST++ break dancing frames.

387
388
389
390
391
392
393
394
395
396
397
398
399



(a) VF prediction.

(b) Motion prediction.

Figure 4: AAUQ approximately lower bounds the guidance function error for both applications.

402
403

measurements to the MAE of the guidance function. In both applications, we observe that AAUQ is indeed proportional to a linear lower bound on the guidance function error.

404
405
406
407
408
409
410
411
412
413
414
415
416
417

Pattern guidance improves predictions. Tables 1 and 2 report the MAE and CRPS of PGDM_{MAE} and PGDM_{GDE} with and without pattern guidance. Compared to unguided predictions, pattern guidance reduces prediction error and CRPS significantly. In particular, on UWHVF, guidance reduces the MAE / CRPS of PGDM_{GDE} by up to 40.67% / 56.26%. On AIST++, guidance achieves a reduction of up to 14.12% / 14.10%. To further study the impact of pattern guidance, we also show in Table 3 the MAE of PGDM_{MAE} and PGDM_{GDE} on UWHVF with $\bar{w} = 1, 2, 3, 4, 5$. Similar results on AIST++ and with CRPS, as well as qualitative examples, are shown in Appendix G. In general, the standard deviation of MAE decreases as the guidance scale increases towards the optimal \bar{w} choice, indicating that guidance improves both prediction consistency and quality. Notably, we also find that excessive pattern guidance may lead to diminishing returns. In practice, an appropriate \bar{w} may be selected by a hyperparameter search.

418
419
420
421
422
423
424
425
426
427
428
429

PGDM outperforms baselines. Tables 1 and 2 also compare the performance of PGDM_{MAE} and PGDM_{GDE} to our baselines. Across the board, PGDM_{MAE} with pattern guidance achieves significantly lower MAE than baselines. On UWHVF, PGDM_{MAE} with guidance surpasses GenViT, TimeGrad, CSDI, and ARMD by up to 65.68%, 29.36%, 7.20%, and 24.30%, respectively. On AIST++, PGDM_{MAE} surpasses TimeGrad, CSDI, and ARMD by up to 82.60%, 36.71%, 93.64%. In terms of CRPS, on AIST++, PGDM_{MAE} with guidance outperforms TimeGrad, CSDI, and ARMD by up to 82.87%, 40.86%, and 92.55%, respectively. On UWHVF, PGDM outperforms GenViT and ARMD by up to 84.83% and 15.38%, but does not outperform TimeGrad and CSDI. We note that all models achieve significantly higher CRPS on UWHVF than AIST++, indicating that this small dataset provides limited signal for learning well-calibrated predictive distributions. We emphasize that PGDM achieves the lowest MAE on UWHVF, demonstrating that explicit pattern modeling improves point prediction even when the distribution itself is challenging to estimate.

430
431

PGDM can outperform baselines with guidance alone. We note that PGDM_{MAE} achieves lower MAE than baselines even without guidance ($\bar{w} = 0$) on some AIST++ genres (e.g., break dancing). In these cases, to better demonstrate that PGDM’s performance comes from pattern guidance, rather than

Table 1: Mean absolute error of PGDM_{MAE}, PGDM_{GDE}, and baselines for both the visual field prediction (UWHVF dataset) and the human motion prediction (10 dance genres from AIST++ dataset) applications. For PGDM_{MAE} and PGDM_{GDE} with $\bar{w} > 0$, we show the guidance and mixing scale result that achieved the lowest error (for \bar{w} and w^* choices, see Appendix D). Mean and standard deviation are taken across five samples.

Visual Field Prediction							
GenViT	TimeGrad	CSDI	ARMD	PGDM _{GDE} $\bar{w} = 0$	PGDM _{GDE} $\bar{w} > 0$	PGDM _{MAE} $\bar{w} = 0$	PGDM _{MAE} $\bar{w} > 0$
8.61±0.0018	4.19±0.0327	3.19±0.0421	3.91±0.0083	5.20±0.0407	3.08±0.0153	3.75±0.0437	2.96±0.0117

Human Motion Prediction							
Genre	TimeGrad	CSDI	ARMD	PGDM _{GDE} $\bar{w} = 0$	PGDM _{GDE} $\bar{w} > 0$	PGDM _{MAE} $\bar{w} = 0$	PGDM _{MAE} $\bar{w} > 0$
Break	2.10±0.0064	0.47±0.0010	4.26±0.0026	0.52±0.0032	0.45±0.0015	0.41±0.0013	0.38±0.0009
House	3.71±0.0159	1.02±0.0045	9.11±0.0035	0.90±0.0013	0.82±0.0014	0.79±0.0017	0.74±0.0014
Ballet Jazz	1.32±0.0098	0.55±0.0054	3.38±0.0011	0.49±0.0004	0.42±0.0009	0.42±0.0008	0.38±0.0002
Street Jazz	1.65±0.0102	0.56±0.0054	7.57±0.0049	0.60±0.0016	0.54±0.0005	0.52±0.0017	0.48±0.0008
Krump	2.37±0.0067	0.77±0.0017	8.90±0.0046	0.88±0.0016	0.79±0.0005	0.77±0.0016	0.70±0.0013
LA Hip Hop	3.30±0.0157	0.78±0.0023	8.51±0.0032	0.90±0.0009	0.81±0.0009	0.80±0.0010	0.74±0.0006
Lock	3.03±0.0086	0.76±0.0028	7.65±0.0031	0.78±0.0017	0.71±0.0023	0.72±0.0011	0.67±0.0006
Mid. Hip Hop	3.35±0.0113	1.04±0.0048	8.98±0.0051	1.05±0.0034	0.95±0.0038	0.88±0.0014	0.82±0.0009
Pop	2.55±0.0105	0.70±0.0053	6.43±0.0019	0.52±0.0013	0.48±0.0018	0.47±0.0008	0.44±0.0018
Wack	1.03±0.0047	0.44±0.0042	3.39±0.0026	0.49±0.0054	0.45±0.0079	0.44±0.0044	0.39±0.0028

Table 2: CRPS_{SUM} (lower is better) of PGDM_{MAE}, PGDM_{GDE}, and baselines for both the visual field prediction (UWHVF dataset) and the human motion prediction (10 dance genres from AIST++ dataset) applications. For PGDM_{MAE} and PGDM_{GDE} with $\bar{w} > 0$, we show the guidance and mixing scale result that achieved the lowest error (for \bar{w} and w^* choices, see Appendix D). Mean and standard deviation are taken across three seeds.

Visual Field Prediction							
GenViT	TimeGrad	CSDI	ARMD	PGDM _{GDE} $\bar{w} = 0$	PGDM _{GDE} $\bar{w} > 0$	PGDM _{MAE} $\bar{w} = 0$	PGDM _{MAE} $\bar{w} > 0$
5.119±0.0004	0.751±0.0020	0.658±0.0025	0.918±0.0029	1.887±0.0119	0.825±0.0032	0.794±0.0122	0.777±0.0040

Human Motion Prediction							
Genre	TimeGrad	CSDI	ARMD	PGDM _{GDE} $\bar{w} = 0$	PGDM _{GDE} $\bar{w} > 0$	PGDM _{MAE} $\bar{w} = 0$	PGDM _{MAE} $\bar{w} > 0$
Break	0.189±0.0010	0.051±0.0001	0.490±0.0006	0.049±0.0001	0.043±0.0001	0.039±0.0001	0.037±<0.0001
House	0.340±0.0011	0.106±0.0002	0.679±0.0008	0.085±0.0005	0.079±0.0003	0.077±0.0004	0.073±0.0002
Ballet Jazz	0.117±0.0005	0.053±0.0002	0.346±0.0005	0.044±0.0002	0.039±0.0002	0.039±0.0001	0.035±0.0002
Street Jazz	0.184±0.0002	0.058±0.0003	0.448±0.0004	0.055±0.0002	0.052±0.0001	0.050±0.0002	0.047±0.0002
Krump	0.298±0.0010	0.093±0.0003	0.911±0.0003	0.091±0.0004	0.078±0.0003	0.079±0.0003	0.071±0.0003
LA Hip Hop	0.331±0.0004	0.088±0.0002	0.752±0.0014	0.083±0.0007	0.077±0.0005	0.075±0.0006	0.071±0.0005
Lock	0.367±0.0004	0.080±0.0002	0.652±0.0014	0.073±0.0001	0.066±0.0001	0.068±0.0002	0.063±0.0001
Mid. Hip Hop	0.368±0.0018	0.109±0.0002	0.832±0.0005	0.097±0.0003	0.089±0.0003	0.085±0.0002	0.081±0.0003
Pop	0.252±0.0004	0.076±0.0005	0.460±0.0001	0.051±0.0002	0.047±0.0002	0.046±0.0002	0.045±0.0002
Wack	0.112±0.0005	0.051±0.0002	0.400±0.0012	0.049±0.0002	0.044±0.0002	0.043±0.0002	0.039±0.0003

model training or architecture design alone, we also emphasize the results for PGDM_{GDE} in Tables 1 and 2. Even when unguided PGDM_{GDE} performs worse than baselines, pattern guidance almost always reduces the error of PGDM_{GDE} to a lower or competitive level compared to baselines. These results demonstrate that pattern guidance is essential for higher quality predictions.

5.2 ABLATIONS

For all ablations, we report full results with additional discussion in Appendix H.

Table 3: Mean absolute error (MAE) of PGDM_{MAE} and PGDM_{GDE} on the VF prediction application (UWHVF dataset) with varying levels of guidance. Percent improvements over baselines are shown in the Δ MAE (%) columns. Mean and standard deviation are taken across five samples.

Model		MAE (dB)	Δ MAE (%) vs. GenViT	Δ MAE (%) vs. TimeGrad	Δ MAE (%) vs. CSDI	Δ MAE (%) vs. ARMD	Δ MAE (%) vs. $\bar{w} = 0$
PGDM _{MAE}	$\bar{w} = 0$	3.75 \pm 0.0437	56.48 \pm 0.50	10.69 \pm 1.17	-17.32 \pm 1.55	4.29 \pm 1.21	-
	$\bar{w} = 1$	2.96\pm0.0117	65.58\pm0.14	29.36\pm0.67	7.20\pm1.25	24.30\pm0.31	20.90 \pm 0.71
	$\bar{w} = 2$	2.97 \pm 0.0112	65.54 \pm 0.13	29.29 \pm 0.67	7.10 \pm 1.27	24.22 \pm 0.29	20.81 \pm 0.73
	$\bar{w} = 3$	2.97 \pm 0.0108	65.51 \pm 0.13	29.23 \pm 0.67	7.03 \pm 1.29	24.15 \pm 0.26	20.75 \pm 0.76
	$\bar{w} = 4$	2.97 \pm 0.0107	65.48 \pm 0.13	29.17 \pm 0.67	6.95 \pm 1.31	24.09 \pm 0.09	20.68 \pm 0.78
	$\bar{w} = 5$	2.97 \pm 0.0104	65.45 \pm 0.12	29.10 \pm 0.66	6.86 \pm 1.32	24.02 \pm 0.02	20.60 \pm 0.80
PGDM _{GDE}	$\bar{w} = 0$	5.20 \pm 0.0407	39.60 \pm 0.47	-23.95 \pm 1.10	-62.83 \pm 2.59	-32.84 \pm 1.05	-
	$\bar{w} = 1$	3.16 \pm 0.0212	63.28 \pm 0.24	24.65 \pm 0.76	1.01 \pm 1.44	19.25 \pm 0.58	39.21 \pm 0.23
	$\bar{w} = 2$	3.13 \pm 0.0201	63.67 \pm 0.23	25.44 \pm 0.76	2.06 \pm 1.43	20.10 \pm 0.55	39.85 \pm 0.23
	$\bar{w} = 3$	3.10 \pm 0.0187	63.93 \pm 0.21	25.99 \pm 0.74	2.77 \pm 1.41	20.68 \pm 0.51	40.29 \pm 0.24
	$\bar{w} = 4$	3.09 \pm 0.0170	64.09 \pm 0.20	26.32 \pm 0.72	3.20 \pm 1.39	21.03 \pm 0.47	40.55 \pm 0.25
	$\bar{w} = 5$	3.08 \pm 0.0153	64.16 \pm 0.18	26.47 \pm 0.70	3.40 \pm 1.39	21.19 \pm 0.43	40.67\pm0.27

Impact of pattern mixing. We study the impact of pattern mixing on the MAE and CRPS_{SUM} of PGDM by varying the maximum mixing scale \bar{w}^* from 0.0 to 1.0 while keeping the maximum guidance scale \bar{w} fixed. On UWHVF, we find that a strong mixing signal substantially improves performance. In contrast, on AIST++, a weaker mixing scale is preferable, and an excessive mixing scale can degrade performance. This indicates that the pattern prediction model itself makes more accurate predictions on the small UWHVF dataset, even without the downstream denoiser. This suggests that explicit pattern modeling is particularly beneficial in low-data regimes.

Impact of dynamic scaling. We additionally perform ablations to demonstrate the benefits of dynamic scaling. We compare the MAE and CRPS_{SUM} of PGDM with a dynamic scale to that with a constant scale, holding \bar{w} and \bar{w}^* fixed. We find that dynamic scaling consistently performs comparably or better than constant scaling. Dynamic scaling likely yields the greatest improvements over constant scaling when novel patterns are seen inference time.

Impact of pattern prediction. Finally, we compare PGDM with and without its pattern prediction model f_A . PGDM *without* pattern prediction inherently follows the assumption that patterns remain constant over time. This ablation therefore serves as a proxy for Diff-MGR (Zhao et al., 2024), which follows this same assumption but cannot be evaluated directly due to unavailable code and implementation details. PGDM with pattern prediction consistently outperforms PGDM without pattern prediction, most significantly when patterns change rapidly over time. This result highlights the importance of the pattern prediction model in accounting for the highly dynamic patterns that commonly appear in realistic temporal data.

6 CONCLUSIONS

In this paper, we proposed Pattern-Guided Diffusion Models (PGDM), which leverage inherent archetypal patterns to forecast future steps from multivariate time series data. PGDM is guided by a pattern guidance function that predicts future patterns within the data. To estimate the trustworthiness of this guidance function, we introduced a novel uncertainty quantification metric that approximately lower bounds the guidance function error. Finally, we proposed to dynamically tune the level to which PGDM follows the pattern guidance based on this uncertainty metric. We found that PGDM outperforms baseline models, and pattern guidance improves the prediction quality of PGDM. Two limitations of PGDM present interesting avenues for future work. First, PGDM has less benefit for out-of-distribution data exhibiting unseen patterns. Second, the use of AAUQ as an approximate lower-bound for guidance function error assumes that the temporal data is relatively continuous and does not rapidly change between the observed history and target prediction window. In some cases, this assumption is violated (e.g., rapidly changing signals sampled with a low frequency). Based on these limitations, PGDM may be further improved by updating the set of extracted patterns at inference time and accounting for signal dynamics when calculating the guidance scale.

REFERENCES

- 540
541
542 Ping Chang, Huayu Li, Stuart F Quan, Shuyang Lu, Shu-Fen Wung, Janet Roveda, and Ao Li. A
543 transformer-based diffusion probabilistic model for heart rate and blood pressure forecasting in
544 intensive care unit. *Computer Methods and Programs in Biomedicine*, 246:108060, 2024.
- 545 Adele Cutler and Leo Breiman. Archetypal analysis. *Technometrics*, 36(4):338–347, 1994.
- 546
547 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances*
548 *in neural information processing systems*, 34:8780–8794, 2021.
- 549 Tobias Elze, Louis R Pasquale, Lucy Q Shen, Teresa C Chen, Janey L Wiggs, and Peter J Bex.
550 Patterns of functional vision loss in glaucoma determined with archetypal analysis. *Journal of The*
551 *Royal Society Interface*, 12(103):20141118, 2015.
- 552
553 Shibo Feng, Chunyan Miao, Zhong Zhang, and Peilin Zhao. Latent diffusion transformer for proba-
554 bilistic time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
555 volume 38, pp. 11979–11987, 2024.
- 556 Jiaxin Gao, Qinglong Cao, and Yuntian Chen. Auto-regressive moving diffusion models for time
557 series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39,
558 pp. 16727–16735, 2025.
- 559 Tianpei Gu, Guangyi Chen, Junlong Li, Chunze Lin, Yongming Rao, Jie Zhou, and Jiwen Lu. Stochas-
560 tic trajectory prediction via motion indeterminacy diffusion. In *Proceedings of the IEEE/CVF*
561 *conference on computer vision and pattern recognition*, pp. 17113–17122, 2022.
- 562
563 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*,
564 2022.
- 565
566 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*
567 *neural information processing systems*, 33:6840–6851, 2020.
- 568 Junfeng Hu, Xu Liu, Zhencheng Fan, Yuxuan Liang, and Roger Zimmermann. Towards unifying
569 diffusion models for probabilistic spatio-temporal graph learning. In *Proceedings of the 32nd ACM*
570 *International Conference on Advances in Geographic Information Systems*, pp. 135–146, 2024.
- 571
572 John L Keltner, Chris A Johnson, Kimberly E Cello, Mary A Edwards, Shannan E Bandermann,
573 Michael A Kass, Mae O Gordon, Ocular Hypertension Treatment Study Group, et al. Classification
574 of visual field abnormalities in the ocular hypertension treatment study. *Archives of Ophthalmology*,
575 121(5):643–650, 2003.
- 576
577 Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned
578 3d dance generation with aist++. In *Proceedings of the IEEE/CVF international conference on*
computer vision, pp. 13401–13412, 2021.
- 579 Yan Li, Xinjiang Lu, Yaqing Wang, and Dejing Dou. Generative time series forecasting with
580 diffusion, denoise, and disentanglement. *Advances in Neural Information Processing Systems*, 35:
581 23009–23022, 2022.
- 582
583 Kai Lv, Liang Yuan, and Xiaoyu Ni. Learning autoencoder diffusion models of pedestrian group
584 relationships for multimodal trajectory prediction. *IEEE Transactions on Instrumentation and*
585 *Measurement*, 2024.
- 586
587 Kedi Lyu, Haipeng Chen, Zhenguang Liu, Beiqi Zhang, and Ruili Wang. 3d human motion prediction:
588 A survey. *Neurocomputing*, 489:345–365, 2022.
- 589
590 Giovanni Montesano, Andrew Chen, Randy Lu, Cecilia S Lee, and Aaron Y Lee. Uwhvf: a real-
591 world, open source dataset of perimetry tests from the humphrey field analyzer at the university of
592 washington. *Translational Vision Science & Technology*, 11(1):2–2, 2022.
- 593
Kashif Rasul, Calvin Seward, Ingmar Schuster, and Roland Vollgraf. Autoregressive denoising
diffusion models for multivariate probabilistic time series forecasting. In *International conference*
on machine learning, pp. 8857–8868. PMLR, 2021.

- 594 Luke J Saunders, Felipe A Medeiros, Robert N Weinreb, and Linda M Zangwill. What rates of
595 glaucoma progression are clinically significant? *Expert review of ophthalmology*, 11(3):227–234,
596 2016.
- 597 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
598 learning using nonequilibrium thermodynamics. In *International conference on machine learning*,
599 pp. 2256–2265. pmlr, 2015.
- 600 Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. Csd: Conditional score-based diffu-
601 sion models for probabilistic time series imputation. *Advances in neural information processing*
602 *systems*, 34:24804–24816, 2021.
- 603 Ye Tian, Mingyang Zang, Anurag Sharma, Sophie Z Gu, Ari Leshno, and Kaveri A Thakoor.
604 Glaucoma progression detection and humphrey visual field prediction using discriminative and
605 generative vision transformers. In *International Workshop on Ophthalmic Medical Image Analysis*,
606 pp. 62–71. Springer, 2023.
- 607 Weizhuo Wang, C Karen Liu, and Monroe Kennedy III. Egonav: Egocentric scene-aware human
608 trajectory prediction. *arXiv preprint arXiv:2403.19026*, 2024.
- 609 Haomin Wen, Youfang Lin, Yutong Xia, Huaiyu Wan, Qingsong Wen, Roger Zimmermann, and
610 Yuxuan Liang. Diffstg: Probabilistic spatio-temporal graph forecasting with denoising diffusion
611 models. In *Proceedings of the 31st ACM International Conference on Advances in Geographic*
612 *Information Systems*, pp. 1–12, 2023.
- 613 Theodor Westny, Björn Olofsson, and Erik Frisk. Diffusion-based environment-aware trajectory
614 prediction. *arXiv preprint arXiv:2403.11643*, 2024.
- 615 Xiulong Yang, Sheng-Min Shih, Yinlin Fu, Xiaoting Zhao, and Shihao Ji. Your vit is secretly a hybrid
616 discriminative-generative diffusion model. *arXiv preprint arXiv:2208.07791*, 2022.
- 617 Yiyuan Yang, Ming Jin, Haomin Wen, Chaoli Zhang, Yuxuan Liang, Lintao Ma, Yi Wang, Chenghao
618 Liu, Bin Yang, Zenglin Xu, et al. A survey on diffusion models for time series and spatio-temporal
619 data. *arXiv preprint arXiv:2404.18886*, 2024.
- 620 Tianlong Zhao, Guangle Song, Xuemei Li, Lizhen Cui, and Caiming Zhang. Diff-mgr: Dynamic
621 causal graph attention and pattern reproduction guided diffusion model for multivariate time series
622 probabilistic forecasting. *Information Sciences*, 675:120742, 2024.

623 A PROOF OF THEOREM 1

624 For convenience, let $X = x_{1:T}$ and $Y = x_{T:T'}$. Then we have

$$\begin{aligned}
 625 L_{f_G}(X) &= \|Y - Af_A(c_A(X))\| \\
 626 &= \|Y - Af_A(c_A(X)) + Ac_A(Y) - Ac_A(Y)\| \\
 627 &\geq \|Y - Ac_A(Y)\| - \|Af_A(c_A(X)) - Ac_A(Y)\| \\
 628 &= L_{c_A}(Y) - L_{f_A}(X).
 \end{aligned}$$

629 The inequality above follows from the reverse triangle inequality.

630 B PROOF OF THEOREM 2

631 First, observe that $u_A(x_t)$ is the distance between x_t and the set $\text{Conv}A$. This can be seen by noting
632 that

$$633 c_A(x_t) = \arg \min_c \|x_t - Ac\| = \arg \min_{\bar{x} \in \text{Conv}A} \|x_t - \bar{x}\|, \quad (9)$$

634 where $c \in \mathbb{R}^p$ has positive elements that sum to one. Let $\hat{x} = Ac_A(x_t)$. Recall from Definition 1 that
635 $u_A(x_t) = \|x_t - \hat{x}\|$.

Similarly, $\text{dist}(x_t, \text{Conv}D)$ is defined as

$$\text{dist}(x_t, \text{Conv}D) = \min_{\bar{d} \in \text{Conv}D} \|x_t - \bar{d}\|.$$

Let d be such that $\text{dist}(x_t, \text{Conv}D) = \|x_t - d\|$.

The remainder of the proof follows from a straightforward application of the reverse triangle inequality:

$$\begin{aligned} \|\hat{x} - d\| &= \|\hat{x} - d + x_t - x_t\| \\ &\geq \| \|x_t - d\| - \|x_t - \hat{x}\| \|. \end{aligned}$$

Then, we have

$$- \|\hat{x} - d\| \leq \|x_t - d\| - \|x_t - \hat{x}\| \leq \|\hat{x} - d\|.$$

With some rearranging, we arrive at Equation equation 4 by letting $\delta = \|\hat{x} - d\|$.

Now note that if $p = n$, then selecting $A = D$ minimizes the archetypal analysis objective equation 1 (Cutler & Breiman, 1994, Proposition 1) with RSS of 0, and A fully expresses D . Then $\text{Conv}A = \text{Conv}D$ and $\hat{x} = d$. Finally, $\delta = 0$. We briefly remark that δ therefore captures the expressiveness of the archetypes.

C PATTERN MIXING

In Lines 9 and 10 of Algorithm 2, we include an additional pattern mixing step in our sequence prediction process. The final PGDM prediction is a linear combination of the raw pattern prediction from the guidance function and the the pattern-guided output of the diffusion model. We include this step to overcome some of the practical challenges of PGDM. In practice, we find that PGDM’s capacity for pattern guidance is highly dependent on appropriate architecture design. We therefore include pattern mixing as an additional step to overcome this challenge.

While pattern mixing improves prediction quality, pattern guidance is still necessary. Figure 5 illustrates the impact of pattern guidance and pattern mixing. Without guidance, the model may make highly varied predictions that are far from the groundtruth. With pattern guidance, PGDM narrows the distribution of predictions and shifts it towards the ground truth. Pattern mixing further shifts the distribution, without affecting sample diversity. Our results demonstrate exactly this. The unguided PGDM prediction has higher error and variance. With guidance and mixing, the error and standard deviation are significantly reduced, demonstrating that both pattern guidance and pattern mixing aid in improving predictions.

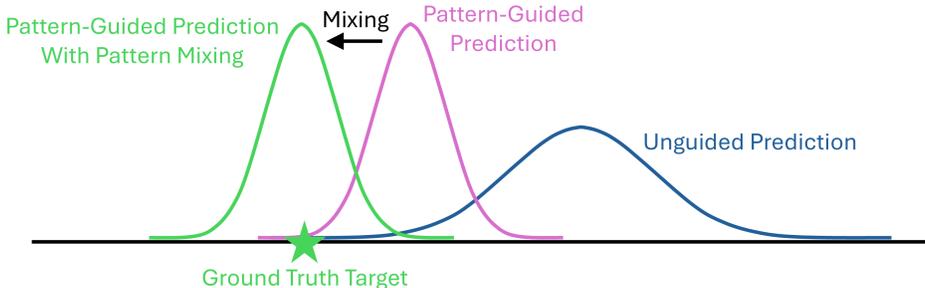


Figure 5: Impacts of pattern guidance and pattern mixing.

D IMPLEMENTATION DETAILS

For both case studies, we split our data into 70% training, 15% validation, and 15% test sets. For the motion capture data, we remove rotations around the vertical axis and supply the isolated rotation angles as additional inputs to PGDM and the baseline models. This normalizes the direction in which the motion capture skeletons are facing, allowing for more straightforward pattern extraction.

Table 4: Hyperparameter choices for pattern prediction model.

	Pattern Prediction Model		PGDM _{MAE}			PGDM _{GDE}		
	Batch Size	LR	Batch Size	LR	Epochs	Batch Size	LR	Epochs
UWHVF	32	1×10^{-4}	32	5×10^{-5}	200	64	1×10^{-5}	100
Break	32	5×10^{-4}	32	1×10^{-3}	300	64	5×10^{-4}	200
House	64	5×10^{-4}	32	1×10^{-3}	300	32	5×10^{-4}	200
Ballet jazz	64	5×10^{-4}	32	1×10^{-3}	300	64	5×10^{-4}	300
Street jazz	64	5×10^{-4}	32	1×10^{-3}	300	64	1×10^{-3}	200
Krump	64	5×10^{-4}	32	1×10^{-3}	300	32	5×10^{-4}	200
LA Hip Hop	64	5×10^{-4}	32	1×10^{-3}	300	32	5×10^{-4}	200
Lock	64	5×10^{-4}	32	1×10^{-3}	300	32	1×10^{-3}	200
Middle Hip Hop	64	5×10^{-4}	32	1×10^{-3}	300	64	5×10^{-4}	200
Pop	32	5×10^{-4}	32	1×10^{-3}	300	32	5×10^{-4}	200
Wack	32	5×10^{-4}	32	1×10^{-3}	300	32	5×10^{-4}	300

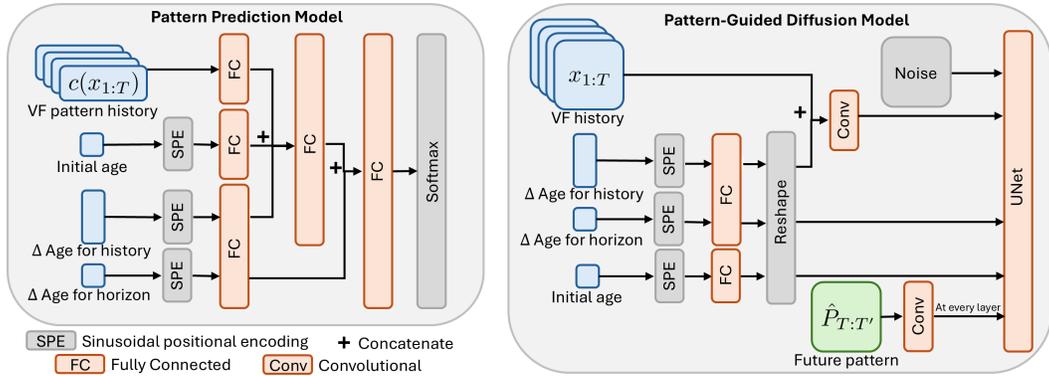


Figure 6: Pattern prediction model and pattern-guided diffusion model for visual field prediction.

Table 5: Selected values of \bar{w} and \bar{w}^* for Table 1.

	PGDM _{MAE}		PGDM _{GDE}	
	\bar{w}	\bar{w}^*	\bar{w}	\bar{w}^*
UWHVF	1	1.0	5	1.0
Break	1	0.2	2	0.2
House	1	0.2	1	0.2
Ballet jazz	1	0.2	1	0.2
Street jazz	2	0.0	2	0.0
Krump	2	0.0	1	0.2
LA Hip Hop	1	0.2	1	0.2
Lock	1	0.2	1	0.2
Middle Hip Hop	1	0.2	1	0.2
Pop	1	0.2	1	0.2
Wack	1	0.2	1	0.2

To train our guidance function, we first extract archetypal patterns from the most recent VF x_T of each sequence in the training set. By extracting archetypes from only a single point in each sequence, we avoid leakage between the training, validation, and test sets. We select the number of archetypes p by a hyperparameter search through $p = 2, \dots, 25$ with selection criterion following Elze et al. (2015). We then train our pattern prediction model (see Figure 6 for architecture) to predict the pattern representation of each sequence. We train the model with the Adam optimizer on a KL-divergence loss function with the hyperparameters shown in Table 4 and patience 20. These hyperparameters were selected over a search of batch size 32 to 64 and learning rate 10^{-4} to 5×10^{-4} , with mean absolute error (MAE) as selection criterion.

Table 6: Parameter counts.

	PGDM Pattern Predictor	PGDM Denoiser	TimeGrad	CSDI	ARMD	GenViT
UWHVF	61,581	394,567	60,103	610,945	375	4,393,874
Break	270,687	418,483	61,251	610,961	613	-
House	282,222					
Ballet jazz	255,307					
Street jazz	251,462					
Krump	282,222					
LA Hip Hop	282,222					
Lock	282,222					
Middle Hip Hop	282,222					
Pop	270,687					
Wack	243,772					

Table 7: Inference latency (ms). Latency is computed by measuring the average wall-clock time of a forward pass with batch size 32 (over 10 repeated runs on a fixed batch after warmup) and dividing the resulting batch latency by 32. Latency on the AIST++ dataset is measured using the break genre.

	PGDM	TimeGrad	CSDI	ARMD	GenViT
UWHVF	742.60	<0.01	10.63	0.04	314.28
AIST++	732.30	<0.01	27.20	0.03	-

We train our diffusion model (see Figure 6 for architecture) with the Adam optimizer on a mean square error loss function with the hyperparameters shown in Table 4. For the VF prediction application, these hyperparameters were selected over a search of batch size 32 to 64, learning rate 10^{-5} to 5×10^{-5} , and 100 to 1000 epochs. For the motion prediction application, the model was trained with learning rate scheduling, and the hyperparameters were selected over a search of batch size 32 to 64, learning rate 5×10^{-4} to 10^{-3} and 200 to 300 epochs. In both applications, for our selection criterion, we measure MAE with maximum guidance scale $\bar{w} = 1, \dots, 10$ and no pattern mixing, and we choose only from models with the highest capacity for pattern guidance (i.e., error continues to reduce with increasing \bar{w}). Of these, we select the models with lowest achievable MAE over the tested range of \bar{w} . To better evaluate the full effect of pattern guidance on model performance, we also select models with the highest impact of pattern guidance over the tested range of \bar{w} (e.g., the greatest achievable percent decrease in error from applying guidance). For all PGDM models, we train with conditioning dropout probability $p_{\text{drop}} = 0.2$. We evaluate with maximum tolerable uncertainty $\gamma = 0.1, 0.03, 0.06, 0.04, 0.04, 0.05, 0.05, 0.06, 0.06, 0.03$, and 0.05 for UWHVF, break, house, ballet jazz, street jazz, krump, LA hip hop, lock, middle hip hop, pop, and wack, respectively. These were chosen based on the range of uncertainties on the validation data. In Table 5, we report the choices of \bar{w} and \bar{w}^* that we use to generate Table 1.

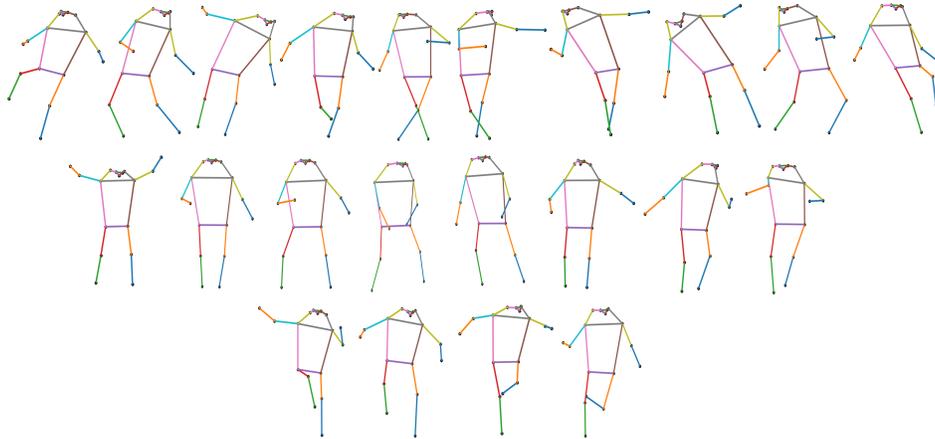
For our baselines TimeGrad, CSDI, and ARMD, we select hyperparameters following the published implementation details. For the GenViT model, we select hyperparameters from a hyperparameter search, as those published in Tian et al. (2023) were for a simpler task with $H = 1$ and $T = 1$. We train with batch size 16, learning rate 10^{-5} , and 50 epochs. These hyperparameters were selected from a search over batch size 8 to 16, learning rate 10^{-5} to 5×10^{-5} , and 10 to 50 epochs, with MAE as selection criterion. These ranges were chosen to remain consistent with the settings of Tian et al. (2023).

We use default architectures for all baseline models except ARMD. The original ARMD implementation uses a single layer linear network, with layer size equal to the prediction horizon, which can be extremely small in our application settings and therefore insufficiently expressive. For a fairer comparison, we instead apply a lightweight two layer MLP with hidden dimension 32 and layer normalization. In addition, ARMD’s evolution (i.e., forward) and devolution (i.e., reverse) processes assume equal history and prediction lengths. To accommodate unequal horizons, we adapt both processes so that a window of length H slides from time step 1 through T . Finally, following ARMD’s DDIM-style sampling procedure, we adopt $\eta = 0.1$ to introduce a small amount of stochasticity. These minimal adjustments preserve the core design of ARMD while enabling its application to our forecasting setting.

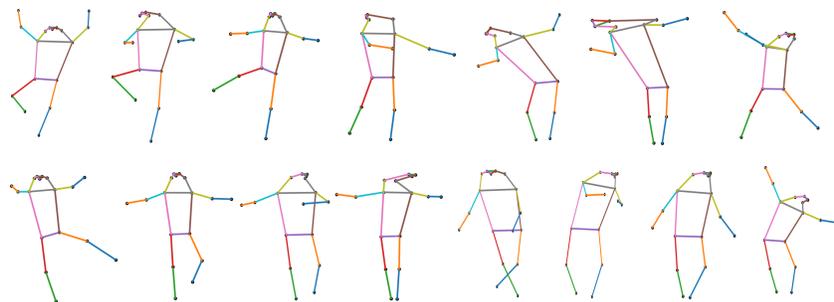
810 All experiments were performed on a machine with 42 GB of GPU memory. Each model requires less
 811 than 1 GB. We report parameter counts for each model in Table 6 and inference latency in Table 7.
 812 The sampling efficiency of our model can be substantially improved using accelerated samplers such
 813 as DDIM. Exploring these optimizations is beyond the scope of this work.

815 E PATTERNS EXTRACTED FROM AIST++

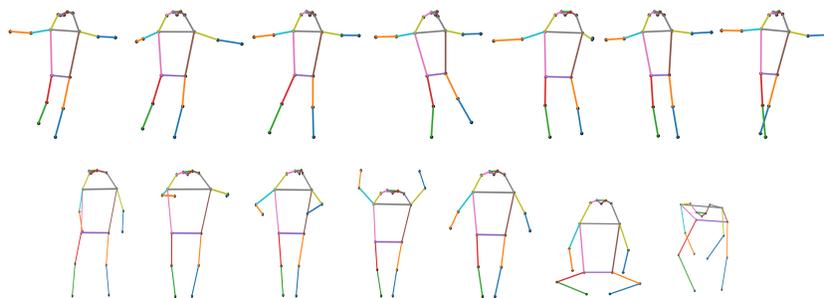
816
 817 Figures 7 to 15 show the patterns extracted from each genre of the AIST++ dataset.



834 Figure 7: Twenty two archetypes extracted from AIST++ house dancing frames.



847 Figure 8: Fifteen archetypes extracted from AIST++ ballet jazz dancing frames.



860 Figure 9: Fourteen archetypes extracted from AIST++ street jazz dancing frames.

864
865
866
867
868
869
870
871
872
873
874
875
876
877

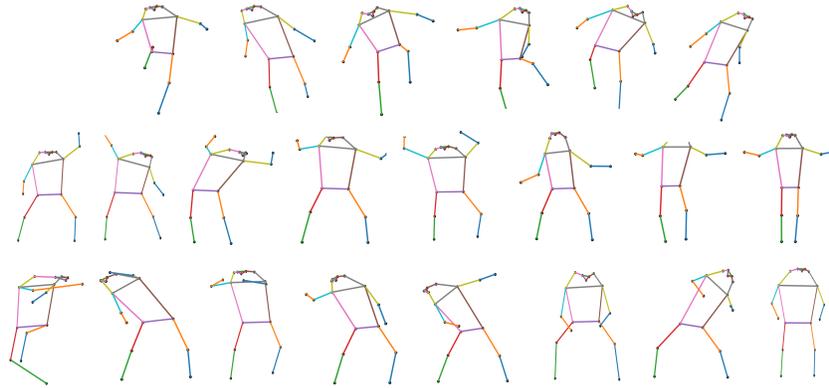


Figure 10: Twenty two archetypes extracted from AIST++ krump dancing frames.

880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896

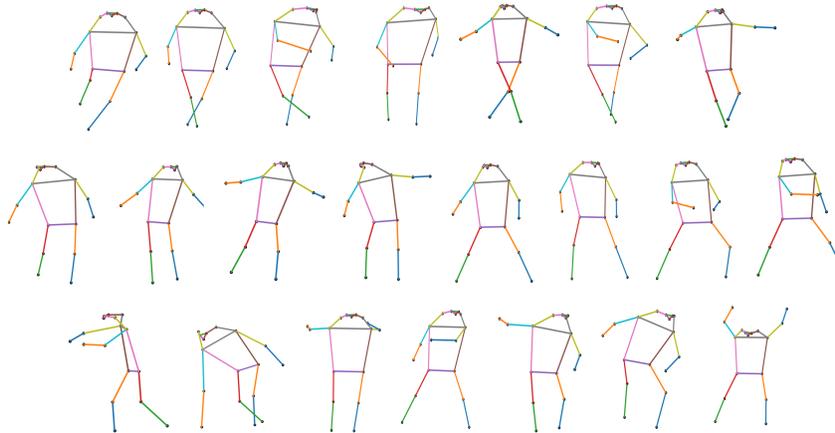


Figure 11: Twenty two archetypes extracted from AIST++ LA hip hop dancing frames.

897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915

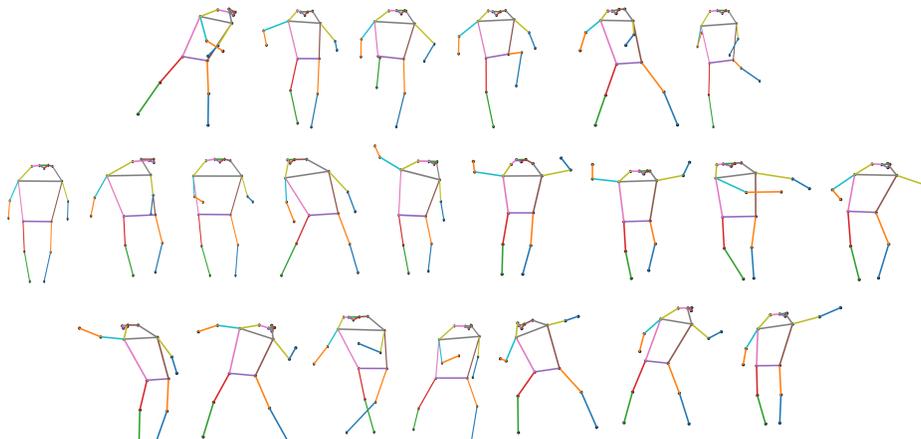


Figure 12: Twenty two archetypes extracted from AIST++ lock dancing frames.

916
917

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

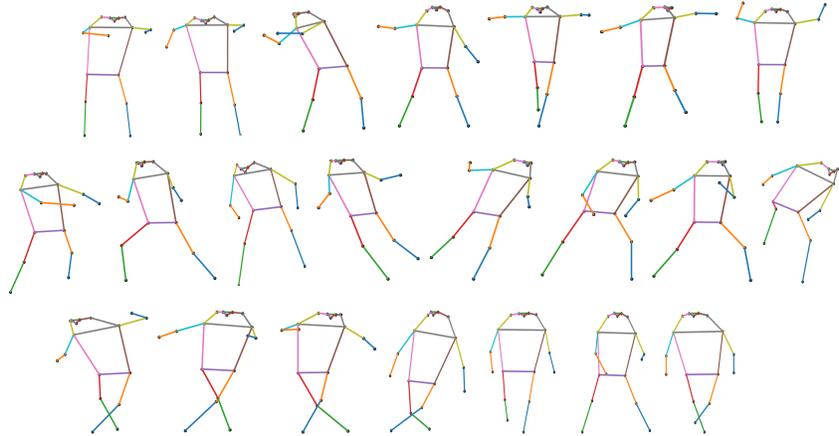


Figure 13: Twenty two archetypes extracted from AIST++ middle hip hop dancing frames.

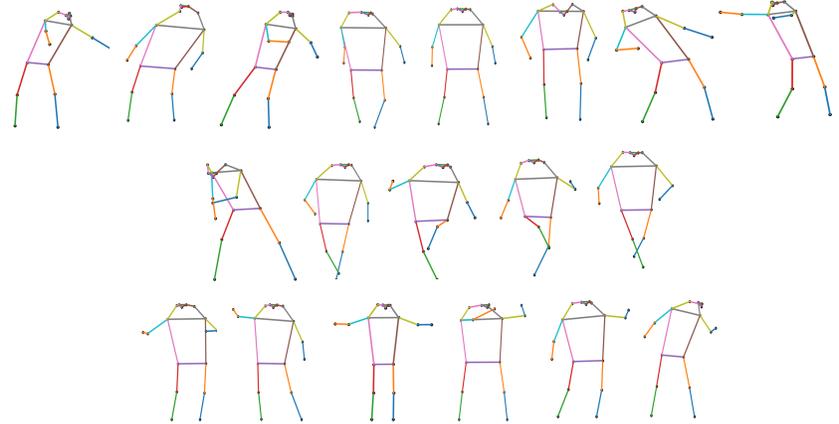


Figure 14: Nineteen archetypes extracted from AIST++ pop dancing frames.

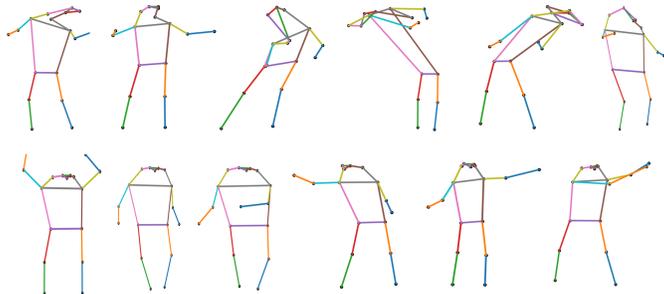


Figure 15: Twelve archetypes extracted from AIST++ wack dancing frames.

F EVALUATION OF PGDM COMPONENTS

Table 8 shows the reconstruction error of the extracted patterns, the error of the pattern prediction model, and the error of the guidance function.

Table 8: Mean absolute error (MAE) of the guidance function and its components. Note that the pattern prediction is performed in the pattern representation space, with range (0, 1).

	Archetypal Analysis	Pattern Prediction	Guidance Function
UWHVF	2.3146	0.0466	3.1560
Break	0.7929	0.0111	0.8440
House	1.6591	0.0111	1.7636
Ballet Jazz	0.6269	0.0156	0.6698
Street Jazz	1.3009	0.0108	1.3438
Krump	1.6768	0.0085	1.7677
LA Hip Hop	1.6707	0.0107	1.7622
Lock	1.3892	0.0102	1.4743
Middle Hip Hop	1.8489	0.0113	1.9682
Pop	0.9886	0.0106	1.0572
Wack	0.5353	0.0166	0.5973

G IMPACT OF PATTERN GUIDANCE

In the main text, we observed that pattern guidance reduces the error of PGDM’s predictions. To further illustrate this point, qualitative examples for both applications are shown in Figure 16. For VF prediction, we show five example $H = 1$ step-ahead predictions from PGDM_{GDE}. When pattern guidance is not used ($\bar{w} = 0$), PGDM makes a noisy prediction based only on the past visual field data. When pattern guidance is added ($\bar{w} = 5$), PGDM incorporates the pattern prediction in its forecast. The outcome resembles a mixture of the pattern prediction and the unguided prediction (see Ex. 2 of 16a). For motion prediction, we show one example $H = 5$ step ahead prediction for PGDM_{GDE}. In this example, we highlight the bent right leg of the skeleton. Without pattern guidance ($\bar{w} = 0$), the model predicts nearly no motion in the leg across the horizon. In contrast, the guidance function predicts a set of patterns that change over time, matching the moving right leg of the ground truth frames. When guidance is used ($\bar{w} = 2$), PGDM incorporates this motion into its prediction and forecasts more accurate future frames.

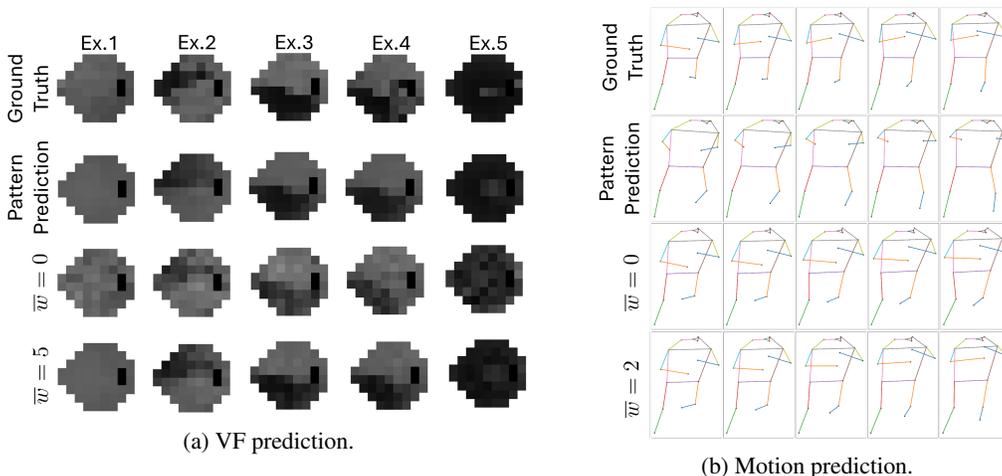


Figure 16: Qualitative examples of pattern guidance for PGDM_{GDE} on a) the visual field prediction application and b) the human motion prediction application.

In Tables 9 to 18, we show the quantitative effect of pattern guidance levels $\bar{w} = 1, 2, 3, 4, 5$ on PGDM’s MAE for the human motion prediction applications, holding the mixing scale w^* constant at its optimal choice. Each table corresponds to one of the dance genres in the AIST++ dataset. Table 19 additionally shows the effect of pattern guidance levels on the CRPS of PGDM for both the visual field and human motion prediction applications.¹ In general PGDM_{MAE} and PGDM_{GDE} achieve their best performances with relatively light guidance. Beyond this point, the pattern guidance has diminishing returns, even increasing the prediction error when the guidance scale is too high. In practice, the appropriate \bar{w} may be selected in a manner similar to a hyperparameter search. We also observe that, in most cases, the standard deviation of the MAE decreases as \bar{w} increases up to the optimal \bar{w} . This indicates that pattern guidance improves both the quality and the consistency of predictions.

Table 9: Mean absolute error (MAE) of PGDM_{MAE}, PGDM_{GDE}, and baselines on the break dancing genre of the AIST++ dataset. Percent improvements over baselines are shown in the Δ MAE (%) columns. Mean and standard deviation are taken across five samples.

Model	MAE (dB)	Δ MAE (%) vs. TimeGrad	Δ MAE (%) vs. CSDI	Δ MAE (%) vs. ARMD	Δ MAE (%) vs. $\bar{w} = 0$	
TimeGrad	2.10±0.0064	-	-	-	-	
CSDI	0.47±0.0010	-	-	-	-	
ARMD	4.26±0.0026	-	-	-	-	
PGDM _{MAE}	$\bar{w} = 0$	0.41±0.0013	80.33±0.10	12.54±0.15	90.31±0.03	-
	$\bar{w} = 1$	0.38±0.0009	81.91±0.07	19.54±0.14	91.08±0.02	8.01±0.14
	$\bar{w} = 2$	0.38±0.0009	81.80±0.07	19.04±0.16	91.03±0.02	7.43±0.16
	$\bar{w} = 3$	0.40±0.0007	81.16±0.06	16.22±0.17	90.72±0.02	4.21±0.28
	$\bar{w} = 4$	0.42±0.0009	80.19±0.09	11.91±0.19	90.24±0.02	-0.72±0.33
	$\bar{w} = 5$	0.44±0.0010	79.02±0.09	6.71±0.14	89.66±0.02	-6.66±0.32
PGDM _{GDE}	$\bar{w} = 0$	0.52±0.0032	75.25±0.18	-10.06±0.76	87.80±0.07	-
	$\bar{w} = 1$	0.45±0.0023	78.32±0.13	3.60±0.64	89.32±0.05	12.41±0.29
	$\bar{w} = 2$	0.45±0.0015	78.75±0.11	5.48±0.45	89.53±0.03	14.12±0.35
	$\bar{w} = 3$	0.46±0.0008	78.24±0.09	3.24±0.28	89.28±0.02	12.08±0.44
	$\bar{w} = 4$	0.48±0.0010	77.34±0.11	-0.79±0.21	88.83±0.02	8.42±0.63
	$\bar{w} = 5$	0.50±0.0012	76.21±0.12	-5.81±0.21	88.27±0.03	3.86±0.74

¹To avoid redundancy, we report CRPS alone and do not report percent improvements over baselines. Trends apparent in the reported MAE results are consistent with those reported in the CRPS results. Compared to unguided predictions, PGDM_{GDE} achieves the greatest reduction in CRPS using guidance (56.26% on UWHVF and 14.10% on the LA hip hop genre of AIST++). Compared to baselines, PGDM_{MAE} can surpass baselines by up to 84.55% on UWHVF and 92.55% on the break dancing genre of AIST++.

1080
1081
1082
1083
1084
1085
1086

Table 10: Mean absolute error (MAE) of PGDM_{MAE}, PGDM_{GDE}, and baselines on the house dancing genre of the AIST++ dataset. Percent improvements over baselines are shown in the Δ MAE (%) columns. Mean and standard deviation are taken across five samples.

Model	MAE (dB)	Δ MAE (%) vs. TimeGrad	Δ MAE (%) vs. CSDI	Δ MAE (%) vs. ARMD	Δ MAE (%) vs. $\bar{w} = 0$	
TimeGrad	3.71 \pm 0.0159	-	-	-	-	
CSDI	1.02 \pm 0.0045	-	-	-	-	
ARMD	9.11 \pm 0.0035	-	-	-	-	
PGDM _{MAE}	$\bar{w} = 0$	0.79 \pm 0.0017	78.63 \pm 0.12	22.11 \pm 0.35	91.30 \pm 0.02	-
	$\bar{w} = 1$	0.74 \pm 0.0014	79.95 \pm 0.11	26.93 \pm 0.34	91.84 \pm 0.02	6.19 \pm 0.10
	$\bar{w} = 2$	0.75 \pm 0.0011	79.75 \pm 0.11	26.20 \pm 0.31	91.76 \pm 0.01	6.96 \pm 0.09
	$\bar{w} = 3$	0.78 \pm 0.0012	79.03 \pm 0.12	23.57 \pm 0.29	91.46 \pm 0.01	1.87 \pm 0.17
	$\bar{w} = 4$	0.82 \pm 0.0014	78.01 \pm 0.12	19.85 \pm 0.34	91.05 \pm 0.02	-2.90 \pm 0.20
$\bar{w} = 5$	0.86 \pm 0.0019	76.80 \pm 0.14	15.43 \pm 0.41	90.56 \pm 0.02	-8.57 \pm 0.28	
PGDM _{GDE}	$\bar{w} = 0$	0.90 \pm 0.0013	75.64 \pm 0.11	11.19 \pm 0.39	90.08 \pm 0.08	-
	$\bar{w} = 1$	0.82 \pm 0.0014	77.88 \pm 0.12	19.38 \pm 0.35	91.00 \pm 0.02	9.22 \pm 0.12
	$\bar{w} = 2$	0.82 \pm 0.0017	77.81 \pm 0.12	19.13 \pm 0.40	90.97 \pm 0.02	8.94 \pm 0.25
	$\bar{w} = 3$	0.85 \pm 0.0017	77.09 \pm 0.13	16.51 \pm 0.42	90.67 \pm 0.02	5.98 \pm 0.23
	$\bar{w} = 4$	0.89 \pm 0.0019	76.09 \pm 0.15	12.84 \pm 0.29	90.27 \pm 0.02	1.85 \pm 0.18
$\bar{w} = 5$	0.93 \pm 0.0015	74.96 \pm 0.14	8.73 \pm 0.28	89.81 \pm 0.02	-2.78 \pm 0.15	

1104
1105
1106
1107
1108
1109
1110

Table 11: Mean absolute error (MAE) of PGDM_{MAE}, PGDM_{GDE}, and baselines on the ballet jazz dancing genre of the AIST++ dataset. Percent improvements over baselines are shown in the Δ MAE (%) columns. Mean and standard deviation are taken across five samples.

Model	MAE (dB)	Δ MAE (%) vs. TimeGrad	Δ MAE (%) vs. CSDI	Δ MAE (%) vs. ARMD	Δ MAE (%) vs. $\bar{w} = 0$	
TimeGrad	1.32 \pm 0.0098	-	-	-	-	
CSDI	0.55 \pm 0.0054	-	-	-	-	
ARMD	3.38 \pm 0.0011	-	-	-	-	
PGDM _{MAE}	$\bar{w} = 0$	0.42 \pm 0.0008	67.94 \pm 0.19	22.83 \pm 0.84	87.53 \pm 0.02	-
	$\bar{w} = 1$	0.38 \pm 0.0002	71.13 \pm 0.21	30.50 \pm 0.67	88.77 \pm 0.01	9.94 \pm 0.16
	$\bar{w} = 2$	0.39 \pm 0.0005	70.53 \pm 0.24	29.07 \pm 0.64	88.54 \pm 0.02	8.08 \pm 0.24
	$\bar{w} = 3$	0.41 \pm 0.0006	68.59 \pm 0.27	24.39 \pm 0.69	87.79 \pm 0.02	2.02 \pm 0.30
	$\bar{w} = 4$	0.45 \pm 0.0011	65.89 \pm 0.29	17.90 \pm 0.82	86.74 \pm 0.03	-6.39 \pm 0.42
$\bar{w} = 5$	0.49 \pm 0.0012	62.85 \pm 0.30	10.58 \pm 0.93	85.55 \pm 0.03	-15.87 \pm 0.45	
PGDM _{GDE}	$\bar{w} = 0$	0.49 \pm 0.0004	62.75 \pm 0.30	10.34 \pm 0.83	85.52 \pm 0.02	-
	$\bar{w} = 1$	0.42 \pm 0.0009	67.74 \pm 0.27	22.34 \pm 0.81	87.45 \pm 0.03	13.39 \pm 0.17
	$\bar{w} = 2$	0.43 \pm 0.0010	67.15 \pm 0.29	20.92 \pm 0.81	87.23 \pm 0.03	11.81 \pm 0.20
	$\bar{w} = 3$	0.46 \pm 0.0010	64.86 \pm 0.32	15.40 \pm 0.89	86.33 \pm 0.03	5.65 \pm 0.24
	$\bar{w} = 4$	0.50 \pm 0.0008	61.66 \pm 0.33	7.70 \pm 0.93	85.09 \pm 0.03	-2.94 \pm 0.21
$\bar{w} = 5$	0.55 \pm 0.0011	57.90 \pm 0.38	-1.34 \pm 1.04	83.63 \pm 0.04	-13.02 \pm 0.26	

1131
1132
1133

Table 12: Mean absolute error (MAE) of PGDM_{MAE}, PGDM_{GDE}, and baselines on the street jazz dancing genre of the AIST++ dataset. Percent improvements over baselines are shown in the Δ MAE (%) columns. Mean and standard deviation are taken across five samples.

Model	MAE (dB)	Δ MAE (%) vs. TimeGrad	Δ MAE (%) vs. CSDI	Δ MAE (%) vs. ARMD	Δ MAE (%) vs. $\bar{w} = 0$	
TimeGrad	1.65 \pm 0.0102	-	-	-	-	
CSDI	0.56 \pm 0.0054	-	-	-	-	
ARMD	7.57 \pm 0.0049	-	-	-	-	
PGDM _{MAE}	$\bar{w} = 0$	0.52 \pm 0.0017	68.58 \pm 0.25	6.58 \pm 0.79	93.14 \pm 0.02	-
	$\bar{w} = 1$	0.48 \pm 0.0010	70.66 \pm 0.23	12.76 \pm 0.76	93.59 \pm 0.01	6.62 \pm 0.19
	$\bar{w} = 2$	0.48 \pm 0.0008	70.87 \pm 0.22	13.39 \pm 0.80	93.64 \pm 0.01	7.29 \pm 0.18
	$\bar{w} = 3$	0.49 \pm 0.0009	70.23 \pm 0.22	11.49 \pm 0.85	93.50 \pm 0.01	5.26 \pm 0.20
	$\bar{w} = 4$	0.51 \pm 0.0009	69.07 \pm 0.23	8.03 \pm 0.91	93.25 \pm 0.01	1.56 \pm 0.22
	$\bar{w} = 5$	0.54 \pm 0.0009	67.53 \pm 0.25	3.45 \pm 0.95	92.91 \pm 0.01	-3.34 \pm 0.25
PGDM _{GDE}	$\bar{w} = 0$	0.60 \pm 0.0016	63.85 \pm 0.26	-7.49 \pm 1.02	92.11 \pm 0.02	-
	$\bar{w} = 1$	0.55 \pm 0.0007	66.66 \pm 0.23	0.88 \pm 0.99	92.72 \pm 0.01	7.79 \pm 0.19
	$\bar{w} = 2$	0.54 \pm 0.0005	67.12 \pm 0.20	2.25 \pm 0.97	92.82 \pm 0.01	9.06 \pm 0.19
	$\bar{w} = 3$	0.55 \pm 0.0005	66.41 \pm 0.21	0.12 \pm 0.96	92.67 \pm 0.01	7.08 \pm 0.17
	$\bar{w} = 4$	0.58 \pm 0.0004	65.08 \pm 0.22	-3.82 \pm 0.97	92.38 \pm 0.01	3.41 \pm 0.21
	$\bar{w} = 5$	0.60 \pm 0.0005	63.39 \pm 0.22	-8.84 \pm 1.00	92.01 \pm 0.01	-1.25 \pm 0.21

Table 13: Mean absolute error (MAE) of PGDM_{MAE}, PGDM_{GDE}, and baselines on the krump dancing genre of the AIST++ dataset. Percent improvements over baselines are shown in the Δ MAE (%) columns. Mean and standard deviation are taken across five samples.

Model	MAE (dB)	Δ MAE (%) vs. TimeGrad	Δ MAE (%) vs. CSDI	Δ MAE (%) vs. ARMD	Δ MAE (%) vs. $\bar{w} = 0$	
TimeGrad	2.37 \pm 0.0067	-	-	-	-	
CSDI	0.77 \pm 0.0017	-	-	-	-	
ARMD	8.90 \pm 0.0046	-	-	-	-	
PGDM _{MAE}	$\bar{w} = 0$	0.77 \pm 0.0016	67.56 \pm 0.13	-0.04 \pm 0.31	91.36 \pm 0.02	-
	$\bar{w} = 1$	0.71 \pm 0.0014	70.27 \pm 0.13	8.30 \pm 0.31	92.08 \pm 0.01	8.33 \pm 0.08
	$\bar{w} = 2$	0.70 \pm 0.0013	70.40 \pm 0.13	8.70 \pm 0.31	92.11 \pm 0.01	8.73 \pm 0.12
	$\bar{w} = 3$	0.73 \pm 0.0011	69.31 \pm 0.13	5.33 \pm 0.30	91.82 \pm 0.01	5.37 \pm 0.12
	$\bar{w} = 4$	0.77 \pm 0.0010	67.53 \pm 0.13	-0.14 \pm 0.28	91.35 \pm 0.01	-0.11 \pm 0.13
	$\bar{w} = 5$	0.82 \pm 0.0008	65.26 \pm 0.12	-7.15 \pm 0.25	90.74 \pm 0.01	-7.11 \pm 0.16
PGDM _{GDE}	$\bar{w} = 0$	0.88 \pm 0.0016	62.85 \pm 0.12	-14.59 \pm 0.26	90.10 \pm 0.02	-
	$\bar{w} = 1$	0.79 \pm 0.0005	66.85 \pm 0.10	-2.23 \pm 0.20	91.17 \pm 0.01	10.79 \pm 0.11
	$\bar{w} = 2$	0.79 \pm 0.0005	66.83 \pm 0.10	-2.32 \pm 0.24	91.16 \pm 0.01	10.71 \pm 0.11
	$\bar{w} = 3$	0.81 \pm 0.0006	65.81 \pm 0.11	-5.46 \pm 0.27	90.89 \pm 0.01	7.97 \pm 0.12
	$\bar{w} = 4$	0.85 \pm 0.0007	64.38 \pm 0.11	-9.86 \pm 0.28	90.51 \pm 0.01	4.13 \pm 0.12
	$\bar{w} = 5$	0.88 \pm 0.0007	62.76 \pm 0.11	-14.84 \pm 0.30	90.08 \pm 0.01	-0.22 \pm 0.13

1188
1189
1190
1191
1192
1193
1194

Table 14: Mean absolute error (MAE) of PGDM_{MAE}, PGDM_{GDE}, and baselines on the LA hip hop dancing genre of the AIST++ dataset. Percent improvements over baselines are shown in the Δ MAE (%) columns. Mean and standard deviation are taken across five samples.

Model	MAE (dB)	Δ MAE (%) vs. TimeGrad	Δ MAE (%) vs. CSDI	Δ MAE (%) vs. ARMD	Δ MAE (%) vs. $\bar{w} = 0$	
TimeGrad	3.30 \pm 0.0157	-	-	-	-	
CSDI	0.78 \pm 0.0023	-	-	-	-	
ARMD	8.51 \pm 0.0032	-	-	-	-	
PGDM _{MAE}	$\bar{w} = 0$	0.80 \pm 0.0010	75.83 \pm 0.13	-1.95 \pm 0.36	90.62 \pm 0.01	-
	$\bar{w} = 1$	0.74 \pm 0.0006	77.58 \pm 0.11	5.46 \pm 0.33	91.30 \pm 0.01	7.26 \pm 0.06
	$\bar{w} = 2$	0.75 \pm 0.0006	77.41 \pm 0.11	4.71 \pm 0.33	91.23 \pm 0.01	6.53 \pm 0.08
	$\bar{w} = 3$	0.77 \pm 0.0007	76.61 \pm 0.11	1.37 \pm 0.35	90.92 \pm 0.01	3.25 \pm 0.10
	$\bar{w} = 4$	0.81 \pm 0.0007	75.48 \pm 0.11	-3.41 \pm 0.36	90.48 \pm 0.01	-1.44 \pm 0.11
	$\bar{w} = 5$	0.85 \pm 0.0007	74.12 \pm 0.12	-9.14 \pm 0.37	89.96 \pm 0.01	-7.05 \pm 0.12
PGDM _{GDE}	$\bar{w} = 0$	0.90 \pm 0.0009	72.82 \pm 0.15	-14.62 \pm 0.34	89.45 \pm 0.01	-
	$\bar{w} = 1$	0.81 \pm 0.0009	75.42 \pm 0.13	-3.65 \pm 0.32	90.46 \pm 0.01	9.57 \pm 0.03
	$\bar{w} = 2$	0.82 \pm 0.0006	75.28 \pm 0.13	-4.27 \pm 0.33	90.41 \pm 0.01	9.03 \pm 0.05
	$\bar{w} = 3$	0.85 \pm 0.0005	74.34 \pm 0.13	-8.23 \pm 0.33	90.04 \pm 0.01	5.58 \pm 0.07
	$\bar{w} = 4$	0.89 \pm 0.0007	73.00 \pm 0.14	-13.87 \pm 0.40	89.52 \pm 0.01	0.66 \pm 0.10
	$\bar{w} = 5$	0.94 \pm 0.0009	71.44 \pm 0.15	-20.47 \pm 0.44	88.92 \pm 0.01	-5.10 \pm 0.10

1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218

Table 15: Mean absolute error (MAE) of PGDM_{MAE}, PGDM_{GDE}, and baselines on the lock dancing genre of the AIST++ dataset. Percent improvements over baselines are shown in the Δ MAE (%) columns. Mean and standard deviation are taken across five samples.

Model	MAE (dB)	Δ MAE (%) vs. TimeGrad	Δ MAE (%) vs. CSDI	Δ MAE (%) vs. ARMD	Δ MAE (%) vs. $\bar{w} = 0$	
TimeGrad	3.03 \pm 0.0086	-	-	-	-	
CSDI	0.76 \pm 0.0028	-	-	-	-	
ARMD	7.65 \pm 0.0031	-	-	-	-	
PGDM _{MAE}	$\bar{w} = 0$	0.72 \pm 0.0011	76.12 \pm 0.10	4.24 \pm 0.23	90.53 \pm 0.01	-
	$\bar{w} = 1$	0.67 \pm 0.0006	78.05 \pm 0.07	11.98 \pm 0.28	91.29 \pm 0.01	8.07 \pm 0.12
	$\bar{w} = 2$	0.68 \pm 0.0003	77.70 \pm 0.06	10.60 \pm 0.32	91.15 \pm 0.01	6.63 \pm 0.16
	$\bar{w} = 3$	0.71 \pm 0.0007	76.67 \pm 0.07	6.47 \pm 0.29	90.75 \pm 0.01	2.32 \pm 0.13
	$\bar{w} = 4$	0.75 \pm 0.0008	75.26 \pm 0.09	0.81 \pm 0.27	90.19 \pm 0.01	-3.59 \pm 0.12
	$\bar{w} = 5$	0.80 \pm 0.0011	73.57 \pm 0.09	-5.98 \pm 0.38	89.51 \pm 0.02	-10.68 \pm 0.24
PGDM _{GDE}	$\bar{w} = 0$	0.78 \pm 0.0017	74.17 \pm 0.09	-3.57 \pm 0.29	89.75 \pm 0.03	-
	$\bar{w} = 1$	0.71 \pm 0.0023	76.69 \pm 0.12	6.54 \pm 0.23	90.75 \pm 0.03	9.76 \pm 0.22
	$\bar{w} = 2$	0.72 \pm 0.0020	76.43 \pm 0.11	5.50 \pm 0.34	90.65 \pm 0.03	8.76 \pm 0.30
	$\bar{w} = 3$	0.75 \pm 0.0015	75.37 \pm 0.11	1.23 \pm 0.22	90.23 \pm 0.02	4.64 \pm 0.18
	$\bar{w} = 4$	0.79 \pm 0.0013	73.83 \pm 0.10	-4.95 \pm 0.26	89.62 \pm 0.02	-1.33 \pm 0.13
	$\bar{w} = 5$	0.85 \pm 0.0008	71.88 \pm 0.09	-12.73 \pm 0.35	88.85 \pm 0.01	-8.85 \pm 0.17

1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Table 16: Mean absolute error (MAE) of PGDM_{MAE}, PGDM_{GDE}, and baselines on the middle hip hop dancing genre of the AIST++ dataset. Percent improvements over baselines are shown in the Δ MAE (%) columns. Mean and standard deviation are taken across five samples.

Model	MAE (dB)	Δ MAE (%) vs. TimeGrad	Δ MAE (%) vs. CSDI	Δ MAE (%) vs. ARMD	Δ MAE (%) vs. $\bar{w} = 0$	
TimeGrad	3.35 \pm 0.0113	-	-	-	-	
CSDI	1.04 \pm 0.0048	-	-	-	-	
ARMD	8.98 \pm 0.0051	-	-	-	-	
PGDM _{MAE}	$\bar{w} = 0$	0.88 \pm 0.0014	73.79 \pm 0.11	15.40 \pm 0.46	90.21 \pm 0.01	-
	$\bar{w} = 1$	0.82 \pm 0.0009	75.40 \pm 0.10	20.61 \pm 0.45	90.82 \pm 0.01	6.15 \pm 0.11
	$\bar{w} = 2$	0.83 \pm 0.0013	75.11 \pm 0.09	19.68 \pm 0.43	90.71 \pm 0.01	5.05 \pm 0.11
	$\bar{w} = 3$	0.87 \pm 0.0008	74.14 \pm 0.09	16.54 \pm 0.45	90.34 \pm 0.01	1.34 \pm 0.15
	$\bar{w} = 4$	0.91 \pm 0.0011	72.74 \pm 0.10	12.03 \pm 0.50	89.82 \pm 0.01	-3.98 \pm 0.20
	$\bar{w} = 5$	0.97 \pm 0.0010	71.07 \pm 0.10	6.64 \pm 0.49	89.20 \pm 0.01	-10.36 \pm 0.17
PGDM _{GDE}	$\bar{w} = 0$	1.05 \pm 0.0034	68.67 \pm 0.13	-1.11 \pm 0.37	88.30 \pm 0.03	-
	$\bar{w} = 1$	0.95 \pm 0.0038	71.69 \pm 0.11	8.64 \pm 0.25	89.43 \pm 0.04	9.64 \pm 0.20
	$\bar{w} = 2$	0.95 \pm 0.0031	71.56 \pm 0.05	8.21 \pm 0.32	89.38 \pm 0.03	9.22 \pm 0.31
	$\bar{w} = 3$	0.99 \pm 0.0026	70.50 \pm 0.06	4.78 \pm 0.36	88.98 \pm 0.03	5.83 \pm 0.25
	$\bar{w} = 4$	1.04 \pm 0.0022	69.06 \pm 0.07	0.16 \pm 0.45	88.45 \pm 0.02	1.25 \pm 0.35
	$\bar{w} = 5$	1.09 \pm 0.0027	67.37 \pm 0.11	-5.30 \pm 0.58	87.82 \pm 0.02	-4.14 \pm 0.36

Table 17: Mean absolute error (MAE) of PGDM_{MAE}, PGDM_{GDE}, and baselines on the pop dancing genre of the AIST++ dataset. Percent improvements over baselines are shown in the Δ MAE (%) columns. Mean and standard deviation are taken across five samples.

Model	MAE (dB)	Δ MAE (%) vs. TimeGrad	Δ MAE (%) vs. CSDI	Δ MAE (%) vs. ARMD	Δ MAE (%) vs. $\bar{w} = 0$	
TimeGrad	2.55 \pm 0.0105	-	-	-	-	
CSDI	0.70 \pm 0.0053	-	-	-	-	
ARMD	6.43 \pm 0.0019	-	-	-	-	
PGDM _{MAE}	$\bar{w} = 0$	0.47 \pm 0.0008	81.57 \pm 0.09	32.95 \pm 0.58	92.69 \pm 0.01	-
	$\bar{w} = 1$	0.44 \pm 0.0018	82.60 \pm 0.08	36.71 \pm 0.68	93.10 \pm 0.03	5.60 \pm 0.38
	$\bar{w} = 2$	0.45 \pm 0.0018	82.40 \pm 0.07	35.97 \pm 0.67	93.02 \pm 0.03	4.49 \pm 0.37
	$\bar{w} = 3$	0.47 \pm 0.0015	81.64 \pm 0.09	33.20 \pm 0.69	92.72 \pm 0.02	0.37 \pm 0.36
	$\bar{w} = 4$	0.50 \pm 0.0009	80.52 \pm 0.10	29.14 \pm 0.66	92.28 \pm 0.01	-5.69 \pm 0.23
	$\bar{w} = 5$	0.53 \pm 0.0007	79.07 \pm 0.10	23.88 \pm 0.67	91.70 \pm 0.01	-13.54 \pm 0.23
PGDM _{GDE}	$\bar{w} = 0$	0.52 \pm 0.0013	79.44 \pm 0.09	25.21 \pm 0.57	91.85 \pm 0.02	-
	$\bar{w} = 1$	0.48 \pm 0.0018	81.01 \pm 0.12	30.92 \pm 0.50	92.47 \pm 0.03	7.64 \pm 0.39
	$\bar{w} = 2$	0.49 \pm 0.0016	80.96 \pm 0.12	30.73 \pm 0.61	92.45 \pm 0.03	7.39 \pm 0.48
	$\bar{w} = 3$	0.50 \pm 0.0014	80.29 \pm 0.13	28.31 \pm 0.59	92.19 \pm 0.02	4.15 \pm 0.43
	$\bar{w} = 4$	0.53 \pm 0.0011	79.16 \pm 0.09	24.20 \pm 0.57	91.74 \pm 0.02	-1.35 \pm 0.45
	$\bar{w} = 5$	0.57 \pm 0.0015	77.76 \pm 0.11	19.11 \pm 0.51	91.18 \pm 0.02	-8.15 \pm 0.50

Table 18: Mean absolute error (MAE) of PGDM_{MAE}, PGDM_{GDE}, and baselines on the wack dancing genre of the AIST++ dataset. Percent improvements over baselines are shown in the Δ MAE (%) columns. Mean and standard deviation are taken across five samples.

Model	MAE (dB)	Δ MAE (%) vs. TimeGrad	Δ MAE (%) vs. CSDI	Δ MAE (%) vs. ARMD	Δ MAE (%) vs. $\bar{w} = 0$	
TimeGrad	1.03 \pm 0.0047	-	-	-	-	
CSDI	0.44 \pm 0.0042	-	-	-	-	
ARMD	3.39 \pm 0.0026	-	-	-	-	
PGDM _{MAE}	$\bar{w} = 0$	0.44 \pm 0.0044	57.55 \pm 0.45	0.61 \pm 1.39	87.15 \pm 0.12	-
	$\bar{w} = 1$	0.39\pm0.0028	61.54\pm0.28	9.96\pm1.37	88.36\pm0.08	9.40\pm0.64
	$\bar{w} = 2$	0.40 \pm 0.0027	60.53 \pm 0.26	7.59 \pm 1.36	88.05 \pm 0.07	7.01 \pm 0.75
	$\bar{w} = 3$	0.43 \pm 0.0018	58.40 \pm 0.20	2.59 \pm 1.19	87.41 \pm 0.05	1.99 \pm 0.69
	$\bar{w} = 4$	0.45 \pm 0.0015	55.70 \pm 0.20	-3.71 \pm 1.06	86.59 \pm 0.05	-4.36 \pm 0.76
	$\bar{w} = 5$	0.48 \pm 0.0019	52.72 \pm 0.25	-10.69 \pm 0.96	85.69 \pm 0.05	-11.38 \pm 0.83
PGDM _{GDE}	$\bar{w} = 0$	0.49 \pm 0.0054	52.11 \pm 0.44	-12.12 \pm 1.81	85.50 \pm 0.16	-
	$\bar{w} = 1$	0.45 \pm 0.0079	56.17 \pm 0.63	-2.63 \pm 2.45	86.73 \pm 0.23	8.47 \pm 0.75
	$\bar{w} = 2$	0.46 \pm 0.0079	55.41 \pm 0.61	-4.40 \pm 2.37	86.50 \pm 0.24	6.89 \pm 0.83
	$\bar{w} = 3$	0.48 \pm 0.0075	53.20 \pm 0.57	-9.58 \pm 2.32	85.83 \pm 0.22	2.28 \pm 0.66
	$\bar{w} = 4$	0.51 \pm 0.0059	50.33 \pm 0.45	-16.29 \pm 2.06	84.96 \pm 0.18	-3.72 \pm 0.28
	$\bar{w} = 5$	0.54 \pm 0.0046	46.95 \pm 0.30	-24.20 \pm 1.83	83.94 \pm 0.14	-10.78 \pm 0.55

Table 19: Continuous ranked probability score (CRPS) of PGDM_{MAE} and PGDM_{GDE} for both the visual field prediction (UWHVF dataset) and the human motion prediction (10 dance genres from the AIST++ dataset) case studies. Mean and standard deviation are taken across three seeds

	PGDM _{MAE}					
	$\bar{w} = 0$	$\bar{w} = 1$	$\bar{w} = 2$	$\bar{w} = 3$	$\bar{w} = 4$	$\bar{w} = 5$
UWHVF	0.794 \pm 0.0122	0.777\pm0.0040	0.781 \pm 0.0040	0.786 \pm 0.0040	0.790 \pm 0.0039	0.794 \pm 0.0038
Break	0.039 \pm 0.0001	0.037\pm<0.0001	0.037\pm0.0001	0.040 \pm 0.0002	0.043 \pm 0.0002	0.046 \pm 0.0002
House	0.077 \pm 0.0004	0.073\pm0.0002	0.075 \pm 0.0001	0.080 \pm 0.0001	0.085 \pm 0.0001	0.092 \pm 0.0001
Ballet Jazz	0.039 \pm 0.0001	0.035\pm0.0002	0.037 \pm 0.0002	0.040 \pm 0.0003	0.046 \pm 0.0003	0.051 \pm 0.0003
Street Jazz	0.050 \pm 0.0002	0.047\pm0.0002	0.047\pm0.0002	0.049 \pm 0.0002	0.052 \pm 0.0003	0.055 \pm 0.0003
Krump	0.079 \pm 0.0003	0.070\pm0.0003	0.071 \pm 0.0003	0.076 \pm 0.0003	0.083 \pm 0.0003	0.090 \pm 0.0003
LA Hip Hop	0.075 \pm 0.0006	0.071\pm0.0005	0.073 \pm 0.0005	0.078 \pm 0.0005	0.083 \pm 0.0005	0.089 \pm 0.0005
Lock	0.068 \pm 0.0002	0.063\pm0.0001	0.066 \pm 0.0002	0.070 \pm 0.0002	0.076 \pm 0.0001	0.083 \pm 0.0001
Middle Hip Hop	0.085 \pm 0.0002	0.081\pm0.0003	0.084 \pm 0.0002	0.089 \pm 0.0002	0.097 \pm 0.0001	0.104 \pm 0.0001
Pop	0.046 \pm 0.0002	0.045\pm0.0002	0.046 \pm 0.0002	0.048 \pm 0.0002	0.052 \pm 0.0002	0.057 \pm 0.0002
Wack	0.043 \pm 0.0002	0.039\pm0.0003	0.042 \pm 0.0002	0.046 \pm 0.0001	0.052 \pm 0.0001	0.057 \pm 0.0001

	PGDM _{GDE}					
	$\bar{w} = 0$	$\bar{w} = 1$	$\bar{w} = 2$	$\bar{w} = 3$	$\bar{w} = 4$	$\bar{w} = 5$
UWHVF	1.887 \pm 0.0119	0.927 \pm 0.0028	0.887 \pm 0.0024	0.857 \pm 0.0020	0.837 \pm 0.0024	0.825\pm0.0032
Break	0.049 \pm 0.0001	0.043\pm0.0001	0.043\pm0.0001	0.046 \pm 0.0001	0.051 \pm 0.0001	0.056 \pm 0.0001
House	0.085 \pm 0.0005	0.079\pm0.0003	0.080 \pm 0.0002	0.085 \pm 0.0001	0.090 \pm 0.0001	0.095 \pm 0.0001
Ballet Jazz	0.044 \pm 0.0002	0.039\pm0.0002	0.040 \pm 0.0002	0.045 \pm 0.0002	0.051 \pm 0.0003	0.058 \pm 0.0003
Street Jazz	0.055 \pm 0.0002	0.051\pm0.0002	0.052 \pm 0.0001	0.054 \pm 0.0001	0.057 \pm 0.0001	0.061 \pm 0.0001
Krump	0.091 \pm 0.0004	0.078\pm0.0003	0.080 \pm 0.0004	0.085 \pm 0.0004	0.091 \pm 0.0004	0.097 \pm 0.0004
LA Hip Hop	0.083 \pm 0.0007	0.077\pm0.0005	0.079 \pm 0.0004	0.084 \pm 0.0004	0.090 \pm 0.0004	0.097 \pm 0.0004
Lock	0.073 \pm 0.0001	0.066\pm<0.0001	0.069 \pm <0.0001	0.074 \pm 0.0001	0.080 \pm <0.0001	0.087 \pm 0.0001
Middle Hip Hop	0.097 \pm 0.0003	0.089\pm0.0003	0.091 \pm 0.0003	0.097 \pm 0.0002	0.104 \pm 0.0002	0.111 \pm 0.0002
Pop	0.051 \pm 0.0002	0.047\pm0.0002	0.048 \pm 0.0002	0.051 \pm 0.0002	0.054 \pm 0.0002	0.059 \pm 0.0002
Wack	0.049 \pm 0.0002	0.044\pm0.0002	0.046 \pm 0.0001	0.050 \pm 0.0002	0.055 \pm 0.0002	0.060 \pm 0.0002

1350 H ABLATIONS

1351
1352 We report ablations evaluating the impact of pattern mixing and dynamic guidance and mixing scales
1353 here. For these ablations, we fix either or both the maximum guidance scale \bar{w} and the maximum
1354 mixing scale \bar{w}^* to the optimal choice reported in Table 5 in Appendix D. We evaluate performance
1355 with both MAE and CRPS_{SUM}.

1356 **Impact of pattern mixing.** In Algorithm 2, the final pattern mixing step mixes the deterministic
1357 prediction from the pattern predictor with the probabilistic prediction from the denoiser. The
1358 maximum mixing scale \bar{w}^* determines the strength of the pattern signal compared to the denoiser
1359 signal in the final forecast. In Tables 20 and 21, we study the effect of pattern mixing by holding
1360 \bar{w} constant and varying the mixing scale \bar{w}^* . On the UWHVF dataset, we find that a strong mixing
1361 signal provides a significant benefit to the prediction. In contrast, on the AIST++ dataset, a weaker
1362 mixing signal of 0.0 or 0.2 is preferable, and an excessive mixing scale can degrade performance. This
1363 suggests that the pattern prediction model itself makes much more accurate predictions on UWHVF.
1364 This observation highlights one benefit of the explicit pattern modeling used by PGDM. On a dataset
1365 as small as UWHVF, it is challenging to reliably learn the underlying distribution, as demonstrated
1366 by the high CRPS across all models on UWHVF in Table 2. In such low-data regimes, predicting in
1367 the low-dimensional pattern space is more sample-efficient, allowing the pattern prediction model to
1368 make strong point forecasts even when distributional estimation is unreliable. In this case, a designer
1369 may opt to use a higher mixing scale, leading to lowered MAE. We accomplish exactly this in Table 1.

1370 **Impact of dynamic scale.** For both the pattern guidance and pattern mixing, PGDM applies a
1371 dynamic scale $w \in (0, \bar{w})$ and $w^* \in (0, \bar{w}^*)$, respectively. In Tables 22 and 23, we compare the
1372 performances of PGDM and predictions made with a constant scale, holding \bar{w} and \bar{w}^* fixed. Across
1373 the UWHVF dataset and most genres of the AIST++ dataset, the dynamic scale reduces the prediction
1374 error of PGDM. On the ballet jazz and wack genres of AIST++, dynamic scaling performs comparably
1375 to constant scaling. It is likely that in these two cases, fewer novel patterns are seen at inference
1376 time. The dynamic scale, which is determined by the uncertainty of the pattern prediction, is most
1377 beneficial when the patterns extracted by archetypal analysis do not fully capture the patterns seen at
1378 inference time (i.e., distribution shift). The dynamic guidance scale ensures that, when novel patterns
1379 occur, the pattern guidance is not followed by the diffusion model. Similarly, the dynamic mixing
1380 scale ensures that pattern predictions on novel patterns are not heavily incorporated into the final
1381 forecast.

1382 **Impact of pattern prediction.** To address the challenge of dynamically evolving patterns, PGDM
1383 relies on a model that predicts the future patterns appearing in the sequence. The most similar related
1384 work to PGDM, Diff-MGR (Zhao et al., 2024), does not account for this realistic and common feature
1385 of temporal data. Diff-MGR assumes that patterns will remain constant over time. As we were able
1386 to obtain neither code nor sufficient implementation details to evaluate Diff-MGR on our applications,
1387 we instead perform an ablation study of PGDM where we assume constant patterns as a proxy for this
1388 baseline. In Tables 22 and 23, we compare the performances of PGDM and predictions conditioned
1389 only on the most recently observed pattern $Ac_A(x_T)$ (i.e., without the pattern prediction model f_A).
1390 Across UWHVF and all ten genres of the AIST++ dataset, PGDM with f_A outperforms PGDM
1391 without f_A . On UWHVF, the pattern prediction model adds a slight performance gain, as visual
1392 fields and their patterns may progress slowly over several decades (Saunders et al., 2016). In other
1393 cases, the performance gap from adding pattern prediction is more significant, e.g., the house, krump,
1394 and middle hip hop genres of AIST++. In these particular dance genres, patterns may change more
1395 rapidly due to dance style or music tempo. For example, the reported music tempo for the house genre
1396 is 110-130 BPM, compared to the 80-130 BPM tempo of the remaining genres (Li et al., 2021). This
1397 study highlights the importance of the pattern prediction model in accounting for dynamic patterns.

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

Table 20: Impact of maximum mixing scale \bar{w}^* on mean absolute error (MAE) of PGDM_{MAE} and PGDM_{GDE} . MAE is reported on UWHVF and all ten genres of the AIST++ dataset with fixed \bar{w} . Mean and standard deviation are taken across five samples.

	PGDM _{MAE} with Maximum Mixing Scale $\bar{w}^* =$					
	0.0	0.2	0.4	0.6	0.8	1.0
UWHVF	3.68±0.0415	3.45±0.0352	3.26±0.0293	3.11±0.0230	3.01±0.0168	2.96±0.0117
Break	0.39±0.0011	0.38±0.0009	0.41±0.0007	0.48±0.0006	0.56±0.0006	0.65±0.0006
House	0.74±0.0012	0.74±0.0014	0.83±0.0012	0.97±0.0010	1.14±0.0009	1.34±0.0008
Ballet jazz	0.39±0.0002	0.38±0.0002	0.40±0.0002	0.45±0.0002	0.52±0.0003	0.59±0.0002
Street jazz	0.48±0.0008	0.50±0.0008	0.57±0.0008	0.68±0.0007	0.80±0.0006	0.94±0.0005
Krump	0.70±0.0013	0.71±0.0007	0.80±0.0004	0.92±0.0002	1.08±0.0002	1.25±0.0002
LA Hip Hop	0.74±0.0007	0.74±0.0006	0.81±0.0005	0.93±0.0004	1.09±0.0004	1.26±0.0003
Lock	0.67±0.0005	0.67±0.0006	0.74±0.0006	0.86±0.0006	1.01±0.0005	1.18±0.0004
Middle Hip Hop	0.82±0.0009	0.82±0.0009	0.91±0.0010	1.06±0.0010	1.24±0.0010	1.45±0.0010
Pop	0.44±0.0017	0.44±0.0018	0.49±0.0017	0.56±0.0017	0.65±0.0015	0.75±0.0015
Wack	0.41±0.0031	0.39±0.0028	0.41±0.0024	0.44±0.0020	0.50±0.0015	0.56±0.0010

	PGDM _{GDE} with Maximum Mixing Scale $\bar{w}^* =$					
	0.0	0.2	0.4	0.6	0.8	1.0
UWHVF	4.00±0.0217	3.73±0.0205	3.49±0.0198	3.30±0.0193	3.16±0.0180	3.08±0.0153
Break	0.46±0.0013	0.45±0.0015	0.47±0.0017	0.52±0.0018	0.59±0.0018	0.68±0.0018
House	0.83±0.0012	0.82±0.0014	0.89±0.0013	1.01±0.0011	1.18±0.0010	1.36±0.0009
Ballet jazz	0.45±0.0010	0.42±0.0009	0.44±0.0008	0.48±0.0007	0.53±0.0006	0.61±0.0006
Street jazz	0.54±0.0005	0.55±0.0006	0.61±0.0007	0.70±0.0007	0.82±0.0006	0.95±0.0005
Krump	0.79±0.0010	0.79±0.0005	0.85±0.0003	0.96±0.0003	1.10±0.0004	1.26±0.0004
LA Hip Hop	0.82±0.0009	0.81±0.0009	0.87±0.0008	0.98±0.0007	1.12±0.0006	1.28±0.0005
Lock	0.71±0.0022	0.71±0.0023	0.77±0.0023	0.89±0.0022	1.03±0.0020	1.19±0.0019
Middle Hip Hop	0.96±0.0040	0.95±0.0038	1.02±0.0036	1.14±0.0034	1.31±0.0032	1.50±0.0030
Pop	0.49±0.0015	0.48±0.0018	0.52±0.0018	0.58±0.0018	0.67±0.0018	0.77±0.0018
Wack	0.47±0.0083	0.45±0.0079	0.46±0.0075	0.49±0.0070	0.55±0.0064	0.61±0.0057

Table 21: Impact of maximum mixing scale \bar{w}^* on continuous ranked probability score (CRPS_{SUM}) of PGDM_{MAE} and PGDM_{GDE}. CRPS_{SUM} (lower is better) is reported on UWHVF and all ten genres of the AIST++ dataset with fixed \bar{w} . Mean and standard deviation are taken across three seeds.

	PGDM _{MAE} with Maximum Mixing Scale $\bar{w}^* =$					
	0.0	0.2	0.4	0.6	0.8	1.0
UWHVF	0.795±0.0102	0.771±0.0096	0.756±0.0086	0.750±0.0071	0.756±0.0054	0.777±0.0040
Break	0.037±<0.0001	0.037±<0.0001	0.041±<0.0001	0.048±<0.0001	0.059±0.0001	0.072±0.0001
House	0.072±0.0003	0.073±0.0002	0.082±0.0002	0.096±0.0001	0.114±0.0001	0.135±<0.0001
Ballet jazz	0.036±0.0002	0.035±0.0002	0.038±0.0002	0.045±0.0001	0.055±0.0001	0.067±0.0001
Street jazz	0.047±0.0002	0.051±0.0002	0.064±0.0002	0.082±0.0002	0.104±0.0002	0.127±0.0001
Krump	0.071±0.0003	0.073±0.0003	0.081±0.0003	0.093±0.0002	0.109±0.0002	0.128±0.0001
LA Hip Hop	0.070±0.0006	0.071±0.0005	0.078±0.0004	0.090±0.0003	0.106±0.0002	0.125±0.0001
Lock	0.062±0.0001	0.063±0.0001	0.071±0.0001	0.084±0.0001	0.102±0.0002	0.122±0.0002
Middle Hip Hop	0.078±0.0002	0.081±0.0003	0.090±0.0003	0.106±0.0002	0.126±0.0002	0.149±0.0001
Pop	0.044±0.0002	0.045±0.0002	0.049±0.0002	0.056±0.0002	0.065±0.0001	0.077±0.0001
Wack	0.039±0.0003	0.039±0.0003	0.042±0.0002	0.050±0.0001	0.060±0.0001	0.072±0.0001

	PGDM _{GDE} with Maximum Mixing Scale $\bar{w}^* =$					
	0.0	0.2	0.4	0.6	0.8	1.0
UWHVF	0.888±0.0065	0.847±0.0045	0.816±0.0031	0.801±0.0024	0.803±0.0023	0.825±0.0032
Break	0.044±0.0001	0.043±0.0001	0.046±0.0001	0.053±0.0001	0.062±0.0001	0.073±0.0001
House	0.078±0.0003	0.079±0.0003	0.086±0.0001	0.098±0.0001	0.116±<0.0001	0.136±<0.0001
Ballet jazz	0.040±0.0002	0.039±0.0002	0.041±0.0001	0.047±0.0002	0.056±0.0002	0.067±0.0001
Street jazz	0.052±0.0001	0.055±0.0002	0.067±0.0002	0.084±0.0002	0.105±0.0001	0.128±0.0001
Krump	0.077±0.0003	0.078±0.0003	0.084±0.0003	0.094±0.0003	0.110±0.0003	0.128±0.0002
LA Hip Hop	0.076±0.0006	0.077±0.0005	0.083±0.0004	0.094±0.0003	0.109±0.0003	0.126±0.0002
Lock	0.065±<0.0001	0.066±0.0001	0.073±0.0001	0.086±0.0001	0.103±0.0002	0.122±0.0002
Middle Hip Hop	0.087±0.0003	0.089±0.0003	0.097±0.0003	0.111±0.0002	0.130±0.0001	0.152±0.0001
Pop	0.047±0.0001	0.047±0.0002	0.051±0.0002	0.057±0.0002	0.066±0.0001	0.077±0.0001
Wack	0.045±0.0002	0.044±0.0002	0.048±0.0002	0.055±0.0002	0.064±0.0002	0.076±0.0002

Table 22: Impact of dynamic scaling and pattern prediction model f_A on mean absolute error (MAE) of PGDM_{MAE} and PGDM_{GDE}. MAE is reported for UWHVF and all ten genres of the AIST++ dataset with fixed \bar{w} and \bar{w}^* . Mean and standard deviation are taken across five samples.

	PGDM _{MAE}			PGDM _{GDE}		
	Const. Scale	Without f_A	PGDM	Const. Scale	Without f_A	PGDM
UWHVF	3.16±0.0000	2.99±0.0101	2.96±0.0117	3.16±0.0000	3.19±0.0119	3.08±0.0153
Break	0.39±0.0010	0.41±0.0008	0.38±0.0009	0.46±0.0012	0.49±0.0025	0.45±0.0015
House	0.77±0.0013	0.84±0.0019	0.74±0.0014	0.84±0.0012	0.94±0.0011	0.82±0.0014
Ballet jazz	0.38±0.0002	0.40±0.0006	0.38±0.0002	0.42±0.0010	0.45±0.0002	0.42±0.0009
Street jazz	0.49±0.0010	0.49±0.0005	0.48±0.0008	0.56±0.0006	0.55±0.0002	0.54±0.0005
Krump	0.74±0.0009	0.92±0.0012	0.70±0.0013	0.81±0.0003	0.94±0.0009	0.79±0.0005
LA Hip Hop	0.77±0.0007	0.81±0.0005	0.74±0.0006	0.84±0.0009	0.89±0.0009	0.81±0.0009
Lock	0.69±0.0002	0.75±0.0011	0.67±0.0006	0.72±0.0015	0.79±0.0019	0.71±0.0023
Middle Hip Hop	0.86±0.0009	0.94±0.0009	0.82±0.0009	0.97±0.0022	1.10±0.0039	0.95±0.0038
Pop	0.47±0.0019	0.46±0.0020	0.44±0.0018	0.50±0.0017	0.51±0.0014	0.48±0.0018
Wack	0.39±0.0025	0.41±0.0020	0.39±0.0028	0.45±0.0074	0.47±0.0070	0.45±0.0079

1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 1523
 1524
 1525
 1526
 1527
 1528
 1529
 1530
 1531
 1532
 1533
 1534
 1535
 1536
 1537
 1538
 1539
 1540
 1541
 1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550
 1551
 1552
 1553
 1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565

Table 23: Impact of dynamic scaling and pattern prediction model f_A on continuous ranked probability score (CRPS_{SUM}) of PGDM_{MAE} and PGDM_{GDE}. CRPS_{SUM} (lower is better) is reported for UWHVF and all ten genres of the AIST++ dataset with fixed \bar{w} and \bar{w}^* . Mean and standard deviation are taken across three seeds.

	PGDM _{MAE}			PGDM _{GDE}		
	Const. Scale	Without f_A	PGDM	Const. Scale	Without f_A	PGDM
UWHVF	0.865±0.0000	0.843±0.0034	0.777±0.0040	0.865±0.0000	0.973±0.0014	0.825±0.0032
Break	0.038±0.0001	0.043±0.0001	0.037±0.0001	0.046±0.0001	0.052±0.0001	0.043±0.0001
House	0.077±0.0002	0.089±0.0003	0.073±0.0002	0.081±0.0001	0.099±0.0002	0.079±0.0003
Ballet jazz	0.036±0.0002	0.039±0.0001	0.035±0.0002	0.039±0.0002	0.042±0.0001	0.039±0.0002
Street jazz	0.049±0.0002	0.051±0.0002	0.047±0.0002	0.054±0.0001	0.056±0.0002	0.052±0.0001
Krump	0.078±0.0003	0.133±0.0007	0.071±0.0003	0.081±0.0004	0.122±0.0005	0.078±0.0003
LA Hip Hop	0.075±0.0004	0.085±0.0006	0.071±0.0005	0.080±0.0004	0.093±0.0005	0.077±0.0005
Lock	0.066±0.0002	0.079±0.0002	0.063±0.0001	0.068±0.0001	0.082±0.0001	0.066±0.0001
Middle Hip Hop	0.086±0.0003	0.107±0.0002	0.081±0.0003	0.093±0.0003	0.122±0.0006	0.089±0.0003
Pop	0.047±0.0002	0.049±0.0002	0.045±0.0002	0.049±0.0002	0.052±0.0003	0.047±0.0002
Wack	0.039±0.0002	0.044±0.0003	0.039±0.0003	0.045±0.0002	0.050±0.0003	0.044±0.0002