

Bring Your Dreams to Life: Continual Text-to-Video Customization

Jiahua Dong^{1*}, Xudong Wang^{2,3*}, Wenqi Liang⁴, Zongyan Han¹, Meng Cao¹, Duzhen Zhang¹, Hanbin Zhao^{5,6}, Zhi Han^{2†}, Salman Khan^{1,7}, Fahad Shahbaz Khan^{1,8}

¹Mohamed bin Zayed University of Artificial Intelligence

²State Key Laboratory of Robotics and Intelligent Systems, Shenyang Institute of Automation, Chinese Academy of Sciences

³University of Chinese Academy of Sciences

⁴University of Trento

⁵Zhejiang University

⁶Zhejiang Key Laboratory of Intelligent High Speed UAV Technology

⁷Australian National University

⁸Linköping University

{dongjiahua1995, liangwenqi0123}@gmail.com, {wangxudong, hanzhi}@sia.cn

{zongyan.han, meng.cao, duzhen.zhang, salman.khan, fahad.khan}@mbzuai.ac.ae, zhaohanbin@zju.edu.cn

Abstract

Customized text-to-video generation (CTVG) has recently witnessed great progress in generating tailored videos from user-specific text. However, most CTVG methods assume that personalized concepts remain static and do not expand incrementally over time. Additionally, they struggle with forgetting and concept neglect when continuously learning new concepts, including subjects and motions. To resolve the above challenges, we develop a novel Continual Customized Video Diffusion (CCVD) model, which can continuously learn new concepts to generate videos across various text-to-video generation tasks by tackling forgetting and concept neglect. To address catastrophic forgetting, we introduce a concept-specific attribute retention module and a task-aware concept aggregation strategy. They can capture the unique characteristics and identities of old concepts during training, while combining all subject and motion adapters of old concepts based on their relevance during testing. Besides, to tackle concept neglect, we develop a controllable conditional synthesis to enhance regional features and align video contexts with user conditions, by incorporating layer-specific region attention-guided noise estimation. Extensive experimental comparisons demonstrate that our CCVD outperforms existing CTVG models.

Introduction

Text-to-video diffusion models (TVDMs) (Yang et al. 2025; Wu et al. 2023; Guo et al. 2025) employ iterative denoising processes to synthesize realistic and dynamic video content directly from textual descriptions. By leveraging large-scale datasets and powerful spatiotemporal architectures (Wang and Yang 2024), they generate high-quality and contextually aligned videos that accurately reflect the semantics of the input text (Khachatryan et al. 2023). Despite significant advancements, TVDMs often struggle with fine-grained control

over object appearance and motion patterns, limiting their applicability in personalized tasks (Lee, Kwon, and Kim 2024; Wang et al. 2024b; Liu et al. 2023). To address this limitation, customized text-to-video generation (CTVG) (Jeong, Park, and Ye 2024; He et al. 2025) has been introduced to personalize video synthesis based on user-specified concepts. This involves integrating personalized subject identities and motion patterns into generated videos via the design of subject and motion adapters (Ren et al. 2025; Chen et al. 2024).

However, the majority of current CTVG methods (Yang et al. 2024; He et al. 2025; She et al. 2025) are built on the premise that personalized concepts defined by users remain static and do not grow progressively over time. Nonetheless, this premise does not align with real-world scenarios, where users often seek to continuously create and synthesize video sequences with new personalized concepts (*i.e.*, new subjects and motion patterns) based on their evolving experiences. To tackle this limitation, CTVG methods (Huang et al. 2025a; Zhou et al. 2024; Gu et al. 2024) generally need to retain all text-image training data associated with previous concepts to finetune pretrained TVDMs (Xing et al. 2024) using the subject and motion adapters (Wei et al. 2024b). Yet, the substantial computational expenses and privacy issues (Van Le et al. 2023) can render them unfeasible as the quantity of personalized concepts gradually increases over time. If they preserve all subject and motion adapters tied to previously customized concepts and merge them to continuously adapt to new concepts (Wu, Huang, and Wei 2024; Zhong et al. 2024), they face a substantial decline in the distinct characteristics of old concepts (*i.e.*, catastrophic forgetting (Lopez-Paz and Ranzato 2017; Ostapenko et al. 2019)). Besides, in practical applications, users typically seek to control the contexts (*i.e.*, subjects and motions) linked to multiple old concepts in the generated videos by supplying additional conditions (*e.g.*, bounding boxes). This imposes substantial strain on the above CTVG methods (Wei et al. 2024b; Zhao et al. 2025), making them particularly vulnerable to the issue of concept neglect (Sun et al. 2024) (*i.e.*, certain old concepts may be

*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

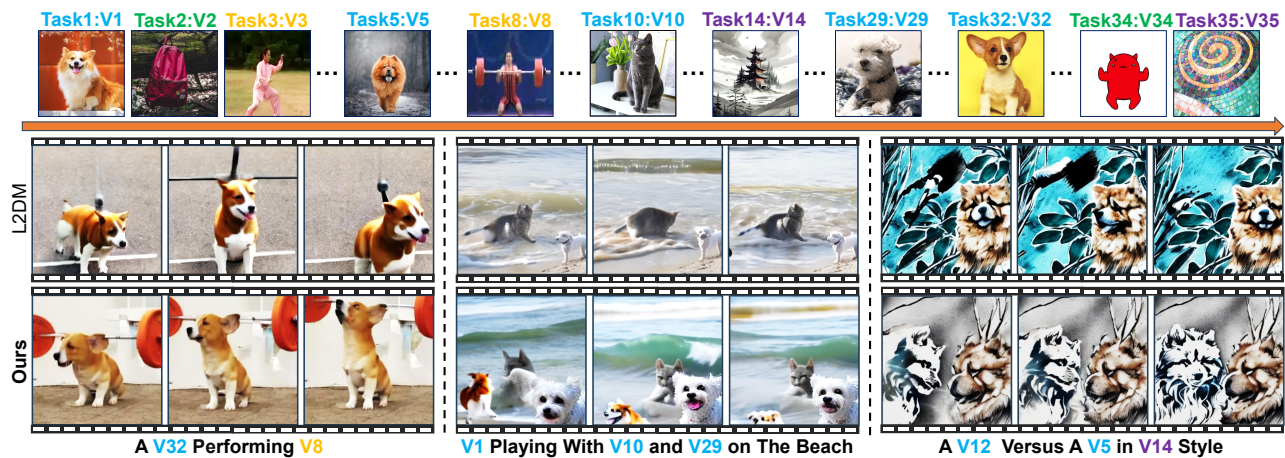


Figure 1: Results of our model in continually learning new concepts, including subjects and motions, under the CTVC setting.

overlooked in the multi-concept video customization).

To address the aforementioned real-world challenges, this paper introduces a novel practical problem referred to as Continual Text-to-Video Customization (CTVC). Under the CTVC setting, as illustrated in Fig. 1, we can consecutively learn new personalized concepts (*i.e.*, subjects and motion patterns) to generate video sequences across a variety of text-to-video generation tasks, including single/multi-concept video customization (Huang et al. 2025b), style transfer (Gao et al. 2020) and video editing (Ceylan, Huang, and Mitra 2023). Moreover, users can control the synthesis of video contexts, comprising both subjects and motion patterns, for multi-concept video customization by specifying certain conditions (*e.g.*, bounding boxes). As previously mentioned, there are two challenges to conduct various text-to-video synthesis tasks in CTVC: **catastrophic forgetting** of previously learned concepts when continually learning new ones, and **concept neglect** when performing multi-concept video customization based on user-specified conditions.

To the CTVC problem, we propose a novel Continual Customized Video Diffusion (CCVD) model, which is the first work to tackle catastrophic forgetting and concept neglect across a variety of CTVC tasks. **First**, to overcome catastrophic forgetting of previously learned personalized concepts, we develop a concept-specific attribute retention (CAR) module for training and a task-aware concept aggregation (TCR) strategy for testing. The CAR module incorporates layer-specific concept tokens and a concept orthogonal loss to explore the unique characteristics and identity of personalized concepts, while the TCR module merges all subject and motion adapters of old concepts, weighted by their relevance to CTVC tasks. **Second**, we design a controllable conditional synthesis module to address concept neglect in the multi-concept video customization task. It employs layer-specific region attention to enhance the representational capacity of latent regional features and utilizes attention-guided noise estimation to control the context of personalized subjects and motions within the generated video, aligning with user-specified conditions. Experiments demonstrate the effectiveness of our CCVD. The key contributions are as follows:

- We introduce a new problem termed Continual Text-to-Video Customization (CTVC), characterized by two key challenges: catastrophic forgetting and concept neglect.
- We develop a Continual Customized Video Diffusion (CCVD) model, capable of continuously learning new concepts while tackling forgetting and concept neglect.
- We propose a concept-specific attribute retention module and a task-aware concept aggregation strategy to mitigate catastrophic forgetting of previously learned concepts.
- We devise a controllable conditional synthesis module to address concept neglect by leveraging layer-specific region attention and attention-guided noise estimation.

Related Work

Customized text-to-video generation (CTVG) integrates user-specific subjects and motion patterns to synthesize personalized videos while maintaining contextual consistency and visual quality. Following previous works (Gal et al. 2023; Dong et al. 2024b), which addresses subject-driven and style-driven concept customization by leveraging a few images to finetune Stable Diffusion (Takagi and Nishimoto 2023; Wang et al. 2024a), recent studies (Wei et al. 2024b; Jiang et al. 2024; Ren et al. 2025) have extended them into the temporal domain for customized text-to-video generation. DreamVideo (Wei et al. 2024b) introduces a two-stage learning framework that decouples the subject learning from motion learning. MotionBooth (Wu et al. 2025a) incorporates subject region loss and video preservation loss to enhance alignment between the input concept and generated visual content. Additionally, (Chen et al. 2024; She et al. 2025; Liang et al. 2025) have addressed the more complex challenge of multi-concept video customization. ConceptMaster (Huang et al. 2025b) tackles multi-concept video customization via decoupled concept embeddings and systematic data construction. However, existing methods (Fei et al. 2025; Jiang et al. 2024) encounter the issues of catastrophic forgetting and concept neglect when continually learning new concepts for video customization.

Continual Learning (CL) (Li and Hoiem 2017; Toldo and Ozay 2022) enables models to incrementally acquire new

knowledge while retaining previously learned information without retraining from scratch. Some regularization methods incorporate constraint terms to limit modifications to network weights associated with previous tasks (Kirkpatrick et al. 2017; Dong et al. 2022), or employ model distillation (Li and Hoiem 2016; Wei et al. 2024a) to transfer knowledge between previous and current models. Approaches based on data replay preserve samples from old categories (Belouadah and Popescu 2019) or utilize generative models to produce synthetic data belonging to old classes (Qi et al. 2025). Structural adaptation methods (Yoon et al. 2018; Douillard et al. 2022) progressively modify their architectures to incorporate new information while maintaining previous knowledge. Besides, feature-based methods leverage contrastive mechanisms (Madaan et al. 2022) or prompt techniques (Masum et al. 2024; Smith et al. 2023) to facilitate ongoing task acquisition. Yet, these approaches (Wei, Xu, and Deng 2025; Dong et al. 2024a) cannot be directly employed to address continual text-to-video customization problem.

Preliminary and Problem Definition

Preliminary: Text-to-video diffusion models (TVDMs) (Wang et al. 2023a; Blattmann et al. 2023; Wang et al. 2023b; Wu et al. 2025b) use textual prompts as input to generate videos that visually match the content described in the prompts. They synthesize videos in the latent feature space by introducing an encoder $\mathcal{E}(\cdot)$ and a decoder $\mathcal{D}(\cdot)$. Customized text-to-video generation (CTVG) (Wei et al. 2024b; Chefer et al. 2024) uses a subject adapter and a motion adapter to generate customized videos by finetuning the pretrained TVDMs (Lee, Kwon, and Kim 2024; Wang et al. 2024b). Given a personalized image \mathbf{x} and its text prompt \mathbf{p} , the encoder $\mathcal{E}(\cdot)$ projects \mathbf{x} to a latent feature \mathbf{z} , and the text encoder $\Phi(\cdot)$ encodes \mathbf{p} to the textual embedding $\mathbf{c} = \Phi(\mathbf{p})$. Here, \mathbf{x} is a personalized pattern. The objective to synthesize a custom video for the given concept $\{\mathbf{x}, \mathbf{p}\}$ is defined as:

$$\mathcal{L}_{\text{CTVG}} = \mathbb{E}_{\mathbf{z} \sim \mathcal{E}(\mathbf{x}), \mathbf{c}, \epsilon \sim \mathcal{N}(0, \mathbf{I}), t} [\|\epsilon - \epsilon_{\theta}(\mathbf{z}_t, \mathbf{c}, t)\|_2^2], \quad (1)$$

where \mathbf{z}_t is the noisy latent feature at the t -th timestep. $\epsilon_{\theta}(\cdot)$ represents a 3D UNet (Wang et al. 2023a) that predicts the noise estimation $\epsilon_{\theta}(\mathbf{z}_t, \mathbf{c}, t)$, guiding Gaussian noise $\mathcal{N}(0, \mathbf{I})$.

In $\epsilon_{\theta}(\cdot)$, $\theta = \theta_0 + \Delta\theta$ and each layer contains a temporal transformer block (TTB) and a spatial transformer block (STB). Here, $\theta_0 = \{\mathbf{W}_l^0\}_{l=1}^L$ denotes the parameters of pretrained TVDMs, and $\Delta\theta = \{\Delta\mathbf{W}_l\}_{l=1}^L$ indicates the parameters of the subject adapter or motion adapter. In the l -th ($l = 1, \dots, L$) layer of $\epsilon_{\theta}(\cdot)$, $\mathbf{W}_l^0 \in \mathbb{R}^{r \times s}$ is the pretrained weight, and $\Delta\mathbf{W}_l \in \mathbb{R}^{r \times s}$ denotes the weight of the subject adapter or motion adapter. r and s are the matrix dimension. As proposed in (Wei et al. 2024b), $\Delta\mathbf{W}_l = \mathbf{R}_l \mathbf{S}_l$ consists of a down-projected layer $\mathbf{R}_l \in \mathbb{R}^{r \times b}$ and an up-projected layer $\mathbf{S}_l \in \mathbb{R}^{b \times s}$, where b is the feature dimension. As shown in Fig. 2(a), if \mathbf{x} is a personalized subject, $\Delta\mathbf{W}_l$ is placed in the STB; otherwise, it is placed in the TTB to learn the motion pattern. However, most existing CTVG methods (Wei et al. 2024b; Huang et al. 2025b) presume that the number of personalized concepts (*i.e.*, subjects and motion patterns) for each user stays fixed over time. This assumption does

not hold in practical scenarios, where users typically want to continuously synthesize new personalized concepts according to their evolving experiences. Furthermore, they face the challenge of catastrophic forgetting on previously learned personalized concepts and experience concept neglect when synthesizing customized videos in concept-incremental way.

Problem Definition: For the above challenges, we propose a new Continual Text-to-Video Customization (CTVC) problem. Define U continuous customized text-to-video generation tasks as $\mathcal{T} = \{\mathcal{T}^u\}_{u=1}^U$, where the u -th task $\mathcal{T}^u = \{\mathbf{x}_i^u, \mathbf{y}_i^u, \mathbf{p}_i^u\}_{i=1}^{n_u}$ comprises n_u images \mathbf{x}_i^u , concept tokens $\mathbf{y}_i^u \in \mathcal{Y}^u$ and their corresponding text prompts \mathbf{p}_i^u . In this paper, \mathbf{p}_i^u refers to the textual description of a subject or motion pattern (*e.g.*, $[V_*]$ $[V_{\text{cat}}]$ swims in the pool or a boy $[V_{\text{plays}}]$ football), and \mathbf{y}_i^u indicates the concept token (*e.g.*, $[V_*]$ $[V_{\text{cat}}]$ or $[V_{\text{plays}}]$) in \mathbf{p}_i^u . \mathcal{Y}^u denotes the concept space of the u -th task, which consists of K^u new concepts $\mathbf{y}^u = \cup_{i=1}^{n_u} \mathbf{y}_i^u$. In CTVC, the concept space \mathcal{Y}^u of the u -th task does not overlap with any previous tasks: $\mathcal{Y}^u \cap (\cup_{j=1}^{u-1} \mathcal{Y}^j) = \emptyset$. It indicates that K^u new concepts in \mathcal{T}^u are different from $\sum_{j=1}^{u-1} K^j$ previously learned concepts from $\cup_{j=1}^{u-1} \mathcal{T}^j$. According to users' preference, the u -th task \mathcal{T}^u may be one of the following versatile video customization tasks: single/multi-concept video synthesis (Huang et al. 2025b), video editing (Feng et al. 2024), and style transfer (Chen et al. 2017). Given practical and privacy concerns, we do not replay or store any data from old tasks, allowing us to learn all video customization tasks in an incremental way. CTVC aims to continually learn a sequence of concepts while alleviating forgetting and concept neglect of old concepts.

The Proposed Model

The architecture of CCVD, designed to resolve the CTVC problem, is illustrated in Fig. 2. It comprises three key components: (a) a concept-specific attribute retention, (b) a task-aware concept aggregation strategy to mitigate catastrophic forgetting, and (c) a controllable conditional synthesis module, which includes layer-specific region attention and attention-guided noise estimation to solve concept neglect.

Concept-Specific Attribute Retention

Motivated by (Wei et al. 2024b), if the K^u new concepts in the u -th task \mathcal{T}^u belong to subjects, we place the subject adapter in the spatial transformer block (STB) to finetune the pretrained 3D UNet $\epsilon_{\theta_0}(\cdot)$ via Eq. (1) for learning \mathcal{T}^u . Otherwise, we place the motion adapter in the temporal transformer block (TTB) to learn personalized motion patterns in \mathcal{T}^u . After learning the u -th ($u = 1, \dots, U$) task, we obtain an updated 3D UNet $\epsilon_{\theta_u}(\cdot)$, where $\theta_u = \theta_0 + \Delta\theta_u$. Here, $\Delta\theta_u = \{\Delta\mathbf{W}_l^u\}_{l=1}^L$ denotes the updated parameters of the subject adapter or motion adapter of the u -th task. $\Delta\mathbf{W}_l^u = \mathbf{R}_l^u \mathbf{S}_l^u$ is the updated weight in the l -th layer of $\epsilon_{\theta_u}(\cdot)$, $\mathbf{R}_l^u \in \mathbb{R}^{r \times b}$ and $\mathbf{S}_l^u \in \mathbb{R}^{b \times s}$ are the down-projected and up-projected layers. Intuitively, $\Delta\theta_u$ is capable of capturing the majority of personalized concept identity in \mathcal{T}^u . In light of this, a straightforward approach to resolving the CTVC problem involves retaining all updated subject and motion adapters $\{\Delta\theta_j\}_{j=1}^{u-1}$ learned so far, and combining

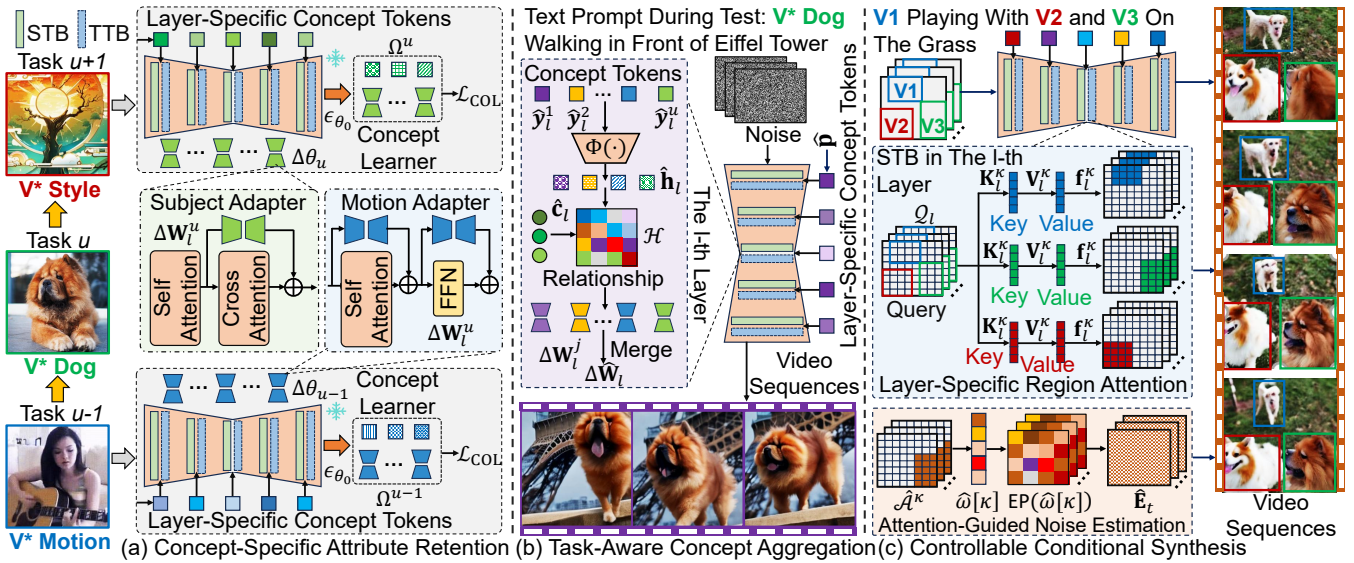


Figure 2: Concept-Specific Attribute Retention. It includes (a) a concept-specific attribute retention module, (b) a task-aware concept aggregation strategy to overcome catastrophic forgetting of previous concepts during training and testing, and (c) a controllable conditional synthesis with layer-specific region attention and attention-guided noise estimation

them linearly by assessing their respective contributions to learning personalized concepts. However, continually learning new concepts in the CTVC setting may cause this trivial approach to suffer a significant loss of unique attributes in certain personalized concepts. This issue is commonly termed catastrophic forgetting (Smith et al. 2024; Sun et al. 2024) of previously learned concepts. To alleviate catastrophic forgetting, as depicted in Fig. 2(a), we propose a concept-specific attribute retention module. This module incorporates layer-specific concept tokens and a concept orthogonal loss to capture unique concept attributes within different tasks.

Layer-Specific Concept Tokens: Unlike DreamVideo (Wei et al. 2024b), which uses textual inversion (Gal et al. 2023; Zhang et al. 2023; Voynov et al. 2023) to apply a uniform text prompt across all layers, we propose layer-specific concept tokens to enhance the preservation of unique attributes associated with old concepts. At the u -th video customization task, L layer-specific concept tokens $\{\mathbf{y}_{i,l}^u\}_{l=1}^L$ are defined for any triplet $\{\mathbf{x}_i^u, \mathbf{y}_i^u, \mathbf{p}_i^u\} \in \mathcal{T}^u$. Specifically, given a textual description \mathbf{p}_i^u ($[V_*] [V_{\text{cat}}]$ swims in the pool), we denote the layer-specific text prompt as $\mathbf{p}_{i,l}^u$ ($[V_*^l] [V_{\text{cat}}^l]$ swims in the pool), and the layer-specific concept token as $\mathbf{y}_{i,l}^u$ ($[V_*^l] [V_{\text{cat}}^l]$) at the l -th layer of $\epsilon_{\theta_u}(\cdot)$. Subsequently, we initialize $\{\mathbf{p}_{i,l}^u\}_{l=1}^L$ using \mathbf{p}_i^u and integrate the textual embedding $\mathbf{c}_{i,l}^u = \Phi(\mathbf{p}_{i,l}^u)$, which is encoded through the text encoder $\Phi(\cdot)$, into the l -th layer of $\epsilon_{\theta_u}(\cdot)$. When we train $\epsilon_{\theta_u}(\cdot)$ via Eq. (1) to synthesize customized videos, the layer-specific concept tokens are effective to explore distinctive and inherent identity attributes of previously learned concepts, thereby overcoming the catastrophic forgetting.

Concept Orthogonal Loss: Under the CTVC setting, the discriminative capacity of unique identity characteristics to differentiate between various personalized concepts may sub-

stantially degrade as the continual learning of a series of video customization tasks progresses, particularly for semantically similar concepts. To resolve this issue, we introduce a concept orthogonal loss \mathcal{L}_{COL} to perform orthogonality constrain among different concept subspaces across video customization tasks. Specifically, the updated weight $\Delta \mathbf{W}_l^u = \mathbf{R}_l^u \mathbf{S}_l^u \in \mathbb{R}^{r \times s}$ at the l -th layer of $\epsilon_{\theta_u}(\cdot)$ can be interpreted as comprising two key components: the concept subspace of the u -th task $\mathbf{R}_l^u = [\mathbf{r}_l^{u,1}, \dots, \mathbf{r}_l^{u,b}] \in \mathbb{R}^{r \times b}$, and its corresponding weighting coefficient $\mathbf{S}_l^u = [s_l^{u,1}, \dots, s_l^{u,b}]^\top \in \mathbb{R}^{b \times s}$. Here, $\mathbf{r}_l^{u,i} \in \mathbb{R}^r$ is the i -th ($i = 1, \dots, b$) basis vector of the concept space \mathbf{R}_l^u and $s_l^{u,i} \in \mathbb{R}^s$ denotes its linear coefficient. When learning the u -th video customization task \mathcal{T}^u , we want the concept subspace of the u -th task to be orthogonal to the concept subspaces of previous tasks $\cup_{j=1}^{u-1} \mathcal{T}^j$: $\sum_{j=1}^{u-1} \sum_{l=1}^L \text{tr}(\mathbf{R}_l^j (\mathbf{R}_l^u)^\top) = 0$. We devise an alternative optimization approach aimed at minimizing the absolute value of the inner product among distinct subspaces. As a result, we formulate the concept orthogonal loss \mathcal{L}_{COL} as follows:

$$\mathcal{L}_{\text{COL}} = \mathbb{E}_{\mathbf{z} \sim \mathcal{E}(\mathbf{x}), \mathbf{c}, \epsilon \sim \mathcal{N}(0, \mathbf{I}), t} [\|\epsilon - \epsilon_{\theta_u}(\mathbf{z}_t, \mathbf{c}, t)\|_2^2 + \lambda \sum_{j=1}^{u-1} \sum_{l=1}^L \text{tr}(\mathbf{R}_l^j (\mathbf{R}_l^u)^\top)], \quad (2)$$

where λ is a hyper-parameter. The loss \mathcal{L}_{COL} helps the layer-specific concept tokens better explore the distinctive attributes of previous concepts, especially for semantically similar ones, by reducing correlations among concept subspaces. Upon finishing the learning of the u -th video customization task via Eq. (2), we obtain the task-wise concept learner $\Omega^u = \{\Delta \mathbf{W}_l^u, \hat{\mathbf{y}}_l^u\}_{l=1}^L$, where $\hat{\mathbf{y}}_l^u = \{\hat{\mathbf{y}}_l^{u,k}\}_{k=1}^{K^u}$ and $\hat{\mathbf{y}}_l^{u,k}$ is the k -th learned concept tokens at the l -th layer.

Task-Aware Concept Aggregation

In order to alleviate catastrophic forgetting of previously learned concept during the testing phase, we retain all task-wise concept learners $\{\Omega^j\}_{j=1}^u$ learned so far, and propose a task-aware concept aggregation strategy to dynamically aggregate these concept learners for customized text-to-video generation tasks. As in Fig. 2(b), for a text prompt $\hat{\mathbf{p}}$ during testing, we first initialize the layer-specific text prompt $\{\hat{\mathbf{p}}_l\}_{l=1}^L$ using $\hat{\mathbf{p}}$, and then generate layer-specific textual embeddings $\hat{\mathbf{c}} = \{\hat{\mathbf{c}}_l \in \mathbb{R}^{n_e \times d}\}_{l=1}^L$ through the text encoder $\Phi(\cdot)$. Here, n_e is the number of tokens and d denotes the embedding dimension. Subsequently, we gather all concept tokens $\{\mathcal{Z}_l\}_{l=1}^L$ from stored concept learners $\{\Omega^j\}_{j=1}^u$, where $\mathcal{Z}_l = \cup_{j=1}^u \hat{\mathbf{V}}_l^j \in \mathbb{R}^{n_t}$ comprises $n_t = \sum_j K^j$ layer-specific concept tokens at the l -th layer. After projecting $\{\mathcal{Z}_l\}_{l=1}^L$ to latent embeddings $\{\mathbf{h}_l \in \mathbb{R}^{n_t \times d}\}_{l=1}^L$ using $\Phi(\cdot)$, we compute the average of these latent embeddings corresponding to the same task, resulting in $\{\hat{\mathbf{h}}_l \in \mathbb{R}^{u \times d}\}_{l=1}^L$. u is the number of tasks learned up until now. After calculating the relationships $\mathcal{H} \in \mathbb{R}^u$ between $\hat{\mathbf{h}}_l$ and $\hat{\mathbf{c}}_l$, we leverage \mathcal{H} to dynamically aggregate $\{\Delta \mathbf{W}_l^j\}_{j=1}^u$ at the l -th layer to obtain $\Delta \widehat{\mathbf{W}}_l$:

$$\mathcal{H} = \max(\hat{\mathbf{c}}_l \cdot (\hat{\mathbf{h}}_l)^\top), \Delta \widehat{\mathbf{W}}_l = \sum_{j=1}^u \Delta \mathbf{W}_l^j \cdot \zeta(\mathcal{H})_j, \quad (3)$$

where $\max(\cdot)$ is the maximization along the row axis. The normalization of \mathcal{H} is achieved by $\zeta(\mathcal{H}) = \mathcal{H}^2 / \|\mathcal{H}^2\|_F \in \mathbb{R}^u$, and $\zeta(\mathcal{H})_j$ corresponds to the j -th element of $\zeta(\mathcal{H})$.

Test Model: After computing $\Delta \theta_* = \{\Delta \widehat{\mathbf{W}}_l\}_{l=1}^L$ via Eq. (3), we derive a new 3D UNet $\epsilon_{\theta_*}(\cdot)$ for testing, where $\theta_* = \theta_0 + \Delta \theta_*$. At the l -th layer of $\epsilon_{\theta_*}(\cdot)$, if the text prompt $\hat{\mathbf{p}}$ belongs to a subject, we position $\Delta \widehat{\mathbf{W}}_l$ in the STB; if $\hat{\mathbf{p}}$ describes a motion pattern, we place $\Delta \widehat{\mathbf{W}}_l$ in TTB; if $\hat{\mathbf{p}}$ encompasses both subject and motion, we position $\Delta \widehat{\mathbf{W}}_l$ in STB and TTB respectively. Particularly, $\epsilon_{\theta_*}(\cdot)$ has incorporated significant unique attributes of all concepts learned up until now, effectively reducing catastrophic forgetting of previously learned concepts during the testing phase.

Controllable Conditional Synthesis

When employing for multi-concept customization, it fails to produce high-fidelity videos that align with user-specified conditions (e.g., bounding boxes), due to concept neglect (i.e., some subjects or motions are omitted in the generated videos). To tackle this issue, as depicted in Fig. 2(c), we develop a layer-specific region attention and an attention-guided noise estimation for customized multi-concept video generation.

Layer-Specific Region Attention: In addition to the initial text prompt $\hat{\mathbf{p}}$, users can utilize n_r regional conditions $\{\hat{\mathbf{p}}^\kappa, \hat{\mathbf{b}}^\kappa\}_{\kappa=1}^{n_r}$ to control the video context according to their preference. Here, $\hat{\mathbf{b}}^\kappa \in \mathbb{R}^{n_v \times 4}$ is the bounding box, which is used to generate personalized concept associated with the κ -th region text prompt $\hat{\mathbf{p}}^\kappa$. n_v is the number of frames in the synthesized videos. After obtaining the layer-specific textual embedding $\hat{\mathbf{c}}^\kappa = \{\hat{\mathbf{c}}_l^\kappa \in \mathbb{R}^{n_e \times d}\}_{l=1}^L$ for $\hat{\mathbf{p}}^\kappa$ using $\Phi(\cdot)$, we leverage $\{\hat{\mathbf{c}}_l^\kappa\}_{l=1}^L$ and the new 3D UNet $\epsilon_{\theta_*}(\cdot)$ to extract the latent features $\{\mathbf{f}_l \in \mathbb{R}^{n_v \times h_l \times w_l \times d}\}_{l=1}^L$ for $\hat{\mathbf{p}}$. w_l and h_l are

the width and height of the feature \mathbf{f}_l . At the l -th layer, we further conduct layer-specific region attention between \mathbf{f}_l and $\hat{\mathbf{c}}_l^\kappa$ to derive the κ -th region feature $\mathbf{f}_l^\kappa \in \mathbb{R}^{n_v \times h_l^\kappa \times w_l^\kappa \times d}$:

$$\mathbf{f}_l^\kappa = \mathcal{A}_l^\kappa \cdot \mathbf{V}_l^\kappa = \sigma(\mathbf{Q}_l(\mathbf{K}_l^\kappa)^\top / \sqrt{d} \cdot \mathbf{V}_l^\kappa), \quad (4)$$

where $\mathcal{A}_l^\kappa = \sigma(\mathbf{Q}_l(\mathbf{K}_l^\kappa)^\top / \sqrt{d}) \in \mathbb{R}^{n_v \times h_l^\kappa \times w_l^\kappa \times n_e}$ denotes the κ -th region attention, $\sigma(\cdot)$ is the sigmoid function, w_l^κ and h_l^κ represent the width and height of the κ -th bounding box $\hat{\mathbf{b}}^\kappa$. Moreover, $\mathbf{Q}_l \in \mathbb{R}^{n_v \times h_l^\kappa \times w_l^\kappa \times d}$, $\mathbf{K}_l^\kappa \in \mathbb{R}^{n_e \times d}$ and $\mathbf{V}_l^\kappa \in \mathbb{R}^{n_e \times d}$ denote the query, key and value matrices:

$$\mathbf{Q}_l = \Gamma(\mathbf{f}_l \mathbf{w}_q \otimes \hat{\mathbf{m}}_l^\kappa), \mathbf{K}_l^\kappa = \hat{\mathbf{c}}_l^\kappa \mathbf{w}_k, \mathbf{V}_l^\kappa = \hat{\mathbf{c}}_l^\kappa \mathbf{w}_v, \quad (5)$$

where $\hat{\mathbf{m}}_l^\kappa \in \mathbb{R}^{n_v \times h_l \times w_l}$ represents the binary mask of the κ -th region, with all values within the κ -th bounding box $\hat{\mathbf{b}}^\kappa$ assigned to 1. \otimes denotes the Hardward product, $\Gamma(\cdot)$ preserves exclusively the features located within $\hat{\mathbf{b}}^\kappa$. $\mathbf{w}_q, \mathbf{w}_k, \mathbf{w}_v \in \mathbb{R}^{d \times d}$ are the projection matrices in $\epsilon_{\theta_*}(\cdot)$.

Subsequently, we substitute the values of \mathbf{f}_l within the bounding boxes $\{\hat{\mathbf{b}}^\kappa\}_{\kappa=1}^{n_r}$ with their corresponding region features $\{\mathbf{f}_l^\kappa\}_{\kappa=1}^{n_r}$, resulting in a new feature map $\hat{\mathbf{f}}_l$ at the l -th layer. The layer-specific region attention is then applied to every layer of $\epsilon_{\theta_*}(\cdot)$ for multi-concept video customization.

Attention-Guided Noise Estimation: Given the initial text prompt $\hat{\mathbf{p}}$, motivated by (Ho and Salimans 2021), we can derive the noise estimation $\mathbf{E}_t = \epsilon_{\theta_*}(\mathbf{z}_t, t) + \tau \cdot (\epsilon_{\theta_*}(\mathbf{z}_t, \hat{\mathbf{c}}, t) - \epsilon_{\theta_*}(\mathbf{z}_t, t)) \in \mathbb{R}^{n_v \times h_L \times w_L \times d_L}$ at the t -th timescale, by incorporating the aforementioned layer-wise region attention into $\epsilon_{\theta_*}(\cdot)$. Here, \mathbf{E}_t is the output of the L -th layer of $\epsilon_{\theta_*}(\cdot)$, and $\tau = 7.5$ is the scaling factor. w_L, h_L and d_L represent the width, height and channel dimension of the noise estimation \mathbf{E}_t . $\epsilon_{\theta_*}(\mathbf{z}_t, t)$ denotes the unconditional noise estimation, while $\epsilon_{\theta_*}(\mathbf{z}_t, \hat{\mathbf{c}}, t)$ represent noise estimation conditioned on $\hat{\mathbf{c}}$. As a result, the κ -th ($\kappa = 1, \dots, n_r$) region noise estimation $\mathbf{E}_t^\kappa \in \mathbb{R}^{n_v \times h_L \times w_L \times d_L}$ is expressed:

$$\mathbf{E}_t^\kappa = \epsilon_{\theta_*}(\mathbf{z}_t, t) + \tau \cdot (\epsilon_{\theta_*}(\mathbf{z}_t, [\hat{\mathbf{c}}^\kappa, \hat{\mathbf{b}}^\kappa], t) - \epsilon_{\theta_*}(\mathbf{z}_t, t)), \quad (6)$$

where $\epsilon_{\theta_*}(\mathbf{z}_t, [\hat{\mathbf{c}}^\kappa, \hat{\mathbf{b}}^\kappa], t)$ represents the noise estimation based on the κ -th regional condition $[\hat{\mathbf{c}}^\kappa, \hat{\mathbf{b}}^\kappa]$.

To address the challenge of concept neglect in multi-concept video customization, we aim to place greater emphasis on personalized concepts that are easily overlooked. In light of this, we propose an attention-guided noise estimation module that leverages region attention to assign larger weights to easily overlooked concepts when aggregating n_r region estimations $\{\mathbf{E}_t^\kappa\}_{\kappa=1}^{n_r}$. Specifically, for the κ -th ($\kappa = 1, \dots, n_r$) region condition, we upsample L layer-specific region attentions $\{\mathcal{A}_l^\kappa \in \mathbb{R}^{n_v \times h_l^\kappa \times w_l^\kappa \times n_e}\}_{l=1}^L$ to $\{\hat{\mathcal{A}}_l^\kappa \in \mathbb{R}^{n_v \times h_L^\kappa \times w_L^\kappa \times n_e}\}_{l=1}^L$, and average them to obtain $\hat{\mathcal{A}}^\kappa = \frac{1}{L} \sum_{l=1}^L \hat{\mathcal{A}}_l^\kappa$. Subsequently, we select the maximum attention value $\omega^\kappa = \max(\hat{\mathcal{A}}^\kappa) \in \mathbb{R}^{n_v}$ to reweight the noise estimation \mathbf{E}_t^κ . Intuitively, if the attention value ω^κ of the κ -th region ($\kappa = 1, \dots, n_r$) is small, it indicates that the corresponding concept is more likely to be overlooked in video customization. Thus, we assign more weight to the noise estimation with a small attention value. To achieve this,

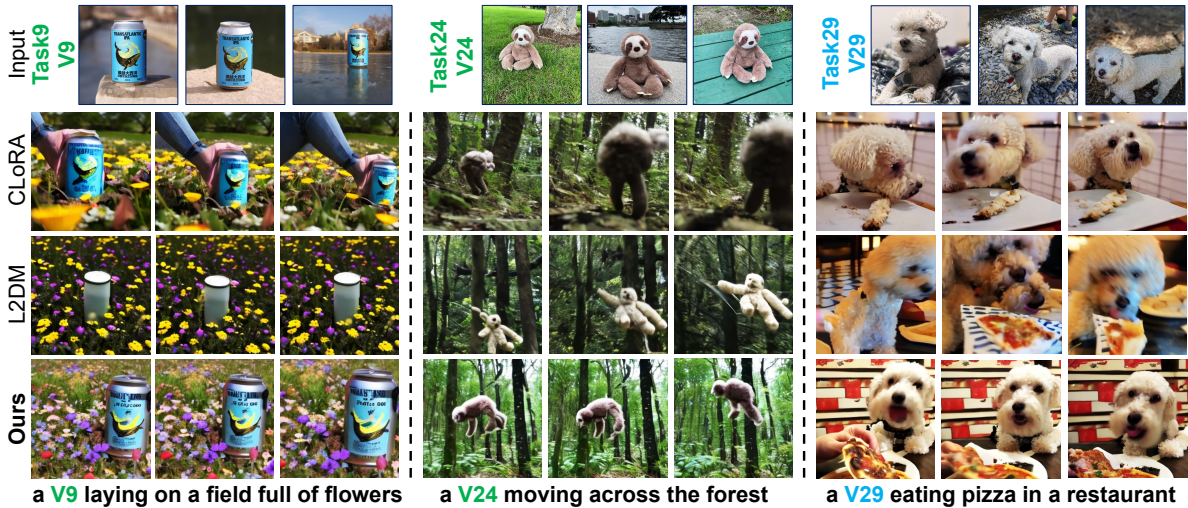


Figure 3: Comparison results of single-concept video customization under the CTVC setting.

we concatenate $\{\omega^\kappa\}_{\kappa=1}^{n_r}$ as $\omega \in \mathbb{R}^{n_v \times n_r}$ and normalize ω as $\hat{\omega} \in \mathbb{R}^{n_v \times n_r}$. To resolve concept neglect, we employ $\hat{\omega}$ to aggregate n_r region noise estimation to obtain $\hat{\mathbf{E}}_t$:

$$\hat{\mathbf{E}}_t = \mathbf{E}_t + \sum_{\kappa=1}^{n_r} (1 - \text{EP}(\hat{\omega}[\kappa])) \otimes \mathbf{E}_t^\kappa \otimes \hat{\mathbf{m}}_L^\kappa, \quad (7)$$

where $\hat{\omega}[\kappa] \in \mathbb{R}^{n_v}$ is the κ -th column of $\hat{\omega}$, and $\text{EP}(\hat{\omega}[\kappa])$ expands $\hat{\omega}[\kappa]$ as $\mathbb{R}^{n_v \times h_L \times w_L \times d_L}$. $\hat{\mathbf{m}}_L^\kappa \in \mathbb{R}^{n_v \times h_L \times w_L}$ represents the binary region mask corresponding to the bounding box $\hat{\mathbf{b}}^\kappa \in \mathbb{R}^{n_v \times 4}$ at the L -th layer of $\epsilon_{\theta_*}(\cdot)$. Finally, we follow (Wei et al. 2024b) and feed $\hat{\mathbf{E}}_t$ into UNet $\epsilon_{\theta_*}(\cdot)$ to perform multi-concept customization via gradual denoising.

Experiments

Implementation Details

Dataset: Inspired by (Chefer et al. 2024; Wei et al. 2024b), we introduce a new continual video customization (CVC) dataset including 35 sequential text-to-video customization tasks to evaluate the effectiveness of our model under the CTVC setting. **Setups:** For fair comparisons, our model and all comparative methods utilize DreamVideo (Wei et al. 2024b) as the baseline backbone. We use the Adam optimizer for training, with an initial learning rate of 1.0×10^{-4} for textual embeddings and 1.0×10^{-5} for the 3D UNet. We set $\lambda = 0.1$ in Eq. (2), and $n_v = 32$ to synthesize customized videos. **Metrics:** We perform comprehensive qualitative and quantitative evaluations across various synthesis tasks including single/multi-concept video customization, style transfer and video editing. For quantitative comparisons, we follow (Wei et al. 2024b) to utilize CLIPT, CLIPI, DINOI, TCons as evaluation metrics. Moreover, we use the forgetting of CLIPT, CLIPI, and DINOI (denoted as F-CLIPT, F-CLIPI, and F-DINOI) (Sun et al. 2024) to assess CTVC performance. Notably, more details about the CVC dataset and the comparison methods are provided in the supplementary material.

Methods	CLIPT	CLIPI	DINOI	TCons	F-CLIPT	F-CLIPI	F-DINO
Finetuning	22.9	52.3	25.1	96.3	9.5	14.5	11.2
EWC	28.1	58.4	29.5	96.1	5.4	7.8	6.3
LWF	24.8	59.3	30.0	96.2	4.3	6.8	5.8
LoRA-M	25.5	57.8	29.0	96.1	5.2	8.4	6.2
CLoRA	28.7	60.7	31.0	96.1	3.5	5.3	4.1
L2DM	28.3	62.3	31.2	96.2	3.1	3.9	3.7
CCVD	29.9	64.3	33.4	96.4	1.1	2.3	1.9

Table 1: Results of single-concept video customization.

Qualitative Comparisons

To demonstrate the superiority of our model, as in Fig. 3–6, we conduct qualitative comparisons on single/multi-concept video customization, style transfer and video editing. 1) In Fig. 3, our model surpasses other methods in single-concept video customization, since it can mitigate forgetting by incorporating concept-specific attribute retention module to explore unique concept identity. 2) In Fig. 4, all comparative methods fail to generate multiple concepts, while our model can leverage layer-specific region attention and attention-guided noise estimation to synthesize videos with multiple user-specified concepts. 3) CCVD achieves superior performance in video style transfer (Fig. 5), as our task-aware concept aggregation module can better maintain style integrity. 4) In Fig. 6, we integrate Video-P2P (Liu et al. 2024) as a plugin for editing. The large improvement of CCVD compared to other methods is attributed to the concept-specific attribute retention module, which can preserve unique concept.

Quantitative Comparisons

To conduct quantitative comparisons between our model and other methods, we follow the evaluation protocol proposed in (Chefer et al. 2024). Specifically, we set 10 prompts per concept and generate 10 videos per prompt, resulting in a total of 100 synthesized videos for performance assessment. As shown in Tab. 1, our proposed model surpasses competing methods across all evaluation metrics, which verifies the

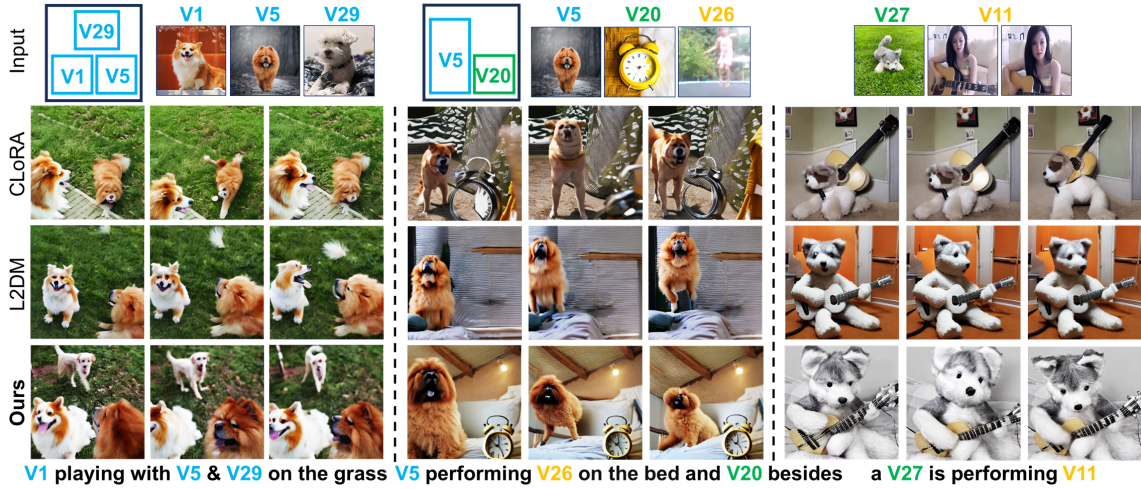


Figure 4: Comparison results of multi-concept video customization under the CTVC setting.



Figure 5: Comparisons of style transfer under CTVC setting.

Methods	TCA	LCT	COL	CLIPT	CLIPi	DINOI	F-CLIPT
Baseline	✗	✗	✗	25.5	57.8	29.0	5.2
Baseline w/ TCA	✓	✗	✗	28.7	62.2	31.3	3.4
Ours w/o COL	✓	✓	✗	29.8	63.6	33.0	1.3
Ours (CCVD)	✓	✓	✓	29.9	64.3	33.4	1.1

Table 2: Ablation of single-concept video customization.

superiority of our model to resolve the CTVC problem. This indicates that our model preserves the unique identity of each learned concept via the proposed concept-specific attribute retention module and task-aware concept aggregation.

Ablation Study

To verify the efficacy of all proposed modules, we conduct comprehensive ablation experiments in terms of layer-specific concept tokens (LCT), concept orthogonal loss (COL), task-aware concept aggregation (TCA), layer-specific region attention (LRA) and attention-guided noise estimation (ANE). As shown in Tab. 2, compared to “Baseline”,

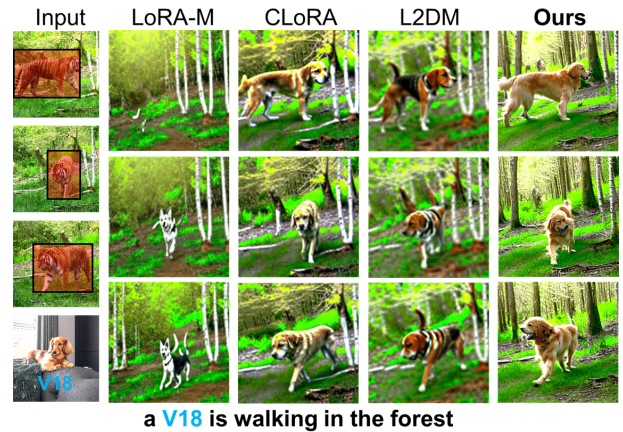


Figure 6: Comparisons of video editing under CTVC setting.

our model achieves significant performance improvements across all evaluation metrics after incorporating the proposed LCT, COL, and TCA modules for single-concept video customization. It illustrates that our model effectively mitigates forgetting by capturing distinctive attributes of old concepts while dynamically aggregating subject and motion adapters.

Conclusion

In this work, we propose a new problem termed Continual Text-to-Video Customization (CTVC), and develop a novel Continual Customized Video Diffusion (CCVD) model to resolve the CTVC problem. Specifically, we introduce a concept-specific attribute retention module during training and a task-aware concept aggregation during testing to alleviate catastrophic forgetting. Additionally, we design a controllable conditional synthesis module to address concept neglect by incorporating layer-specific region attention and attention-guided noise estimation. Comprehensive comparisons show that our CCVD has superiority performance.

Acknowledgments

This work was supported by the National Key Research and Development Program of China under Grant 2024YFB4707700, and the National Natural Science Foundation of China under Grant U23A20343, and CAS Project for Young Scientists in Basic Research, Grant No. YSBR-041, and Liaoning Provincial “Selecting the Best Candidates by Opening Competition Mechanism” Science and Technology Program under Grant 2023JH1/10400045.

References

- Belouadah, E.; and Popescu, A. 2019. IL2M: Class Incremental Learning With Dual Memory. In *ICCV*, 583–592.
- Blattmann, A.; Rombach, R.; Ling, H.; Dockhorn, T.; Kim, S. W.; Fidler, S.; and Kreis, K. 2023. Align Your Latents: High-Resolution Video Synthesis with Latent Diffusion Models. In *CVPR*, 22563–22575.
- Ceylan, D.; Huang, C.-H. P.; and Mitra, N. J. 2023. Pix2Video: Video Editing using Image Diffusion. In *ICCV*, 23149–23160.
- Chefer, H.; Zada, S.; Paiss, R.; Ephrat, A.; Tov, O.; Rubinstein, M.; Wolf, L.; Dekel, T.; Michaeli, T.; and Mosseri, I. 2024. Still-Moving: Customized Video Generation without Customized Video Data. *ACM Trans. Graph.*, 43(6).
- Chen, D.; Liao, J.; Yuan, L.; Yu, N.; and Hua, G. 2017. Coherent Online Video Style Transfer. In *ICCV*, 1114–1123.
- Chen, H.; Wang, X.; Zhang, Y.; Zhou, Y.; et al. 2024. DisenStudio: Customized Multi-subject Text-to-Video Generation with Disentangled Spatial Control. In *ACM MM*.
- Dong, J.; Li, H.; Cong, Y.; Sun, G.; Zhang, Y.; and Van Gool, L. 2024a. No One Left Behind: Real-World Federated Class-Incremental Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(4): 2054–2070.
- Dong, J.; Liang, W.; Li, H.; Zhang, D.; Cao, M.; Ding, H.; et al. 2024b. How to Continually Adapt Text-to-Image Diffusion Models for Flexible Customization? In *NeurIPS*.
- Dong, J.; Wang, L.; Fang, Z.; Sun, G.; Xu, S.; Wang, X.; and Zhu, Q. 2022. Federated Class-Incremental Learning. In *CVPR*.
- Douillard, A.; Ramé, A.; Couairon, G.; and Cord, M. 2022. DyTox: Transformers for Continual Learning With DYnamic Token eXpansion. In *CVPR*, 9285–9295.
- Fei, Z.; Li, D.; Qiu, D.; Wang, J.; Dou, Y.; Wang, R.; Xu, J.; Fan, M.; Chen, G.; Li, Y.; et al. 2025. Skyreels-a2: Compose anything in video diffusion transformers. *arXiv preprint arXiv:2504.02436*.
- Feng, R.; Weng, W.; Wang, Y.; Yuan, Y.; Bao, J.; Luo, C.; et al. 2024. CCEdit: Creative and Controllable Video Editing via Diffusion Models. In *CVPR*, 6712–6722.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-or, D. 2023. An Image is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion. In *ICLR*.
- Gao, W.; Li, Y.; Yin, Y.; and Yang, M.-H. 2020. Fast Video Multi-Style Transfer. In *WACV*, 3211–3219.
- Gu, Y.; Zhou, Y.; Wu, B.; Yu, L.; Liu, J.-W.; Zhao, R.; Wu, J. Z.; Zhang, D. J.; et al. 2024. VideoSwap: Customized Video Subject Swapping with Interactive Semantic Point Correspondence. In *CVPR*, 7621–7630.
- Guo, Y.; Yang, C.; Rao, A.; Agrawala, M.; Lin, D.; and Dai, B. 2025. SparseCtrl: Adding Sparse Controls to Text-to-Video Diffusion Models. In *ECCV*, 330–348.
- He, H.; Xu, Y.; Guo, Y.; Wetzstein, G.; Dai, B.; Li, H.; and Yang, C. 2025. CameraCtrl: Enabling Camera Control for Video Diffusion Models. In *ICLR*.
- Ho, J.; and Salimans, T. 2021. Classifier-Free Diffusion Guidance. In *NeurIPS Workshop*.
- Huang, C.-P.; Wu, Y.-S.; Chung, H.-K.; Chang, K.-P.; Yang, F.-E.; and Wang, Y.-C. F. 2025a. Videomage: Multi-subject and motion customization of text-to-video diffusion models. In *CVPR*, 17603–17612.
- Huang, Y.; Yuan, Z.; Liu, Q.; Wang, Q.; Wang, X.; Zhang, R.; Wan, P.; Zhang, D.; and Gai, K. 2025b. ConceptMaster: Multi-Concept Video Customization on Diffusion Transformer Models Without Test-Time Tuning. *arxiv preprint arxiv:2501.04698*.
- Jeong, H.; Park, G. Y.; and Ye, J. C. 2024. VMC: Video Motion Customization Using Temporal Attention Adaption for Text-to-Video Diffusion Models. In *CVPR*, 9212–9221.
- Jiang, Y.; Wu, T.; Yang, S.; Si, C.; Lin, D.; Qiao, Y.; Loy, C. C.; and Liu, Z. 2024. Videobooth: Diffusion-based video generation with image prompts. In *CVPR*, 6689–6700.
- Khachatryan, L.; Movsisyan, A.; Tadevosyan, V.; Henschel, R.; Wang, Z.; Navasardyan, S.; and Shi, H. 2023. Text2Video-Zero: Text-to-Image Diffusion Models are Zero-Shot Video Generators. In *ICCV*, 15908–15918.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13): 3521–3526.
- Lee, T.; Kwon, S.; and Kim, T. 2024. Grid Diffusion Models for Text-to-Video Generation. In *CVPR*, 8734–8743.
- Li, Z.; and Hoiem, D. 2016. Learning Without Forgetting. In *ECCV*, 614–629.
- Li, Z.; and Hoiem, D. 2017. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12): 2935–2947.
- Liang, F.; Ma, H.; He, Z.; Hou, T.; Hou, J.; Li, K.; Dai, X.; Juefei-Xu, F.; Azadi, S.; Sinha, A.; et al. 2025. Movie Weaver: Tuning-Free Multi-Concept Video Personalization with Anchored Prompts. *arXiv preprint arXiv:2502.07802*.
- Liu, J.; Wang, W.; Liu, W.; He, Q.; and Liu, J. 2023. EDT2V: An Efficient Training Framework for Diffusion-based Text-to-Video Generation. In *IJCNN*, 1–8.
- Liu, S.; Zhang, Y.; Li, W.; Lin, Z.; and Jia, J. 2024. VideoP2P: Video Editing with Cross-Attention Control. In *CVPR*, 8599–8608.
- Lopez-Paz, D.; and Ranzato, M. 2017. Gradient episodic memory for continual learning. In *NeurIPS*, 6470–6479.

- Madaan, D.; Yoon, J.; Li, Y.; Liu, Y.; and Hwang, S. J. 2022. Representational Continuity for Unsupervised Continual Learning. In *ICLR*.
- Masum, M. A.; Pratama, M.; Ramasamy, S.; Liu, L.; et al. 2024. PIP: Prototypes-Injected Prompt for Federated Class Incremental Learning. In *CIKM*, 1670–1679.
- Ostapenko, O.; Puscas, M.; Klein, T.; Jähnichen, P.; and Nabi, M. 2019. Learning to Remember: A Synaptic Plasticity Driven Framework for Continual Learning. In *CVPR*, 11313–11321.
- Qi, Z.; Tang, Y.-P.; Meng, L.; Yu, H.; Li, X.; and Meng, X. 2025. Class-wise Balancing Data Replay for Federated Class-Incremental Learning. In *NeurIPS*.
- Ren, Y.; Zhou, Y.; Yang, J.; Shi, J.; Liu, D.; Liu, F.; Kwon, M.; and Shrivastava, A. 2025. Customize-A-Video: One-Shot Motion Customization of Text-to-Video Diffusion Models. In *ECCV*, 332–349.
- She, D.; Liu, M.; Pang, J.; Wang, J.; Yang, Z.; He, W.; Zhang, G.; et al. 2025. CustomVideoX: 3D Reference Attention Driven Dynamic Adaptation for Zero-Shot Customized Video Diffusion Transformers. *arXiv preprint arXiv:2502.06527*.
- Smith, J. S.; Hsu, Y.-C.; Zhang, L.; Hua, T.; Kira, Z.; Shen, Y.; and Jin, H. 2024. Continual diffusion: Continual customization of text-to-image diffusion with c-lora. *Transactions on Machine Learning Research*.
- Smith, J. S.; Karlinsky, L.; Gutta, V.; Cascante-Bonilla, P.; Kim, D.; et al. 2023. CODA-Prompt: COntinual Decomposed Attention-Based Prompting for Rehearsal-Free Continual Learning. In *CVPR*, 11909–11919.
- Sun, G.; Liang, W.; Dong, J.; Li, J.; Ding, Z.; and Cong, Y. 2024. Create Your World: Lifelong Text-to-Image Diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(9): 6454–6470.
- Takagi, Y.; and Nishimoto, S. 2023. High-resolution image reconstruction with latent diffusion models from human brain activity. In *CVPR*, 14453–14463.
- Toldo, M.; and Ozay, M. 2022. Bring Evanescent Representations to Life in Lifelong Class Incremental Learning. In *CVPR*, 16711–16720.
- Van Le, T.; Phung, H.; Nguyen, T. H.; Dao, Q.; Tran, N. N.; and Tran, A. 2023. Anti-DreamBooth: Protecting users from personalized text-to-image synthesis. In *ICCV*, 2116–2127.
- Voynov, A.; Chu, Q.; Cohen-Or, D.; and Aberman, K. 2023. P+: Extended Textual Conditioning in Text-to-Image Generation. *arxiv preprint arxiv:2303.09522*.
- Wang, C.; Yu, M.; Wu, L.; Meng, L.; Li, X.; and Meng, X. 2024a. InstantAS: Minimum Coverage Sampling for Arbitrary-Size Image Generation. In *ACM MM*, 3180–3188.
- Wang, J.; Yuan, H.; Chen, D.; Zhang, Y.; Wang, X.; and Zhang, S. 2023a. ModelScope Text-to-Video Technical Report. *arXiv preprint arXiv:2308.06571*.
- Wang, W.; and Yang, Y. 2024. VidProM: A Million-scale Real Prompt-Gallery Dataset for Text-to-Video Diffusion Models. In *NeurIPS*.
- Wang, X.; Yuan, H.; Zhang, S.; Chen, D.; Wang, J.; et al. 2023b. VideoComposer: Compositional Video Synthesis with Motion Controllability. In *CIKM*.
- Wang, X.; Zhang, S.; Yuan, H.; Qing, Z.; Gong, B.; Zhang, Y.; Shen, Y.; Gao, C.; and Sang, N. 2024b. A Recipe for Scaling up Text-to-Video Generation with Text-free Videos. In *CVPR*, 6572–6582.
- Wei, K.; Xu, Z.; and Deng, C. 2025. Compress to One Point: Neural Collapse for Pre-Trained Model-Based Class-Incremental Learning. In *AAAI*, 20, 21465–21473.
- Wei, K.; Yang, X.; Xu, Z.; and Deng, C. 2024a. Class-incremental unsupervised domain adaptation via pseudo-label distillation. *IEEE Transactions on Image Processing*, 33: 1188–1198.
- Wei, Y.; Zhang, S.; Qing, Z.; Yuan, H.; Liu, Z.; Liu, Y.; Zhang, Y.; Zhou, J.; and Shan, H. 2024b. Dream Video: Composing Your Dream Videos with Customized Subject and Motion. In *CVPR*, 6537–6549.
- Wu, J.; Li, X.; Zeng, Y.; Zhang, J.; Zhou, Q.; Li, Y.; Tong, Y.; and Chen, K. 2025a. Motionbooth: Motion-aware customized text-to-video generation. *NeurIPS*, 37: 34322–34348.
- Wu, J. Z.; Ge, Y.; Wang, X.; Lei, S. W.; Gu, Y.; Shi, Y.; Hsu, W.; Shan, Y.; Qie, X.; and Shou, M. Z. 2023. Tune-A-Video: One-Shot Tuning of Image Diffusion Models for Text-to-Video Generation. In *ICCV*, 7589–7599.
- Wu, T.; Zhang, Y.; Wang, X.; Zhou, X.; Zheng, G.; et al. 2025b. CustomCrafter: Customized Video Generation with Preserving Motion and Concept Composition Abilities. *AAAI*, 39: 8469–8477.
- Wu, X.; Huang, S.; and Wei, F. 2024. Mixture of LoRA Experts. In *ICLR*.
- Xing, Z.; Dai, Q.; Hu, H.; Wu, Z.; and Jiang, Y.-G. 2024. SimDA: Simple Diffusion Adapter for Efficient Video Generation. In *CVPR*, 7827–7839.
- Yang, S.; Hou, L.; Huang, H.; Ma, C.; Wan, P.; Zhang, D.; Chen, X.; and Liao, J. 2024. Direct-a-Video: Customized Video Generation with User-Directed Camera Movement and Object Motion. In *SIGGRAPH*.
- Yang, Z.; Teng, J.; Zheng, W.; Ding, M.; Huang, S.; Xu, J.; Yang, Y.; Hong, W.; et al. 2025. CogVideoX: Text-to-Video Diffusion Models with An Expert Transformer. In *ICLR*.
- Yoon, J.; Yang, E.; Lee, J.; and Hwang, S. J. 2018. Lifelong Learning with Dynamically Expandable Networks. In *ICLR*.
- Zhang, Y.; Huang, N.; Tang, F.; Huang, H.; Ma, C.; Dong, W.; and Xu, C. 2023. Inversion-Based Style Transfer With Diffusion Models. In *CVPR*, 10146–10156.
- Zhao, R.; Gu, Y.; Wu, J. Z.; Zhang, D. J.; Liu, J.-W.; et al. 2025. MotionDirector: Motion Customization of Text-to-Video Diffusion Models. In *ECCV*, 273–290.
- Zhong, M.; Shen, Y.; Wang, S.; Lu, Y.; Jiao, Y.; Ouyang, S.; Yu, D.; Han, J.; and Chen, W. 2024. Multi-LoRA Composition for Image Generation. *arxiv preprint arxiv:2402.16843*.
- Zhou, Y.; Zhang, R.; Gu, J.; Zhao, N.; Shi, J.; and Sun, T. 2024. SUGAR: Subject-Driven Video Customization in a Zero-Shot Manner. *arxiv preprint arxiv:2412.10533*.