

Learning Temporally and Semantically Consistent Unpaired Video-to-video Translation Through Pseudo-Supervision From Synthetic Optical Flow

Kaihong Wang^{1*}, Kumar Akash², Teruhisa Misu²

¹ Boston University

² Honda Research Institute USA, Inc.

kaiwkh@bu.edu, kakash@honda-ri.com, tmisu@honda-ri.com

Abstract

Unpaired video-to-video translation aims to translate videos between a source and a target domain without the need of paired training data, making it more feasible for real applications. Unfortunately, the translated videos generally suffer from temporal and semantic inconsistency. To address this, many existing works adopt spatiotemporal consistency constraints incorporating temporal information based on motion estimation. However, the inaccuracies in the estimation of motion deteriorate the quality of the guidance towards spatiotemporal consistency, which leads to unstable translation. In this work, we propose a novel paradigm that regularizes the spatiotemporal consistency by synthesizing motions in input videos with the generated optical flow instead of estimating them. Therefore, the synthetic motion can be applied in the regularization paradigm to keep motions consistent across domains without the risk of errors in motion estimation. Thereafter, we utilize our unsupervised recycle and unsupervised spatial loss, guided by the pseudo-supervision provided by the synthetic optical flow, to accurately enforce spatiotemporal consistency in both domains. Experiments show that our method is versatile in various scenarios and achieves state-of-the-art performance in generating temporally and semantically consistent videos. Code is available at: https://github.com/wangkaihong/Unsup_Recycle_GAN/.

Introduction

Generative Adversarial Network (Goodfellow et al. 2014) has shown its versatility in various computer vision applications, especially in cross-domain image-to-image translation (Isola et al. 2017) tasks. However, such tasks normally require a huge amount of pixel-level annotation in datasets and generating such dataset with aligned image pairs is not only expensive but sometimes also unnecessary when expected output is not uniquely paired with input. Consequently, unpaired image-to-image translation (Zhu et al. 2017a) addresses this task in the absence of paired images. On this basis, unpaired video-to-video translation (Bashkurova, Usman, and Saenko 2018; Bansal et al. 2018) is further introduced to translate a source domain video to its target domain counterpart.

*Work done as an intern at Honda Research Institute USA, Inc. Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

A major challenge in extending synthesis from image-level to video-level is temporal inconsistency; it originates from the subtle difference in temporally adjacent frames, which is further magnified in the unstable translation process. As a result, translating every frame independently leads to flickering effects even if every single frame is correctly translated to the target domain from an image translation perspective. On the other hand, semantic inconsistency is another common problem highly correlated with temporal inconsistency, which refers to the changes in the semantic labels during the unstable translation process.

To address these problems, Recycle-GAN (Bansal et al. 2018) introduces spatiotemporal consistency constraint that utilizes temporal information in videos to regularize the temporal consistency in the translation, and is widely adopted by subsequent works. Practically, the temporal information is obtained from the estimated motions between temporally adjacent frames to synchronize motions before and after the translation, thereby enforcing cyclic spatiotemporal consistency. Meanwhile, semantic consistency is also improved when temporally contextual factors are incorporated as additional spatial smoothing. As a result, generators are guided to perform temporally and semantically consistent video translation at the same time.

Intuitively, the effect of regularization of the spatiotemporal consistency greatly relies on the accuracy of motion estimation. However, most existing works fail to accurately estimate motion and therefore weaken the benefits of the regularization. For instance, the pixel-level temporal predictor P in Recycle-GAN (Bansal et al. 2018), as illustrated in Figure 1(a), and the optical flow estimators in (Chen et al. 2019; Park et al. 2019), as illustrated in Figure 1(b), are very likely to make inaccurate estimations of the motions due to capacity limitation of the models or domain gap problems. This impairs the consistency of motions across domains, and the enforcement of the spatiotemporal consistency based on the inaccurately estimated motions will inadvertently be included as the target of optimization of the generators for the translation model. Therefore, generators will be misled when they are forced to fix the mistakes made by motion estimation in the paradigm and synthesize consistent images given inconsistent motions. As a result, the temporal and semantic stability of translated videos is undermined.

In this paper, we propose a novel paradigm that learns

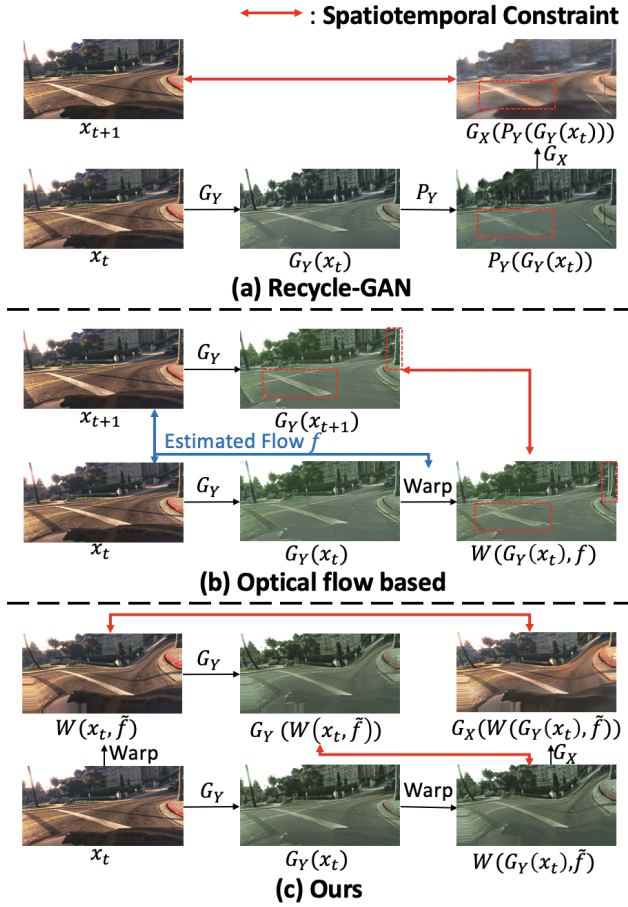


Figure 1: Comparison of spatiotemporal consistency constraints. Given an input image x_t and the future frame x_{t+1} , spatiotemporal consistency can be enforced in various paradigms as in (a) Recycle-GAN, (b) optical flow based methods and (c) ours through the cross-domain generator G_X and G_Y . Both (a) and (b) suffer from inaccuracies in motion estimation, e.g., the blurriness in $P_Y(G_Y(x_t))$ originated from the future frame predictor P in (a) and mismatched regions highlighted in red rectangles originated from the warping operation W with the estimated optical flow f in (b), which lead to inaccurate guidance towards spatiotemporal consistency. In comparison, our method avoids such errors in motion estimation by using synthetic optical flow to ensure accurate regularization of spatiotemporal consistency.

temporally and semantically consistent unpaired video-to-video translation through spatiotemporal consistency regularization via pseudo-supervision from synthetic optical flow. Specifically, given a static frame from a source video, we synthesize a random optical flow, instead of estimating one, to simulate its temporal motion and synthesize the future frame through warping. Similarly, for the target domain counterpart translated from the source frame, we simulate its future frame with the same synthetic optical flow. After that, spatiotemporal consistency regularization is enforced

between the pair of future frames with the consistent motion across domains without risk of incorrect motion estimation; the synthetic optical flow acts as a flawless motion representation for both source and target domains. Note that the quality of our pseudo-supervision does not rely on the reality of the synthetic optical flow that can simulate true motions in the domain. Instead, it focuses on ensuring the accuracy of the regularization by keeping motions across domains consistent while enforcing spatiotemporal consistency. Besides the improvements in accuracy, our method eliminates the need for motion estimation and thereby achieves better computational efficiency as compared with the existing methods.

To construct a more reliable and efficient spatiotemporal consistency regularization paradigm, we introduce two unsupervised losses: unsupervised recycle loss and unsupervised spatial loss. Given a pair of synthetic future frames in the source and target domain generated from one input static frame, our unsupervised recycle loss enforces spatiotemporal cyclic consistency in the source domain. Meanwhile, the unsupervised spatial loss regularizes spatiotemporal consistency in the target domain. To summarize, our main contributions in this work include:

1. We propose a novel paradigm that learns from the reliable guidance toward spatiotemporal consistency provided by the pseudo-supervision from synthetic optical flow and thereby mitigating the inaccuracies and computational costs originating from the motion estimation process.
2. We introduce the unsupervised recycle loss and the unsupervised spatial loss to conduct more accurate and efficient spatiotemporal consistency regularization.
3. We validate our method through extensive experiments and demonstrate that our method achieves state-of-the-art performance in generating temporally and semantically consistent videos.

Related Works

Image-to-image & Video-to-video Translation

Inspired by recent progress in generative adversarial networks, several early image-to-image translation methods (Denton et al. 2015; Radford, Metz, and Chintala 2016; Zhu et al. 2016; Zhao, Mathieu, and LeCun 2017; Isola et al. 2017; Shetty, Fritz, and Schiele 2018; Choi et al. 2018; Zhu et al. 2017b) have been proposed to tackle problems in various applications such as image generating and editing. However, all these methods require costly input-output image pairs. Many methods (Liu and Tuzel 2016; Taigman, Polyak, and Wolf 2017; Shrivastava et al. 2017) have also been proposed to address the unpaired version of the problem. Independently, cycle-consistency constraint has been introduced (Zhu et al. 2017a; Liu, Breuel, and Kautz 2017; Kim et al. 2017) that significantly improves the quality of translation by forcing cyclic consistency of an image before and after translations and inspires many following works (Yi et al. 2017; Hoffman et al. 2018; Huang et al. 2018; Yang, Xie, and Wang 2018; Liu et al. 2019; Baek et al. 2020).

On the basis of image-to-image translation, the video-to-video translation problem is firstly explored

in vid2vid (Wang et al. 2018). The quality of synthesized videos is further optimized in (Wei et al. 2018; Mallya et al. 2020) by incorporating long-term information. Few-shot vid2vid (Wang et al. 2019) alleviates the dependency of the model on a huge dataset and enhances the generalization capability of the model. Similar to its image counterpart, these methods also require paired annotated videos, which is expensive to acquire. Therefore, to address unpaired video-to-video translation, a spatiotemporal 3D translator is proposed in (Bashkirova, Usman, and Saenko 2018) on the basis of existing image-to-image translation networks. To regularize spatiotemporal consistency, Recycle-GAN (Bansal et al. 2018) integrates temporal factors into its Cycle-GAN-based model by predicting future frames. Mocycle-GAN (Chen et al. 2019) and STC-V2V (Park et al. 2019) adopt optical flow to improve the accuracy of their spatiotemporal consistency constraints. To further improve the coherence of the generated video, a spatiotemporal discriminator unit, as well as a self-supervised ping-pong loss, is proposed in TecoGAN (Chu et al. 2020) for long-term temporal consistency regularization. UVIT (Liu et al. 2021) applies an RNN-based translator as well as a self-supervised video interpolation loss to incorporate temporal information and generates videos in a recurrent manner. Another branch of works (Bonneel et al. 2015; Lai et al. 2018; Lei, Xing, and Chen 2020) seek to improve the temporal consistency of videos in a post-processing manner that is independent from the translating process, which can limit the quality of the improvement.

Although spatiotemporal consistency constraints in these methods improve the quality of the generated videos, these constraints rely on some form of motion estimation that is not perfect. As a result, errors in motion estimation undermine the spatiotemporal consistency constraints and thereby deteriorates the quality of the generated videos. To investigate the effect of the accuracy in the estimation of motions on the quality of spatiotemporal consistency constraint, we conducted a pilot experiment on STC-V2V (Park et al. 2019) with different optical flow estimating methods and calculate the warping error. The definition of the warping error will be discussed in later sections. Besides the original method that adopts FlowNet2 (Ilg et al. 2017) to estimate flow between two temporally adjacent frames, we also ran an experiment using the same model but with intentionally reversed order of the input frames to simulate a less accurate estimation, and another experiment with a more advanced flow estimator RAFT (Teed and Deng 2020). As presented in the results in Table 1, the improvement of the optical flow estimating method contributes to the decrease in the warping error, which indicates the enhancement of the temporal stability of the resulting videos. Therefore, we conclude that regularizing the spatiotemporal consistency with the absence of accurate motion estimation limits the performance of the generator to conduct stable video translation.

Compared with existing unpaired video-to-video translation methods, our work focuses on providing more accurate and efficient spatiotemporal consistency regularization through the pseudo-supervision from synthetic optical flow instead of estimated motions between frames.

Optical flow estimator	EPE (\downarrow)	Warping Error (\downarrow)
Reversed FlowNet2	-	0.050356
FlowNet2	2.02	0.039703
RAFT	1.43	0.037891
Oracle	-	0.029555

Table 1: Pilot experiments with different optical flow estimation methods. End-point error (EPE) represents the error of the flow estimator on the train set of Sintel, indicating the error rate of the optical flow estimator, while warping error evaluates the temporal stability of the generated videos. Oracle represents the warping error in the source video.

Video Style Transfer

Different from unpaired video-to-video translation, video style transfer methods are mainly built on the basis of existing non-adversarial image style transfer and focus on the improvement of temporal consistency of the resulting videos given content videos and style images. Efficient feed-forward networks for fast video style transfer with different variants of temporal constraints (Ruder, Dosovitskiy, and Brox 2016; Huang et al. 2017; Gupta et al. 2017; Chen et al. 2017; Gao et al. 2018; Chen et al. 2020) are adopted to optimize temporal consistency. However, all these methods either require ground truth optical flow from the dataset or suffer from the incorrect estimation of optical flow, which limits their performance. To address these problems, Compound-VST (Wang et al. 2020) proposes a single-frame based video style transfer method that adopts compound regularization to learn temporal consistency through synthetic optical flow while LinearVST (Li et al. 2019) and MCCNet (Deng et al. 2021) seek to enhance the learning of arbitrary style transfer by modifying the architecture of models.

Noticing the drawbacks of existing spatiotemporal constraints and inspired by the recent progress in CompoundVST (Wang et al. 2020), we integrate the pseudo-supervision from synthetic optical flow into our Cycle-GAN-based paradigm. Note that, compared to CompoundVST, our work focuses on the cross-domain unpaired video-to-video translation problem and not just video style transfer. Also, we propose a novel regularization paradigm that more effectively enforces spatiotemporal consistency in both domains and therefore facilitates not only temporally stable but also semantically consistent translation. Recent methods (Eilertsen, Mantiuk, and Unger 2019; Zhang et al. 2021) have also focused on enhancing video quality in different specific tasks through the regularization of consistency across frames based on motion. Nevertheless, we note that our method addresses the more general problem of unpaired video-to-video translation.

Method

In this work, we propose a paradigm that tackles the temporal and semantic inconsistency problem in unpaired video-to-video translation tasks. Given a source video $x_{1:T} = \{x_t\}_{t=1}^T$ with T frames from the source domain X , and a target video $y_{1:S} = \{y_s\}_{s=1}^S$ with S frames from the target




Figure 2: Illustration of the comparison between existing consistency constraint methods and ours. A pair of input images x_t and y_s are translated through the cross-domain generators G_X and G_Y to the other domain. P_X and P_Y are the future frame predictor in the source and target domain, while $W(x, \tilde{f})$ represents the warping operation on the image x with the synthetic optical flow \tilde{f} . (a) Recycle-GAN: Temporal information is incorporated into the recycle loss by predicting future frames in the target domain through the frame predictor P and the consistency is regularized between the real future frame and the reconstructed estimated future frame in the original domain. (b) Ours: Our paradigm replaces the future frame predictor with a warping operation given a synthetic optical flow \tilde{f} to simulate future frames. With the ensured accuracy in the motion simulation, we regularize the consistency of the pair of future frames in both domains with our unsupervised recycle and spatial loss.

domain Y , we aim to learn a mapping $G_Y : X \rightarrow Y$ as well as the opposite mapping $G_X : Y \rightarrow X$ implemented on two generators without any paired videos or frames as input. With the learnt generators, target videos are synthesized as $\hat{x}_{1:T} = \{\hat{x}_t\}_{t=1}^T$ where $\hat{x}_t = G_Y(x_t)$, and $\hat{y}_{1:S} = \{\hat{y}_s\}_{s=1}^S$ where $\hat{y}_s = G_X(y_s)$.

Compared with most existing methods that translate videos in a recurrent manner, our paradigm achieves stable video-to-video synthesis through a pair of translators taking only a single frame as input at a time. Our method consists of two objectives: an adversarial loss that guides the translation of an image to the other domain and a pair of unsupervised losses that keep videos temporally and semantically consistent across translation. The illustration of our paradigm compared with Recycle-GAN is shown in Fig. 2.

Adversarial Loss

To translate videos to the other domain, we train the generators and discriminators adversarially in a zero-sum game as in (Goodfellow et al. 2014). Practically, we optimize our discriminators D_X/D_Y so that they can correctly distinguish synthetic images \hat{x}_t/\hat{y}_s from real images x_t/y_s . At the same time, generators are trained to synthesize realistic images that can fool the discriminator. Therefore, the adversarial loss \mathcal{L}_{adv} is given by

$$\min_G \max_D \mathcal{L}_{adv} = \sum_s \log D_Y(y_s) + \sum_t \log(1 - D_Y(\hat{y}_s)) + \sum_t \log D_X(x_t) + \sum_s \log(1 - D_X(\hat{x}_t)). \quad (1)$$

Unsupervised Losses

Given a source frame x_t as the t -th frame from source video $x_{1:T}$, we simulate the synthetic future frame \tilde{x}_{t+1} by warping the current frame with a synthetic optical flow \tilde{f} as $\tilde{x}_{t+1} = W(x_t, \tilde{f}) = F(x_t, \tilde{f}) + \Delta$ as in Compound-VST (Wang et al. 2020). Here F represents the warping operation with the synthetic optical flow and Δ represents pixel-level random Gaussian noise. Refer to Appendix A for more details regarding the generation of the synthetic optical flow.

Unsupervised Recycle Loss: Given an original frame x_t/y_s and its translated frame \hat{x}_t/\hat{y}_s , we simulate their synthetic following frame $\tilde{x}_{t+1}/\tilde{y}_{s+1}$ and $W(\hat{x}_t, \tilde{f})/W(\hat{y}_s, \tilde{f})$ through a warping operation with an identical synthetic optical flow \tilde{f} . After that, we map $W(\hat{x}_t, \tilde{f})/W(\hat{y}_s, \tilde{f})$ back to the original domain as in Recycle-GAN and enforce its consistency with $\tilde{x}_{t+1}/\tilde{y}_{s+1}$ using unsupervised recycle loss \mathcal{L}_{ur} given by

$$\mathcal{L}_{ur} = \sum_t \|\tilde{x}_{t+1} - G_X(W(\hat{x}_t, \tilde{f}))\|_1 + \sum_s \|\tilde{y}_{s+1} - G_Y(W(\hat{y}_s, \tilde{f}))\|_1. \quad (2)$$

Compared with Recycle-GAN (Bansal et al. 2018) that applies a pixel-level frame predictor to estimate the future frame, our paradigm uses the synthetic optical flow that ensures the consistency of the motion across domains. Consequently, the error from motion estimation becomes irrelevant to the translation process resulting in guaranteed accuracy of spatiotemporal consistency guidance.

Unsupervised Spatial Loss: In parallel with \mathcal{L}_{ur} , our method also applies an unsupervised spatial loss \mathcal{L}_{us} that enforces the consistency between $\tilde{x}_{t+1}/\tilde{y}_{s+1}$ and $W(\hat{x}_t, \tilde{f})/W(\hat{y}_s, \tilde{f})$ in the opposite direction (see Equation (3)). Practically, $\tilde{x}_{t+1}/\tilde{y}_{s+1}$ is mapped to the other domain to be consistent with $W(\hat{x}_t, \tilde{f})/W(\hat{y}_s, \tilde{f})$ so that the warping and the translation is commutative to each other.

$$\mathcal{L}_{us} = \sum_t \|G_Y(\tilde{x}_{t+1}) - W(\hat{x}_t, \tilde{f})\|_1 + \sum_s \|G_X(\tilde{y}_{s+1}) - W(\hat{y}_s, \tilde{f})\|_1 \quad (3)$$

Guided by the pseudo-supervision from synthetic optical flow, our unsupervised losses include errors only from inconsistent translation processes in G_X and G_Y , and minimizing the loss value leads to direct improvement in the quality of translation without misleading the generators.




Figure 3: Synthesis of optical flow and future frame.

Full Objective

Combining all loss functions we listed above, we present the full objective of our paradigm to optimize a single-frame based unpaired video-to-video translation system as

$$\mathcal{L} = \mathcal{L}_{adv} + \lambda_{ur}\mathcal{L}_{ur} + \lambda_{us}\mathcal{L}_{us}. \quad (4)$$

The weight of our unsupervised losses λ_{ur} and λ_{us} are set to 10 for all experiments. Note that additional losses can be added to the objective as needed for the translation task.

Experiments

We evaluate and compare our method with existing methods on various aspects including temporal and semantic consistency as well as human preference. Experiments were conducted in several video-to-video translation tasks that include the intra-modality Viper-to-Cityscapes translation between driving scenes in Viper (Richter, Hayder, and Koltun 2017) and Cityscapes (Cordts et al. 2016) along with the inter-modality video-to-labels translation between the driving scenes and the label maps in Viper. We also ran experiments on a series of face-to-face translation tasks and present the results in Appendix C to showcase the versatility of our method. Refer to Appendix A for further details regarding the datasets, architectures, training protocols, and other specific settings applied in the experiments. Following our baselines (Zhu et al. 2017a; Bansal et al. 2018), we set the weights for both of our losses as 10 for generality.

Optical flow synthesis: We follow the protocol in CompoundVST (Wang et al. 2020) to synthesize optical flow, which consists of motions and shifts. For each of the x and y direction, the motion optical flow \tilde{f}_m is randomly generated as a grid, and each element of the grid corresponds to a 100×100 block in the input image. The grid is sampled from $\mathcal{N}(0, \sigma_m^2 I)$ with $\sigma_m = 8$, and upsampled through bilinear interpolation to the original size of the input image. The shift optical flow \tilde{f}_s is generated as a constant shift sampled from $\mathcal{N}(0, \sigma_s^2)$ for each of x and y direction with $\sigma_s = 10$. The generated synthetic optical flow $\tilde{f} = \tilde{f}_m + \tilde{f}_s$ is finally smoothed using a box filter of size 100×100 . After that, a future frame can be simulated by warping an input image with the synthetic flow and adding Gaussian noise $\Delta \sim \mathcal{N}(0, \sigma_n^2 I)$ and $\sigma_n \sim \mathcal{U}(0.01, 0.02)$. The overall process of the optical flow and future frame synthesis is illustrated in Figure 3.

Method	Warping Error (\downarrow)
Cycle-GAN	0.059346
Recycle-GAN	0.052738
STC-V2V	0.039703
Ours	0.035984
Oracle(Source)	0.029555

Table 2: Comparison of warping error between our method with Cycle-GAN, Recycle-GAN, and STC-V2V as well as oracle results from the input source video for Viper-to-Cityscapes translation.

Viper-to-Cityscapes Experiments

In this experiment, we validate the temporal consistency and visual quality of the translated videos from our method. Cycle-GAN (Zhu et al. 2017a), Recycle-GAN (Bansal et al. 2018), and STC-V2V (Park et al. 2019) are included as existing state-of-the-art methods for comparison. Additionally, we adopt the content loss from STC-V2V during the training of our method and set λ_{cont} as 1 to constraint the similarity of content in images during translation. Qualitative results are presented in Figure 4.

Temporal consistency First, we quantitatively evaluate the temporal consistency in the Viper-to-Cityscapes translation using warping error as the metric as presented in STC-V2V. Practically, the error \mathcal{L}_{warp} is calculated as the occlusion-mask-weighted Manhattan distance between adjacent frames in videos, given ground truth optical flow $f_{t \rightarrow t-1}$ from Viper and is given by

$$\mathcal{L}_{warp} = \sum_{t=2}^K O_{t \rightarrow t-1} \|\hat{x}_t - F(\hat{x}_{t-1}, f_{t \rightarrow t-1})\|_1. \quad (5)$$

Here $O_{t \rightarrow t-1} = \exp(-\alpha \|x_t - F(x_{t-1}, f_{t \rightarrow t-1})\|_2)$ is the occlusion mask estimated based on the difference in adjacent source frames x_t and x_{t-1} , and α is set to 50. Intuitively, lower warping error indicates a lower difference in the adjacent generated frames, implying stronger video temporal stability. We compare our method with three baselines as well as the oracle warping error in source videos, and the results are presented in Table 2. We observe that our method reduces the warping error compared with existing methods and approximates closer to the oracle level. This shows that our paradigm achieves state-of-the-art performance in producing temporally consistent videos.

User studies To more comprehensively evaluate the visual quality of the generated videos, we conducted user studies deployed on Amazon Mechanical Turk (AMT) and tested our method against the other baselines. Twenty one participants were invited to review pairs of videos generated by our method versus each of the baselines in both the Viper-to-Cityscapes and the Cityscapes-to-Viper translation for the videos in the validation set of Viper and Cityscapes. Each participant was given 15 pairs of videos consisting of comparisons between ours against the three baselines from 5 randomly sampled videos from both domains. In every comparison, they were asked to review a pair of videos for at least




Figure 4: Comparison of existing methods with ours in the Viper-to-Cityscapes translation experiments. Based on the highlighted regions, we observe significantly fewer semantic errors. For example, in (a), the regions of the sky are more detailed and reasonable for Viper-to-Cityscapes translation. Similarly, in (b), the regions of buildings in the first and second rows as well as the regions of sky in the third row are better translated for Cityscapes-to-Viper translation.

7 seconds and select their preference in terms of consistency and quality. We calculate the user preference in percentage according to the frequency of their selections for all three comparisons against baselines, and report the results in Table 3. Results indicate that our method produces more visually pleasing and temporally stable videos in both directions compared with existing baselines.

Video-to-Labels Experiments

To further verify that our method performs semantically consistent video-to-video translation, we conducted the video-to-labels and labels-to-video experiments following protocols in previous works (Bansal et al. 2018; Chen et al. 2019; Liu et al. 2021). In this scenario, we translate the videos to their corresponding pixel-level semantic segmentation labels in Viper on the basis of the architecture of Cycle-GAN (Zhu et al. 2017a) and MUNIT (Huang et al. 2018). Besides our unsupervised losses, we also applied the cycle loss following (Zhu et al. 2017a; Bansal et al. 2018) and set λ_{cyc} as 10 during our training process to further enhance the regularization for the sake of fair comparison. The quality of the translation, in terms of semantic consistency, was evaluated quantitatively in both directions. We incorporated Cycle-GAN (Zhu et al. 2017a), Recycle-GAN (Bansal et al. 2018), and STC-V2V (Park et al. 2019) as baseline methods in the following experiments and compare their performance with ours. Additionally, we suppressed the regularization for the video domain in our paradigm to avoid generating vague contents from the regions in the label domain that have the same semantic labels. Further discussions for this setting are available in ablation studies and Appendix A. Qualitative re-

Method	Forward	Backward
Ours / Cycle-GAN	90.48 / 9.52	92.38 / 7.62
Ours / Recycle-GAN	79.05 / 20.95	75.24 / 24.76
Ours / STC-V2V	75.24 / 24.76	80.95 / 19.05

Table 3: Comparison of user preference between our method with others. Forward indicates user preference in percentage for the translation from Viper to Cityscapes, while backward indicates that for the translation from Cityscapes to Viper.

Criterion	method	Day	Sunset	Rain	Snow	Night	All
MP	Cycle-GAN	29.68	32.89	34.87	28.65	17.92	28.29
	MUNIT	39.92	66.54	70.46	64.20	50.74	56.40
	Recycle-GAN	46.89	69.77	60.10	56.98	31.88	52.10
	STC-V2V	40.98	67.51	62.99	61.86	36.07	52.49
	Ours	56.31	76.90	71.75	67.86	48.79	63.20
	Ours _{munit}	58.69	78.58	74.08	71.06	44.82	64.36
AC	Cycle-GAN	7.54	6.15	6.08	6.45	4.05	6.56
	MUNIT	13.38	14.66	16.76	15.40	11.70	14.36
	Recycle-GAN	13.62	13.80	11.47	13.62	6.91	13.40
	STC-V2V	12.23	13.54	13.34	13.66	6.69	12.32
	Ours	16.58	17.22	15.79	16.14	8.27	16.13
	Ours _{munit}	18.37	19.83	19.07	18.50	10.12	18.34
mIoU	Cycle-GAN	4.11	3.66	3.34	3.56	1.94	3.76
	MUNIT	7.58	10.17	12.05	11.04	7.40	9.73
	Recycle-GAN	9.02	10.90	8.07	9.60	3.58	9.53
	STC-V2V	8.06	10.61	9.22	10.19	3.97	9.09
	Ours	11.90	13.49	11.75	12.60	5.14	12.29
	Ours _{munit}	13.29	15.64	13.97	14.17	6.49	13.71

Table 4: Comparison of baseline methods with ours in segmentation score for video-to-label translation experiments on Viper.

method	MP	AC	mIoU
Recycle-GAN	60.1	14.9	11.0
Mocycle-GAN	64.9	17.7	13.2
Improvement (%)	7.99	18.79	20.00
Recycle-GAN	52.10	13.40	9.53
Ours	63.20	16.13	12.29
Improvement (%)	21.31	20.37	28.96
Recycle-GAN	53.65	14.84	10.11
UVIT	65.20	17.59	13.07
Improvement (%)	21.53	18.53	29.28
Recycle-GAN	52.10	13.40	9.53
Ours _{munit}	64.36	18.34	13.71
Improvement (%)	23.53	36.87	43.86

Table 5: Comparison of state-of-the-art methods with ours in video-to-label translation experiments.

sults are shown in Figure 5.




Figure 5: Comparison of existing methods with ours in the video-to-labels translation experiments. We can notice more accurate and detailed segmentation results for our method around the vegetation in the first, second, and fourth row; poles in the third row; and the trees in the fifth row. Similarly, in (b), we can also observe more realistic, natural, and detailed scenes, such as roads and sky in the third and fourth rows as well as vegetation at the first, second, and fifth rows from our results.

Video-to-Labels We conducted experiments of the translation from videos to labels, which is equivalent to video semantic segmentation, and evaluated the performance by comparing the generated labels with ground truth in the held-out validation set. We report the segmentation scores, including mean pixel accuracy (MP), average class accuracy (AC), and intersection over union (mIoU), for all baseline methods and ours in Table 4. Additionally, we also integrated our unsupervised spatiotemporal consistency constraints into a multimodal translation model MUNIT and report the baseline MUNIT result as well as our result as $Ours_{munit}$ in the table. Further details regarding the integration are available in Appendix A. Based on the results, we can observe the improvement of Recycle-GAN and STC-V2V over Cycle-GAN, revealing the effectiveness of the spatiotemporal consistency regularization. Meanwhile, our method significantly outperforms all the baseline methods in three metrics under all weather conditions, which shows the benefit of learning accurate spatiotemporal consistency through pseudo-supervision from synthetic optical flow. We further compare our method with the state-of-the-arts including Mocycle-GAN (Chen et al. 2019), a Cycle-GAN-based model learning from motion-guided spatiotemporal constraints, and UVIT (Liu et al. 2021), an RNN-based method performing a multimodal translation. To be more specific, we take Recycle-GAN as the common baseline and compare their improvements in percentage over it with ours in Table 5. Since UVIT adopts a multimodal paradigm, we compare our method combined with MUNIT against it for the sake of fair comparison. As presented in Table 5, our paradigm brings in greater improvements in all three segmentation metrics over the Recycle-GAN without relying on any extra architectures except for the basic discriminators and generators in Cycle-GAN, indicating the necessity and effectiveness of ensuring the accuracy in spatiotemporal consistency constraint using synthetic optical flow. Besides, it shows the generality and flexibility of our unsupervised losses that enable effective combination with other translation models to further enhance the translation quality.

Labels-to-Video Besides video-to-labels translation, we also compare the performance of existing methods with ours in the opposite direction that generates videos given labels. Following previous works (Isola et al. 2017; Bansal et al.

Criterion	method	Day	Sunset	Rain	Snow	Night	All
MP	Cycle-GAN	26.70	47.25	40.26	42.61	26.02	35.77
	MUNIT	30.74	70.48	53.41	59.14	45.74	50.56
	Recycle-GAN	23.88	51.34	35.32	45.45	28.32	36.35
	STC-V2V	29.73	52.79	42.42	47.91	37.25	41.34
	Ours	34.90	66.05	46.46	55.49	37.22	47.42
	$Ours_{munit}$	35.95	66.41	47.62	57.05	48.49	50.52
AC	Cycle-GAN	11.03	9.83	10.34	11.09	9.49	11.20
	MUNIT	11.36	16.31	14.53	14.17	14.14	14.36
	Recycle-GAN	10.94	12.11	11.64	11.38	10.99	11.56
	STC-V2V	12.65	13.44	12.57	12.33	13.01	13.10
	Ours	14.86	16.55	15.01	15.41	14.14	15.73
	$Ours_{munit}$	14.86	18.98	16.03	16.42	17.07	17.29
mIoU	Cycle-GAN	5.76	6.70	5.70	6.53	4.83	6.36
	MUNIT	5.95	12.36	7.92	9.25	7.73	8.66
	Recycle-GAN	5.23	8.33	6.22	7.01	5.61	6.67
	STC-V2V	7.15	9.16	7.30	7.59	7.55	8.06
	Ours	8.48	12.25	8.96	10.36	8.43	10.08
	$Ours_{munit}$	8.55	13.75	9.09	11.07	10.37	11.07

Table 6: Comparison of baseline methods with ours in FCN-score for labels-to-video translation experiments on Viper.

2018; Chen et al. 2019), we adopted FCN-score to evaluate the quality of the generated frames. The intuition behind this metric is that an FCN model, pretrained on the Viper dataset, should generally perform better on the videos translated from labels if they are close enough to the original videos. Therefore, higher FCN-scores indicate higher quality in the generated videos. Table 6 compares the FCN-scores for all baseline methods with ours as well as its combination with MUNIT. We can observe from the results that with the refinement of the spatiotemporal consistency constraint, the quality of the generated images gradually improves, and our method achieves the best performance in all three metrics under all weather conditions. Similar to the video-to-labels experiments, we also measure our improvement of FCN-score in percentage of over the common baseline, Recycle-GAN, with Mocycle-GAN and report the comparison in Table 7. The comparisons of our improvement over Recycle-GAN against current state-of-the-art methods indicate the advantage of our unsupervised paradigm. Without relying on any external architecture, our unsupervised spatiotemporal consistency learns from pseudo-supervision

method	MP	AC	mIoU
Recycle-GAN	43.5	12.2	7.0
Mocycle-GAN	47.6	16.0	10.1
Improvement (%)	9.43	31.15	44.29
Recycle-GAN	36.35	11.56	6.67
Ours	47.42	15.73	10.08
Improvement (%)	30.45	36.07	51.12

Table 7: Comparison of Mocycle-GAN with ours in labels-to-video translation experiments.

UR	US	VDRS	Cyc	V2L			L2V		
				mIoU	AC	MP	mIoU	AC	MP
✓				3.08	5.54	23.88	3.69	7.18	27.66
	✓			0.16	3.23	4.86	4.75	8.55	27.27
✓	✓			9.89	13.39	57.12	5.45	10.20	29.50
✓	✓	✓		9.53	12.96	58.54	7.00	12.19	37.93
✓	✓	✓	✓	12.29	16.13	63.20	10.08	15.73	47.42

Table 8: Ablation studies on video-to-labels translation. UR represents unsupervised recycle loss, US represents unsupervised spatial loss, VDRS represents video domain regularization suppression and Cyc represents cycle loss.

through synthetic optical flow and therefore achieves new state-of-the-art performance. Additionally, further comparison with UVIT in this setting is available in Appendix C.

Ablation Studies

To verify the effectiveness of our unsupervised spatiotemporal consistency constraints and other relevant components, we ran ablation studies on the video-to-labels translation task. More ablation studies conducted on Viper-to-Cityscapes as well as to explore the effect of the synthetic optical flow are available in Appendix B.

For video-to-labels translation, we ran experiments with each of our loss functions and their meaningful combinations. Additionally, we also turned off the video domain regularization suppression to investigate its effects in video generation. The performance of these experiments was evaluated in both segmentation scores for the video-to-labels translation and FCN-scores for the labels-to-video translation. Based on the results shown in Table 8, we can observe that each of our unsupervised losses cannot achieve semantically consistent translation independently, while their combination performs significantly better in the video-to-labels translation. On that basis, enabling video domain regularization suppression further improves the quality of the videos translated from labels, as it prevents local details suppression. Finally, the application of the cycle loss contributes another leap in the quality of the generated videos, showing the effectiveness of the cyclic spatial constraints in preserving semantic consistency. Qualitative results are presented in Appendix C for a more intuitive understanding.

To verify the effectiveness of using the proposed method to generate synthetic optical flow, we compare three optical flow generation methods with our paradigm: 1) optical

λ_{UR}	λ_{US}	Warping Error
1	10	0.040817
5	10	0.037563
10	1	0.059230
10	5	0.035434
10	10	0.035984

Table 9: Ablation studies regarding hyper-parameters.

Optical Flow Method	Warping Error
Estimated	0.048983
Translation only	0.050842
Scaling only	0.037798
Ours	0.035984

Table 10: Ablation studies on different optical flow generation methods.

flow estimated through FlowNet2, 2) optical flow synthesized with only translation motions, and 3) optical flow synthesized with only scaling motions. We compare the warping error on the Viper-to-Cityscapes translation experiment. As presented in Table 10, estimated optical flow deteriorates the translation possibly due to artifacts of discontinuity as compared to our smooth synthetic flow. The artifacts combined with our strong consistency constraints undermines adversarial training and thereby the quality of translation. On the other hand, optical flow with only low-level motions, such as translation or scaling, cannot provide rich information for temporal and semantic consistency for the training process. In contrast, our synthetic flow ensures the quality of the optical flow with no flow discontinuity, providing stronger support to our effective regularization paradigm.

Finally, we tested hyper-parameter sensitivity in the Viper-to-Cityscapes translation with $(\lambda_{UR}, \lambda_{US})$ as $(1, 10)$, $(5, 10)$, $(10, 1)$, and $(10, 5)$. The resulting warping error are presented in Table 9. The results show that when either of the weights is not large enough, the effectiveness of the consistency constraint will be compromised. On the other hand, when the ratio of the weights gets closer to 1, our method yields stable results, revealing the robustness of our method against hyper-parameters.

Conclusion

In this work, we propose a novel paradigm for unpaired video-to-video translation using pseudo-supervision via synthetic optical flow resulting in improved temporal and semantic consistency. Compared with typical approaches, our method simulates motions by synthesizing optical flow instead of estimating them to ensure the consistency of motions across domains and therefore guarantees the accuracy of the spatiotemporal consistency regularization in the cross-domain translation process. With the guidance of the unsupervised recycle loss as well as the unsupervised spatial loss, our paradigm effectively facilitates stable video-to-video translation. We finally demonstrate the state-of-the-art performance of our method through extensive experiments.

References

- Baek, K.; Choi, Y.; Uh, Y.; Yoo, J.; and Shim, H. 2020. Rethinking the Truly Unsupervised Image-to-Image Translation. *arXiv preprint arXiv:2006.06500*.
- Bansal, A.; Ma, S.; Ramanan, D.; and Sheikh, Y. 2018. Recycle-GAN: Unsupervised Video Retargeting. In *ECCV*.
- Bashkirova, D.; Usman, B.; and Saenko, K. 2018. Unsupervised Video-to-Video Translation. *arXiv preprint arXiv:1806.03698*.
- Bonneel, N.; Tompkin, J.; Sunkavalli, K.; Sun, D.; Paris, S.; and Pfister, H. 2015. Blind video temporal consistency. *ACM Trans. Graph.*
- Chen, D.; Liao, J.; Yuan, L.; Yu, N.; and Hua, G. 2017. Coherent Online Video Style Transfer. In *ICCV*.
- Chen, X.; Zhang, Y.; Wang, Y.; Shu, H.; Xu, C.; and Xu, C. 2020. Optical Flow Distillation: Towards Efficient and Stable Video Style Transfer. In *ECCV*.
- Chen, Y.; Pan, Y.; Yao, T.; Tian, X.; and Mei, T. 2019. Mocycle-GAN: Unpaired Video-to-Video Translation. In *ACM Multimedia*.
- Choi, Y.; Choi, M.; Kim, M.; Ha, J.; Kim, S.; and Choo, J. 2018. StarGAN: Unified Generative Adversarial Networks for Multi-Domain Image-to-Image Translation. In *CVPR*.
- Chu, M.; Xie, Y.; Mayer, J.; Leal-Taixé, L.; and Thuerey, N. 2020. Learning temporal coherence via self-supervision for GAN-based video generation. *ACM Transactions on Graphics*.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *CVPR*.
- Deng, Y.; Tang, F.; Dong, W.; Huang, H.; Ma, C.; and Xu, C. 2021. Arbitrary Video Style Transfer via Multi-Channel Correlation. In *AAAI*.
- Denton, E. L.; Chintala, S.; Szlam, A.; and Fergus, R. 2015. Deep Generative Image Models using a Laplacian Pyramid of Adversarial Networks. In *NIPS*.
- Eilertsen, G.; Mantiuk, R. K.; and Unger, J. 2019. Single-Frame Regularization for Temporally Stable CNNs. In *CVPR*.
- Gao, C.; Gu, D.; Zhang, F.; and Yu, Y. 2018. ReCoNet: Real-Time Coherent Video Style Transfer Network. In *ACCV*.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A. C.; and Bengio, Y. 2014. Generative Adversarial Nets. In *NIPS*.
- Gupta, A.; Johnson, J.; Alahi, A.; and Fei-Fei, L. 2017. Characterizing and Improving Stability in Neural Style Transfer. In *ICCV*.
- Hoffman, J.; Tzeng, E.; Park, T.; Zhu, J.; Isola, P.; Saenko, K.; Efros, A. A.; and Darrell, T. 2018. CyCADA: Cycle-Consistent Adversarial Domain Adaptation. In Dy, J. G.; and Krause, A., eds., *ICML*.
- Huang, H.; Wang, H.; Luo, W.; Ma, L.; Jiang, W.; Zhu, X.; Li, Z.; and Liu, W. 2017. Real-Time Neural Style Transfer for Videos. In *CVPR*.
- Huang, X.; Liu, M.; Belongie, S. J.; and Kautz, J. 2018. Multimodal Unsupervised Image-to-Image Translation. In *ECCV*.
- Ilg, E.; Mayer, N.; Saikia, T.; Keuper, M.; Dosovitskiy, A.; and Brox, T. 2017. FlowNet 2.0: Evolution of Optical Flow Estimation with Deep Networks. In *CVPR*.
- Isola, P.; Zhu, J.; Zhou, T.; and Efros, A. A. 2017. Image-to-Image Translation with Conditional Adversarial Networks. In *CVPR*.
- Kim, T.; Cha, M.; Kim, H.; Lee, J. K.; and Kim, J. 2017. Learning to Discover Cross-Domain Relations with Generative Adversarial Networks. In *ICML*.
- Lai, W.; Huang, J.; Wang, O.; Shechtman, E.; Yumer, E.; and Yang, M. 2018. Learning Blind Video Temporal Consistency. In *ECCV*.
- Lei, C.; Xing, Y.; and Chen, Q. 2020. Blind Video Temporal Consistency via Deep Video Prior. In *NeurIPS*.
- Li, X.; Liu, S.; Kautz, J.; and Yang, M. 2019. Learning Linear Transformations for Fast Image and Video Style Transfer. In *CVPR*.
- Liu, K.; Gu, S.; Romero, A.; and Timofte, R. 2021. Unsupervised Multimodal Video-to-Video Translation via Self-Supervised Learning. In *WACV*.
- Liu, M.; Breuel, T.; and Kautz, J. 2017. Unsupervised Image-to-Image Translation Networks. In *NIPS*.
- Liu, M.; Huang, X.; Mallya, A.; Karras, T.; Aila, T.; Lehtinen, J.; and Kautz, J. 2019. Few-Shot Unsupervised Image-to-Image Translation. In *ICCV*.
- Liu, M.; and Tuzel, O. 2016. Coupled Generative Adversarial Networks. In *NIPS*.
- Mallya, A.; Wang, T.; Sapiro, K.; and Liu, M. 2020. World-Consistent Video-to-Video Synthesis. In *ECCV*.
- Park, K.; Woo, S.; Kim, D.; Cho, D.; and Kweon, I. S. 2019. Preserving Semantic and Temporal Consistency for Unpaired Video-to-Video Translation. In *ACM Multimedia*.
- Radford, A.; Metz, L.; and Chintala, S. 2016. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. In *ICLR*.
- Richter, S. R.; Hayder, Z.; and Koltun, V. 2017. Playing for Benchmarks. In *ICCV*.
- Ruder, M.; Dosovitskiy, A.; and Brox, T. 2016. Artistic Style Transfer for Videos. In Rosenhahn, B.; and Andres, B., eds., *GCPR*.
- Shetty, R.; Fritz, M.; and Schiele, B. 2018. Adversarial Scene Editing: Automatic Object Removal from Weak Supervision. In *NIPS*.
- Shrivastava, A.; Pfister, T.; Tuzel, O.; Susskind, J.; Wang, W.; and Webb, R. 2017. Learning from Simulated and Unsupervised Images through Adversarial Training. In *CVPR*.
- Taigman, Y.; Polyak, A.; and Wolf, L. 2017. Unsupervised Cross-Domain Image Generation. In *ICLR*.
- Teed, Z.; and Deng, J. 2020. RAFT: Recurrent All-Pairs Field Transforms for Optical Flow. In *ECCV*.

- Wang, T.; Liu, M.; Tao, A.; Liu, G.; Catanzaro, B.; and Kautz, J. 2019. Few-shot Video-to-Video Synthesis. In *NIPS*.
- Wang, T.; Liu, M.; Zhu, J.; Yakovenko, N.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. Video-to-Video Synthesis. In *NIPS*.
- Wang, W.; Xu, J.; Zhang, L.; Wang, Y.; and Liu, J. 2020. Consistent Video Style Transfer via Compound Regularization. In *AAAI*.
- Wei, X.; Zhu, J.; Feng, S.; and Su, H. 2018. Video-to-Video Translation with Global Temporal Consistency. In *ACM Multimedia*.
- Yang, X.; Xie, D.; and Wang, X. 2018. Crossing-Domain Generative Adversarial Networks for Unsupervised Multi-Domain Image-to-Image Translation. In *ACM Multimedia*.
- Yi, Z.; Zhang, H. R.; Tan, P.; and Gong, M. 2017. DualGAN: Unsupervised Dual Learning for Image-to-Image Translation. In *ICCV*.
- Zhang, F.; Li, Y.; You, S.; and Fu, Y. 2021. Learning Temporal Consistency for Low Light Video Enhancement From Single Images. In *CVPR*.
- Zhao, J. J.; Mathieu, M.; and LeCun, Y. 2017. Energy-based Generative Adversarial Networks. In *ICLR*.
- Zhu, J.; Krähenbühl, P.; Shechtman, E.; and Efros, A. A. 2016. Generative Visual Manipulation on the Natural Image Manifold. In *ECCV*.
- Zhu, J.; Park, T.; Isola, P.; and Efros, A. A. 2017a. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *ICCV*.
- Zhu, J.; Zhang, R.; Pathak, D.; Darrell, T.; Efros, A. A.; Wang, O.; and Shechtman, E. 2017b. Toward Multimodal Image-to-Image Translation. In *NIPS*.