
RETARGETING VIDEO WITH AN END-TO-END FRAMEWORK

Thi-Ngoc-Hanh Le

National Cheng-Kung University
Taiwan
ngochanh.le1987@gmail.com

HuiGuang Huang

National Cheng-Kung University
Taiwan
604300468@qq.com

Yi-Ru Chen

National Cheng-Kung University
Taiwan
y2133256@gmail.com

Tong-Yee Lee*

National Cheng-Kung University
Taiwan
tonylee@mail.ncku.edu.tw

ABSTRACT

Video holds significance in computer graphics applications. Because of the heterogeneous of digital devices, retargeting videos becomes an essential function to enhance user viewing experience in such applications. In the research of video retargeting, preserving the relevant visual content in videos, avoiding flicking, and processing time are the vital challenges. Extending image retargeting techniques to the video domain is challenging due to the high running time. Prior work of video retargeting mainly utilizes time-consuming preprocessing to analyze frames. Plus, being tolerant of different video content, avoiding important objects from shrinking, and the ability to play with arbitrary ratios are the limitations that need to be resolved in these systems requiring investigation. In this paper, we present an end-to-end RETVI method to retarget videos to arbitrary aspect ratios. We eliminate the computational bottleneck in the conventional approaches by designing RETVI with two modules, content feature analyzer (CFA) and adaptive deforming estimator (ADE). The extensive experiments and evaluations show that our system outperforms previous work in quality and running time.

Keywords video retargeting, RETVI, analyze video, deforming, grid movement, pixel movement

1 Introduction

Video is a widely used media form that holds significant importance in computer graphics applications due to its ability to convey motion, simulate reality, and engage viewers. Because of the heterogeneous of digital devices, adapting videos to different display size, resolution, or aspect ratios (referred to as “*Video Retargeting*”) has become an essential function in these applications. For example, video retargeting can be employed to (1) adapt video content in real-time based on user actions or analytical queries [4]; (2) ensure that the video content is visually pleasing, legible, and seamlessly integrated with the augmented reality environment [5]; or (3) dynamically adjusts the content of virtual reality (VR) videos to match the specific characteristics




Figure 1: Comparison on running time (second) of the existing SOTA image retargeting methods and our RETVI.

of different VR headsets [6]. Additionally, resizing videos has become increasingly popular with the advent of smartphones equipped with video capabilities and the rise of social media platforms. Whether you choose a standard aspect ratio for content to be streamed on a laptop or a cell phone, it is crucial to ensure that the video presents clearly without unsightly cropping, providing the best experience for your audience. By leveraging video retargeting in these applications, it could ultimately enhance users overall viewing experience. To drive a video to a target size, conventional video retargeting methods require a time-consuming and expensive process to analyze input videos prior to resizing. In contrast, we have developed a deep learning-based framework allows for retargeting videos to arbitrary aspect ratios in an end-to-end manner. Our approach eliminates the computational bottleneck present in conventional methods while delivering higher quality results.

Researchers explore the problems of video retargeting with various approaches, including conventional and state-of-the-art (SOTA) techniques. The retargeting image could be the background research of video domain. Several attempts have resolved the image retargeting’s problem. Each method usually consists of two main steps: (1) determining the importance of image pixels by extracting an importance map and (2) performing an image retargeting operator to obtain the retargeting image [7]. Deep learning-based techniques can produce good visual results. However, it requires equipment with high computation power and comprehensive datasets. We can see a comparison on the running time of the most recent SOTA image retargeting systems in Fig. 1. Extending a deep learning-based image retargeting technique to the video is tough to operate.

Unlike retargeting images, working on the video domain is more challenging. Along with preserving important content after resizing sessions, producing temporally coherent retargeting videos is also vital to judging the performance of such a video retargeting system. Researchers have explored this research domain in various ways [50, 35, 55, 33, 27]. However, they share an identical workflow as those in the image domain. They need to utilize some off-the-shelf content analysis, i.e., saliency map and segmentation, to pre-process video frames in advance. This approach leads to two significant downfalls. First, the quality of resized videos depends on the performance of these auxiliary functions. This may not only downgrade the capability of the proposed method but also prevent the method from playing with diverse video content. Second, analyzing video frames is more expensive than single images. It may take minutes to process on a single frame. Processing time is an important factor when retargeting is used in services like video on demand or live streams, especially when the ability to retarget videos with high resolutions in real-time is required [21]. Therefore, using such a time-consuming system to retarget video in the high demand of the digital age could be impractical.

To this end, the challenges in video resizing are not only retaining important content in such temporal coherent results but also the demand for a scheme that can adapt to diverse video content, tolerate various aspect ratios, and be fast in running time. Motivated by this observation, we propose a novel end-to-end framework, abbreviated as **RETVI**, to address the above challenges. We aim to retarget videos to arbitrary aspect ratios without using hand-craft pre-processing or image/video retargeting annotations. Our designed framework

attempts to capture the critical information of video and preserve them after retargeting session without distortion and shrinking artifacts. Our essential contribution is reducing the running time cost to process video retargeting in an end-to-end manner. To achieve these, we design our framework consisting of a content feature analyzer (CFA) and an adaptive deforming estimator (ADE). The CFA is responsible for learning the content information of the input frames. Analyzed features serve as guidance for deforming the frame to a target size. In this manner, we propose a neural network, which can be considered an alternative to the time-consuming pre-processing in prior work. With the knowledge learned from CFA and a given resizing ratio, our ADE module tries to teach the network how to deform the input frame with minimum distortion. We test the performance of our proposed method through various videos. The experimental results and evaluations demonstrate that our system outperforms previous works in quality and running time.

In summary, our main contributions are as follows:

- We investigate a novel end-to-end video retargeting based on high-level representation of video frames.
- We design a network that can effectively demonstrate the important content of the video, which performs better the expensive computation of the preprocessing in prior video resizing systems.
- Our method can adapt to various video contents and arbitrary aspect ratios.
- End-to-end and fast running time enable our system to be potentially embedded in other services or applications.

2 Related work

2.1 Image Retargeting

Image Retargeting has received much attention during the last decade. Researchers approach this domain with different techniques, from conventional to state-of-the-art. Most of the classical image resizing algorithms [8, 36, 18, 31, 19, 12] share an identical workflow with two steps. In the first step, an importance map is generated via visual attention analysis methods, like saliency detection. The main goal is to preserve the regions with high importance as well as possible. In the second step, an operator retargets the image [21]. Cropping, scaling, warping, browsing, seam carving, or combining them are the typical operators selected and used in these resizing algorithms. Most recently, Kim *et al.* [22] investigated a novel approach, a grid encoding model for image retargeting, which takes each horizontal/vertical distance between two adjacent vertices as an optimization variable. The results demonstrate that their method consistently outperforms previous methods on qualitative and quantitative perspectives [22].

With the revolution of deep learning technologies, several works have applied these new techniques to resolve the problems in image retargeting [57, 47, 48, 37, 10, 30]. CarvingNet [47] uses an encoder-decoder CNN to develop an importance map based on a learning model. WSSDCNN [10], a network learns the image content via an attention map, guiding pixel-wise mapping during retargeting. Cycle-IR [48] solves the problem of image retargeting through unsupervised learning. Conceptually, they build the network based on reverse mapping from the retargeted images to the given input images. Formulating the multi-operator retargeting upon reinforcement learning technique is a novel approach presented in SAMIR [57]. This approach can produce results with lower computational costs. Nevertheless, the running time is relatively high, and the results still suffer from cropping seriously.

2.2 Video Retargeting

Preserving the relevant visual content in videos while avoiding flicking is the most crucial challenge of video retargeting [21]. Using image retargeting techniques on individual video frames does not provide satisfactory

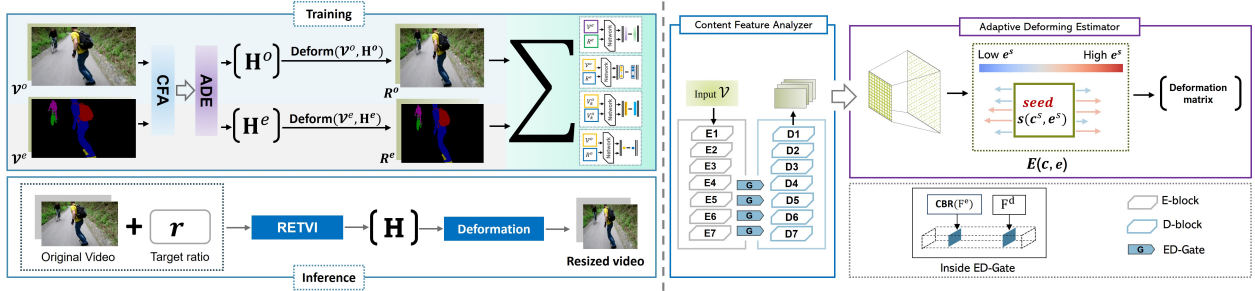


Figure 2: Left: the workflow of our RETVI in training and inference; right: visualization of the CFA and ADE modules.

results as they might change entirely different areas in adjacent frames [24]. Video retargeting becomes an exciting research topic recently when the explosive growth of social platforms and digital devices demands the videos to be resized to display them nicely.

The early attempts to retarget videos use a straightforward technique, cropping [35, 9, 53]. Later, researchers investigate more algorithms to resolve the problems in this research domain. Extending the typical techniques in image retargeting to video is commonly used by researchers, particularly seam carving and warping operators. Extending seam carving to video retargeting might lead to high processing time. Therefore, researchers in [14, 17, 50] focus on speeding up the processing time and saving memory space when extending this technique to videos. Applying warping to retargeting videos may cause undesirable artifacts. It is because of a temporal motion of an image region in one direction followed by motion in the opposite direction. Therefore, warping-based research applied to videos focuses on image stability and run-time performance [21]. Gallea et al. [15] use a 2D grid for image retargeting and add a third dimension to handle changes over time. The authors in [28, 51] extend the axis-aligned image retargeting to video by computing the deformation for selected keyframes and interpolating the other frames. A contrasting approach to these extensions is introduced by Lin et al. [33]. The authors use a uniform grid mesh for the warping in this work. They focus on preserving the important objects in a video while warping the non-important regions in a way similar to linear scaling [21]. The most recent work in this research domain is introduced by Lee et al. [27], which uses a deep neural network for video retargeting. Their concept uses object detection to allocate the bounding boxes of the main objects in a video. Then, the remaining background areas are resized to preserve the content of the bounding boxes without deformation. Along with retargeting traditional images/videos, the research domain on stereo image/video retargeting also receives attention. The studies proposed in [32, 34, 45, 13] demonstrate that analyzing the importance in input is an essential factor to gain good quality resized stereo images/videos.

The sharp contrast between our framework and theirs [50, 55, 33, 27] is that we develop an end-to-end system for retargeting videos to arbitrary aspect ratios. Our proposed method is faster than the conventional techniques, effectively preserving important information in videos after retargeting session without distortion or shrinking phenomenon, and tolerant to videos with dense content and multiple moving objects alike.

3 Methodology

3.1 System overview

We propose a novel framework RETVI for retargeting videos to the arbitrary aspect ratio. Our approach is to retarget videos without hand-craft pre-processing or image/video retargeting annotations. We convert the video retargeting problem to an unsupervised learning problem of conditional transformation regression without explicitly computing a transformation. As diagrammed in Fig. 2, a pair of video frames ($\mathcal{V}^o, \mathcal{V}^e$) is

required in the training process, in which \mathcal{V}^o is the input video frame and \mathcal{V}^e is the extracted foreground of \mathcal{V}^o . In the inference phase, only the video with frames $\{\mathcal{V}^o\}$ is required as input in our end-to-end retargeting process. In the following, in the context that we do not need to differentiate \mathcal{V}^o and \mathcal{V}^e , we use \mathcal{V} to denote such an input frame in our system. Specifically, we model our RETVI with two modules, a content feature analyzer (CFA) and an adaptive deforming estimator (ADE). The specific procedure is outlined in Fig. 2(b).

As named, the CFA is responsible for analyzing the input content. It accepts video frames as the input. As a result, CFA converts the input frames \mathcal{V} to high-level representation. Conceptually, ADE receives feature representation of \mathcal{V} and a random resizing ratio as inputs, and tries to teach the network how to deform \mathcal{V} with minimum distortion. Afterwards, we train the network with the weighted sum of four loss functions. Once trained, given a video and target size, our network can efficiently and effectively produce the retargeting result in an end-to-end manner. Instead of using an expensive pre-processing for visual attention analysis as in the aforementioned conventional retargeting systems (*e.g.*, saliency map, segmentation, optical flow), we analyze the input frames by the CFA module. The features obtained from CFA are more robust to noise and occlusions than saliency maps or segmentations. This is because they rely on higher-level features, such as shape, texture, and context, which are less affected by noise and occlusions than low-level features such as color or intensity. Besides, CFA can be more computationally efficient and generalizable to new images once trained.

3.2 Network architecture

3.2.1 Content Feature Analyzer

This module, dubbed as **CFA**, shoulders the task of learning the contextual information of frames during encoding them into latent space. This process can be considered as an alternative to the time-consuming pre-processing, *i.e.*, saliency and segmentation, in those described above conventional resizing systems. To define the important regions in frames, such a well-known object detection also could be used. Nonetheless, detection focuses on detecting the bounding boxes surrounding the detected objects. Meanwhile, we aim at estimating at the pixel level to capture more semantic information in the frame rather than using the bounding boxes. Consequently, we structure CFA such that it can produce pixel-wise output of input frame from encoded features. We design CFA with the inspiration from the U-Net model [41], which has been a successful network in medical image segmentation. It is later varied to be used in other applications, such as natural image analysis, natural language processing, and image classification. U-Net is shaped in various structures depending on the goal of a particular application. Similarly, we design a variation of this standard model to meet the purpose mentioned earlier. The tweaks of our design are as follows.

Given a video frame \mathcal{V} in size $H \times W$, we generate a pyramid of spatial features from coarse to fine granularities through a so-called E-Blocks, which is structured by $Conv3 \times 3 \rightarrow Normalization \rightarrow Tanh$. Formally, the feature maps produced by an E-Block is formulated as:

$$\mathbf{E}_i = \mathcal{T}(\mathcal{B}(C^3(\mathbf{E}_{i-1}, \kappa_i))), \quad (1)$$

here $C^3(\cdot)$ performs the 3×3 convolution, $\mathcal{B}(\cdot)$ represents the batch normalization, and $\mathcal{T}(\cdot)$ indicates the Tanh activation function. \mathbf{E}_i is a matrix in $\mathbb{R}^{H_i \times W_i \times \kappa}$ with κ_i is the number of kernels used in the i^{th} E-Block, $i = 1 \dots 7$. The contrary path is a symmetric form of those obtained from E-Blocks. It is worth noting that we deepen the number of layers to seven since we aim to detect both the more extensive features (*i.e.*, the main objects in the foreground) and the more fine-grained features (*i.e.*, the small objects in the background). In the pure UNet and its variants, skip connections play a crucial role in facilitating information flow between different levels of the network, enabling better feature propagation and learning. They mostly use skip connections at every upsample layer. In contrast, we employ the skip connections half-way, *i.e.*, at

layer 4 to layer 7, as features are with more fine granularity dwell deeper layers. This strategy could aid in discriminating objects from the background regions or multiple objects distributed entire frame. It eventually boosts the capability of the decoder. Specifically, we embed an *ED-Gate* (\mathbf{G}) at each skip connection to summarize the fine features and construct the decoded features. The feature maps at each symmetric layer are expressed as follows:

$$\mathbf{D}_i = \begin{cases} \mathbf{G}(\mathbf{E}_i, \mathbf{D}_{i+1}), & \text{if } 4 \leq i \leq 6 \\ \Omega(\mathbf{D}_{i+1}), & \text{if } 1 \leq i < 4 \end{cases}, \quad (2)$$

here $\Omega(\cdot)$ represents for a so-called D-Block, which is formulated as:

$$\Omega(\mathcal{F}^{out}) = \mathcal{R}(\mathcal{B}(C^T(\mathcal{F}^{in}))), \quad (3)$$

where \mathcal{R} is the ReLu activation function, and $C^T(\cdot)$ is the transposed convolution operator. In essence, the $C^T(\cdot)$ operation forms the same connectivity as the normal convolution but in the backward direction. Moreover, the weights in C^T are learnable, we accordingly do not need a pre-defined interpolation method.

In the ED-Gate, given two feature maps \mathcal{F}^e and \mathcal{F}^d , respectively from encoding and decoding layers, \mathcal{F}^e is first fed to a CBR block, which is structured as $Conv3 \times 3 \rightarrow Normalization \rightarrow ReLu$. It is then concatenated with \mathcal{F}^d to result in the output feature maps at a certain skip connection. Theoretically, the output of an ED-Gate is expressed as:

$$\mathcal{F}^g = \text{CBR}(\mathcal{F}^e) \boxtimes \mathcal{F}^d, \quad (4)$$

with \boxtimes is the concatenation operation.

We can obtain certain advantages with the above design. First, increasing the depth of layers can capture more complex features in the image, which can lead to improve the accuracy of understanding the image content. This is especially important for analyzing the image having complex textures. Second, the ED-Gate boosts that the information can flow more directly between the encoder and decoder. Third, this design can help improve the parameter efficiency of our model by allowing it to capture more information with fewer parameters. This can help prevent overfitting and improve the generalization performance of the model.

3.2.2 Adaptive Deforming Estimator

With the decoded feature map \mathbf{D}_1 obtained from the CFA, the question here is how to thread it to learn the appropriate deformation matrix. In other words, we need to define a function $f : \mathbb{R}^{h \times w \times k} \rightarrow \mathbb{R}^{S_h \times S_w}$, where h, w, k is the height, width, and kernel size of \mathbf{D}_1 . The matrix $\mathbf{H} \in \mathbb{R}^{S_h \times S_w}$ is shaped in the same size with the input frame, *i.e.*, $S_h \times S_w$. Finding the mapping function f is a game that several prior work has challenged with, such as, geometric warping with industrial style transfer [52], video stabilization [56]. Optionally, depending on the goal of a certain application, the function f is designed in different ways. In our current application, we aim to estimate a deformation to drive the input frame from a source size \mathbf{S} to a target one \mathbf{T} , $\mathbf{S}^{S_h \times S_w} \rightarrow \mathbf{T}^{T_h \times T_w}$, with minimum distortion and content awareness. For this goal, we design the Adaptive Deforming Estimator (**ADE**). In this stage, we consider the advantage of both pixel-based and grid-based strategies when formulating the function f . Pixel-based is simple and can be computationally efficient. However, it can lead to loss the detail and resolution, particularly when downsampling an image. Meanwhile, grid-based methods can preserve more detail and resolution in the resized frame than pixel-based methods. However, it can be more computationally expensive than pixel-based methods. To alleviate the burden of the trade-offs between computational efficiency and image quality, we devise an in-between strategy, which can use the advantage of each method.

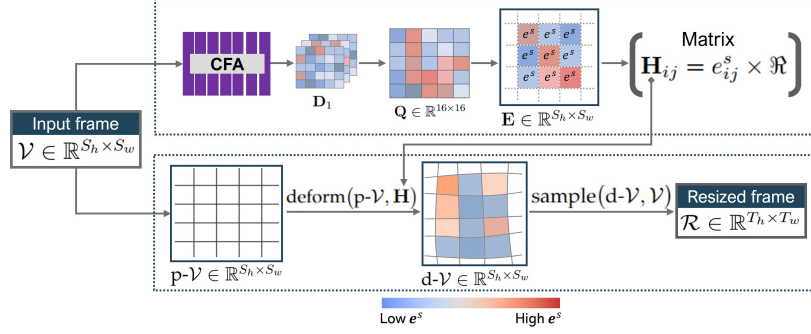


Figure 3: The workflow of our resizing strategy

Now, we detail how we utilize feature maps \mathbf{D}_1 to deform frames. The workflow of this phase is outlined in Fig. 3. First, we transform \mathbf{D}_1 to 2D-grid form via an activation function:

$$Q = \frac{1 - e^{-2a}}{1 + e^{-2a}}, \quad (5)$$

here, $a \in \mathbf{D}_1$, a tensor with 16 channels; hence, $Q \in \mathbb{R}^{2 \times 16 \times 16}$. Through equation (5), each element in the feature map \mathbf{D}_1 is returned values in range $(-1, 1)$, which indicates the fine features of the input frame. To define the mapping value in Q to pixel in \mathcal{V} , we interpolate Q to a regular $S_h \times S_w$ grid pixel coordinate: $\mathbf{E} = \nabla_b(Q)$, here ∇_b denotes the upsampling operator with bilinear mode. The resultant dense map \mathbf{E} represents for the information of the corresponding pixel in the input frame. As a result, each pixel in \mathbf{E} is also a grid cell, and we assume each of them to be a seed. A seed $s(c^s, e^s) \in \mathbf{E}$ has a coordinate $c^s(x, y)$ and energy $e^s(x, y)$. The energy of a seed demonstrates how important the pixel is. A pixel belonging to important regions should have a smaller energy compared to those belonging to the reverse one. In other words, pixels are with higher energy will be more deformed than others. We call \mathbf{E} an *energy map*, which is considered as the guidance for the later deforming estimation.

Given a target size $\mathbf{T}(T_h, T_w)$, to construct the resizing form (denoted as \mathcal{R}) of frame \mathcal{V} , instead of deforming \mathcal{V} directly, we make a tweak to manipulate this process. First, the energy of seeds in \mathbf{E} is extracted and re-formulated to define a matrix \mathbf{H} , in which value at coordinate (i, j) is a tuple with:

$$\begin{cases} \mathbf{H}_x^{ij} = e_x^s \times \mathfrak{R} \\ \mathbf{H}_y^{ij} = e_y^s \times \mathfrak{R} \end{cases} \quad (6)$$

here $i = 0, \dots, S_h - 1$; $j = 0, \dots, S_w - 1$; and

$$\mathfrak{R} = \begin{cases} (1 - r)^2, & \text{if reducing} \\ -(1 - r)^2, & \text{if enlarging} \end{cases} \quad (7)$$

where $r = T_w/S_w$, i.e., the resizing ratio. To make our model be tolerant with arbitrary target sizes, r is randomized in the range of $[0.25, \dots, 1.25]$ in our training. The matrix \mathbf{H} represents the seed's flow intensity when the current space is retargeted. Thereafter, we initialize an empty pseudo frame $p\text{-}\mathcal{V}$ in the same size with \mathcal{V} . Then, we partition \mathcal{V} and $p\text{-}\mathcal{V}$ into a regular grid of size $(S_h \times S_w)$. Consequently, we deform $p\text{-}\mathcal{V}$ using the matrix \mathbf{H} , resulting in a deformed grid $d\text{-}\mathcal{V}$. The coordinate of a cell in $d\text{-}\mathcal{V}$ is defined as: $(c_x^s + \mathbf{H}_x^{ij}, c_y^s + \mathbf{H}_y^{ij})$. By applying different deformation weights to the x-axis and y-axis separately, we can achieve non-linear changes in the shape of the resized frames. Once deforming $p\text{-}\mathcal{V}$ to $d\text{-}\mathcal{V}$, we use $d\text{-}\mathcal{V}$ as

a mapping to sample pixel values from \mathcal{V} and generate resizing frame \mathcal{R} . The frame \mathcal{R} is constructed by interpolating the pixel values according to $d\mathcal{V}$. The matrix \mathbf{H} is a learnable parameter, and \mathcal{R} is optimized at each iteration during training by four loss functions. The energy e^s is contravariant with the pixel property at a particular seed. That is, the more important the pixel is, the smaller the energy is.

Our above tweak serves the following benefits. First, the deformation is applied on the coordinate system itself, without considering the content of the frame. Hence, it results in a smooth and continuous transformation. Second, we have more control over the transformation of the coordinate space. This allows us to apply various aspect ratios on the deformation. Besides, by formulating \mathfrak{R} during the training, the model can learn to adapt to various resizing ratios under the control of loss functions. As a result, we do not need to re-train the network whenever it plays with a new target size. Also, such a time-consuming pre-processing mentioned above is eliminated and videos are retargeted in an end-to-end process.

3.3 Loss function

Let the generator network be denoted by f parameterized by weights μ , it transforms an input frame $\mathcal{V} \in \mathbb{R}^{H \times W \times C}$ into an output deformation matrix \mathbf{H} via the mapping $\mathbf{H} = f_\mu(\mathcal{V})$. Network learning adjusts the parameters μ through minimizing four loss functions \mathcal{L}_{cri} , \mathcal{L}_{glo} , \mathcal{L}_{tem} , and \mathcal{L}_{fid} . Each loss function computes a scalar value $\mathcal{L}(\cdot)$ measuring the difference between the input frame and the retargeting frame corresponding to four loss functions. The network is trained to minimize a weighted combination of the loss function:

$$\mu^* = \arg \min_{\mu} E_{\mathcal{V}} [\lambda_c \mathcal{L}_{cri} + \lambda_g \mathcal{L}_{glo} + \lambda_t \mathcal{L}_{tem} + \lambda_f \mathcal{L}_{fid}]. \quad (8)$$

The parameters λ_c , λ_g , λ_t and λ_f are all set to 1 in our training. In the following, we delve in detail the figuration of each loss function. The ablated results on each component in our optimization are presented in later Sec.4.5.

Critical region loss (\mathcal{L}_{cri}). In reconstructing the frame to a new ratio, the goal is to recover the impaired frame to match the pristine undistorted counterpart visually. Thus, we need to design the loss that would adhere to that goal. We are inspired by the perceptual loss [20] to supervise feature changes. Perceptual loss is expressed in various forms depending on specific application, e.g., image style transfer, image restoration, image colorization. Specifically, in our current application, we formulate this loss to distinguish important objects in a frame and maintain their proportional shape between the ground truth foreground and the estimated foreground frame, written as:

$$\mathcal{L}_{cri} = \sum_i \frac{1}{C_i H_i W_i} \| (\Gamma(\Phi_i(\mathcal{V}^e)) - \Gamma(\Phi_i(\mathcal{R}^e))) \|_2^2, \quad (9)$$

where Γ is an off-the-shelf feature extractor. In our experiment, we use VGG-19 as a feature extractor. However, other pre-trained models, such as ResNet, could result in equivalent effect. $\Phi_i(\cdot)$ is the feature maps in the i^{th} layer of the corresponding input parameter \mathcal{V}^e or \mathcal{R}^e . Here, \mathcal{V}^e is annotated foreground of the input frame \mathcal{V}^o and \mathcal{R}^e is resized form of \mathcal{V}^e . With this loss function, the regions with low energy in the input frame can be well preserved.

Global integrity loss (\mathcal{L}_{glo}). In addition to the critical regions, we also consider the overall presentation of the output image. This objective function is particularly effective when the content is dense and distributed entire frame. To preserve the frame information as much as possible, making resized frames more harmonious and natural, we rely on the advantage of an image classification network (shortened by ICFNet). In a certain ICFNet, for instance, VGG or ResNet, the architecture consists of feature extraction blocks and ends with a fully connected layer, which is then used for producing predictions. Inspired by this, we formulate \mathcal{L}_{glo}

to define how the content of the input frame is preserved in the retargeted one. Let denote x be the one-dimensional latent vector of an input frame \mathcal{V} after the fully connected layer of the ICFNet. The probabilities in vector x for all possible N_c classes is expressed as:

$$P(x|N_c) = \frac{\exp(x)}{\sum_{j=1}^{N_c} \exp(x_j)} \quad (10)$$

The objective of \mathcal{L}_{glo} can be simplified as:

$$\mathcal{L}_{glo} = \| P(x^o) - P(x^r) \|_2^2, \quad (11)$$

here P is defined by Eq.(10); arguments x^o, x^r are respectively the latent vectors of $\mathcal{V}^o, \mathcal{R}^o$ after feeding them to a ICFNet. In our experiment, the ICFNet we use is ResNet. Using VGG models may also result in equivalent effect. With Eq.(11), we can measure the different degree of the pair (x^o, x^r) , i.e., defining whether they are classified into the same classification. In this way, the content information can be compared without being limited by the size difference between frames \mathcal{V}^o and \mathcal{R}^o .

Temporal consistency loss (\mathcal{L}_{tem}). Temporal consistency should be specifically considered in our video retargeting task. The temporal consistency loss is used for evaluating the coherence between adjacent resized frames. To obtain this coherency, several works in the video generation domain, such as video style transfer, video stabilization, video segmentation, *etc.*, usually calculate the pixel correspondence between adjacent frames with optical flows. However, the time cost of optical flow estimation is expensive. It could affect the efficiency of the model. Besides, in our video retargeting application, the sizes of the input and retargeted frames are different, so finding the mapping optical flows in this context is impractical. Instead, we base on the knowledge of the correlations between the features of the adjacent frames to infer the temporal consistency degree in the retargeted videos. For any triplet of the adjacent frames with the respective supper script $t - 1$, t , and $t + 1$, the temporal consistency loss is formulated as:

$$\mathcal{L}_{tem} = \sum \| \gamma(\mathcal{V}_{t-1}^o, \mathcal{V}_t^o, \mathcal{V}_{t+1}^o) - \gamma(\mathcal{R}_{t-1}^o, \mathcal{R}_t^o, \mathcal{R}_{t+1}^o) \|_2^2 \quad (12)$$

where $\gamma(\cdot)$ is to calculate the correlation of features between adjacent frames. Assuming with a certain triplet of three adjacent frames $(\mathbf{F}_{t-1}, \mathbf{F}_t, \mathbf{F}_{t+1})$, $\gamma(\cdot)$ is expressed as:

$$\gamma(\cdot) = \frac{\sum (\Gamma(\Phi_i(\mathbf{F}_{t-1})) \times \Gamma(\Phi_i(\mathbf{F}_t))) \times \Gamma(\Phi_i(\mathbf{F}_{t+1}))}{H_i W_i C_i} \quad (13)$$

with $\Gamma(\cdot)$ and $\Phi(\cdot)$ are defined in the same way with Eq.(9); $i \in [1 \dots 4]$. Using the correlation of features between adjacent frames could be more appropriate to maintain temporal coherence than optical flow in video retargeting applications. Video retargeting involves changing the resolution of a video, which can cause distortion and loss of detail. To avoid these issues, it's important to maintain temporal coherence between frames, so that the motion and appearance of objects in the video remain consistent. Compared to other alternatives, such as optical flow, our formulation is simplicity, robustness, and useful for our task when changing the size.

Fidelity loss (\mathcal{L}_{fid}). Apart from minimizing geometry distortions, the quality of the resized frames is also a vital aspect we take into account in our current application. The matrix estimated by our network plays as the weight for deforming frames. However, the retargeted results might be blurry due to the relative coordinates we used to reconstruct the pixels in the ADE module. To alleviate this phenomenon, we propose the fidelity loss. Our idea in this manner comes from the concept of Generative Adversarial Networks (GANs). Unlike the discriminator in such a GAN-based network which tries to classify examples as either real (from the domain) or fake (generated), our formulation aims to discriminate the fidelity of retargeted frame to its

corresponding input in terms of pixel color value. We note here that without \mathcal{L}_{fid} , our network alone is capable to estimate the deformation matrix efficiently. Besides, our model is not supervised by the annotated retargeting data, the optimization plays as auto-alignment manner. To this end, we propose this loss to boost the performance of our model in terms of “clean quality.”

A pair of frames ($\mathcal{V}^o, \mathcal{R}^o$) is taken as input in our formulation. They are first fed to a so-called *Fid-Disc* to analyze them in feature space. Unlike other loss functions, \mathcal{L}_{cri} , \mathcal{L}_{glo} , \mathcal{L}_{tem} , which use an off-the-shelf feature extractor, we design a lightweight module to configure \mathcal{L}_{fid} . The *Fid-Disc* contains 6 CBR blocks, as used in the ED-Gate, and 1 fully-connected hidden layer (64 channels). Accordingly, the fidelity loss is expressed as:

$$\mathcal{L}_{fid} = \frac{1}{2} \left(\xi(\mathcal{D}(\mathcal{V}^o), 1) + \xi(\mathcal{D}(\mathcal{R}^o), 0) \right), \quad (14)$$

where

$$\xi(\cdot) = -\frac{1}{n} \sum ((\tau[i] * \log(\eta[i])) + (1 - \tau[i]) * \log(1 - \eta[i])). \quad (15)$$

Here, \mathcal{D} is the *Fid-Disc*; η is the score output by \mathcal{D} with the parameters \mathcal{V}^o or \mathcal{R}^o ; and τ is the score we want it to discriminate, i.e., $\tau \in \{0, 1\}$. The training goal of \mathcal{D} is to distinguish the input frame from its resized frame by the ADE module. Our network tries to generate a high score (i.e., close to 1) from \mathcal{D} , so that the output of ADE module can be closer to the quality of the input frame.

4 Experimental Results

4.1 Implementation Details

We implement our RETVI in Pytorch [39]. All experiments were performed on a PC equipped with Intel Core i7-770 CPU, 16GB RAM and a GeForce RTX 2080 Ti GPU. We train RETVI using batch size 2 within 150 epochs. The training takes approximately 10 hours. For the loss function, we use Adam optimizer [23] with a learning rate of 1×10^{-3} .

In terms of training data, we train our RETVI on DAVIS dataset [40], which consists of 90 videos in size 854×480 with single and multiple moving objects. In the dataset, each video has a corresponding foreground annotation. We use it as \mathcal{V}^e and the original video as \mathcal{V}^o . Preparing training data by collecting videos using existing tools/methods to segment foreground for them is also an alternative. However, this way is not recommended since the performance of the foreground extractor may cause occlusion effect and make the accuracy of importance region unstable.

To produce final resizing frame \mathcal{R} after sampling d- \mathcal{V} to input frame \mathcal{V} , we reconstruct the content in a window (\mathbf{W}) of $\left[\left(\frac{S_w}{2} - \frac{T_w}{2} \right) : \left(\frac{S_w}{2} + \frac{T_w}{2} \right) \right]$. Since our RETVI is trained to play with arbitrary ratios while controlling several aspects of video quality (such as global content, critical region content preservation, temporal coherence, etc.), the content after sampling may not always fit with the window \mathbf{W} . More specifically, when resizing a frame of width S_w to a target width T_w , we try to estimate matrix \mathbf{H} such that the total movement of pixels is as tight as possible to T_w . However, due to the trace-off in our afore discussion, the total movement would be lightly greater than T_w in some cases. As a result, cropping may occur. See detail visualization in the supplementary file.

4.2 Our Results and Discussion

To demonstrate the abilities of our method, we test it on various video contents. Fig.4 exhibits the results of three typical showcases in our experiments, each video has different attributes. To be specific, the video “*Dancing*” is with people, video “*Talkshow*” is with a single object and text attributes, and the video “*Packing gift*”



Figure 4: Performance of our method on different video contents by retargeting to ratio from 0.9 to 0.3 of width.




Figure 5: Our method challenges on images with different retargetability levels. The input images and results by F-MultiOp and AAD are quoted from the source paper [49].

is with multiple objects distributed in entire frame. Among the three videos, the video “Dancing” is a camera-move-video. Readers are encouraged to explore our project website (<http://graphics.csie.ncku.edu.tw/RETVI>) to see more results. We can summarize the aspects that enable our results to advance previous works as follows.

We can alleviate the shrinking phenomenon of important objects when reducing sizes effectively. This aspect is gained by benefiting from the CFA module and the critical region loss (\mathcal{L}_{cri}). We observe that this phenomenon often involves reducing videos to a considerably smaller size, *i.e.*, less than 0.6 of width. In the demonstration of Fig.4, we put the results generated by linear scaling to facilitate inferring the quality of our results. It can be seen the dancers (video “Dancing”), the lady (video “Talkshow”), and the gift-box shape (video “Packing gift”) are preserved quite well without shrinking artifacts. It is obvious to observe the video “Talkshow”, the shape of the lady is shrunk and the text is distorted when the video is reduced to 0.3 of width by linear scaling. However, they still appear in a visually pleasing manner in our results.

Being tolerated of various video contents and arbitrary aspect ratios without suffering from distortion could be an advantage of our work. As can be seen in Fig. 4, the important contents in three sample videos are not damaged when the resizing ratios varies from 0.9 to 0.3. Particularly in the video ‘‘Packing gift’’, it is such a challenging video since the important contents are dense and distributed in entire frame. Thanks to the performance of our CFA module, our network can fully understand such dense content. Besides, with the aid from global integrity loss \mathcal{L}_{glo} , all information of this challenging video is preserved harmoniously and integratively with the source.

Being extendable to image retargeting is a plus of our proposed system. Several deep learning-based algorithms have been investigated to resize image recently. Detail comparisons with these methods are presented in the supplementary file. Here, we exhibit the typical cases mentioned as the challenge in this research domain to demonstrate the capability of our RETVI. The research by Tang *et al.* [49] concludes that not all images can be equally resized. This property is defined via a retargetability score, ranged from 0 to 1. They show that the images with low retargetability scores are difficult to preserve visually and semantically important content after resizing. For this challenge, we test our RETVI on the images with different retargetability levels. Fig.5 presents two samples, one is with medium level (score: 0.57), and the other is with low retargetability (score: 0.1). As shown in the figure, in image (a), the compared method F-MultiOp generates result with serious cropping (*i.e.*, the hand and the mug). In image (b), the shape of the main object (*i.e.*, the clock) is distorted by deformation in AAD [38]. In contrast, these phenomena do not occur in our results. Our RETVI produces plausible results in both cases. The results of this experiment demonstrate that our method is potentially used for retargeting images with various image content.

Table 1: Comparison on execution time per frame

Method	Pre-processing		Retargeting
Ref. [33]	<i>Saliency map-4(s)</i>	<i>Segmentation-72(s)</i>	1.32(s)
Ours	No	No	0.016(s)

Finally, running time could be an essential advantage of our method over prior work. Table 1 reports the running time (excluding video reading/writing) that our RETVI competes with a conventional video resizing system [33]. We run two methodologies on the same PC configuration, as reported in Sec.4.1. An input video consisting of 144 frames in size of 1280×720 is retargeted to 50% of the width in this experiment. As can be seen, in Lin *et al.* [33], a time-consuming process falls in generating the saliency map and segmentation, approximately more than 1 minute. And the optimization in the retargeting phase spends about 1.32 seconds. As a result, this conventional system spends minutes to retarget a frame. Meanwhile, our proposed method is an end-to-end system, which only needs approximately 2.3 seconds to resize a video with 144 frames (*i.e.*, 0.016 seconds per frame). It’s worth noting that it could be unfair when comparing the running time between a conventional system versus a deep learning-based one since our RETVI also uses significant cost for training process. Nonetheless, the time cost for training our RETVI is 10 hours, which is such a relatively normal training time for a deep learning model. The workflow and cost make the system [33] cumbersome to be incorporated in demanding services/devices. On contrary, the advantage of running time of RETVI once trained reveals that it is feasible to embed our scheme into a resizing application/service.

Potential application. With the abilities of our method in the above discussion, we can apply it to resize videos to make them perform well in the 9:16 aspect ratio. This aspect ratio became popular when smartphones were created with video capabilities. For example, the optimal measure for an Instagram story is 1080px by 1920px, which means its ratio is 9:16. The same goes for other popular social platforms, *e.g.*, Facebook



Figure 6: Our RETVI (the first row) competes with Adobe Express (the second row) on resizing video to 9:16 aspect ratio.



Figure 7: Predictable results for tall video.

Reels, Youtube Short videos, Tiktok, Instagram story, *etc.* Several existing commercial applications, Adobe Express [2], Veed [3], and Flexier [1] have been providing streamers with a function to edit the ratio of their videos before uploading. Nevertheless, these applications share an identical technique, *i.e.*, a window of size 9:16 is allocated at the middle of video and they manually crop two sides of the video. Therefore, they can keep the ratio of object in the same as in the input video. Yet, in the cases that videos with multiple objects, retaining the important objects in a 9:16 ratio is challenging. We provide a showcase in Fig. 6, in which we compare with Adobe Express. As can be seen, our system does not crop out the tasteful content in this example. The shape of objects may be smaller than those by Adobe Express, but it can capture the tasteful moment on each frame and make it visually pleasing. The visual video can be seen here ¹.

¹<http://graphics.csie.ncku.edu.tw/RETVI/CompareAdobe.mp4>

To be specific, in order to change a video to a tall one, we apply a parameter to control the movement weight in equation (6). As a result, this equation becomes:

$$\begin{cases} \mathbf{H}_x^{ij} = e_x^s \times \mathfrak{R} \times \vartheta \\ \mathbf{H}_y^{ij} = e_y^s \times \mathfrak{R} \times \vartheta \end{cases} \quad (16)$$

here, ϑ is a constant greater than 1. This parameter can be adjusted by users such that the preserved content in the tall video can catch the user’s expectation. This extension enables users to have more predictable results when editing videos with multiple objects to 9:16 aspect ratio. We provide the results generated by different values of ϑ in Fig.7. As can be seen, varying value of ϑ yields different retargeting results. Since the estimated values in matrix \mathbf{H} is capable to avoid distortion, a constant ϑ produces linear increasement. The bigger ϑ can retain object in the tight ratio with those in the input video, but it may not retain much content as the smaller ϑ . Depending on what users interest to preserve, ϑ can be adjusted to achieve the expected results.

4.3 Evaluation metrics

To evaluate the performance of the proposed method, we use four metrics. First, we estimate the distortion degree of retargeted videos. Second, we measure the stability of the retargeted videos. Third, we use an image quality measurement to evaluate the performance of our method on a benchmark image retargeting dataset. And for the last metric, we base on human visual perception. In this evaluation session, we use a set of nine videos and an image dataset. We collect videos data from Youtube such that the videos have diverse content, e.g., single-moving objects, multiple-moving objects, complex backgrounds, or important content distributed in the entire frame. Four of them are the camera-move-videos, and the remainders are object-move-videos. Since few methods focus on retargeting videos, particularly deep learning-based approach, the competitor we quantitatively compare is a typical video resizing system Lin et al. [33]. Moreover, the source code of this work is provided by the authors, hence it is reliable to use and fair for comparison/evaluation. The image dataset we use in our assessment is RetargetMe [43], which consists of 80 images. We evaluate the image data in comparison with four recent SOTA image resizing methodologies.

E.1. Video quality measurement

As the ground truth for video retargeting is not available, it is challenging to define the quality of videos after retargeting process. Therefore, in this regards, we elaborate as follows.

We first adopt the bidirectional similarity measure [46] to evaluate the quality on single video frame. Simakov et al. [46] propose this measurement to describe the coherence and completeness between input and output images [29]. It is widely used for quantitative analysis retargeting results in several works. Given a video with n frames, we have two sets: a set of the source video frames $\mathcal{V}^s = \{F_1^s, \dots, F_n^s\}$ and the other is those retargeted by a certain video retargeting method $\mathcal{V}^t = \{F_1^t, \dots, F_n^t\}$. On each pair (F_k^s, F_k^t) , the error of F_k^t over F_k^s is expressed as:

$$E_k(F_k^s, F_k^t) = \frac{1}{N} \left(\sum_{p \in F_k^s} \min_{q \in F_k^t} \delta(p, q) + \sum_{q \in F_k^t} \min_{p \in F_k^s} \delta(q, p) \right), \quad (17)$$

where $k \in [1 \dots n]$, N is the number of patches on F_k^s and F_k^t ; $\delta(\cdot)$ is defined by sum of squared distance of two patches. Afterwards, we calculate the mean (\mathcal{M}_E) of E_k to define the error degree of a retargeted video. The lower is better.

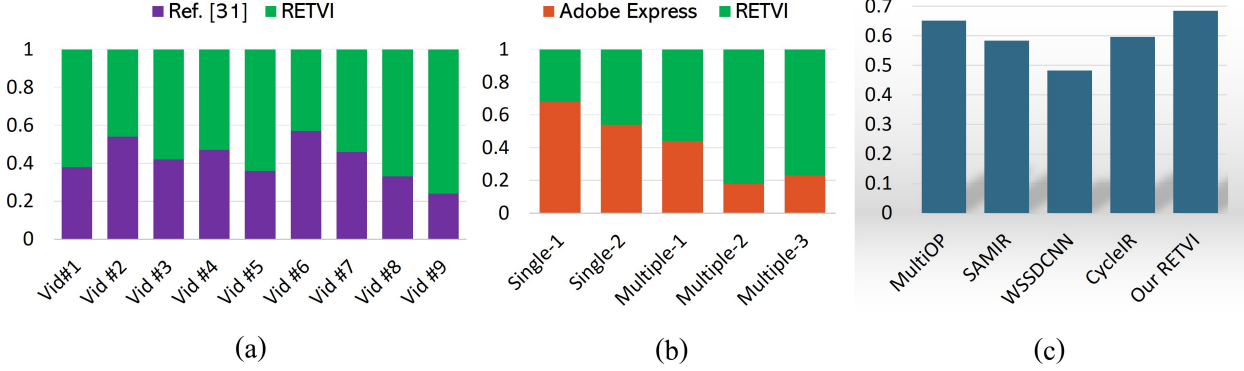


Figure 8: Analysis of human perception on group G-1 (a) and group G-2 (b). (c) is the comparison on ARS score.

E.2. Video stability measurement

To measure the stability of resized videos without annotated data, we estimate the differences of adjacent frames in the retargeted video and compare them against those in the source one. This concept is used in various video generation applications, *e.g.*, video stylization [11, 25], video resequencing [26]. The calculation is based on the fact that the source videos are temporally coherent, and the retargeted results are rendered from the same frame set with them but in different ratios. We, therefore, treat the value in the source as the ground truth to judge the stability degree of the resulting videos. Accordingly, a generated video with a score that is tightly asymptotic to the ground truth would be in good stability.

Given a video with frame set $\mathcal{V} = \{F_t\}, t \in [1 \dots n]$, n is the total frames of \mathcal{V} , the difference degree of two adjacent frames is formulated as:

$$\mathcal{D}_{t \rightarrow t-1} = \frac{1}{100} \times \frac{|F_t - F_{t-1}|}{H \times W}, \tag{18}$$

where H and W are the height and width of frames in set \mathcal{V} . $|F_t - F_{t-1}|$ returns a residual of two frames, which can be seen in the supplementary. Accordingly, the stability of the video is:

$$STB = \frac{1}{n} \sum_{t=1}^n \mathcal{D}_{t \rightarrow t-1}. \tag{19}$$

For this metric, the lower score represents a better stability.

In Table 2, we analyze the quantitative evaluation and comparison using the above two metrics, \mathcal{M}_E and STB . With nine videos we prepared, we used Lin et al. [33]’s method and our RETVI to generate results. As can be seen, our method achieves a lower distortion on average compared to Lin’s. The highest \mathcal{M}_E in our results is 3.17; two videos have relatively low distortion (*i.e.*, video #4 and video #9). Inferring to Lin’s results, the distortion level in theirs is higher than ours in most cases. The lowest \mathcal{M}_E of their results is reported at 2.97, which is higher than our average score. This analysis reveals that the retargeting results generated by our system are less distorted than Lin et al. [33]. In terms of stability, Lin’s method and our RETVI have comparable scores. However, we can see that their values of STB of the camera-move-videos are relatively higher than those of the reverse cases. The analysis results reveal that our model is stable when working on diverse input videos.

E.3. Image quality measurement

Although our focus is not retargeting images, as mentioned above, our method effectively applies to single images. To quantitatively evaluate this performance, we adopt the Aspect ratio similarity (ARS) [54]. This algorithm evaluates the visual quality of retargeted images by exploiting the local block changes with a visual importance pooling strategy. The ARS score represents the geometric changes during retargeting [54]. For this metric, all of the images in the RetargetMe are used. We compare the results of these images generated by our method against a typical conventional method, multi-operator [42], and three deep learning-based techniques [10, 48, 57]. The results of multi-operator [42] are encompassed in the RetargetMe dataset, while those of WSSDCNN [10], Cycle-IR [48], and SAMIR [57] are generated from the source code released by the authors.

Table 2: Analysis on the quality of video results - The last four rows are camera-move videos

Methods	Source video	Lin et al. [33]		Our method	
Testing data	$STB \downarrow$	$\mathcal{M}_E \downarrow$	$STB \downarrow$	$\mathcal{M}_E \downarrow$	$STB \downarrow$
1-Talkshow	0.057	3.24	0.062	2.81	0.063
2-Packing gift	0.041	3.71	0.058	2.63	0.055
3-Holding mug	0.052	4.15	0.066	3.16	0.061
4-Cap movie	0.063	2.97	0.071	1.97	0.06
5-Shewing	0.047	3.62	0.051	2.84	0.048
6-Masha	0.081	3.52	0.09	3.17	0.082
7-Dancing	0.075	4.83	0.082	2.69	0.078
8-Moana Movie	0.089	4.29	0.1	2.24	0.09
9-Uptown dance	0.084	3.74	0.094	1.86	0.087
Average	0.065	3.78	0.74	2.59	0.069

The quantitative comparison of the ARS score is shown in Fig. 8(c). As can be seen, our RETVI achieves a higher score comparing all the competitors. Although the deviation is not significant, it is sufficient to demonstrate our method’s performance. It is because the RetargetMe dataset encompasses challenging images. A good performance on this dataset could be assessed as a reliable method.

E.4. User study

We further estimate the performance of our RETVI based on human visual perception. We conduct a user study on two groups. In one group (G-1), we recruit 13 users (nine of them with graphics-related backgrounds). For the other (G-2), we invite nine users who are either tiktokers or senior users on the existing social platforms. In G-1, we use nine videos mentioned in evaluations E.1, E.2 to generate results retargeting to a half size of width by our RETVI and Lin et al. [33]. In G-2, we prepare five videos (two with a single object and three with multiple objects) and resize them to the 9:16 aspect ratio by Adobe Express [2] and our RETVI. In both groups, participants are shown two retargeted videos at one time and asked to choose the one




Figure 9: Comparison with Lin et al. [33]

they prefer. We receive 117 responses on G-1 and 45 responses on G-2. We then compute the percentage of votes for each video.

Fig. 8(a)-(b) shows the statistics of users’ preferences. It can be seen that our method receives majority votes from the participants in G-1 and G-2 alike. There are two cases in G-1 (e.g., vid#2 and #6), in which users judge Lin’s results are better than ours. However, the difference is not significant (0.08 and 0.14, respectively), and the score of our RETVI is relatively higher in the remaining seven cases (77.78% in total showcases). The data in G-2 reveals that users prefer the tall videos by Adobe Express to ours if the videos have a single object. Yet, our RETVI wins Adobe Express in cases with multiple objects.




Figure 10: Comparison between our results (b) with Lee et al. [27] (c). Results in (c) are quoted from the source paper [27].

4.4 Visual comparisons to prior work

Fig. 9 visualizes the comparison between our results and Lin *et al.* [33]’s. The visual results reveal that Lin’s results suffer from some shortcomings. First, several noticeable artifacts occur in the main object and line structure of the background region (highlighted by yellow rectangles). The reason is that their scheme relies on a saliency map [16] to estimate the movement weight of quad vertices. As a result, if the content frame is dense, the saliency method may not be tolerated to analyze the frame effectively. For example, the shape of the lady in frame #30 and frame #87 is plausible, but the door is shrunk significantly. It is even smaller than linear scaling. Besides, in frame #199, the line structure of the door appears quite well, but the head of the lady is distorted. In contrast, our method produces more visually pleasing results without these distortions. Second, the generated video by Lin *et al.* [33] in this case still has the noticeable flicking artifact. See the video here² for better visualization.

In Fig. 10, we provide a visual comparison with Lee *et al.* [27], an object detection approach for retargeting video. As can be observed, there is no artifact or distortion in their results and ours alike. The noticeable point here is the ratio of objects in the retargeted frames. With the goal that the contents inside the bounding boxes must remain intact in the retargeted frame, Lee *et al.* [27] do a good job, but their appearance is in a close ratio with those in the input frames. For example, visually inspecting this figure, the background contents (e.g., the door on the left, the wall on the right, or the plants) are scaled up when the width increases in this experiment. However, the shape of the lady is hard to recognize the changes in the ratio over the source frame. These may reduce the harmonization of the results. In contrast, our method does not suffer as they do. Our results are more harmonious in both background contents and important objects. The visual video can be seen here³.

Apart from the above comparisons with video retargeting methodologies, we further discuss the ability of our method via visual comparison with four recent state-of-the-art image retargeting systems, WSSDCNN [10], SAMIR [57], grid encoding model [22] and Cycle-IR [48]. Detail of discussion and visualization on this comparison is presented in the supplementary file.

4.5 Ablation Study

A.1. Verify the effectiveness of CFA module

The CFA module is proposed to alternate the time-consuming preprocessing phase in such a conventional video retargeting system. The CFA is structured as an improvement from a plain UNet design. Our proposed CFA can capture more fine-grained feature, which is an essential key for our retargeting system to tolerate with diverse content videos. We demonstrate the effectiveness of CFA module by removing it from our training and train our RETVI with the plain UNet [41]. We show these ablation analyses in Fig. 11, in which the visualization of Grad-CAMs [44] and energy maps are obtained from our network with and without the CFA module. The results reveal that with the CFA module, our network has much larger attended regions. This enables our RETVI to play with more complex video contents without suffering from distortion and linear-like appearance.

A.2. Study on the impact of loss functions

Performance of our RETVI is affected by the optimization of the loss functions. Here, we discuss the impact of each component we configured in the objective function. To verify this, we remove in turn each loss function in Eq. (8) from our training. The ablated results are presented in Fig. 12. It is obvious that without

²<http://graphics.csie.ncku.edu.tw/RETVI/CompareLin.mp4>

³<http://graphics.csie.ncku.edu.tw/RETVI/CompareLee.mp4>

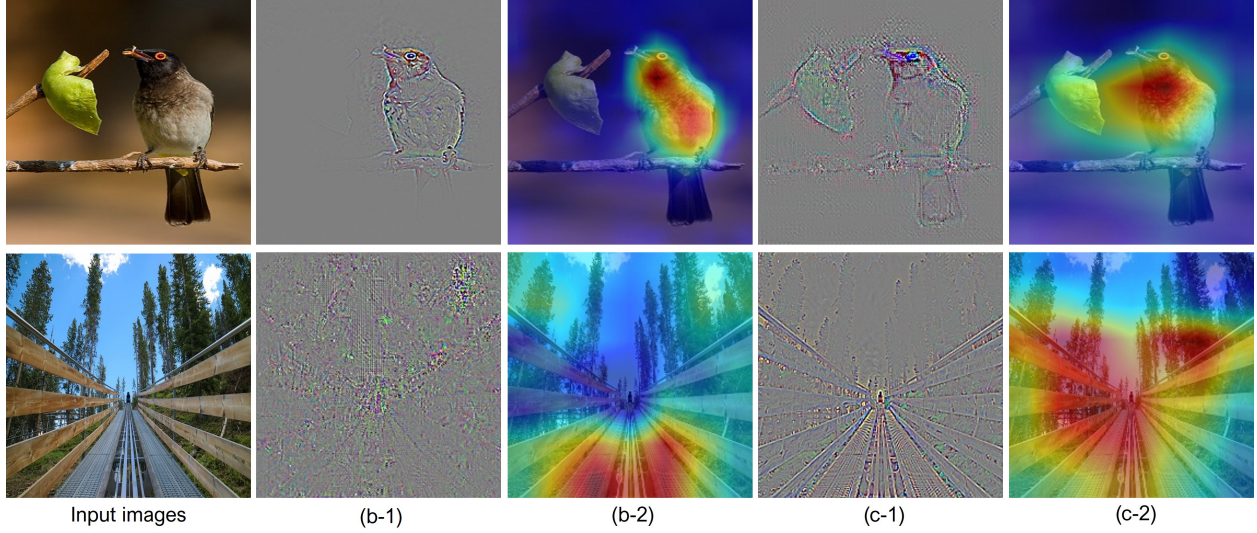


Figure 11: The Grad-Cam visualization (b-2, c-2) and energy map (b-1, c-1) of our CFA module (c- columns) versus the standard UNet (b- columns).

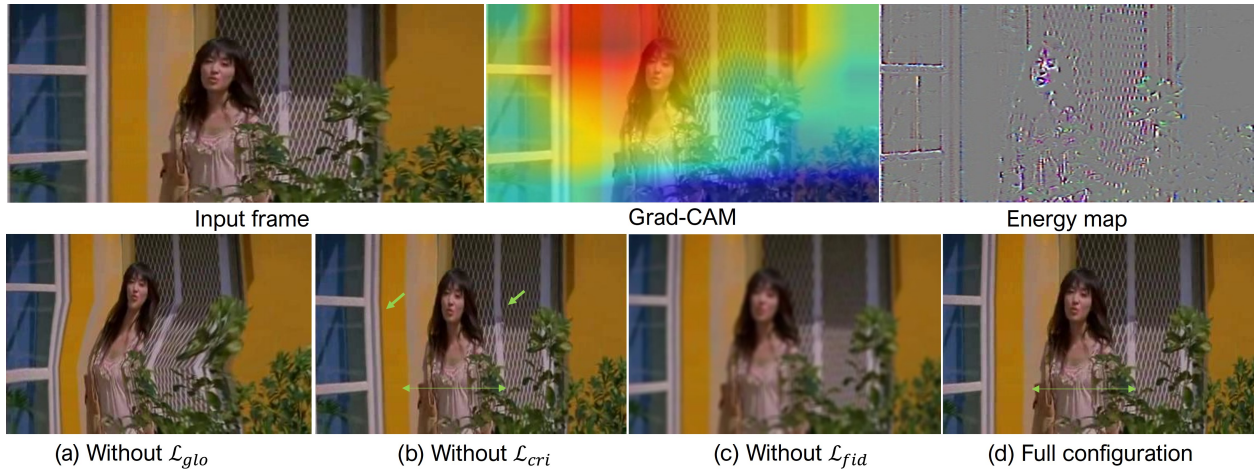


Figure 12: Ablated results for loss function. The resized results (second row) in this experiment are resized to ratio 0.7.

each component in the total loss function, it is challenging to obtain ideal retargeting results. To be specific, we can see that without the guidance of \mathcal{L}_{cri} , the important region in the frame is damaged. On the one hand, the object in the foreground is shrunk linearly, on the other hand, the background captured as critical region by our CFA is distorted. We also discuss the effect of the global loss \mathcal{L}_{glo} in Fig.12(b). Without \mathcal{L}_{glo} , we fail to preserve the initiative of the content frame. Meanwhile, removing \mathcal{L}_{fid} could yield resizing form without damaging the content but blurry effect is a negative side. The full configuration facilitates our model producing visual pleasing result. Also, the absence of \mathcal{L}_{tem} causes serious flickering artifact in the generated videos. The aid of this loss in our objective serves resulting videos in good stability. See the ablated result of \mathcal{L}_{tem} here⁴.

Limitation. In the cases that the input videos are the landscape scenes or too dense, our method may not perform well. The failure phenomenon here is that the retargeted videos are quite similar to linear scale. We

⁴<http://graphics.csie.ncku.edu.tw/RETVI/TemporalLoss.mp4>

note here that distortion does not occur in these cases. It is because our ADE module fails to differentiate which object/region is important. Thus, the estimated energy of pixels is relatively identical. Another limitation falls into the failure of the loss \mathcal{L}_{glo} . Since we rely on performance of ResNet to define the integrity of estimated frame versus the input one, the failure of ResNet leads to \mathcal{L}_{glo} to be disabled. The \mathcal{L}_{glo} performs well in most of the cases, but it is not guaranteed to be stable in the videos with multiple objects distributed in entire frame. Another limitation is that if the main objects located in the leftmost of rightmost, the resized results may not appealing due to the cropping effect as we discussed in Sec.4.1. An example in this case is visualized in the supplementary file.

5 Conclusion

In this paper, we propose a new RETVI framework for retargeting videos. With two modules configured in our method, our RETVI presents high performance in handling videos with diverse contents and produces visually pleasing results when retargeting to arbitrary aspect ratios. The analysis and experimental results demonstrate that our method substantially advances prior works. With the fast running time of our end-to-end RETVI, our system is potentially embedded into a video resizing application/service. We perceive that our system can bypass the computational bottlenecks in conventional methods. And it is potential to extend for stereo image/video retargeting. For the shortcoming we discussed, we plan to investigate techniques that configure the loss function to be independent from the existing feature extractor. In terms of cropping effect, a possible way can improve is automatically define physical region of the important content. This knowledge could serve us to shift the rendering window more appropriately. Besides, developing a text-driven framework to consider semantic issue in the retargeting videos and investigating technique to retarget videos with enlarging and reducing two dimensions simultaneously are also a possibly extension in our near future. This could be a potential way to visualize users' expectation in such a video retargeting system.

References

- [1] Flexier, 2021. URL <https://flixier.com/tools/resize-video/resize-video-for-instagram>. Accessed: March 24, 2023.
- [2] Adobe express, 2022. URL <https://www.adobe.com/express/feature/video/resize/instagram>. Accessed: March 24, 2023.
- [3] Veed, version 2021. URL <https://www.veed.io/tools/resize-video/resize-video-for-instagram-story>. Accessed: March 24, 2023.
- [4] Adobe after effects. <https://helpx.adobe.com/after-effects>, version 2022. Accessed: 2023-09-07.
- [5] Arcore. <https://developers.google.com/ar>, version 2023. Accessed: 2023-09-07.
- [6] Unity. <https://www.trustradius.com/products/unity>, version 2023. Accessed: 2023-09-07.
- [7] B. Asheghi, P. Salehpour, A. M. Khiavi, and M. Hashemzadeh. A comprehensive review on content-aware image retargeting: From classical to state-of-the-art methods. *Signal Processing*, page 108496, 2022.
- [8] S. Avidan and A. Shamir. Seam carving for content-aware image resizing. In *ACM SIGGRAPH 2007 papers*, pages 10–es. 2007.
- [9] W.-H. Cheng, C.-W. Wang, and J.-L. Wu. Video adaptation for small display based on content recomposition. *IEEE Transactions on Circuits and Systems for Video Technology*, 17(1):43–58, 2006.

- [10] D. Cho, J. Park, T.-H. Oh, Y.-W. Tai, and I. So Kweon. Weakly-and self-supervised learning for content-aware deep image retargeting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4558–4567, 2017.
- [11] Y. Deng, F. Tang, W. Dong, H. Huang, C. Ma, and C. Xu. Arbitrary video style transfer via multi-channel correlation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1210–1217, 2021.
- [12] W. Dong, F. Wu, Y. Kong, X. Mei, T.-Y. Lee, and X. Zhang. Image retargeting by texture-aware synthesis. *IEEE Transactions on Visualization and Computer Graphics*, 22(2):1088–1101, 2015.
- [13] X. Fan, J. Lei, J. Liang, Y. Fang, N. Ling, and Q. Huang. Stereoscopic image retargeting based on deep convolutional neural network. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(12):4759–4770, 2021.
- [14] R. Furuta, I. Tsubaki, and T. Yamasaki. Fast volume seam carving with multipass dynamic programming. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(5):1087–1101, 2016.
- [15] R. Gallea, E. Ardizzone, and R. Pirrone. Physical metaphor for streaming media retargeting. *IEEE Transactions on Multimedia*, 16(4):971–979, 2014.
- [16] S. Goferman, L. Zelnik-Manor, and A. Tal. Context-aware saliency detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(10):1915–1926, 2011.
- [17] M. Grundmann, V. Kwatra, M. Han, and I. Essa. Discontinuous seam-carving for video retargeting. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 569–576. IEEE, 2010.
- [18] M. Hashemzadeh, B. Asheghi, and N. Farajzadeh. Content-aware image resizing: An improved and shadow-preserving seam carving method. *Signal Processing*, 155:233–246, 2019.
- [19] Y. Jin, L. Liu, and Q. Wu. Nonhomogeneous scaling optimization for realtime image resizing. *The Visual Computer*, 26:769–778, 2010.
- [20] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [21] J. Kiess, S. Kopf, B. Guthier, and W. Effelsberg. A survey on content-aware image and video retargeting. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 14(3):1–28, 2018.
- [22] Y. Kim, H. Eun, C. Jung, and C. Kim. A quad edge-based grid encoding model for content-aware image retargeting. *IEEE Transactions on Visualization and Computer Graphics*, 25(12):3202–3215, 2018.
- [23] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [24] S. Kopf, J. Kiess, H. Lemelson, and W. Effelsberg. Fscav: Fast seam carving for size adaptation of videos. In *Proceedings of the 17th ACM international conference on multimedia*, pages 321–330, 2009.
- [25] T.-N.-H. Le, Y.-H. Chen, and T.-Y. Lee. Structure-aware video style transfer with map art. *ACM Transactions on Multimedia Computing, Communications and Applications*.
- [26] T.-N.-H. Le, S.-Y. Yao, C.-T. Wu, and T.-Y. Lee. Regenerating arbitrary video sequences with distillation path-finding. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–14, 2023. doi: 10.1109/TVCG.2023.3237739.
- [27] S. J. Lee, S. Lee, S. I. Cho, and S.-J. Kang. Object detection-based video retargeting with spatial-temporal consistency. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(12): 4434–4439, 2020.

- [28] B. Li, L.-Y. Duan, J. Wang, R. Ji, C.-W. Lin, and W. Gao. Spatiotemporal grid flow for video retargeting. *IEEE Transactions on Image Processing*, 23(4):1615–1628, 2014.
- [29] Y. Liang, Y.-J. Liu, and D. Gutierrez. Objective quality prediction of image retargeting algorithms. *IEEE Transactions on Visualization and Computer Graphics*, 23(2):1099–1110, 2016.
- [30] J. Lin, T. Zhou, and Z. Chen. Deepir: A deep semantics driven framework for image retargeting. In *2019 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 54–59. IEEE, 2019.
- [31] S.-S. Lin, I.-C. Yeh, C.-H. Lin, and T.-Y. Lee. Patch-based image warping for content-aware retargeting. *IEEE Transactions on Multimedia*, 15(2):359–368, 2012.
- [32] S.-S. Lin, C.-H. Lin, S.-H. Chang, and T.-Y. Lee. Object-coherence warping for stereoscopic image retargeting. *IEEE transactions on circuits and systems for video technology*, 24(5):759–768, 2013.
- [33] S.-S. Lin, C.-H. Lin, I.-C. Yeh, S.-H. Chang, C.-K. Yeh, and T.-Y. Lee. Content-aware video retargeting using object-preserving warping. *IEEE Transactions on Visualization and Computer Graphics*, 19(10):1677–1686, 2013.
- [34] S.-S. Lin, C.-H. Lin, Y.-H. Kuo, and T.-Y. Lee. Consistent volumetric warping using floating boundaries for stereoscopic video retargeting. *IEEE Transactions on Circuits and Systems for Video Technology*, 26(5):801–813, 2015.
- [35] D. Liu, Z. Wu, X. Lin, and R. Ji. Towards perceptual video cropping with curve fitting. *Multimedia Tools and Applications*, 75:12465–12475, 2016.
- [36] H. Liu, X. Xie, W.-Y. Ma, and H.-J. Zhang. Automatic browsing of large pictures on mobile devices. In *Proceedings of the eleventh ACM international conference on Multimedia*, pages 148–155, 2003.
- [37] S. Liu, Z. Wei, Y. Sun, X. Ou, J. Lin, B. Liu, and M.-H. Yang. Composing semantic collage for image retargeting. *IEEE Transactions on Image Processing*, 27(10):5032–5043, 2018.
- [38] D. Panozzo, O. Weber, and O. Sorkine. Robust image retargeting via axis-aligned deformation. In *Computer Graphics Forum*, volume 31, pages 229–236. Wiley Online Library, 2012.
- [39] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019.
- [40] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017.
- [41] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [42] M. Rubinstein, A. Shamir, and S. Avidan. Multi-operator media retargeting. *ACM Transactions on Graphics (TOG)*, 28(3):1–11, 2009.
- [43] M. Rubinstein, D. Gutierrez, O. Sorkine, and A. Shamir. A comparative study of image retargeting. In *ACM SIGGRAPH Asia 2010 papers*, pages 1–10. 2010.
- [44] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.

- [45] F. Shao, W. Lin, W. Lin, Q. Jiang, and G. Jiang. Qoe-guided warping for stereoscopic image retargeting. *IEEE Transactions on Image Processing*, 26(10):4790–4805, 2017.
- [46] D. Simakov, Y. Caspi, E. Shechtman, and M. Irani. Summarizing visual data using bidirectional similarity. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- [47] E. Song, M. Lee, and S. Lee. Carvingnet: content-guided seam carving using deep convolution neural network. *IEEE Access*, 7:284–292, 2018.
- [48] W. Tan, B. Yan, C. Lin, and X. Niu. Cycle-ir: Deep cyclic image retargeting. *IEEE Transactions on Multimedia*, 22(7):1730–1743, 2019.
- [49] F. Tang, W. Dong, Y. Meng, C. Ma, F. Wu, X. Li, and T.-Y. Lee. Image retargetability. *IEEE Transactions on Multimedia*, 22(3):641–654, 2019.
- [50] B. Wang, H. Xiong, Z. Ren, and C. W. Chen. Deformable shape preserving video retargeting with salient curve matching. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 4(1): 82–94, 2014.
- [51] B. Yan, B. Yuan, and B. Yang. Effective video retargeting with jittery assessment. *IEEE Transactions on Multimedia*, 16(1):272–277, 2013.
- [52] J. Yang, F. Guo, S. Chen, J. Li, and J. Yang. Industrial style transfer with large-scale geometric warping and content preservation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7834–7843, 2022.
- [53] J. Zhang, S. Li, and C.-C. J. Kuo. Compressed-domain video retargeting. *IEEE Transactions on Image Processing*, 23(2):797–809, 2013.
- [54] Y. Zhang, W. Lin, X. Zhang, Y. Fang, and L. Li. Aspect ratio similarity (ars) for image retargeting quality assessment. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1080–1084. IEEE, 2016.
- [55] Y.-F. Zhang, S.-M. Hu, and R. R. Martin. Shrinkability maps for content-aware video resizing. In *Computer Graphics Forum*, volume 27, pages 1797–1804. Wiley Online Library, 2008.
- [56] M. Zhao and Q. Ling. Pwstabilenet: Learning pixel-wise warping maps for video stabilization. *IEEE Transactions on Image Processing*, 29:3582–3595, 2020.
- [57] Y. Zhou, Z. Chen, and W. Li. Weakly supervised reinforced multi-operator image retargeting. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(1):126–139, 2020.