Guided Interactive Video Object Segmentation Using Reliability-Based Attention Maps

Yuk Heo
Korea University
yukheo@mcl.korea.ac.kr

Yeong Jun Koh
Chungnam National University
yjkoh@cnu.ac.kr

Chang-Su Kim
Korea University
changsukim@korea.ac.kr

Abstract

We propose a novel guided interactive segmentation (GIS) algorithm for video objects to improve the segmentation accuracy and reduce the interaction time. First, we design the reliability-based attention module to analyze the reliability of multiple annotated frames. Second, we develop the intersection-aware propagation module to propagate segmentation results to neighboring frames. Third, we introduce the GIS mechanism for a user to select unsatisfactory frames quickly with less effort. Experimental results demonstrate that the proposed algorithm provides more accurate segmentation results at a faster speed than conventional algorithms. Codes are available at https://github.com/yuk6heo/GIS-RAmap.

1. Introduction

Video object segmentation (VOS) is a task to cut out objects of interest in a video. It is useful in various applications such as video editing, video summarization, video inpainting, and self-driving cars. VOS is challenging since it should deal with multiple objects, object deformation, and object occlusion. Because of this difficulty, semi-supervised VOS, which uses a fully annotated segmentation mask in the first frame, has been widely researched. This approach can improve the segmentation performance but requires a lot of time and effort for annotations (*e.g.* around 79 seconds per instance [4]). Also, it does not have a fallback mechanism when unsatisfactory results are obtained.

Interactive VOS adopts user-friendly annotations, *e.g.* scribbles, which are simple enough to provide repeatedly. Figure 1 shows the round-based interactive VOS process. First, a user selects a target frame and draws scribble annotations on it. After extracting query object information from the scribbles, the algorithm obtains segmentation results for all frames. Second, the user finds a frame with unsatisfactory results and then provides additional scribbles. The algorithm then exploits both sets of scribbles to refine the VOS results. This is repeated until the user is satisfied.
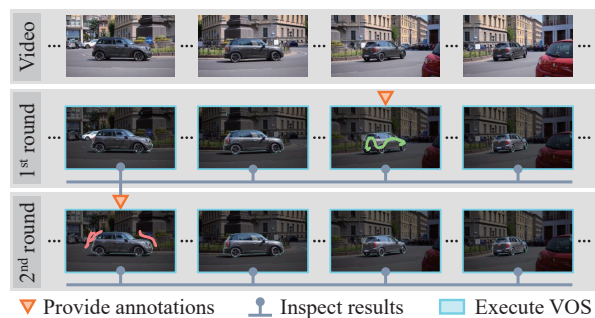

Figure 1: Illustration of the round-based interactive VOS. The proposed algorithm greatly reduces the inspection time, by guiding users to select a frame for annotations efficiently and effectively. It is recommended to watch the supplementary video with a real-time demo of the proposed guided interactive system.

Recently, an automatic simulation scheme of the round-based interactive VOS was designed in [4]. However, the simulation is significantly different from real applications in that it immediately determines segmentation results to correct, by comparing them with the ground-truth. In contrast, a real user should spend considerable time to inspect the results and select poorly segmented regions. Since conventional interactive VOS algorithms [7, 20, 22, 24] have been developed based on the simulation in [4], they do not consider the time for finding unsatisfactory results in practice. In contrast, we propose a guided interactive segmentation (GIS) algorithm for video objects, which guides users to find poorly segmented regions quickly and effectively.

Moreover, although interactive VOS can use the information in N annotated frames in the Nth round, the conventional algorithms [7, 20] do not exploit those multiple annotated frames thoroughly. Heo *et al.* [7] simply average the features from multiple annotated frames. Miao *et al.* [20] use only the best matching result between a target frame and multiple annotated frames. On the contrary, we analyze the reliability of each annotated frame to refine segmentation results in a target frame more accurately.

In this paper, we propose the GIS algorithm using reliability-based attention (R-attention) maps. First, we transfer query object information from annotated frames to a target frame using R-attention maps, which represent pixel-wise reliability of the annotated frames. Next, we perform intersection-aware propagation to propagate segmentation results to neighboring frames sequentially. Third, we compute a guidance score, called R-score, to reduce or remove the processing time for selecting the frame to be annotated in each round. Experimental results demonstrate that the proposed GIS algorithm outperforms recent state-of-the-arts in both the interactive VOS simulation in [4] and real-world applications.

This paper has three main contributions:

1. Two novel operators to exploit multiple annotations and neighboring results are developed for VOS: R-attention and intersection-aware propagation modules.

2. We propose the notion of guidance in interactive VOS.

3. The proposed GIS algorithm outperforms the state-of-the-arts significantly in both speed and accuracy.

2. Related Work

Unsupervised VOS: It is a task to find primary objects [14] without any user annotations in video sequences. Traditional approaches [13,15,16,25,34] use motion, object proposals, or saliency to solve this problem. Recently, with the availability of big VOS datasets [26,38], many deep-learning-based unsupervised methods [11,18,30,35,39–41] have been proposed.

Semi-supervised VOS: In semi-supervised VOS, a user provides fully annotated masks for target objects at the first frame. Many algorithms have been developed to extract significant features for target objects using user annotations. Early deep learning methods [1, 3, 19, 32] focused on fine-tuning networks using annotation masks at the first frames. Instead of the computationally demanding fine-tuning, some algorithms employ optical flow for initial segment propagation [12] or motion feature extraction [9]. Also, networks to refine segmentation results in previous frames without using motion have been developed in [17,21]. Instead of the first frame, optimal frames to be annotated in semi-supervised VOS were determined in [6]. Matching-based algorithms [5,10,23,24,31] perform pixel-wise feature matching between annotated and target frames to segment out target objects. For example, key-value memory operations based on non-local networks [36] are performed in [23,24] to perform the matching between a target frame and already segmented frames. In [37], the selection network, which predicts scores of previously segmented frames, is used to choose the frames for segmentation propagation.

Interactive VOS: It aims at achieving satisfactory VOS results through an iterative process of drawing simple annotations, such as scribbles, point clicks, or bounding boxes. Early interactive VOS algorithms [27, 28, 33] constructed graph models, by connecting pixels with edges and then assigning edge weights using hand-craft features. Segmentation results for query objects were then obtained by graph optimization techniques.

Recently, deep-learning methods have been developed for interactive VOS. Benard and Gygli [2] used point clicks to extract an object mask in a single frame and apply a semi-supervised VOS algorithm to propagate the mask. Chen *et al*. [5] employed pixel-wise metric learning to cut out a query object using a few point clicks. Caelles *et al*. [4] introduced a round-based interactive VOS process and the automatic simulation algorithm to mimic human interactions in real applications. Many recent interactive algorithms [7,20,22,24] follow this round-based process.

Oh *et al*. [22] developed two segmentation networks for interactive VOS: the first one estimates target object regions from user interactions and the second one propagates the segmentation results to neighboring frames. Miao *et al*. [20] proposed networks to obtain segmentation results through both interaction and propagation. They employed global and local distance maps in [31] to match a target frame to an annotated frame and the previous frame, respectively. Heo *et al*. [7] designed global and local transfer modules to effectively transfer features in annotated and previous frames to a target frame. Oh *et al*. [24] encoded annotation regions into keys and values in a non-local manner. These interactive VOS algorithms [7, 20, 22, 24], however, have limitations. First, they do not consider the processing time to select the poorest segmentation results, on which additional annotations are provided. Second, they do not fully exploit the property that scribble data in multiple annotated frames have different reliability and different relevance to a target frame.

3. Proposed Algorithm

The proposed algorithm cuts query objects off in a video $\mathcal{I} = \{I_1, \ldots, I_T\}$ with T frames interactively using sparse annotations (scribbles or points) in each segmentation round. Let I_{a_i} at time instance a_i be the annotated frame in the ith round. First, the segmentation is performed bidirectionally starting from I_{a_1}. Subsequently, in the N-th round, it is also done bidirectionally, but using all previously annotated frames $\mathcal{I}_a = \{I_{a_1}, \ldots, I_{a_N}\}$.

Figure 2 shows how the proposed algorithm segments a target frame I_t. First, we encode I_t into the frame feature $\mathbf{F}_t$. Second, we obtain the interfused object feature $\mathbf{G}_t$ by combining query object information in all annotated frames in $\mathcal{I}_a$ using the R-attention module. Third, using the neighbor frame $I_n \in \{I_{t-1}, I_{t+1}\}$ and its segmentation result Y_n, we perform the intersection-aware propagation to yield the
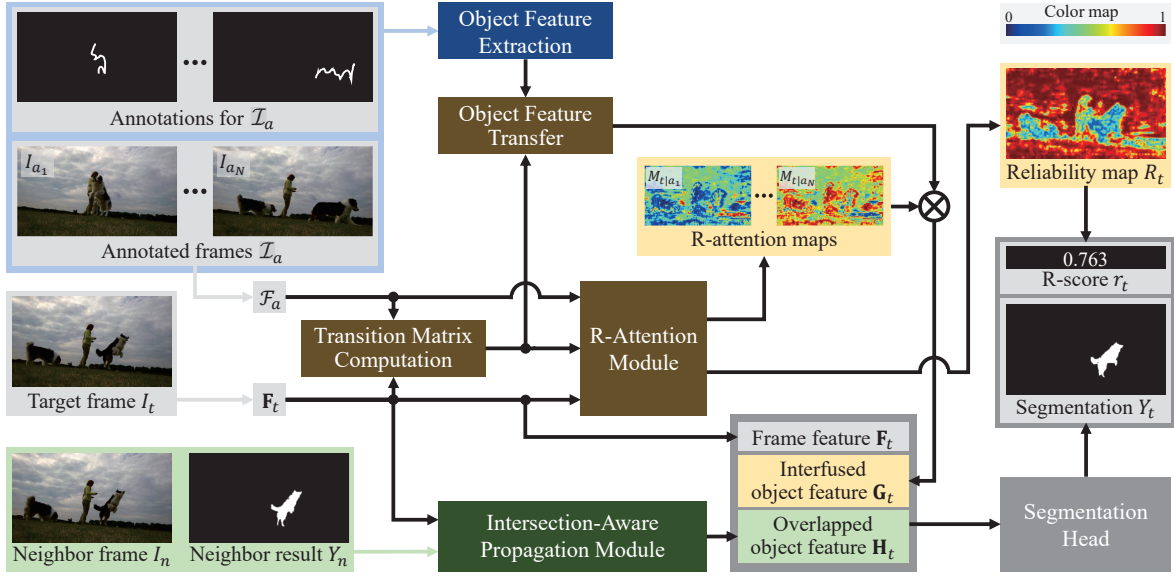
Figure 2: An overview of the proposed GIS algorithm.

overlapped object feature $\mathbf{H}_t$. Last, the segmentation head decodes these three features $\mathbf{F}_t$, $\mathbf{G}_t$, and $\mathbf{H}_t$ to generate the segmentation result Y_t of the target frame I_t.

Moreover, we propose a novel guidance mechanism for interactive VOS. From the R-attention module, we extract the reliability map R_t to represent pixel-wise reliability of the segmentation result Y_t. Also, by averaging these pixel-wise scores, we obtain the R-score r_t. These guidance data R_t and r_t enable a user to select a less reliable frame and provide annotations on it more quickly and more effectively.

3.1. Interfused Object Feature

As in [7], we transfer segmentation information in each annotated frame into a target frame. However, whereas [7] combines the information from multiple annotated frames simply through averaging, we fuse transferred object features based on their reliability levels. To this end, we develop the R-attention module.

Transition matrix computation: We encode each annotated frame $I_{a_i} \in \mathcal{I}_a$ into $\mathbf{F}_{a_i}$ to obtain the annotated frame feature set $\mathcal{F}_a = \{\mathbf{F}_{a_1}, \ldots, \mathbf{F}_{a_N}\}$. Each frame feature is an $HW \times C_1$ matrix, in which each row contains the C_1-dimensional feature vector for a pixel. Here, $H \times W$ is the spatial resolution of the feature. Then, using the ith annotated feature $\mathbf{F}_{a_i}$ and the target feature $\mathbf{F}_t$, we obtain the transition matrix

$$\mathbf{A}_{a_i \to t} = \text{softmax}\left(\phi_A(\mathbf{F}_t) \times \phi_A(\mathbf{F}_{a_i})^T\right) \quad (1)$$

of size $HW \times HW$. Here, ϕ_A is a feature transform, implemented by a learnable 1×1 convolution, to reduce the dimension of each row vector from C_1 to C_2. Note in Fig-

ure 2 that a frame feature is used by different modules. To adapt the same feature for different purposes, we employ multiple feature transforms, including ϕ_A.

In (1), the softmax operation is applied to each column. Thus, $\mathbf{A}_{a_i \to t}$ is a positive matrix with each column adding to 1. It is hence the transition matrix [29], whose entry in row r and column c represents the probability that the cth pixel in I_{a_i} is mapped to the rth pixel in I_t. We compute the transition matrices from all annotated frames to I_t to yield the transition matrix set $\mathcal{A}_{a \to t} = \{\mathbf{A}_{a_1 \to t}, \ldots, \mathbf{A}_{a_N \to t}\}$.

Object feature transfer: We use annotations on I_{a_i} to generate an object saliency map via a sparse-to-dense network in Figure 3. We adopt A-Net [7], the encoder of which is based on SE-ResNet50 [8], as the sparse-to-dense network. To form the feature of the query object, we combine three intermediate features: R3 and R5 features from the encoder and the context feature of the penultimate layer of the decoder. After making their spatial resolutions identical, we convolve and then concatenate them. Then, the concatenated feature passes through another convolution layer to form the object feature $\mathbf{E}_{a_i}$ of size $HW \times C_3$. Consequently, the object feature set $\mathcal{E}_a = \{\mathbf{E}_{a_1}, \ldots, \mathbf{E}_{a_N}\}$ is obtained from the N annotated frames.

We transfer all object features in $\mathcal{E}_a$ to the target frame I_t by

$$\mathbf{E}_{t|a_i} = \mathbf{A}_{a_i \to t} \times \mathbf{E}_{a_i} \quad (2)$$

using the transition matrix $\mathbf{A}_{a_i \to t}$ in (1). Thus, we have the transferred object feature set $\mathcal{E}_{t|a} = \{\mathbf{E}_{t|a_1}, \ldots, \mathbf{E}_{t|a_N}\}$, which encodes the object information in I_t approximately. Since the reliability of the transition matrix $\mathbf{A}_{a_i \to t}$ is different for each i, it is unreasonable to exploit $\mathbf{E}_{t|a_i}$ equally
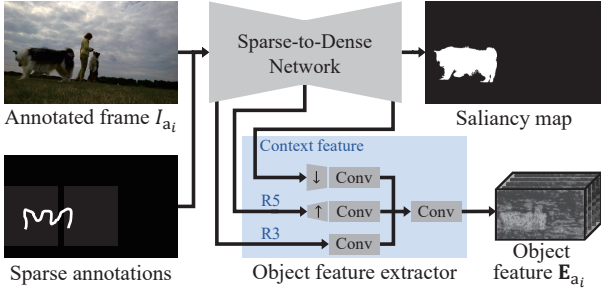
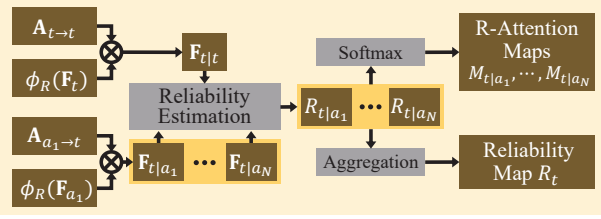Figure 3: A diagram of the object feature extraction.

Figure 4: A diagram of the R-attention module.

for the query object segmentation. To address this issue, we propose the R-attention mechanism.

R-attention: Figure 4 shows how to generate R-attention maps. Similar to (2), let $\mathbf{F}_{t_2|t_1} = \mathbf{A}_{t_1 \to t_2} \times \phi_R(\mathbf{F}_{t_1})$ denote the transferred frame feature from I_{t_1} to I_{t_2}, where ϕ_R is a feature transform. We obtain the feature difference matrix

$$\mathbf{D}_{t|a_i} = [\mathbf{F}_{t|a_i} - \mathbf{F}_{t|t}]^{\circ 2} \qquad (3)$$

where $\circ 2$ is the entry-wise power operator. Note that its entry $\mathbf{D}_{t|a_i}(p, c)$ equals the squared distance between the cth feature components of pixel p in $\mathbf{F}_{t|a_i}$ and $\mathbf{F}_{t|t}$. Ideally, all entries in $\mathbf{D}_{t|a_i}$ should be near zero, because both $\mathbf{F}_{t|a_i}$ and $\mathbf{F}_{t|t}$ represent the same frame I_t. However, they are not in practice, since the transition matrix $\mathbf{A}_{a_i \to t}$ in (2) is imperfect. For the same reason, the transferred object feature $\mathbf{E}_{t|a_i}$ in (2) may be unreliable. Hence, we define the reliability map $R_{t|a_i}$ for $\mathbf{E}_{t|a_i}$ as

$$R_{t|a_i}(p) = \frac{1}{\max_c \mathbf{D}_{t|a_i}(p, c) + \epsilon} \quad \text{for each } p, \quad (4)$$

where ϵ is a small positive number to prevent division by zero. A large value of $R_{t|a_i}(p)$ indicates that the pth row vector (*i.e.* feature vector for pixel p) in $\mathbf{E}_{t|a_i}$ is reliable.

By applying the softmax function over the N reliability maps, we generate the R-attention map $M_{t|a_i}$ for the transferred object feature $\mathbf{E}_{t|a_i}$, which is given by

$$M_{t|a_i}(p) = \frac{\exp R_{t|a_i}(p)}{\sum_{k=1}^{N} \exp R_{t|a_k}(p)} \quad \text{for each } p. \quad (5)$$
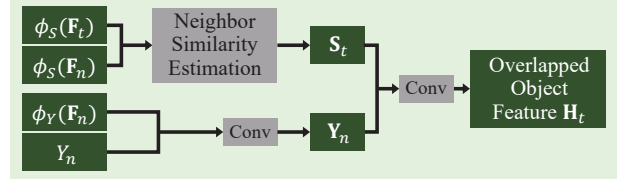

Figure 5: A diagram of the intersection-aware propagation.

Next, we obtain the interfused object feature $\mathbf{G}_t$ by fusing all transferred object features using the R-attention maps,

$$\mathbf{G}_t = \sum_{i=1}^{N} M_{t|a_i} \otimes \mathbf{E}_{t|a_i} \qquad (6)$$

where $\otimes$ means that $M_{t|a_i}$ is multiplied entry-wise to each column in $\mathbf{E}_{t|a_i}$. Through this R-attention mechanism, the interfused feature $\mathbf{G}_t$ contains more reliable information about the query object than each individual $\mathbf{E}_{t|a_i}$ does.

Furthermore, by aggregating $R_{t|a_1}, \ldots, R_{t|a_N}$, we generate the overall reliability map R_t by

$$R_t(p) = \max_i \exp \left(R_{t|a_i}(p) - \tfrac{1}{\epsilon} \right). \qquad (7)$$

Note from (4) and (7) that each $R_t(p)$ is in the range $[0, 1]$, with 1 indicating the maximum reliability level. A high $R_t(p)$ means that the interfused feature vector for pixel p in $\mathbf{G}_t$ is reliable. Because $\mathbf{G}_t$ plays an essential role in segmenting the target frame I_t, $R_t(p)$ also represents the reliability of the segmentation result. Section **??** describes how to use R_t to guide the interactive VOS process.

3.2. Overlapped Object Feature

When segmenting the target frame I_t, the segmentation result Y_n of the neighbor frame $I_n \in \{I_{t-1}, I_{t+1}\}$ is available according to the segmentation direction. We exploit this neighbor information to delineate the query object in I_t more accurately. As shown in Figure 5, we first obtain the neighbor similarity

$$\mathbf{S}_t = \exp \left(- [\phi_S(\mathbf{F}_t) - \phi_S(\mathbf{F}_n)]^{\circ 2} \right) \qquad (8)$$

using a feature transform ϕ_S. $\mathbf{S}_t$ represents how similar the features of I_t and I_n are to each other. A row vector in $\mathbf{S}_t$ tends to have entries near one, when the corresponding pixel belongs to the intersection of the same object between the adjacent frames, as illustrated in Figure 6.

Next, we convert $\mathbf{F}_n$ via another feature transform ϕ_Y and concatenate it with Y_n. We convolve this concatenated signal to yield $\mathbf{Y}_n$ that represents the query object in I_n. We then concatenate $\mathbf{S}_t$ and $\mathbf{Y}_n$ and use another convolution layer to obtain the overlapped object feature $\mathbf{H}_t$. Note that $\mathbf{S}_t$ indicates the overlapped region (or intersection) of the query object in I_n and I_t. Therefore, we combine $\mathbf{S}_t$ and
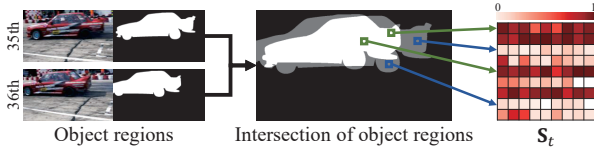
Figure 6: An example of the neighbor similarity $\mathbf{S}_t$. The green pixels belong to the intersection, whereas the blue ones do not.

Figure 7: Illustration of the interactive VOS process.

$\mathbf{Y}_n$ to recognize the query object feature in the overlapped region, and hand over this information to the target frame I_t via $\mathbf{H}_t$. This intersection-aware propagation of object information enables the segment head to exploit the neighbor information selectively and reliably.

Last, we use the frame feature $\mathbf{F}_t$, the interfused object feature $\mathbf{G}_t$, and the overlapped object feature $\mathbf{H}_t$ to obtain the segmentation mask Y_t.

3.3. Guided Interactive VOS Process

In the first round, to segment the target frame $I_t = I_{a_1}$, we set $I_n = I_{a_1}$ and substitute Y_n with the saliency map of I_t generated by the sparse-to-dense network. This is because there is no neighbor frame already segmented. After performing the segmentation of I_{a_1}, we propagate the segmentation mask Y_{a_1} bidirectionally to segment the other frames. In the second round, sparse annotations are given on I_{a_2} to refine its segmentation result Y_{a_2}, which is also propagated bidirectionally until another annotated frame is met. This is repeated until the user is satisfied with the VOS result. Figure 7 illustrates this process.

Suppose that there are K query objects with their annotations. Then, for each query object in each target frame I_t, the segmentation head uses the three features $\mathbf{F}_t$, $\mathbf{G}_t$, $\mathbf{H}_t$ to estimate the object probability map. Consequently, we have $\hat{Y}_{t,1}, \ldots, \hat{Y}_{t,K}$, where $\hat{Y}_{t,k}$ denotes the probability map for the kth query object. By applying the soft aggregation scheme [23] to these maps and then allocating each pixel to the background or the query object with the highest probability, we yield the binary segmentation masks $Y_{t,1}, \ldots, Y_{t,K}$ at the target frame.

In interactive VOS, it is important to enable the user to provide annotations quickly with less effort. Therefore, after performing the segmentation in each round, we compute the R-score r_t of each frame I_t,

$$r_t = \frac{\alpha}{\|\mathcal{U}_t\|} \sum_{p \in \mathcal{U}} R_t(p) + \frac{1-\alpha}{\|\mathcal{O}_t\|} \sum_{p \in \mathcal{O}} R_t(p) \qquad (9)$$

where $\mathcal{U}_t$ is the set of pixels in the entire frame and $\mathcal{O}_t$ is the union set of pixels in the segmented object regions. When $\alpha = 0$, the pixel-wise reliability in (7) is averaged over the foreground segments only. When $\alpha = 1$, it is averaged over t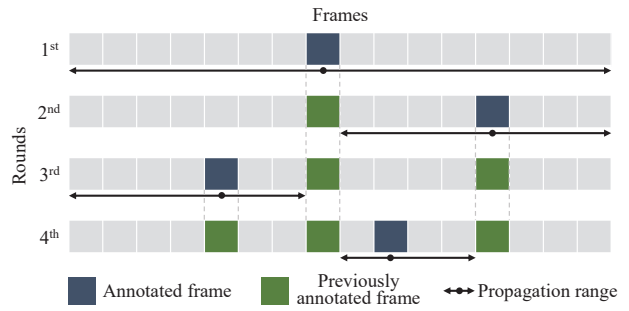he entire frame. In this work, α is set to 0.5 to consider all pixels but also to emphasize the foreground segments.

Guided selection of annotated frames: In practice, it takes considerable time to find the most poorly segmented frame and provide annotations. To alleviate this problem, from the second round, the R-scores $\{r_1, ..., r_T\}$ are used to guide the user to provide additional annotations for the next round. Instead of a time-consuming search over the entire video, the user can select the frames for annotations in two ways.

- **RS1:** The single frame with the lowest R-score is chosen for next annotations.

- **RS4:** The four frames with the lowest R-scores are determined subject to the constraint that their time distances are at least $T/10$. In the interactive VOS simulation, the most poorly segmented frame is selected among the four frames, by comparing the segmentation results with the ground-truth. In real applications, users are provided with the segmentation results of these four guided frames only. Then, the user chooses a frame among them and provides annotations. The interactive process terminates, when the user is satisfied with the guided frames.

3.4. Implementation Details

Network details: To encode each frame I_t to the frame feature $\mathbf{F}_t$, we employ SE-ResNet50 [8] from the first layer to R4 with an output stride 8. Thus, H and W in Section 3.1 are $\frac{1}{8}$ of the height and width of an input frame, respectively. The dimension C_1 of each frame feature vector is 1,024, while the dimensions C_2 of output vectors of the four feature transforms ϕ_A, ϕ_R, ϕ_S, and ϕ_Y are equally set to 128. Also, C_3 for $\mathbf{E}_t$, $\mathbf{G}_t$, and $\mathbf{H}_t$ is set to 256. For the segmentation head, we adopt the decoder architecture in [7].

Training: We use the training sets of DAVIS2017 [26] and YouTube-VOS [38]. To emulate the first round, we randomly form a mini-sequence by taking five consecutive frames (one annotated frame and four target frames) from a video sequence. To emulate the second round, we pick one additional frame as the second annotated frame. In the training, we proceed up to the second round due to limited

Table 1: Comparative assessment of the proposed algorithm with the state-of-the-art interactive VOS algorithms on the DAVIS2017 validation set. The best results are boldfaced.

	AUC-J	J@60s	AUC-J&F	J&F@60s
Oh *et al.* [22]	0.691	0.734	0.778	0.787
Miao *et al.* [20]	0.749	0.761	0.787	0.795
Heo *et al.* [7]	0.771	0.790	0.809	0.827
Oh *et al.* [24]	-	-	0.839	0.848
Proposed-GT	0.817	0.826	0.853	0.863
Proposed-RS1	0.818	0.827	0.855	0.864
Proposed-RS4	**0.820**	**0.829**	**0.856**	**0.866**

GPU memories. To imitate sparse annotations, we use two types: 1) random points and 2) scribble generation in [4]. More implementation details are in the supplementary document.

4. Experimental Results

First, we compare the proposed GIS algorithm with the state-of-the-art interactive VOS algorithms. Second, we analyze the proposed algorithm through various ablation studies and visualization of feature maps. Third, we perform a user study to demonstrate the effectiveness of the proposed algorithm in real applications.

4.1. Comparative Assessment

Interactive VOS simulation is conducted on two datasets: DAVIS2017 [26] and YouTube-IVOS.

DAVIS2017: In the DAVIS interactive VOS simulation [4], human interactions are emulated up to 8 rounds by an algorithm. In each round, after VOS is performed, the algorithm determines the frame with the poorest performance, by comparing the segmentation results with the ground-truth, and provides additional annotations on it. We follow this procedure for the comparison with conventional algorithms. We also follow the two guided procedures RS1 and RS4 in Section 3.3 to confirm the effectiveness of R-scores. The validation set of 30 video sequences in DAVIS2017 is used for the assessment. For each video sequence, three distinct initial scribbles are provided, which means that the performance is averaged over 90 interactive VOS trials.

We quantify the segmentation performance using the region similarity (J) and the contour accuracy (F). We measure the area under the curve (AUC) of a performance-versus-time graph from 0 to 488 seconds for J score (AUC-J) or for joint J and F scores (AUC-J&F). Also, we measure the performance at 60 seconds for J score (J@60s) or for joint J and F scores (J&F@60s) to assess how accurately the segmentation is carried out within 60 seconds. Table 1 compares the proposed GIS algorithm with the recent state-of-the-art algorithms [7, 20, 22, 24], in which
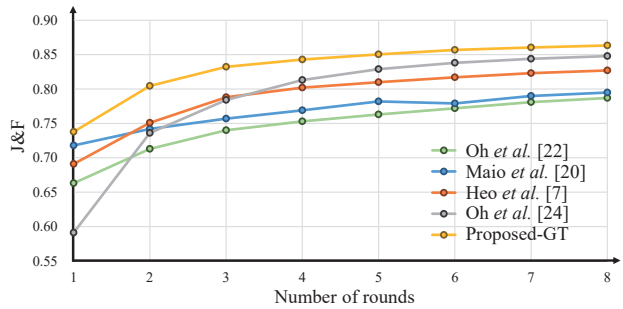

Figure 8: Comparison of J&F scores on the DAVIS2017 validation set according to the rounds.

Proposed-GT, Proposed-RS1, and Proposed-RS4 denote the settings when the ground-truth, RS1, and RS4 are used to choose frames to be annotated in next rounds, respectively. Note that Proposed-GT has the same experimental conditions as the conventional algorithms [7, 20, 22, 24]. In Proposed-RS4, the DAVIS algorithm selects the poorest frame among the four guided frames using the ground-truth. In all settings, the DAVIS algorithm provides annotations. The scores of the conventional algorithms are provided by the respective authors. It can be observed from Table 1 that the proposed algorithm outperforms the state-of-the-arts by significant margins in all metrics. In other words, the proposed algorithm performs the best in both accuracy and speed. Also, Proposed-RS1 and Proposed-RS4 perform better than Proposed-GT, by employing R-scores and selecting annotated frames effectively. Figure 8 shows the J&F scores according to the rounds. The proposed algorithm yields the best score in every round with no exception.

YouTube-IVOS: For extensive experiments, we construct the YouTube-IVOS dataset from YouTube-VOS [38], which is the largest VOS dataset. For the DAVIS algorithm to emulate user interactions, ground-truth segmentation masks are needed. Since the validation set in YouTube-VOS does not provide the ground-truth, we sample 200 videos from its training set to compose YouTube-IVOS. For each video, we generate four different initial annotations by varying the number of point clicks. Specifically, we randomly pick 5, 10, 20, and 50 point clicks from the ground-truth mask for each query object and then use those clicks as annotations in the first round. The interactive VOS is performed up to 4 rounds, since the performance is saturated in early rounds due to the short lengths of the YouTube-VOS videos.

Table 2 compares the average J&F scores of the proposed algorithm with those of Miao *et al.* [20] and Heo *et al.* [7] according to the rounds. Notice that [22] and [24] are not compared in this test, because their full source codes are unavailable. In this test, the proposed GIS network and the Heo *et al.*'s network are trained without the 200 videos in YouTube-IVOS. We see that the proposed algorithm outperforms the other algorithms meaningfully in all rounds.

Table 2: Comparative assessment of the proposed algorithm with the state-of-the-art interactive VOS algorithms on the Youtube-IVOS dataset.

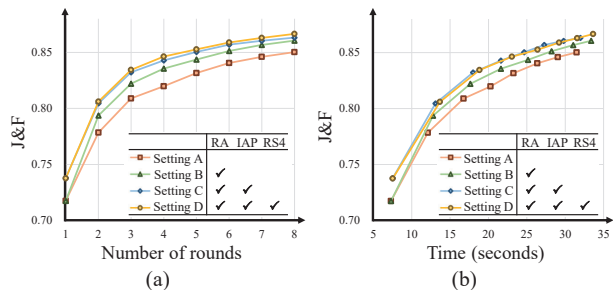|	J&F-1st	J&F-2nd	J&F-3rd	J&F-4th
Miao *et al.* [20]	0.525	0.620	0.674	0.706
Heo *et al.* [7]	0.643	0.721	0.768	0.797
Proposed-GT	**0.672**	**0.754**	**0.806**	**0.830**

Figure 9: Ablation study results on the DAVIS2017 validation set.

4.2. Analysis

Ablation study: We analyze the effectiveness of three components in the proposed algorithm:

(1) R-attention (RA)

(2) Intersection-aware propagation (IAP)

(3) R-score guidance with four candidate frames (RS4)

Figure 9 plots the J&F scores of four settings A, B, C, and D, combining these three components, on the DAVIS2017 validation set. First, using R-attention maps, the overall performance in every round increases significantly (gaps between A and B). Also, the intersection-aware propagation improves the performance especially in early rounds (gaps between B and C). The R-score guidance affects the performance only slightly (gaps between C and D), because in this test the computer searches the frames to be annotated using the ground-truth. In real applications, the R-score guidance enables human users to search the frames efficiently and thus reduces the overall segmentation time, as will be verified in later experiments.

Intersection-aware propagation: We compare the proposed IAP module with two existing propagation methods: the local distance map (LDM) in [20] and the local transfer module (LTM) in [7]. LDM matches each pixel in a target frame to a local region in a neighbor frame at the feature level to estimate a local distance map, while LTM transfers the segmentation result of a neighbor frame based on the local affinity. In Table 3, the baseline means the proposed network without IAP. In other words, the baseline uses only the frame feature $\mathbf{F}_t$ and the interfused object feature $\mathbf{G}_t$

Table 3: Comparison of the proposed intersection-aware propagation (IAP) module with conventional local propagation methods on the DAVIS2017 validation dataset.

	J&F-1st	J&F-3rd	J&F-5th	FPS
Baseline	0.717	0.821	0.839	9.26
Baseline+LDM [20]	0.727	0.824	0.844	8.23
Baseline+LTM [7]	0.730	0.828	0.846	8.69
Baseline+IAP	**0.737**	**0.832**	**0.850**	**8.72**

for the segmentation. We plug LDM or LTM into the baseline. Table 3 compares the J&F scores in 1st, 3rd, and 5th rounds and the segmentation speeds (frames per second, FPS). Compared with LDM and LTM, the proposed IAP helps the baseline network to achieve higher segmentation accuracies and a faster speed.

R-attention maps: Figure 10 illustrates R-attention maps $M_{t|a_i}$ in (5). Specifically, the first and second rows in Figure 10 show the R-attention maps of the first annotation at I_{a_1} and the second annotation at I_{a_2}, respectively. We sample time instances $t_1 \sim t_5$ uniformly between a_1 and a_2. Note that R-attention values depend on the pixel-wise feature similarities between annotated and target frames. For instance, $M_{t_4|a_1}$ has low R-attention values on the query objects (cart and two people), which have significantly different appearance and sizes between I_{a_1} and I_{t_4}. Also, $M_{t_3|a_2}$ has low R-attention values around the plywood ramp that almost disappears in I_{a_2}. This example indicates that R-attention module faithfully provides the reliability information of the transferred object feature from the annotated frame to the target frame.

Reliability maps and R-scores: Examples of reliability maps and R-scores are in Figure 11. In the first round, a poor segmentation result for a woman marked in yellow is obtained, because she is partly occluded by a man riding a bike. Thus, the proposed algorithm yields the reliability map that has low values within the black box containing the woman. As the segmentation result is refined in subsequent rounds, the reliability map has higher values and the R-score gets larger. This means that both the reliability map R_t in (4) and the R-score r_t in (9) are good indicators of how accurately the target frame I_t is segmented.

4.3. User Study

We conducted a user study to assess the proposed GIS algorithm in real applications. We recruited 12 volunteers, who provided scribbles iteratively until they were satisfied. We measured the average time in seconds per video (SPV), including the time for providing scribbles, the running time of the algorithm, and the time for finding unsatisfactory frames. Also, we measured the average rounds per video (RPV) and the average J&F score over all video sequences.
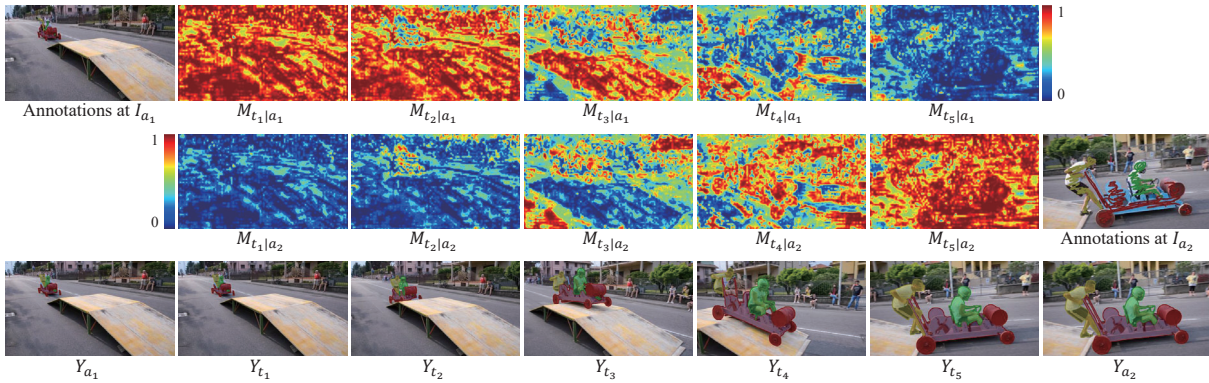
Figure 10: Visualization of R-attention maps and segmentation results on the "soapbox" sequence in the second round.
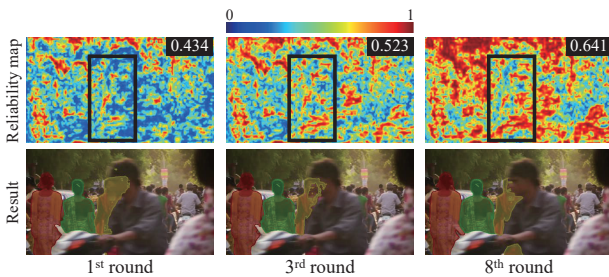

Figure 11: The reliability maps and R-scores for the 33rd frame of the "India" video in the first, third, and eighth rounds. The number in the upper right corner of each reliability map is the R-score.

Table 4 compares these user study results on the validation set in DAVIS2017 [26]. 'Proposed w/o RS' denotes the proposed algorithm without using R-scores. All three settings of the proposed algorithm outperform the state-of-the-art algorithm [7] in all metrics. This indicates that the proposed algorithm requires less running time and less interaction, while providing better segmentation results. Proposed-RS1 and Proposed-RS4 require less time to complete the process than 'Proposed w/o RS' with only negligible performance degradation, by removing or reducing the time for selecting unsatisfactory frames as shown in Figure 12. Especially, the total time for segmenting a video is significantly reduced in Proposed-RS1, which needs no time for inspection.

5. Conclusions

We proposed the novel GIS algorithm for video objects based on R-attention and intersection-aware propagation. First, the interfused object feature is extracted by transferring query object information from annotated frames to a target frame using the R-attention module. Second, the overlapped object feature is obtained via the intersection-aware propagation using a neighbor frame. Then, the segmentation is performed using the frame feature, interfused

Table 4: User study results.

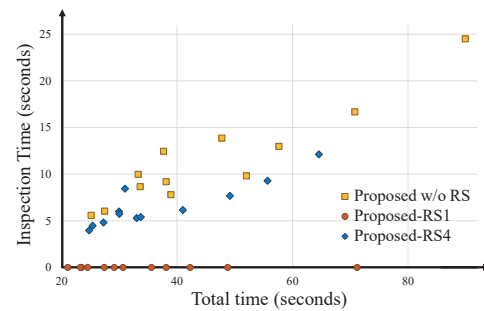|	SPV(s)	RPV	J&F
Heo *et al.* [7]	66.83	2.42	0.769
Proposed w/o RS	46.01	1.89	**0.794**
Proposed-RS1	**34.59**	1.82	0.789
Proposed-RS4	37.10	**1.69**	**0.794**

Figure 12: The total time and the inspection time for selecting unsatisfactory frames per video. Each user is represented by each mark.

object feature, and overlapped object feature. Moreover, we developed the GIS mechanism that enables users to determine next annotated frames quickly. Experimental results showed that the proposed algorithm outperforms the conventional algorithms significantly.

Acknowledgements

References

[1] L. Bao, B. Wu, and W. Liu. CNN in MRF: Video object segmentation via inference in a CNN-based higher-order spatio-temporal MRF. In *CVPR*, 2018.

[2] A. Benard and M. Gygli. Interactive video object segmentation in the wild. 2017.

[3] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool. One-shot video object segmentation. In *CVPR*, 2017.

[4] S. Caelles, A. Montes, K.-K. Maninis, Y. Chen, L. Van Gool, F. Perazzi, and J. Pont-Tuset. The 2018 DAVIS challenge on video object segmentation. In *arXiv:1803.00557*, 2018.

[5] Y. Chen, J. Pont-Tuset, A. Montes, and L. Van Gool. Blazingly fast video object segmentation with pixel-wise metric learning. In *CVPR*, 2018.

[6] B. A. Griffin and J. J. Corso. BubbleNets: Learning to select the guidance frame in video object segmentation by deep sorting frames. In *CVPR*, 2019.

[7] Y. Heo, Y. J. Koh, and C.-S. Kim. Interactive video object segmentation using global and local transfer modules. In *ECCV*, 2020.

[8] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.

[9] P. Hu, G. Wang, X. Kong, J. Kuen, and Y.-P. Tan. Motion-guided cascaded refinement network for video object segmentation. In *CVPR*, 2018.

[10] Y.-T. Hu, J.-B. Huang, and A. G. Schwing. Videomatch: Matching based video object segmentation. In *ECCV*, 2018.

[11] S. D. Jain, B. Xiong, and K. Grauman. Fusionseg: Learning to combine motion and appearance for fully automatic segmentation of generic objects in videos. In *CVPR*, 2017.

[12] W.-D. Jang and C.-S. Kim. Online video object segmentation via convolutional trident network. In *CVPR*, 2017.

[13] W.-D. Jang, C. Lee, and C.-S. Kim. Primary object segmentation in videos via alternate convex optimization of foreground and background distributions. In *CVPR*, 2016.

[14] Y. J. Koh, W.-D. Jang, and C.-S. Kim. POD: Discovering primary objects in videos based on evolutionary refinement of object recurrence, background, and primary object models. In *CVPR*, 2016.

[15] Y. J. Koh and C.-S. Kim. Primary object segmentation in videos based on region augmentation and reduction. In *CVPR*, 2017.

[16] Y. J. Koh, Y.-Y. Lee, and C.-S. Kim. Sequential clique optimization for video object segmentation. In *ECCV*, 2018.

[17] H. Lin, X. Qi, and J. Jia. AGSS-VOS: Attention guided single-shot video object segmentation. In *ICCV*, 2019.

[18] X. Lu, W. Wang, C. Ma, J. Shen, L. Shao, and F. Porikli. See more, know more: Unsupervised video object segmentation with co-attention Siamese networks. In *CVPR*, 2019.

[19] K.-K. Maninis, S. Caelles, Y. Chen, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool. Video object segmentation without temporal information. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(6):1515–1530, 2018.

[20] J. Miao, Y. Wei, and Y. Yang. Memory aggregation networks for efficient interactive video object segmentation. In *CVPR*, 2020.

[21] S. W. Oh, J.-Y. Lee, K. Sunkavalli, and S. J. Kim. Fast video object segmentation by reference-guided mask propagation. In *CVPR*, 2018.

[22] S. W. Oh, J.-Y. Lee, N. Xu, and S. J. Kim. Fast user-guided video object segmentation by interaction-and-propagation networks. In *CVPR*, 2019.

[23] S. W. Oh, J.-Y. Lee, N. Xu, and S. J. Kim. Video object segmentation using space-time memory networks. In *ICCV*, 2019.

[24] S. W. Oh, J.-Y. Lee, N. Xu, and S. J. Kim. Space-time memory networks for video object segmentation with user guidance. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020.

[25] A. Papazoglou and V. Ferrari. Fast object segmentation in unconstrained video. In *ICCV*, 2013.

[26] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool. The 2017 DAVIS Challenge on Video Object Segmentation. In *arXiv:1704.00675*, 2017.

[27] B. L. Price, B. S. Morse, and S. Cohen. LIVEcut: Learning-based interactive video segmentation by evaluation of multiple propagated cues. In *ICCV*, 2009.

[28] N. Shankar Nagaraja, F. R. Schmidt, and T. Brox. Video segmentation with just a few strokes. In *ICCV*, 2015.

[29] G. Strang. Linear transformations. In *Introduction to Linear Algebra, 5th Ed.*, 2016.

[30] P. Tokmakov, K. Alahari, and C. Schmid. Learning motion patterns in videos. In *CVPR*, 2017.

[31] P. Voigtlaender, Y. Chai, F. Schroff, H. Adam, B. Leibe, and L.-C. Chen. FEELVOS: Fast end-to-end embedding learning for video object segmentation. In *CVPR*, 2019.

[32] P. Voigtlaender and B. Leibe. Online adaptation of convolutional neural networks for video object segmentation. In *BMVC*, 2017.

[33] J. Wang, P. Bhat, R. A. Colburn, M. Agrawala, and M. F. Cohen. Interactive video cutout. *ACM Trans. Graphics*, 24(3):585–594, 2005.

[34] W. Wang, J. Shen, and F. Porikli. Saliency-aware geodesic video object segmentation. In *CVPR*, 2015.

[35] W. Wang, H. Song, S. Zhao, J. Shen, S. Zhao, S. C. Hoi, and H. Ling. Learning unsupervised video object segmentation through visual attention. In *CVPR*, 2019.

[36] X. Wang, R. Girshick, A. Gupta, and K. He. Non-local neural networks. In *CVPR*, 2018.

[37] R. Wu, H. Lin, X. Qi, and J. Jia. Memory selection network for video propagation. In *ECCV*, 2020.

[38] N. Xu, L. Yang, Y. Fan, D. Yue, Y. Liang, J. Yang, and T. Huang. YouTube-VOS: A large-scale video object segmentation benchmark. In *arXiv:1809.03327*, 2018.

[39] Z. Yang, Q. Wang, L. Bertinetto, W. Hu, S. Bai, and P. H. S. Torr. Anchor diffusion for unsupervised video object segmentation. In *ICCV*, 2019.

[40] M. Zhen, S. Li, L. Zhou, J. Shang, H. Feng, T. Fang, and L. Quan. Learning discriminative feature with crf for unsupervised video object segmentation. *ECCV*, 2020.

[41] T. Zhou, J. Li, S. Wang, R. Tao, and J. Shen. MATNet: Motion-attentive transition network for zero-shot video object segmentation. *IEEE Trans. Image Process.*, 29:8326–8338, 2020.