



HAL
open science

Unsupervised learning-based long-term superpixel tracking

Pierre-Henri Conze, Florian Tilquin, Mathieu Lamard, Fabrice Heitz, Gwéno­lé Quellec

► **To cite this version:**

Pierre-Henri Conze, Florian Tilquin, Mathieu Lamard, Fabrice Heitz, Gwéno­lé Quellec. Unsupervised learning-based long-term superpixel tracking. *Image and Vision Computing*, 2019, 89, pp.289-301. 10.1016/j.imavis.2019.06.011 . hal-02189986

HAL Id: hal-02189986

<https://imt-atlantique.hal.science/hal-02189986>

Submitted on 20 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Unsupervised learning-based long-term superpixel tracking

Pierre-Henri Conze^{a,b,*}, Florian Tilquin^c, Mathieu Lamard^{b,d}, Fabrice Heitz^c,
Gwenolé Quéllec^b

^a*IMT Atlantique, Technopôle Brest-Iroise, 29238 Brest, France*

^b*LaTIM UMR 1101, Inserm, 22 av Camille Desmoulins, 29238 Brest, France*

^c*ICube UMR 7357, CNRS, Unistra, 300 bd Sébastien Brant, 67412 Illkirch, France*

^d*Université de Bretagne Occidentale, 2 av Foch, 29609 Brest, France*

Abstract

Finding correspondences between structural entities decomposing images is of high interest for computer vision applications. In particular, we analyze how to accurately track superpixels - visual primitives generated by aggregating adjacent pixels sharing similar characteristics - over extended time periods relying on unsupervised learning and temporal integration. A two-step video processing pipeline dedicated to long-term superpixel tracking is proposed. First, unsupervised learning-based superpixel matching provides correspondences between consecutive and distant frames using new context-rich features extended from greyscale to multi-channel and forward-backward consistency constraints. Resulting elementary matches are then combined along multi-step paths running through the whole sequence with various inter-frame distances. This produces a large set of candidate long-term superpixel pairings upon which majority voting is performed. Video object tracking experiments demonstrate the accuracy of our elementary estimator against state-of-the-art methods and proves the ability of multi-step integration to provide accurate long-term superpixel matches compared to usual direct and sequential integration.

Keywords: superpixel matching, unsupervised learning, superpixel tracking, multi-step integration, random forests, forward-backward consistency

*corresponding author, pierre-henri.conze@imt-atlantique.fr

1. Introduction

Finding correspondences between multiple images is a fundamental problem in computer vision tasks including scene segmentation [1], 3D reconstruction [2], visual tracking [3], trajectory analysis [4] or video editing like 2D-to-3D video conversion [5] and graphic elements propagation [6]. Established via local or global search, correspondences are usually either **sparse through key feature points** [7] or for the dense pixel grid as with optical flow [8]. Patch-based approximate nearest neighbor (ANN) search methods such as PatchMatch (PM) [9] and its extension to multi-resolution [10] are mainly followed to find correspondences between patches. Alternatively, finding associations between structural entities decomposing images by grouping pixels enables a semi-dense coverage of **whole images** while drastically reducing the cost of correspondence. Contrary to a regular decomposition of the image grid into patches which would not offer enough consistent support regions for matching, superpixels - visual primitives generated by aggregating adjacent pixels sharing similar characteristics into **visually meaningful entities** [11] - provide more reliable areas **preserving both** image geometry and object contours. Moreover, the hypothesis that motion discontinuities are a subset of photometric contours is usually used to preserve boundaries between objects exhibiting different motion. **In particular,** these findings have motivated recent optical flow algorithms using image data and smoothness terms adapted to the superpixel level [12, 13]. Conversely, we claim that image matching relying on superpixels could benefit from these advantages to offer more consistent associations than pixel or patch matches while providing a better management of motion discontinuities. **More precisely,** this paper focuses on how to accurately find correspondences between superpixels over extended time periods.

In this context, superpixel correspondences have been already employed for visual tracking through superpixel-based discriminative appearance models [3] or object-background confidence maps [14]. However, these works perform superpixel matching based on comparisons of intrinsic superpixel features only,

without taking full advantage of neighborhood information. Conversely, [15] exploits a structure of superpixel neighborhood called SuperPatch involved in a superpixel PM framework. Superpixel neighborhood information greatly improves correspondences since it alleviates some matching failures due to irregular decomposition of the same image content, not directly comparable between
35 images. However, even with incorporated neighborhood information, directly computing a matching distance between irregular structures can be tedious, especially when images are divided into a large collection of superpixels.

A prior pixel-to-superpixel mapping can drive the matching at the superpixel
40 level to provide more precise correspondences. In this direction, [16, 17] uses random forests (RF) [18] to establish supervoxel correspondences between two 3D images in an unsupervised fashion. RF is trained on one image by using supervoxel indexes as voxel-wise class labels and robust context-rich features to describe the extended neighborhood. Applying RF on the other image yields a
45 voxel labelling which is then regularised using majority voting within supervoxel boundaries. First validated on medical image registration, we explore the use of such learning-based superpixel matching for accurate superpixel matching across long video sequences.

Despite recent advances related to optical flow integration [19, 6], the tem-
50 poral tracking of superpixels over long-term video sequences has received little attention in the literature. [20] uses a constrained graph where nodes denote superpixels and edges encode spatial, temporal, and appearance constraints. However, temporal constraints only model short-term smoothness between consecutive frames. **Generative probabilistic model [21] and proximity-weighted
55 patch matching [22] employed to generate temporal superpixels only exploit short-term correspondences from one frame to the subsequent one.** The same finding arises in [3] whose tracker is conducted sequentially and therefore prone to motion drift. Establishing long-term superpixel correspondences requires to perform superpixel matching between consecutive and distant frames and
60 therefore to handle simultaneously small and large displacements. To address this challenge, we exploit the concept of multi-step integration introduced for

long-term motion estimation using optical flow [6]. The idea is to generate a large set of elementary displacement estimations performed between consecutive frames or with larger inter-frame distances. Once combined, elementary
65 multi-step estimations result in a large set of long-term correspondences which are significant enough to be fused through statistical processing. We are not aware of any studies that have recovered this concept for long-term superpixel tracking while it could bring many benefits in **this** context. Indeed, it can alleviate matching errors during superpixel trajectory estimation since new steps
70 can give a chance to match with a correct location again compared to sequential processing whose tracks may be lost. Moreover, statistical processing upon large representative long-term superpixel candidates can solve the uncertainty component present for matching tasks.

It must be reported that deep learning has become popular for object track-
75 ing relying on convolutional networks to learn discriminative features to encode the target appearance [23, 24] or recurrent networks trained with reinforcement learning to learn how to predict object locations across videos [25]. Despite their high performance, these methods only provide very sparse bounding box tracking and do not describe how the boundaries of an irregular shaped object
80 evolves in time as expected through long-term superpixel tracking.

In summary, two main contributions are proposed towards accurate long-term superpixel tracking. First, unsupervised learning-based superpixel matching is generalized and adapted from medical image processing [16, 17] to computer vision in order to find associations along **video** sequences between consecutive and distant images decomposed into **superpixels** (Sect.2). The approach is carried out using classifiers such as k -nearest neighbors (k NN) or RF [18], incorporates new forward-backward consistency constraints and fully exploits dedicated context-rich features we extended from greyscale [26, 16, 17] to multi-channel to incorporate neighborhood information on RGB frames. Second, based on this learning-based matching approach used as an elementary
90 displacement estimator, we propose a multi-step integration strategy for long-term superpixel tracking (**Sect.3**). It combines multiple elementary superpixel

matches obtained for some intermediate images following randomly selected multi-step paths. This produces a large set of candidate long-term superpixel pairings upon which a majority voting selection is performed. Based on object tracking experiments **extending a preliminary study presented in [27]**, Sect.4 assesses the accuracy of the proposed elementary estimator against state-of-the-art methods and proves the ability of multi-step integration to provide an efficient long-term superpixel tracking compared to **both** standard direct and sequential integration. We end with conclusions and perspectives in Sect.5.

2. Superpixel matching with unsupervised learning

Let \mathcal{V} be a video sequence of RGB images. **In this section**, unsupervised learning-based superpixel matching is addressed between two consecutive or distant images I_f and I_s of \mathcal{V} . Each image $I_q : \Omega_q \subset \mathbb{N}^2 \rightarrow \mathbb{N}^3$ associates a RGB color vector $I_q(\mathbf{x}_q)$ to each pixel \mathbf{p}_q located at $\mathbf{x}_q \in \Omega_q$ with $q \in \{f, s\}$ ¹.

2.1. Problem formulation

Let $\mathcal{F} = \{\mathbf{f}_i\}_{i \in \{1, \dots, |\mathcal{F}|\}}$ and $\mathcal{S} = \{\mathbf{s}_j\}_{j \in \{1, \dots, |\mathcal{S}|\}}$ be respectively the set of $|\mathcal{F}|$ and $|\mathcal{S}|$ connected superpixels partitioning I_f and I_s . The superpixel decomposition **can be performed using any superpixel algorithm [11] such as Simple Linear Iterative Clustering (SLIC) [28]** which aggregates neighboring pixels \mathbf{p}_q based on spatial and intensity proximity criteria. Forward superpixel matching from I_f to I_s ($f < s$) consists in automatically learning a matching function $h_{f,s}$ that maps each superpixel $\mathbf{f}_i \in \mathcal{F}$ of I_f to a given superpixel $\mathbf{s}_j \in \mathcal{S}$ of I_s [16] such that:

$$\forall i \in \{1, \dots, |\mathcal{F}|\}, \exists j \in \{1, \dots, |\mathcal{S}|\} \mid h_{f,s}(\mathbf{f}_i) = \mathbf{s}_j \quad (1)$$

Backward matching from I_s to I_f can be similarly considered by estimating $h_{s,f}$ mapping each superpixel $\mathbf{s}_j \in \mathcal{S}$ to a given superpixel $\mathbf{f}_i \in \mathcal{F}$. In what follows, learning-based superpixel matching is described in forward from I_f to I_s .

¹stands for *first* and *second*

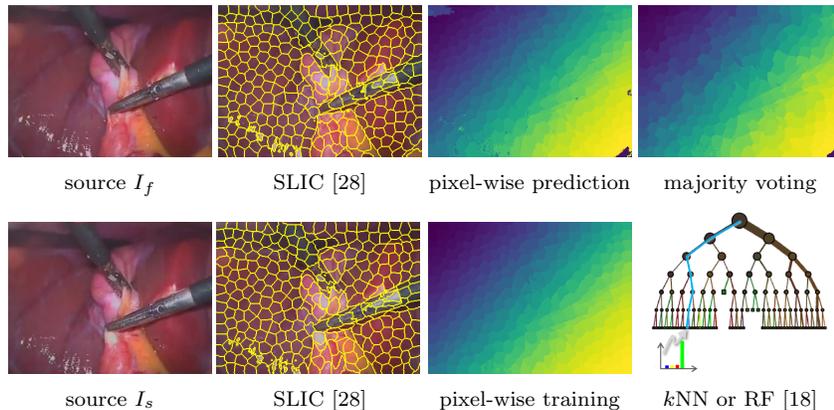


Figure 1: Superpixel matching between I_f and I_s using unsupervised learning applied with SLIC [28] superpixel indexes as label entities followed by majority voting following [16]. This example is produced for image pair $\{I_{50}, I_{70}\}$ of the `lapa` sequence [29] using RF [18].

2.2. Overall strategy

Instead of relying on nearest neighbor search at the superpixel level through
 120 superpixel feature comparisons [14, 15], which is prone to ambiguity due to possible severe overlaps in feature space, we explore the use of pixel-wise k -nearest neighbors (k NN) or random forests (RF) [18] to establish correspondences between superpixels over-segmenting I_f and I_s , as formulated in Eq.1. Usually employed with success for multi-class classification or regression, we show that
 125 such classifiers can also be used with profit for accurate superpixel correspondences. The overall learning-based superpixel matching strategy illustrated Fig.1 is carried out in an unsupervised manner, *i.e.* **without any previously collected training data**. To give a powerful representation of the global context, RF or k NN is considered with new pixel-wise context-rich features extended
 130 from greyscale [26, 16, 17] to multi-channel RGB and described in Sect.2.3.

The key idea is to perform training on the target image (I_s) by using superpixel indexes as pixel-wise class labels and testing on the **first** image (I_f) to get a pixel-to-superpixel mapping, as introduced in [16]. In particular, the classifier aims at assigning a superpixel $\mathbf{s}_j \in \mathcal{S}$ to each pixel $\mathbf{p}_f \in \Omega_f$. A training set
 135 is thus built by considering all pixels \mathbf{p}_s of Ω_s with their associated superpixel

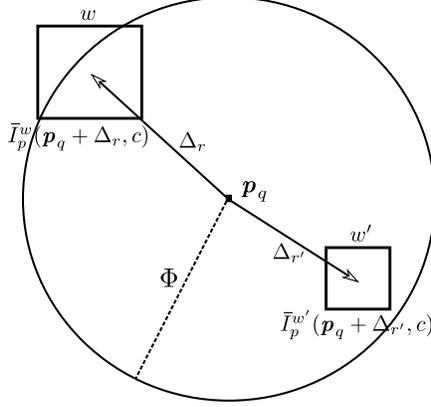


Figure 2: Pixel-wise context-rich multi-channel features provide a description of the extended spatial context (see Sect.2.3 for further details).

index, i.e. the index of the superpixel s_j they belong to. Once trained, the classifier is applied to I_f to predict for each \mathbf{p}_f the index of a superpixel of \mathcal{S} . This pixel-to-superpixel mapping is addressed in detail in Sect.2.4.

Mapping results are further regularized following superpixel boundaries to reach robust superpixel matches. Within each superpixel \mathbf{f}_i of I_f , the most represented superpixel index among all pixel-wise predictions indicates the best superpixel match $h_{f,s}(\mathbf{f}_i)$. This final superpixel matching step is detailed in Sect.2.5 with new forward-backward (FW-BW) consistency constraints.

2.3. Context-rich multi-channel features

Extracting contextual information from near size-variable patches is key to discriminate pixels towards robust pixel-to-superpixel mapping. Let $\bar{I}_q^w(\mathbf{p}_q, c)$ be the average intensity on a local box of size w centered on \mathbf{p}_q for channel $c \in \{r, g, b\}$. Pixel-wise context-rich features $\theta(\mathbf{p}_q) = \{\theta_m(\mathbf{p}_q)\}_{m \in \{0, \dots, K_a - 1\}}$ assigned to pixels \mathbf{p}_q are extended from greyscale [26, 16, 17] to multi-channel as follows:

$$\theta_m(\mathbf{p}_q) = \bar{I}_p^w(\mathbf{p}_q + \Delta_r, c) - \beta \times \bar{I}_p^{w'}(\mathbf{p}_q + \Delta_{r'}, c) \quad (2)$$

where displacements $\Delta_{\{r, r'\}}$ are randomly defined starting from \mathbf{p}_q in a disc of maximal radius Φ (Fig.2). $\beta \in \{0, 1\}$ is a binary parameter which focuses

whether on intensity differences between two boxes randomly located in the extended neighborhood ($\beta = 1$) or on the value obtained from one single box only ($\beta = 0$). Color intensities around \mathbf{p}_q are included in the feature vector $\theta(\mathbf{p}_q)$ by forcing $\Delta_r = \beta = 0$ for all possible pre-defined box sizes w and channels c .

By randomly generating many different box sizes w and offsets Δ_r , we obtain a large set of K_a features describing the extended spatial context for all color channels c . Parameters $\{w, w', \Delta_r, \Delta'_r, \beta, c\}$ are randomly generated once and remain similar whatever the image I_f or I_s under consideration.

2.4. Pixel-to-superpixel mapping

Pixel-to-superpixel mapping relies on machine learning to compute pixel-to-superpixel mapping probabilities denoted as $\mathbf{p}(h_{f,s}(\mathbf{p}_f) = \mathbf{s}_n)$ for each pixel $\mathbf{p}_f \in \Omega_f$ with respect to all superpixels $\mathbf{s}_n \in \mathcal{S}$ with $n \in \{1, \dots, |\mathcal{S}|\}$.

The procedure with random forests (RF) [18] is conducted as follows. The forest is formed by T uncorrelated trees made of both internal nodes splitting data according to binary tests Ψ and leaf nodes which reach all together a final data partition. At each internal node, the split sends pixels \mathbf{p}_q to left and right child nodes during training ($q = f$) and prediction ($q = s$). The associated binary test Ψ focuses on a random subset $\hat{\theta}(\mathbf{p}_q)$ of context-rich multi-channel visual features $\theta(\mathbf{p}_q)$ assigned to \mathbf{p}_q (Sect.2.3) and divides the input pixel set based on the following split rule:

$$\Psi(\mathbf{p}_q, \theta(\mathbf{p}_q)) = \begin{cases} \text{true, if } \hat{\theta}(\mathbf{p}_q) > \tau \\ \text{false, otherwise} \end{cases} \quad (3)$$

where $\hat{\theta}(\mathbf{p}_q)$ is compared to a threshold τ and $q \in \{f, s\}$.

Internal node parameters ($\{\tau, \hat{\theta}(\mathbf{p}_q)\}$) are optimized *via* information gain maximization with respect to the training dataset $\mathcal{L} = \{\mathbf{p}_s, c(\mathbf{s}_j)\}$ combining pixels \mathbf{p}_s belonging to \mathbf{s}_j with their associated superpixel index $c(\mathbf{s}_j) = j$ with $j \in \{1, \dots, |\mathcal{S}|\}$. **Any additional training data is required additionally to \mathcal{L} which makes this strategy unsupervised.** After optimization, each leaf node l_t of the t^{th}

180 tree receives a partition \mathcal{L}_{l_t} of \mathcal{L} and produces the class probability distribution $\mathbf{p}_{l_t}(c(\mathbf{s}_j)|\mathcal{L})$ for all superpixels \mathbf{s}_j .

To predict the corresponding superpixel index $c(\mathbf{s}_n)$ of a given pixel $\mathbf{p}_f \in \Omega_f$ with associated visual features $\theta(\mathbf{p}_f)$ during testing, \mathbf{p}_f is injected into each optimized tree and finally reaches a leaf node l_t per tree following the successive
 185 split rules (Eq.3). The pixel-to-superpixel mapping probability $\mathbf{p}(h_{f,s}(\mathbf{p}_f) = \mathbf{s}_n)$ denoting the probability that \mathbf{s}_n is assigned to \mathbf{p}_f is obtained for each \mathbf{s}_n by:

$$\mathbf{p}(h_{f,s}(\mathbf{p}_f) = \mathbf{s}_n) = \frac{1}{T} \sum_{t=1}^T \mathbf{p}_{l_t}(c(\mathbf{s}_n)|\mathcal{L}) = \frac{1}{T} \sum_{t=1}^T \frac{|\{\mathbf{p}_s, c(\mathbf{s}_l)\} \in \mathcal{L}_{l_t} \mid l = n|}{|\mathcal{L}_{l_t}|} \quad (4)$$

Contrary to RF, the k NN classifier simply stores instances of training data
 190 **instead** of building a general internal model. Pixel-to-superpixel mapping probabilities $\mathbf{p}(h_{f,s}(\mathbf{p}_f) = \mathbf{s}_n)$ are computed by looking at the class (superpixel index) distribution among the k nearest neighbors \mathbf{p}_s of each pixel \mathbf{p}_f in feature space. Nearest neighbors are estimated using Euclidean distance on context-rich **fea-**
tures assigned to pixels \mathbf{p}_q with $q \in \{f, s\}$.

195 2.5. Superpixel-to-superpixel matching

Once pixel-to-superpixel mapping probabilities $\mathbf{p}(h_{f,s}(\mathbf{p}_f) = \mathbf{s}_n)$ are computed for each superpixel $\mathbf{s}_n \in \mathcal{S}$ using context-rich features (Sect.2.3) involved in RF or k NN (Sect.2.4), two steps are required to get final superpixel pairings. First, the final pixel-to-superpixel mapping for each \mathbf{p}_f of I_f can be found using:

200

$$h_{f,s}(\mathbf{p}_f) = \mathbf{s}_j = \arg \max_{\mathbf{s}_n \in \mathcal{S}} \mathbf{p}(h_{f,s}(\mathbf{p}_f) = \mathbf{s}_n) \quad (5)$$

Second, majority voting among all pixels of a given superpixel $\mathbf{f}_i \in \mathcal{F}$ can be performed by selecting the most represented superpixel index. The final
 205 matching $h_{f,s}(\mathbf{f}_i) = \mathbf{s}_j$ is defined such that $c(\mathbf{s}_j)$ satisfies:

$$c(\mathbf{s}_j) = \arg \max \mathbf{hist}(\{c(h_{f,s}(\mathbf{p}_f)) \mid \mathbf{p}_f \in \mathbf{f}_i\}) \quad (6)$$

An alternative consists in averaging the pixel-to-superpixel mapping probabilities at the superpixel level instead of making hard decision for each \mathbf{p}_f as performed in Eq.5:

$$\mathbf{p}(h_{f,s}(\mathbf{f}_i) = \mathbf{s}_n) = \frac{1}{|\mathbf{f}_i|} \sum_{\mathbf{p}_f \in \mathbf{f}_i} \mathbf{p}(h_{f,s}(\mathbf{p}_f) = \mathbf{s}_n) \quad (7)$$

210

We keep at this point all possible outcomes between candidate matches. Decisions are postponed to the superpixel level by finding the superpixel $\mathbf{s}_j \in \mathcal{S}$ which maximizes $\mathbf{p}(h_{f,s}(\mathbf{f}_i) = \mathbf{s}_n)$:

$$h_{f,s}(\mathbf{f}_i) = \mathbf{s}_j = \arg \max_{\mathbf{s}_n \in \mathcal{S}} \mathbf{p}(h_{f,s}(\mathbf{f}_i) = \mathbf{s}_n) \quad (8)$$

215

Forward-backward consistency can be enforced in the context where two mapping functions are learned: $h_{f,s}$ (resp. $h_{s,f}$) that maps each supervoxel $\mathbf{f}_i \in \mathcal{F}$ ($\mathbf{s}_j \in \mathcal{S}$) to a given \mathbf{s}_j (\mathbf{f}_i) belonging to \mathcal{S} (\mathcal{F}) in forward (backward).

220

Thus, we extend Eq.8 with a new consistency constraint that guides the mutual matching between \mathbf{f}_i and \mathbf{s}_j :

$$h_{f,s}(\mathbf{f}_i) = \arg \max_{\mathbf{s}_n \in \mathcal{S}} \mathbf{p}(h_{f,s}(\mathbf{f}_i) = \mathbf{s}_n) \times \mathbf{p}(h_{s,f}(\mathbf{s}_n) = \mathbf{f}_i) \quad (9)$$

225

The whole unsupervised learning-based strategy described above can be performed all along the video \mathcal{V} to match superpixels decomposing consecutive or distant images, both in forward and backward directions.

3. Long-term superpixel tracking using multi-step integration

230

We address at this stage long-term superpixel tracking for sequence \mathcal{V} composed of $N + 1$ RGB frames $I_n : \Omega_n \subset \mathbb{N}^2 \rightarrow \mathbb{N}^3$ using the learning-based superpixel matching strategy, described Sect.2 for a given pair of consecutive or distant frames, as elementary estimator. Each frame I_n is decomposed into a set of superpixels **with the same amount of visual primitives for each**. One particular frame (usually the first one) of \mathcal{V} is defined as the reference frame and denoted I_{ref} . In this context, we aim at finding correspondences between
235 superpixels over-segmenting I_{ref} and superpixels defined in frames I_n with

$n \in \{0, \dots, N\} \neq ref$. Let $\mathcal{F} = \{\mathbf{f}_i\}_{i \in \{1, \dots, |\mathcal{F}|\}}$ and $\mathcal{S} = \{\mathbf{s}_j\}_{j \in \{1, \dots, |\mathcal{S}|\}}$ be the set of $|\mathcal{F}|$ and $|\mathcal{S}|$ connected superpixels partitioning I_{ref} and I_n .

Both superpixel trajectory estimation between the reference frame and all the images of the sequence and superpixel matching to match each image to the reference frame can be considered, as in [19, 6]. From-the-reference estimation is useful for information pushing from superpixels of I_{ref} whereas to-the-reference estimation allows information propagation over superpixels of each frame I_n by pulling it from I_{ref} . The description below focuses on a given pair $\{I_{ref}, I_n\}$ where I_n is located far away from I_{ref} . Correspondences for the whole sequence are obtained by processing each pair $\{I_{ref}, I_n\}$ independently $\forall n \neq ref$.

Starting from learning-based superpixel matching (Sect.2) as elementary motion estimator, two temporal integration schemes can be considered at first glance to find $h_{ref,n}$ mapping each superpixel $\mathbf{f}_i \in \mathcal{F}$ to a given superpixel \mathbf{s}_j of \mathcal{S} such that $h_{ref,n}(\mathbf{f}_i) = \mathbf{s}_j$. First, sequential integration can be employed passing through all intermediate frames, similarly to dense point tracking algorithms [30]. This step-by-step strategy can gradually apprehend appearance changes and large displacements but may lead to large error accumulation, resulting in a substantial drift over extended time periods. This drawback is further enhanced when using superpixels since superpixel decompositions across the sequence may result in an irregular partitioning of the image content. Second, to avoid error accumulations, direct matching [31] can be applied between superpixels of I_{ref} and I_n , exactly as in Sect.2. However, this ignores that \mathcal{V} consists of inter-related images with redundant and smoothly evolving content, which makes large displacement and aspect changes challenging to handle.

Issues related to both sequential and direct superpixel tracking could be partially compensated by complexifying the superpixel matching models and criteria, but an uncertainty component remains. This argues in favor of a statistical processing (Sect.3.2) which takes into account a large set of candidate long-term superpixel matches obtained using multi-step combination (Sect.3.1) of elementary superpixel pairings previously established through unsupervised learning following Sect.2.

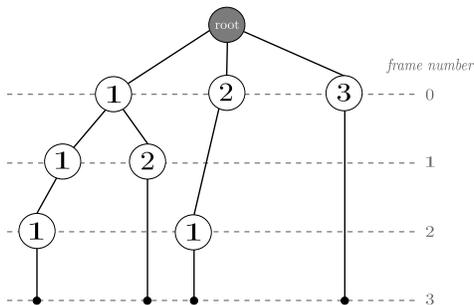


Figure 3: Generation of step sequences from I_0 to I_3 with steps 1, 2, and 3 by creating a tree structure: $\Gamma_{0,3} = \{\{1, 1, 1\}, \{1, 2\}, \{2, 1\}, \{3\}\}$.

3.1. Multi-step integration of elementary matches

Multi-step integration aims at producing a large set of candidate long-term superpixel pairings between I_{ref} and I_n using intermediate superpixel correspondences to form a significant set of samples upon which statistical selection (Sect.3.2) is relevant. Formely introduced for optical flow integration [6], we show that this heuristic can be extended towards accurate long-term superpixel tracking. As inputs, we take a set of superpixel match fields pre-estimated from each frame of \mathcal{V} including I_{ref} . These matches are computed between consecutive frames or with larger inter-frame distances [6] using learning-based superpixel matching (Sect.2). Let $\mathcal{A}_n = \{\alpha_1, \alpha_2, \dots, \alpha_{Q_n}\} \subset \{1, \dots, N-n\}$ be the set of Q_n possible steps at instant n which means that $\{h_{n,n+\alpha_1}, h_{n,n+\alpha_2}, \dots, h_{n,n+\alpha_{Q_n}}\}$ has been previously learned using k NN or RF. Thus, for each step $\alpha_q \in \mathcal{A}_n$, we have a superpixel match in $I_{n+\alpha_q}$ for each superpixel of I_n through the mapping function $h_{n,n+\alpha_q}$, and this for each frame.

The starting point of multi-step integration consists in initially generating all possible step sequences (Fig.3), i.e. combinations of steps, to join I_n from I_{ref} . Then, each generated step sequence defines a multi-step path (Fig.4) linking each superpixel \mathbf{f}_i of I_{ref} to a superpixel \mathbf{s}_j in I_n passing through superpixels of some intermediate frames.

Let $\Gamma_{ref,n} = \{\gamma_0, \gamma_1, \dots, \gamma_{K-1}\}$ be the set of the K possible step sequences

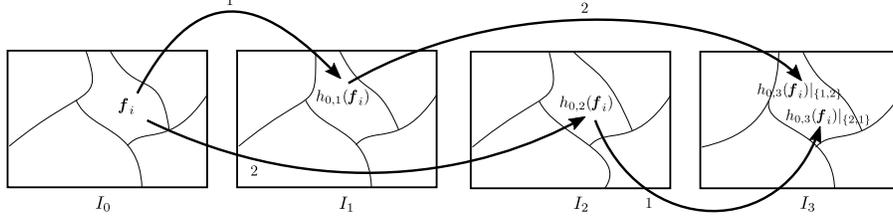


Figure 4: Generation of multi-step paths corresponding to step sequences $\{\{1, 2\}, \{2, 1\}\} \subset \Gamma_{0,3}$ from I_0 to I_3 .

γ_k between I_{ref} and I_n . A step sequence $\gamma_k = \{\alpha_1^k, \alpha_2^k, \dots, \alpha_{K_{\gamma_k}}^k\}$ is defined by a set of K_{γ_k} steps which once cascaded join I_n from I_{ref} . $\Gamma_{ref,n}$ is computed
 290 by building a tree structure (Fig.3) where each node corresponds to a field of superpixel matches assigned to a given frame for a given step value (node value). Going from the root node to leaf nodes of this tree structure gives the possible step sequences which are stacked into $\Gamma_{ref,n}$. For instance, the tree displayed
 in Fig.3 indicates the 4 possible step sequences from I_0 to I_3 with steps 1, 2,
 295 and 3: $\Gamma_{0,3} = \{\{1, 1, 1\}, \{1, 2\}, \{2, 1\}, \{3\}\}$.

Once all the K possible step sequences γ_k between I_{ref} and I_n are generated, the corresponding multi-step paths are constructed (Fig.4). For step sequence $\gamma_k = \{\alpha_1^k, \alpha_2^k, \dots, \alpha_{K_{\gamma_k}}^k\} \in \Gamma_{ref,n}$ composed of K_{γ_k} steps, superpixel matching between I_{ref} and I_n is performed via:

$$h_{ref,n}(\mathbf{f}_i)|_{\gamma_k} = h_{ref + \sum_{p=1}^{K_{\gamma_k}-1} \alpha_p^k, n} \circ \dots \circ h_{ref + \alpha_1^k, ref + \alpha_1^k + \alpha_2^k} \circ h_{ref, ref + \alpha_1^k}(\mathbf{f}_i) \quad (10)$$

300 with $ref + \sum_{p=1}^{K_{\gamma_k}} \alpha_p^k = n$. Once all the steps $\alpha_j^k \in \gamma_k$ have been run through, one gets $h_{ref,n}(\mathbf{f}_i)|_{\gamma_k}$, the superpixel in I_n corresponding to $\mathbf{f}_i \in I_{ref}$ obtained with step sequence γ_k . For $\gamma_k = \{1, 2\} \in \Gamma_{0,3}$ for instance (Fig.4), we have:

$$h_{0,3}(\mathbf{f}_i)|_{\{1,2\}} = h_{1,3}(h_{0,1}(\mathbf{f}_i)) = h_{1,3} \circ h_{0,1}(\mathbf{f}_i) \quad (11)$$

A large set of candidate superpixels in I_n is finally reached by considering all the step sequences of $\Gamma_{ref,n}$ and this for each superpixel \mathbf{f}_i defined in I_{ref} .

305 Thus, to each \mathbf{f}_i is associated a large set of K candidate superpixels in I_n defined as $\mathcal{T}_{ref,n}(\mathbf{f}_i) = \{h_{ref,n}(\mathbf{f}_i)|_{\gamma_k}\}_{k \in \{0,1,\dots,K\}}$.

Multi-step integration has been previously presented as an exhaustive candidate generation process. In practice, selecting only a subset of all possible step sequences and therefore associated multi-step paths is required to be able
 310 to build and keep in memory the multi-step integration stage outputs growing exponentially [6]. For instance, 5877241 multi-step paths can be generated for a distance of 30 frames using steps 1, 2, 5 and 10. Up to a few thousands can be actually considered to avoid computational and memory issues. The selection of L step sequences $\mathbf{\Gamma}_{ref,n}^* = \{\gamma_0, \gamma_1, \dots, \gamma_{L-1}\}$ among the K possible step
 315 sequences $\mathbf{\Gamma}_{ref,n}$ is therefore necessary, with $L \ll K$.

Two complexity reduction rules are taken from [6]. We start by removing the largest step sequences in terms of number of constituting steps. A threshold of K_{max} number of steps is thus set and only step sequences $\gamma_k = \{\alpha_1^k, \alpha_2^k, \dots, \alpha_{K_{\gamma_k}}^k\}$ for which $K_{\gamma_k} \leq K_{max}$ are kept. Indeed, too many steps
 320 may induce an important drift due to multiple intermediates. Then, random selection of L step sequences among remaining ones is performed.

3.2. Long-term match selection

Once step selection is performed, we obtain for each superpixel \mathbf{f}_i of I_{ref} a set of L candidate superpixels $\mathcal{T}_{ref,n}(\mathbf{f}_i) = \{h_{ref,n}(\mathbf{f}_i)|_{\gamma_l \in \mathbf{\Gamma}_{ref,n}^*}\}$ defined in I_n
 325 with $l \in \{0, \dots, L-1\}$. The final candidate selection is performed via majority voting among $\mathcal{T}_{ref,n}(\mathbf{f}_i)$, i.e. the final matching $h_{ref,n}(\mathbf{f}_i) = \mathbf{s}_j$ is defined such that $c(\mathbf{s}_j)$ satisfies:

$$c(\mathbf{s}_j) = \arg \max \text{hist}(\{c(h_{ref,n}(\mathbf{f}_i)|_{\gamma_l}) \mid \gamma_l \in \mathbf{\Gamma}_{ref,n}^*\}) \quad (12)$$

Thank to the random step sequence selection (Sect.3.1), the set of generated superpixel candidates is both significative and uncorrelated enough to assume
 330 that the most represented superpixel provides an accurate superpixel match.

Forward-backward consistency can be also considered in this context by providing to-the-reference multi-step paths additionally to from-the-reference ones.

We thus incorporate in $\mathcal{T}_{ref,n}(\mathbf{f}_i)$ superpixels \mathbf{s}_j such that $h_{n,ref}(\mathbf{s}_j)|_{\gamma_l \in \mathbf{\Gamma}_{n,ref}^*} = \mathbf{f}_i$ where $\mathbf{\Gamma}_{n,ref}^*$ is the set of L selected step sequences in the to-the-reference
 335 direction. The resulting additional superpixel candidates are referred as *reverse* candidates in opposition to *direct* ones, i.e. those which were formerly stacked into $\mathcal{T}_{ref,n}(\mathbf{f}_i)$. To further guide mutual matching between \mathbf{f}_i and \mathbf{s}_j , one can apply majority voting (Eq.12) only on **superpixel candidates** generated in both from/to-the-reference directions.

340 Superpixel correspondences with respect to I_{ref} are provided for the whole sequence relying on multi-step integration applied independently for each pair $\{I_{ref}, I_n\} \forall n \neq ref$ and based on unsupervised learning-based superpixel matching as an elementary estimator.

4. Application to video object tracking

345 Different aspects of the proposed methodology are evaluated through video object tracking experiments. First, the ability of unsupervised learning-based superpixel matching to provide a reliable accurate elementary estimator between consecutive and distant frames is proven with comparisons to state-of-the-art methods (Sect.4.1). Second, the capacity of the proposed multi-step integration
 350 stage to perform robust long-term superpixel tracking is shown using both k NN and RF-based multi-step elementary superpixel matches (Sect.4.2). Moreover, multi-step integration results are assessed with respect to straightforward direct and sequential integration outputs. Third, multi-step integration is further analyzed by studying the impact of different candidate generation strategies in
 355 terms of tracking accuracy (Sect.4.3).

To provide a generic evaluation while ensuring content diversity and representativity, video object tracking is performed over 10 sequences (Tab.1) extracted from 4 databases: **bag**, **fish3** (denoted **fsh3**) and **octopus** (**octo**) from the Visual Object Tracking (VOT) database [32], **sleep1** (**sle1**) with albedo
 360 from MPI Sintel [33], **lapa** from the laparoscopy dataset [29] as well as **swan**, **bear**, **camel** (**cam1**), **cows** and **flamingo** (**flam**) from the Densely Annotated

	img	size	obj	NR	LD	BC	DB	SV	PO	TS	IC
bag [32]	101	480×360	1	x	x		x				x
fsh3 [32]	101	480×270	1		x		x				x
octo [32]	51	640×360	1	x		x	x	x			
lapa [29]	81	364×270	1			x	x				x
sle1 [33]	50	512×218	3			x		x		x	
swan [34]	50	427×240	1		x		x				x
bear [34]	82	427×240	1	x	x	x	x		x		x
caml [34]	90	427×240	1	x	x	x	x		x	x	x
cows [34]	104	427×240	1	x	x		x	x	x		x
flam [34]	80	427×240	1	x	x	x	x			x	x

Table 1: Overview of sequences extracted from [32, 29, 33, 34] and used for object tracking experiments with associated sequence length, tracked object number and video attributes including complex non-rigid motion (NR), large displacement (LD), background clutter (BC), dynamic background (DB), scale variations (SV), partial occlusions (PO), thin structures (TS), illuminations changes and shadows (IC).

Video Segmentation (DAVIS) database [34]. As detailed in Tab.1, these sequences cover altogether many challenging situations such as complex non-rigid motion (NR), large displacement (LD), background clutter (BC), i.e. color similarities with background or between objects, dynamic background (DB) including moving background objects and camera viewpoint changes, scale variations (SV), partial occlusions (PO), thin structures (TS), illuminations changes and shadows (IC). Image sizes vary from 364×270 to 640×360 (Tab.1). Except for lapa whose ground-truth (GT) masks have been created from our own, all sequences were provided with associated GT masks indicating exact object delineations.

Video object tracking, also called semi-supervised video object segmentation task, consists in estimating for the whole sequence the exact location of a semantically meaningful free-shape region of interest (ROI) manually defined in one single image referred as reference frame. Once produced, tracking results are assessed for each pair $\{I_{ref}, I_n\}$ with $n \neq ref$ based on three complementary measures. First, Dice scores [35] measure the region-based segmentation similarity between estimated X and GT Y masks by computing $\frac{2|X \cap Y|}{|X| + |Y|}$. Then, contour-based precision P_c and recall R_c between estimated and GT masks can be estimated relying on bipartite graph matching to be robust to small inac-

380 curacies [36]. In practice, we focus on the F-measure combining precision and recall using $\mathcal{F} = \frac{2P_c R_c}{P_c + R_c}$. Bi-partite matching is approximated using morphology operators, as in [34]. Finally, consistency-based assessment is performed relying on the percentage of pixels of I_{ref} located inside the ROI and whose belonging
 385 superpixel \mathbf{f}_{ref} is consistent in terms of forward-backward binary consistency:

$$h_{n,ref}(h_{ref,n}(\mathbf{f}_{ref})) = \mathbf{f}_{ref} \quad (13)$$

In terms of computation time, performing RF-based matching followed by multi-step integration using steps $\{1, 2, 5, 10, 20\}$ on a sequence of 640×360
 390 frames such as `octo` with 500 **SLIC** [28] superpixels takes approximately 6 minutes per frame using a 3.1GHz Intel Xeon CPU processor and Python implementation, without extensive code optimization. Processing time is reduced about 17% when relying on k NN for unsupervised learning.

4.1. Elementary superpixel matching

395 Our first experiments consist in evaluating the proposed unsupervised learning-based superpixel matching (Sect.2) between consecutive and distant frames against state-of-the-art methods. In this direction, ROI tracking is performed through direct integration (DIR) in the to-the-reference direction, i.e. relying on direct processing of image pairs $\{I_n, I_{ref}\}$ without any sequential or multi-
 400 step combinations of pre-estimated superpixel matches. Unsupervised learning-based matching using both k NN and RF classifiers is compared to three other methodologies: superpixel-to-superpixel matching using superpixel-wise average color (RGBm) and color histogram (RGBh) features, PatchMatch (PM) [9], as well as optical flow through Farnebäck [8] and SIFT Flow [37]. Unsupervised
 405 learning-based matching works with $|\mathcal{F}| = |\mathcal{S}| = 500$ **SLIC** superpixels per frame with the same compactness parameter (i.e. same weighting between spatial and intensity proximity for each frame) and employs $K_a = 80$ context-rich multi-channel features computed with $\Phi = 40$ as maximal radius and $w \in \{3, 5, 7\}$ as possible box sizes (Sect.2.4). RF is made of $T = 100$ trees whereas k NN relies
 410 on 5 neighbors for queries. RGBm and RGBh use respectively average RGB

colors and RGB histograms (using 10 bins) as superpixel-wise features to give correspondences in I_{ref} for each superpixel of I_n in a nearest neighbor manner. As for unsupervised learning-based matching, RGBm and RGBh exploit images decomposed into 500 **SLIC** superpixels. PM [9] is looking for the best patch matches using 9×9 windows with 6 iterations including both propagation and random refinement steps. Farnebäck [8] and SIFT Flow [37] estimators are used using by-default parameters. Learning-based and superpixel-to-superpixel matching are performed once a groundtruth label is assigned to each superpixel of I_{ref} to indicate its belonging to the ROI to be tracked. **In practice**, 50% of the constituting pixels must be included into the ROI to label a superpixel as part of the object in I_{ref} . Label propagation can be then easily done at the superpixel level once to-the-reference superpixel pairings are obtained. Conversely, PM and optical flow estimators use dense to-the-reference fields to propagate labels at the pixel level from I_{ref} to the whole sequence.

Dice and F-measure scores temporally averaged across each of the previously described sequences are given Tab.2 for each method. Bold and underline results indicate first and second best scores. Results indicate a good matching accuracy reached using the proposed unsupervised learning-based strategy for both consecutive and distant frames. On average, RF-based superpixel pairings provide the best direct tracking results with **Dice** and F-measure of 86.9 and 80.3, followed by k -NN-based results which reach 74.7 and 70.3. Both methods are significantly superior to the other **existing approaches**. Averaged **Dice** (F-measure) goes down to 67.6 (57.8) and 61.6 (53.0) for RGBh and RGBm respectively. Despite fairly good scores for **octo**, **cam1**, **cows** and **flam**, PM and optical flow methods do not globally outperform unsupervised learning-based and superpixel-to-superpixel matching with averaged **Dice** (F-measure) of 57.9 (53.2), 54.2 (48.4) and 51.9 (46.8) for PM, Farnebäck and SIFT Flow.

These findings are illustrated visually Fig.5 for the pair $\{I_1, I_{34}\}$ of **swan**. Red and green boundaries denote propagated and GT ROI location. We can notice that PM, Farnebäck and SIFT Flow under-estimate the area covered by the swan, especially for the neck and near the water. RGBm, RGBh as

	Dice							F-measure						
	spx matching		PM[9]	optical flow		proposed DIR		spx matching		PM[9]	optical flow		proposed DIR	
	RGBm	RGBh		Far[8]	SF[37]	kNN	RF	RGBm	RGBh		Far[8]	SF[37]	kNN	RF
bag	67.1	87.5	11.6	11.8	11.0	96.9	<u>92.8</u>	59.3	82.6	10.7	10.1	10.6	97.8	<u>89.0</u>
fsh3	89.8	<u>89.7</u>	23.8	16.7	15.2	64.8	89.4	91.9	<u>91.4</u>	27.4	21.4	18.7	67.3	81.5
octo	25.6	47.6	98.6	<u>96.5</u>	96.4	84.8	85.5	17.2	32.1	95.2	<u>92.9</u>	<u>92.9</u>	75.6	74.9
lapa	65.8	67.7	70.4	61.1	48.5	89.2	<u>87.9</u>	57.1	57.7	53.0	51.2	51.5	86.1	<u>85.3</u>
sl1.1	<u>84.5</u>	84.2	41.7	35.2	32.5	82.9	94.9	81.2	<u>83.3</u>	53.1	42.2	43.1	81.8	94.4
sl1.2	24.5	51.2	24.3	13.4	11.8	<u>82.0</u>	89.4	21.6	44.6	50.4	34.1	24.4	<u>82.4</u>	90.3
sl1.3	38.5	51.2	13.5	7.95	7.75	<u>66.7</u>	88.3	39.5	49.1	24.7	12.1	11.4	<u>72.9</u>	77.5
swan	<u>91.0</u>	89.2	78.5	84.9	83.9	90.5	91.3	<u>87.4</u>	83.9	62.6	70.7	69.2	85.7	88.0
bear	81.5	73.8	79.7	<u>85.6</u>	84.0	84.3	89.1	57.7	46.5	56.3	64.5	<u>65.9</u>	65.2	76.3
camel	58.5	49.3	85.2	85.2	80.5	67.5	70.3	45.6	38.2	72.4	64.3	<u>67.2</u>	56.4	59.3
cows	64.2	67.5	89.6	89.6	84.4	16.8	87.1	32.2	36.9	77.3	68.3	60.4	15.5	<u>69.8</u>
flam	48.2	52.4	77.5	<u>77.2</u>	66.3	70.0	76.2	45.6	46.7	55.6	48.8	45.9	<u>56.3</u>	67.2
avg	<u>61.6</u>	<u>67.6</u>	<u>57.9</u>	<u>54.2</u>	<u>51.9</u>	<u>74.7</u>	86.9	<u>53.0</u>	<u>57.8</u>	<u>53.2</u>	<u>48.4</u>	<u>46.8</u>	<u>70.3</u>	80.3

Table 2: Dice and F-measure scores for ROI tracking across 10 sequences using direct integration (DIR), i.e. direct processing of image pairs $\{I_{ref}, I_n\}$. Four methodologies are compared: superpixel-to-superpixel matching using superpixel-wise average color (RGBm) and color histogram (RGBh) features, PatchMatch (PM) [9], optical flow through Farnebäck [8] and SIFT Flow [37] as well as the proposed unsupervised learning-based superpixel matching using kNN/RF classifiers. Bold and underline results indicate first and second best scores.

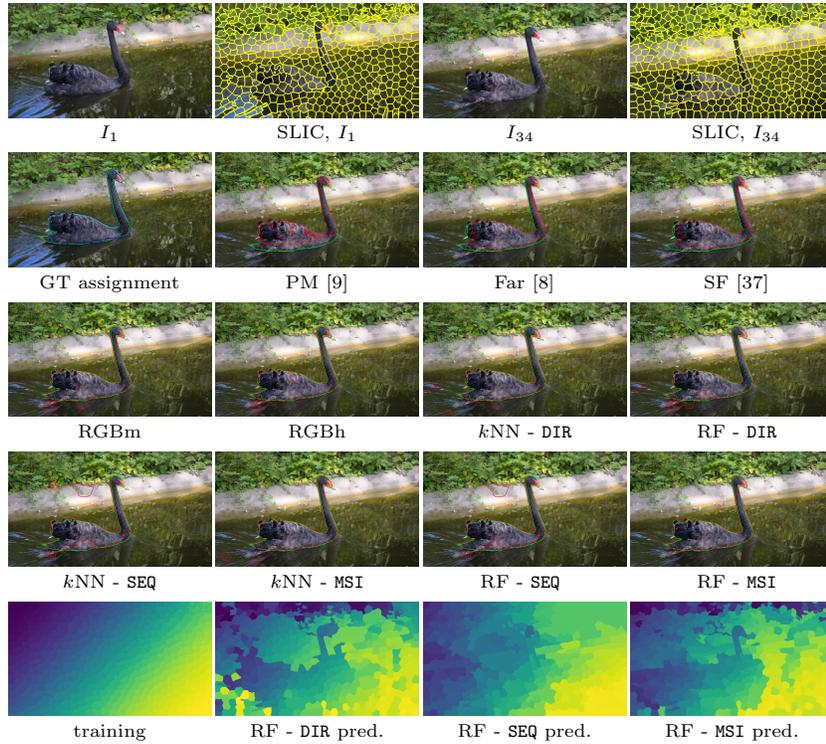


Figure 5: ROI propagation for *swan* ($\{I_1, I_{34}\}$) [34] with DIR, SEQ and MSI (steps $\{1, 2, 5, 10, 20, 30\}$) integrations using k NN and RF. Results are compared with: superpixel-to-superpixel matching with average color (RGBm) and color histogram (RGBh) features, PatchMatch (PM) [9], optical flow through Farnebäck (Far) [8] and SIFT Flow (SF) [37]. Blue boundaries in I_1 indicate superpixel labelling resulting from GT assignment. Green and red boundaries correspond to groundtruth (GT) and estimated tracking results. The last row displays training (I_1) and prediction (I_{34}) masks resulting from DIR, SEQ and MSI integrations of RF-based elementary pairings.

well as RF and k NN-based direct (DIR) superpixel pairings provide clearly better contours despite the tendency to propagate the ROI outside the swan area due to shadows and color similarities with background. One drawback with
445 superpixel-based methods is the lack of boundary adherence which may not suit perfectly the object to be tracked. This aspect is revealed for I_1 by the blue boundaries which indicate the superpixel labelling resulting from GT assignment. However, Tab.2 demonstrates that this limitation is compensated by robust superpixel matching heuristics compared to straightforward pixel-wise
450 matching and mask propagation which rely on a regular decomposition of the image grid without enough context considerations. Finally, Fig.5 shows that a more accurate propagation can be achieved with multi-step integration of learning-based superpixel pairings, especially with RF. It tends to indicate the ability of unsupervised learning-based superpixel matching to provide a reliable
455 and accurate enough elementary estimator towards efficient long-term multi-step matching and tracking. The performance achieved with multi-step integration is more deeply demonstrated in the next section.

4.2. Long-term superpixel tracking

Long-term ROI tracking resulting from direct (DIR), sequential (SEQ) and
460 multi-step (MSI) integrations are compared based on unsupervised learning-based SLIC [28] superpixel matching whose accuracy against state-of-the-art methods has been demonstrated in Sect.4.1 with both k NN and RF classifiers. MSI is applied with $L=200$ maximal step sequences per image pairs. Only step sequences whose length is less than or equal to $K_{max} = 7$ are kept to prevent from
465 motion drift (Sect.3.1). k NN and RF-based elementary multi-step superpixel matches are obtained with steps $\{1, 2, 5, 10, 20\}$ for octo, swan and sle1 and $\{1, 2, 5, 10, 20, 30, 50\}$ for longer sequences (fsh3, bag, lapa, bear, cam1, cows and flam). Context-rich features are estimated using the same parameters as in Sect.4.1. Majority voting (Eq.12) focuses only on superpixel candidates generated in both to/from-the-reference directions to improve forward-backward
470 consistency (see Sect.4.3 for further details).

	Dice						F-measure						consistency					
	<i>k</i> NN			RF			<i>k</i> NN			RF			<i>k</i> NN			RF		
	DIR	SEQ	MSI	DIR	SEQ	MSI	DIR	SEQ	MSI	DIR	SEQ	MSI	DIR	SEQ	MSI	DIR	SEQ	MSI
bag	96.9	<u>97.7</u>	<u>97.7</u>	92.8	74.9	92.9	97.8	<u>99.7</u>	99.4	89.0	65.9	88.3	<u>39.2</u>	12.1	32.1	30.2	12.2	27.8
fsh3	64.8	90.5	85.7	89.4	91.1	<u>92.0</u>	67.3	92.6	81.5	91.8	93.9	<u>95.4</u>	56.5	37.6	74.3	69.5	38.8	58.7
octo	84.8	91.3	93.0	85.5	91.3	92.8	75.6	83.6	<u>86.3</u>	74.9	83.6	86.0	<u>84.6</u>	67.5	76.6	82.9	67.1	75.7
lapa	89.2	86.5	92.8	87.9	88.3	92.8	86.1	86.2	<u>95.0</u>	85.3	86.8	94.8	96.1	68.9	96.8	90.2	64.6	91.5
sl1.1	82.9	95.3	95.6	94.9	95.5	<u>95.9</u>	81.8	93.2	94.8	94.4	94.0	<u>95.5</u>	89.2	63.4	<u>96.8</u>	94.5	76.1	93.4
sl1.2	82.0	80.9	90.3	89.4	80.5	<u>90.9</u>	82.4	77.9	93.0	90.3	73.1	<u>93.9</u>	76.0	62.4	<u>91.6</u>	84.9	61.6	91.2
sl1.3	66.7	92.4	92.2	88.3	92.4	<u>94.2</u>	72.9	76.7	78.1	77.5	76.7	<u>81.6</u>	86.7	78.1	95.0	87.7	75.6	100
swan	90.5	85.0	<u>93.5</u>	91.3	86.9	93.4	85.7	73.9	93.6	88.0	78.4	<u>93.8</u>	<u>86.1</u>	57.7	83.0	82.3	59.3	77.7
bear	84.3	87.8	87.7	89.1	87.0	<u>92.5</u>	65.2	68.2	72.2	76.3	68.9	<u>82.1</u>	65.6	61.2	<u>68.5</u>	62.7	62.1	67.2
camel	67.5	76.6	76.9	70.3	77.7	<u>79.6</u>	56.4	64.7	67.2	59.3	63.3	<u>69.1</u>	58.7	41.5	51.8	<u>63.8</u>	42.3	54.7
cows	16.8	80.1	87.3	87.1	79.1	<u>89.1</u>	15.5	59.6	66.8	69.8	61.4	<u>74.1</u>	15.5	39.6	<u>68.5</u>	75.8	39.3	67.6
flam	70.0	63.5	76.7	76.2	66.5	80.8	56.3	62.5	67.1	67.2	62.8	<u>71.9</u>	55.4	34.9	63.4	43.0	39.7	55.0
avg	<i>74.7</i>	<i>85.6</i>	<i>89.1</i>	<i>86.9</i>	<i>84.3</i>	<i>90.6</i>	<i>70.3</i>	<i>78.2</i>	<i>82.9</i>	<i>80.3</i>	<i>75.7</i>	<i>85.5</i>	<i>67.5</i>	<i>52.1</i>	<i>74.9</i>	<i>72.3</i>	<i>53.2</i>	<i>71.7</i>

Table 3: **Dice**, F-measure and consistency scores for ROI tracking across 10 sequences. We compare direct (DIR), sequential (SEQ) and multi-step (MSI) integration based on unsupervised learning-based superpixel matching using *k*NN and RF. Bold results indicate the best performance between DIR, SEQ and MSI. Underline scores highlight best results between *k*NN and RF-based methods.

Tab.3 presents temporally averaged metrics (Dice, F-measure and consistency scores) obtained by DIR, SEQ and MSI across all sequences using k NN and RF. Except for consistency scores when relying on RF, Tab.3 confirms that MSI is the best integration strategy towards long-term superpixel tracking compared to DIR and SEQ. For instance, RF and k NN-based MSI reach the highest Dice scores with 90.6 and 89.1 in comparison to 84.3 (86.9) and 85.6 (74.7) obtained with RF and k NN-based SEQ (DIR). Second and third positions in terms of Dice and F-measure vary depending on the classifier. SEQ outperforms DIR for k NN whereas RF exhibits the opposite behavior. Except for MSI in terms of consistency and SEQ for F-measure, another main finding is that RF-based elementary matches usually make better long-term tracking than k NN-based pairings. This suggests that RF is more able to deal with a large amount of classes in this context. Despite a slightly poorer management of context-rich features, k NN remains an interesting and faster alternative, depending on expected tracking quality requirements. Since the balance between processing time and tracking quality is guided by the application, we can imagine that k -NN could be chosen against RF in some contexts.

Temporal evolutions of Dice and F-measure scores are displayed Fig.6 along `lapa`, `s1e1.2` (`s1e1` for object 2), `octo` and `swan` sequences with both classifiers. As already confirmed, best tracking results are reached with MSI compared to DIR and SEQ, especially for distant pairs. Contrary to SEQ whose performance decreases across sequences due to error accumulations (`lapa` and `swan` especially), multi-step estimations involved in MSI allow to fix uncorrect superpixel tracks as we can notice for `s1e1.2` from frame I_{20} . Moreover, DIR is not prone to motion drift as SEQ but direct matching becomes tedious when inter-frame distances increase as shown for `octo` starting from I_{109} . Finally, it can be noticed that the temporal behavior remains almost the same regardless of the classifier.

Finally, quantitative results are illustrated by series of ROI selection and visual tracking examples for several pairs of `lapa` (Fig.7), `s1e1` (Fig.8), `bag`, `fsh3`, and `octo` (Fig.9) sequences. Fig.7 shows that k NN-based MSI provides a very good delineation of the surgical tool for all image pairs, which suggests that

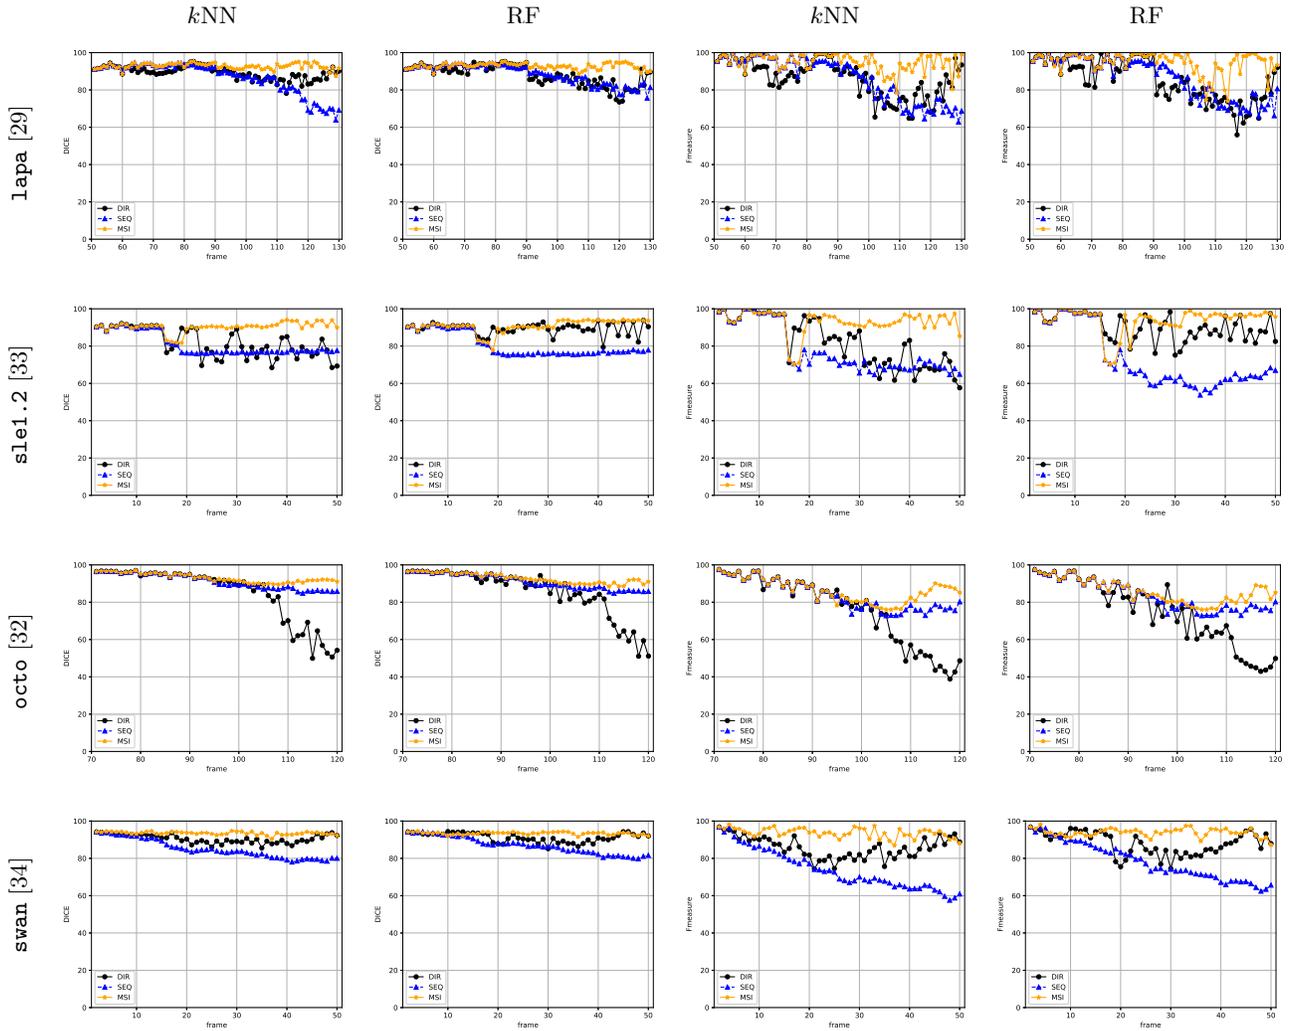


Figure 6: Temporal evolution of **Dice** (left) and **F-measure** (right) scores during ROI tracking across `lapa` [29], `sle1.2` [33], `octo` [32] and `swan` [34] sequences. We compare direct (DIR), sequential (SEQ) and multi-step (MSI) integration based on k NN and RF-based elementary pairings.

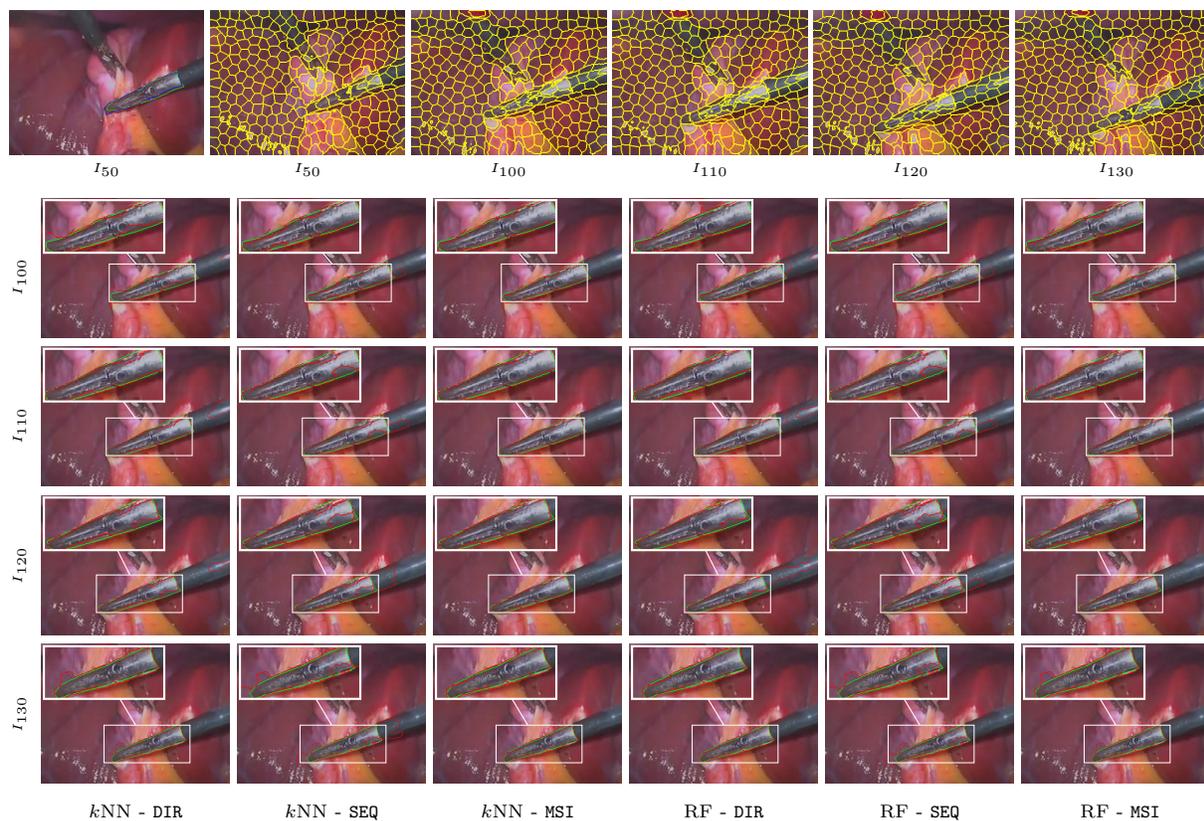


Figure 7: ROI selection and tracking across lapa sequence [29] from I_{50} to I_{130} . We compare direct (DIR), sequential (SEQ) and multi-step (MSI, steps $\{1, 2, 5, 10, 20, 30, 50\}$) integration (Sect.3) of superpixel matches obtained through unsupervised learning-based superpixel matching (Sect.2) with k NN and RF [18]. Superpixel decompositions are obtained *via* SLIC [28]. Blue boundaries (I_{50}) indicate superpixel labelling resulting from GT assignment. Green and red boundaries correspond resp. groundtruth (GT) and estimated tracking results.

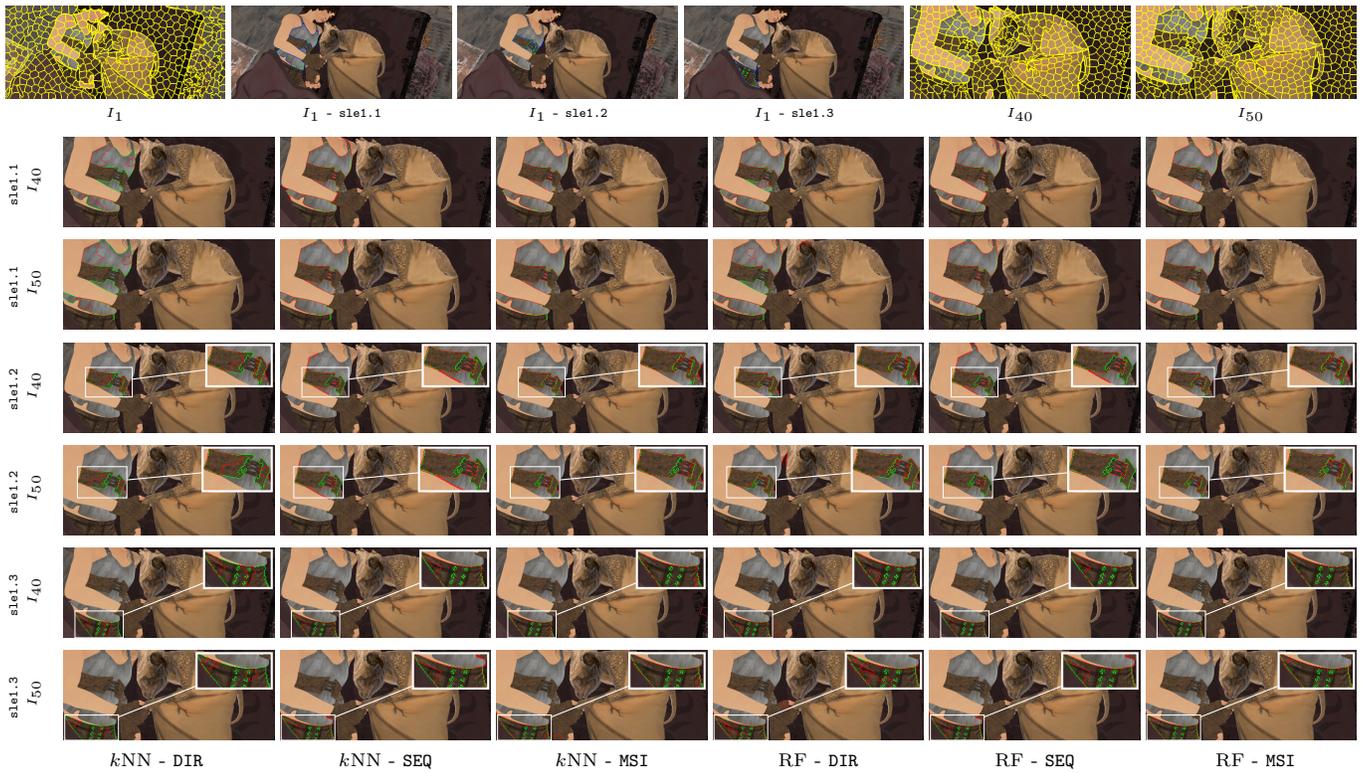


Figure 8: ROI selections and tracking across `sle1` sequence [33] from I_1 to I_{50} . We compare direct (DIR), sequential (SEQ) and multi-step (MSI, steps $\{1, 2, 5, 10, 20\}$) integration based on unsupervised learning-based superpixel matching with k NN and RF.

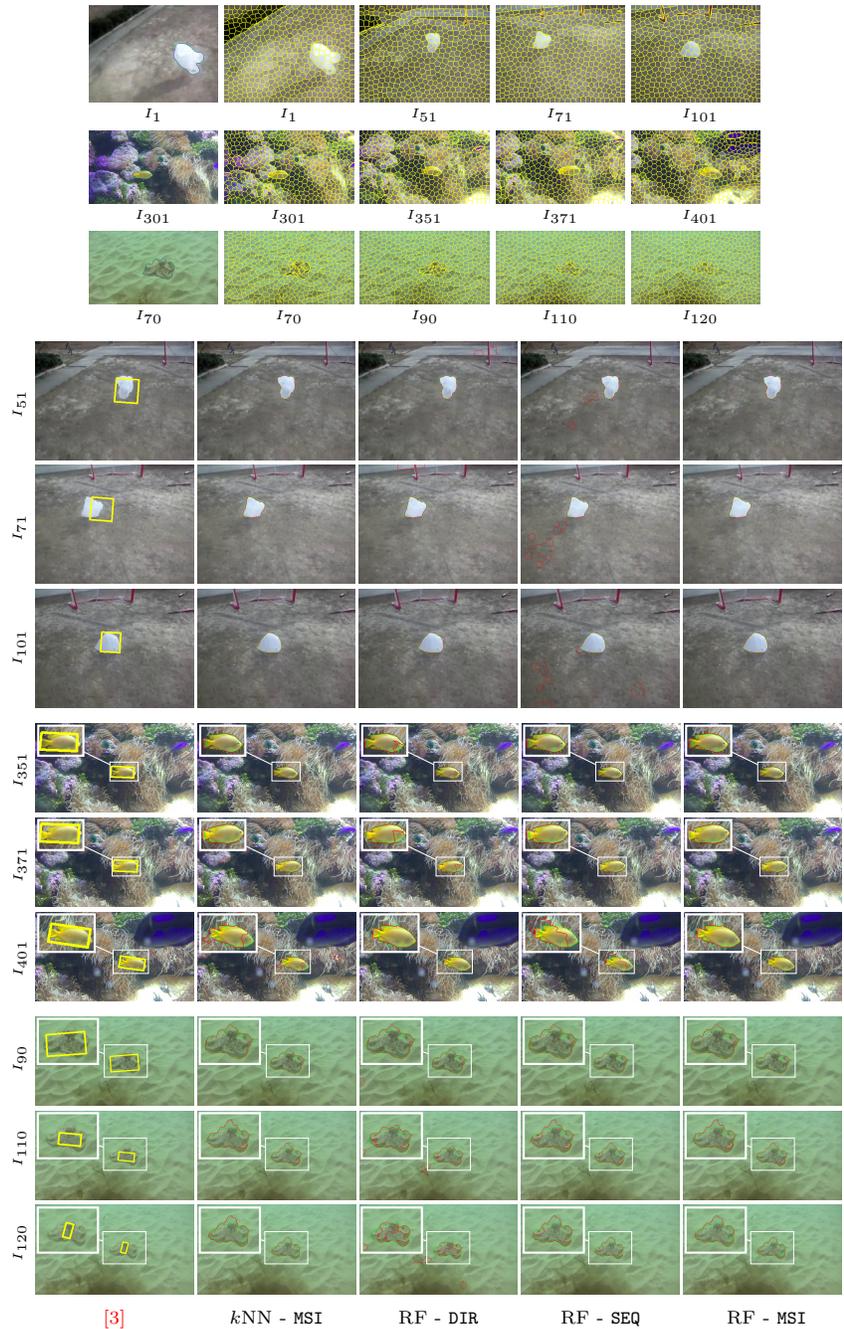


Figure 9: ROI selection and tracking across **bag** (resp. **fsh3**, **octo**) sequences [32] for segments $[I_1, I_{101}]$ ($[I_{301}, I_{401}]$, $[I_{70}, I_{120}]$). We **evaluate** DIR, SEQ and MSI (steps $\{1, 2, 5, 10, 20, 30, 50\}$ for **bag** and **fsh3**, $\{1, 2, 5, 10, 20\}$ for **octo**) integration based on unsupervised learning-based superpixel matching with k NN and RF [18]. **Results from [3] are also reported with bounding boxes displayed in yellow.**

series of medical images can be also processed with the proposed methodology. Reliable ROI tracking through MSI is also shown on synthetic images (Fig.8) despite strong scale variations. Propagation of matching errors with SEQ is clearly illustrated Fig.7 for `1apa` (k NN) and Fig.9 for `bag`. Tracking failures with DIR are temporally uncorrelated but strong enough to damage the propagation task for the fish (Fig.9) due to color variations of its right part. ROI tracking for `octo` (Fig.9) gives correct results both with k NN and RF despite significant color similarities with the dynamic background. Experiments for `swan` (Fig.5) and `sle1` (Fig.8) videos demonstrate that RF-based MSI outperforms k NN-based MSI as well as RF-based DIR and SEQ. More globally, results suggest that the SLIC superpixel algorithm [28] is enough accurate to provide visually meaningful entities to be accurately tracked along video sequences.

Comparisons with [3] are provided for `bag`, `fsh3`, and `octo` sequences (Fig.9). The first finding is that [3] only provides sparse bounding box tracking and do not describe how the boundaries evolves in time, as performed by the proposed superpixel tracker. Despite an accurate tracking of the foreground fish (`fsh3` sequence), [3] results in long-term tracking failures for `bag` (especially in I_{71}). Furthermore, it does not show robustness to strong scale variations and color similarities with the dynamic background (I_{110} and I_{120} for `octo`) contrary to RF-based MSI which offers an efficient pixel-wise long-term object delineation.

Another visualization through prediction masks given Fig.5 for the $\{I_1, I_{34}\}$ `swan` pair confirms the ability of RF-based MSI to reach accurate long-term correspondences (see for instance the swan beak). In such maps, a given object part must keep the same color between training and prediction in case of correct matching. Improvements from RF-based DIR to MSI can be noticed for both low-textured areas (as in the water) and highly-textured ones.

4.3. Long-term candidate generation

We propose to perform a more in-depth study of multi-step integration by comparing different long-term candidate generation strategies in terms of tracking accuracy. As described Sect.3.2, long-term superpixel candidates can be

	Dice			F-measure			consistency		
	MSId	MSIr	MSIm	MSId	MSIr	MSIm	MSId	MSIr	MSIm
bag	84.5	91.9	92.9	80.3	88.3	88.3	19.2	27.8	27.8
fsh3	92.0	92.0	92.0	94.9	95.4	95.4	52.8	59.6	58.7
octo	91.8	92.2	92.8	83.8	84.4	86.0	71.6	75.8	75.7
lapa	92.6	92.9	92.8	94.2	95.2	94.8	85.4	91.5	91.5
sl1.1	95.7	95.8	95.9	94.5	94.8	95.5	91.2	93.7	93.4
sl1.2	90.3	91.0	90.9	91.8	94.5	93.9	78.6	90.4	91.2
sl1.3	94.2	94.3	94.2	81.3	81.5	81.6	90.0	99.5	100
swan	93.2	93.1	93.4	93.4	93.0	93.8	70.6	77.6	77.7
bear	92.2	92.5	92.5	83.3	82.9	82.1	54.7	67.0	67.2
cam1	79.6	79.3	79.6	69.6	69.3	69.1	44.2	54.1	54.7
cows	90.0	89.3	89.1	75.5	74.4	74.1	59.8	67.2	67.6
flam	79.8	80.7	80.8	64.0	69.0	71.9	40.9	54.3	55.0
avg	89.7	90.4	90.6	83.9	85.2	85.5	63.3	71.5	71.7

Table 4: **Dice**, F-measure and consistency scores for ROI tracking across 10 sequences. Based on RF-based superpixel elementary matches, we compare three different superpixel candidate generation strategies for multi-step (MSI) integration using: only direct candidates (MSId), direct and reverse candidates (MSIr), only candidates generated in both direct and reverse directions (MSIm). Bold results indicate the best performance.

generated using direct candidates only (MSId), both direct and reverse candidates (MSIr) or only candidates generated in both direct and reverse directions (MSIm). Note that the previously given MSI results corresponded to MSIm where only superpixel duplicates are taken into account for majority voting (Eq.12). MSId, MSIr and MSIm are comparatively evaluated in terms of tracking accuracy based on RF-based superpixel elementary matches. **Dice**, F-measure and consistency scores are reported across the 10 sequences used for ROI tracking.

Results from Tab.4 bring two main findings. First, we observe that tracking accuracy is improved when reverse candidates are used additionally to direct ones. Consistency ratios are clearly improved (from 63.3 to 71.5% when comparing MSId/MSIr) as expected but **Dice** and F-measure improvements can be also observed with gains of 0.7 and 1.3 between MSId/MSIr. Second, relying on superpixel duplicates only (MSIm) brings a slightly better ROI tracking compared to MSIr. Average results slightly increase from 90.4 to 90.6, 85.2 to 85.5

and 71.5 to 71.7 for **Dice**, F-measure and consistency which shows that extensive mutual matching guidance as performed with **MSIm** is the best **way** to perform long-term superpixel tracking from multi-step elementary correspondences.

550 **5. Conclusion**

In this work, we proposed a two-step pipeline dedicated to long-term superpixel tracking. **Unsupervised** learning-based superpixel matching is firstly considered as an elementary displacement estimator to provide correspondences between consecutive and distant images using either nearest neighbors or random forests with robust context-rich features we extended from greyscale to
555 multi-channel and forward-backward consistency constraints. Resulting elementary matches are then combined along multi-step paths running through the sequence with various inter-frame distances to produce a large set of candidate long-term superpixel pairings upon which majority voting selection is performed. Compared to state-of-the-art methods including pixel or patch-based
560 strategies which may suffer from regular support regions, video object tracking experiments demonstrate that unsupervised learning can produce reliable correspondences between **visually meaningful entities**. Moreover, the ability of multi-step integration to combine these pairings towards accurate long-term
565 superpixel tracking has been shown against usual direct and sequential integrations. Extending this work from single to hierarchical multi-scale superpixel decomposition would deserve further investigation for future research since dealing with multiple spatial extends can drive the matching process in a coarse-to-fine fashion. Other **superpixel algorithms as well as other** features such as spectral
570 features could be employed to further improve unsupervised learning-based matching while reducing processing time. In addition, very long-term superpixel tracking could be reached by combining superpixel pairings estimated with respect to multiple reference frames. Our contributions also give new insights for optical flow and registration initialization, in particular to provide a better
575 management of large displacements, appearance and illumination changes. More

generally, the proposed framework could be easily extended to other imaging modalities including series of medical images for anatomical structure tracking.

References

- [1] J. Lezama, K. Alahari, J. Sivic, I. Laptev, Track to the future: Spatio-temporal video segmentation with long-range motion cues, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2011, pp. 3369–3376. 580
- [2] S. M. Seitz, B. Curless, J. Diebel, D. Scharstein, R. Szeliski, A comparison and evaluation of multi-view stereo reconstruction algorithms, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2006, pp. 519–528. 585
- [3] F. Yang, H. Lu, M.-H. Yang, Robust superpixel tracking, *IEEE Transactions on Image Processing* 23 (4) (2014) 1639–1651.
- [4] H. Wang, A. Kläser, C. Schmid, C.-L. Liu, Dense trajectories and motion boundary descriptors for action recognition, *International Journal of Computer Vision* 103 (1) (2013) 60–79. 590
- [5] X. Cao, Z. Li, Q. Dai, Semi-automatic 2D-to-3D conversion using disparity propagation, *IEEE Transactions on Broadcasting* 57 (2) (2011) 491–499.
- [6] P.-H. Conze, P. Robert, T. Crivelli, L. Morin, Multi-reference combinatorial strategy towards longer long-term dense motion estimation, *Computer Vision and Image Understanding* 150 (2016) 66–80. 595
- [7] J. Shi, C. Tomasi, Good features to track, in: *IEEE International Conference on Computer Vision and Pattern Recognition*, 1994, pp. 593–600.
- [8] G. Farnebäck, Two-frame motion estimation based on polynomial expansion, *Image Analysis* (2003) 363–370.
- [9] C. Barnes, E. Shechtman, A. Finkelstein, D. B. Goldman, Patchmatch: A randomized correspondence algorithm for structural image editing, *ACM Transaction on Graphics* 28 (3) (2009) 24–1. 600

- [10] C. Barnes, E. Shechtman, D. Goldman, A. Finkelstein, The generalized patchmatch correspondence algorithm, *European Conference on Computer Vision* (2010) 29–43.
605
- [11] D. Stutz, A. Hermans, B. Leibe, Superpixels: An evaluation of the state-of-the-art, *Computer Vision and Image Understanding* 166 (2018) 1–27.
- [12] H.-S. Chang, Y.-C. F. Wang, Superpixel-based large displacement optical flow, in: *IEEE International Conference on Image Processing*, 2013, pp. 3835–3839.
610
- [13] S. Donné, J. Aelterman, B. Goossens, W. Philips, Fast and robust variational optical flow for high-resolution images using SLIC superpixels, in: *International Conference on Advanced Concepts for Intelligent Vision Systems*, 2015, pp. 205–216.
- [14] H. Fan, J. Xiang, Z. Chen, Visual tracking by local superpixel matching with markov random field, in: *Pacific Rim Conference on Multimedia*, 2016.
615
- [15] R. Giraud, V. T. Ta, A. Bugeau, P. Coupé, N. Papadakis, Superpatchmatch: An algorithm for robust correspondences using superpixel patches, *IEEE Transactions on Image Processing* 26 (8) (2017) 4068–4078.
- [16] F. Kanavati, T. Tong, K. Misawa, M. Fujiwara, K. Mori, D. Rueckert, B. Glocker, Supervoxel classification forests for estimating pairwise image correspondences, *Pattern Recognition* 63 (2017) 561–569.
620
- [17] P.-H. Conze, F. Tilquin, V. Noblet, F. Rousseau, F. Heitz, P. Pessaux, Hierarchical multi-scale supervoxel matching using random forests for automatic semi-dense abdominal image registration, in: *IEEE International Symposium on Biomedical Imaging*, 2017, pp. 490–493.
625
- [18] L. Breiman, Random Forests, *Machine learning* 45 (1) (2001) 5–32.
- [19] T. Crivelli, M. Fradet, P.-H. Conze, P. Robert, P. Pérez, Robust optical flow integration, *IEEE Transactions on Image Processing* 24 (1) (2015) 484–498.
630

- [20] L. Wang, H. Lu, M. H. Yang, Constrained superpixel tracking, *IEEE Transactions on Cybernetics* (99) (2017) 1–12.
- [21] J. Chang, D. Wei, J. W. Fisher, A video representation using temporal superpixels, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 2051–2058.
- 635 [22] S.-H. Lee, W.-D. Jang, C.-S. Kim, Temporal superpixels based on proximity-weighted patch matching, in: *IEEE International Conference on Computer Vision*, 2017, pp. 3610–3618.
- [23] C. Ma, J.-B. Huang, X. Yang, M.-H. Yang, Hierarchical convolutional features for visual tracking, in: *IEEE International Conference on Computer Vision*, 2015, pp. 3074–3082.
- 640 [24] H. Li, Y. Li, F. Porikli, Deeptack: Learning discriminative feature representations online for robust visual tracking, *IEEE Transactions on Image Processing* 25 (4) (2016) 1834–1848.
- [25] D. Zhang, H. Maei, X. Wang, Y.-F. Wang, Deep reinforcement learning for visual object tracking in videos, arXiv preprint 1701.08936.
- 645 [26] B. Glocker, D. Zikic, D. R. Haynor, Robust registration of longitudinal spine CT, in: *Medical Image Computing and Computer-Assisted Intervention*, 2014, pp. 251–258.
- [27] P.-H. Conze, F. Tilquin, M. Lamard, F. Heitz, G. Quellec, Long-term superpixel tracking using unsupervised learning and multi-step integration, in: *IEEE International Conference on Advanced Technologies for Signal and Image Processing*, 2018.
- 650 [28] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, S. Susstrunk, SLIC superpixels compared to state-of-the-art superpixel methods, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34 (2012) 2274–2282.
- 655

- [29] R. Sznitman, K. Ali, R. Richa, R. H. Taylor, G. D. Hager, P. Fua, Data-driven visual tracking in retinal microsurgery, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, 2012, pp. 568–575.
- [30] T. Brox, J. Malik, Large displacement optical flow: descriptor matching in variational motion estimation, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33 (3) (2011) 500–513.
- [31] S. Roth, V. Lempitsky, C. Rother, Discrete-continuous optimization for optical flow estimation, in: *Statistical and Geometrical Approaches to Visual Motion Analysis*, 2009, pp. 1–22.
- [32] M. Kristan, et al., The visual object tracking VOT2016 challenge results, in: *Workshop of European Conference on Computer Vision*, 2016.
- [33] D. J. Butler, J. Wulff, G. B. Stanley, M. J. Black, A naturalistic open source movie for optical flow evaluation, in: *European Conference on Computer Vision*, 2012, pp. 611–625.
- [34] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, A. Sorkine-Hornung, A benchmark dataset and evaluation methodology for video object segmentation, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 724–732.
- [35] A. P. Zijdenbos, B. M. Dawant, R. A. Margolin, A. C. Palmer, Morphometric analysis of white matter lesions in MR images: method and validation, *IEEE Transactions on Medical Imaging* 13 (4) (1994) 716–724.
- [36] D. R. Martin, C. C. Fowlkes, J. Malik, Learning to detect natural image boundaries using local brightness, color, and texture cues, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26 (5) (2004) 530–549.
- [37] C. Liu, J. Yuen, A. Torralba, Sift flow: Dense correspondence across scenes and its applications, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33 (5) (2011) 978–994.