

## Visualizing Skiers' Trajectories in Monocular Videos

Matteo Dunnhofer      Luca Sordi      Christian Micheloni  
Machine Learning and Perception Lab, University of Udine, Udine, Italy



Figure 1. **Visualization of skiers' trajectories in monocular videos.** We present `SkiTraVis`, an algorithm to compute and visualize the trajectory performed by an alpine skier in videos acquired by a single monocular camera. `SkiTraVis` takes in input a streaming video capturing a skier and visualizes its past trajectory in each frame. In this picture, the first and third rows show the input frames, while the second and fourth the trajectory computed by `SkiTraVis`. The interval in seconds from the start of the video is also reported.

### Abstract

*Trajectories are fundamental to winning in alpine skiing. Tools enabling the analysis of such curves can enhance the training activity and enrich broadcasting content. In this paper, we propose `SkiTraVis`, an algorithm to visualize the sequence of points traversed by a skier during its performance. `SkiTraVis` works on monocular videos and constitutes a pipeline of a visual tracker to model the skier's motion and of a frame correspondence module to estimate the camera's motion. The separation of the two motions enables the visualization of the trajectory according to the moving camera's perspective. We performed experiments on videos of real-world professional competitions to quantify the visualization error, the computational efficiency, as well as the applicability. Overall, the results achieved demonstrate the potential of our solution for broadcasting media enhancement and coach assistance.*

### 1. Introduction

Due to the popularity of alpine skiing [60], the competitive form of such a sport discipline is one of the most important markets of the winter sports industry [12]. The international regulations of alpine skiing competitions [54] require the definition of a path of so-called turning gates down a hill. Differing in the overall length of the course and in the distance between the gates, four popular alpine skiing sub-disciplines are present nowadays: Slalom (SL), Giant Slalom (GS), Super-G (SG), and Downhill (DH). In each of these, the winning goal is the same: ski down through the gates of the designed course in the shortest time possible.

To do so, an athlete should optimize its skiing trajectories along the track. Indeed, in alpine skiing the turns performed by an athlete form a trajectory of points on the slope surface that if optimized to reduce the distance with the turning gates can lead to saved time during the descent [6, 7]. Toward such an objective, it becomes important to measure

and analyze a trajectory in order to determine the specific points of the performance affecting the overall time taken. Machines able to perceive such evidence are valuable tools to obtain increased knowledge of the skiing performance. The currently available solutions to measure a trajectory of an alpine skier make use of wearable devices (e.g. GNSS trackers, IMUs) put on the body or skis, and of 3D terrain models to position the data coming from the sensors at different time steps [24, 35]. The drawback of such an approach is the repeated time-consuming installation and calibration of the sensor device on the athlete’s body, and the acquisition of a precise terrain model. Moreover, the contextual information influencing the skiing performance (e.g. positions of the turning gates, conditions of slope surface, weather, and visibility) is completely lost by such kinds of system. Research on training process improvement for skiing pointed out the opportunity for a post-performance video review to overcome such limitations [44, 58].

The importance of video-based performance analysis paves the way for training-assistive tools based on computer vision. Indeed, computer vision applied to videos capturing the athlete’s performance is a valid option to measure skiing performance without the need for GNSS sensors or ground surface models [44, 55]. Previous work on this topic focused on 3D skier pose estimation [1] on a motion capture system comprising six manually-operated mutually-calibrated pan-tilt-zoom cameras capturing a Giant Slalom track section of three gates. By means of multi-view human pose estimation algorithms [47, 63, 64], it was shown that the 3D pose of athletes can be reconstructed at different time steps, and from that time series, the skiing trajectory could be computed inside a virtual 3D environment. Despite proving the suitability of computer vision for trajectory analysis, such work has limited applicability in practice because of the complex multi-camera setup.

In this paper, we aim for a different goal: to visualize, directly inside the frames of a capturing monocular video, the trajectory of a skier. We describe the attempt to design and implement a computer vision algorithm to compute and visualize the trajectory in real competition conditions. Our approach, we refer to as *SkiTraVis*, enables the analysis of a trajectory in a qualitative way, meaning that the trajectory can be visualized in relation to the visual appearance of the context appearing in the video. *SkiTraVis* is inputted with just a single video with sufficient contextual details and works in a streaming fashion, i.e. in each video frame it outputs the perspective-correct trajectory exploiting the pixel information of the frame and the one occurring in the previous frames. Our solution decouples the skier’s and the camera’s motions between consecutive frames through a visual object tracker [67] and an image key-point detection [53] and tracking algorithm [39]. The corresponding key-points are used by RANSAC [23] to estimate the

perspective transformation between the frames, which in turn is used to map the previous motion of the skier to the camera’s changed perspective. We evaluated *SkiTraVis* on multiple clips taken from broadcasting videos capturing skier performance in real-world conditions. The qualitative and quantitative results achieved suggest the applicability of the method for augmented performance review and analysis during training sessions, and also for enriched broadcasting content. *SkiTraVis* works not only for alpine skiing but even for snowboarding and for flight trajectory visualization in ski jumping and freestyle skiing.

## 2. Related Work

### 2.1. Computer Vision Applications in Skiing

Thanks to the advancements in computer vision methodologies [9, 26, 34, 46], recently several applications have been possible for vision-based performance analysis in skiing. In [56], a dataset was proposed to evaluate object detection and tracking algorithms to localize recreational alpine skiers in images and videos. The work described in [68] evaluated object detection and human pose estimation algorithms to recognize falls of alpine skiers. On a similar topic, it has been discussed about the combination of 2D pose detectors with kinematics models to distinguish normal skiing situations from out-of-balance and fall ones in pictures capturing competitive alpine skiers [69]. Štepec et al. designed an algorithm based on 2D human pose trajectories to score the style of jumps in ski jumping from a monocular video [62]. Concerning the same skiing discipline, Ludwig et al. designed improved vision transformer architectures [37] to detect arbitrary key points on the human body [40] and on the skis [41] in still frames of the jumps. For freestyle skiing and snowboarding, Wang et al. developed a systematic approach composed of a visual tracker, a human pose estimator, and a pose classifier, to evaluate the quality of aerial jumps in monocular videos [65]. In [43], research focused on the design of an algorithm to synchronize, spatially and temporally, two videos capturing snowboarders to compare the timing and spatial extent of their maneuvers.

The proposed *SkiTraVis* shares some similarities with [62, 65] in the systematic design based on a visual object tracker [16, 17, 19, 67] to model the motion of a skier. However, differently from all the other works, here we tackle a different task that, to the best of our knowledge, no other solution addressed before.

### 2.2. Vision-based Trajectory Analysis in Sports

Computer vision-based techniques have been successfully used in different sports disciplines to reconstruct the trajectory of various elements [8, 10, 11, 30, 42]. Chen et al. developed an algorithm to compute the 3D trajectory of a volleyball ball in a static-camera video by transforming

the output of the ball’s visual tracker with prior knowledge of the volleyball court [10]. In similar capturing settings, Calandre et al. reconstructed the trajectory of table-tennis balls with a visual detection and tracking module, a registration between image and table coordinates, and a kinematic ball motion model [8]. In [30], the tracking-by-deblatting paradigm was introduced for better localize of – hence compute better trajectories – fast-moving objects such as table-tennis or tennis balls in static camera videos. In [42], Mixed Integer Programming was used to track, under motion blur conditions, the position of volleyball, basketball, and soccer balls, with the goal of computing their motion trajectories.

In basketball analytics, different solutions have been proposed to reconstruct the trajectories of moving players [11, 28, 38, 45]. All these research endeavors presented a pipeline composed of a visual tracking algorithm to get player-specific localizations across video frames, and the application of a homography transformation between frame and court coordinates. Their differences lie in the tracking principles used, such as histogram-based condensation [45], Cam-Shift [28], deformable part models and kalman filters [38], or deep learning methods [11], and in the homography estimation, based either on manual annotation [45] or on the matching of court lines and boundaries [11, 28, 38]. Similar approaches have been also exploited in computer vision-based applications for soccer [27] and ice-hockey [61].

The methodology described in this paper is similar to [11, 28, 38, 45] in the design of a system to track an athlete (in our case a skier), and to compute key-point correspondence from which to estimate a homography. Differently from such solutions, we use a more recent transformer-based approach [67] to implement the athlete tracking module, and we do not compute key-point matches based on some form of prior knowledge of the playing field. This is because alpine skiing is performed on a course that not only varies across competition locations but even across the different sections of the same mountain slope (e.g. in the width, steepness, and snow appearance). Hence, it is difficult to define a model of the playing field as in soccer, basketball, or ice-hockey. For this motivation, in this work, we explore the usage of generic key-point correspondence algorithms [39, 51, 53, 57] to compute the homography that enables the computation of the skier trajectory.

### 3. Methodology

The goal of this paper is to develop an algorithm for augmenting a video capturing the performance of an alpine skier with the visualization of the trajectory he/she skied from his/her position in the first frame up to his/her position in each other frame. Qualitative examples of our intention are presented in Figure 1.

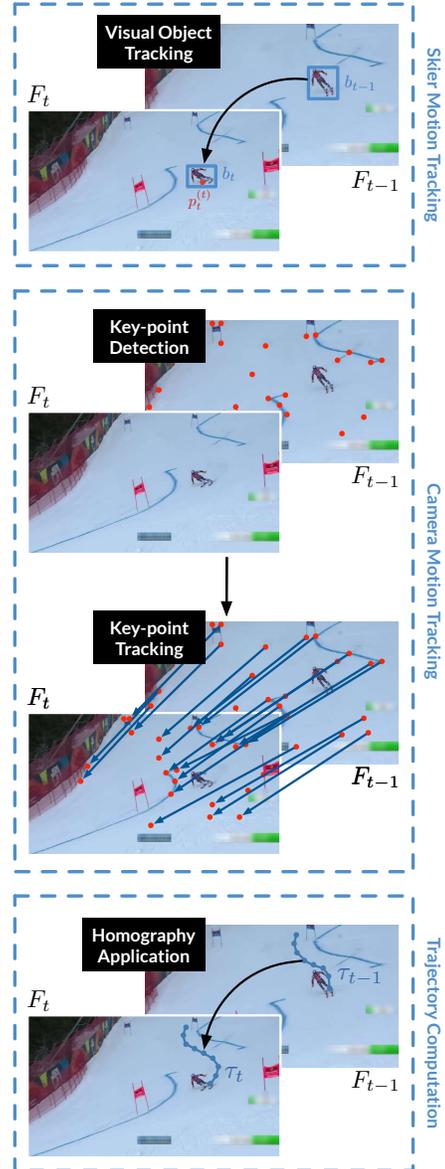


Figure 2. **Pipeline of SkiTraVis.** The algorithm processes each frame  $F_t$  in an online fashion. A visual tracker is run to model the skier’s motion between the previous frame  $F_{t-1}$  and  $F_t$ . Detection and tracking of **static key-points** in the frames are run to estimate the camera’s motion. Then, the homography between the two frames is computed and used to transform the previous trajectory  $\tau_{t-1}$  into  $\tau_t$  according to the perspective change.

#### 3.1. Preliminaries

The videos given as input to our SkiTraVis are considered to capture the performance of an individual skier while he/she is constantly visible (at least partially) in the scene. We do not put any constraints on the configuration (intrinsic and extrinsic parameters) of the camera that captured the videos, if not the presence of a sufficient amount of visual features in the scene (more details about this will

be given in Section 4). Formally, we consider a video  $\mathcal{V} = \{F_t \in \mathcal{I}\}_{t=0}^T$  as a sequence of  $T \in \mathbb{N}$  frames  $F_t$ , where  $\mathcal{I} = \{0, \dots, 255\}^{w \times h \times 3}$  is the space of RGB images. We use  $p_t^{(t)} = (x_t^{(t)}, y_t^{(t)})$  to denote the coordinates of the point that summarizes the position of the athlete in the image coordinate system of  $F_t$  (e.g. the point of contact between the skier and the snow surface).

**Problem Formulation.** The goal is to compute, at each frame  $F_t$ , a trajectory

$$\tau_t = \{p_i^{(t)}\}_{i=0}^t, \quad (1)$$

which is the sequence of points  $p_i^{(t)}$  traversed by the athlete in the 2D space of previous frames, from  $F_0$  up to  $F_t$ . To respect an online processing modality, the algorithm is not allowed to process  $F_{\hat{t}}$ ,  $\hat{t} > t$  but just the pixel information available in  $F_t$  and eventual computations performed in preceding frames. After its computation, the points in  $\tau_t$  are used to highlight the pixels of  $F_t$ , hence showing the trajectory performed by the skier. The online setting aligns with the real-time requirements of different applications (e.g. broadcasting and coach assistance) because it does not require waiting for the athlete’s exercise to be terminated for the trajectory to be produced.

### 3.2. Pipeline of SkiTraVis

Figure 2 depicts the pipeline of the proposed SkiTraVis. The algorithm processes each  $F_t$  online according to their temporal order.  $F_t$  is first given to a visual object tracking algorithm designed to model the motion of the skier and to provide its position  $p_t^{(t)}$ . Then, the pipeline estimates the camera’s motion between the consecutive frames  $F_{t-1}, F_t$  by computing the homography transformation  $\mathbf{H}_t \in \mathbb{R}^{3 \times 3}$ . This is found by detecting and tracking static key-points present in the two frames and by using RANSAC on the corresponding key-points.  $\mathbf{H}_t$  is used to map the points of the trajectory available in the previous frame,  $\tau_{t-1}$ , into the coordinate system of  $F_t$ . After that,  $p_t^{(t)}$  is appended to such a transformed trajectory to obtain  $\tau_t$ , the list of all the points traversed by the athlete with respect to  $F_t$ ’s perspective. We now describe the different components of SkiTraVis in more detail.

**Skier Motion Tracking.** The first step of the pipeline is to exploit a visual object tracker [16, 18, 31] to track the motion of the athlete across all the frames up to  $F_t$ . For this, we used a tracker outputting a bounding-box  $b_t = (x_t^{(b)}, y_t^{(b)}, w_t^{(b)}, h_t^{(b)}) \in \mathbb{R}^4$  enclosing the athlete’s appearance at every  $F_t$ .  $x_t^{(b)}, y_t^{(b)}$  represent the coordinates of the top-left corner of the box while  $w_t^{(b)}, h_t^{(b)}$  estimate the



Figure 3. **The points forming a trajectory.** The points  $p_i^{(t)}$  that compose the trajectory  $\tau_t$  (represented here as red dots) are extracted from the bounding-box  $b_t$  (represented by the blue rectangles) outputted by the skier tracking algorithm. The points roughly denote the skiers’ feet’ position.

width and height. We consider the position of the skier  $p_t^{(t)} = (x_t^{(t)}, y_t^{(t)})$  in  $F_t$  as

$$x_t^{(t)} = x_t^{(b)} + w_t^{(b)} \cdot 0.5 \quad (2)$$

$$y_t^{(t)} = y_t^{(b)} + h_t^{(b)} \cdot k \quad (3)$$

$$k \in [0, 1]. \quad (4)$$

Given an accurate bounding-box and  $k = 0.9$ , we assume  $p_t^{(t)}$  to approximate the closest point of contact between the skier’s feet and the snow, as shown by Figure 3.<sup>1</sup> The tracker is initialized in the first frame  $F_0$  of the video with the bounding-box  $b_0$  that outlines the appearance of the target skier. Such a piece of information can be obtained by asking a human operator to provide the bounding-box for the athlete of interest via some user-friendly annotation system or by a skier-specific detection algorithm. Based on experiments, we found the state-of-the-art transformer-based single object tracker STARK [67] to provide a great balance between accuracy and efficiency. We fine-tuned the pre-trained publicly-available model executing the original training methodology on video data specifically designed for alpine skier tracking.

**Camera Motion Tracking.** To render the trajectory of the skier  $\tau_{t-1}$  into  $F_t$ , we need to compute the perspective

<sup>1</sup>More information is given in Appendix A.1 of the supplementary document.

transformation occurring between  $F_{t-1}$  and  $F_t$ . To implement this step we make use of the homography matrix  $\mathbf{H}_t$  mapping the pixel coordinates of  $F_{t-1}$  into  $F_t$ . To compute  $\mathbf{H}_t$ , we match static key-points in the fields of view of  $F_{t-1}$  and  $F_t$  through an off-the-shelf generic key-point detector and tracker. Key-points are first detected in  $F_{t-1}$ , and then the correspondences of such points are found in  $F_t$  by local tracking. This procedure computes an alignment between the same semantic points of the images that express how static objects in the scene have apparently moved between the frames. Since the objects are static and do not actually move, the camera’s motion is quantified [5, 25]. Based on experiments, we have found traditional methods [39, 53] to work the best. The method of Shi-Tomasi [53] was used to detect visual features in  $F_{t-1}$ , and the Lucas-Kanade’s optical flow method [39] to track such visual descriptors in  $F_t$ . The usage of such off-the-shelf instances is motivated by the fact that alpine skiing is performed on courses that change a lot, across competition locations and even across the different sections of the same mountain slope. Hence, it is not trivial to define a generally applicable model of the playing field, which we leave for future investigations.

**Trajectory Computation.** With the alignments of corresponding key-points, we compute the homography matrix  $\mathbf{H}_t$  through the RANSAC algorithm [23], which applies an iterative optimization procedure on the correspondences in order to find the best homography matrix that explains them. We excluded the key-points in  $F_t$  within the bounding-box  $b_t$  because they belong to the skier which is a non-static object. If present, we also discarded all the key-points lying on the superimposed virtual graphics showing the characteristics of the athlete’s performance (e.g. the running time). After  $\mathbf{H}_t$  is computed, it is used to map the points  $p_i^{(t-1)}$  of  $\tau_{t-1}$  in the new frame  $F_t$ . In more detail, at each  $F_t$ ,  $\tau_{t-1}$  consists of all the skier positions tracked in the preceding  $t - 1$  frames and transformed according to the perspective of  $F_{t-1}$ . The trajectory  $\tau_t$  for the latest  $F_t$  is obtained by the multiplication of the homogeneous coordinates of each  $p_i^{(t-1)} \in \tau_{t-1}$  with the homography matrix, such that

$$\tau_t = \{p_i^{(t)}\}_{i=0}^t, \quad (5)$$

$$p_i^{(t)} = [x_i^{(t)}, y_i^{(t)}] : 0 \leq x_i^{(t)} < w, 0 \leq y_i^{(t)} < h, \quad (6)$$

$$[x_i^{(t)}, y_i^{(t)}, 1] = [x_i^{(t-1)}, y_i^{(t-1)}, 1]^T \cdot \mathbf{H}_t, \quad (7)$$

$$[x_i^{(t-1)}, y_i^{(t-1)}] = p_i^{(t-1)}. \quad (8)$$

As can be noticed by Eq. (5),  $\tau_t$  also includes the point  $p_t^{(t)}$  extracted by the tracker’s bounding-box for  $F_t$ . This is achieved by appending such a point to  $\tau_t$  after the application of  $\mathbf{H}_t$  to  $\tau_{t-1}$ . Eq. (6) specifies that only the

transformed points falling inside the  $F_t$ ’s image coordinate system are retained in  $\tau_t$ .

More concretely, at the first time step in which the frame matching procedure is executed, i.e.  $t = 1$ ,  $\tau_0 = \{p_0^{(0)}\}$  consists only of the point extracted by the bounding-box (detected or manually labeled) highlighting the target skier in  $F_0$ . To obtain the transformation of  $p_0^{(0)}$  in  $F_1$ , its homogeneous representation is multiplied by  $\mathbf{H}_1$  to compute the homogeneous vector  $[x_0^{(1)}, y_0^{(1)}, 1]$ , from which  $p_0^{(1)}$  is extracted.  $\tau_1 = \{p_0^{(1)}, p_1^{(1)}\}$  is obtained by appending  $p_1^{(1)}$ . The process is repeated until the end of the video.

## 4. Experimental Settings<sup>2</sup>

### 4.1. Data and Annotations

To determine the quality of SkiTraVis’s trajectories, we make use of a dedicated evaluation dataset representing real-world application scenarios. The videos belonging to this dataset appear in the test-set of SkiTB, a video dataset designed for the development of athlete tracking methods in different skiing disciplines.<sup>3</sup> Such a dataset features 100 multi-camera broadcasting videos (60 videos for training and 40 for testing) that capture completely, from the start to the finish line, the performance of professional alpine skiers. Each frame of each video is manually annotated with a bounding-box containing the appearance of the athlete’s body and equipment. Manually annotating the trajectories performed by skiers in each frame of the video clips is a challenging task other than very time-consuming. Indeed, it is hard to visually identify the specific points at which the skier passed since the whitish appearance of the snow’s surface does not show easily identifiable reference points to localize them. We hence exploited an automatic annotation procedure followed by manual verification. We executed the state-of-the-art LOFTR image matching algorithm [57] to detect and match key-points of consecutive frames of SkiTB’s monocular test videos. RANSAC [23] was then used to compute homographies on the matches. The homographies were used to map the skier localization points extracted from the manual bounding-box annotations into trajectories for each frame, in a similar fashion as described for SkiTraVis. The videos with the generated trajectories have been visually assessed and selected by our team, after being instructed by two professional alpine skiing coaches on meaningful trajectory visualizations. Overall, this procedure allowed us to obtain 233 monocular clips whose characteristics are reported in Table 1.

Leveraging the proposed semi-automatic annotation protocol limits the validity of the proposed SkiTraVis to the videos whose frames are successfully processed by LOFTR

<sup>2</sup>Detailed information is provided in the Appendix B.

<sup>3</sup>SkiTB is developed by our research team and it will be released soon with a separate publication.

Table 1. **Statistics of the dataset used to evaluate** SkiTraVis. The first block of rows gives details about the characteristics of the video clips, while the second provides information about the featured real-world application scenarios.

# monocular video clips	233
# frames	25783
Video clip length in # frames (min, avg, max)	30, 111, 354
Video clip length in seconds (min, avg, max)	1.0, 3.7, 11.8
Resolution, frame-rate	720p, 30 FPS
$\tau_t$ length in # points (min, avg, max)	1, 25, 229
<hr/>	
# clips for alpine skiing sub-discipline (SL, GS, SG, DH)	35, 55, 71, 72
# athletes (men, women)	23 (13, 10)
# locations	29
# weather conditions (clear, overcast, fog, snowing)	200, 30, 2, 1

[57]. We tried to quantify the conditions in which such an algorithm works by analyzing the difference between used and discarded video clips. As can be noticed by the width of the RGB histograms in Figure 4, the suitable videos include frames with a generally larger image contrast. We also evaluated the number of features per-frame discovered by the Shi-Tomasi algorithm [53], resulting in an average number of 1430 (median 1219) for suitable clips and 1126 (median 811) for unsuitable ones. These differences are mostly due to the appearance of winter scenarios in images, which provide whitish textures containing few visual anchors that feature detection algorithms can exploit.

## 4.2. Evaluation Protocol

The protocol used to execute SkiTraVis for evaluation is similar to the One-Pass Evaluation for visual tracking [66]. The algorithm is run on each video separately, and it is given the first frame and the respective ground-truth bounding-box to initialize the visual tracker on the skier. Then, SkiTraVis is inputted with each of the remaining frames  $F_t$  iteratively, it processes the frames according to the described methodology and returns each  $\tau_t$ . To compute quantitative results, each  $\tau_t$  is compared with the respective reference trajectory  $\sigma_t$ , created as described in the previous sub-section. In addition, we evaluate the bounding-boxes  $b_t$  predicted by the visual tracker and the homographies  $\mathbf{H}_t$  generated after the frame matching step, in order to understand the impact of different implementations.

## 4.3. Evaluation Measures

To quantify the distance between  $\tau_t$  and  $\sigma_t$ , we compute the average Euclidean pixel distance between the points of the reference  $\sigma_t$  with the respective, ordered based on the insertion time step, of  $\tau_t$ . We refer to this measure as Mean Per Point Trajectory Error (MPPTe  $\downarrow$ ). In addition, we report the pixel values of Dynamic Time Warping (DTW  $\downarrow$ ) [50], a popular measure in trajectory analysis. To evaluate the skier tracking, we employ the Area Under the Curve (AUC  $\uparrow$ ) of the success plot which is equivalent to the average bounding-box overlap [66]. For the

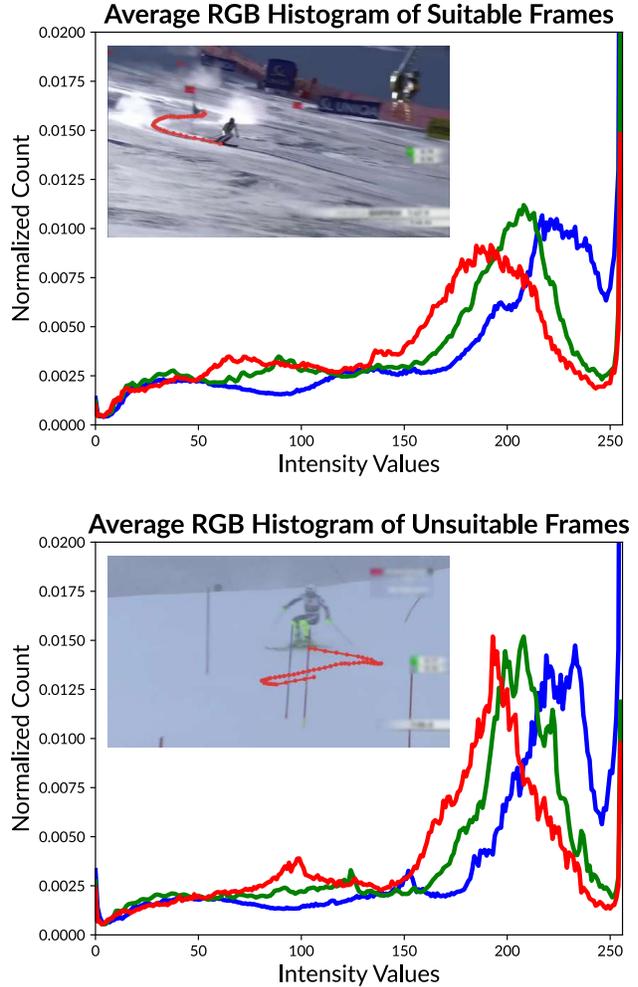


Figure 4. **RGB color distributions for the applicability of** SkiTraVis. The top plot shows the average RGB values distribution and an example (with the **generated reference trajectory**) of the frames belonging to the evaluation set used for SkiTraVis, while the bottom one shows the RGB values distributions and an example (with a **wrongly generated trajectory**) of unsuitable frames. As can be noticed, suitable clips include frames that have a larger image contrast.

frame matching module, we evaluate the mean squared error (MSE  $\downarrow$ ) between the estimated homographies and the respective LOFTR-validated ones.

Besides the error, we also rated the efficiency of SkiTraVis under different configurations of visual trackers and frame matching methods. For both categories, we report the average module-specific running speed in milliseconds (ms  $\downarrow$ ) and in Frames-Per-Second (FPS  $\uparrow$ ), and the average delay in seconds (Delay  $\downarrow$ ) between the time instant of the video clip's end and the end instant of the whole SkiTraVis pipeline's processing time. This value measures the average overall delay in visualizing the trajectory for the entire video clip, and it can be used to understand

Table 2. *SkiTraVis*'s performance with different visual trackers for skier motion modeling. The STARK tracker [67] fine-tuned for skier tracking presents great accuracy coupled with a good running time. In all of these experiments, the ST + LK key-point detector and tracker were used.

Tracker	MPPTE ↓	DTW ↓	AUC ↑	ms ↓	FPS ↑	Delay ↓
Oracle	9.8	278	100	-	-	-
STARK-ft	12.1	313	90.1	29	34	6
STARK	19.3	484	73.2	29	34	6
SuperDiMP	21.2	479	67.6	89	12	11
SiamRPN++	31.6	791	62.8	20	49	5
MOSSE	76.1	1075	37.0	9	117	3

how much time has to be waited to obtain the trajectory after the observation of the skier's performance.

#### 4.4. Evaluated Implementations

To the best of our knowledge, *SkiTraVis* is the first algorithm for skier trajectory visualization in monocular videos, and there is no comparable solution present in the literature. Thus, to provide useful insights, we evaluate the performance of our algorithm under different configurations of state-of-the-art visual trackers and frame correspondence algorithms. For the visual tracker, we evaluated: the deep learning-based methodologies STARK [67], SuperDiMP [2, 32], SiamRPN++ [36], trained for generic object tracking; and the traditional method MOSSE [3]. We also fine-tuned STARK with the aforementioned *SkiTB* training data. For key-point detection and matching, we evaluated: the learning-based method of SuperPoint (SP) [13] and SuperGlue (SG) [51], and the Recurrent All-pairs Field Transform (RAFT) dense optical flow method [59]; and the traditional methods as the Shi-Tomasi (ST) [53] detector with the optical flow algorithm Lucas-Kanade [39], and the Oriented FAST and Rotated BRIEF (ORB) descriptor [49] with the brute force (BF) matcher. Code was implemented in Python and run on a machine with an Intel Xeon E5-2690 v4 @ 2.60GHz CPU, 320 GB of RAM, and an NVIDIA TITAN V GPU.

## 5. Discussion

In this section, we discuss the results achieved by *SkiTraVis*. Overall, by looking at the numbers in Table 2 and 3, we observe that the configuration with the fine-tuned STARK-ft tracker and the ST + LK key-point detector and tracker has the best balance between trajectory error and processing speed. A visualization of the MPPTE ↓ and DTW ↓ errors, as presented in Figure 5 demonstrate that MPPTE ↓ and DTW ↓ values lower than 15 and 400 pixels lead to a consistent perception of the trajectory. Some demonstrative videos are available at [tinyurl.com/2raemvf5](http://tinyurl.com/2raemvf5). The MPPTE ↓ score is less than the 10% of the average skier appearance's height which is 123 pixels (computed from the manually annotated

Table 3. *SkiTraVis*'s performance with different frame matching algorithms. The ST + LK [39, 53] combination shows the greatest balance between accuracy and processing speed. In all of these experiments, the STARK-ft visual tracker was used.

Detection & Matching	MPPTE ↓	DTW ↓	MSE ↓	ms ↓	FPS ↑	Delay ↓
LOFTR	4.8	101	0	801	1	90
ST + LK	12.1	313	9952	42	24	6
SP + SG	15.1	352	25.7	135	7	15
ORB + BF	28.7	836	111732	783	1	84
RAFT	39.8	807	22	641	2	73

bounding-boxes). The running time of the STARK-ft ST + LK configuration has an average Delay ↓ of 6 seconds which is promising for usage in broadcasting applications (e.g. for replays) and during training activities (as the real-time requirement is not so strict).

**Skier Motion Tracking.** Table 2 reports the accuracy of *SkiTraVis* when different visual trackers are employed. Of the models trained for generic object tracking (shown in rows 3-6), STARK results in the best performing under the MPPTE ↓ and AUC ↑ measures, while SuperDiMP achieves the lowest DTW ↓ score. SiamRPN++ and MOSSE fall behind in all the scores, apart from the processing speed where they result better. Considering this aspect, it is worth noticing that STARK presents a great balance between accuracy and efficiency. For this motivation, we fine-tuned it (as STARK-ft, second row of Table 2) with the alpine skier-specific training-set of *SkiTB*. Such a version reduces the MPPTE ↓ and DTW ↓ of 37% and 35%, respectively, and presents an MPPTE ↓ error higher only of 2.3 pixels with respect to the trajectories extracted by the manually labeled bounding-boxes (denoted as Oracle, first row).

**Camera Motion Tracking.** Table 3 presents the performance of *SkiTraVis* under different frame matching strategies. The ST detector and the LK optical-flow method (ST + LK, row two) results in the closest to LOFTR in the quality of the trajectories, and the best in the computational efficiency. The SG strategy (SP + SG, row three) follows shortly in MPPTE ↓ and DTW ↓, while it is much slower in the computation. We observe a good trajectory computation by ST + LK even though its computed homographies present a larger MSE ↓ to those of the latter strategy. The remaining tested methods, ORB + BF and RAFT, show a much larger error in all the considered measures. Table 3 also reports the results achieved with the LOFTR method (row one) used to generate the reference trajectories. When run in this configuration with STARK-ft as the skier tracker, the MPPTE ↓ and DTW ↓ decrease by 60% and 68% respectively. By comparing these improvements with those given by the Oracle tracker, we can conclude that most of the *SkiTraVis*'s error is due to the frame matching module. Due to its good MPPTE ↓ and DTW ↓ values, as well as

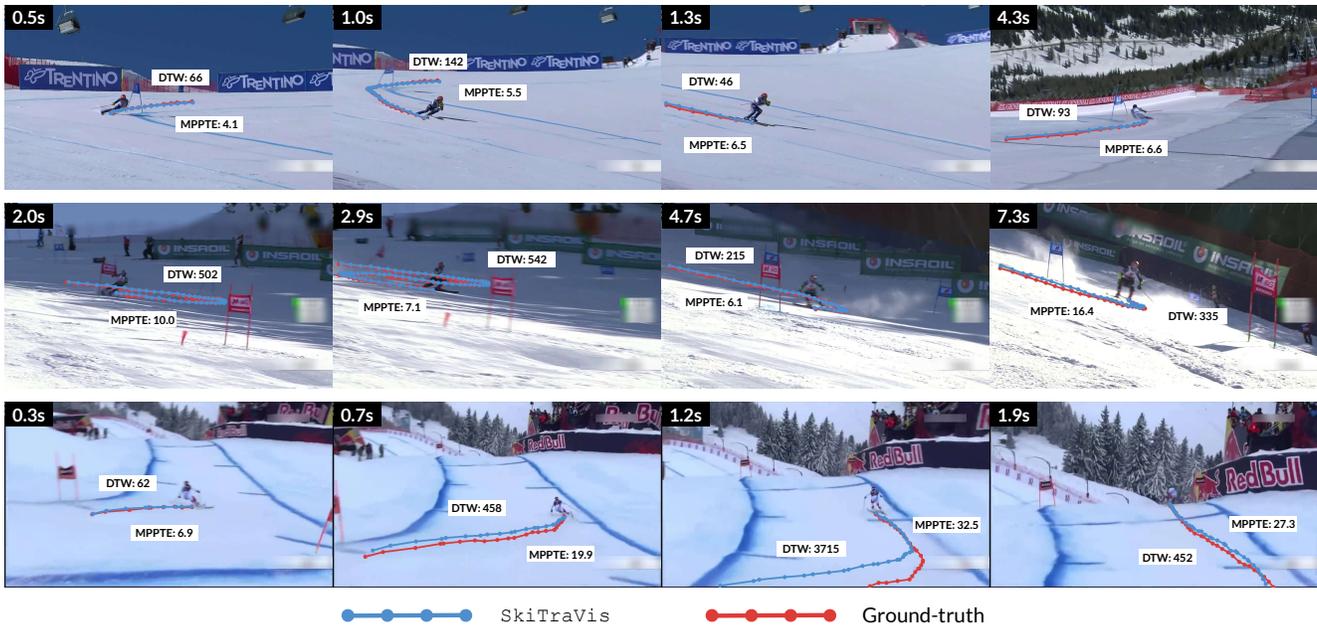


Figure 5. **Visualizations of trajectory errors.** Here we provide visual demonstrations of how the  $MPPE \downarrow$  and  $DTW \downarrow$  values represent the quality of the [trajectory](#) visualizations. The first two rows show frames where the reconstruction is visually consistent, while the last row presents video frames with a large error.

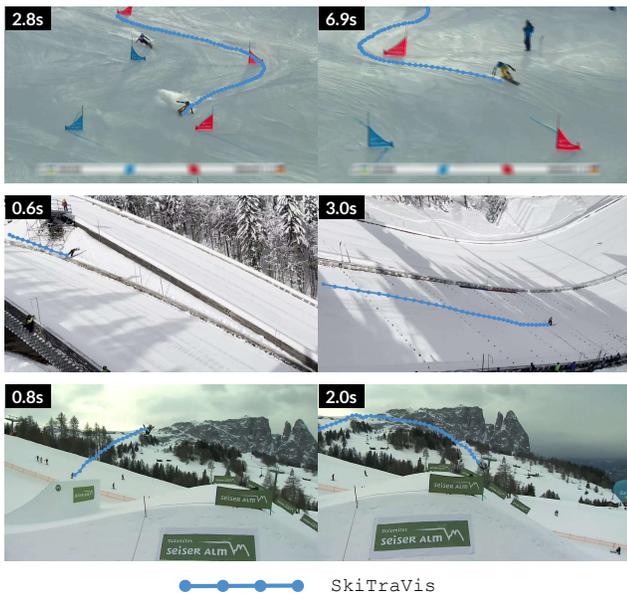


Figure 6. **Generalization to other skiing disciplines.** We found [SkiTraVis](#) to work well for the generation of [trajectories](#) in other skiing disciplines such as snowboarding, and for the visualization of the [flight trajectory](#) in ski jumping and freestyle skiing.

considering the discussion in Section 4.1, the results of the STARK-ft and LOFTR configuration demonstrate its applicability in practice, but when the processing time is not an issue since LOFTR is the slowest frame matching strategy.

**Generalization.** Our algorithm generalizes to other skiing disciplines. Figure 6 shows qualitative examples of [SkiTraVis](#) applied to the generation of the trajectory performed by alpine snowboarders, and to the visualization of the flight trajectory in ski jumping and freestyle skiing.

## 6. Conclusions

In this paper, we present [SkiTraVis](#), an algorithm to visualize the trajectories performed on an alpine skier in a monocular video. The solution uses a skier-specific visual tracking algorithm to model its motion, and an image-generic key-point-based frame matching algorithm to estimate the camera motion and the consequent transformation between the video frames. Thanks to the separation of the two motions, it is possible to relocate the previous skier positions according to the captured perspective, hence obtaining a visualization of the trajectory performed by the athlete. The experimental campaign focused on analyzing the algorithm’s accuracy and efficiency, as well as the conditions of applicability. Under some circumstances, we believe that [SkiTraVis](#) can be used in broadcasting and coach assistive applications, and more generally, it will serve as a baseline for future research on trajectory visualization for computer vision-based skiing analytics.

**Acknowledgments.** Research supported by the project between the University of Udine and the organizing committee of EYOF 2023 Friuli-Venezia Giulia.

## References

- [1] Roman Bachmann, Jörg Spörri, Pascal Fua, and Helge Rhodin. Motion capture from pan-tilt cameras with unknown orientation. In *2019 International Conference on 3D Vision (3DV)*, pages 308–317. IEEE, 2019. [2](#), [12](#), [13](#)
- [2] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6182–6191, 2019. [7](#), [14](#)
- [3] David S Bolme, J Ross Beveridge, Bruce A Draper, and Yui Man Lui. Visual object tracking using adaptive correlation filters. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 2544–2550. IEEE, 2010. [7](#), [14](#)
- [4] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000. [14](#)
- [5] Matthew Brown and David G Lowe. Automatic panoramic image stitching using invariant features. *International journal of computer vision*, 74:59–73, 2007. [5](#)
- [6] Congying Cai and Xiaolan Yao. Dynamic analysis and trajectory optimization for the nonlinear ski-skier system. *Control Engineering Practice*, 114:104868, 2021. [1](#)
- [7] Cong-ying Cai and Xiao-lan Yao. Trajectory optimization with constraints for alpine skiers based on multi-phase nonlinear optimal control. *Frontiers of Information Technology & Electronic Engineering*, 21(10):1521–1534, 2020. [1](#)
- [8] Jordan Calandre, Renaud Péteri, Laurent Mascariilla, and Benoit Tremblais. Table tennis ball kinematic parameters estimation from non-intrusive single-view videos. In *2021 International Conference on Content-Based Multimedia Indexing (CBMI)*, pages 1–6, 2021. [2](#), [3](#)
- [9] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. [2](#), [12](#)
- [10] Hua-Tsung Chen, Chien-Li Chou, Wen-Jiin Tsai, and Suh-Yin Lee. 3d ball trajectory reconstruction from single-camera sports video for free viewpoint virtual replay. In *2011 Visual Communications and Image Processing (VCIP)*, pages 1–4, 2011. [2](#), [3](#)
- [11] Liang-Hua Chen, Chih-Wen Su, and Hsiang-An Hsiao. Player trajectory reconstruction for tactical analysis. *Multimedia Tools and Applications*, 77(23):30475–30486, 2018. [2](#), [3](#)
- [12] The Nielsen Company. Audi FIS Ski World Cup 2021/22 - TV Media Evaluation Event Summary, 2022. [1](#)
- [13] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 224–236, 2018. [7](#), [14](#)
- [14] Matteo Dunnhofer, Antonino Furnari, Giovanni Maria Farinella, and Christian Micheloni. Is first person vision challenging for object tracking? In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 2698–2710, 2021. [12](#)
- [15] Matteo Dunnhofer, Antonino Furnari, Giovanni Maria Farinella, and Christian Micheloni. Visual object tracking in first person vision. *International Journal of Computer Vision*, 131(1):259–283, 2023. [12](#)
- [16] Matteo Dunnhofer and Christian Micheloni Machine. Colocolot: Combining complementary trackers in long-term visual tracking. In *2022 26th International Conference on Pattern Recognition (ICPR)*, pages 5132–5139. IEEE, 2022. [2](#), [4](#)
- [17] Matteo Dunnhofer, Niki Martinel, Gian Luca Foresti, and Christian Micheloni. Visual Tracking by means of Deep Reinforcement Learning and an Expert Demonstrator. In *Proceedings of The IEEE/CVF International Conference on Computer Vision Workshops*, 2019. [2](#)
- [18] Matteo Dunnhofer, Niki Martinel, and Christian Micheloni. Tracking-by-Trackers with a Distilled and Reinforced Model. In *Asian Conference on Computer Vision*, 2020. [4](#)
- [19] Matteo Dunnhofer, Niki Martinel, and Christian Micheloni. Weakly-supervised domain adaptation of deep regression trackers via reinforced knowledge distillation. *IEEE Robotics and Automation Letters*, 6(3):5016–5023, 2021. [2](#)
- [20] Heng Fan, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Hexin Bai, Yong Xu, Chunyuan Liao, and Haibin Ling. Lasot: A high-quality benchmark for large-scale single object tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5374–5383, 2019. [12](#)
- [21] Hao-Shu Fang, Jiefeng Li, Hongyang Tang, Chao Xu, Haoyi Zhu, Yuliang Xiu, Yong-Lu Li, and Cewu Lu. Alphapose: Whole-body regional multi-person pose estimation and tracking in real-time. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [12](#)
- [22] fengyang95. pyCFTrackers, 2019. [14](#)
- [23] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. [2](#), [5](#), [14](#)
- [24] Matthias Gilgien, Jörg Spörri, Julien Chardonens, Josef Kröll, and Erich Müller. Determination of external forces in alpine skiing using a differential global navigation satellite system. *Sensors*, 13(8):9821–9835, 2013. [2](#)
- [25] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, New York, NY, USA, 2 edition, 2003. [5](#)
- [26] Kaiping He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#)
- [27] Yutaro Honda, Rei Kawakami, Ryota Yoshihashi, Kenta Kato, and Takeshi Naemura. Pass receiver prediction in soccer using video and players' trajectories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3503–3512, June 2022. [3](#)
- [28] Min-Chun Hu, Ming-Hsiu Chang, Ja-Ling Wu, and Lin Chi. Robust camera calibration and player tracking in broadcast basketball video. *IEEE Transactions on Multimedia*, 13(2):266–279, 2011. [3](#)

- [29] Glenn Jocher, Ayush Chaurasia, Alex Stoken, Jirka Borovec, NanoCode012, Yonghye Kwon, Kalen Michael, TaoXie, Jiacong Fang, imyhxy, Lorna, Zeng Yifu, Colin Wong, Abhiram V, Diego Montes, Zhiqiang Wang, Cristi Fati, Je-bastin Nadar, Laughing, UnglvKitDe, Victor Sonck, tkianai, yxNONG, Piotr Skalski, Adam Hogan, Dhruv Nair, Max Strobel, and Mrinal Jain. ultralytics/yolov5: v7.0 - YOLOv5 SOTA Realtime Instance Segmentation, Nov. 2022. [14](#)
- [30] Jan Kotera, Denys Rozumnyi, Filip Sroubek, and Jiri Matas. Intra-frame object tracking by deblatting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. [2](#), [3](#)
- [31] Matej Kristan, Aleš Leonardis, Jiří Matas, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Hyung Jin Chang, Martin Danelljan, Luka Čehovin Zajc, Alan Lukežič, Ondrej Drbohlav, et al. The tenth visual object tracking vot2022 challenge results. In *Computer Vision – ECCV 2022 Workshops*, pages 431–460, Cham, 2023. Springer Nature Switzerland. [4](#), [12](#)
- [32] Matej Kristan, Aleš Leonardis, Jiří Matas, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Martin Danelljan, Luka Čehovin Zajc, Alan Lukežič, Ondrej Drbohlav, et al. The eighth visual object tracking vot2020 challenge results. In *Computer Vision–ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 547–601. Springer, 2020. [7](#), [14](#)
- [33] Matej Kristan, Jiří Matas, Aleš Leonardis, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Hyung Jin Chang, Martin Danelljan, Luka Čehovin, Alan Lukežič, Ondrej Drbohlav, Jani Kämpylä, Gustav Häger, Song Yan, Jinyu Yang, Zhongqun Zhang, and Gustavo Fernández. The ninth visual object tracking vot2021 challenge results. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 2711–2738, October 2021. [12](#)
- [34] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. [2](#)
- [35] Andreas Krüger and Jürgen Edelmann-Nusser. Application of a full body inertial measurement system in alpine skiing: A comparison with an optical video based system. *Journal of applied biomechanics*, 26(4):516–521, 2010. [2](#)
- [36] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4282–4291, 2019. [7](#), [14](#)
- [37] Yanjie Li, Shoukui Zhang, Zhicheng Wang, Sen Yang, Wankou Yang, Shu-Tao Xia, and Erjin Zhou. Tokenpose: Learning keypoint tokens for human pose estimation. In *Proceedings of the IEEE/CVF International conference on computer vision*, pages 11313–11322, 2021. [2](#)
- [38] Wei-Lwun Lu, Jo-Anne Ting, James J. Little, and Kevin P. Murphy. Learning to track and identify players from broadcast sports videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(7):1704–1716, 2013. [3](#)
- [39] Bruce D Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *IJCAI’81: 7th international joint conference on Artificial intelligence*, volume 2, pages 674–679, 1981. [2](#), [3](#), [5](#), [7](#), [14](#)
- [40] Katja Ludwig, Philipp Harzig, and Rainer Lienhart. Detecting arbitrary intermediate keypoints for human pose estimation with vision transformers. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, pages 663–671, January 2022. [2](#)
- [41] Katja Ludwig, Daniel Kienzle, Julian Lorenz, and Rainer Lienhart. Detecting arbitrary keypoints on limbs and skis with sparse partly correct segmentation masks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, pages 461–470, January 2023. [2](#)
- [42] Andrii Maksai, Xinchao Wang, and Pascal Fua. What players do with the ball: A physically constrained interaction modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. [2](#), [3](#)
- [43] Seiji Matsumura, Dan Mikami, Naoki Saijo, and Makio Kashino. Spatiotemporal motion synchronization for snowboard big air, 2021. [2](#)
- [44] Mirela Ostrek, Helge Rhodin, Pascal Fua, Erich Müller, and Jörg Spörri. Are existing monocular computer vision-based 3d motion capture approaches ready for deployment? a methodological study on the example of alpine skiing. *Sensors*, 19(19):4323, 2019. [2](#)
- [45] Matej Perše, Matej Kristan, Stanislav Kovačič, Goran Vučkovič, and Janez Perš. A trajectory-based analysis of coordinated team activity in a basketball game. *Computer Vision and Image Understanding*, 113(5):612–621, 2009. Computer Vision Based Analysis in Sport Environments. [3](#)
- [46] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. [2](#)
- [47] Helge Rhodin, Jörg Spörri, Isinsu Katircioglu, Victor Constantin, Frédéric Meyer, Erich Müller, Mathieu Salzmann, and Pascal Fua. Learning monocular 3d human pose estimation from multi-view images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8437–8446, 2018. [2](#)
- [48] E. Riba, D. Mishkin, D. Ponsa, E. Rublee, and G. Bradski. Kornia: an open source differentiable computer vision library for pytorch. In *Winter Conference on Applications of Computer Vision*, 2020. [14](#)
- [49] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *2011 International conference on computer vision*, pages 2564–2571. Ieee, 2011. [7](#), [14](#)
- [50] H. Sakoe and S. Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1):43–49, 1978. [6](#), [14](#)
- [51] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4938–4947, 2020. [3](#), [7](#), [14](#)

- [52] Abraham Savitzky and Marcel JE Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry*, 36(8):1627–1639, 1964. [15](#)
- [53] Jianbo Shi and Tomasi. Good features to track. In *1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 593–600, 1994. [2](#), [3](#), [5](#), [6](#), [7](#), [14](#)
- [54] International Ski and Snowboard Federation (FIS). The International Ski Competition Rules (ICR) Book IV Joint Regulations For Alpine Skiing, 2022. [1](#)
- [55] Jörg Spörri. Research dedicated to sports injury prevention—the sequence of prevention on the example of alpine ski racing. *Habilitation with Venia Docendi in Biomechanics*, 1(2):7, 2016. [2](#), [12](#), [13](#)
- [56] Philip Steinkellner and Klaus Schöffmann. Evaluation of object detection systems and video tracking in skiing videos. In *2021 International Conference on Content-Based Multimedia Indexing (CBMI)*, pages 1–6, 2021. [2](#)
- [57] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. LoFTR: Detector-free local feature matching with transformers. *CVPR*, 2021. [3](#), [5](#), [6](#), [14](#)
- [58] Matej Supej, Jörg Spörri, and Hans-Christer Holmberg. Methodological and practical considerations associated with assessment of alpine skiing performance using global navigation satellite systems. *Frontiers in Sports and Active Living*, 1:74, 2020. [2](#), [12](#)
- [59] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 402–419. Springer, 2020. [7](#)
- [60] Laurent Vanat. 2022 International Report on Snow & Mountain Tourism, 2022. [1](#)
- [61] Kanav Vats, Pascale Walters, Mehrnaz Fani, David A. Clausi, and John S. Zelek. Player tracking and identification in ice hockey. *Expert Systems with Applications*, 213:119250, 2023. [3](#)
- [62] Dejan Štepec and Danijel Škočaj. Video-based ski jump style scoring from pose trajectory. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, pages 682–690, January 2022. [2](#)
- [63] Bastian Wandt, James J Little, and Helge Rhodin. Elepose: Unsupervised 3d human pose estimation by predicting camera elevation and learning normalizing flows on 2d poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6635–6645, 2022. [2](#)
- [64] Bastian Wandt, Marco Rudolph, Petriša Zell, Helge Rhodin, and Bodo Rosenhahn. Canonpose: Self-supervised monocular 3d human pose estimation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13294–13304, 2021. [2](#)
- [65] Jianbo Wang, Kai Qiu, Houwen Peng, Jianlong Fu, and Jianke Zhu. Ai coach: Deep human pose estimation and analysis for personalized athletic training assistance. In *Proceedings of the 27th ACM international conference on multimedia*, pages 374–382, 2019. [2](#)
- [66] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Object tracking benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1834–1848, 2015. [6](#), [12](#), [14](#)
- [67] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10448–10457, October 2021. [2](#), [3](#), [4](#), [7](#), [12](#), [14](#)
- [68] Yulin Zhu and Wei Qi Yan. Ski fall detection from digital images using deep learning. In *Proceedings of the 5th International Conference on Control and Computer Vision, ICCCV '22*, page 70–78, New York, NY, USA, 2022. Association for Computing Machinery. [2](#)
- [69] Michael Zwölfer, Dieter Heinrich, Kurt Schindelwig, Bastian Wandt, Helge Rhodin, Joerg Spoerri, and Werner Nachbauer. Improved 2d keypoint detection in out-of-balance and fall situations – combining input rotations and a kinematic model, 2021. [2](#)