



# FootApp: An AI-powered system for football match annotation

Silvio Barra<sup>1</sup> · Salvatore M. Carta<sup>2</sup> · Alessandro Giuliani<sup>2</sup> · Alessia Pisu<sup>2</sup> · Alessandro Sebastian Podda<sup>2</sup>  · Daniele Riboni<sup>2</sup>

Received: 26 February 2021 / Revised: 22 March 2022 / Accepted: 3 June 2022 /  
Published online: 4 July 2022  
© The Author(s) 2022

## Abstract

In the last years, scientific and industrial research has experienced a growing interest in acquiring large annotated data sets to train artificial intelligence algorithms for tackling problems in different domains. In this context, we have observed that even the market for football data has substantially grown. The analysis of football matches relies on the annotation of both individual players' and team actions, as well as the athletic performance of players. Consequently, annotating football events at a fine-grained level is a very expensive and error-prone task. Most existing semi-automatic tools for football match annotation rely on cameras and computer vision. However, those tools fall short in capturing team dynamics and in extracting data of players who are not visible in the camera frame. To address these issues, in this manuscript we present FootApp, an AI-based system for football match annotation. First, our system relies on an advanced and mixed user interface that exploits both vocal and touch interaction. Second, the motor performance of players is captured and processed by applying machine learning algorithms to data collected from inertial sensors worn by players. Artificial intelligence techniques are then used to check the consistency of generated labels, including those regarding the physical activity of players, to automatically recognize annotation errors. Notably, we implemented a full prototype of the proposed system, performing experiments to show its effectiveness in a real-world adoption scenario.

**Keywords** Intelligent user interfaces · Pattern recognition · Artificial intelligence

## 1 Introduction

In the era of Big Data, data annotation activities are very popular, mainly for two reasons: the first is that they can provide a semantic structure to the data, like videos [12], images [19],

---

✉ Alessandro Sebastian Podda  
sebastianpodda@unica.it

<sup>1</sup> Department of Information Technology and Electric Engineering, University of Naples "Federico II", Via Claudio 21, 80125 Napoli, Italy

<sup>2</sup> Department of Mathematics and Computer Science, University of Cagliari, Via Ospedale 72, 09124 Cagliari, Italy

time-series [7], and audio file [37]. This allows the fast recovery of specific events from a large amount of information. Secondly, with the ever-increasing use of data-driven and learning-based techniques, like neural networks, deep and shallow learning approaches, the need for labelled data has become fundamental [15]. At the state of the art, there are several tools which try to fasten such activities, like LabelImg<sup>1</sup>, but these have two main drawbacks. The first is that they mainly work on images and the process of scaling the tagging over the frames of a video is quite effort expensive. The second is that the annotation involves a continuous interaction with the user, which needs to manually select the region to annotate, give it a label, save the work, and go on with the next frame. Obviously, this activity could be feasible for annotating a limited number of images, but it is not the best choice for annotating video events.

Football event annotation has the primary objective of labelling all the events happening in a football match [28]. The final goal of the activity is not only the production of a detailed report about the actions of each player (acting in its own role as defender, midfielder, forward, winger, or goalkeeper) during the match, but, in parallel, it aims at collecting team dynamics, in order to analyse forwarding, game construction, and defending tactic movements. Therefore, it is often needed to label a single action under two different points of view: (i) single player actions, expressed as *tag combinations* (e.g., MIDFIELDER → KILLER PASS → POSITIVE), and (ii) team actions (i.e., actions which involve more players, such as COUNTERATTACK). Moreover, in addition to the active plays (i.e., those whose object is the ball), there are plenty of passive plays which need to be labelled as well, in order to properly track the team dynamics, such as double teaming, defending actions, and man-to-man/zone marking. Often, this activity must be done for both teams, in order to provide a full set of tags.

It is easy to understand that the aforementioned annotation practices involve a considerable workload for those who are in charge of manual annotating the events within a football match. It is indeed estimated that the tagging activities of an entire football match (excluding additional minutes and extra times) may take from 6 to 8 hours for about 2,000 events. Annotation of soccer videos is an activity which is very popular in the market, and in the last fifteen years several systems have been proposed in this field. Most approaches are based on a computer vision based analysis over the video stream, with all the pros and cons of this solution.

However, the widespread diffusion of learning techniques has enabled a novel generation of approaches and contributions to this research field. Indeed, given the large amount of data yielded by soccer videos, together with the ease of recovering these data, the need for flexible and automatic tools for analysing, validating, testing, and hence predicting/classifying events and potential annotation of soccer videos has arisen.

In line with the previous discussions, we can summarize the motivations of our proposal with the following research questions:

- As annotating an entire match requires manually annotating thousands of events, how can the effort of annotators be lightened?
- To overcome the limits of current approaches, based mainly on video information, what methodologies can be adopted or developed to capture more complex football events?
- To provide more accurate information, how can annotation errors be detected?

---

<sup>1</sup><https://pypi.org/project/labelImg/>

In order to address these challenges, in this paper, we propose *FootAPP*, a system that aims at supporting operators in football match annotation. The user is supported by means of three main modules:

- a mixed user interface exploiting touch and voice;
- a sensor-based football activity recognition module;
- an artificial intelligence (AI) algorithm for detecting annotation errors.

The work relies on two different types of annotations. The first one is produced by experts who manually annotate the plays happened in a soccer match. The second annotation is automatically extracted from the data stream captured from sensors worn by players. Besides enriching the manual annotations with further details, the latter information is also used to detect annotation errors. In this regard, an AI-based tool is also proposed which compares the event information with those acquired by sensors in order to warn the expert operator in case of conflicting labels.

The main contributions of the paper are:

- a twofold soccer videos annotation system is proposed, which helps the synchronization between wide view events and single player activities;
- a detector for annotation errors is also developed, which helps to reveal possible conflicts between play events and players activities.

The paper is organized as follows. Section 2 reports the related work. In Section 3, an overview of the proposed system is reported, together with details on the sensor data acquisition process. In Section 4, we describe the football activity recognition process applied to wearable sensor data. Section 5 details the AI-driven technique for recognizing annotation errors. In Section 6, we describe the event tagging and sensor data infrastructure of *FootAPP*, together with the database organization, the experimental evaluation, and the achieved results. Section 7 concludes the paper.

## 2 Related work

Many studies have addressed annotation of football matches [2, 38], but most of them rely on video analysis using computer vision techniques [31, 36].

In [4] the authors have proposed a semantic annotation approach for automatically extracting significant clips from a video stream to create a set of highlights of the match. In particular, their solution detects the main highlights of the match using finite state machines that model the temporal evolution of a set of observable cues, such as ball motion or player movements. Using computer vision techniques, the visual cues are estimated, and the events detected. For an analogous purpose, in [32] a rule-based tagging system is proposed, which supports filtering for specific actions and automatic annotation of soccer events that video analysts can enhance. Another rule-based system is presented in [17]. The proposed solution is a fuzzy-based reasoning system that aims to enhance broadcasting video streams with semantic event annotations. This system is used as a classifier that takes audiovisual features in input and produces semantic concepts in output. In [5], the authors created an ontology-based system for video annotation that follows the principles of the Dynamic Pictorially Enriched Ontology model. Also, mainly in the last years, many techniques have been invented, which exploit computer graphics, thus giving birth to 3D soccer video understanding, like proposed in [27]. In this work, a tool has been introduced that allows users to perform many operations, such as displaying actions on the ball and identifying players,

while also allowing incorporation with advanced visualisation systems. Another example is proposed in [23] where a solution is illustrated that allows to automatically generate annotations and provide the spatio-temporal positions of the players on the field.

However, the widespread diffusion of learning techniques has enabled a novel generation of approaches and contributions to this research field. Indeed, given the large amount of data yielded by soccer videos, together with the ease of recovering these data, the need for flexible and automatic tools for analysing, validating, testing, and hence predicting/classifying events and potential annotation of soccer videos has arisen. Different works explored deep learning approaches, like in [30], which offers a system that focuses on identifying passes made during matches using LSTMs. In [33] CNNs were used to predict the team of players and identify who is controlling the ball. Other techniques were also experimented on, such as GANs, which were used in [25], and attention-aware based approaches were implemented in [11, 22].

All the cited papers face the annotation problem only using videos (either processed as an online or offline stream). Unfortunately, in many cases, the semantic annotation cannot be precise enough to describe a particular event, mainly due to two reasons:

- it is hard to catch information about team plays, like offsides or counterattacks actions, because such events count on the participation of more than one player, besides needing a deep knowledge about football tactics. Often indeed, the movement of a singular player makes the difference between an action rather than another;
- no information about the single movement of the player is used; this is a drawback since, without details about how the player is moving, it is not possible to measure the involvement of a player in a certain action, neither to catch information about those players which are outside the camera frame.

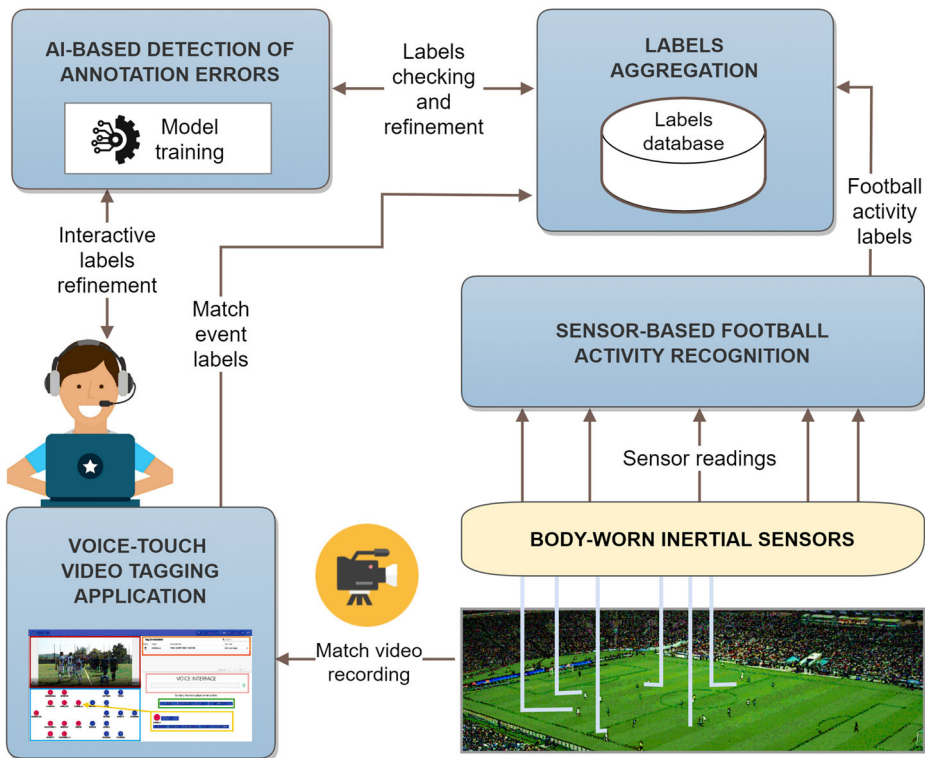
### 3 System overview

We hereafter illustrate the overall architecture of the proposed infrastructure. A high-level overview of our architecture is provided in Fig. 1.

As shown in Fig. 1, the main component of FootApp consists of the *Voice-touch Video Tagging Application* [6]. This application, briefly outlined in Section 6.1, provides a video player with the recording of the match to analyse and a user-friendly interface for vocal and touch interaction for match tagging. Through it, a human annotator deals with the labeling of football match events in FootApp. The generated labels are then stored in the *Labels database*.

In addition to the annotation activity, the involved players are required to wear inertial sensors during the game. Those sensors record all movements made by the players. Then, the *Sensor-based Football Activity Recognition* module of FootApp automatically processes offline the data collected by sensors, thus generating additional labels. The features and capabilities of this module are described in detail in Section 4. Roughly, its purpose is to recognize the athletic activity of each player (e.g., running, jumping, falling...) throughout the match, regardless of whether they are taking an active part in the development of the game.

Therefore, at the end of the match, the system sends the football activity labels to the *Labels aggregation* module. Specifically, this module owns a shared model of labels, arranged in a hierarchy, and is in charge of integrating the automatically-generated labels with those manually produced by the human annotator. The resulting integrated labels are



**Fig. 1** Overall system architecture of FootApp

then stored in the labels database, too. The latter is organized in relational tables and also includes data about football teams, players, matches, leagues, and competitions.

At the highest layer, the system features an *AI-based detection of annotation errors* module, responsible for refining the annotations by discovering – as its name reveals – potential annotation errors in stored labels. In particular, it aims at detecting a wide spectrum of errors, including wrong and missing labels. More in detail, it is trained on a large dataset of correct annotations; the trained model is hence used to check novel generated annotations. Whenever the module identifies a possible annotation error, it displays a warning. The human annotator then manually checks such a warning, which exploits the FootApp interface to inspect the video in the time interval affected by the issue and possibly fix the annotation error. Finally, annotation corrections are applied to the labels database.

In the light of the above, the described approach is not just a simple experimental proposal, but it is actually applicable to real-world contexts. In fact, the studies carried out for the realization of the prototype have involved semi-professional teams, both in training contexts and in official matches. However, many teams of the major football leagues in Europe (e.g., Liverpool FC and Manchester United FC, to name a few) nowadays adopt solutions based on inertial sensors (through appropriately engineered under-clothing, sometimes referred as *sports bras* [13]), currently to measure the performance of athletes. The solution presented in this manuscript intends to extend the use of such tools to the application domains of match analysis and match annotation, to be intended as an additional logical level that can be applied to existing solutions.

## 4 Sensor-based football activity recognition

Physical activity annotations represent an important source of information for assessing the performance of teams and players [14]. Indeed, there is a consensus in the literature on the fact that the performance of sports teams depends not only on tactical and strategic efficiency but also on the motor fitness and athletic skills of players [9]. Several studies showed that it is possible to automatically recognize physical activities by means of machine learning algorithms and sensor data [18, 35].

Moreover, in the considered scenario, information about the activity carried out by players can be matched with manually-entered labels to detect inconsistencies due to possible annotation errors. For instance, if the labels database includes a *dummy run*<sup>2</sup> for the player  $P$  at time  $t$  of the match, but the activity of  $P$  at  $t$ , acquired through the inertial sensors, was *standing still*, the module for the detection of annotation errors presented in Section 5 may issue an alert.

Hence, our system includes a specific module designated to automatically annotate the physical activities carried out by players during the match. To recognize the activities, such a module relies on wearable sensors and supervised machine learning. Figure 2 shows the information flow of our football activity recognition framework. As said, it requires players to wear inertial sensing devices on different parts of their bodies during the match. More specifically, each sensing device consists of three sensors: a tri-axial accelerometer, a gyroscope, and a magnetometer. The sensors continuously record data at a given frequency, e.g., 120 Hz. Thus, each sensor provides three streams of readings, computed on the  $x$ ,  $y$ , and  $z$  axis, respectively, for a total of nine measurements at a time. We represent a sensor reading  $r$  with the tuple  $(p, t, s, v)$ , where  $p$  is the unique identifier of the player,  $t$  is the absolute timestamp of sensor reading,  $s$  is the identifier of the sensed value (e.g., '*x-axis acceleration of the sensor worn on the left upper arm*'), and  $v$  is the sensed value. The data stream acquired from the sensors is then communicated to an *aggregation layer*, which stores the collected data in a relational database and translates the absolute time of sensor readings to the relative time of the match. We also notice that, for our application, the activity recognition stage is performed offline. Consequently, sensor data are not necessarily communicated in real-time to the aggregation layer but may be downloaded in batch mode from the sensing devices.

Since raw inertial sensor data cannot be directly used for activity recognition [10, 26], a *feature extraction* layer extracts feature vectors from the raw data by exploiting a sliding windows approach. Each resulting feature then corresponds to statistics data about some measure (e.g., '*maximum x-axis acceleration*'), computed in a single sliding window interval. From the signal produced by a sensor on a given axis during a sliding window of duration  $d$ , the extraction module computes a series of features generated from the statistical processing of the data. The features are the followings:

- *minimum* value,
- *maximum* value,
- *average* value,
- *variance*,
- *asymmetry*,
- *kurtosis*,

---

<sup>2</sup>A so-called *dummy run* occurs when a player performs an off-the-ball run to create space for his teammate with the ball.

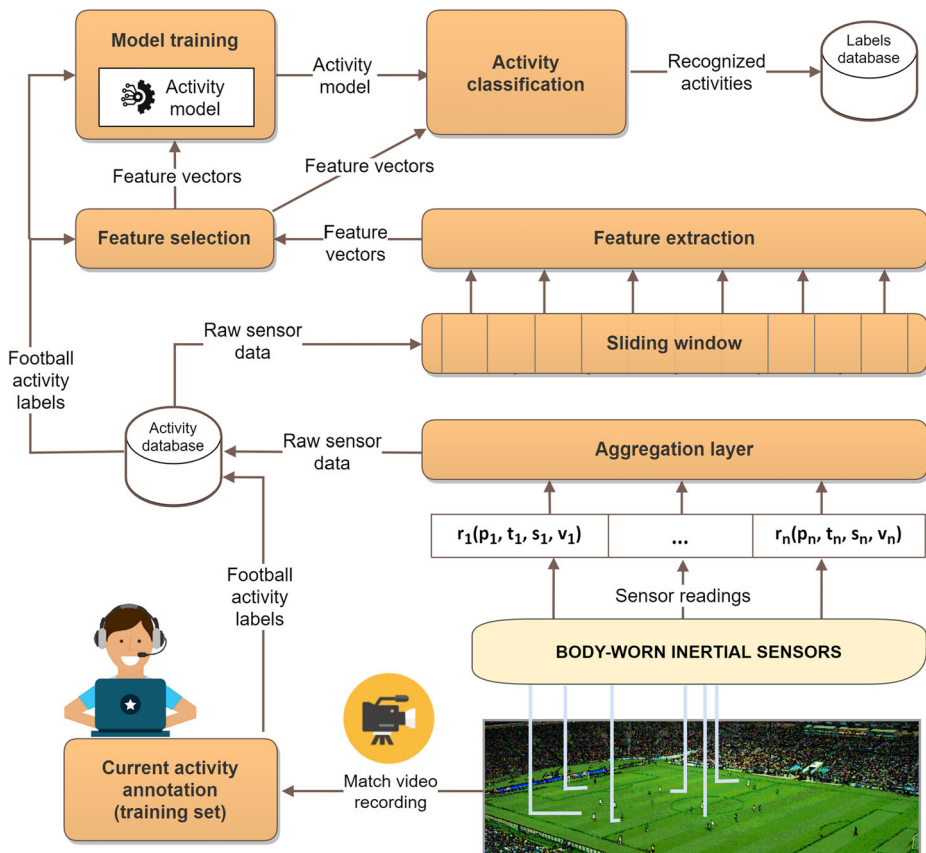


Fig. 2 The proposed football activity recognition framework

- ten equidistant samples from the *auto-correlation sequence*,
- the first five *peaks of the Discrete Fourier Transform (DFT)*, and the corresponding five *frequencies*.

Thus, globally, for each signal, the module computes 26 features. Since, as said before, each sensing device produces nine signals, the total number of features for each device is 234. However, when a player wears  $n$  sensing devices, the number of features is consequently  $234 \cdot n$ . Such a large number of features may confuse the classifier and may result in overfitting activity recognition to the training set samples [20]. To address this problem, we added to our framework an additional *feature selection layer*. This layer, that exploits a *chi-square selection*, aims to find the features that actually discriminate between activity classes. The features are then ranked according to their discriminating power and the best  $k$  are then selected. Only the selected features are used to train the classifier and recognize novel instances of activity.

For the training and activity classification processes, a supervised strategy is adopted through the use of a dataset of manually annotated activities. As explained in Section 6, in our implementation, we use a Random Forest classifier, but different machine learning algorithms could be used. The model training module uses a set of feature vectors that are

manually annotated with the current activity of the player to train a model of activities. The feature vectors correspond to those obtained from our sensors. The machine learning classifier recognizes the current activity class of novel feature vectors using the trained model. Recognized activity instances are finally stored in the labels database. Of course, the set of activities that our system can recognize depends on the activity classes considered in the training set. At the time of writing, we consider 19 activities, including walking, running, lying, jumping. However, the set of activities can be seamlessly extended using a more comprehensive training set.

## 5 AI-based detection of annotation errors

As explained before, match analysis is the task of annotating and analyzing football matches based on the observed events. Each annotation describes an event that occurs during a football match, e.g., ball possession, fouls, or goals. Each event can be characterized by the match name, a set of tags, a description label, the time when the event started, the time when it ended, the player being the active part, and his team. Moreover, players wear inertial sensors that continuously acquire data about their movements. Discovering annotation errors and reporting them to the human annotator is useful to improve the quality of the analysis. To this end, the proposed system embeds an automated error detection module based on AI algorithms, as reported in Fig. 1. As explained below, our algorithm relies on Frequent Itemset Mining (FIM) [21].

In order to identify annotation errors, our algorithm mines a dataset of correct annotation for extracting relationships between match events and activities of the players. To this end, FIM is one of the most widely used techniques to extract knowledge from data. FIM is currently used for many tasks, and it consists of extracting events, patterns, or items that frequently (or rarely) appear in the data. The most used algorithm is *Apriori* [8]. It identifies frequent individual items on datasets and extends them to larger itemsets. It is based on the Apriori paradigm, which states that if an itemset is frequent, all its subsets are frequent. Conversely, if an itemset is not frequent, also its supersets are not frequent. As FIM can extract relationships between elements, it can also be used to infer meaningful relationships between match events and activity instances.

A set of temporally close events and activities is represented as a unique entry using a sliding window on the list of annotations acquired during a match. Each entry describes a temporal segment of the match. In this way, the entire match is described as a set of entries. We assume that elements frequently occurring together in entries may be significantly correlated. Applying the Apriori algorithm allows the extraction of frequent itemsets representing which combinations of items frequently appear in close temporal proximity. The Apriori algorithm identifies single frequent items in a dataset and extends the identification to larger and larger sets of items, until the found sets appear thick enough. The sets of frequent items determined by Apriori can be used to determine association rules that highlight general trends in the dataset to which it has been applied.

Our algorithm applies Apriori to generate a set of *association rules*. An association rule describes correlations among items belonging to a frequent itemset. Each rule defines an implication, in the form of  $X \rightarrow Y$ , in which  $X$  is a subset called *antecedent itemset* and  $Y$  is a subset named *consequent itemset*. If  $X$  belongs to an entry, it is probable that also  $Y$  appears in the same entry. Each rule can be characterized using the following metrics:



– **Support**

$$\text{support}(X \rightarrow Y) = \frac{\text{number of entries containing both } X \text{ and } Y}{\text{total number of entries}}.$$

This measure is used to indicate the percentage of entries in the dataset which respect the rule.

– **Confidence**

$$\text{confidence}(X \rightarrow Y) = \frac{\text{number of entries containing } X \text{ and } Y}{\text{total number of entries containing } X}.$$

This measure estimates the likelihood that  $Y$  belongs to an entry given that  $X$  belongs to that entry.

– **Conviction**

$$\text{conviction}(X \rightarrow Y) = \frac{1 - \text{support}(Y)}{1 - \text{confidence}(X \rightarrow Y)}.$$

This measure indicates to what extent  $Y$  depends on  $X$ . If  $X$  and  $Y$  are independent, the value of conviction is 1.

By applying our algorithm, we can find significant patterns of annotations from the data. For instance, we may derive that, with high confidence, if an entry includes an event of type *pass* by a player  $p$ , the same entry also includes an activity *kicking* by  $p$ . That pattern is represented as: {Pass}  $\rightarrow$  {Kicking}. Since some of these patterns can be easily derived from domain knowledge, we complement the association rules automatically extracted from the data with other ones manually defined by domain experts.

Based on the level of support, confidence, and conviction, our system analyses new annotations and reports possible errors to the human annotator. The annotator can choose the level of sensitivity of our system by specifying thresholds for the support, confidence, and conviction values, which would issue alerts of possible errors.

## 6 Prototype and experimental evaluation

As reported in Section 1, our proposal claims to address the following research questions:

1. How can the manual effort of annotators be lightened?
2. What methodologies can be adopted or developed to capture more complex football events?
3. How can annotation errors be detected?

We have developed a full prototype of FootApp, which faces all the aforementioned questions, as described below. We have released the implemented prototype, together with the code used for experiments, on a public repository<sup>3</sup>. As shown in Fig. 1, our system takes into account two sources of information:

- *human annotators*, which watch the match and, thank the mixed touch-vocal User Interface, tags all the events happening in the scene. The mixed interface addresses question 1, as using voice or touch interfaces permits to simplify the annotation process;
- *body-worn sensors*, which acquire data on players activities, as already reported in Section 4. The integration of this data source faces question 2, proposing an innovative

<sup>3</sup><https://github.com/aleessiap/FootApp>

solution for capturing more complex events that classical video-based systems usually do not recognize. Furthermore, AI techniques are applied to the acquired data to annotate each event.

These infrastructures act as shown in Fig. 3. The match events are acquired based on the process described in [6], in which a vocal interface is exploited for fastening the annotation of all the events in the scene. More details on hardware and implementation of the integrated vocal/touch interface and of the sensor-based activity recognition module are given in Sections 6.1 and 6.2, respectively. The labeled contextual data are organized into a Relational database (see Section 6.3), which gathers and categorizes all the information. The database also contains missing event labels representing the annotation errors required for experimenting with the error detection algorithms. Furthermore, to address research question 3 (i.e., to detect annotation errors), AI algorithms are applied to the stored data (as described in Section 5).

Let us point out that all the AI algorithms are applied offline to the acquired data, and no particular computational resources (e.g., using GPUs or specific hardware) are needed. In particular, training the Random Forest model on our dataset takes less than 2 minutes using a CPU-based machine.

Experiments on both activity recognition and errors detection are reported in Section 6.4.

## 6.1 Event tagging infrastructure

The event tagging infrastructure interacts with the user by means of a twofold modality. Details are reported in our previous work [6]. The system exploits a *Voice User Interface* (VUI), integrated with a touch interface, to enhance the user experience and usability of the FootApp web application.

The underlying idea that guided its design consists in reducing the time needed for tagging matches, with a twofold advantage: i) first, the exploitation of a multi-modal interface, which combines the benefits of both voice interaction (e.g., for longer and more articulated commands) and touch interaction (for the direct player selection or click on a suggested tag, etc.); ii) a greater accessibility of the application for people who have difficulty in using traditional interfaces (mouse and keyboard), so as to choose the most suited mode to them, while maintaining a high level of efficiency in the execution of the task.

**User interface** Currently, FootApp GUI is organized as shown in Fig. 4. We describe its components in what follows:

- the top left block (contoured in red in Fig. 4) contains a video player, in which the match is reproduced; directional arrows ease the forward and backward skips of the video;
- the bottom left block (contoured in cyan) shows a virtual field in which the team lineups and shape are shown; the buttons with player numbers are clickable, to be inserted in a tag combination;
- the top right block (contoured in orange) contains a summary of the tag combination records already registered for the current match;
- the middle right block (contoured in pink) contains the voice interface, which automatically activates when the user casts a command;
- the block contoured in green shows the keyboard which activates as the user selects a player from the virtual field or spells his number (first level tags);

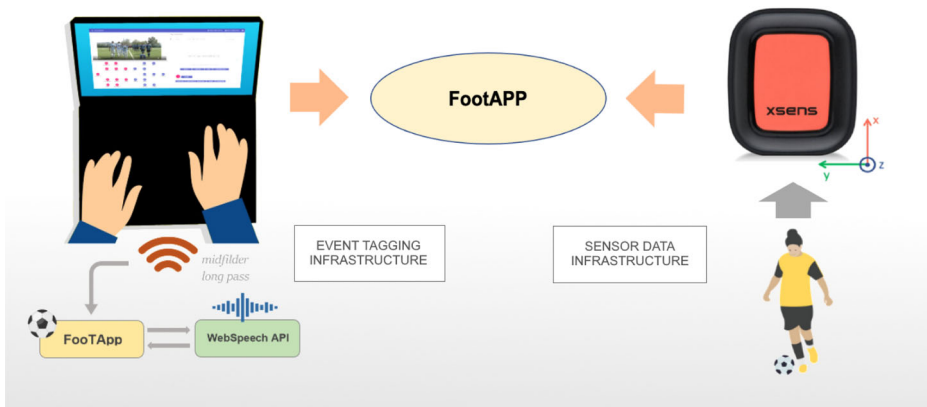


Fig. 3 Event tagging and Sensor Data infrastructures

- the block contoured in yellow contains additional details associated with the selected event. This block pops up when a main event is activated, and could potentially provide a large number of second/third level options.

**VUI integration and empirical tests** We developed the voice user interface by exploiting the Web Speech API [1], which defines a complex interface (*SpeechRecognition*) that provides a set of methods for transforming speech into text. The implemented VUI structure is shown in Fig. 4. From preliminary empirical tests conducted with professional annotators of football match events, we noticed that total migration to a VUI would have partially obtained the expected benefits. Then, we adapted the original interface to enable a combined tagging mode (voice and touch), resulting in a significant reduction for annotating a single football match, empirically estimated at around -25% than the time required by a

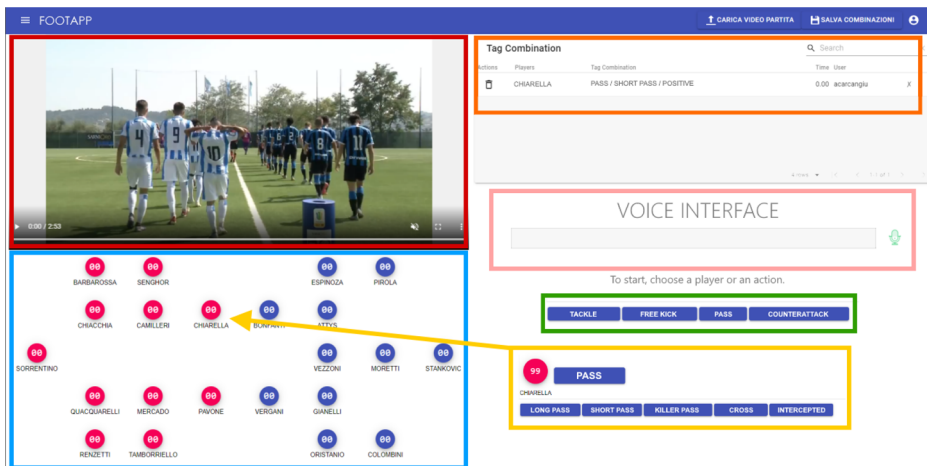


Fig. 4 The Vocal User Interface of FootAPP

traditional system. In particular, the introduction of touch functionalities, coupled with the vocal interface, had a positive effect for two main reasons. First, the processing time of long voice commands (which, despite the robustness of the exploited API, can last for a few seconds) could negatively impact the annotation performance. Second, it is quicker to execute some tagging patterns (e.g., the selection of the player from the virtual field in the bottom left block of Fig. 4) in touch modality than in vocal mode. On the contrary, other tagging practices (as an example, finding second/third level tags among many options) can obtain concrete benefits from the use of the voice interface.

## 6.2 Sensor data infrastructure

In our prototype, for the acquisition of inertial data from body-worn sensors, we used the XSSENS DOT SENSOR instruments, shown in Fig. 3. XSens Dot was chosen because it allows the use of non-invasive and lightweight sensors, which do not affect the movements of players during a match, how shown in [24], in which it is used for monitoring the cardiorespiratory activity for long periods. The sensors detect and store the following data at 120 Hz frequency:

- Orientation, computed by means the Euler's angles on the  $x$ ,  $y$  and  $z$  axes (in degrees);
- Acceleration on  $x$ ,  $y$  and  $z$  axes, both with and without gravity (in  $m/s^2$ );
- Angular speed on  $x$ ,  $y$  and  $z$  axes (in degrees/ $s^2$ );
- Magnetic field on  $x$ ,  $y$  and  $z$  axes, comprising of normalization factor.

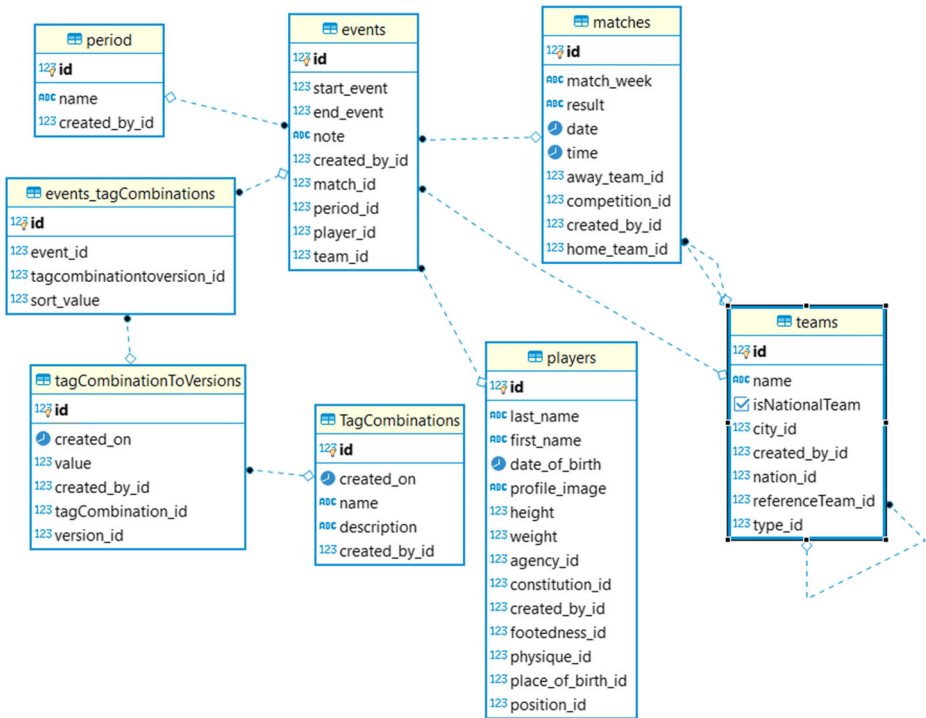
Such data are then downloaded (in offline mode) into a `csv` file by means of an application running on an Android device. A synchronization phase is needed in this phase, given the time difference between the Android application, which is synchronized to the international time, and the sensors, which count the time since the moment it is powered on. This operation is quite simple, and it is done by registering the moment in which a specific sensor is powered on, thus synchronizing this timestamp with the international measure followed by the Android application. The prototype, which is implemented in Python, follows three main steps:

- Data acquisition and feature extraction;
- Classifier training and activity prediction;
- Prediction storing into the database.

In Section 4, the 234 features extraction process is described. Given the high number of features, in order to avoid overfitting, a feature selection process takes place, in which the 30 best features are extracted. This operation is done using a *chi-squared* based selection process which removes all the features but the best  $k$ , with  $k = 30$  [16]. The selected features feed a machine learning classifier. In the current implementation, we use a Random Forest classifier. The prediction obtained for each activity is then inserted into the database of annotations (developed in PostgreSQL), which is described below.

## 6.3 Database description

The database consists of about 90 tables, which allow the management of the entire set of data, including information about players, teams, coaches, matches, and predicted activities. Part of the database diagram is shown in Fig. 5. In the table `EVENTS` all the activities of the single players are recorded. The tables contain several foreign keys which allow the recovery of the information regarding the player (`PLAYER_ID`), the match (`MATCH_ID`), the



**Fig. 5** The XSens Dot Platform used for registering players’ body movement on the pitch

half of the match (PERIOD\_ID) and the team of the player (TEAM\_ID). The single activity is described in the table TAGCOMBINATIONS which contains its name and description.

For the most common queries, prepared statements have been defined as the one in the following, in which by defining MATCH\_ID, PERIOD\_ID and PLAYER\_ID, all the events happened in a specific match, in a specific period and involving a specific player are returned in chronological order.

```

SELECT start_event, end_event, "TagCombinations".name
FROM "TagCombinations"
JOIN "tagCombinationToVersions"
ON "tagCombination_id" = "TagCombinations".id
JOIN "event_tagCombinations"
ON "tagcombinationtoversion_id" = "tagCombinationToVersion".id
JOIN events on events.id = "events_tagCombinations".event_id
JOIN players on players.id = events.players_id
WHERE match_id = ? AND player_id = ? AND period_id = ?
ORDER BY start_event
    
```

### 6.4 Experimental results

In the following, we report experimental results on our modules for activity recognition and detection of annotation errors, respectively.

### 6.4.1 Football activity recognition

In order to evaluate the effectiveness of our activity recognition method presented in Section 4, we have used a large dataset<sup>4</sup> of physical and sport activities. The dataset considers 19 activities, including walking (at different speeds), running, lying, jumping, standing. For each activity, the data have been acquired from 8 different subjects (4 males and 4 females, aged 20 to 30) for 5 minutes per subject. Hence, the total time duration of the dataset is 760 minutes. Sensor data have been acquired at 25 Hz sampling frequency. Five sensing devices were worn by each subject: one on each knee, one on each wrist, and one on the chest. Each sensing device included one tri-axial accelerometer, one tri-axial gyroscope, and one tri-axial magnetometer. The data have been acquired in natural conditions. The dataset is described in detail in [3].

We have developed the algorithms presented in Section 4 in Python, using the *sklearn* machine learning libraries for implementing the feature selection and classification functions. We have used a *leave-one-subject-out* cross-validation approach, using one subject's data for the test set and the other subjects' data for training. For the classifier, we have used two algorithms: Random Forest [34], and Support Vector Machines [29], since in the literature they are considered among the most effective classifiers for sensor-based activity recognition [10].

Results are shown in Table 1. We use the standard metrics of precision, recall, and F-score, the latter being the harmonic mean of precision and recall. Overall, our results are compliant with the best results achieved by state-of-the-art methods on the same dataset [3]. By closely inspecting the results shown in the confusion matrix, reported in Table 2 for the Random Forest classifier and in Table 3 for the Support Vector Machines classifier, respectively, we can notice that a few activity classes (in particular, A9, A10, and A11) achieve relatively values of F-score. Those activity classes represent different kinds of walking, which are difficult to distinguish based on sensor data. By merging those classes in a unique class, we could improve the overall recognition rate without disrupting the data utility for our application.

### 6.4.2 Detection of annotation errors

We have used the Python language to implement and experiment the frequent itemset mining technique explained in Section 5. The libraries that we used are *pandas*, *numpy* and *mlxtend*. The dataset of the manually annotated events is structured in the following way:

- *Episode*: an integer that identifies the event or activity in the match;
- *Match*: a string that identifies the name of the match;
- *Team*: a string that identifies the name of the team having an active part in the episode;
- *Start*: the time (in the format minutes:seconds) of the episode start;
- *End*: the time (in the format minutes:seconds) of the episode end;
- *Half*: an integer that indicates in which half of the match the episode occurred;
- *Description*: a string that describes the type of the episode;
- *Tags*: a list of strings that characterize the episode and its details;
- *Player*: a string that indicates the name of the player who had an active part in the episode;
- *Notes*: a string that contains comments or notes about the episode.

<sup>4</sup><https://archive.ics.uci.edu/ml/datasets/Daily+and+Sports+Activities>

**Table 1** Precision, recall, and F-score for each activity class using Random Forest (left) and Support Vector Machines (right)

Precision	Recall	F-score	Activity	Precision	Recall	F-score	Activity
1	1	1	A1	1	1	1	A1
0.811	0.627	0.707	A2	0.403	0.755	0.53	A2
1	0.875	0.933	A3	1	0.992	0.996	A3
0.889	1	0.941	A4	0.996	1	0.998	A4
0.699	0.915	0.792	A5	0.666	0.888	0.761	A5
0.95	0.902	0.925	A6	0.805	0.863	833	A6
0.623	0.729	0.672	A7	0.504	0.263	0.346	A7
0.747	0.795	0.77	A8	0.706	0.34	0.459	A8
0.57	0.625	0.596	A9	0.584	0.615	0.599	A9
0.57	0.631	0.599	A10	0.575	0.744	0.649	A10
0.791	0.425	0.553	A11	0.792	0.61	0.689	A11
0.996	1	0.998	A12	1	1	1	A12
0.867	0.831	0.849	A13	0.86	0.654	0.743	A13
0.768	0.746	0.757	A14	0.719	0.625	0.669	A14
0.966	1	0.983	A15	1	1	1	A15
0.998	0.965	0.981	A16	1	1	1	A16
1	1	1	A17	0.996	1	0.998	A17
1	0.965	0.982	A18	0.996	992	0.994	A18
0.851	0.931	0.89	A19	0.873	0.898	0.934	A19
0.847	0.84	0.838	Average	0.82	0.803	0.8	Average

For the sake of annotation error detection, we decided to reduce the modeling of the episode to its Description field, since it is the most representative data. Using the Apriori algorithm, we mined the frequent itemsets of episodes from a dataset of matches labeled by professional annotators. The dataset included annotations from 67 football matches: among them, 58 were matches of the Campionato Primavera 1 competition (i.e., the premier youth soccer league in Italy), while 6 were international under-17 matches and 3 involved matches that took place during the Chinese championship. Totally, the dataset included 103,926 annotations. Then, the itemsets were used to construct the association rules using the method provided by the *mlxtend* library. A domain expert analysed each rule to check whether it was representative of a real-world behaviour of football match events and activities. Other rules were manually added to specify behaviours useful in the domain but not captured by our data-driven technique. As explained in Section 5, manually added rules have a qualitative estimation of their confidence level, but lack values in the other metrics (i.e., support and conviction). Table 4 shows a sample of the mined rules.

The association rules describe the typical pattern of a game that they were able to extract from the dataset. If an event does not fulfill the rules, it is safe to assume it may be an error. For example, the first rule describes the relationship between the events *Reception*, *Construction* and *Pass*. This associative rule obtained a confidence value of more than 0.90, indicating that *Reception* and *Construction* appear in more than 90% of the events that have *Pass* as a nearby event in the time window. This rule has a support value of only

**Table 2** Confusion matrix obtained with the Random forest classifier

Activity	Classified as																			
	A1	A2	A3	A4	A5	A6	A7	A8	A9	A10	A11	A12	A13	A14	A15	A16	A17	A18	A19	
A1	480	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
A2	0	301	0	0	0	0	155	24	0	0	0	0	0	0	0	0	0	0	0	0
A3	0	0	420	60	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
A4	0	0	0	480	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
A5	0	0	0	0	439	3	0	2	0	0	5	0	21	10	0	0	0	0	0	0
A6	0	0	0	0	11	433	0	8	4	0	0	0	1	0	0	1	0	0	0	22
A7	0	54	0	0	0	0	349	76	0	0	0	0	0	0	0	0	0	0	0	0
A8	0	16	0	0	2	1	56	381	14	0	0	0	0	0	0	0	0	0	0	9
A9	0	0	0	0	29	9	0	1	300	120	11	0	3	0	0	0	0	0	0	7
A10	0	0	0	0	0	0	0	0	158	303	18	0	0	1	0	0	0	0	0	0
A11	0	0	0	0	52	0	0	0	16	105	204	0	12	91	0	0	0	0	0	0
A12	0	0	0	0	0	0	0	0	0	0	0	480	0	0	0	0	0	0	0	0
A13	0	0	0	0	51	8	0	0	7	4	1	0	399	6	0	0	0	0	0	4
A14	0	0	0	0	39	0	0	1	20	0	19	0	24	358	0	0	0	0	0	19
A15	0	0	0	0	0	0	0	0	0	0	0	0	0	0	480	0	0	0	0	0
A16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	17	463	0	0	0	0
A17	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	480	0	0	0
A18	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	463	0	17
A19	0	0	0	0	5	2	0	17	7	0	0	2	0	0	0	0	0	0	0	447



**Table 3** Confusion matrix obtained with the SVM classifier

Activity	Classified as																			
	A1	A2	A3	A4	A5	A6	A7	A8	A9	A10	A11	A12	A13	A14	A15	A16	A17	A18	A19	
A1	480	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
A2	0	372	0	0	0	0	79	29	0	0	0	0	0	0	0	0	0	0	0	0
A3	0	0	476	2	0	0	0	0	0	0	0	0	0	0	0	0	2	0	0	0
A4	0	0	0	480	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
A5	0	0	0	0	426	12	0	0	0	0	0	0	22	20	0	0	0	0	0	0
A6	0	0	0	0	30	414	0	17	17	0	0	0	1	0	0	0	0	0	0	1
A7	0	352	0	0	0	0	126	1	0	0	0	0	0	0	0	0	0	0	0	0
A8	0	200	0	0	0	43	45	163	20	0	0	0	1	0	0	0	0	0	0	7
A9	0	0	0	0	2	13	0	1	295	144	10	0	1	14	0	0	0	0	0	0
A10	0	0	0	0	0	0	0	0	121	357	2	0	0	0	0	0	0	0	0	0
A11	0	0	0	0	10	0	0	0	14	96	293	0	1	66	0	0	0	0	0	0
A12	0	0	0	0	0	0	0	0	0	0	0	480	0	0	0	0	0	0	0	0
A13	0	0	0	0	83	14	0	4	24	20	4	0	314	17	0	0	0	0	0	0
A14	0	0	0	0	87	0	0	0	7	2	61	0	23	300	0	0	0	0	0	0
A15	0	0	0	0	0	0	0	0	0	0	0	0	0	0	480	0	0	0	0	0
A16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	480	0	0	0	0
A17	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	480	0	0	0
A18	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	476	4	0
A19	0	0	0	0	2	18	0	16	2	2	0	0	2	0	0	0	0	2	0	431

**Table 4** Association rules for annotation error detection

Rule	Confidence	Support	Conviction
{Reception, Construction} → {Pass}	0.924	0.017	2.747
{Unmarking} → {Pass}	0.897	0.053	2.032
{Interception, Reception} → {Pass}	0.911	0.015	2.343
{Reception, Unmarking} → {Pass}	0.896	0.014	2.005
{Interception} → {Pass}	0.870	0.164	1.601
{Reception} → {Pass}	0.881	0.119	1.755
{Offensive moves} → {Pass}	0.877	0.041	1.699
{Lose the ball} → {Recover the ball}	0.794	0.055	4.510
{Transmission} → {Construction}	0.631	0.035	2.230
{Shielding} → {Reception}	0.519	0.011	1.799
{Cross Interception} → {Pass refinement}	0.536	0.023	1.891
{Offensive moves} → {Reception}	0.460	0.022	1.601
{Pass, Offensive moves} → {Reception}	0.462	0.019	1.607
{Shielding} → {Defensive duel}	0.448	0.010	1.184
{Cross Interception} → {Line Attack}	0.327	0.014	1.272
{Pass} → {Kicking}	High	Manually defined rule	
{Header} → {Jumping}	Medium	Manually defined rule	
{Offensive moves} → {Running}	Medium	Manually defined rule	
{Construction} → {Running}	Low	Manually defined rule	
{Shielding} → {Standing still}	Low	Manually defined rule	

0.017; this indicates that the set of itemsets containing {*Reception, Construction, Pass*} is not very frequent *per se*, this is not surprising considering the large number of episode types and the high variability of plays in football. The last metric computed is Conviction which is equal to 2.747, meaning that the two itemsets are significantly correlated.

## 7 Conclusion

In this paper, we have presented FootApp, an innovative interface for football match event annotation. FootApp relies on a mixed interface exploiting vocal and touch interaction. Thanks to the use of body-worn inertial sensors and supervised learning, our system can automatically acquire labels regarding the activity carried out by players during the match. An AI module is in charge of detecting inconsistencies in the labels, including those describing the players' activity, for recognizing possible annotation errors. We have implemented a full prototype of our system and experimented with the system. The results indicate the effectiveness of our algorithms. The current configuration shows a three-tier architecture, in which the subsystems consist respectively of a mixed interaction User Interface, an engine which achieves annotation errors detection and a Data Layer, composed of a Database and a activity recognition logic layer which gathers information from the analysed matches and fills the database. Such an architecture opens to several enhancement solutions, e.g. a cloud/grid based configuration, in which the database is located on another machine and the activity recognition task is achieved with a cloud computing paradigm. This possibility

could allow the development of a simple mobile-based front-end both for the tagging activities (in this case, a tablet -at least 11.4 inches- is suggested for obtaining a full immersion between the touch gestures and the vocal command), and for the off-line tagging recovery. In the last scenario, it would be useful to use already built-in speech recognition engines. Such enhancements are already object of study and developments within our research group.

Future work will focus on the possible implementation of a scalable cloud architecture and the refinement of error detection algorithms to more accurately capture possible errors.

**Acknowledgements** This work was partially funded by the POR FESR Sardegna 2014-2020 project “MISTER: Match Information System and Technologies for the Evaluation of the Performance”.

**Funding** Open access funding provided by Università degli Studi di Cagliari within the CRUI-CARE Agreement.

## Declarations

**Conflict of interests** The authors declare that they have no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Adorf J (2013) Web speech api. KTH Royal Institute of Technology
2. Alan O, Akpinar S, Sabuncu O, Cicekli N, Alpaslan F (2008) Ontological video annotation and querying system for soccer games. In: 2008 23rd International Symposium on Computer and Information Sciences, pp 1–6. <https://doi.org/10.1109/ISCIS.2008.4717936>
3. Altun K, Barshan B, Tunçel O (2008) Comparative study on classifying human activities with miniature inertial and magnetic sensors. *Pattern Recogn* 43(10):3605–3620
4. Assfalg J, Bertini M, Colombo C, Del Bimbo A, Nunziati W (2003) Semantic annotation of soccer videos: automatic highlights identification. *Comput Vis Image Underst* 92(2):285–305. <https://doi.org/10.1016/j.cviu.2003.06.004>. Special issue on video retrieval and summarization
5. Ballan L, Bertini M, Del Bimbo A, Serra G (2010) Semantic annotation of soccer videos by visual instance clustering and spatial/temporal reasoning in ontologies. *Multimed Tools Appl* 48(2):313–337
6. Barra S, Carcangiu A, Carta S, Podda AS, Riboni D (2020) A voice user interface for football event tagging applications. In: *Proceedings of the International Conference on Advanced Visual Interfaces, AVI '20*. Association for Computing Machinery, New York, NY. <https://doi.org/10.1145/3399715.3399967>
7. Barra S, Carta SM, Corriga A, Podda AS, Recupero DR (2020) Deep learning and time series-to-image encoding for financial forecasting. *IEEE/CAA J Automat Sin* 7(3):683–692. <https://doi.org/10.1109/JAS.2020.1003132>
8. Borgelt C, Kruse R (2002) Induction of association rules: Apriori implementation. In: *Compstat*, pp 395–400. Springer
9. Carling C, Reilly T, Williams AM (2008) *Performance assessment for field sports*. Routledge, Evanston
10. Chen L, Hoey J, Nugent CD, Cook DJ, Yu Z (2012) Sensor-based activity recognition. *IEEE Trans Syst Man Cybern Part C Appl Rev* 42(6):790–808
11. Cioppa A, Deliege A, Giancola S, Ghanem B, Droogenbroeck MV, Gade R, Moeslund TB (2020) A context-aware loss function for action spotting in soccer videos. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*

12. Fernandez D, Varas D, Espadaler J, Masuda I, Ferreira J, Woodward A, Rodriguez D, Giro-i Nieto X, Carlos Riveiro J, Bou E (2017) Vits: Video tagging system from massive web multimedia collections. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops
13. Goal.com: Footballers don't wear bras - sporting reasons for under-shirt clothing explained. <https://www.goal.com/en-au/news/footballers-dont-wear-bras-sporting-reasons-under-shirt-clothing-/1aak15v6271f814s624c5ws52t>. Accessed 30 Nov 2021
14. Grehaigne J-F, Godbout P, Bouthier D (1997) Performance assessment in team sports. *J Teach Phys Educ* 16(4):500–516
15. Hao D, Zhang L, Sumkin J, Mohamed A, Wu S (2020) Inaccurate labels in weakly-supervised deep learning: Automatic identification and correction and their impact on classification performance. *IEEE J Biomed Health Inform* 24(9):2701–2710. <https://doi.org/10.1109/JBHI.2020.2974425>
16. Haryanto AW, Mawardi EK et al (2018) Influence of word normalization and chi-squared feature selection on support vector machine (svm) text classification. In: 2018 International seminar on application for technology of information and communication, pp 229–233. IEEE
17. Hosseini M-S, Eftekhari-Moghadam A-M (2013) Fuzzy rule-based reasoning approach for event detection and annotation of broadcast soccer video. *Appl Soft Comput* 13(2):846–866. <https://doi.org/10.1016/j.asoc.2012.10.007>
18. Khan AM, Lee Y-K, Lee SY, Kim T-S (2010) A triaxial accelerometer-based physical-activity recognition via augmented-signal features and a hierarchical recognizer. *IEEE Trans Inf Technol Biomed* 14(5):1166–1172
19. Li X, Xu C, Wang X, Lan W, Jia Z, Yang G, Xu J (2019) Coco-cn for cross-lingual image tagging, captioning, and retrieval. *IEEE Trans Multimedia* 21(9):2347–2360. <https://doi.org/10.1109/TMM.2019.2896494>
20. Liu H, Dougherty ER, Dy JG, Torkkola K, Tuv E, Peng H, Ding C, Long F, Berens M, Parsons L et al (2005) Evolving feature selection. *IEEE Intell Syst* 20(6):64–76
21. Luna JM, Fournier-Viger P, Ventura S (2019) Frequent itemset mining: A 25 years review. *Wiley Interdiscip Rev Data Min Knowl Disc* 9(6):e1329
22. Ma S, Shao E, Xie X, Liu W (2020) Event detection in soccer video based on self-attention. In: 2020 IEEE 6th International Conference on Computer and Communications (ICCC), pp 1852–1856. <https://doi.org/10.1109/ICCC51575.2020.9344896>
23. Morra L, Manigrasso F, Lamberti F (2020) Soccer: Computer graphics meets sports analytics for soccer event recognition. *SoftwareX* 12:100612. <https://doi.org/10.1016/j.softx.2020.100612>
24. Presti DL, Massaroni C, Caponero M, Formica D, Schena E (2021) Cardiorespiratory monitoring using a mechanical and an optical system. In: 2021 IEEE International Symposium on Medical Measurements and Applications (MeMeA), pp 1–6. <https://doi.org/10.1109/MeMeA52024.2021.9478750>
25. Qian X, Cheng X, Cheng G, Yao X, Jiang L (2021) Two-stream encoder gan with progressive training for co-saliency detection. *IEEE Signal Process Lett* 28:180–184. <https://doi.org/10.1109/LSP.2021.3049997>
26. Riboni D, Bettini C (2011) Cosar: hybrid reasoning for context-aware activity recognition. *Pers Ubiquit Comput* 15(3):271–289
27. Sámano A, Ocegueda-Hernández VC, Guerrero-Carrizales F, Fuentes JM, Mendizabal-Ruiz G Gvr: An intuitive tool for the visualization and easy interpretation of advanced exploration methods for the analysis of soccer matches
28. Sharma RA, Gandhi V, Chari V, Jawahar CV (2017) Automatic analysis of broadcast football videos using contextual priors. *SIViP* 11(1):171–178
29. Shawe-Taylor J, Cristianini N (2000) Support vector machines, vol. 2. Cambridge University Press, Cambridge
30. Sorano D, Carrara F, Cintia P, Falchi F, Pappalardo L (2020) Automatic pass annotation from soccer videostreams based on object detection and lstm. [arXiv:2007.06475](https://arxiv.org/abs/2007.06475)
31. Sorano D, Carrara F, Cintia P, Falchi F, Pappalardo L (2021) Automatic pass annotation from soccer video streams based on object detection and lstm. In: Dong Y, Ifrim G, Mladenčić D, Saunders C, Van Hoecke S (eds) Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track, pp 475–490. Springer International Publishing, Cham
32. Stein M, Janetzko H, Breitkreutz T, Seebacher D, Schreck T, Grossniklaus M, Couzin ID, Keim DA (2016) Director's cut: Analysis and annotation of soccer matches. *IEEE Comput Graph Appl* 36(5):50–60. <https://doi.org/10.1109/MCG.2016.102>
33. Theagarajan R, Bhanu B (2021) An automated system for generating tactical performance statistics for individual soccer players from videos. *IEEE Trans Circuits Syst Video Technol* 31(2):632–646. <https://doi.org/10.1109/TCSVT.2020.2982580>
34. Verikas A, Gelzinis A, Bacauskiene M (2011) Mining data with random forests: A survey and results of new tests. *Pattern Recogn* 44(2):330–349
35. Wang J, Chen Y, Hao S, Peng X, Hu L (2019) Deep learning for sensor-based activity recognition: A survey. *Pattern Recogn Lett* 119:3–11

36. Wang Z, Yu J, He Y (2017) Soccer video event annotation by synchronization of attack-defense clips and match reports with coarse-grained time information. *IEEE Trans Circuits Syst Video Technol* 27(5):1104–1117. <https://doi.org/10.1109/TCSVT.2016.2515280>
37. Xu Y, Kong Q, Huang Q, Wang W, Plumbley MD (2017) Convolutional gated recurrent neural network incorporating spatial features for audio tagging. In: 2017 International Joint Conference on Neural Networks (IJCNN), pp 3461–3466. <https://doi.org/10.1109/IJCNN.2017.7966291>
38. Yan F, Christmas W, Kittler J (2005) A tennis ball tracking algorithm for automatic annotation of tennis match. In *British machine vision conference* 2:619–628

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.