



(51) International Patent Classification:

G06T 7/20 (2017.01) G06N 3/04 (2006.01)
G06K 9/00 (2006.01)

(21) International Application Number:

PCT/US2020/020442

(22) International Filing Date:

28 February 2020 (28.02.2020)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/811,889 28 February 2019 (28.02.2019) US

(71) Applicant: STATS LLC [US/US]; 203 North LaSalle St., Suite 2200, Chicago, IL 60601 (US).

(72) Inventors: SHA, Long; c/o STATS LLC, 203 North LaSalle St., Suite 2200, Chicago, IL 60601 (US). GAN-

GULY, Sujoy; c/o STATS LLC, 203 North LaSalle St., Suite 2200, Chicago, IL 60601 (US). WEI, Xinyu; c/o STATS LLC, 203 North LaSalle St., Suite 2200, Chicago, IL 60601 (US). LUCEY, Patrick Joseph; c/o STATS LLC, 203 North LaSalle St., Suite 2200, Chicago, IL 60601 (US). CHERUKUMUDI, Aditya; c/o STATS LLC, 203 North LaSalle St., Suite 2200, Chicago, IL 60601 (US).

(74) Agent: WOLFE, Joseph et al.; DLA Piper LLP US, 33 Arch Street, 26th Floor, Boston, MD 02110 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,

(54) Title: SYSTEM AND METHOD FOR GENERATING PLAYER TRACKING DATA FROM BROADCAST VIDEO




FIG. 1

(57) Abstract: A system and method of generating a player tracking prediction are described herein. A computing system retrieves a broadcast video feed for a sporting event. The computing system segments the broadcast video feed into a unified view. The computing system generates a plurality of data sets based on the plurality of trackable frames. The computing system calibrates a camera associated with each trackable frame based on the body pose information. The computing system generates a plurality of sets of short tracklets based on the plurality of trackable frames and the body pose information. The computing system connects each set of short tracklets by generating a motion field vector for each player in the plurality of trackable frames. The computing system predicts a future motion of a player based on the player's motion field vector using a neural network.

SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR,
TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

System and Method for Generating Player Tracking Data from Broadcast Video

Cross-Reference to Related Applications

[0001] This application claims priority to U.S. Provisional Application No. 62/811,889, filed February 28, 2019, which is hereby incorporated by reference in its entirety.

Field of the Disclosure

[0002] The present disclosure generally relates to system and method for generating player tracking data from broadcast video.

Background

[0003] Player tracking data has been implemented for a number of years in a number of sports for both team and player analysis. Conventional player tracking systems, however, require sports analytics companies to install fixed cameras in each venue in which a team plays. This constraint has limited the scalability of player tracking systems, as well as limited data collection to currently played matches. Further, such constraint provides a significant cost to sports analytics companies due to the costs associated with installing hardware in the requisite arenas, as well as maintaining such hardware.

Summary

[0004] In some embodiments, a method of generating a player tracking prediction is disclosed herein. A computing system retrieves a broadcast video feed for a sporting event. The broadcast video feed includes a plurality of video frames. The computing system segments the broadcast video feed into a unified view. The unified view includes a plurality of trackable frames. The plurality of trackable frames is a subset of the plurality of video frames. The computing system generates a plurality of data sets based on the plurality of trackable frames. The plurality of data sets includes playing surface segmentation information, ball tracking information, and body pose information for each player in each trackable frame. The computing system calibrates a camera associated with each trackable frame based on the playing surface segmentation information and the body pose information. The computing system generates a plurality of sets of short tracklets based on the plurality of trackable frames and the body pose information. The computing system connects each set of short tracklets by generating a motion field vector for each player in the plurality of trackable frames. The computing system predicts a future motion of a player based on the player's motion field vector using a neural network.

[0005] In some embodiments, a system for generating a player tracking prediction is disclosed herein. The computing system includes a processor and a memory. The memory has programming instructions stored thereon, which, when executed by the processor, performs one or more operations. The one or more operations include retrieving a broadcast video feed for a sporting event. The broadcast video feed includes a plurality of video frames. The one or more operations further include segmenting the broadcast video feed into a unified view. The unified view includes a plurality of trackable frames, the plurality of trackable frames is a subset of the plurality of video frames. The one or more operations further include generating a plurality of data sets based on the plurality of trackable frames. The plurality of data sets includes playing surface segmentation information, ball tracking information, and body pose information for each player in each trackable frame. The one or more operations further include calibrating a camera associated with each trackable frame based on the playing surface segmentation information and the body pose information. The one or more operations further include generating a plurality of sets of short tracklets based on the plurality of trackable frames and the body pose information. The one or more operations further include connecting each set of short tracklets by generating a motion field vector for each player in the plurality of trackable frames. The one or more operations further include predicting a future motion of a player based on the player's motion field vector using a neural network.

[0006] In some embodiments, a non-transitory computer readable medium is disclosed herein. The non-transitory computer readable medium includes one or more sequences of instructions that, when executed by one or more processors, perform one or more operations. The one or more operations include retrieving a broadcast video feed for a sporting event. The broadcast video feed includes a plurality of video frames. The one or more operations further include segmenting the broadcast video feed into a unified view. The unified view includes a plurality of trackable frames, the plurality of trackable frames is a subset of the plurality of video frames. The one or more operations further include generating a plurality of data sets based on the plurality of trackable frames. The plurality of data sets includes playing surface segmentation information, ball tracking information, and body pose information for each player in each trackable frame. The one or more operations further include calibrating a camera associated with each trackable frame based on the playing surface segmentation information and the body pose information. The one or more operations further include generating a plurality of sets of short tracklets based on the plurality of

trackable frames and the body pose information. The one or more operations further include connecting each set of short tracklets by generating a motion field vector for each player in the plurality of trackable frames. The one or more operations further include predicting a future motion of a player based on the player's motion field vector using a neural network.

Brief Description of the Drawings

[0007] So that the manner in which the above recited features of the present disclosure can be understood in detail, a more particular description of the disclosure, briefly summarized above, may be had by reference to embodiments, some of which are illustrated in the appended drawings. It is to be noted, however, that the appended drawings illustrated only typical embodiments of this disclosure and are therefore not to be considered limiting of its scope, for the disclosure may admit to other equally effective embodiments.

[0008] Figure 1 is a block diagram illustrating a computing environment, according to example embodiments.

[0009] Figure 2 is a block diagram illustrating a computing environment, according to example embodiments.

[0010] Figure 3 is a block diagram illustrating aspects of operations discussed above and below in conjunction with Figure 2 and Figures 4-10, according to example embodiments.

[0011] Figure 4 is a flow diagram illustrating a method of generating player tracks, according to example embodiments.

[0012] Figure 5 is a flow diagram illustrating a method of generating trackable frames, according to example embodiments.

[0013] Figure 6 is a block diagram illustrating aspects of operations discussed above in conjunction with Figure 5, according to example embodiments.

[0014] Figure 7 is a flow diagram illustrating a method of calibrating a camera for each trackable frame, according to example embodiments.

[0015] Figure 8 is a block diagram illustrating aspects of operations discussed above in conjunction with Figure 7, according to example embodiments.

[0016] Figure 9 is a flow diagram illustrating a method of tracking players, according to example embodiments.

[0017] Figure 10 is a flow diagram illustrating a method of tracking players, according to example embodiments.

[0018] Figure 11 is a block diagram illustrating aspects of operations discussed above in conjunction with Figure 10, according to example embodiments.

[0019] Figure 12 is a block diagram illustrating architecture of re-identification agent, according to example embodiments.

[0020] Figure 13A is a block diagram illustrating a computing device, according to example embodiments.

[0021] Figure 13B is a block diagram illustrating a computing device, according to example embodiments.

[0022] To facilitate understanding, identical reference numerals have been used, where possible, to designate identical elements that are common to the figures. It is contemplated that elements disclosed in one embodiment may be beneficially utilized on other embodiments without specific recitation.

Detailed Description

[0023] Player tracking data has been an invaluable resource for leagues and teams to evaluate not only the team itself, but the players on the team. Conventional approaches to harvesting or generating player tracking data are limited, however, relied on installing fixed cameras in a venue in which a sporting event would take place. In other words, conventional approaches for a team to harvest or generate player tracking data required that team to equip each venue with a fixed camera system. As those skilled in the art recognize, this constraint has severely limited the scalability of player tracking systems. Further, this constraint also limits player tracking data to matches played after installation of the fixed camera system, as historical player tracking data would simply be unavailable.

[0024] The one or more techniques describe herein provide a significant improvement over conventional systems by eliminating the need for a fixed camera system. Instead, the one or more techniques described herein are directed to leveraging the broadcast video feed of a sporting event to generate player tracking data. By utilizing the broadcast video feed of the sporting event, not only is the need for a dedicated fixed camera system in each arena eliminated, but generating historical player tracking data from historical sporting events would now be possible.

[0025] Leveraging the broadcast video feed of the sporting event is not, however, a trivial task. For example, included in a broadcast video feed may a variety of different camera angles, close-ups of players, close-ups of the crowd, close-ups of the coach, video of the commentators, commercials, halftime shows, and the like. Further, due to the pace of a sporting event, it is inevitable that some players be either occluded from the broadcast video feed cameras or entirely out of the line of sight of the broadcast video feed cameras.

[0026] Accordingly, the one or more techniques describe herein provide various mechanisms directed to overcoming these challenges so that the broadcast video feed of a sporting event may be used to generate player tracking data.

[0027] Figure 1 is a block diagram illustrating a computing environment 100, according to example embodiments. Computing environment 100 may include camera system 102, organization computing system 104, and one or more client devices 108 communicating via network 105.

[0028] Network 105 may be of any suitable type, including individual connections via the Internet, such as cellular or Wi-Fi networks. In some embodiments, network 105 may connect terminals, services, and mobile devices using direct connections, such as radio frequency identification (RFID), near-field communication (NFC), Bluetooth™, low-energy Bluetooth™ (BLE), Wi-Fi™, ZigBee™, ambient backscatter communication (ABC) protocols, USB, WAN, or LAN. Because the information transmitted may be personal or confidential, security concerns may dictate one or more of these types of connection be encrypted or otherwise secured. In some embodiments, however, the information being transmitted may be less personal, and therefore, the network connections may be selected for convenience over security.

[0029] Network 105 may include any type of computer networking arrangement used to exchange data or information. For example, network 105 may be the Internet, a private data network, virtual private network using a public network and/or other suitable connection(s) that enables components in computing environment 100 to send and receive information between the components of environment 100.

[0030] Camera system 102 may be positioned in a venue 106. For example, venue 106 may be configured to host a sporting event that includes one or more agents 112. Camera system 102 may be configured to capture the motions of all agents (i.e., players) on the playing surface, as well as one or more other objects of relevance (e.g., ball, referees, etc.). In some embodiments, camera

system 102 may be an optically-based system using, for example, a plurality of fixed cameras. For example, a system of six stationary, calibrated cameras, which project the three-dimensional locations of players and the ball onto a two-dimensional overhead view of the court may be used. In another example, a mix of stationary and non-stationary cameras may be used to capture motions of all agents on the playing surface as well as one or more objects of relevance. As those skilled in the art recognize, utilization of such camera system (e.g., camera system 102) may result in many different camera views of the court (e.g., high sideline view, free-throw line view, huddle view, face-off view, end zone view, etc.). Generally, camera system 102 may be utilized for the broadcast feed of a given match. Each frame of the broadcast feed may be stored in a game file 110.

[0031] Camera system 102 may be configured to communicate with organization computing system 104 via network 105. Organization computing system 104 may be configured to manage and analyze the broadcast feed captured by camera system 102. Organization computing system 104 may include at least a web client application server 114, a data store 118, an auto-clipping agent 120, a data set generator 122, a camera calibrator 124, a player tracking agent 126, and an interface agent 128. Each of auto-clipping agent 120, data set generator 122, camera calibrator 124, player tracking agent 126, and interface agent 128 may be comprised of one or more software modules. The one or more software modules may be collections of code or instructions stored on a media (e.g., memory of organization computing system 104) that represent a series of machine instructions (e.g., program code) that implements one or more algorithmic steps. Such machine instructions may be the actual computer code the processor of organization computing system 104 interprets to implement the instructions or, alternatively, may be a higher level of coding of the instructions that is interpreted to obtain the actual computer code. The one or more software modules may also include one or more hardware components. One or more aspects of an example algorithm may be performed by the hardware components (e.g., circuitry) itself, rather as a result of the instructions.

[0032] Data store 118 may be configured to store one or more game files 124. Each game file 124 may include the broadcast data of a given match. For example, the broadcast data may a plurality of video frames captured by camera system 102.

[0033] Auto-clipping agent 120 may be configured parse the broadcast feed of a given match to identify a unified view of the match. In other words, auto-clipping agent 120 may be configured

to parse the broadcast feed to identify all frames of information that are captured from the same view. In one example, such as in the sport of basketball, the unified view may be a high sideline view. Auto-clipping agent 120 may clip or segment the broadcast feed (e.g., video) into its constituent parts (e.g., difference scenes in a movie, commercials from a match, etc.). To generate a unified view, auto-clipping agent 120 may identify those parts that capture the same view (e.g., high sideline view). Accordingly, auto-clipping agent 120 may remove all (or a portion) of untrackable parts of the broadcast feed (e.g., player close-ups, commercials, half-time shows, etc.). The unified view may be stored as a set of trackable frames in a database.

[0034] Data set generator 122 may be configured to generate a plurality of data sets from the trackable frames. In some embodiments, data set generator 122 may be configured to identify body pose information. For example, data set generator 122 may utilize body pose information to detect players in the trackable frames. In some embodiments, data set generator 122 may be configured to further track the movement of a ball or puck in the trackable frames. In some embodiments, data set generator 122 may be configured to segment the playing surface in which the event is taking place to identify one or more markings of the playing surface. For example, data set generator 122 may be configured to identify court (e.g., basketball, tennis, etc.) markings, field (e.g., baseball, football, soccer, rugby, etc.) markings, ice (e.g., hockey) markings, and the like. The plurality of data sets generated by data set generator 122 may be subsequently used by camera calibrator 124 for calibrating the cameras of each camera system 102.

[0035] Camera calibrator 124 may be configured to calibrate the cameras of camera system 102. For example, camera calibrator 124 may be configured to project players detected in the trackable frames to real world coordinates for further analysis. Because cameras in camera systems 102 are constantly moving in order to focus on the ball or key plays, such cameras are unable to be pre-calibrated. Camera calibrator 124 may be configured to improve or optimize player projection parameters using a homography matrix.

[0036] Player tracking agent 126 may be configured to generate tracks for each player on the playing surface. For example, player tracking agent 126 may leverage player pose detections, camera calibration, and broadcast frames to generate such tracks. In some embodiments, player tracking agent 126 may further be configured to generate tracks for each player, even if, for example, the player is currently out of a trackable frame. For example, player tracking agent 126 may utilize body pose information to link players that have left the frame of view.

[0037] Interface agent 128 may be configured to generate one or more graphical representations corresponding to the tracks for each player generated by player tracking agent 126. For example, interface agent 128 may be configured to generate one or more graphical user interfaces (GUIs) that include graphical representations of player tracking each prediction generated by player tracking agent 126.

[0038] Client device 108 may be in communication with organization computing system 104 via network 105. Client device 108 may be operated by a user. For example, client device 108 may be a mobile device, a tablet, a desktop computer, or any computing system having the capabilities described herein. Users may include, but are not limited to, individuals such as, for example, subscribers, clients, prospective clients, or customers of an entity associated with organization computing system 104, such as individuals who have obtained, will obtain, or may obtain a product, service, or consultation from an entity associated with organization computing system 104.

[0039] Client device 108 may include at least application 132. Application 132 may be representative of a web browser that allows access to a website or a stand-alone application. Client device 108 may access application 132 to access one or more functionalities of organization computing system 104. Client device 108 may communicate over network 105 to request a webpage, for example, from web client application server 114 of organization computing system 104. For example, client device 108 may be configured to execute application 132 to access content managed by web client application server 114. The content that is displayed to client device 108 may be transmitted from web client application server 114 to client device 108, and subsequently processed by application 132 for display through a graphical user interface (GUI) of client device 108.

[0040] Figure 2 is a block diagram illustrating a computing environment 200, according to example embodiments. As illustrated, computing environment 200 includes auto-clipping agent 120, data set generator 122, camera calibrator 124, and player tracking agent 126 communicating via network 205.

[0041] Network 205 may be of any suitable type, including individual connections via the Internet, such as cellular or Wi-Fi networks. In some embodiments, network 205 may connect terminals, services, and mobile devices using direct connections, such as radio frequency identification (RFID), near-field communication (NFC), Bluetooth™, low-energy Bluetooth™ (BLE), Wi-Fi™,

ZigBee™, ambient backscatter communication (ABC) protocols, USB, WAN, or LAN. Because the information transmitted may be personal or confidential, security concerns may dictate one or more of these types of connection be encrypted or otherwise secured. In some embodiments, however, the information being transmitted may be less personal, and therefore, the network connections may be selected for convenience over security.

[0042] Network 205 may include any type of computer networking arrangement used to exchange data or information. For example, network 205 may be the Internet, a private data network, virtual private network using a public network and/or other suitable connection(s) that enables components in computing environment 200 to send and receive information between the components of environment 200.

[0043] Auto-clipping agent 120 may include principal component analysis (PCA) agent 202, clustering model 204, and neural network 206. As recited above, when trying to understand and extract data from a broadcast feed, auto-clipping agent 120 may be used to clip or segment the video into its constituent parts. In some embodiments, auto-clipping agent 120 may focus on separating a predefined, unified view (e.g., a high sideline view) from all other parts of the broadcast stream.

[0044] PCA agent 202 may be configured to utilize a PCA analysis to perform per frame feature extraction from the broadcast feed. For example, given a pre-recorded video, PCA agent 202 may extract a frame every X-seconds (e.g., 10 seconds) to build a PCA model of the video. In some embodiments, PCA agent 202 may generate the PCA model using incremental PCA, through which PCA agent 202 may select a top subset of components (e.g., top 120 components) to generate the PCA model. PCA agent 202 may be further configured to extract one frame every X seconds (e.g., one second) from the broadcast stream and compress the frames using PCA model. In some embodiments, PCA agent 202 may utilize PCA model to compress the frames into 120-dimensional form. For example, PCA agent 202 may solve for the principal components in a per video manner and keep the top 100 components per frame to ensure accurate clipping.

[0045] Clustering model 204 may be configured to cluster the top subset of components into clusters. For example, clustering model 204 may be configured to center, normalize, and cluster the top 120 components into a plurality of clusters. In some embodiments, for clustering of compressed frames, clustering model 204 may implement k-means clustering. In some

embodiments, clustering model 204 may set $k = 9$ clusters. K-means clustering attempts to take some data $x = \{x_1, x_2, \dots, x_n\}$ and divide it into k subsets, $S = \{S_1, S_2, \dots, S_k\}$ by optimizing:

$$\arg \min_S \sum_j^k \sum_{x \in S_j} \|x - \mu_j\|^2$$

where μ_j is the mean of the data in the set S_j . In other words, clustering model 204 attempts to find clusters with the smallest inter-cluster variance using k-means clustering techniques. Clustering model 204 may label each frame with its respective cluster number (e.g., cluster 1, cluster 2, . . . , cluster k).

[0046] Neural network 206 may be configured to classify each frame as trackable or untrackable. A trackable frame may be representative of a frame that includes captures the unified view (e.g., high sideline view). An untrackable frame may be representative of a frame that does not capture the unified view. To train neural network 206, an input data set that includes thousands of frames pre-labeled as trackable or untrackable that are run through the PCA model may be used. Each compressed frame and label pair (i.e., cluster number and trackable/untrackable) may be provided to neural network 206 for training.

[0047] In some embodiments, neural network 206 may include four layers. The four layers may include an input layer, two hidden layers, and an output layer. In some embodiments, input layer may include 120 units. In some embodiments, each hidden layer may include 240 units. In some embodiments, output layer may include two units. The input layer and each hidden layer may use sigmoid activation functions. The output layer may use a SoftMax activation function. To train neural network 206, auto-clipping agent 120 may reduce (e.g., minimize) the binary cross-entropy loss between the predicted label for sample $\hat{y}_j$ and the true label y_j by:

$$H = -\frac{1}{N} \sum_j^N y_j \log y + (1 - y_j) \log(1 - \hat{y}_j)$$

[0048] Accordingly, once trained, neural network 206 may be configured to classify each frame as untrackable or trackable. As such, each frame may have two labels: a cluster number and trackable/untrackable classification. Auto-clipping agent 120 may utilize the two labels to determine if a given cluster is deemed trackable or untrackable. For example, if auto-clipping agent 120 determines that a threshold number of frames in a cluster are considered trackable (e.g.,

80%), auto-clipping agent 120 may conclude that all frames in the cluster are trackable. Further, if auto-clipping agent 120 determines that less than a threshold number of frames in a cluster are considered untrackable (e.g., 30% and below), auto-clipping agent 120 may conclude that all frames in the cluster are untrackable. Still further, if auto-clipping agent 120 determines that a certain number of frames in a cluster are considered trackable (e.g., between 30% and 80%), auto-clipping agent 120 may request that an administrator further analyze the cluster. Once each frame is classified, auto-clipping agent 120 may clip or segment the trackable frames. Auto-clipping agent 120 may store the segments of trackable frames in database 205 associated therewith.

[0049] Data set generator 122 may be configured to generate a plurality of data sets from auto-clipping agent 120. As illustrated, data set generator 122 may include pose detector 212, ball detector 214, and playing surface segmenter 216. Pose detector 212 may be configured to detect players within the broadcast feed. Data set generator 122 may provide, as input, to pose detector 212 both the trackable frames stored in database 205 as well as the broadcast video feed. In some embodiments, pose detector 212 may implement Open Pose to generate body pose data to detect players in the broadcast feed and the trackable frames. In some embodiments, pose detector 212 may implement sensors positioned on players to capture body pose information. Generally, pose detector 212 may use any means to obtain body pose information from the broadcast video feed and the trackable frame. The output from pose detector 212 may be pose data stored in database 215 associated with data set generator 122.

[0050] Ball detector 214 may be configured to detect and track the ball (or puck) within the broadcast feed. Data set generator 122 may provide, as input, to ball detector 214 both the trackable frames stored in database 205 and the broadcast video feed. In some embodiments, ball detector 214 may utilize a faster region-convolutional neural network (R-CNN) to detect and track the ball in the trackable frames and broadcast video feed. Faster R-CNN is a regional proposal based network. Faster R-CNN uses a convolutional neural network to propose a region of interest, and then classifies the object in each region of interest. Because it is a single unified network, the regions of interest and the classification steps may improve each other, thus allowing the classification to handle objects of various sizes. The output from ball detector 214 may be ball detection data stored in database 215 associated with data set generator 122.

[0051] Playing surface segmenter 216 may be configured to identify playing surface markings in the broadcast feed. Data set generator 122 may provide, as input, to playing surface segmenter

216 both trackable frames stored in database 205 and the broadcast video feed. In some embodiments, playing surface segmenter 216 may be configured to utilize a neural network to identify playing surface markings. The output from playing surface segmenter 216 may be playing surface markings stored in database 215 associated with data set generator 122.

[0052] Camera calibrator 124 may be configured to address the issue of moving camera calibration in sports. Camera calibrator 124 may include spatial transfer network 224 and optical flow module 226. Camera calibrator 124 may receive, as input, segmented playing surface information generated by playing surface segmenter 216, the trackable clip information, and posed information. Given such inputs, camera calibrator 124 may be configured to project coordinates in the image frame to real-world coordinates for tracking analysis.

[0053] Keyframe matching module 224 may receive, as input, output from playing surface segmenter 216 and a set of templates. For each frame, keyframe matching module 224 may match the output from playing surface segmenter 216 to a template. Those frames that are able to match to a given template are considered keyframes. In some embodiments, keyframe matching module 224 may implement a neural network to match the one or more frames. In some embodiments, keyframe matching module 224 may implement cross-correlation to match the one or more frames.

[0054] Spatial transformer network (STN) 224 may be configured to receive, as input, the identified keyframes from keyframe matching module 224. STN 224 may implement a neural network to fit a playing surface model to segmentation information of the playing surface. By fitting the playing surface model to such output, STN 224 may generate homography matrices for each keyframe.

[0055] Optical flow module 226 may be configured to identify the pattern of motion of objects from one trackable frame to another. In some embodiments, optical flow module 226 may receive, as input, trackable frame information and body pose information for players in each trackable frame. Optical flow module 226 may use body pose information to remove players from the trackable frame information. Once removed, optical flow module 226 may determine the motion between frames to identify the motion of a camera between successive frames. In other words, optical flow module 226 may identify the flow field from one frame to the next.

[0056] Optical flow module 226 and STN 224 may work in conjunction to generate a homography matrix. For example, optical flow module 226 and STN 224 may generate a homography matrix for each trackable frame, such that a camera may be calibrated for each frame. The homography

matrix may be used to project the track or position of players into real-world coordinates. For example, the homography matrix may indicate a 2-dimensional to 2-dimensional transform, which may be used to project the players' locations from image coordinates to the real world coordinates on the playing surface.

[0057] Player tracking agent 126 may be configured to generate a track for each player in a match. Player tracking agent 126 may include neural network 232 and re-identification agent 232. Player tracking agent 126 may receive, as input, trackable frames, pose data, calibration data, and broadcast video frames. In a first phase, player tracking agent 126 may match pairs of player patches, which may be derived from pose information, based on appearance and distance. For example, let H_j^t be the player patch of the j^{th} player at time t , and let $I_j^t = \{x_j^t, y_j^t, w_j^t, h_j^t\}$ be the image coordinates x_j^t, y_j^t , the width w_j^t , and the height h_j^t of the j^{th} player at time t . Using this, player tracking agent 126 may associate any pair of detections using the appearance cross correlation $C_{ij}^t = H_i^t * H_j^{t+1}$ and $L_{ij}^t = \|I_i^t - I_j^{t+1}\|_2^2$ by finding:

$$\arg \max_{ij} (C_{ij}^t + L_{ij}^t)$$

where I is the bounding box positions (x, y) , width w , and height h ; C is the cross correlation between the image patches (e.g., image cutout using a bounding box) and measures similarity between two image patches; and L is a measure of the difference (e.g., distance) between two bounding boxes I .

[0058] Performing this for every pair may generate a large set of short tracklets. The end points of these tracklets may then be associated with each other based on motion consistency and color histogram similarity. In some embodiments, a tracklet may be referred to as a subset of a track.

[0059] For example, let v_i be the extrapolated velocity from the end of the i^{th} tracklet and v_j be the velocity extrapolated from the beginning of the j^{th} tracklet. Then $c_{ij} = v_i \cdot v_j$ may represent the motion consistency score. Furthermore, let $p(h)_i$ represent the likelihood of a color h being present in an image patch i . Player tracking agent 126 may measure the color histogram similarity using Bhattacharyya distance:

$$D_B(p_i, p_j) = -\ln(BC(p_i, p_j)) \text{ with } BC(p_i, p_j) = \sum_h \sqrt{p_i(h)p_j(h)}$$

[0060] Recall, tracking agent 120 finds the matching pair of tracklets by finding:

$$\arg \max_{ij} (c_{ij} + D_B(p_i, p_j)).$$

[0061] Solving for every pair of broken tracklets may result in a set of clean tracklets, while leaving some tracklets with large, i.e., many frames, gaps. To connect the large gaps, player tracking agent may augment affinity measures to include a motion field estimation, which may account for the change of player direction that occurs over many frames.

[0062] The motion field may be a vector field that represents the velocity magnitude and direction as a vector on each location on the playing surface. Given the known velocity of a number of players on the playing surface, the full motion field may be generated using cubic spline interpolation. For example, let $X_i = \{x_i^t\}_{t \in (0, T)}$ to be the court position of a player i at every time t . Then, there may exist a pair of points that have a displacement $d_i^\lambda(x_i^t) = x_i^t - x_i^{t+\lambda}$ if $\lambda < T$. Accordingly, the motion field may then be:

$$V(x, \lambda) = G(x, 5) * \sum_i d_i^\lambda(x_i^t),$$

where $G(x, 5)$ may be a Gaussian kernel with standard deviation equal to about five feet. In other words, motion field may be a Gaussian blur of all displacements.

[0063] Neural network 232 may be used to predict player motion fields given ground truth player trajectories. Given a set of ground truth player trajectories, X_i , the velocity of each player at each frame may be calculated, which may provide the ground truth motion field for neural network 232 to learn. For example, given a set of ground truth player trajectories X_i , player tracking agent 126 may be configured to generate the set $\hat{V}(x, \lambda)$, where $\hat{V}(x, \lambda)$ may be the predicted motion field. Neural network 232 may be trained, for example, to minimize $\|V(x, \lambda) - \hat{V}(x, \lambda)\|_2^2$. Player trajectory agent may then generate the affinity score for any tracking gap of size λ by:

$$K_{ij} = V(x, \lambda) \cdot d_{ij}$$

where $d_{ij} = x_i^t - x_j^{t+\lambda}$ is the displacement vector between all broken tracks with a gap size of λ .

[0064] Re-identification agent 234 may be configured to link players that have left the field of view. Re-identification agent 234 may include track generator 236, conditional autoencoder 240, and Siamese network 242.

[0065] Track generator 236 may be configured to generate a gallery of tracks. Track generator 236 may receive a plurality of tracks from database 205. For each track X , there may include a player identity label y , and for each player patch I , pose information p may be provided by the pose detection stage. Given a set of player tracks, track generator 236 may build a gallery for each

track where the jersey number of a player (or some other static feature) is always visible. The body pose information generated by data set generator 122 allows track generator 236 to determine a player's orientation. For example, track generator 236 may utilize a heuristic method, which may use the normalized shoulder width to determine the orientation:

$$S_{orient} = \frac{\|l_{Lshoulder} - l_{Rshoulder}\|_2}{\|l_{Neck} - l_{Hip}\|_2}$$

where l may represent the location of one body part. The width of shoulder may be normalized by the length of the torso so that the effect of scale may be eliminated. As two shoulders should be apart when a player faces towards or backwards from the camera, track generator 236 may use those patches whose S_{orient} is larger than a threshold to build the gallery. After this stage, each track X_n , may include a gallery:

$$G_n = \{l_i | S_{orient,i} > thresh\} \forall l_i \in X_n$$

[0066] Conditional autoencoder 240 may be configured to identify one or more features in each track. For example, unlike conventional approaches to re-identification issues, players in team sports may have very similar appearance features, such as clothing style, clothing color, and skin color. One of the more intuitive differences may be the jersey number that may be shown at the front and/or back side of each jersey. In order to capture those specific features, conditional autoencoder 240 may be trained to identify such features.

[0067] In some embodiments, conditional autoencoder 240 may be a three-layer convolutional autoencoder, where the kernel sizes may be 3x3 for all three layers, in which there are 64, 128, 128 channels respectively. Those hyper-parameters may be tuned to ensure that jersey number may be recognized from the reconstructed images so that the desired features may be learned in the autoencoder. In some embodiments, $f(I_i)$ may be used to denote the features that are learned from image i .

[0068] Use of conditional autoencoder 240 improves upon conventional processes for a variety of reasons. First, there is typically not enough training data for every player because some players only play a very short time in each game. Second, different teams can have the same jersey colors and jersey numbers, so classifying those players may be difficult.

[0069] Siamese network 242 may be used to measure the similarity between two image patches. For example, Siamese network 242 may be trained to measure the similarity between two image patches based on their feature representations $f(I)$. Given two image patches, their feature

representations $f(I_i)$ and $f(I_j)$ may be flattened, connected, and input into a perception network. In some embodiments, L_2 norm may be used to connect the two sub-networks of $f(I_i)$ and $f(I_j)$. In some embodiments, perception network may include three layers, which may include 1024, 512, and 216 hidden units, respectively. Such network may be used to measure the similarity $s(I_i, I_j)$ between every pair of image patches of the two tracks that have no time overlapping. In order to increase the robustness of the prediction, the final similarity score of the two tracks may be the average of all pairwise scores in their respective galleries:

$$S(x_n, x_m) = \frac{1}{|G_n||G_m|} \sum_{i \in G_n, j \in G_m} s(I_i, I_j)$$

[0070] This similarity score may be computed for every two tracks that do not have time overlapping. If the score is higher than some threshold, those two tracks may be associated.

[0071] Figure 3 is a block diagram 300 illustrating aspects of operations discussed above and below in conjunction with Figure 2 and Figures 4-10, according to example embodiments. Block diagram 300 may illustrate the overall workflow of organization computing system 104 in generating player tracking information. Block diagram 300 may include set of operations 302-308. Set of operations 302 may be directed to generating trackable frames (e.g., Method 500 in Figure 5). Set of operations 304 may be directed to generating one or more data sets from trackable frames (e.g., operations performed by data set generator 122). Set of operations 306 may be directed to camera calibration operations (e.g., Method 700 in Figure 7). Set of operations 308 may be directed to generating and predicting player tracks (e.g., Method 900 in Figure 9 and Method 1000 in Figure 10).

[0072] Figure 4 is a flow diagram illustrating a method 400 of generating player tracks, according to example embodiments. Method 400 may begin at step 402.

[0073] At step 402, organization computing system 104 may receive (or retrieve) a broadcast feed for an event. In some embodiments, the broadcast feed may be a live feed received in real-time (or near real-time) from camera system 102. In some embodiments, the broadcast feed may be a broadcast feed of a game that has concluded. Generally, the broadcast feed may include a plurality of frames of video data. Each frame may capture a different camera perspective.

[0074] At step 404, organization computing system 104 may segment the broadcast feed into a unified view. For example, auto-clipping agent 120 may be configured to parse the plurality of

frames of data in the broadcast feed to segment the trackable frames from the untrackable frames. Generally, trackable frames may include those frames that are directed to a unified view. For example, the unified view may be considered a high sideline view. In other examples, the unified view may be an endzone view. In other examples, the unified view may be a top camera view.

[0075] At step 406, organization computing system 104 may generate a plurality of data sets from the trackable frames (i.e., the unified view). For example, data set generator 122 may be configured to generate a plurality of data sets based on trackable clips received from auto-clipping agent 120. In some embodiments, pose detector 212 may be configured to detect players within the broadcast feed. Data set generator 122 may provide, as input, to pose detector 212 both the trackable frames stored in database 205 as well as the broadcast video feed. The output from pose detector 212 may be pose data stored in database 215 associated with data set generator 122.

[0076] Ball detector 214 may be configured to detect and track the ball (or puck) within the broadcast feed. Data set generator 122 may provide, as input, to ball detector 214 both the trackable frames stored in database 205 and the broadcast video feed. In some embodiments, ball detector 214 may utilize a faster R-CNN to detect and track the ball in the trackable frames and broadcast video feed. The output from ball detector 214 may be ball detection data stored in database 215 associated with data set generator 122.

[0077] Playing surface segmenter 216 may be configured to identify playing surface markings in the broadcast feed. Data set generator 122 may provide, as input, to playing surface segmenter 216 both trackable frames stored in database 205 and the broadcast video feed. In some embodiments, playing surface segmenter 216 may be configured to utilize a neural network to identify playing surface markings. The output from playing surface segmenter 216 may be playing surface markings stored in database 215 associated with data set generator 122.

[0078] Accordingly, data set generator 120 may generate information directed to player location, ball location, and portions of the court in all trackable frames for further analysis.

[0079] At step 408, organization computing system 104 may calibrate the camera in each trackable frame based on the data sets generated in step 406. For example, camera calibrator 124 may be configured to calibrate the camera in each trackable frame by generating a homography matrix, using the trackable frames and information. The homography matrix allows camera calibrator 124 to take those trajectories of each player in a given frame and project those trajectories into

real-world coordinates. By projection player position and trajectories into real world coordinates for each frame, camera calibrator 124 may ensure that the camera is calibrated for each frame.

[0080] At step 410, organization computing system 104 may be configured to generate or predict a track for each player. For example, player tracking agent 126 may be configured to generate or predict a track for each player in a match. Player tracking agent 126 may receive, as input, trackable frames, pose data, calibration data, and broadcast video frames. Using such inputs, player tracking agent 126 may be configured to construct player motion throughout a given match. Further, player tracking agent 126 may be configured to predict player trajectories given previous motion of each player.

[0081] Figure 5 is a flow diagram illustrating a method 500 of generating trackable frames, according to example embodiments. Method 500 may correspond to operation 404 discussed above in conjunction with Figure 4. Method 500 may begin at step 502.

[0082] At step 502, organization computing system 104 may receive (or retrieve) a broadcast feed for an event. In some embodiments, the broadcast feed may be a live feed received in real-time (or near real-time) from camera system 102. In some embodiments, the broadcast feed may be a broadcast feed of a game that has concluded. Generally, the broadcast feed may include a plurality of frames of video data. Each frame may capture a different camera perspective.

[0083] At step 504, organization computing system 104 may generate a set of frames for image classification. For example, auto-clipping agent 120 may utilize a PCA analysis to perform per frame feature extraction from the broadcast feed. Given, for example, a pre-recorded video, auto-clipping agent 120 may extract a frame every X-seconds (e.g., 10 seconds) to build a PCA model of the video. In some embodiments, auto-clipping agent 120 may generate the PCA model using incremental PCA, through which auto-clipping agent 120 may select a top subset of components (e.g., top 120 components) to generate the PCA model. Auto-clipping agent 120 may be further configured to extract one frame every X seconds (e.g., one second) from the broadcast stream and compress the frames using PCA model. In some embodiments, auto-clipping agent 120 may utilize PCA model to compress the frames into 120-dimensional form. For example, auto-clipping agent 120 may solve for the principal components in a per video manner and keep the top 100 components per frame to ensure accurate clipping. Such subset of compressed frames may be considered the set of frames for image classification. In other words, PCA model may be used to compress each frame to a small vector, so that clustering can be conducted on the frames more

efficiently. The compression may be conducted by selecting the top N components from PCA model to represent the frame. In some examples, N may be 100.

[0084] At step 506, organization computing system 104 may assign each frame in the set of frames to a given cluster. For example, auto-clipping agent 120 may be configured to center, normalize, and cluster the top 120 components into a plurality of clusters. In some embodiments, for clustering of compressed frames, auto-clipping agent 120 may implement k-means clustering. In some embodiments, auto-clipping agent 120 may set $k = 9$ clusters. K-means clustering attempts to take some data $x = \{x_1, x_2, \dots, x_n\}$ and divide it into k subsets, $S = \{S_1, S_2, \dots, S_k\}$ by optimizing:

$$\arg \min_S \sum_j^k \sum_{x \in S_j} \|x - \mu_j\|^2$$

where μ_j is the mean of the data in the set S_j . In other words, clustering model 204 attempts to find clusters with the smallest inter-cluster variance using k-means clustering techniques. Clustering model 204 may label each frame with its respective cluster number (e.g., cluster 1, cluster 2, . . . , cluster k).

[0085] At step 508, organization computing system 104 may classify each frame as trackable or untrackable. For example, auto-clipping agent 120 may utilize a neural network to classify each frame as trackable or untrackable. A trackable frame may be representative of a frame that includes captures the unified view (e.g., high sideline view). An untrackable frame may be representative of a frame that does not capture the unified view. To train the neural network (e.g., neural network 206), an input data set that includes thousands of frames pre-labeled as trackable or untrackable that are run through the PCA model may be used. Each compressed frame and label pair (i.e., cluster number and trackable/untrackable) may be provided to neural network for training. Accordingly, once trained, auto-clipping agent 120 may classify each frame as untrackable or trackable. As such, each frame may have two labels: a cluster number and trackable/untrackable classification.

[0086] At step 510, organization computing system 104 may compare each cluster to a threshold. For example, auto-clipping agent 120 may utilize the two labels to determine if a given cluster is deemed trackable or untrackable. In some embodiments, f auto-clipping agent 120 determines that a threshold number of frames in a cluster are considered trackable (e.g., 80%), auto-clipping agent

120 may conclude that all frames in the cluster are trackable. Further, if auto-clipping agent 120 determines that less than a threshold number of frames in a cluster are considered untrackable (e.g., 30% and below), auto-clipping agent 120 may conclude that all frames in the cluster are untrackable. Still further, if auto-clipping agent 120 determines that a certain number of frames in a cluster are considered trackable (e.g., between 30% and 80%), auto-clipping agent 120 may request that an administrator further analyze the cluster.

[0087] If at step 510 organization computing system 104 determines that greater than a threshold number of frames in the cluster are trackable, then at step 512 auto-clipping agent 120 may classify the cluster as trackable.

[0088] If, however, at step 510 organization computing system 104 determines that less than a threshold number of frames in the cluster are trackable, then at step 514, auto-clipping agent 120 may classify the cluster as untrackable.

[0089] Figure 6 is a block diagram 600 illustrating aspects of operations discussed above in conjunction with method 500, according to example embodiments. As shown, block diagram 600 may include a plurality of sets of operations 602-608.

[0090] At set of operations 602, video data (e.g., broadcast video) may be provided to auto-clipping agent 120. Auto-clipping agent 120 may extract frames from the video. In some embodiments, auto-clipping agent 120 may extract frames from the video at a low frame rate. An incremental PCA algorithm may be used by auto-clipping agent to select the top 120 components (e.g., frames) from the set of frames extracted by auto-clipping agent 120. Such operations may generate a video specific PCA model.

[0091] At set of operations 604, video data (e.g., broadcast video) may be provided to auto-clipping agent 120. Auto-clipping agent 120 may extract frames from the video. In some embodiments, auto-clipping agent 120 may extract frames from the video at a medium frame rate. The video specific PCA model may be used by auto-clipping agent 120 to compress the frames extracted by auto-clipping agent 120.

[0092] At set of operations 606, the compressed frames and a pre-selected number of desired clusters may be provided to auto-clipping agent 120. Auto-clipping agent 120 may utilize k-means clustering techniques to group the frames into one or more clusters, as set forth by the pre-selected number of desired clusters. Auto-clipping agent 120 may assign a cluster label to each compressed

frames. Auto-clipping agent 120 may further be configured to classify each frame as trackable or untrackable. Auto-clipping agent 120 may label each respective frame as such.

[0093] At set of operations 608, auto-clipping agent 120 may analyze each cluster to determine if the cluster includes at least a threshold number of trackable frames. For example, as illustrated, if 80% of the frames of a cluster are classified as trackable, then auto-clipping agent 120 may consider the entire cluster as trackable. If, however, less than 80% of a cluster is classified as trackable, auto-clipping agent may determine if at least a second threshold number of frames in a cluster are trackable. For example, is illustrated if 70% of the frames of a cluster are classified as untrackable, auto-clipping agent 120 may consider the entire cluster trackable. If, however, less than 70% of the frames of the cluster are classified as untrackable, i.e., between 30% and 70% trackable, then human annotation may be requested.

[0094] Figure 7 is a flow diagram illustrating a method 700 of calibrating a camera for each trackable frame, according to example embodiments. Method 700 may correspond to operation 408 discussed above in conjunction with Figure 4. Method 700 may begin at step 702.

[0095] At step 702, organization computing system 104 may retrieve video data and pose data for analysis. For example, camera calibrator 124 may retrieve from database 205 the trackable frames for a given match and pose data for players in each trackable frame. Following step 702, camera calibrator 124 may execute two parallel processes to generate homography matrix for each frame. Accordingly, the following operations are not meant to be discussed as being performed sequentially, but may instead be performed in parallel or sequentially.

[0096] At step 704, organization computing system 104 may remove players from each trackable frame. For example, camera calibrator 124 may parse each trackable frame retrieved from database 205 to identify one or more players contained therein. Camera calibrator 124 may remove the players from each trackable frame using the pose data retrieved from database 205. For example, camera calibrator 124 may identify those pixels corresponding to pose data and remove the identified pixels from a given trackable frame.

[0097] At step 706, organization computing system 104 may identify the motion of objects (e.g., surfaces, edges, etc.) between successive trackable frames. For example, camera calibrator 124 may analyze successive trackable frames, with players removed therefrom, to determine the motion of objects from one frame to the next. In other words, optical flow module 226 may identify the flow field between successive trackable frames.

[0098] At step 708, organization computing system 104 may match an output from playing surface segmenter 216 to a set of templates. For example, camera calibrator 124 may match one or more frames in which the image of the playing surface is clear to one or more templates. Camera calibrator 124 may parse the set of trackable clips to identify those clips that provide a clear picture of the playing surface and the markings therein. Based on the selected clips, camera calibrator 124 may compare such images to playing surface templates. Each template may represent a different camera perspective of the playing surface. Those frames that are able to match to a given template are considered keyframes. In some embodiments, camera calibrator 124 may implement a neural network to match the one or more frames. In some embodiments, camera calibrator 124 may implement cross-correlation to match the one or more frames.

[0099] At step 710, organization computing system 104 may fit a playing surface model to each keyframe. For example, camera calibrator 124 may be configured to receive, as input, the identified keyframes. Camera calibrator 124 may implement a neural network to fit a playing surface model to segmentation information of the playing surface. By fitting the playing surface model to such output, camera calibrator 124 may generate homography matrices for each keyframe.

[0100] At step 712, organization computing system 104 may generate a homography matrix for each trackable frame. For example, camera calibrator 124 may utilize the flow fields identified in step 706 and the homography matrices for each key frame to generate a homography matrix for each frame. The homography matrix may be used to project the track or position of players into real-world coordinates. For example, given the geometric transform represented by the homography matrix, camera calibrator 124 may use this transform to project the location of players on the image to real-world coordinates on the playing surface.

[0101] At step 714, organization computing system 104 may calibrate each camera based on the homography matrix.

[0102] Figure 8 is a block diagram 800 illustrating aspects of operations discussed above in conjunction with method 700, according to example embodiments. As shown, block diagram 800 may include inputs 802, a first set of operations 804, and a second set of operations 806. First set of operations 804 and second set of operations 806 may be performed in parallel.

[0103] Inputs 802 may include video clips 808 and pose detection 810. In some embodiments, video clips 808 may correspond to trackable frames generated by auto-clipping agent 120. In some

embodiments, pose detection 810 may correspond to pose data generated by pose detector 212. As illustrated, only video clips 808 may be provided as input to first set of operations 804; both video clips 804 and post detection 810 may be provided as input to second set of operations 806.

[0104] First set of operations 804 may include semantic segmentation 812, keyframe matching 814, and STN fitting 816. At semantic segmentation 812, playing surface segmenter 216 may be configured to identify playing surface markings in a broadcast feed. In some embodiments, playing surface segmenter 216 may be configured to utilize a neural network to identify playing surface markings. Such segmentation information may be performed in advance and provided to camera calibration 124 from database 215. At keyframe matching 814, keyframe matching module 224 may be configured to match one or more frames in which the image of the playing surface is clear to one or more templates. At STN fitting 816, STN 226 may implement a neural network to fit a playing surface model to segmentation information of the playing surface. By fitting the playing surface model to such output, STN 224 may generate homography matrices for each keyframe.

[0105] Second set of operations 806 may include camera motion estimation 818. At camera flow estimation 818, optical flow module 226 may be configured to identify the pattern of motion of objects from one trackable frame to another. For example, optical flow module 226 may use body pose information to remove players from the trackable frame information. Once removed, optical flow module 226 may determine the motion between frames to identify the motion of a camera between successive frames.

[0106] First set of operations 804 and second set of operations 806 may lead to homography interpolation 816. Optical flow module 226 and STN 224 may work in conjunction to generate a homography matrix for each trackable frame, such that a camera may be calibrated for each frame. The homography matrix may be used to project the track or position of players into real-world coordinates.

[0107] Figure 9 is a flow diagram illustrating a method 900 of tracking players, according to example embodiments. Method 900 may correspond to operation 410 discussed above in conjunction with Figure 4. Method 900 may begin at step 902.

[0108] At step 902, organization computing system 104 may retrieve a plurality of trackable frames for a match. Each of the plurality of trackable frames may include one or more sets of metadata associated therewith. Such metadata may include, for example, body pose information

and camera calibration data. In some embodiments, player tracking agent 126 may further retrieve broadcast video data.

[0109] At step 904, organization computing system 104 may generate a set of short tracklets. For example, player tracking agent 126 may match pairs of player patches, which may be derived from pose information, based on appearance and distance to generate a set of short tracklets. For example, let H_j^t be the player patch of the j^{th} player at time t , and let $I_j^t = \{x_j^t, y_j^t, w_j^t, h_j^t\}$ be the image coordinates x_j^t, y_j^t , the width w_j^t , and the height h_j^t of the j^{th} player at time t . Using this, player tracking agent 126 may associated any pair of detections using the appearance cross correlation $C_{ij}^t = H_i^t * H_j^{t+1}$ and $L_{ij}^t = \|I_i^t - I_j^{t+1}\|_2^2$ by finding:

$$\arg \max_{ij} (C_{ij}^t + L_{ij}^t).$$

[0110] Performing this for every pair may generate a set of short tracklets. The end points of these tracklets may then be associated with each other based on motion consistency and color histogram similarity.

[0111] For example, let v_i be the extrapolated velocity from the end of the i^{th} tracklet and v_j be the velocity extrapolated from the beginning of the j^{th} tracklet. Then $c_{ij} = v_i \cdot v_j$ may represent the motion consistency score. Furthermore, let $p(h)_i$ represent the likelihood of a color h being present in an image patch i . Player tracking agent 126 may measure the color histogram similarity using Bhattacharyya distance:

$$D_B(p_i, p_j) = -\ln(BC(p_i, p_j)) \text{ with } BC(p_i, p_j) = \sum_h \sqrt{p_i(h)p_j(h)}$$

[0112] At step 906, organization computing system 104 may connect gaps between each set of short tracklets. For example, recall that tracking agent 120 finds the matching pair of tracklets by finding:

$$\arg \max_{ij} (c_{ij} + D_B(p_i, p_j)).$$

[0113] Solving for every pair of broken tracklets may result in a set of clean tracklets, while leaving some tracklets with large, i.e., many frames, gaps. To connect the large gaps, player tracking agent 126 may augment affinity measures to include a motion field estimation, which may account for the change of player direction that occurs over many frames.

[0114] The motion field may be a vector field which measures what direction a player at a point on the playing surface x would be after some time λ . For example, let $X_i = \{x_i^t\}_{t \in (0, T)}$ to be the court position of a player i at every time t . Then, there may exist a pair of points that have a displacement $d_i^\lambda(x_i^t) = x_i^t - x_i^{t+\lambda}$ if $\lambda < T$. Accordingly, the motion field may then be:

$$V(x, \lambda) = G(x, 5) * \sum_i d_i^\lambda(x_i^t),$$

where $G(x, 5)$ may be a Gaussian kernel with standard deviation equal to about five feet. In other words, motion field may be a Gaussian blur of all displacements.

[0115] At step 908, organization computing system 104 may predict a motion of an agent based on the motion field. For example, player tracking system 126 may use a neural network (e.g., neural network 232) to predict player trajectories given ground truth player trajectory. Given a set of ground truth player trajectories X_i , player tracking agent 126 may be configured to generate the set $\hat{V}(x, \lambda)$, where $\hat{V}(x, \lambda)$ may be the predicted motion field. Player tracking agent 126 may train neural network 232 to reduce (e.g., minimize) $\|V(x, \lambda) - \hat{V}(x, \lambda)\|_2^2$. Player tracking agent 126 may then generate the affinity score for any tracking gap of size λ by:

$$K_{ij} = V(x, \lambda) \cdot d_{ij}$$

[0116] where $d_{ij} = x_i^t - x_j^{t+\lambda}$ is the displacement vector between all broken tracks with a gap size of λ . Accordingly, player tracking agent 126 may solve for the matching pairs as recited above. For example, given the affinity score, player tracking agent 126 may assign every pair of broken tracks using a Hungarian algorithm. The Hungarian algorithm (e.g., Kuhn-Munkres) may optimize the best set of matches under a constraint that all pairs are to be matched.

[0117] At step 910, organization computing system 104 may output a graphical representation of the prediction. For example, interface agent 128 may be configured to generate one or more graphical representations corresponding to the tracks for each player generated by player tracking agent 126. For example, interface agent 128 may be configured to generate one or more graphical user interfaces (GUIs) that include graphical representations of player tracking each prediction generated by player tracking agent 126.

[0118] In some situations, during the course of a match, players or agents have the tendency to wander outside of the point-of-view of camera. Such issue may present itself during an injury, lack of hustle by a player, quick turnover, quick transition from offense to defense, and the like.

Accordingly, a player in a first trackable frame may no longer be in a successive second or third trackable frame. Player tracking agent 126 may address this issue via re-identification agent 234.

[0119] Figure 10 is a flow diagram illustrating a method 1000 of tracking players, according to example embodiments. Method 1000 may correspond to operation 410 discussed above in conjunction with Figure 4. Method 1000 may begin at step 1002.

[0120] At step 1002, organization computing system 104 may retrieve a plurality of trackable frames for a match. Each of the plurality of trackable frames may include one or more sets of metadata associated therewith. Such metadata may include, for example, body pose information and camera calibration data. In some embodiments, player tracking agent 126 may further retrieve broadcast video data.

[0121] At step 1004, organization computing system 104 may identify a subset of short tracks in which a player has left the camera's line of vision. Each track may include a plurality of image patches associated with at least one player. An image patch may refer to a subset of a corresponding frame of a plurality of trackable frames. In some embodiments, each track X may include a player identity label y . In some embodiments, each player patch I in a given track X may include pose information generated by data set generator 122. For example, given an input video, pose detection, and trackable frames, re-identification agent 234 may generate a track collection that includes a lot of short broken tracks of players.

[0122] At step 1006, organization computing system 104 may generate a gallery for each track. For example, given those small tracks, re-identification agent 234 may build a gallery for each track. Re-identification agent 234 may build a gallery for each track where the jersey number of a player (or some other static feature) is always visible. The body pose information generated by data set generator 122 allows re-identification agent 234 to determine each player's orientation. For example, re-identification agent 234 may utilize a heuristic method, which may use the normalized shoulder width to determine the orientation:

$$S_{orient} = \frac{\|l_{Lshoulder} - l_{Rshoulder}\|_2}{\|l_{Neck} - l_{Hip}\|_2}$$

where l may represent the location of one body part. The width of shoulder may be normalized by the length of the torso so that the effect of scale may be eliminated. As two shoulders should be apart when a player faces towards or backwards from the camera, re-identification agent 234

may use those patches whose S_{orient} is larger than a threshold to build the gallery. Accordingly, each track X_n , may include a gallery:

$$G_n = \{I_i | S_{orient,i} > thresh\} \forall I_i \in X_n$$

[0123] At step 1008, organization computing system 104 may match tracks using a convolutional autoencoder. For example, re-identification agent 234 may use conditional autoencoder (e.g., conditional autoencoder 240) to identify one or more features in each track. For example, unlike conventional approaches to re-identification issues, players in team sports may have very similar appearance features, such as clothing style, clothing color, and skin color. One of the more intuitive differences may be the jersey number that may be shown at the front and/or back side of each jersey. In order to capture those specific features, re-identification agent 234 may train conditional autoencoder to identify such features.

[0124] In some embodiments, conditional autoencoder may be a three-layer convolutional autoencoder, where the kernel sizes may be 3x3 for all three layers, in which there are 64, 128, 128 channels respectively. Those hyper-parameters may be tuned to ensure that jersey number may be recognized from the reconstructed images so that the desired features may be learned in the autoencoder. In some embodiments, $f(I_i)$ may be used to denote the features that are learned from image i .

[0125] Using a specific example, re-identification agent 234 may identify a first track that corresponds to a first player. Using conditional autoencoder 240, re-identification agent 234 may learn a first set of jersey features associated with the first track, based on for example, a first set of image patches included or associated with the first track. Re-identification agent 234 may further identify a second track that may initially correspond to a second player. Using conditional autoencoder 240, re-identification agent 234 may learn a second set of jersey features associated with the second track, based on, for example, a second set of image patches included or associated with the second track.

[0126] At step 1010, organization computing system 104 may measure a similarity between matched tracks using a Siamese network. For example, re-identification agent 234 may train Siamese network (e.g., Siamese network 242) to measure the similarity between two image patches based on their feature representations $f(I)$. Given two image patches, their feature representations $f(I_i)$ and $f(I_j)$ may be flattened, connected, and fed into a perception network. In some embodiments, L_2 norm may be used to connect the two sub-networks of $f(I_i)$ and $f(I_j)$. In some

embodiments, perception network may include three layers, which include 1024, 512, and 216 hidden units, respectively. Such network may be used to measure the similarity $s(I_i, I_j)$ between every pair of image patches of the two tracks that have no time overlapping. In order to increase the robustness of the prediction, the final similarity score of the two tracks may be the average of all pairwise scores in their respective galleries:

$$S(x_n, x_m) = \frac{1}{|G_n||G_m|} \sum_{i \in G_n, j \in G_m} s(I_i, I_j)$$

[0127] Continuing with the aforementioned example, re-identification agent 234 may utilize Siamese network 242 to compute a similarity score between the first set of learned jersey features and the second set of learned jersey features.

[0128] At step 1012, organization computing system 104 may associate the tracks, if their similarity score is higher than a predetermined threshold. For example, re-identification agent 234 may compute a similarity score be computed for every two tracks that do not have time overlapping. If the score is higher than some threshold, re-identification agent 234 may associate those two tracks may.

[0129] Continuing with the above example, re-identification agent 234 may associated with first track and the second track if, for example, the similarity score generated by Siamese network 242 is at least higher than a threshold value. Assuming the similarity score is higher than the threshold value, re-identification agent 234 may determine that the first player in the first track and the second player in the second track are indeed one in the same.

[0130] Figure 11 is a block diagram 1100 illustrating aspects of operations discussed above in conjunction with method 1000, according to example embodiments.

[0131] As shown block diagram 1100 may include input video 1102, pose detection 1104, player tracking 1106, track collection 1108, gallery building and pairwise matching 1110, and track connection 1112. Block diagram 1100 illustrates a general pipeline of method 1000 provided above.

[0132] Given input video 1102, pose detection information 1104 (e.g., generated by pose detector 212), and player tracking information 1106 (e.g., generated by one or more of player tracking agent 126, auto-clipping agent 120, and camera calibrator 124), re-identification agent 234 may generate track collection 1108. Each track collection 1108 may include a plurality of short broken tracks

(e.g., track 1114) of players. Each track 1114 may include one or more image patches 1116 contained therein. Given the tracks 1114, re-identification agent 234 may generate a gallery 1110 for each track. For example, gallery 1110 may include those image patches 1118 in a given track that include an image of a player in which their orientation satisfies a threshold value. In other words, re-identification agent 234 may generate gallery 1110 for each track 1114 that includes image patches 1118 of each player, such that the player's number may be visible in each frame. Image patches 1118 may be a subset of image patches 1116. Re-identification agent 234 may then pairwise match each frame to compute a similarity score via Siamese network. For example, as illustrated, re-identification agent 234 may match a first frame from track 2 with a second frame from track 1 and feed the frames into Siamese network.

[0133] Re-identification agent 234 may then connect tracks 1112 based on the similarity scores. For example, if the similarity score of two frames exceed some threshold, re-identification agent 234 may connect or associate those tracks.

[0134] Figure 12 is a block diagram illustrating architecture 1200 of Siamese network 242 of re-identification agent 234, according to example embodiments. As illustrated, Siamese network 242 may include two sub-networks 1202, 1204, and a perception network 1205.

[0135] Each of two sub-networks 1202, 1204 may be configured similarly. For example, sub-network 1202 may include a first convolutional layer 1206, a second convolutional layer 1208, and a third convolutional layer 1210. First sub-network 1202 may receive, as input, a player patch I_1 and output a set of features learned from player patch I_1 (denoted $f(I_1)$). Sub-network 1204 may include a first convolutional layer 1216, a second convolutional layer 1218, and a third convolutional layer 1220. Second sub-network 1204 may receive, as input, a player patch I_2 and may output a set of features learned from player patch I_2 (denoted $f(I_2)$). The output from sub-network 1202 and sub-network 1204 may be an encoded representation of the respective player patches I_1, I_2 . In some embodiments, the output from sub-network 1202 and sub-network 1204 may be followed by a flatten operation, which may generate respective feature vectors $f(I_1)$ and $f(I_2)$, respectively. In some embodiments, each feature vector $f(I_1)$ and $f(I_2)$ may include 10240 units. In some embodiments, the L2 norm of $f(I_1)$ and $f(I_2)$ may be computed and used as input to perception network 1205.

Perception network 1205 may include three layers 1222-1226. In some embodiments, layer 1222 may include 1024 hidden units. In some embodiments, layer 1224 may include 512 hidden units.

In some embodiments, layer 1226 may include 256 hidden units. Perception network 1205 may output a similarity score between image patches I_1 and I_2 .

[0136] Figure 13A illustrates a system bus computing system architecture 1300, according to example embodiments. System 1300 may be representative of at least a portion of organization computing system 104. One or more components of system 1300 may be in electrical communication with each other using a bus 1305. System 1300 may include a processing unit (CPU or processor) 1310 and a system bus 1305 that couples various system components including the system memory 1315, such as read only memory (ROM) 1320 and random access memory (RAM) 1325, to processor 1310. System 1300 may include a cache of high-speed memory connected directly with, in close proximity to, or integrated as part of processor 1310. System 1300 may copy data from memory 1315 and/or storage device 1330 to cache 1312 for quick access by processor 1310. In this way, cache 1312 may provide a performance boost that avoids processor 1310 delays while waiting for data. These and other modules may control or be configured to control processor 1310 to perform various actions. Other system memory 1315 may be available for use as well. Memory 1315 may include multiple different types of memory with different performance characteristics. Processor 1310 may include any general purpose processor and a hardware module or software module, such as service 1 1332, service 2 1334, and service 3 1336 stored in storage device 1330, configured to control processor 1310 as well as a special-purpose processor where software instructions are incorporated into the actual processor design. Processor 1310 may essentially be a completely self-contained computing system, containing multiple cores or processors, a bus, memory controller, cache, etc. A multi-core processor may be symmetric or asymmetric.

[0137] To enable user interaction with the computing device 1300, an input device 1345 may represent any number of input mechanisms, such as a microphone for speech, a touch-sensitive screen for gesture or graphical input, keyboard, mouse, motion input, speech and so forth. An output device 1335 may also be one or more of a number of output mechanisms known to those of skill in the art. In some instances, multimodal systems may enable a user to provide multiple types of input to communicate with computing device 1300. Communications interface 1340 may generally govern and manage the user input and system output. There is no restriction on operating on any particular hardware arrangement and therefore the basic features here may easily be substituted for improved hardware or firmware arrangements as they are developed.

[0138] Storage device 1330 may be a non-volatile memory and may be a hard disk or other types of computer readable media which may store data that are accessible by a computer, such as magnetic cassettes, flash memory cards, solid state memory devices, digital versatile disks, cartridges, random access memories (RAMs) 1325, read only memory (ROM) 1320, and hybrids thereof.

[0139] Storage device 1330 may include services 1332, 1334, and 1336 for controlling the processor 1310. Other hardware or software modules are contemplated. Storage device 1330 may be connected to system bus 1305. In one aspect, a hardware module that performs a particular function may include the software component stored in a computer-readable medium in connection with the necessary hardware components, such as processor 1310, bus 1305, display 1335, and so forth, to carry out the function.

[0140] Figure 13B illustrates a computer system 1350 having a chipset architecture that may represent at least a portion of organization computing system 104. Computer system 1350 may be an example of computer hardware, software, and firmware that may be used to implement the disclosed technology. System 1350 may include a processor 1355, representative of any number of physically and/or logically distinct resources capable of executing software, firmware, and hardware configured to perform identified computations. Processor 1355 may communicate with a chipset 1360 that may control input to and output from processor 1355. In this example, chipset 1360 outputs information to output 1365, such as a display, and may read and write information to storage device 1370, which may include magnetic media, and solid state media, for example. Chipset 1360 may also read data from and write data to RAM 1375. A bridge 1380 for interfacing with a variety of user interface components 1385 may be provided for interfacing with chipset 1360. Such user interface components 1385 may include a keyboard, a microphone, touch detection and processing circuitry, a pointing device, such as a mouse, and so on. In general, inputs to system 1350 may come from any of a variety of sources, machine generated and/or human generated.

[0141] Chipset 1360 may also interface with one or more communication interfaces 1390 that may have different physical interfaces. Such communication interfaces may include interfaces for wired and wireless local area networks, for broadband wireless networks, as well as personal area networks. Some applications of the methods for generating, displaying, and using the GUI disclosed herein may include receiving ordered datasets over the physical interface or be generated

by the machine itself by processor 1355 analyzing data stored in storage 1370 or 1375. Further, the machine may receive inputs from a user through user interface components 1385 and execute appropriate functions, such as browsing functions by interpreting these inputs using processor 1355.

[0142] It may be appreciated that example systems 1300 and 1350 may have more than one processor 1310 or be part of a group or cluster of computing devices networked together to provide greater processing capability.

[0143] While the foregoing is directed to embodiments described herein, other and further embodiments may be devised without departing from the basic scope thereof. For example, aspects of the present disclosure may be implemented in hardware or software or a combination of hardware and software. One embodiment described herein may be implemented as a program product for use with a computer system. The program(s) of the program product define functions of the embodiments (including the methods described herein) and can be contained on a variety of computer-readable storage media. Illustrative computer-readable storage media include, but are not limited to: (i) non-writable storage media (e.g., read-only memory (ROM) devices within a computer, such as CD-ROM disks readably by a CD-ROM drive, flash memory, ROM chips, or any type of solid-state non-volatile memory) on which information is permanently stored; and (ii) writable storage media (e.g., floppy disks within a diskette drive or hard-disk drive or any type of solid state random-access memory) on which alterable information is stored. Such computer-readable storage media, when carrying computer-readable instructions that direct the functions of the disclosed embodiments, are embodiments of the present disclosure.

[0144] It will be appreciated to those skilled in the art that the preceding examples are exemplary and not limiting. It is intended that all permutations, enhancements, equivalents, and improvements thereto are apparent to those skilled in the art upon a reading of the specification and a study of the drawings are included within the true spirit and scope of the present disclosure. It is therefore intended that the following appended claims include all such modifications, permutations, and equivalents as fall within the true spirit and scope of these teachings.

What is Claimed:

1. A method of generating a player tracking prediction, comprising:
 - retrieving, by a computing system, a broadcast video feed for a sporting event, the broadcast video feed comprising a plurality of video frames;
 - segmenting, by the computing system, the broadcast video feed into a unified view, wherein the unified view comprises a plurality of trackable frames, the plurality of trackable frames is a subset of the plurality of video frames;
 - generating, by the computing system, a plurality of data sets based on the plurality of trackable frames, wherein the plurality of data sets comprises playing surface segmentation information, ball tracking information, and body pose information for each player in each trackable frame;
 - calibrating, by the computing system, a camera associated with each trackable frame based on the playing surface segmentation information and the body pose information;
 - generating, by the computing system, a plurality of sets of short tracklets based on the plurality of trackable frames and the body pose information;
 - connecting, by the computing system, each set of short tracklets by generating a motion field vector for each player in the plurality of trackable frames; and
 - predicting, by the computing system, a future motion of a player based on the player's motion field vector using a neural network.
2. The method of claim 1, wherein segmenting, by the computing system, the broadcast video feed into a unified view comprises:
 - parsing the broadcast video feed to identify a first subset of video frames corresponding to a same view of the sporting event; and
 - discarding a second subset of video frames corresponding to a different view of the sporting event.
3. The method of claim 1, wherein generating, by the computing system, the plurality of data sets based on the plurality of trackable frames comprises:

generating body pose information for each player in each trackable frame of the plurality of trackable frames.

4. The method of claim 3, wherein calibrating, by the computing system, the camera associated with each trackable frame based on the playing surface segmentation information and the body pose information comprises:

identifying a pattern of motion between two successive trackable frames by identifying players in each frame using the body pose information and removing, from each trackable frame, each player.

5. The method of claim 1, wherein generating, by the computing system, the plurality of data sets based on the plurality of trackable frames comprises:

identifying, in each trackable frame of the plurality of trackable frames, playing surface markings using a trained neural network configured to identify the playing surface markings.

6. The method of claim 5, wherein calibrating, by the computing system, the camera associated with each trackable frame based on the playing surface segmentation information and the body pose information comprises:

retrieving the playing surface markings generated by the trained neural network; and
comparing the playing surface markings to a template of playing surface markings to generate one or more keyframes.

7. The method of claim 1, wherein predicting, by the computing system, the future motion of the player based on the player's motion field vector using a neural network comprises:

constructing player motion throughout the sporting event using the plurality of trackable frames, the body pose information, and the calibrated camera.

8. A system for generating a player tracking prediction, comprising:
a processor; and

a memory having programming instructions stored thereon, which, when executed by the processor, performs one or more operations comprising:

retrieving a broadcast video feed for a sporting event, the broadcast video feed comprising a plurality of video frames;

segmenting the broadcast video feed into a unified view, wherein the unified view comprises a plurality of trackable frames, the plurality of trackable frames is a subset of the plurality of video frames;

generating a plurality of data sets based on the plurality of trackable frames, wherein the plurality of data sets comprises playing surface segmentation information, ball tracking information, and body pose information for each player in each trackable frame;

calibrating a camera associated with each trackable frame based on the body pose information;

generating a plurality of sets of short tracklets based on the plurality of trackable frames and the body pose information;

connecting each set of short tracklets by generating a motion field vector for each player in the plurality of trackable frames; and

predicting a future motion of a player based on the player's motion field vector using a neural network.

9. The system of claim 8, wherein segmenting the broadcast video feed into a unified view comprises:

parsing the broadcast video feed to identify a first subset of video frames corresponding to a same view of the sporting event; and

discarding a second subset of video frames corresponding to a different view of the sporting event.

10. The system of claim 8, wherein generating the plurality of data sets based on the plurality of trackable frames comprises:

generating body pose information for each player in each trackable frame of the plurality of trackable frames.

11. The system of claim 10, wherein calibrating the camera associated with each trackable frame based on the playing surface segmentation information and the body pose information comprises:

identifying a pattern of motion between two successive trackable frames by identifying players in each frame using the body pose information and removing, from each trackable frame, each player.

12. The system of claim 8, wherein generating the plurality of data sets based on the plurality of trackable frames comprises:

identifying, in each trackable frame of the plurality of trackable frames, playing surface markings using a trained neural network configured to identify the playing surface markings.

13. The system of claim 12, wherein calibrating the camera associated with each trackable frame based on the playing surface segmentation information and the body pose information comprises:

retrieving the playing surface markings generated by the trained neural network; and
comparing the playing surface markings to a template of playing surface markings to generate one or more keyframes.

14. The system of claim 8, wherein predicting the future motion of the player based on the player's motion field vector using a neural network comprises:

constructing player motion throughout the sporting event using the plurality of trackable frames, the body pose information, and the calibrated camera.

15. A non-transitory computer readable medium including one or more sequences of instructions that, when executed by one or more processors, perform one or more operations comprising:

retrieving, by a computing system, a broadcast video feed for a sporting event, the broadcast video feed comprising a plurality of video frames;

segmenting, by the computing system, the broadcast video feed into a unified view, wherein the unified view comprises a plurality of trackable frames, the plurality of trackable frames is a subset of the plurality of video frames;

generating, by the computing system, a plurality of data sets based on the plurality of trackable frames, wherein the plurality of data sets comprises playing surface segmentation information, ball tracking information, and body pose information for each player in each trackable frame;

calibrating, by the computing system, a camera associated with each trackable frame based on the playing surface segmentation information and the body pose information;

generating, by the computing system, a plurality of sets of short tracklets based on the plurality of trackable frames and the body pose information;

connecting, by the computing system, each set of short tracklets by generating a motion field vector for each player in the plurality of trackable frames; and

predicting, by the computing system, a future motion of a player based on the player's motion field vector using a neural network.

16. The non-transitory computer readable medium of claim 15, wherein segmenting, by the computing system, the broadcast video feed into a unified view comprises:

parsing the broadcast video feed to identify a first subset of video frames corresponding to a same view of the sporting event; and

discarding a second subset of video frames corresponding to a different view of the sporting event.

17. The non-transitory computer readable medium of claim 15, wherein generating, by the computing system, the plurality of data sets based on the plurality of trackable frames comprises:

generating body pose information for each player in each trackable frame of the plurality of trackable frames.

18. The non-transitory computer readable medium of claim 17, wherein calibrating, by the computing system, the camera associated with each trackable frame based on the playing surface segmentation information and the body pose information comprises:

identifying a pattern of motion between two successive trackable frames by identifying players in each frame using the body pose information and removing, from each trackable frame, each player.

19. The non-transitory computer readable medium of claim 15, wherein generating, by the computing system, the plurality of data sets based on the plurality of trackable frames comprises:

identifying, in each trackable frame of the plurality of trackable frames, playing surface markings using a trained neural network configured to identify the playing surface markings.

20. The non-transitory computer readable medium of claim 19, wherein calibrating, by the computing system, the camera associated with each trackable frame based on the playing surface segmentation information and the body pose information comprises:

retrieving the playing surface markings generated by the trained neural network; and

comparing the playing surface markings to a template of playing surface markings to generate one or more keyframes.

100




FIG. 1





FIG. 2




400

FIG. 4

5/13



500

FIG. 5



7/13

700




FIG. 7




FIG. 8

9/13


900

FIG. 9

10/13



1000

FIG. 10






FIG. 12

FIG. 13A





FIG. 13B



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US20/20442

A. CLASSIFICATION OF SUBJECT MATTER

IPC - G06T 7/20; G06K 9/00; G06N 3/04 (2020.01)

CPC - G06K 9/00342; G06T 7/20, 7/292; G06Q 10/0639; G06T 2207/30221, 2207/30224

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

See Search History document

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

See Search History document

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 2017/210564 A1 (PILLAR VISION, INC.) 07 December 2017; entire document.	1-20
A	WO 2006/103662 A2 (SPORTVU LTD.) 05 October 2006; entire document.	1-20
A	US 2012/0316843 A1 (BENO, T. et al.) 13 December 2012; entire document.	1-20

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"D" document cited by the applicant in the international application

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

21 April 2020 (21.04.2020)

Date of mailing of the international search report

07 MAY 2020

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Shane Thomas

Telephone No. PCT Helpdesk: 571-272-4300