



Toward Data Driven Development in RoboCup

Heinrich Mellmann^(✉), Benjamin Schlotter, and Philipp Strobel

Adaptive Systems Group, Humboldt-Universität zu Berlin,
Unter den Linden 6, 10099 Berlin, Germany
mellmann@informatik.hu-berlin.de

Abstract. Conducting games in RoboCup incurs high cost in terms of effort, time, and money. The scientific outcome, however, is quite limited and often not very conclusive. Especially, analyzing and drawing conclusions about the performance of complex processes like decision making of an individual robot or the behavior on the team level poses a considerable challenge. Collecting more data during the competition games will help to analyze the performance of algorithms, identify errors and areas for improvement, and make more significant statements regarding the performance of the robots. In this work we investigate the possibilities for collection of the large scale RoboCup data and its analysis. We present a system for automatic recording of synchronized videos of RoboCup games and an application for exploration and annotation of large sets of RoboCup-related data. We also present data sets collected during the competitions in 2018 and an algorithm for visual detection and tracking of robots in the RoboCup videos. A first empirical evaluation shows promising results and demonstrates how such data can be integrated and used to validate robot's behavior.

Keywords: Data driven development · Robot detection · Visual robot tracking · Camera localization

1 Introduction

An empirical scientific discipline requires a set of methods and practices for evaluation and comparison of proposed models and solutions. RoboCup provides a unique *common test scenario* for robotics, but its potential is, by far, not realized. Conducting games incurs high cost in terms of effort, time, and money. The scientific outcome, however, is quite limited and often not very conclusive. In most cases only the final score of the games provides feedback about the performance of a team.

Our aim is to shift the focus of the games more towards scientific experiments. This requires a comprehensive data collection infrastructure in order to enable empirical analysis and to generate quantifiable results. This might provide a basis for analysis of complex high level algorithms like team strategy. This might in

turn enable for a quantifiable analysis of the overall progress of the league over the years.

Collecting more data during the competition games will help to analyze the performance of algorithms, and identify errors and areas for improvement. In particular, analysis of the team behavior, like role change, require global synchronized data. This kind of data can be more easily collected in a controlled experiment in a lab environment. However, the algorithms running on a robot tend to behave differently in the isolated environment of the lab compared to real conditions during a competition. From this perspective each RoboCup competition can be seen as a large scale collective experiment evaluating the abilities of the robots. In order to build understanding based on this large scale experiment, we need to capture data more effectively and to make this rich data comparable across the community and over time. For human soccer extensive data collected from games is available which allows researchers to learn behavior models of teams from tracking data. For example in [5] the authors use such models to answer *what if* type of questions.

Teams in the Standard Platform League already collect data during the games. Common team communication and data sent by the game control computer during the game is recorded and made publicly available.¹ Individual teams are collecting the data they need to analyze their team's performance and to solve bugs, e.g., images for ball detection. Many teams record videos of their games. All this data is however unorganized and there are no shared tools and common practices to effectively analyze and evaluate large amounts of data specific to RoboCup. At the same time many teams do not have access to extensive logging infrastructure, this includes in particular new teams joining RoboCup.

So far there have been only few attempts to analyze the game data on a large scale within the Standard Platform League (SPL). A noteworthy example is the technical report [3] by the team Nao Devils from 2016, where the authors analyze the final outdoor game. And in [8] the team B-Human analyses their robot positions in final game against Team HTWK at the RoboCup 2017. Their analysis shows that the both teams do not utilize the corners of the field which might be used for future tactic improvements. In [7] we extracted and annotated events where the robots kicked the ball to draw conclusions about the performance of the proposed decision making mechanism. In the continuation of the work [6], more complex decisions were introduced and evaluated in simulated experiment. A rigorous empirical analysis in this case would require more extensive data collection and more sophisticated analysis methods. These publications outline the motivation and the direction of our work.

The work presented in this paper is the result of an ongoing effort to push forward the development of an ecosystem of tools and practices to support collection, organization and analysis of large amounts of RoboCup-specific data enabling detailed analysis and promoting data driven research and development in RoboCup. Although our current work focuses on the SPL, the developed tools and methods are applicable to other leagues.

¹ <https://spl.robocup.org/downloads/>.

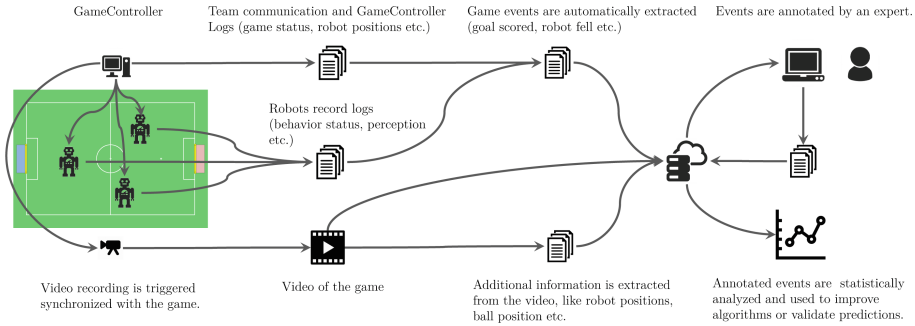


Fig. 1. Overview over the data flow of the implemented data processing ecosystem.

The remainder of the paper is structured as follows. The Sect. 2 gives an overview over the different components of the data collection system and the flow of the data. The Sect. 3 describes the automatic recording system for automatic collection of synchronized RoboCup videos. In the Sect. 4 we discuss the data collected during German Open and RoboCup in 2018. Section 5 describes the *RoboCup Explorer*, a server based application for inspection and annotation of RoboCup data. In the Sect. 6 we present a method for the detection of robots in the game videos based on deep learning. The position of the camera which recorded the video can be determined through alignment of the field lines in the video. This allows to track the detected robots in the global field coordinates resulting in a low cost visual tracking system. The details of the tracking are discussed in Sect. 7 alongside first empirical results.

Technical information about the system as well as links to the open source repository and the collected data can be found on the project page [2].

2 System Overview and Data Flow

At the current point we distinguish between three types of sources from which data can be gathered during a RoboCup soccer game: The *video* of the game, recorded by cameras on the side of the field; The *network communication* between the robots and data sent by the *game control computer* to the robots; and *log data* recorded by each individual robot.

Figure 1 depicts the used components and the data flow of the implemented system. For the collection of the data we utilize existing infrastructure within the SPL. In particular we make use of the *GameController*[1], a specialized software, which is used to control the games within SPL. The network communication of the GameController and all the robots is recorded and made publicly available. The data communicated by the robots contains such information like the estimated position on the field, perceived ball, but also information about the robot’s role. We extend the data collection with the automatic recording system of synchronized game videos, which will be discussed in more details in the Sect. 6.

After a game is completed, the video as well as the log files of the Game-Controller and the log files of each individual robot are collected and stored on a server. The collected data is organized by a defined scheme and named by the date, time and participating teams. On the server, a pipeline of different processing tasks are executed. In particular, spatial and temporal information is extracted from communication data and individual log files. Spatial information may contain robot’s estimated position on the field, perceptions of the ball etc. Temporal information may contain game events like changes of the game state (*READY*, *PLAY*, etc.) or penalties, but also other events like cases where a robot fell or a particular decision was made or a particular action executed, e.g., a kick (cf. Sect. 5). Further information can be extracted from the recorded video. At the current moment we focus on detecting and tracking robots in the video, as discussed in Sect. 5 in more detail. In the future, the detected robot trajectories could be integrated with the data from the log files to reconstruct a full and reliable view of the robot’s behavior.

The collected and processed data can be inspected and further enriched with annotations with the *RoboCup Explorer*, which will be discussed in Sect. 5 in more details. The annotations can be used for empirical evaluation of the behavior patterns and decisions similarly to [9], and perhaps as input for machine learning approaches in the future.

3 Collection of Videos and Data

A central piece of the collected data are videos of the games taken from a global perspective. In order to generate and retrieve the ground truth for the whole game, those videos need to capture the whole field with all robots visible. It is also crucial for the videos to be synchronized with other data collected from the network and from the individual robots, to easily compare them with the generated ground truth. A low cost system, that is quickly deployable and largely automatic, can be best adopted and used by other teams throughout the league. In this section we give a brief overview over a low cost system for automatic synchronized recording of RoboCup videos.

In the Standard Platform League (SPL) the games are controlled by a dedicated computer with a specialized software called *GameController*. The Game-Controller is operated by a human and acts effectively as a translator of the commands given by the human main referee to the robots. The changes in the state of the game, like *READY*, *PLAY* or *FINISH*, and other relevant information are broadcast to the robots through WiFi. Our system is utilizing these messages to synchronize the data from different sources. They are recorded by each individual robot, they are part of the recorded network communication, and are used to start and stop the recording of the video of the game.

For the recording of the videos we currently use GoPro cameras (tested with GoPro Session and GoPro Hero 5 Black). They have a wide angle and are capable of capturing the whole field when placed on the side of the field. These cameras are widely available and provide high video quality. The recording can be triggered over a WiFi interface of the camera.

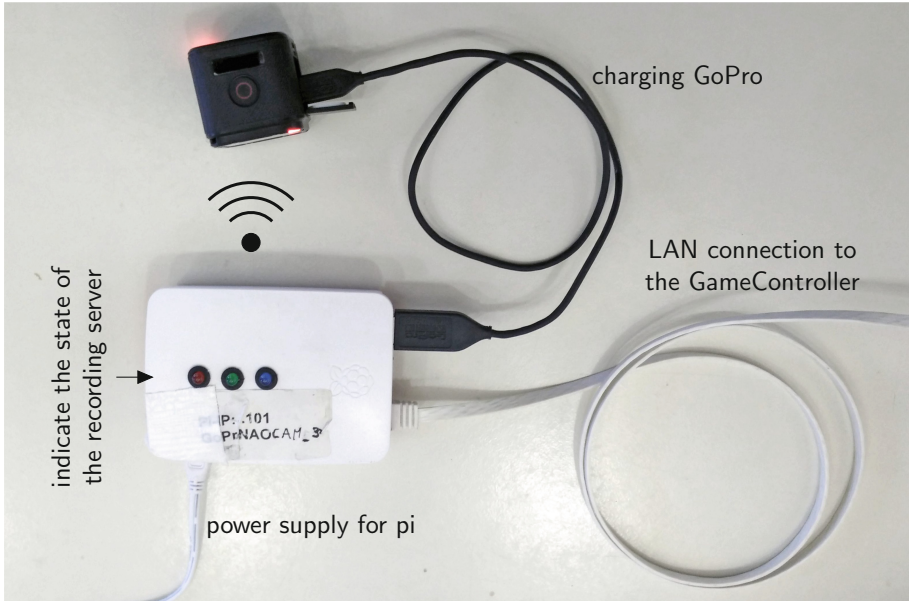


Fig. 2. System for automatic recording of synchronized RoboCup videos. The recording server is based on raspberry pi with the colored indicator LEDs. The GoPro is charged via USB from the raspberry. A LAN cable connects the pi to the GameController network. The Recording server sends commands to the GoPro via WiFi.

The central piece of the video recording system is a *recording server*. The code for the recording server software has been made open source and can be found through the project page [2]. The task of the recording server is to control the recording process of the video, in particular, to start and stop recording of the video based on the messages received from the GameController. The current version of the recording server is running on a *Raspberry Pi 3 Model B*. It is very compact and is equipped with a LAN, WiFi, USB and GPIO pins where additional components can be attached. The recording server is connected to the GameController network through LAN interface and simultaneously to the GoPro camera through the WiFi. It has been equipped with indicator LED-lights in order to enable a fast and robust monitoring of the system's health during a competition. The lights indicate whether the connection to the camera could be established successfully, whether a GameController is currently online or if a recording is running. The overview over the recording system can be seen in the Fig. 2.

The system requires to be set up only once at the beginning of the competition and ensures that each game is recorded and *synchronized* with the corresponding network data. At the current stage, the videos are recorded on the memory card (microSD) of the GoPro camera and can be retrieved in the breaks between the games or at the end of the day. Depending on the available Internet



Fig. 3. Left: An example image from a recording at the German Open 2018. Right: An example image from a recording at the RoboCup 2018

connection, collected data can be immediately uploaded to a server or directly copied by the teams to be used *during* the competition.

The system was tested and used to collect synchronized communication data and video at several events during 2018 and 2019. Most notable, the data gathered during the German Open 2018 and the RoboCup 2018 will be discussed in more details the Sect. 4.

4 Data Sets from GO18 and RC18

The recording system described in the Sect. 3 was successfully deployed during the German Open 2018 and at the RoboCup 2018 in Montreal. The videos and data from both competitions were made publicly available and can be found on our project website [2]. In the following we give a brief overview over the collected data sets.

German Open was played on a single field. For the game recording one camera, GoPro Session, was installed on the side of the field in about 4 m height. With that, the entire playing field could be recorded with only small occasional occlusions by the referees or other persons on the field. An example can be seen in Fig. 3. All of 23 played competition games were completely and correctly recorded. The system was running completely autonomously without errors. All recorded video were made available to interested teams on site and can be found on our project website [2].

RoboCup 2018 in Montreal was played on 5 fields in parallel (numbered A–E), which posed a significant infrastructural challenge. For each of the fields a separate video recording system was deployed. In this case, two different versions of the GoPro cameras were used: *GoPro Session* on the fields A, B and E, and *GoPro Hero 5 Black* on the fields C and D. All Cameras were installed on the side of the field in 2.5 m height.

During the RoboCup 2018 we collected approximately 500 GB of data. The GoPro Videos constitute about 450 GB. The rest is mostly the log data collected by the individual robot of our team, Berlin United. The GameController logs have a negligible size with only 20 MB per halftime. During the competition a number of noteworthy issues have occurred which led to some videos being

Table 1. Evaluation of the video recording during the RoboCup 2018 in Montreal

Field	Camera	Missing	Partially	Complete	Sum
A	GoPro Hero 4 Session	3		19	22
B	GoPro Hero 4 Session	4		12	16
C	GoPro Hero 5 Black	2	5	9	16
D	GoPro Hero 5 Black	2	4	8	14
E	GoPro Hero 4 Session	2		6	8
Sum		13	9	54	76

incomplete, corrupted or missing. In the following we briefly outline those. The evaluation of the recorded games reveals, that of the totally 76 games played

- 54 games were completely recorded
- 9 games were only partially recorded, with only few seconds missing in some cases, but also a complete half
- 13 games were not recorded at all.

The Table 1 shows the statistics of the collected videos in more detail.

The main reason for missing recordings turned out to be the teams unplugging the power supply of the recording system or the LAN connection to the GameController, and re-purposing the ports for their needs, because of the limited number of ports. Two of the incomplete videos were fractured in a large number of segments. Which was most likely caused, by a second instance of a GameController on the same network, probably accidentally started for testing by one of the teams. The two competing instances would start and stop recording in repeated succession. This behavior was verified in an isolated experiment. To prevent this issue we plan to add a configuration option to the raspberry pi server, in order to set the IP address of the gamecontroller and ignore messages of unofficial instances. In another approach, once a recording has started, messages of other instances of the gamecontroller are discarded and therefore prevent fractured video files.

Another minor issue was the stability of the WiFi connection between the raspberry pi and the GoPro Hero 5. The reason for the disconnections could not be fully explained. Probably the interference of other WiFi networks or on a similar frequency sending devices was a factor. Another issue observed with the GoPro Hero 5 was, that the camera draws more power during the video capturing than could be provided by the Raspberry Pi. In some cases, this resulted in incomplete or missing videos after the battery was depleted and the camera shut down. After some research we found that the GoPro Hero 5 can be charged via USB-PD 3.0. With an USB-C charger we were able to provide sufficient power to charge both, the raspberry pi and the GoPro Hero 5.

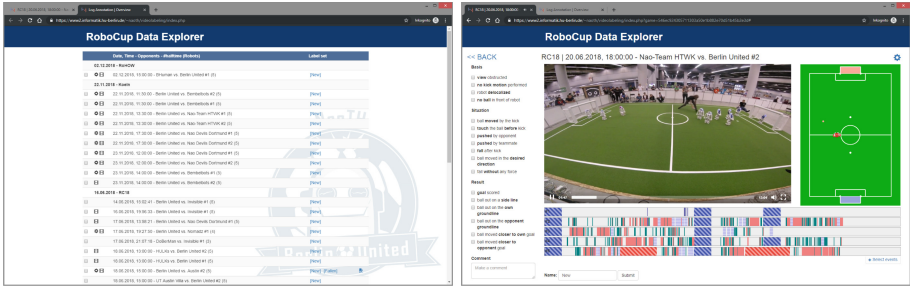


Fig. 4. Left: overview over the data collected during different RoboCup events. Right image: example session for inspection and event annotation for the first half of the game with the team *Nao Team HTWK* at the RoboCup 2018. Main components are timelines with events represented by colored buttons (bottom); visualization of the robots state, position on the field and perceived ball (right) and possible labels for the events (left).

The recorded matches were made available to the interested teams during and after the competition. After the preparation and evaluation the videos were uploaded to YouTube. The links can be found on the project website [2].

5 RoboCup Data Explorer

The *RoboCup Data Explorer* is a web based server application which has been developed for organization, inspection and annotation of large amounts of RoboCup data (videos, network communication and individual logs).

The data can be stored on the same server or provided by another, e.g., the game video are stored on YouTube. The RoboCup Data Explorer provides two different views on the data. The data set view shows a list of *data units* sorted by the event and the game half as illustrated in Fig. 4 (left). Each of these data units contains the data collect during a contiguous game sequence, which is usually a half of an SPL game. Such a data unit can be inspected in a inspection and annotation view as illustrated in Fig. 4 (right).

As can be seen in the Fig. 4 (right), the inspection and annotation view has four main components: video of the game (center), global field view (right), game events (bottom) and annotations panel (left). The communication data as well as data collected by the individual robots can be presented in spatial manner in the global field view, e.g., positions of the robots and the ball, and in form of events in the event view, e.g., game states, penalties, fall downs. In order to be displayed the events and spatial information is automatically extracted from the log files and saved in a unified event log format. The events and spatial information can also be extracted from the videos as discussed in Sect. 6.

The human operator can navigate the game data by selecting the events of interest and jumping to the corresponding position in the video. Selected event can also be annotated to gather human feedback for quantitative analysis or as

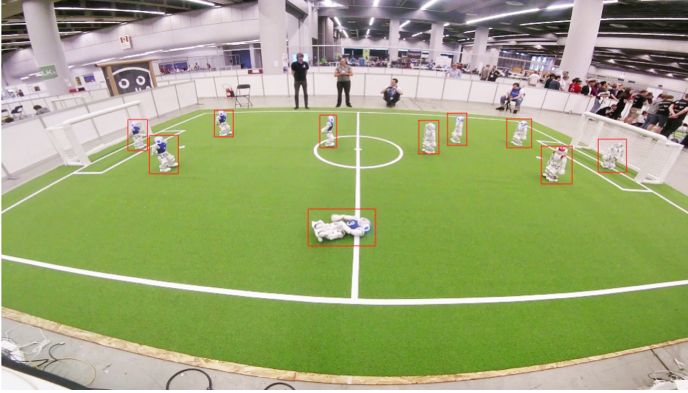


Fig. 5. Example of detected robots in a video recorded with a GoPro camera at the RoboCup 2018 in Montreal/Canada.

training data for machine learning. The annotation tool was used to annotate different kick actions executed by the robots in the videos recorded during the games at the RoboCup in 2015. The kick events were automatically extracted from the log files recorded by the individual robots and synchronized with the video. The results were used to evaluate the performance of the kick decision algorithm and were published in [7].

6 Robot Detection in Videos

Additional information can be automatically extracted from the video using machine vision techniques. Detected robots and ball can be used to verify the perception of the robots or as input for automatic referee system in a future project [10]. On standard datasets like *COCO* or *PASCAL VOC* the state of the art object detection algorithms already achieve impressive results. Since there is no publicly available detector trained to detect Nao robots from that view we used transfer learning to retrain a state of the art detector to detect the Nao robots.

For our experiments we used DC-SPP-YOLO [4] which is based on the YOLO object detection algorithm. To speed up the training performance pretrained weights for the convolutional layers provided by the authors of the YOLO papers were used². Figure 5 illustrates an example for an SPL game scene with detected robots. The training was performed on the video recording of the quarterfinals game³ between the teams *Nao Team HTWK* and *Berlin United* from the RoboCup 2018 as described in Sect. 4. In total 3300 labeled images from the first half of a game were used. Additional 200 images were used for validation during training and further 2500 images were used for testing.

² <https://pjreddie.com/media/files/darknet53.conv.74>.

³ <https://youtu.be/KKtQmRtP06A>.

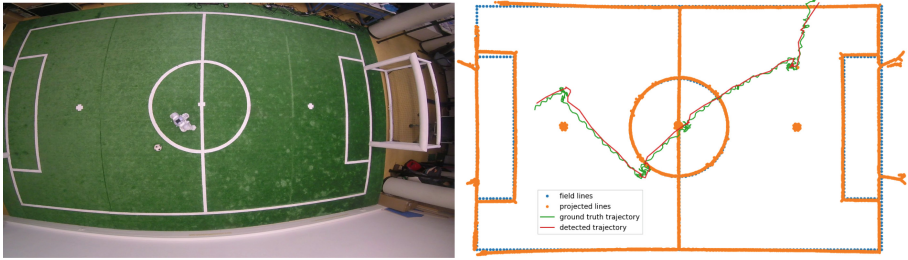


Fig. 6. Tracking of the robot in the laboratory conditions. Left: a frame from the recorded video. Right: trajectory of the robot determined based on the video in comparison to the ground truth recorded by the OptiTrack system. Projected line points are also shown in comparison to the model of the field. Note the goal posts being treated as lines.

On the 2500 test images the mean average precision is 0.999786. The precision and recall values are both 1.0 which means that the correct number of robots were detected in each images and the average IOU for the robot detection is 87.41%. So far the robot detection was tested only on a very limited dataset taken from one game. In recordings of other games the performance of the detection might be impacted by different lighting, different jersey colors and different height and angle of the camera. The results show however that accurate robot detection is possible given enough labeled data. The next section shows that based on the presented robot detection a fairly accurate localization of the robots is possible as well.

7 Visual Robot Tracking

In this section we explore the feasibility of visual tracking of the positions of the robots during the game in RoboCup videos as recorded in Sect. 4. For this we localize the robots detected in the Sect. 6 in relation to the global coordinates of the field. In order to do so, we need to estimate the position and rotation of the camera as well as distortion parameter of the lens.

The distortion of the lens can be determined with standard methods. Here it is important to note, that the GoPro-Videos are recorded in a so called *Super-View* format, which distorts the image additionally in the post-processing and needs to be corrected separately.

In order to localize the camera we use the field lines. The field lines are detected by a simple thresholding. The line points in the image are projected on the field and matched with the expected field lines in a nearest neighbor fashion. The location and rotation of the camera are determined by a classical minimization routine, such that the matching distance error is minimized.

We conducted an isolated experiment in order to evaluate the accuracy of the visual tracking. The robot NAO is dribbling the ball into a goal while being

tracked both in the recorded video and by a dedicated tracking system OptiTrack⁴. The used OptiTrack system consists of 10 infrared cameras being able to track an object marked by reflective markers with a precision of less than 1 mm and frame rate of 120 frames per second. The markers were placed on the head of the robot as can be seen in the Fig. 6 (left).

The GoPro camera was localized with a mean error of 66 mm on the projected line points. This camera position was used to project the detected robot in the video. The robot was detected in the video as described in the Sect. 6. The center of the bounding box was used as an approximation of the center of the robots body and projected with the assumed height of 260 mm.

The resulting projected trajectory of the robot is shown alongside the ground truth recorded by the OptiTrack system in Fig. 6 (right). The error between the both trajectories is 14.93 mm, while maximal deviation is 98 mm. Considering the comparably simple and rudimentary setup of the system and calibration of the camera, the precision is surprisingly high.

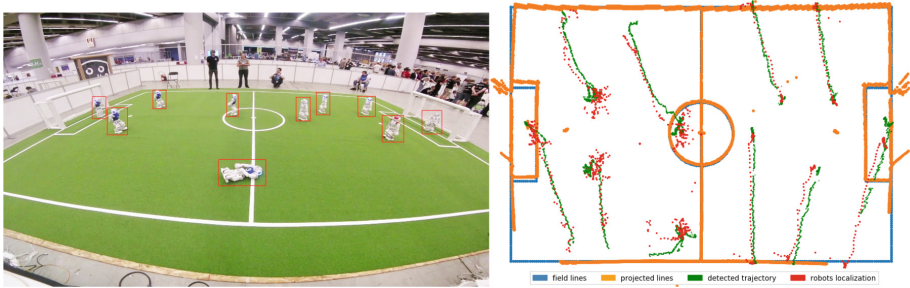


Fig. 7. Tracking of the robots on field conditions during the *Ready* phase. Left: final positions of the robots in the *READY* phase. Right: trajectories of the robots determined based on the video [green] in comparison to the communicated positions of the robots [red]. Projected line points are also shown in comparison to the model of the field. Note the goal posts being treated as lines (Color figure online).

For the second experiment we used a game recording of the quarterfinals between the teams *Nao Team HTWK* and *Berlin United* from the RoboCup 2018 as described in Sect. 6. For this video the camera was localized with an estimated precision of 66 mm on the projected line points. Figure 7 (right) illustrates the trajectories of the robots in the *READY*-phase, i.e., the first 45 s of the game, where robots assume their kickoff positions. The detected positions (green) are shown alongside the positions determined by the robots (red), which have been extracted from the network communication. Estimated deviation between both positions for the localization of the HTWK robots is mean 87 mm with standard deviation of 82 mm. For the Berlin United robots the mean deviation is 55 mm and standard deviation 44 mm.

⁴ <https://optitrack.com/>.

The results of both experiments constitute a basic proof of concept for the feasibility of visual robot tracking in the videos recorded with a wide angle camera. Merged with the communicated localization of the robots a reliable motion trajectory for all robots can be reconstructed. Possible future applications include verification of the robot localization, live tracking, and analysis of the behavior patterns.

At the current stage the algorithm contains a number of heuristics and simplifying assumptions, which may introduce errors. Most notably, the assumed fixed height of the robots center used for projection, might vary depending on the walking height of the robot and doesn't hold for fallen robots. Replacing those heuristics with more accurate estimates might significantly improve the accuracy of the final position estimation.

Note that the identification of a robot's player number and its belonging to the team is currently not implemented. To solve this the player numbers could be assigned manually at the beginning and tracked through the game. Positions estimated by the robots' own self localization and communicated through WiFi can be used to estimate their numbers as well. The classification of a robot's team could be done based on the color of the jerseys.

8 Conclusions

We presented the current progress of our ongoing work towards a data driven development and analysis in RoboCup. The key components of the ecosystem presented in the paper are the automatic recording system for recording synchronized videos of RoboCup games and the RoboCup Explorer - an application for organization and annotation of the RoboCup data. We discussed the data collected during the German Open and RoboCup in 2018. Based on this data we developed a method for detection of the robots in the videos and their tracking on the field.

The results discussed in the paper show a clear path for the future development. On one hand we will continue to improve the tools and the process for collection and organization of the data, e.g., in order to reduce the risk of game video not being recorded. A more extensive analysis of the robot detection is needed in order to evaluate the generalization of the presented work. Furthermore visual tracking methods can be used to increase the robustness of the detection. On the other hand the ecosystem reached enough maturity allowing us to push forward to work on methods for analysis of complex behavior patterns of the robots and bring machine learning techniques to the level of decision making and planing. For example the effect of behavior decisions such as *how many robots stay near the own goal when the ball is in the opponent half* on the game score or other metrics can be analyzed over multiple teams and competitions.

Acknowledgements. This work was supported by the RCF grants in 2017 and 2018.

References

1. B-Human: GameController 2018 - RoboCup Edition (2018). <https://spl.robocup.org/downloads/>. Accessed 22 Apr 2019
2. Berlin United: Tools for Data Driven Research and Development in RoboCup (2019). <http://www.robocup.tools>. Accessed 22 Apr 2019
3. Hofmann, M., Moos, A., Rensen, F., Schwarz, I., Urbann, O.: Playing robot soccer outdoor. In: The 11th Workshop on Humanoid Soccer Robots at 16th IEEE-RAS International Conference on Humanoid Robots (2016). <http://d-fence.sytes.net/research/files/Playing%20Soccer%20Outdoors%20with%20Humanoid%20Robots.pdf>
4. Huang, Z., Wang, J.: DC-SPP-YOLO: dense connection and spatial pyramid pooling based YOLO for object detection. CoRR (2019). [arXiv:1903.08589](https://arxiv.org/abs/1903.08589)
5. Le, H.M., Carr, P., Yue, Y., Lucey, P.: Data-driven ghosting using deep imitation learning (2017)
6. Mellmann, H., Schlotter, B.: Advances on simulation based selection of actions for a humanoid soccer-robot. In: Proceedings of the 12th Workshop on Humanoid Soccer Robots, 17th IEEE-RAS International Conference on Humanoid Robots (Humanoids), Madrid, Spain (2017)
7. Mellmann, H., Schlotter, B., Blum, C.: Simulation based selection of actions for a humanoid soccer-robot. In: Behnke, S., Sheh, R., Sariel, S., Lee, D.D. (eds.) RoboCup 2016. LNCS (LNAI), vol. 9776, pp. 193–205. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68792-6_16
8. Röfer, T., Laue, T., Hasselbring, A., Richter-Klug, J., Röhrig, E.: B-Human 2017 – team tactics and robot skills in the standard platform league. In: Akiyama, H., Obst, O., Sammut, C., Tonidandel, F. (eds.) RoboCup 2017. LNCS (LNAI), vol. 11175, pp. 461–472. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00308-1_38
9. Zhu, D., Veloso, M.: Virtually adapted reality and algorithm visualization for autonomous robots. In: Behnke, S., Sheh, R., Sariel, S., Lee, D.D. (eds.) RoboCup 2016. LNCS (LNAI), vol. 9776, pp. 452–464. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68792-6_38
10. Zhu, D., Veloso, M.: Event-based automated refereeing for robot soccer. *Auton. Robots* **41**(7), 1463–1485 (2017). <https://doi.org/10.1007/s10514-016-9607-8>