

Recognizing Attack Patterns: Clustering of Optical Flow Vectors in RoboCup Soccer

Auke J. Wiggers
6036163

Bachelor thesis
Credits: 18 ECTS

Bachelor Opleiding Kunstmatige Intelligentie

University of Amsterdam
Faculty of Science
Science Park 904
1098 XH Amsterdam

Supervisor
Dr. A. Visser

Faculty of Science
University of Amsterdam
Science Park 904
1098 XH Amsterdam

June 26th, 2012

Abstract

In contrast to human soccer players, autonomous robot soccer players often move according to a limited set of predefined behavioural rules. This knowledge can be used advantageously: If the opponent's behavioural rules are learned, it will be possible to detect these during a match and react accordingly. A method for autonomous activity mining in videos, called Probabilistic Latent Sequential Motifs, is used to find optical flow patterns in video's of a robot soccer player during a penalty shootout. The found patterns are used by a humanoid goalkeeper to predict and anticipate opponent behaviour. Effectiveness of the method is tested by comparing the performance of this goalkeeper, i.e., the ratio of number of goals to number of goals prevented, to that of an existing goalkeeper that only reacts when the ball approaches at sufficient speed. Results show that the found goalkeeper performs fairly well, but that it loses to the existing goalkeeper in terms of performance. Methods that may improve performance are discussed.

Contents

1	Introduction	2
2	Related work	2
3	Approach overview	4
3.1	Probabilistic Latent Sequential Motifs	4
3.2	From images to documents	5
3.3	Prediction task: anticipation	7
4	Experiments and results	8
5	Discussion and future work	13
5.1	Location quantization and scaling	13
5.2	Mirroring	14
5.3	Motion	14
5.4	Bias due to difference in moves	14
5.5	Future work	15
6	Conclusion	16
7	Acknowledgements	16

1 Introduction

RoboCup is an international research and education initiative, attempting to foster Artificial Intelligence and robotics research by providing a standard problem where a wide range of technologies can be integrated and examined. One of these problems is teaching a robot to play a soccer match. During these matches, the goalkeeping robot will often only react if a ball approaches at sufficient speed. Other features in the scene, such as presence of opponent players or proximity of the ball, are often not taken into account. The reason for this is the following: If a logical representation is used to describe the preconditions for a reaction, a new set of logical rules is necessary when a new opponent strategy is encountered.

It would be more efficient if instead, the goalkeeper finds the preconditions for these logical rules autonomously and in an unsupervised manner. A method for unsupervised activity mining in videos called Probabilistic Latent Sequential Motifs (PLSM) is introduced by Varadarajan et al [13], which makes use of topic models to find temporal activity patterns in image data. These patterns represent an activity in a video. Their relation to video data can be compared to the relation of combinations of syllabi to a word. Using this method, a model describing recurring patterns in a dataset can be found by learning offline, i.e., before a match. The detection of these patterns during a match will enable the goalkeeper to anticipate the opponent's behaviour. The effectiveness of the method is tested by comparing performance in a penalty shootout between a regular rule-based goalkeeper, that only reacts when a ball approaches, and the goalkeeper that has found rules using PLSM.

2 Related work

Activity mining is a field of research that is often associated with surveillance scenarios [14], e.g., to detect violent behaviour in a crowd or to analyze busy traffic scenes. In activity mining, the use of topic models [9] is an approach that has proven to be quite successful [4][7]. This model is a statistical type that enables discovery of so-called 'topics', which are abstract occurrences in a document. Often, topic model-based approaches first convert a video to documents containing bags-of-words, where each word is a representation for quantized pixel motion at certain locations in the image [4][7][13][15]. Note that the 'words', the visual features, are only a tag corresponding to

the correct optical flow vector, and do not contain the vector itself. One document is thus the result of analysis of optical flow data in one video. An activity pattern in a video can be represented as a set of relative movement vectors along with their respective starting positions and starting times. These activity patterns are the topics that are to be learned. To remain faithful to the terminology used by Varadarajan et al, these topics will be referred to as ‘motifs’ in this thesis.

When a regular bag-of-words approach is used, the temporal information is lost in the process [15]. To take this information into account as well, methods have been proposed where both the motifs and their starting times are jointly learned by complementing the visual features with their respective timestamps [13] [4]. A method introduced by Emonet et al [4] both the total number of motifs and the motifs themselves are learned in an unsupervised manner using the Hierarchical Dirichlet Process (HDP) [12], which allows for an infinite amount of motifs on multiple levels. In the problem addressed in this thesis, two levels can be used: On the first, the total number of motifs that is shared by all documents; on the second, the individual motifs and their starting times in each document. This in contrast to PLSM, which requires some a priori knowledge (e.g., the number of motifs) to be set beforehand. Although the HDP-based approach may be more complete than PLSM, it would seem that the latter is sufficient to solve the problem addressed in this thesis: The number of strategies for robot soccer may be theoretically infinite, in practice, it is always limited to a certain number of re-occurring actions. For example: Shooting at the goal can be described as a motion in an infinite number of directions and can occur on every position on the field. However, it only occurs when at least one robot is moving and there are only as many possible shooting directions as there are quantization categories. If the number of possible activities is learned in an automatic manner, it is likely that this number will be too high to be useful. Instead, we will use a small number of motifs and let PLSM determine what they are.

PLSM has been used in pedestrian and traffic analysis [13], in scenarios in which the camera is stationary, to find recurring activities in video data. In this thesis, PLSM will be applied in a new area. Instead of a camera placed above the scene, all video data will come from the goalkeeper’s point of view.

3 Approach overview

Several steps are performed to solve the addressed problem. The first step is the creation of temporal documents, derived from video data. From these documents, a model that describes the motifs can be derived using PLSM [13]. Next, a reaction that enables the goalkeeper to anticipate the opponent's activity has to be created. The final step is the prediction of an opponent robot's activity pattern, based on detection of the first frames of a learned motif.

3.1 Probabilistic Latent Sequential Motifs

To find recurring activity patterns, a temporal document d , that represents the image data as a bag-of-words, has to be derived. This document is of size $V \times T$, where V is the vocabulary size and T is the number of timesteps the document covers. At each timestep, the document is filled based on the presence of visual words w (See figure 1). PLSM is applied on this document to find motifs z , and their starting times t_s .

The main assumption of the PLSM model is that given a motif z , the occurrence of words within the document is independent of the time of occurrence. Note that there is a deterministic relation between the time variables $t_a = t_s + t_r$, enabling the use of variable t_r to denote the relative time since the start of a motif: In this model, occurrence of a word depends only on the motif and the time it occurs in the topic, not on the absolute time of occurrence t_a . The joint distribution of the model is given by:

$$p(w, t_a, d, z, t_s) = p(d) p(z|d) p(t_s|z, d) p(w|z) p(t_a - t_s|w, z)$$

The Expectation-Maximization algorithm can be used to estimate the set of model parameters Θ , by maximizing the log-likelihood of the model for the observed data. It is an iterative algorithm, that is initialized using random values as parameters, and is stopped when the increase in log-likelihood is too small. The log-likelihood is given by:

$$E[L] = \sum_{d=0}^D \sum_{w=0}^{N_w} \sum_{t_a=0}^{T_d} \sum_{z=0}^{N_w} \sum_{t_s=0}^{T_{ds}} n(w, t_a, d) p(z, t_s|w, t_a, d) \log p(w, t_a, d, z, t_s)$$

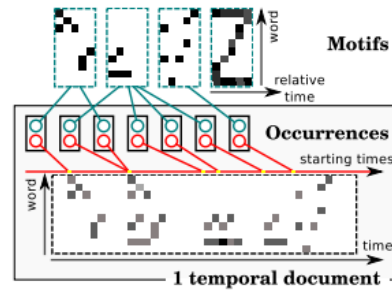


Figure 1: Graphical representation of a temporal document. Image courtesy of Emonet et al [4].

The posterior distribution of variables t_s and z is calculated in the E-step of the algorithm, given w , the absolute time t_a and the document d :

$$p(z, t_s | w, t_a, d) = \frac{p(w, t_a, d, z, t_s)}{p(w, t_a, d)} \text{ with } p(w, t_a, d) = \sum_{z=1}^{N_w} \sum_{t_s=1}^{T_{ds}} p(w, t_a, d, z, t_s)$$

Next, in the M-step, Θ is updated accordingly:

$$p(z|d) \propto \sum_{t_s=1}^{T_{ds}} \sum_{t_r=0}^{T_z-1} \sum_{w=1}^{N_w} n(w, t_s + t_r, d) p(z, t_s | w, t_s + t_r, d) \quad (1)$$

$$p(t_s|z, d) \propto \sum_{w=1}^{N_w} \sum_{t_r=0}^{T_z-1} n(w, t_s + t_r, d) p(z, t_s | w, t_s + t_r, d) \quad (2)$$

$$p_w(w|z) \propto \sum_{d=1}^D \sum_{t_s=1}^{T_s} \sum_{t_r=0}^{T_z-1} n(w, t_s + t_r, d) p(z, t_s | w, t_s + t_r, d) \quad (3)$$

$$p_{t_r}(t_r|w, z) \propto \sum_{d=1}^D \sum_{t_s=0}^{T_{ds}} n(w, t_s + t_r, d) p(z, t_s | w, t_s + t_r, d) \quad (4)$$

The motifs and their starting times can be found in the optimized distributions $p(z|d)$ and $p(t_s|z, d)$.

3.2 From images to documents

The content of temporal documents, that will serve as input for PLSM, depends on how the visual words w that form the vocabulary are defined. These words will be based on optical flow in the image, so first, optical flow is computed in areas of interest, namely the areas containing the ball and opponent robot (See figure 2). If we set the robot so its camera is stationary, the majority of the noise due to camera motion will be eliminated. The resulting optical flow vectors are quantized into four general directions (up, down, left and right) or marked as static if the norm is not sufficiently large. Based on their respective starting locations in the image, they are also quantized into non-overlapping cells of 10×10 pixels. It is possible to use these categorized vectors and their timestamp as low-level features in the temporal document. However, the resulting documents are quite large, as the vocabulary would consist of $64 \times 48 \times 5 = 15360$ words.




Figure 2: Left: Found features in frame f indicated by thin yellow circles. Only found in areas of interest, which are determined by location of the ball (thick yellow circle) and waistband (thick purple circle). Right: Optical flow vectors between frame f and frame $f + 1$, indicated by red arrows.

Instead, Probabilistic Latent Semantic Analysis (PLSA) [5], a dimensionality reduction method that makes use of latent classes, will be used to reduce the size of this vocabulary. This method was used by Varadarajan et al as well, in order to reduce computation time [13]. Its generative model, describing how each variable in the distribution can be sampled, is given by figure 3. PLSA models the probability of co-occurrence of words and documents as a mixture $P(\omega, d)$, creating the assumption that the occurrence of a word ω is independent of the video document d it belongs to, given a latent class c :

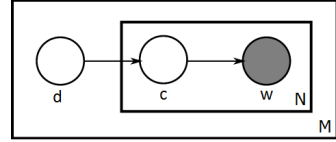


Figure 3: The generative model of PLSA.

$$P(\omega, d) = \sum_c P(c)P(d|c)P(\omega|c)$$

The documents d are defined as word count matrices of size $f \times V$ extracted from overlapping clips of f frames. The parameters of the model $P(\omega|z)$, $P(c)$ and $P(d|c)$ are estimated using the maximum likelihood principle. Given a set of training documents \mathcal{D} , the log-likelihood of the parameters Θ is given by:

$$\mathcal{L}(\Theta|\mathcal{D}) = \sum_{d \in \mathcal{D}} \sum_{\omega} n(d, \omega) \log(P(\omega|d))$$

Similar to the PLSM model, the parameters are optimized using the EM algorithm [6], which allows to learn distributions $P(\omega|c)$ for every latent class. The found distributions are then used to define words w for PLSM, meaning that $w = c$ and $N_w = N_c$. By specifying the desired number of latent classes N_c beforehand, we are able to reduce the size of the vocabulary to any given number. To keep the runtime low, $N_c = 25$ was chosen. The presence of each word w at an absolute time t_a in a document d is then defined as the normalized correlation:

$$n(d, t_a, w) = \frac{1}{\sum_{\omega \in W_c} n(d, \omega)} \sum_{\omega} n(d, \omega) p(\omega|c)$$

In this equation, W_c indicates the set of all words in the distribution $P(\omega|c)$ with a non-zero probability. This ensures us that the found presence is independent of activities elsewhere in the scene.

A resulting temporal document, that is to be used by PLSM, is of size $N_w \times T$, where N_w is the vocabulary size and T the range of absolute times (see figure 1). The documents are used as input for the PLSM method. As stated before, each word w in the vocabulary is a representation of a temporal pattern from the distribution $P(\omega|c)$, as found by PLSA.

3.3 Prediction task: anticipation

The final step is detection of activities, and reacting if one has been detected. The optical flow is calculated in areas containing the ball and the waistband of the opponent robot at one frame per second. The resulting vectors are quantized and converted to the bag-of-words representation. The found PLSA model is used to calculate presence for every latent class c , resulting in a document of exactly one timestep. The probability of d_{new} being a part of a motif is calculated as follows, with $0 \leq \alpha \leq 1$:

$$p(z, t_n | d_{\text{new}}) = (1 - \alpha) * p(z, t_{n-1} | d_{\text{new}}) + \alpha * \sum_w^{N_w} n(w, d_{\text{new}}) p(z|w, d_{\text{new}})$$

In this equation, the probability that a new document of one timestep is part of a motif z depends on previous probabilities $p(z, t_{1...n-1} | d_{\text{new}})$. If, for several timesteps, this probability is higher than a predefined threshold, we assume that the rest of the activity will be similar to the rest of the motif, and the goalkeeper will react accordingly. Note that the goalkeeper's reactions are intended to anticipate the end of their respective corresponding motifs: How the reaction is executed is independent of which of the timesteps in the motif is perceived.

Robot soccer behaviour is often limited to several different strategies. Therefore, it is highly likely that mirroring will enable us to reduce the number of motifs even further: The reaction to an activity perceived on the right side of the goalkeeper can be mirrored, this will result in a reaction that can be used if the same activity, but mirrored, is perceived on the left side of the goalkeeper. However, in practice, mirroring the patterns results in more motifs, which increases computation time (also see 5.2). This renders it useless in real-time applications for the Nao, which only has a single ATOM processor. It will be possible to apply this method for humanoids with greater computing power.


4 Experiments and results

The used robot is the Aldebaran Nao 4.0 [11], which is a humanoid robot with an Intel ATOM 1,6ghz CPU. It has two cameras, capable of taking snapshots of 640×480 pixels. The dataset used for training consists of 2370 images, taken at one frame per second, from one of these cameras. The set of images is split into 43 groups, each of which is between 20 and 128 frames long and describes a single recording of a penalty shootout by one opponent robot, as seen from the goalkeeper’s point of view. PLSA was applied on the optical flow found in these images as described in subsection 3.2, with $N_c = 25$, resulting in 43 temporal documents. PLSM was applied on these documents to find recurring motifs. For comparison, activity patterns were found in the data for various numbers of motifs N_z and its maximum duration T_z , as activities in the image sequence may vary in length (see figure 4).

A fitting reaction was created manually for each motif after analyzing the pattern of optical flow that it describes. The reaction consists of at least one action: Either walking, diving or a combination of the two. The walking direction and distance and the diving direction (left or right) were adjusted as well. As the reactions themselves are not found automatically and are not optimized, there is no guarantee that these are the best anticipating reactions to a specific motif.

The optical flow vectors that correspond to a motif were found by taking distribution $p(\omega|z)$ and, for each motif, calculating the presence of a word in the scene at a timestep t , given the motif, as:

$$n(\omega, t, z) = p(\omega|z) n(\omega, t)$$



(a) Found temporal motifs for number of motifs $N_z = 5$ and maximum number of timesteps $T_z = 10$.



(b) Found temporal motifs for number of motifs $N_z = 10$ and maximum number of timesteps $T_z = 20$.

Figure 4: Graphical representations of two of the found sets of motifs, for different numbers of motifs N_z and maximum duration T_z . Motifs are of size $T_z \times N_c$, with the number of PLSA patterns $N_c = 25$.

During the anticipation task, $p(z)$ is calculated for real scenes once every second. If, for any motif, this probability exceeds a predefined threshold for five successive seconds, the goalkeeper assumes this motif is perceived and executes the corresponding reaction regardless of what happens after the start of the reaction. The effectiveness of these reactions was compared to performance of a ‘regular’ goalkeeper, that only reacts if the ball approaches at sufficient speed by diving, through a total of 15 penalty shootouts. These shootouts were conducted by a Nao with a basic behaviour model: Find the ball, walk towards it, locate the goal, shoot. Figure 5 depicts the starting positions of both robots and the ball in a penalty shootout.



Figure 5: A typical start of a penalty shootout. The opposing player (left) starts at a distance of 3 meters to the goal, walks towards the ball, then kicks it.

We say that a goalkeeper interferes if the ball would have gone into the goal, were it not for the goalkeeper. A miss is when the ball does not reach the goal at all. A hit is, obviously, when the ball passes the keeper, in between the poles of the goal. If the keeper does not react but prevents the hit by sitting still, this is counted as a hit. Interestingly, the goalkeepers that were the result of PLSM did not react in case of a miss, whereas the regular goalkeeper did, although this may be due to chance. There is no time limit for a single shootout: It ends when the player scores, the ball is kicked out of the field (behind the goal line), or the goalkeeper has made a save. Details on specifications the field and goals can be found in [3]. The results can be seen in the following table:

N_z / T_z	Hit	Miss	Goalkeeper interferes
5 / 10	8	3	4
5 / 20	9	4	2
10 / 10	9	3	3
10 / 20	11	4	0

From these results, we can conclude that the most effective goalkeeper that was found using PLSM uses a small number of motifs. A likely cause for this is that PLSM is forced to map a large number of found patterns to a predefined number of motifs: The smaller the number of motifs, the more general they will be. Results also indicate that effectiveness is increased if T_z is low. This seems counterintuitive, as the longer motif contains more timesteps that could match the new document. However, the activities that are represented by motifs of $T_z = 10$ are forced to represent short and concise patterns that match typical recurring actions, whereas the motifs of $T_z = 20$ may represent an entire penalty shootout. An indication for this is that, when shown a real-life scene, $p(z)$ is large for almost every motif, for $T_z = 20$, whereas there is more contrast for $T_z = 10$ (see figure 7).

The goalkeepers that were found using PLSM were compared it to the regular goalkeeper model. This model bases its actions on the speed of the ball. If this speed exceeds a certain threshold, the direction is calculated and the keeper dives left or right. The direction is based on the expected location where the ball will cross the goal line. Results are shown below:

	Hit	Miss	Goalkeeper interferes
Regular goalkeeper	2	4	9

Clearly, results of the goalkeeper found by PLSM are still poor when compared to those of the regular goalkeeper. It is likely that this is due to the noise and redundant information that each motif contains, as figure 6 indicates. The noise may render a motif too general to be useful, as features that do not contribute to the action represented by the motif are included in it as well. Such a motif generally has a high $p(z)$, even if the activity it represents is not perceived. The opposite can also be true: Figure 7 shows us that the fourth motif is, compared to the other motifs, a motif with very low probability of occurring, given the document d_{new} that describes the scene. There are two possible causes: Either the motif is the result of a specific activity that does not occur in the scene, or the motif is the result of noisy optical flow and will therefore always have low probability. In the second case, this motif will not be of any use, which would influence goalkeeper performance.



Figure 6: Three consecutive timesteps of a noisy motif z for $N_z = 5$, $T_z = 10$. Circles indicate cells where optical flow is found, with the thickness of the circle corresponding to presence $p(w|z)$. Obviously, this motif is too general to be of use.

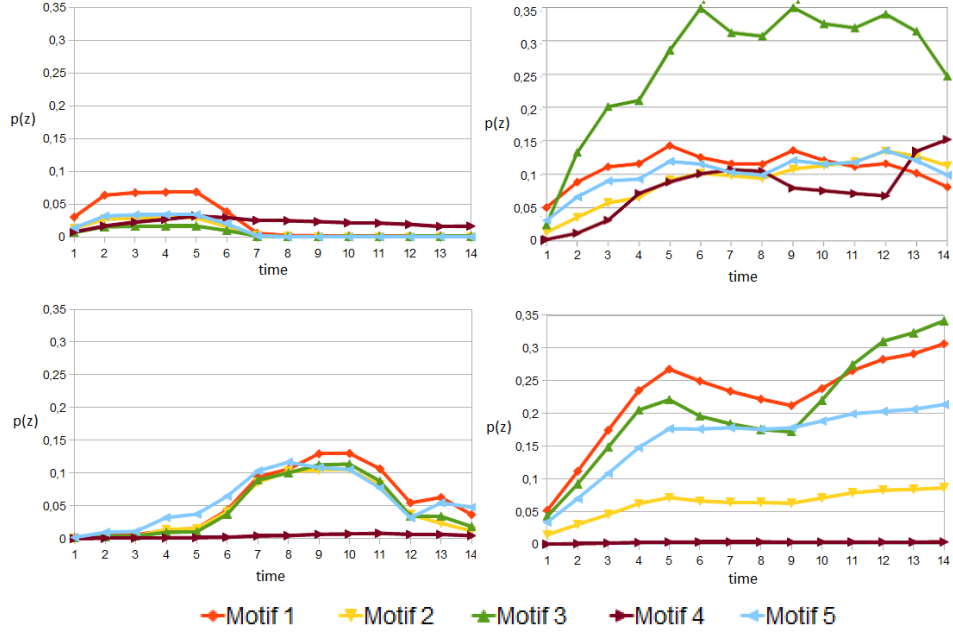


Figure 7: Probability $p(z)$ for $N_z = 5$, for two identical static scenes (left) and two similar non-static scenes (right), for 14 timesteps. Top: Results for $T_z = 10$. Bottom: Results for $T_z = 20$.

5 Discussion and future work

In this thesis, PLSM is applied in an environment in which an activity can occur at every position on the field, with different orientations. This in contrast to, for example, traffic analysis scenarios, in which cars are always on the road and can move in a limited number of directions. It seems that for this problem, the location cue can be discarded: Optical flow vectors on which the activity patterns are based can be translated and scaled to find the pattern that describes the activity when it is perceived at a different location. If this technique was applied, an action can be represented by optical flow alone. Whether this will improve results, and several other factors that are likely to influence results, are discussed in this section.

5.1 Location quantization and scaling

As mentioned in 3.2, the possibility exists that the same activity is perceived at different locations. As a result, a set of perceived patterns, in reality consisting of only one activity but seen at different locations, will not necessarily be mapped to the same motif. One reason for this phenomenon is difference in distance: An opponent robot that is close to the goalkeeper appears larger in the image than when it is far away, and will thus occupy a larger part of the grid. Possible solutions for this problem involve use of a cell grid that uses a logarithmic scale, or dividing the image into larger cells, both of which reduce the number of possible features. However, each activity pattern represents a possible activity of an opponent robot, and these activities affect the game differently if they occur on unusual positions on the field. Therefore, it is impossible to group patterns that are a result of optical flow found on an unusual position, with the already learned motifs, even if they describe the same action by the robot player.

For example: When a player shoots the ball to the left of the goalkeeper and is close, the goalkeeper may want to move left as well to defend the goal, while the same action, when perceived at a greater distance, may not pose a threat at all and not require a reaction. A second reason is that the number of features may be reduced, but it is hard to tell whether we can afford to lose these features: This method may cause PLSM to find motifs that are too general to be used.

5.2 Mirroring

If we assume that a perceived pattern is the result of an action in arbitrary reaction, it should be possible to reduce the number of motifs by mapping patterns that are mirrored versions of each other to a single motif. The disadvantage of this approach is that it increases computation time for the prediction, as each motif that is mirrored requires calculation of $p(z)$. Even if the mirrored versions are mapped to the same motif, it will be necessary to check which of the two versions is perceived before a fitting reaction can be given. Additionally, there may be exceptions to the rule, or the mirrored motif may be too similar to a different motif, causing the goalkeeper to react wrongly.

5.3 Motion

In this thesis, the difficulty of the problem is greatly reduced by removing camera motion. The calculation of optical flow is quite different for a non-stationary camera, even if the absolute camera motion is known. However, the Nao does not stand perfectly still due to the servos not being able to keep a pose exactly the same. As a result, there is always some noise present in calculation of optical flow in a static scene. This slight movement also makes it hard to use background subtraction [8], a method which may improve results due to the removal of objects that are not of interest. If the Nao is able to stand perfectly still during motif detection, it will still be hard to use background subtraction: Camera calibration is needed after the robot has moved (e.g., walking towards an opponent, diving to stop the ball) and this will cost time. Additionally, the Nao's camera moves inside the head after an impact or sudden movement, making any method that assumes that the camera is stationary very susceptible to noise. It is imperative that the goalkeeper reacts as fast as possible, and therefore, background subtraction was not applied.

5.4 Bias due to difference in moves

It can be argued that the experiment is biased: There is no guarantee that the set of possible moves of the goalkeeper found using PLSM is 'better' than that of the existing goalkeeper model. It is also impossible to say if these are put to better use in either of the models, as the reactions for both goalkeepers were created manually. However, the goal of the conducted research was to surpass the existing goalkeeper model: There is no uncertainty to whether the method is effective or not, although results indicate that better

behavioural rules for the goalkeepers found by PLSM will improve effectiveness.

5.5 Future work

The RoboCup consists of several leagues, each league having its own rules and fields of research. The league in which the Nao is used is the Standard Platform League, a soccer competition for humanoid robots. Although optimization of soccer behaviour is one field of research in the league, detection of opponent's behavioural patterns has not been attempted before. The main reason for this is that participating teams are constantly adjusting behaviour; it would be futile to analyze the opponent to create opponent-specific strategies manually.

Automatic activity mining has been used in the Small Size League [2], for example in [1], where Bayesian approaches are used to find patterns in opponent behaviour, and [10], in which a method is described that finds patterns in combinations of actions (e.g., passing, shooting, dribbling) rather than in general movement. Similar to the goal of the research in this thesis, the patterns are used to predict the opponent's behaviour, and effective strategies are created to exploit the opponent's weak points. An important difference between this league and the Standard Platform League is that data is retrieved from two stationary camera's that are placed above the playing surface, whereas in the humanoid leagues, the camera feed of the robots is the only available source of visual information. It is likely that results for PLSM will improve when, instead of point-of-view image data from the goalkeeper, an external camera that gives a top is used. However, as the Standard Platform League rules do not allow use of sensor data other than that of the robots [3], any method that uses external sensors cannot be applied in an official match.

The use of similar automated activity mining methods in official competitions may prove to be advantageous: Detection and prediction of an opponent's movement will make complex actions such as defending the goal by blocking, or passing the ball without it being intercepted, a lot more simple. It is likely that eventually, the detection of an opponent's behavioural patterns will become a necessary part of any robot soccer team, as adapting to the opponent's strategy is a basic element of soccer.

6 Conclusion

This thesis describes an implementation for automatic unsupervised activity mining in video's for a humanoid soccer robot, and its effectiveness when used for the prediction and anticipation of opponent actions. A topic model-based method called Probabilistic Latent Sequential Motifs [13] is used to find recurring patterns of optical flow, referred to as motifs, in a dataset of short image sequences. The number of motifs and the maximum length of such a motif is specified beforehand. The found patterns are then used by the goalkeeper to predict an opponent's action and react accordingly. The effectiveness of the used method was tested by comparing performance (i.e., the ratio of the number of prevented goals to the number of scored goals) of the resulting goalkeeper to that of a simpler goalkeeper model, that only reacts when the ball is approaching, through a penalty shootout. Results indicate that setting the desired number of motifs relatively low contributes to performance, but that a goalkeeper with a predefined set of behavioural rules performs better overall. Nevertheless, the thesis shows that automatic activity mining is a promising field of research in robot soccer.

7 Acknowledgements

The author would like to thank Jagannadan Varadarajan, Remi Emonet and Jean-Marc Odobez for the distribution of their Matlab source code for Probabilistic Latent Sequential Motifs.

References

- [1] D. Ball and G. Wyeth. Classifying an opponents behaviour in robot soccer. In *Proceedings of the 2003 Australasian Conference on Robotics and Automation*, 2003.
- [2] Universität Bremen. Small size league homepage. <http://small-size.informatik.uni-bremen.de/>, 2012.
- [3] RoboCup Technical Committee. Robocup standard platform league (nao) rule book. <http://www.tzi.de/spl/pub/Website/Downloads/Rules2012.pdf>, May 2012.
- [4] R. Emonet, J. Varadarajan, and J.M. Odobez. Extracting and locating temporal motifs in video scenes using a hierarchical non parametric

- bayesian model. In *IEEE Conference on Computer Vision and Pattern Recognition*, June 2011.
- [5] T. Hofmann. Probabilistic latent semantic indexing. In *In Proceedings of the 22th International Conference on Research and Development in Information Retrieval (SIGIR)*, 1999.
 - [6] T. Hofmann. Unsupervised learning by probability latent semantic analysis. *Machine Learning*, 42:177–196, 2001.
 - [7] D. Kuettel, M. D. Breitenstein, L. van Gool, and V. Ferrari. What’s going on? discovering spatio-temporal dependencies in dynamic scenes. In *IEEE Conference on Computer Vision and Pattern Recognition*, June 2010.
 - [8] A. McIvor. Background subtraction techniques. In *Image and Vision Computing New Zealand*, 2000.
 - [9] C. Papadimitriou, P. Raghavan, H. Tamaki, and S. Vempala. Latent semantic indexing: A probabilistic analysis. In *Proceedings of ACM PODS*, 1998.
 - [10] Morteza Mousakhani Ramin Fathzadeh, Vahid Mokhtari and Fariborz Mahmoudi. Mining opponent behavior: A champion of robocup coach competition. *Robotics Symposium, 2006. LARS 06. IEEE 3rd Latin American*, pages 80–83, 2006.
 - [11] Aldebaran Robotics. Aldebaran nao hardware specifications. <http://www.aldebaran-robotics.com/en/Discover-NAO/Key-Features/hardware-platform.html>, 2012.
 - [12] Yee Whye Teh, M.I. Jordan, M.J. Beal, and D.M. Blei. Hierarchical dirichlet processes. *Journal of the American Statistical Association*, 101(476), 2006.
 - [13] J. Varadarajan, R. Emonet, and J.M. Odobez. Probabilistic latent sequential motifs: Discovering temporal activity patterns in video scenes. In *Proceedings of the British Machine Vision Conference*, 2010.
 - [14] J. Varadarajan and J.M. Odobez. Topic models for scene analysis and abnormality detection. In *ICCV 12th International Workshop on Visual Surveillance*, 2009.

- [15] Xiaogang Wang, Xiaoxu Ma, and W. Eric L. Grimson. Unsupervised activity perception in crowded and complicated scenes using hierarchical bayesian models. *IEEE Transactions on pattern analysis and machine intelligence*, 31(3), 2000.