In *Proc. 2002 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*
pages 2770-2775, EPFL, Switzerland, September 30 - October 4, 2002

General Spatial Features for Analysis of Multi-robot and Human Activities from Raw Position Data

Helen Yan, Maja J. Matarić

University of Southern California, Los Angeles, USA, {helenyan, mataric}@usc.edu

Abstract

In this paper, we present a method of quantifying ideas from proxemics into a set of spatial features that can be used for recognizing, classifying, and modeling multi-robot and human activities. We validate the features first with position data from multi-robot activities such as flocking, random walk, and boundary following, and then with human activity data. These features are simple and general, and can also be used as performance measures of robot controllers and in learning robot behaviors.

1 Introduction

Human spatial behavior is a well-studied area in social psychology [1]. One influential discipline of these studies is *proxemics*, the study of the social use of physical space. According to Hall [7], proxemics is a means of non-verbal communication: people use the space around them differently depending on their culture, their activities, and the structure of their surrounding environment. Important ideas in proxemics include the notion of personal space and the concept of distance zones [8]. Personal space is an invisible boundary that an individual maintains during social interactions; if this space is intruded, the individual reacts by moving backward, reestablishing the territory. Different types of social interactions establish stable distance zones at various ranges.

In this paper, we use proxemics ideas to derive a set of features from raw position data of individuals. Our goal is to quantify the spatial interactions in our features so that we can eventually develop activity models that closely resemble human behavior. We also show that feature extraction techniques, routinely used in computer vision [13], are effective here as well. They can be used in evaluating performance of robot controllers as well as in understanding and learning patterns of activity.

We begin with a description of the three types of spatial features we used (pair-based, cluster-based, and environment-based) and their social interpretations. Next, we validate these features with multi-robot and multi-person activity data and present an analysis of the results. Finally, we conclude with an application of the derived features on the classification of multi-robot flocking, random walk, and boundary tracing.

2 Approach

A first step of modeling multi-robot and human activity is to find a set of features that can characterize behavioral patterns. Our approach to deriving such features is to apply ideas from proxemics [7] - the study of space. We believe that position data embed significant spatial interaction dynamics which can be accounted for through the use of proxemic (or spatial) features. We categorize spatial features into three main classes based on the information they encode: pair-based, cluster-based, and environment-based.

2.1 Pair-Based Features

Pair-based features encode information about pairs of individuals. One such feature is *inter-individual distance*, the distance between a pair of data points using a distance metric. In the work presented here, we computed this distance for all pairs of data points, where points were positions of individuals in each captured frame of data. We sorted the results and clustered them using simple Euclidean distance as the distance metric.

2.2 Cluster-Based Features

A cluster is a subset of individuals that are "close". To cluster, we assumed that position data points represent centers of individuals. If the distance between any pair of these falls below a predefined cluster threshold $Th_{cluster}$, the pair is labeled as *equivalent*. $Th_{cluster}$ defines formally the meaning of "closeness"; all equivalent points are then labeled as belonging to the same cluster. This results in clusters of arbitrary shapes, so no *a priori* bias was provided toward particular cluster types. Once clusters are formed, we can extract cluster-based features, described next, which encode information about subsets of individuals and their interactions.

Total Number of Clusters. This is the number of subsets of individuals that are identified as being "close" together in a given frame. By studying this feature under different parameters, we can discover grouping tendencies in a particular data set.

Size of Cluster. This is the number of data points constituting a cluster. This feature provides information about the grouping tendencies of the data set and can be used as a bias to refine or automate clustering.

Cluster Area. We chose to define cluster area as the area of the convex polygon formed by the cluster points, computed using Graham's algorithm [12]. The area of the polygon is computed using the equation below, providing information about group compactness tendencies in the data.

$$A = \frac{\left[\sum_{i=0}^{n-2} (x_i y_{i+1} - y_i x_{i+1})\right] + (x_{n-1} y_0 - y_{n-1} x_0)}{2} \tag{1}$$

2.3 Environment-Based Features

These features encode information about the interaction between individuals and the environment. The most obvious feature is the use of space, i.e., the distribution of individuals in the environment, best described with occupancy grids [11], which effectively reflect instantaneous distribution of individuals within an area.

3 Experimental Design

3.1 Data Sources

Our activity data consist of a collection of Cartesian coordinates, *(x, y)*, for each person or robot. To gather robot activity data, we used Player and Stage [1] [6]. Player is a robot device server for sensor and motor control, and Stage is a realistic simulator for robot devices. The system allows robot positions to be logged throughout the experiments. To gather human activity data, we used a planar laser-based people tracking system [4] that combines range data from multiple laser range finders to effectively track and log positions of people over time.

3.2 Robot Activity

To validate the proposed features with robot activity data, we collected robot position data for three activities with different group sizes and varying environment scales, as shown in Table 1. We gathered 3 sets of robot position data in 2-hour durations for each of these experiment sets.

Table 1: *Robot activity experiments*

Controllers	random walk, flocking, boundary following
Group Sizes	4, 6, 8
Environments	$6{\times}7m^2$, $12{\times}14m^2$

The robot controllers used for these three activities were as follows. For *random walk*, the robot moved forward and turned by a random degree at a randomly chosen interval while avoiding collisions with obstacles and other robots. *Flocking* was accomplished with the simplifying assumption that the robots had access to global positions of all other robots. A leader was assigned initially and all others moved toward it. A count, *stuckcount*, was kept for each robot so when the current leader was stuck for more than a certain duration, a new leader, one with the lowest *stuckcount*, was assigned. *Boundary following* was accomplished as a combination of collision avoidance and the tendency to move toward closest obstacles and align to them.

3.3 Human Activity

We collected a total of 5 hours of unscripted human activity data in our laboratory using 4 lasers and the algorithm described in [4]. The data provide planar blobs corresponding to human positions over time.

4 Experimental Validation with Robot Activity Data

In this section, we demonstrate the experimental validation of the described features with robot position data using different parameters (Table 1).

Inter-Robot Distances. Figure 1 shows the percentages of distances, normalized by the maximum distance of the environment, that fall into the ten equally spaced bins for the different experiment sets. Several observations can be drawn. Flocking robots tend to stay very close to each other relative to the size of the room. Random-walking and boundary-following robots tend to be farther apart from one another. Regardless of the size of the room and the robot group size, the pattern associated with the specific activity stays more or less constant.

Total Number of Clusters. The percentage of number of clusters formed for the experiment sets at cluster threshold 1.1m is given in Figure 2. We can

[1] Player and Stage were developed jointly at the USC Robotics Research Labs and HRL Labs and are freely available under the GNU General Public License from `http://playerstage.sourceforge.net`.

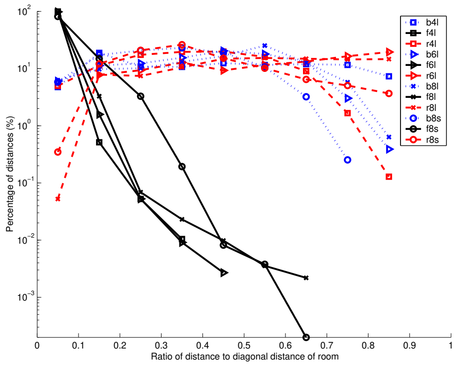

Figure 1: *Inter-robot distances for different experiment sets (l=12×14m^{2}room, s=6×7m^{2}room; 4,6,8=number of robots; b=boundary following, f=flocking, r=randomwalk)*

see that flocking largely forms one cluster whereas random walk largely forms as many clusters as there are robots, as would be expected if the robots' controllers are working properly and the environmental conditions did not get in the way. These trends remain more or less constant even under changes in group sizes and environment scales. Boundary following tends to form many clusters like random walk, but the behavior varies as the environment or group size changes. This is because when the environment scale decreases proportionally, the boundary of the environment also decreases, causing more close contacts between the robots. Similar reasoning can be applied when robot group size changes. Figure 3 shows how the average number of clusters varies using different thresholds. Except for boundary following in the 6×7m^2room (overcrowding), the curves stay fairly flat. This suggests that the results do not vary significantly with slight changes in thresholds.

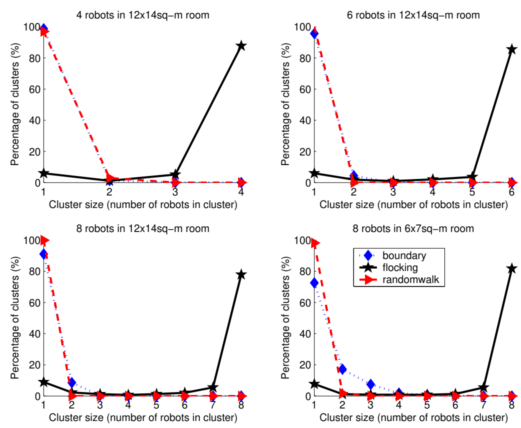

Figure 2: *Number of clusters formed under different experiment parameters with 1.1m as cluster threshold*

Size of Cluster Formed. The percentage of clusters formed at different cluster sizes at threshold 1.1m is given in Figure 4. It shows that flocking tends to give clusters much the same as the group size or with size 1. This can be expected from the rather idealized flocking algorithm we used, which could keep all robots very closely together. Combined with the number of clusters feature, we see that flocking ideally tends to form mostly one cluster of size that equals the group size. In contrast, random walk and boundary following tend to form very small clusters, mostly of sizes 1 and 2.

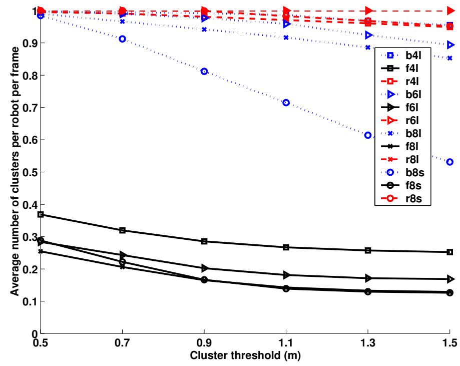

Figure 3: *Average number of clusters per frame normalized by the number of robots for different cluster thresholds*

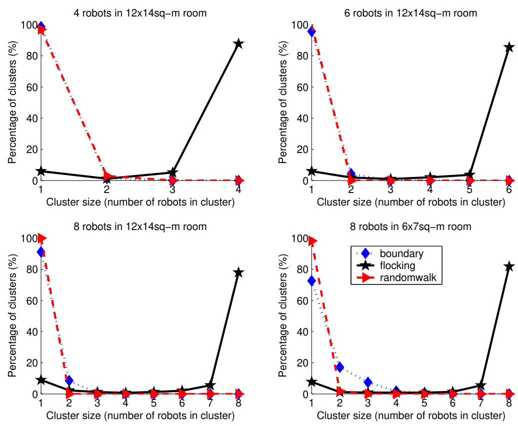

Figure 4: *Cluster sizes formed under different experiment parameters with 1.1m as cluster threshold*

Occupancy Grids. We divided the 6×7m^2 and 12×14m^2 rooms into 3×4 and 6×8 grids, respectively. Figure 5 shows the resulting occupancy. We see that, in general, the robots tend to stay along the edges of the room when boundary following (as the activity suggests) whereas random walking robots tend to distribute more or less evenly throughout the space, as expected. This even distribution minimizes the frequencies of collisions with other robots. In contrast, flocking robots tend to cluster in some parts of the room more than others.

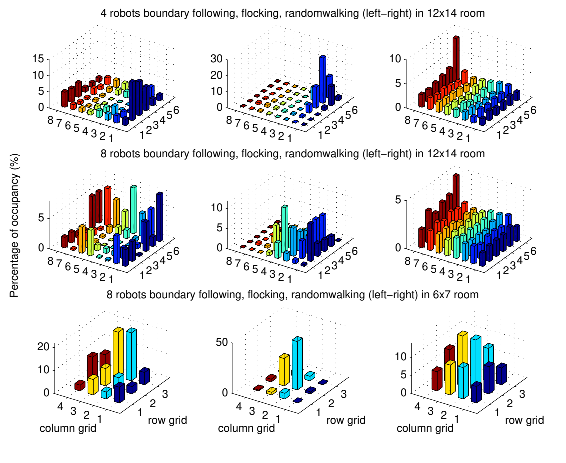

Figure 5: *Occupancy grids under different experimental parameters*

As shown above, the three types of features yield distinct distributions that are roughly invariant to environment scales and robot group sizes for flocking, random-walk, and boundary following. Therefore, we conclude that these features are useful in discriminating robot activities, as we demonstrate empirically in Section 6. In addition, one can apply these features in evaluating the performance of a robot controller or comparing performances between different controllers for the same behavior, as used, for example, by [5] for evaluating robot formations.

5 Experimental Validation with Human Activity Data

From the validation with robot activity data, we observe that the features produce spatial characteristics that match our expectation of these robot controllers. This implies that the features successfully capture the spatial interaction dynamics of these activities. In this section, we validate our features with more complex spatial behavior - unscripted human activity. This validation is important because people, unlike robots, do not maintain specific spatial patterns. ¿From position data, we extract the spatial features in a similar manner as with robot activity data (Figures 6 and 7).

Inter-personal Distances. Figure 6c shows the distribution of inter-personal distances. We observe that people tend to stay close to each other on average. In addition, most people do not stay more than 6m away. This means that people tend not to completely isolate themselves from social interactions in an office environment.

Total number of clusters. We investigated the number of clusters and cluster sizes formed for thresholds 0.5m to 3m. Thresholds 0.5m to 1.0m tend to form many clusters, i.e., there are more 1-person clusters, meaning that people tend to keep at least 0.5m to 1.0m away from each other. This agrees with the notion of personal space [7] - a minimum distance maintained by an individual during social interactions. Thresholds 1.5m to 2.5m yield many fewer clusters per frame, indicating that this zone is what most people are comfortable with when having personal interactions. The data agree with the theory of distance zones for different types of interactions (intimate zone, personal zone, social zone and public zone) suggested in [8].

Cluster Size. Figure 6b shows that people tend to be alone or form only small groups. A few formations of moderate sized groups (5 to 7 members) are observed because of the occurrence of a meeting during a portion of the experiment. These observations suggest that small group interaction is dominant in an office environment.

Cluster Area. Even in the cases where clusters are formed, the clusters mostly have small areas compared to the size of the environment independent of the cluster threshold used (Figure 6d). This demonstrates that, for this domain, small-sized clustering is a more stable form for group interactions than large-sized clustering.

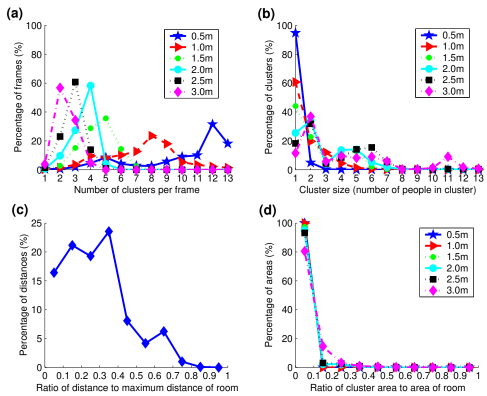

Figure 6: *(a) Number of clusters, (b) Cluster sizes, (c) Distributions of interpersonal distances and (d) cluster areas for human activity data (maximum distance = 16.3m, background area = 103m^2)*

Occupancy Grids. Figure 7 shows the occupancy grids computed using the human activity data. As depicted in the figure, people tend to cluster along the edges of the environment (brighter regions). Medium occupancy is observed in the smaller room where the meeting took place. The occupancy in that room is more or less evenly distributed, indicating that people spread out throughout this region during the meeting. This agrees with the idea suggested by Hall [7] that people tend to distribute themselves around the walls when given a large enough room.

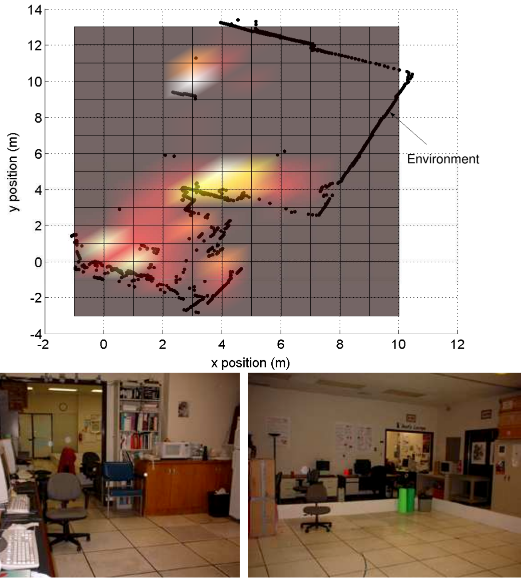

Figure 7: (Top) Occupancy grids computed using unscripted human activity data superimposed on background data points, (Bottom) Test environment

6 Application to Multi-robot Activity Recognition

In this section, we validate the derived spatial features in a simplified multi-robot activity recognition task. In particular, we focus on investigating the effectiveness of features in discriminating flocking, random walk, and boundary tracing (Test1), the effect of varying cluster thresholds (Test2), and experimental parameters (Test3) on classification. We gathered feature vectors from the position data collected from our validation experiments above to construct our training and test sets; the position data used in building the training set were obtained from separate experiments than those used in the test sets (see Table 2 for the feature extraction and experimental parameters). Each feature vector consisted of 67 features extracted from a frame of position data: average number of clusters per robot (1 feature), percentages of different cluster sizes (8 features), percentages of normalized inter-individual distances in ten equally spaced bins (10 features), and percentages of counts in the occupancy grids (6x8 features). Using the 5-nearest neighbor classifier with Euclidean distance as the distance metric for the feature vectors, we obtained the classification results for flocking, random walk, and boundary tracing shown in Table 3.

High recognition rates from Test1 demonstrate that the derived spatial features discriminate the activities very well. Of all three activities, the features work best in classifying random walk and flocking. The overall recognition rates do not vary significantly even when we use a different number of robots or different cluster thresholds in extracting the feature vectors in Test2 and Test3.

Table 2: *Parameters for training and testing sets*

	$Th_{cluster}$	**# Robots**	**Environment**
Training	1.1m	8	$12 \times 14m^2$
Test1	1.1m	8	$12 \times 14m^2$
Test2	1.5m	8	$12 \times 14m^2$
Test3	1.1m	6	$12 \times 14m^2$

Table 3: *Recognition rates for multi-robot activity recognition using 5-nearest neighbor classifier ($num_{training}$≈14200, $num_{testing}$≈7000)*

	Test1	**Test2**	**Test3**
Random walk	99.0%	98.2%	96.1%
Flocking	100.0%	100.0%	99.9%
Boundary Tracing	89.0%	93.7%	77.8%
Overall	96.0%	97.3%	91.3%

7 Related Work

The work presented in this paper is inspired by some of the feature extraction techniques used in computer vision and pattern recognition [13]. Robotics researchers have used some similar techniques. In [9], Han used features like relative object location as observations in a hidden Markov model to represent and recognize behaviors of robots in the soccer domain. Balch [3] proposed a method for spatially analyzing foraging activities of real ants that is in principle similar to the occupancy grids used here. In an earlier paper [2], he introduced an information theoretic metric measuring group diversity based on robot heterogeneity.

8 Summary and Conclusions

In this paper, we have discussed some simple and general spatial features that can be used as a first step in multi-robot and human activity modeling. We have demonstrated the use of inter-individual distances, number of clusters, cluster sizes and occupancy grids in the analysis of robot behaviors under different parameters. As discussed earlier, flocking robots tend to be close to each other forming mostly one cluster, with the same size as the number of robots. They tend to cluster in certain parts of the environment more than others. Random walking robots, on the other hand, tend to distribute themselves evenly throughout space, leading to clusters with mostly one robot. Lastly, boundary following robots cluster along the edge of the room forming clusters with mostly one or two robots. The spatial characteristics exploited by these behaviors are mostly independent of robot group sizes and environment scales. These spatial patterns match what we expect from these robot controllers, demonstrating that the features are indeed relevant and successful in capturing the social interactions dynamics of these behaviors. Hence, one can apply these features to analyze more complex activities that do not exhibit obvious spatial characteristics, such as human activity, which lead to our second validation experiment.

Analysis of the spatial features with unscripted human activity data collected in an office environment analytically supports proxemics theories [7] [8] that people tend to maintain personal distances while not being completely isolated from others. They tend to cluster in small groups that form small areas compared to the environment. They also tend to gather along the edges of the environment, except when there is a meeting. Because of their strong correlation with proxemics theory, we are confident that the spatial statistics shown in this paper are not a coincidence or likely to be environment-specific. Our next goal is to apply these spatial features to the development of activity models that more closely resemble human behavior.

Finally, we demonstrated a successful application of the derived features for the classification of multi-robot flocking, random walk, and boundary tracing, which yielded a recognition rate of 96.0%. We also showed that the features work well even when the number of robots or the clustering threshold is different in the training and test sets. Thus, the proposed features can serve as performance metrics in evaluating how well a robot controller accomplishes its tasks [5] and can be applied in developing activity recognition systems, such as video surveillance systems [10] and narrating agents [9].

Acknowledgments

This work is supported in part by the DARS MARS Program grant DABT63-99-1-0015 and ONR MURI grant N00014-01-1-0890. The authors wish to thank Tørbjorn S. Dahl, Jakob Fredslund and Esben H. Østergaard for their help with the robot controllers.

References

[1] J. Aiello, "Human Spatial Behavior", In *Handbook of Environmental Psychology*, Volume 1, D. Stokols, and I. Altman, eds., John Wiley & Sons, New York, 1987, pp. 389-504.

[2] T. Balch, "Hierarchic social entropy: An information theoretic measure of robot group diversity", *Autonomous Robots*, 8(3), Kluwer Academic Publishing, June 2000, pp. 209-237.

[3] T. Balch, Z. Khan, and M. Veloso, "Automatically tracking and analyzing the behavior of live insect colonies", In *Proceedings of the Fifth International Conference on Autonomous Agents (Agents'01)*, ACM Press, Montreal, Canada, June 2001. pp. 521-528.

[4] A. Fod, A. Howard, and M.J. Matarić, "Laser-Based People Tracking", To appear in *IEEE International Conference on Robotics and Automation (ICRA'02)*, Washington D.C., May 11-15, 2002.

[5] J. Fredslund, M.J. Matarić, "Robot Formations Using Only Local Sensing and Control", *IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA'2001)*, Banff, Alberta, Canada, July 29 - August 1, 2001, pp. 308-313.

[6] B.P. Gerkey, R.T. Vaughan, K. Støy, A. Howard, G.S. Sukhatme, and M.J. Matarić, "Most Valuable Player: A Robot Device Server for Distributed Control". *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'01)*, Wailea, Hawaii, October 29 - November 3, 2001, pp. 1226-1231.

[7] E.T. Hall, *The Silent Language*, Doubleday & Company, Inc., New York, 1959.

[8] E.T. Hall, *The Hidden Dimension*, Doubleday & Company, Inc., New York, 1966.

[9] K. Han, and M. Veloso, "Automated Robot Behavior Recognition Applied to Robotic Soccer", In *Proceedings of the Ninth International Symposium of Robotics Research (ISSR'99)*, Snowbird, Utah, October 1999, pp. 199-204.

[10] Y. Ivanov, C. Stauffer, A. Bobick, and W.E.L. Grimson, "Video Surveillance of Interactions". In *Second IEEE Workshop on Visual Surveillance (VS'99)*, Fort Collins, Colorado, June 1999, pp. 82-89.

[11] L. Matthies, and A. Elfes, "Integration of sonar and stereo range data using a grid-based representation", In *Proceedings of the 1988 IEEE International Conference on Robotics and Automation (ICRA'88)*, Philadelphia, Pennsylvania, April 1988, pp. 727-733.

[12] J. O'Rourke, *Computational Geometry in C, 2nd Edition*, Cambridge University Press, New York, 1998.

[13] M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis, and Machine Vision, 2nd Edition*, PWS Publishing, Pacific Grove, 1999.