# Group Affiliation Detection Using Model Divergence for Wearable Devices

**Dawud Gordon**
TwoSense Labs
Karlsruhe, Germany
dawud@twosense-labs.com

**Martin Wirz**
ETH Zürich
Zürich, Switzerland
wirz@ife.ee.ethz.ch

**Daniel Roggen**
University of Sussex
Sussex, United Kingdom
daniel.roggen@ieee.org

**Gerhard Tröster**
ETH Zürich
Zürich, Switzerland
troester@ife.ee.ethz.ch

**Michael Beigl**
KIT / TecO
Karlsruhe, Germany
michael.beigl@kit.edu

## ABSTRACT

Methods for recognizing group affiliations using mobile devices have been proposed using centralized instances to aggregate and evaluate data. However centralized systems do not scale well and fail when the network is congested. We present a method for distributed, peer-to-peer (P2P) recognition of group affiliations in multi-group environments, using the divergence of mobile phone sensor data distributions as an indicator of similarity. The method assesses pairwise similarity between individuals using model parameters instead of sensor observations, and then interprets that information in a distributed manner. An experiment was conducted with 10 individuals in different group configurations to compare P2P and conventional centralized approaches. Although the output of the proposed method fluctuates, we can still correctly detect 93% of group affiliations by applying a filter. We foresee applications in mobile social networking, life logging, smart environments, crowd situations and possibly crowd emergencies.

## Author Keywords

Group affiliation detection; computational social sciences; peer-to-peer; wearable computing; mobile computing;

## ACM Classification Keywords

I.2.11 Distributed Artificial Intelligence: Miscellaneous

## 1. INTRODUCTION AND MOTIVATION

Around 70% of the time we spend in public areas is done together with other people [16]. In general we are social creatures and spend a great deal of our time in groups of one form or another [9]. Groups are better than individuals at accomplishing tasks, which is often why they are formed in the first place [9]. Understanding group behavior and context is then crucial for systems which are trying to assist these

groups in some fashion [10]. Before an understanding of the group's context can be reached, group and individual affiliations must be identified through the precess of group affiliation detection (**GAD**). Often times several groups can occupy the same space at once [16], making it important to detect non-affiliation as well as affiliation.

A group is two or more individuals who are connected to each other by social relationships [9]. Humans have an innate ability to visually recognize these groups quickly [16], using unconscious processes which can be described using the Gestalt Laws [9]. Our minds automatically observe and group objects together based on proximity, similarity and interaction. It is this perception process of detecting groups and affiliations which GAD proposes to emulate [17]. Since human-like perception is the goal, we are therefore bound to that perception as it defines correct and incorrect affiliation decisions. The problem is then to differentiate inter-group similarity from intra-group similarity.

Members of the same group have similar physical behavior because group members often perform activities together [17], adopt behavioral norms of the group [9], and mimic behavior of interaction partners [7]. By sensing these behavioral similarities, or "**social proximity**," we can effectively detect groups, and group affiliations [14]. Information from wearable sensors is centralized, features indicative of affiliation are extracted, and the result is clustered to identify groups and affiliations [17]. However in situations where centralized aggregation is not practical [11], such as emergencies [4], new methods for evaluating group affiliation using P2P analysis systems must be explored.

We present a method for P2P assessment of group affiliation by modeling the data as a distribution and then calculating the disparity (or similarity) as the Jeffrey's divergence between models from different individuals. We call this method divergence-based affiliation detection (**DBAD**). We compare DBAD with centralized and distributed approaches using signal correlation which is the basis for previous approaches [17, 14]. We present 2 methods for accomplishing GAD, one where nodes exchange Gaussian probability density functions (DBAD-P) of sensor data, and another where they exchange histograms of observations (DBAD-H). We evaluate these methods with an experiment involving 10 individuals

with varying group numbers, sizes and affiliations, conducting a homogenous activity: a scenario with high difficulty.

The DBAD methods perform significantly worse in terms of identifying inter and intra-group similarities at any given instant with a maximum of 63% compared to a 74% for centralized correlation. However, filtering similarity values over time improves recognition to 93%, approximate with centralized values. Compared to a centralized instance, DBAD-H and DBAD-P reduce energy consumption by 24% and 43% respectively, due mainly to reduced communication. DBAD-H reduces response time by 7%, but DBAD-P doubles it, where both approaches increase the mount of memory required, though is still well within tolerances for modern devices (under 4.5 kB). In total, the contribution is a method to estimate **P2P social proximity between individuals** by exchanging only model parameters, and a method for filtering this value over time to **estimate group affiliation**.

## 2. RELATED WORK
GAD differs from group activity recognition in that we are not interested in understanding the behavior, but rather in assessing if the behaviors are affiliated. Marin-Perianu et al. [15] proposed detecting groups of smart goods in supply chains by the degree of correlation between the sensor signals. This approach was later applied to human beings, where the correlation of acceleration signal variance was used to identify group affiliation [17]. There graph analysis methods are required to fuse input from multiple sensors [14].

Brdiczka et al. [5] recognized changes in group configurations by calculating the Jeffrey's divergence over histograms of multi-modal sensor data. There divergence indicates differences between group configurations, where here we investigate divergence as indicator of pair-wise group affinity. Since probability density functions over human trajectories characterize them well [6], it follows that these models could be useful for detecting similarities in that behavior. Blue-Tooth has also been used as a sensing modality to recognize device proximity [8], as have microphone sensors [18]. However the question remains the same: if a sensor is available which senses individual behavior, how can we judge social proximity, and classify affiliation in a peer-to-peer fashion, without exchanging timelines of measurements?

## 3. DIVERGENCE-BASED AFFILIATION DETECTION
The previous work on centralized approaches [17] describes GAD as following. Sensor data streams from devices monitoring potential group members are analyzed and behavior-relevant information is extracted, e.g. acceleration variance, as indicators of individual activity cues [17]. A cross-correlation ($\rho$) analysis of a given time window of these extracted signals is conducted in a pair-wise fashion, resulting in a disparity matrix $\overline{\mathcal{M}}$ in which index $i, j$ indicates the strength of the correlation between the observational data $\mathcal{D}$ of subject $i$ and subject $j$ over a period of time $t$.

$$\overline{\mathcal{M}}_{ij}^t = \rho(\mathcal{D}_i^t, \mathcal{D}_j^t) = \frac{\gamma(\mathcal{D}_i^t, \mathcal{D}_j^t)}{\sigma(\mathcal{D}_i^t)\sigma(\mathcal{D}_j^t)} \quad (1)$$

Here $\gamma$ is the covariance and $\sigma$ the variance over the windows. The multi-dimensional similarity graph represented by $\overline{\mathcal{M}}$ can then be clustered, resulting in an assignment of group affiliation. Cross-correlation requires both signals to be available at the same location ($\gamma$ in the numerator of Eq. (1)) for point-wise operations, requiring communication of raw signals in fully distributed systems.

We present a model-based approach to this problem called divergence-based affiliation detection (DBAD). Each device computes a model of local data based on the sensor signals it has collected over a specified time window. Devices evaluate similarity (social proximity) to each other based on the parameters of these models which are communicated in a P2P fashion. Here we demonstrate this using probability density functions (PDF) for modeling windows of local data, and applying the Jeffrey's divergence as an indicator of behavioral proximity to deduce affiliation. The approach works for **any sensing modality** which delivers similar values for similar inter-individual behavior (social proximity), and can in theory **combine several sensing modalities** into one decision.

### 3.1 Distributed Modeling
DBAD is then as follows. For each sample window, nodes extract relevant activity cues from the sensors. In this case the variance of acceleration signal magnitude, indicative of walking speed [17], and the circular mean of the orientation azimuth, indicative of walking heading. The circular mean of a vector of angles $\bar{\theta}$ consisting of $N$ angles $\theta$ is given by [2]:

$$\mu(\bar{\theta}) = \operatorname{atan2}\left(\frac{\operatorname{imag}(\bar{r})}{\operatorname{real}(\bar{r})}\right), \text{ where } \bar{r} = \frac{1}{N}\sum_{j}^{N} e^{i\theta_j} \quad (2)$$

Acceleration is modeled as a mixture of Gaussians, and orientation as a mixture of von Mises distributions [3] due to the circular nature of the data [6], given by:

$$\operatorname{vonMises}(\theta|\mu, m) = \frac{1}{2\pi I_0(m)} e^{m\cos(\theta-\mu)} \quad (3)$$

where the circular variance $\sigma$ is given by $\sigma(\bar{\theta}) = 1 - \bar{r}$ and $I_0(m)$ is a normalization coefficient, given the zeroth-order modified Bessel function of the first kind [3]. For both models, the number of components is identified using subtractive clustering, with expectation maximization for parameter fitting [3, 6]. The result is a mixture model consisting of $K$ Gaussian components:

$$P(x) = \sum_{k=1}^{K} \pi_k Distr_k(\mathcal{D}) \quad (4)$$

where the type of distribution $Distr_k(\mathcal{D})$ used depends on the data being modeled, using standard Gaussians $\mathcal{N}(x|\mu_k, \sigma_k)$ for acceleration, or $\operatorname{vonMises}(\theta|\mu_k, m_k)$ for orientation data.

### 3.2 Distributed Social Proximity Analysis
Once these mixture models have been fit, nodes (belonging to individuals) exchange the parameters with their single-hop neighbors. Each node $n_i$ in the set of all nodes with dimension $N$ can now calculate their disparity to neighboring nodes based on the Jeffrey's divergence. The Jeffrey's divergence
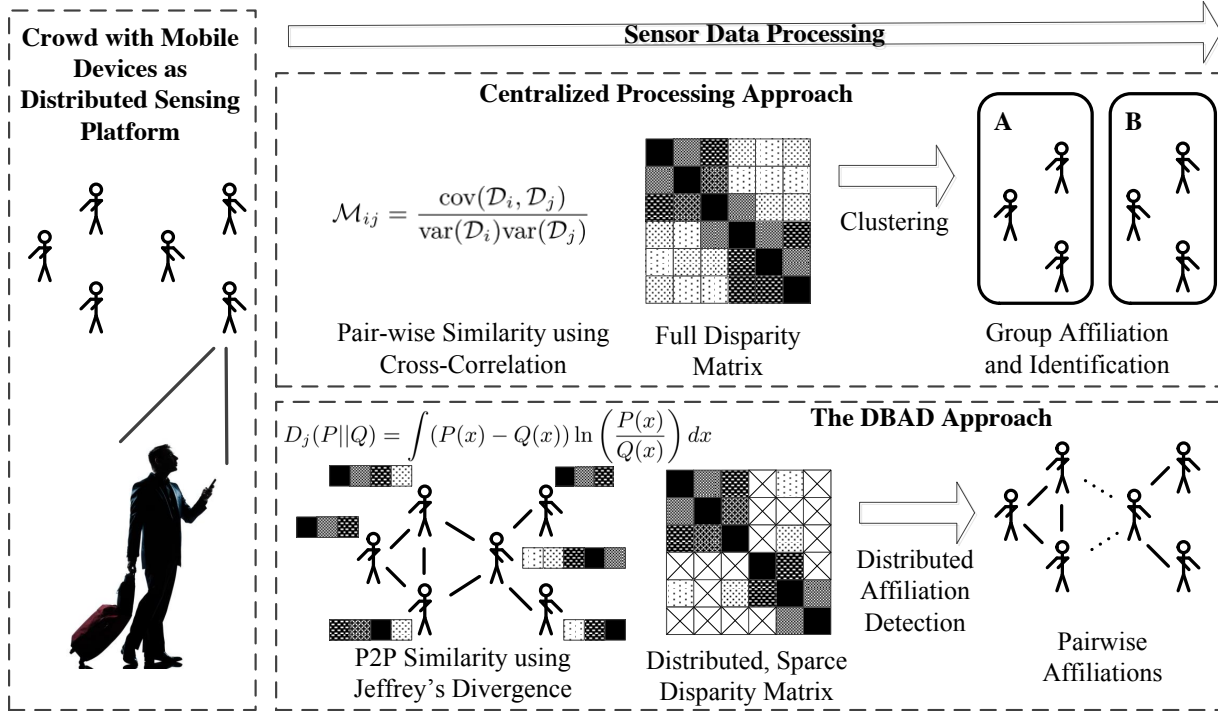
**Figure 1. The centralized and the novel distributed approach to group affiliation detection**

is an extension of the Kullback-Leibler divergence, selected because it is numerically stable and symmetric [5]. The Jeffrey's divergence $D_j$ between two distributions $P$ and $Q$ is given by:

$$D_J(P||Q) = \int (P(x) - Q(x)) \ln\left(\frac{P(x)}{Q(x)}\right) dx \quad (5)$$

Each node calculates its pairwise disparity to all other nodes within its single-hop communication neighborhood $\mathcal{V}^t$ at time $t$. Which nodes are in this neighborhood is dependent on the range of communication $\psi$ (complexities in wireless communication are not modeled), and the physical Euclidean distance between two nodes at the time:

$$\mathcal{V}^t = \{[n_i, n_j]\}|\text{dist}^t(n_i, n_j) \leq \psi \quad (6)$$

The behavioral distance between neighboring nodes can then be acquired as the value of the Jefferey's divergence between distributions of the sensor data of the two nodes.

$$\forall_{[n_i, n_j] \in \mathcal{V}^t} |\overline{\mathcal{M}}_{ij}^t = D_J(Dist(\mathcal{D}_i^t)||Dist(\mathcal{D}_j^t)) \quad (7)$$

In this way the $D_J$ is commutative and both nodes will conclude the same similarity based on the same models. In the centralized approach, the results of the complete pairwise metrics are centrally calculated, yielding a complete similarity matrix for all nodes as shown Fig. 1. Clustering this matrix to find affinity is a relatively straight-forward task, requiring only parameter fitting for clustering thresholds [17]. In a distributed approach this is not the case.

Each mobile device can only communicate with other nodes within reach of local P2P communication, which has 2 important repercussions. First, the similarity matrix is distributed

across the complete set of user devices and is not available to any single device. Since the assumption is that global communication is either unavailable, intermittently unavailable, or cannot be used for cost reasons (e.g. Bandwidth), it also implies that this distributed data entity cannot be directly queried by any single device. Second, its distributed nature also means that the disparity matrix is incomplete or sparse, as disparity is not measured between devices which are not within communication range (see Fig. 1). This presents a challenge of evaluating a distributed, sparse disparity matrix across multiple devices. Each row of the disparity matrix is located on a different device, and several positions contain no data (when $[n_i, n_j] \notin \mathcal{V}$). Since individuals are mobile over time, the vacancy of a position in the disparity matrix $\mathcal{M}_{ij}^t$ at time $t$ also varies over time as well.

A moving average of the disparity matrices is used as a low-pass filter to smooth the social proximity output over time. A buffer of length $b$ is used, where $b$ specifies the number of consecutive windows which are used in the filter. The social proximity between two individuals are averaged across the disparity matrices from those windows:

$$\widetilde{\mathcal{M}}_{ij}^t = \frac{1}{b} \sum_{\tau=0}^{b-1} \overline{\mathcal{M}}_{ij}^{t-\tau} \quad (8)$$

### 3.2.1 Distributed Affiliation detection
Once a social proximity indicator has been assessed, nodes must then convert this into affiliation information. A threshold-based approach was followed where each device makes a decision based on locally observed disparity values
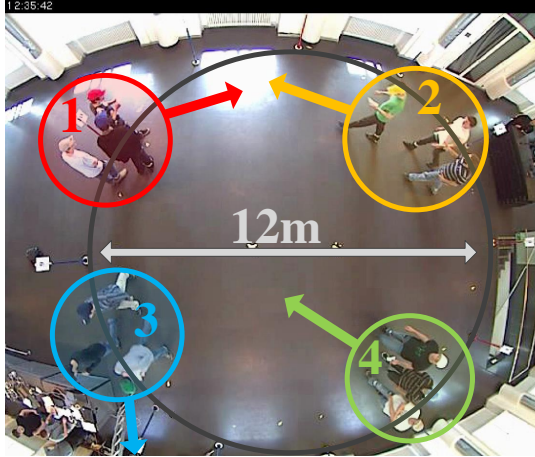
**Figure 2. A still image from the experiment video showing a four group configuration, with annotated group affiliating and heading**



**Figure 3. The Android devices used for subject monitoring (left) and the on-body position of the devices (right)**

and a predefined threshold $\phi$. For each node $n_i$, the classification is made using only the information in $\overline{\mathcal{M}}_{ij}^t$ where $[n_i, n_j] \in \mathcal{V}^t$, or the information local to the node at time $t$. The result is a subset $\mathcal{V}_{\text{affil}_i}^t \in \mathcal{V}$ of nodes which are affiliated withe node $n_i$ at time $t$, based on their disparity:

$$\mathcal{V}_{\text{affil}_i}^t := [n_i, n_j] \in \mathcal{V}^t | (\overline{\mathcal{M}}_{ij}^t \leq \phi) \qquad (9)$$

Where the converse is true for local non-affiliation decision:

$$\mathcal{V}_{\text{non-affil}_i}^t := [n_i, n_j] \in \mathcal{V}^t | (\overline{\mathcal{M}}_{ij}^t > \phi) \qquad (10)$$

The optimal value used for $\phi$ is dependent on the physical activity of the subject, as well as the sensors used to monitor that behavior. For practical purposes, the threshold can be experimentally obtained by maximizing the accuracy.

## 4. GROUP BEHAVIOR EXPERIMENT

Previous experiments with centralized group behavior detection [17] were conducted with groups performing various heterogeneous activities and acceleration sensors. However it is quite possible that the activity performed by all participants is homogeneous, e.g. walking, queuing in crowds [13]. To evaluate performance under these more difficult conditions, an experiment and data set was created using homogeneous activity behavior, namely walking, of several individuals in different group configurations.

The experiment was conducted in a large open room in a university setting. 12 subjects walked through the room in various group configurations while being monitored by wearable android mobile sensing devices attached to the hip of each subject as shown in Fig. 3. The devices monitored a single subject each using 3D accelerometers and magnetic field sensors (orientation), as well as ambient audio. For each subject, the data set contains 51 minutes of data, although 2 devices contained faulty motion sensors, leaving 10 usable subjects.

The experiment was recorded using a wide-angle lens on the ceiling of the room, and each subject was given head-gear of a different color to enable offline individual identification as shown in Fig. 2. 12 labeled posts where set up in a circle with

a diameter of 12 meters inside of a large room, where each post displayed a unique number clearly on a sign in clockwise order. A single member from each group was given a list of numbers, and each group then followed that member from post to post in the randomly assigned order on the list. Between experiments, group affiliations where reassigned and the experiment was repeated in the following configurations: one group (all together), 2 groups, 3 groups, 4 groups, no groups (each subject was given a separate list)[1]. Before each group experiment, subjects hopped in unison 3 times which was used to synchronize data by aligning the periods of free-fall (zero acceleration) across subjects.

Location data for each subject was annotated after the fact using a mixture of manual and automated color tracking software. For this purpose the video of the experiment was taken and the pixel coordinates of the subject's hat was tracked throughout the experiments. The location is given in pixel coordinates from the top left of the video. We converted these coordinates into meters using the diameter of the circle (12 meters = 430 pixels) as a reference. These coordinates contain the elliptical distortion of the wide-angle lens, but can theoretically be transposed into spacial coordinates using the known dimensions of the room and the location of the camera. We argue that for the purpose of this research, this approximation suffices.

The performance of both centralized and the DBAD algorithms was implemented in MATLAB and then simulated using this data set. The simulation was performed on both the accelerometer and orientation data respectively. For this purpose, the data from the experiments was cut up into windows whose length was varied. The variance of the acceleration data was calculated over a 15 second moving window, as this was shown to be effective for centralized forms of group affiliation detection in other scenarios [17]. Using the magnetometer, the azimuth was taken around the vertical axis of the subject, and a moving average of one second was used as an indicator of walking direction.
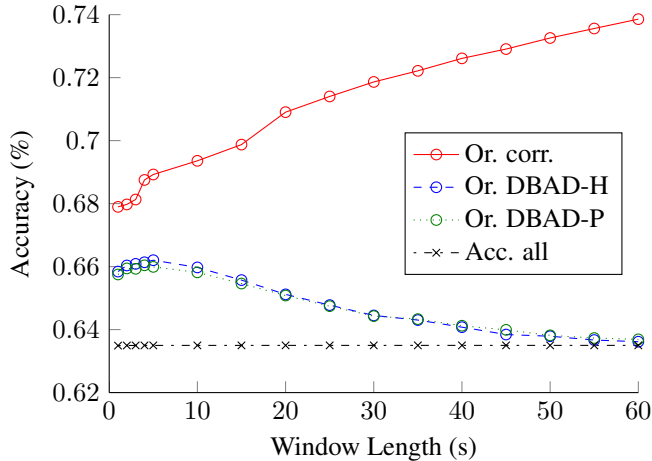
---

[1]`http://www.teco.kit.edu/~gordon/GAD/data_set.zip`

**Figure 4. Performance for cross-correlation (corr.) and DBAD over orientation (or.) and acceleration (acc.) signals**



**Figure 5. Performance for cross-correlation (corr.) and DBAD over orientation (or.) and acceleration (acc.) signals filtering over a 5s window**

GAD was then performed using the centralized approach based on the signal cross-correlation [17], as well as the DBAD algorithms. Numerical integration of a PDF is carried out by estimating a histogram of the PDF. In order to evaluate the effect of modeling error on performance, the same process was also conducted using histograms of the individual sample windows constructed using the data windows directly as well. The resulting sparse, distributed similarity matrices (see Fig. 1 were then classified using $\phi$ for both the PDF-based and histogram-based data, and the results where evaluated in terms of correct and incorrect pairwise affiliation detections. Pairwise affiliations are binary in nature, either indicating affiliation or non-affiliation of two subjects. However for a given group configuration, the distribution of affiliation and non-affiliation is not independent and identical.

## 5. EVALUATION

All algorithms were evaluated using a sliding window whose length was varied between 1 to 60s. The results of this simulation are shown in Fig. 4. The results using the accelerometer remain constant at around 63.5%, at the level of noise when accounting for the imbalance between $X_{\text{affil}}$ and $X_{\text{non-affil}}$. The centralized approach performed best of the three algorithms and improves monotonically with the length of the window, achieving just under 74% for a window length of 60 seconds. Using DBAD-H on the histograms and classifying the resulting complete similarity matrices yields an optimum of around 66.2% at 5 seconds, indicating weak representation of social proximity. Further increasing the sample window reduces the accuracy of the algorithm, as it asymptotically approaches the noise level at 60 seconds. DBAD-P performs only slightly worse than using a histogram, behaving similarly with an optimum of 66.0% at a window length of 5 seconds which then drops off into noise.

While not necessarily a negative result, the recognition rates achieved are not high enough to be useful. The filter in Eq. (8) was then applied to the optimum window length for the distri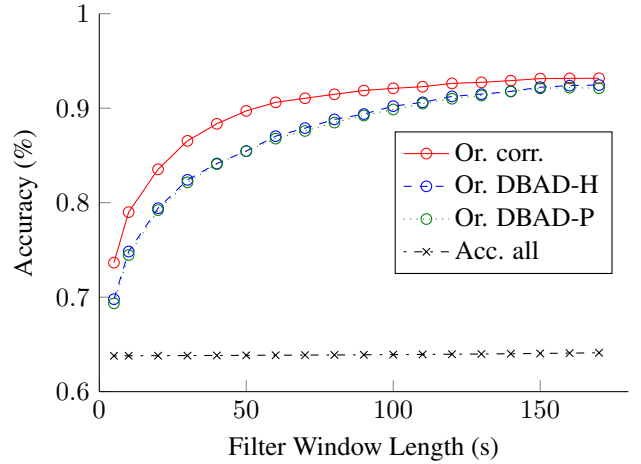buted algorithms of 5 seconds, where the resulting accu-racy over the length of the filter are shown in Fig. 5. Here the accuracy using the acceleration sensor remains almost constant showing only slight increases with filtering, indicating that noisy data is not causing the low values. Using orientation data however, the centralized approach as well as both distributed approaches benefit from filtering, eventually all converging to values of around 93.3%. Optimum values are reached after about 250 seconds of monitoring, or a filter of length 50 classifications over 5 second windows.

One of the goals of the proposed methods is to reduce communication volumes, thereby alleviating stress on the network and reducing battery life of the individual devices. We monitored the rate of exchange of data during the course of the simulations for the different algorithms, the results of which can be seen in Fig. 6. The centralized approach requires each node to exchange the entire sample window's worth of sensor data, in this case sampled at 50 Hz. Regardless of window length, 50 Hz of sensor data must be transferred per second, requiring 4 bytes of data per measurement, or 200 B/s.

For demonstration, this data was also compressed before transmission using a two-step differential encoding followed by the DEFLATE algorithm[2]. For smaller window sizes, the compression overhead reduces the advantages of compression (orientation) or even makes it counter-productive (acceleration), where as window size increases the savings become more pronounced, at about 175 B/s for acceleration and 150 B/s for orientation, being able to save around 12.5% and 25% respectively. The distributed algorithms however greatly outperform the centralized approaches. At their optimal window length of 5 seconds, communicating histograms between nodes (in this case 20 buckets) requires only 8 Bytes/second of communication, and communicating models a factor of 10 less. Concretely these are either $\pi, \mu$ and $\sigma$ values for acceleration data, $\pi, \theta$ and $m$ values for orientation data respectively), as shown in Eq. (4). This is 94.7% and 99.5% reduction when compared even to the centralized approach with
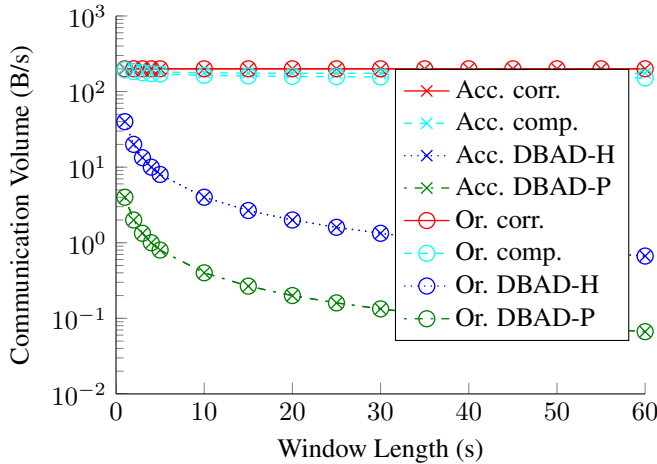
---

[2]https://tools.ietf.org/html/rfc1951

**Figure 6.** Communication volumes for cross-correlation (corr.), cross-correlation using compressed values (comp.) and DBAD over orientation (or.) and acceleration (acc.) signals



**Figure 7.** Performance over communication range for cross-correlation (corr.) and DBAD over orientation (or.) and acceleration (acc.) signals for a 5s window

lossless data compression for the histogram and model-based distributed methods respectively.

One major difference between the distributed approaches and the centralized approach is the use of P2P communication which has a limited communication range. We evaluated the effect of this by varying the effective communication range of individual nodes using the location information annotated from the video. For a given range, nodes are able to only communicate with other nodes which are within a circle with radius equal to the range.

Fig. 7 shows the accuracy results when the communication range of the devices is limited in simulation. At maximum range all nodes can communicate with each other across all experiments. As the range is decreased, the accuracy of the all methods increases to an optimum at 4.5m of 83.1% for the centralized approach, 79.6% for the histogram-based approach, and 81.2% for the approach using model divergence. Decreasing the communication range further incurs a sharp drop, with accuracy eventually dropping off to noise as the distance approaches 0. The optimum of 4.5m is there length where affiliated links are maximized and non affiliated links are minimized within the neighborhood of each node.

The results are demonstrated in Fig. 8, where similarity matrices are displayed instead of disparity for visibility reasons. Each row and column are subjects from 1 to 10, and index $i, j$ is the similarity between subject $i$ and $j$. In Fig. 8a) a typical clustering of a 5 second window by DBAD-P algorithm is shown for two groups. The difference in the similarity between subjects can be seen, but two groups can be identified, one in the upper left and one on the lower right. This also leads to noise in the identification of group affiliation in the same column of Tab. 1. In Fig. 8b) both are in different locations but the heading is similar, as is the case with groups 2 and 4 in Fig. 2. This leads to a drop in precision in Tab. 1 for that window. In Fig. 8c), a communication range of
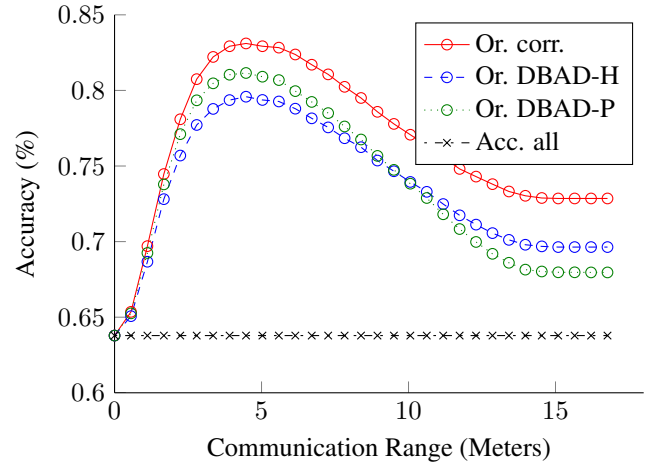
5m greatly increases precision as most inter-group links are removed, but recall lags, as intra-group similarity fails to correlate group affinity. Filtering over the entire experiment Fig. 8d) improves all values, but errors are still caused by intra-group similarity values. The problem with intra-group similarities is demonstrated by Fig. 2(1), where the heading of the individuals in the group differs dramatically. Note that here we use precision and recall for demonstration purposes, but for experiments where with 1 or no groups, F-score and either one or the other of these metrics is undefined.

Finally, we ran a simulation to compare the local resource footprints of the various approaches. In Tab. 2 are simplified approximations, calculated from the bitrate and power consumption of different communication technologies [1][3], and processing times, memory consumption, and data transmission volumes from the previous simulations. This is modeled on an Android Nexus 4 device where processing occurs on a single core which has a consumption of 0.5W. Detecting affiliation using distributed cross-correlation is impractical due to the high response time and total energy cost of classification. The costs are due to the high communication volumes and consumptions caused by communicating raw sensor data over P2P channels. The centralized approach however has an expensive communicator, but the high bandwidth means low communication times. Processing is also offloaded, therefore processing time is low, and total energy is low as well.

DBAD-H has low processing time because model fitting is avoided, and P2P communication reduces the energy even with the reduced bitrate. The total cost of energy of DBAD-H is therefore 24% lower than for centralized cross-correlation. DBAD-P has a more processing for model fitting and analysis than DBAD-P, and therefore increased response time as well, but the total energy required drops due to reduced communication. Nonetheless, DBAD-P reduces total energy con-

---

[3]`http://www.csr.com/sites/default/files/`
`white-papers/comparisons_between_low_power_`
`wireless_technologies.pdf`

**Table 1. Confusion matrices of affiliation (AF) and non-affiliation (NAF) with ground truth (GT) and classification (CL) and resulting metrics accuracy (acc.), precision (prec.), recall (rec.) and F-score, corresponding to the disparity matrices in Fig 8.**

| ↓ GT, CL → | AF | NAF | | ↓ GT, CL → | AF | NAF | | ↓ GT, CL → | AF | NAF | | ↓ GT, CL → | AF | NAF |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| AF | 38 | 12 | | AF | 40 | 10 | | AF | 38 | 12 | | AF | 42 | 8 |
| NAF | 6 | 44 | | NAF | 44 | 6 | | NAF | 8 | 42 | | NAF | 0 | 50 |

| Acc. | Prec. | Rec. | F-sc. | Acc. | Prec. | Rec. | F-sc. | Acc. | Prec. | Rec. | F-sc. | Acc. | Prec. | Rec. | F-sc. |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 0.82 | 0.86 | 0.76 | 0.6 | 0.46 | 0.48 | 0.8 | 0.81 | 0.8 | 0.83 | 0.76 | 0.79 | 0.92 | 1 | 0.84 | 0.91 |
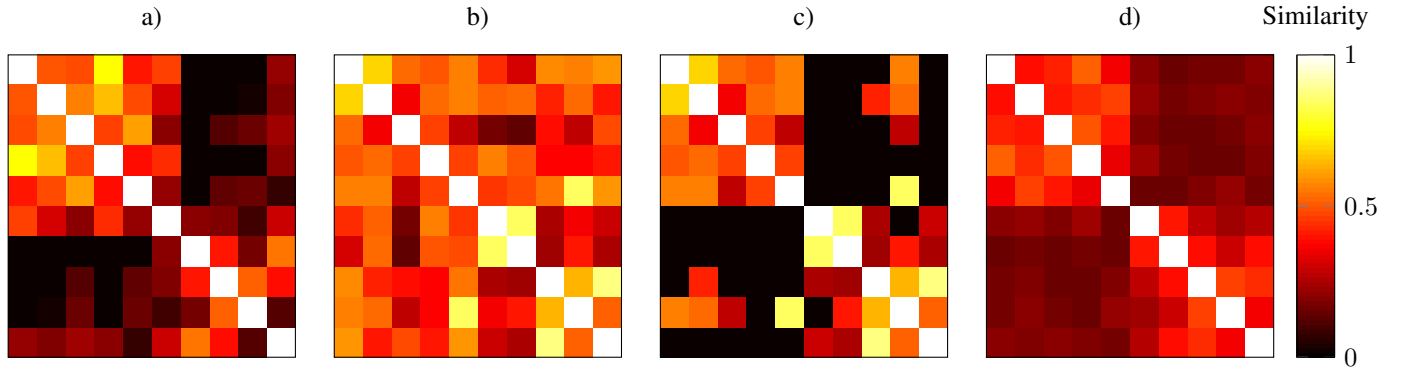


a)  b)  c)  d)  Similarity

**Figure 8. Similarity between subjects in a two-group experiment for a window size of 5s using orientation and the DBAD-P method. a) under normal conditions, b) when both groups have similar headings, c) when the communication range is 5m and d) when averaged over the whole experiment.**

sumption with respect to DBAD-H by a further 24% or by 43% compared to centralized cross-correlation.

## 6. DISCUSSION

Due to the nature of the problem, subjects who are in the same group generate similar sensor patterns for reasons discussed in Sec. 1. However, subjects in different groups may appear to be similar for periods of time, e.g. when both groups walk in the same direction, as is the case with groups 2 and 4 in Fig. 2. By observing subjects for a long enough period (extending window size), the centralized approach can make these temporary phenomena irrelevant as demonstrated in Sec. 5. For the distribution-based approaches however, extending the window size reduces effectiveness as the characteristics of the signal disappear into a flat distribution after enough directional changes (see Fig. 4) [12]. This effect is also compounded by a weakness in the distributed methods themselves, as PDFs and histograms both ignore the time component of the signals.

The filter is so effective because it allows the p2p methods to deal with short-term similarity between non-affiliated subjects by extending the observation range for any given affiliation decision. Reducing communication range however can remove these ambiguities entirely, as the members of different two groups are often not compared with each other if they are outside the communication range $\psi$ (again observe groups 2 and 4 in Fig. 2).

Applications for P2P GAD include life-logging systems which could document who we spent time with, and social network applications for sharing that information. Also DBAD reduces privacy concerns as data is only shared in an aggregated form (models). It is also a useful tool for sup-

porting group activity recognition where the behavior of the group can be identified once the constituents are known [10]. Applications are also foreseen in emergency situations, where groups display affiliative behaviors and should be treated as a single unit [17]. DBAD also has a great potential for crowd emergencies where infrastructure is usually the first casualty. Advantageous is the fact that the effort required by each device is dependent on the density of the crowd (number of neighbors) and not crowd size. This is however an avenue of further research as scalability is not directly evaluated here.

## 7. CONCLUSION

Humans often build groups for social reasons, and because groups can be better at reaching goals than the individuals separately [9]. However, often several different groups have different goals and occupy the same space, and must therefore be differentiated. Contemporary differentiation methods require the central aggregation of sensor information, however this is not feasible when network communication is too expensive, either due to the scale or the environment.

We present a method for distributed, P2P recognition of group affiliations using the divergence of sensor data distributions as an indicator of similarity (DBAD). Divergences can either be calculated using models of individual behavior (DBAD-P) or using histograms of sensor data (DBAD-H). The requirement is that the sensor used is an indicator of social proximity for the kinds of social connections which define the particular group. The results show that the output of the proposed method fluctuates with instantaneous recognition rates only slightly over random. However group affiliations can still be detected 93% of the time by applying a low-pass filter to that output signal.

**Table 2. Resource Consumption Analysis for a centralized correlation approach, distributed correlation approaches and novel DBAD-based methods for 1 Classification of 5 Seconds**

| Approach | Memory Used (B) | Comm. Tech. | Comm. Per Class. (B) | Comm. Time (ms) | Comm. En. (mJ) | Proc. Time (ms) | Proc. En. (mJ) | Total Time (ms) | Total En. (mJ) |
|---|---|---|---|---|---|---|---|---|---|
| Cent. Corr. | 2000 | 3G | 1500 | 29.8 | 13.53 | 61.95 | 0.03 | 91.75 | 13.56 |
| Acc. Comp. | 2000 | BT 4.0 | 18000 | 417.66 | 191.85 | 110.4 | 0.06 | 528.05 | 191.91 |
| Or. Comp. | 2000 | BT 4.0 | 21000 | 487.27 | 223.83 | 110.4 | 0.06 | 597.66 | 223.89 |
| DBAD-Hist | 2160 | BT 4.0 | 960 | 22.28 | 10.23 | 62.98 | 0.03 | 85.26 | 10.26 |
| DBAD-PDF | 4280 | BT 4.0 | 720 | 16.71 | 7.67 | 168.47 | 0.08 | 185.17 | 7.76 |

We show that only having a limited range of communication actually improves system performance, by allowing the devices to implicitly use location information without requiring a further sensor. Analysis of resource consumption indicates that time-series analysis approaches in the network are infeasible due to time and energy required for communication. DBAD-H and DBAD-P reduce energy consumption by 24% and 43% respectively, where DBAD-H reduces response time by 7%, but DBAD-P doubles it, indicating there is a trade-off between energy consumption and response time. Both distributed methods increase the amount of memory used, although usage remains under 4.5 kB. None the less, both methods are independent of centralized resources and can be applied in distributed P2P systems.

## 8. REFERENCES

1. Balasubramanian, N., Balasubramanian, A., and Venkataramani, A. Energy consumption in mobile phones: a measurement study and implications for network applications. In *Proceedings of the 9th Internet measurement conference*, IMC '09, ACM (2009), 280–293.

2. Berens, P. Circstat: A matlab toolbox for circular statistics. *Journal of Statistical Software 31*, 10 (9 2009).

3. Bishop, C. M. *Pattern Recognition and Machine Learning*, vol. 4 of *Information science and statistics*. Springer, 2006.

4. Boin, A., and McConnell, A. Preparing for Critical Infrastructure Breakdowns: The Limits of Crisis Management and the Need for Resilience. *Journal of Contingencies and Crisis Management 15*, 1 (Mar. 2007), 50–59.

5. Brdiczka, O., Maisonnasse, J., Reignier, P., and Crowley, J. L. Detecting small group activities from multimodal observations. *Applied Intelligence 30*, 1 (July 2007), 47–57.

6. Calderara, S., Prati, A., and Cucchiara, R. Mixtures of von mises distributions for people trajectory shape analysis. *Circuits and Systems for Video Technology, IEEE Transactions on 21*, 4 (2011), 457–471.

7. Chartrand, T. L., and Bargh, J. a. The chameleon effect: the perception-behavior link and social interaction. *Journal of personality and social psychology 76*, 6 (June 1999), 893–910.

8. Do, T. M. T., and Gatica-Perez, D. Groupus: Smartphone proximity data and human interaction type mining. In *Wearable Computers (ISWC), 2011 15th Annual International Symposium on* (2011), 21–28.

9. Forsyth, D. *Group Dynamics*, 4th ed. International student edition. Thomson/Wadsworth, 2006.

10. Gordon, D., Scholz, M., and Beigl, M. Group activity recognition using belief propagation for wearable devices. In *Wearable Computers (ISWC), International Symposium on* (2014), 1–10.

11. Gordon, D., Scholz, M., Ding, Y., and Beigl, M. Global peer-to-peer classification in mobile ad-hoc networks: a requirements analysis. In *Proceedings of the 7th conference on Modeling and using context (CONTEXT)*, Springer (2011), 108–114.

12. Gordon, G. R., and Surkis, J. A general simulation methodology for optical bar code evaluation. *Journal of Statistical Computation and Simulation 8*, 2 (1978), 85–101.

13. Helbing, D., and Mukerji, P. Crowd Disasters as Systemic Failures: Analysis of the Love Parade Disaster. *Disasters* (2012).

14. Kjærgaard, M. B., Wirz, M., Roggen, D., and Tröster, G. Detecting pedestrian flocks by fusion of multi-modal sensors in mobile phones. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, UbiComp '12, ACM (New York, NY, USA, 2012), 240–249.

15. Marin-Perianu, R., Marin-Perianu, M., Havinga, P., and Scholten, H. Movement-based group awareness with wireless sensor networks. In *Proceedings of the 5th international conference on Pervasive computing*, PERVASIVE'07, Springer-Verlag (Berlin, Heidelberg, 2007), 298–315.

16. Moussaïd, M., Perozo, N., Garnier, S., Helbing, D., and Theraulaz, G. The Walking Behaviour of Pedestrian Social Groups and Its Impact on Crowd Dynamics. *PLoS ONE 5*, 4 (2010), 7.

17. Roggen, D., Helbing, D., Tröster, G., and Wirz, M. Recognition of crowd behavior from mobile sensors with pattern analysis and graph clustering methods. *Networks and Heterogeneous Media 6*, 3 (Aug. 2011), 521–544.

18. Wirz, M., Roggen, D., and Troster, G. A wearable, ambient sound-based approach for infrastructureless fuzzy proximity estimation. In *Wearable Computers (ISWC), 2010 International Symposium on* (2010), 1–4.