Learning to identify leaders in crowd

Francesco Solera Simone Calderara Rita Cucchiara
University of Modena and Reggio Emilia
{name.surname}@unimore.it

Abstract

Leader identification is a crucial task in social analysis, crowd management and emergency planning. In this paper, we investigate a computational model for the individuation of leaders in crowded scenes. We deal with the lack of a formal definition of leadership by learning, in a supervised fashion, a metric space based exclusively on people spatiotemporal information. Based on Tarde's work on crowd psychology, individuals are modeled as nodes of a directed graph and leaders inherits their relevance thanks to other members references. We note this is analogous to the way websites are ranked by the PageRank algorithm. During experiments, we observed different feature weights depending on the specific type of crowd, highlighting the impossibility to provide a unique interpretation of leadership. To our knowledge, this is the first attempt to study leader identification as a metric learning problem.

1. Introduction

There is no single, agreed, computational definition of a crowd. Still, it is beneficial to distinguish between a gathering of people who simply share a location and a psychological crowd, whose members also share a social identity. Reicher's *Social Identity Theory* [16] explains the ability of a crowd to spontaneously behave in a socially coherent manner without any apparent or explicit exchange of information. The theory proposes that people shift from a personal and private identity to a shared one, resulting in strongly influenced behaviors. A sense of shared social identity is often created in the event of an emergency, enabling crowd members to act as a source of strength for one another, when dealing with regularly recurring crowds, such as football matches and it is proportional to the extent of organization within the crowd itself. A more organized crowd is more likely to exhibit pre-planned, antisocial behavior such as in a demonstration or a protest [3].

Nevertheless, not all the members of a crowd undergo the same level of identity shift [7, 4, 14]. People who define the norms and the values which then become shared among all the other members are recognized as *leaders*. Leader identification is a crucial task in crowd management, emergency planning and sociological analysis. By finding and disconnecting the leaders from the rest of the crowd, efficient containment can be accomplished. On the other hand, an influential voice of non-violence in a crowd can lead to a mass sit-in and a strong leader can take control of an emergency situation, initiate movement and guide suitable crowd behavior avoiding panic [3]. Either way, leaders are key subjects to pay attention to when dealing with otherwise unmanageable crowds.

1.1. The leader in crowd psychology

There is a fundamentally problematic relationship between the leader and the crowd. Crowd phenomena emerged as a democratic repulsive response to traditional models of leadership; it almost seems paradoxical to question the need of a crowd to reestablish a leader. In the last century crowd psychology has tried to understand the true role of leaders in crowds and the mechanisms by which they are selected.

In his book *The Crowd* [7], Le Bon justifies the existence of a leader through the need of the crowd to place trust in someone able to provide orientation and contribute to its overall stability. A leader serves as a guide and members of the crowd are led by direct imitation of the leader's will. Certainly, the leader is an important role to preserve the crowd stability, but he is neither a founding figure nor is he permanently established, see Fig. 1a. Instead, the crowd formation is an emergent process that uses the leader as a stabilizer. While for Le Bon the leader is an elected anonymous figure, in Freud's *Group Psychology* [4] the leader comes to play a constitutive role and every member of the crowd identifies the leader as their "I"-ideal. With respect to Le Bon self-organizing and emergent notion of crowds, Freud provides a highly centralized model of the social community, appreciable in Fig. 1b. The relation between the leader and the crowd is thus radically asymmetric as the members are submitted under the leader; relations between crowd members are secondary. Eventually, and to even a greater extent than Le Bon, Tarde [14] emphasizes the self-

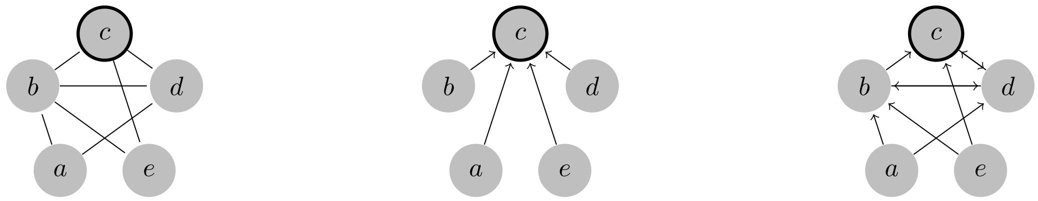

(a) Le Bon's theory (b) Freud's theory (c) Tarde's theory

Figure 1: Different interpretation of relationships between a leader and other members in different crowd psychology theories.

referential emergence of crowd phenomena but characterizing the leader as the "spark" behind the organizational patterns. Unlike Freud, Tarde was not interested in the leader's foundational role. Instead, he analyzes how the leader contributes to the flow of imitation in a society and builds its theory assuming that every process of imitation begins with asymmetry, as highlighted in Fig. 1c. Nevertheless, as individuals initiative converge with leadership, the leader's identity may change without altering the crowd stability.

1.2. Related works

In crowd modeling and simulation, the problem of leader identification is typically cast as the definition of a set of geometrical rules that agents can follow at simulation time to select a leader inside their visual cone, providing virtual individuals with additional attraction forces [10, 12]. Examples of these rules include thresholds on distance and relative velocity with respect to anyone with whom they share their walking direction. Leaders identification becomes thus an individual task, possibly inconsistent across different members of the same crowd, strictly direction-oriented and unaware of social or behavioral information.

On the other hand, in visual crowd analysis leader identification is still an emerging topic. Andersson *et al*. [1] state an individual can be said to be a leader if it does not follow anyone and is followed by sufficiently many others at a proximate distance for a minimum amount of time. The concept of following is again based on geometrical and static rules only, denying the complex social dynamics. Yu *et al*. [17] provide a solution based on iterative modularity clustering over distance information. They simultaneously retrieve social groups and, for each group, the leader is defined as the member providing the dominant contribution to the clustering bisecting eigenvector. By construction, their leader selection procedure is based on a centrality measure computed over the distance information graph, which is only a small part of the behavioral clues that could be exploited in the identification task. As a consequence, this method applicability is mostly limited to standing scenarios.

A more complex approach is presented in Carmi *et al*. [2], where a Bayesian Network is built upon members' position and velocity information and the causality parameters are estimated by regression. By summing over the incoming and outgoing causality links over each member, it is possible to measure individual's contribution in shaping the group's collective behavior. More interestingly, Kjærgaard, *et al*. [5] and Sanchez-Cortes *et al*. [13] complete the task by ranking group members. The former builds a graph based on time-lag spatiotemporal features and employs PageRank to score the nodes by their importance. This time-lag approach mimics the concept of temporal causality introduced in Carmi *et al*. [2]. The latter instead, exploit non-verbal features to obtain sets or ranking scores that are uniformly combined to obtain the leader. Nevertheless, both [2, 5, 13] do not propose a method for weighting the contribution of the different features, which may vary from scene to scene. For this reason, in our method we introduce a metric learning approach over the feature space.

2. Method overview

Building on Tarde's theory [14] for the individuation of leaders in crowds, in Sec. 3 we introduce a set of pairwise features to evaluate social bonds among group members. Our approach takes as input individuals spatiotemporal information (*i.e*. people trajectories) and their partitioning into social groups, as shown in Fig. 2. During training, each feature is used to produce a separate ranking of the members of the considered group and Structural SVM [15] is employed to combine different features contribution, Sec. 4. At test time, for each group, the algorithm returns a ranking of the members, among which the highest will be predicted as the leader. This ranking is computed through the PageRank algorithm [11], able to take advantage of the referential and asymmetric structure of the crowd model (Fig. 1c) to assess the importance of each member. The process can be generalized to the whole crowd by considering larger groups.

Notation Given a video, the sets of trajectories from each group $\mathbf{x} = (T_a, T_b, \dots)$ form the input to our algorithm, where the generic trajectory $T_a = \{(t, x_a^t, y_a^t)\}_t$ contains ground plane metric information. The training set is $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_i^n$ where $\mathbf{y}_i$ is the correct ordering of the members of the i-th group in the space of all possible orderings $\mathcal{Y}(\mathbf{x}_i)$. The leader of $\mathbf{x}_i$ is the j-th member if $\mathbf{y}_i(j) = 1$.

3. Social features for leader identification

Following Tarde's perspective, leadership is a concept that is established through an agreement among group members. This agreement reflects in the way members influence each others path and is consequently visually graspable. We refer to these features as time-lagged features, f_T, $f_{\dot{T}}$, $f_{\ddot{T}}$ and f_I, detailed in Sec. 3.1. Additionally, following relations have also been modeled by means of empirical models built upon psychological experimental findings. We consider these model based features, f_s, f_d and f_r in Sec. 3.2. Eventually, in groups with more than two members, individuals position inside group f_C and the size of the group itself f_M have proven useful in unveiling leadership [12], Sec. 3.3. Consequently, the feature vector $\mathbf{f}$ between members a and b is described as $\mathbf{f}(a,b) = [f_T, f_{\dot{T}}, f_{\ddot{T}}, f_I, f_s, f_d, f_r, f_M, f_C]$.

3.1. Time-lagged features

A generic time-lagged feature $f_\psi(a,b)$ is computed evaluating a feature function ψ on a pair of trajectories T_a and T_b under varying discrete temporal shifts (*i.e.* time-lag) $z \in [-K, +K]$, where K is a fixed number of frames. In order to obtain a scalar value, $f_\psi(a,b)$ is a inverse weighted average over all the shifts:

$$f_\psi(a,b) = \frac{1}{\mathrm{M}} \sum_z z \cdot \psi(a,b;z)^{-1} \tag{1}$$

where values $\psi(a,b;z)$ are obtained by fixing T_a and shifting T_b of z frames and $\mathrm{M} = \sum_z \psi(a,b;z)^{-1}$ is a normalization factor. Eq. (1) evaluates the amount of influence (in terms of feature ψ) of T_b over T_a; precisely if $f_\psi(a,b)$ is positive, we say T_b leads T_a with certainty $|f_\psi(a,b)|$. In this work we consider as feature function ψ the Dynamic Time Warping (DTW) distance between two time series considering their spatial coordinates ($\psi \equiv T$), velocity ($\psi \equiv \dot{T}$), acceleration vectors ($\psi \equiv \ddot{T}$) and mutual influence ($\psi \equiv I$). DTW is useful to mitigate the small measurement noise that occurs when observing real world trajectories and to neglect moderate variations that naturally occur in moving groups.

Members have the tendency to follow the leader in order to keep the group compact and reduce the risk of falling too far from other members as well. This behavior reflects in the way spatiotemporal information are correlated among members of the same group. The distance feature f_T is computed by applying DTW to trajectories as a sequences of point coordinates. In the same way, $f_{\dot{T}}$ and $f_{\ddot{T}}$ are computed by employing velocity and acceleration discrete measurements. Conversely, the mutual influence feature f_I is an unsigned f_T, more formally $I(a,b;z) = T(a,b;z)\ \mathrm{sign}(z)$. This feature measure only the strength of the interaction but ignores the causality direction, highly penalizing members who do not interact inside groups.

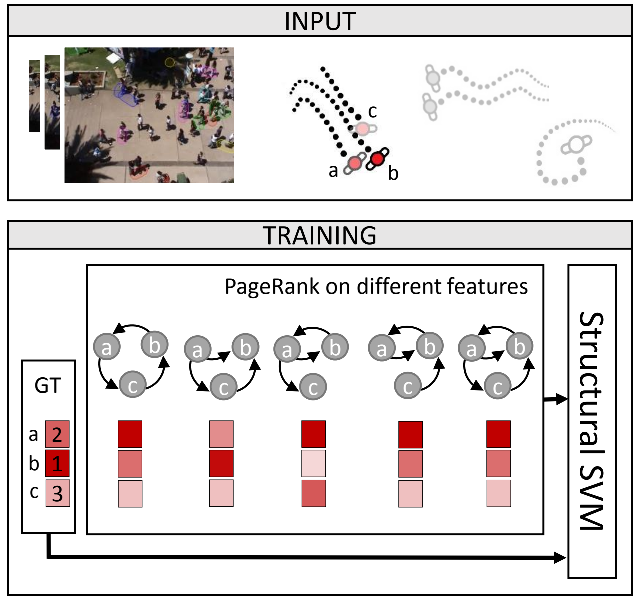

Figure 2: Method overview: for each group, a set of PageRank is computed. Given the correct labeling (GT), the Structural SVM learns how to aggregate ranks from different features. Color intensity indicates leadership probability.

3.2. Model based features

Group dynamics literature has established different collective locomotion patterns to characterize dyadic interaction among members of a groups. These models consider features such as speed, position, acceleration or a combination of them and, through formal empirical equation, describe how an individual a should move if it were to be considered as follower of b. The speed model [8] states that the follower acceleration $\ddot{T}_a$ aims at compensating speed variations between he and the leader b:

$$\ddot{T}_a^s(t) = c_d[\dot{T}_b(t) - \dot{T}_a(t)]. \tag{2}$$

A different strategy is for the follower a to keep its distance to the group leader b fixed [6]. Accordingly, the follower will accelerate (decelerate) to reduce (increse) this distance:

$$\ddot{T}_a^d(t) = c_s[T_b(t) - T_a(t) - \gamma], \tag{3}$$

where γ is the average distance among group members in the considered time window. Eventually, Lemercier *et al.* [9] observed that in walking scenarios the relative speed between a leader and a follower goes to zero but modulated by both the follower current speed and its distance from the leader. This feature is referred to as the speed ratio [5]:

$$\ddot{T}_a^r(t) = c_r\ \dot{T}_b(t) \frac{\dot{T}_b(t) - \dot{T}_a(t)}{T_b(t) - T_a(t)} \tag{4}$$

In order to average out the noise in measurements, the final features f_s, f_d, f_r between a leader b and a follower a are computed as the negative exponential of the mean difference between the observed acceleration vector $\ddot{T}_a(t)$ of the follower a and the one obtained using the different models $\ddot{T}^s_a(t)$, $\ddot{T}^d_a(t)$ and $\ddot{T}^r_a(t)$.

3.3. Personal and group level features

It is known from literature [12] that people tend to dispose in a specific set of predefined configurations when moving inside a crowd. These configurations, which depend on the group size, usually encourage social interactions by placing the leader in a central position. As a consequence, group size and individual position of a member w.r.t. all the others in the group may result in a discriminant information for leadership detection. The group size f_M is a group level feature as it is equal across all the members of a group. Still it is important as the normalization induced by the PageRank produces different values for groups of different sizes. The centrality feature f_C of a member inside the group is computed as the average distance of that member to all the others. This is an individual feature and can be thought, in terms of PageRank, as the transition probability of reaching a destination node independently of the starting one.

4. Learning to rank leaders

To what extent the chosen features may contribute to identify the leader is a matter of scenes and context. In dense crowds leaders will lead the way through other groups and distance won't be a peculiar factor, as opposed to the time to contact information. On the other hand, when the crowd is sparse, groups will tend to reorganize in formations easing social interactions among members of the same group and distance could here be a dominant factor. Other than crowd density, cultural habits and environmental constraints may also play an important role in the way leaders act. For these reasons a learning approach becomes mandatory for this task.

4.1. Features PageRank

Given a group $\mathbf{x}_i$, we build a graph of following relations $G^i_f = (\mathbf{x}_i, f)$ for each feature $f \in \mathbf{f}$, where nodes represents the members and the edges contains the pairwise feature information. Note that this graph is directed and asymmetric, similar to the one reported in Fig. 1c. At this point a $[0, 1]$-normalization is required to guarantee each element of G^i_f lays between 0 and 1 and the sum of elements in each row is 1. These conditions make the graph a row stochastic matrix and we can now interpret it as a transition graph, where each edge $(T_a, T_b) \in (\mathbf{x}_i \times \mathbf{x}_i)$ indicates the probability of considering T_b a leader for T_a. Under these premises, the PageRank algorithm is used to score each node. Due to the properties of the graph G^i_f, the eigensystem $(G^i_f)^T \pi^i_f = \pi^i_f$ has a unique solution, being π^i_f the dominant left eigenvector of G^i_f. The j^{th} entry of π^i_f is the score computed for the group member j. The algebraic solution of π^i_f is given by

$$\pi^i_f = (\mathbf{I} - d\mathcal{M}^i_f)^{-1} \frac{1-d}{|\mathbf{x}_i|} \mathbf{1}, \tag{5}$$

where d is the PageRank damping factor and $|\mathbf{x}_i|$ is the size of group $\mathbf{x}_i$. $\mathcal{M}^i_f = (D^{-1} G^i_f)^T$ is a column stochastic matrix and D is the outdegrees diagonal matrix of G^i_f.

4.2. Structural SVM and optimization

We can now formulate the prediction problem as finding the groups member which obtains the best combined ranking score across all different feature graphs. We can lower bound the problem as a supervised rank aggregation, aiming to find the shared ranking across all feature graphs G_f which maximizes the probabilities in highest positions:

$$F(\mathbf{x}; \mathbf{w}) = \arg\operatorname{sort}(\mathbf{w}^T \Pi^T) = \arg \max_{\mathbf{y} \in \mathcal{Y}(\mathbf{x})} \mathbf{w}^T \Pi^T \bar{\mathbf{y}}, \tag{6}$$

where the sort is considered in descending order.
Here $\Pi = [\pi_{f_T}, \pi_{f_{\dot{T}}}, \pi_{f_{\ddot{T}}}, \pi_{f_I}, \pi_{f_s}, \pi_{f_d}, \pi_{f_r}, \pi_{f_M}, \pi_{f_C}]$ is the PageRank scores concatenation for group $\mathbf{x}$ and $\bar{\mathbf{y}} = |\mathbf{x}| - \mathbf{y} + 1$ for notation convenience. By observing Eq. 6, we can write the problem as a linear combination of the weight vector and some combined representation $\mathbf{\Psi}(\mathbf{x}, \mathbf{y}) = \Pi^T \bar{\mathbf{y}}$. The $\mathbf{w}$-parameterization can now be learned through structured learning. Structural SVM [15] casts the problem of learning in complex and interdependent output spaces as a maximum margin problem:

$$\begin{aligned} \min_{\mathbf{w}, \boldsymbol{\xi}} \quad & \frac{1}{2}\|\mathbf{w}\|^2 + \frac{C}{n}\sum_{i=1}^{n} \xi_i \\ \text{s.t.} \quad & \forall i : \xi_i \geq 0, \\ & \forall i, \forall \mathbf{y} \in \mathcal{Y}(\mathbf{x}_i) \backslash \mathbf{y}_i : \mathbf{w}^T \delta\mathbf{\Psi}_i(\mathbf{y}) \geq \Delta(\mathbf{y}_i, \mathbf{y}) - \xi_i, \end{aligned} \tag{7}$$

where $\delta\mathbf{\Psi}_i(\mathbf{y}) = \mathbf{\Psi}(\mathbf{x}_i, \mathbf{y}_i) - \mathbf{\Psi}(\mathbf{x}_i, \mathbf{y})$, ξ_i are the slack variables introduced in order to accommodate for margin violations and $\Delta(\mathbf{y}_i, \mathbf{y})$ is the loss function. Intuitively, we want to maximize the margin and jointly guarantee that for a given input, every possible output result is considered worst than the correct one by at least a margin of $\Delta(\mathbf{y}_i, \mathbf{y}) - \xi_i$, where the loss is bigger when $\mathbf{y}_i$ and $\mathbf{y}$ are known to be more different. In particular, to measure how far a member was classified w.r.t. its proper positions, we employ a simple squared norm $\Delta(\mathbf{y}_i, \mathbf{y}) = \|\mathbf{y}_i - \mathbf{y}\|^2$.

The quadratic program QP (7) introduces a constraint for every wrong ranking of the group members. In order to deal with this high number of constraints, we replaced them by n piecewise-linear ones by defining the structured hinge-loss:

$$\widetilde{H}(\mathbf{x}_i) \stackrel{\text{def}}{=} \max_{\mathbf{y} \in \mathcal{Y}(\mathbf{x}_i)} \Delta(\mathbf{y}_i, \mathbf{y}) - \mathbf{w}^T \delta\mathbf{\Psi}_i(\mathbf{y}). \tag{8}$$

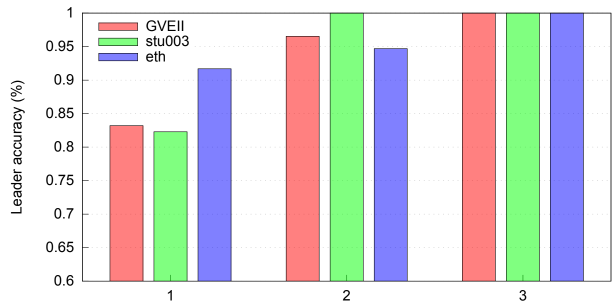

Figure 3: Cumulative accuracy when the leader is allowed to occupy positions up to 3 in the predicted group ordering.

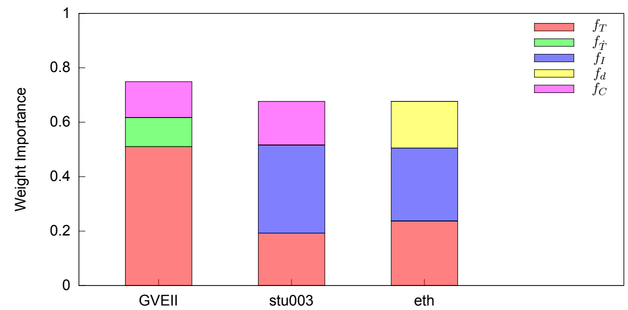

Figure 4: Learned weight importance for top 3 features for every dataset.

	our	[5]	[13]	SVM
`GVEII` groups 117 - 75 - 11	83.2 ± 0.01	73.5	69.1	67.4
`stu003` groups 87 - 20 - 8	82.3± 0.08	71.6	70.3	78.3
`eth` groups 37 - 10 - 10	92.4± 0.20	84.1	83.2	80.1
average results	85.98±0.1	76.4	74.2	75.2

Table 1: Leader detection comparative results. Number of pairs, triplets and groups with more than three members are specified under the dataset name.

The computation of the structured hinge-loss for each element i of the training set amounts to finding the most "violating" output $\mathbf{y}_i^*$ for a given training pair $(\mathbf{x}_i, \mathbf{y}_i)$. Now, we only have n constraints of the form $\xi_i \geq \widetilde{H}(\mathbf{x}_i)$ and QP (7) reduces to $\min_{\mathbf{w}} \frac{1}{2}\|\mathbf{w}\|^2 + \frac{C}{n}\sum_{i=1}^{n} \widetilde{H}(\mathbf{x}_i)$.

By disposing of a maximization oracle, *i.e.* a solver for Eq. (8), and a computed solution $\mathbf{y}_i^*$ given a generic example $\mathbf{x}_i$, cutting plane or subgradient methods (*e.g.* [15]) can easily be applied to the reduced quadratic program, being $\partial_{\mathbf{w}} \widetilde{H}(\mathbf{x}_i) = -\delta \mathbf{\Psi}_i(\mathbf{y}_i^*)$. The choice of the loss, which is linear w.r.t. the maximization argument $\mathbf{y}$, let us search for the most violating constraint as follows:

$$\begin{aligned} \mathbf{y}_i^* &= \arg \max_{\mathbf{y} \in \mathcal{Y}(\mathbf{x}_i)} \|\mathbf{y}_i - \mathbf{y}\|^2 + \mathbf{w}^T \Pi_i^T \bar{\mathbf{y}} \\ &= \arg \max_{\mathbf{y} \in \mathcal{Y}(\mathbf{x}_i)} -2\mathbf{y}_i^T \mathbf{y} + \mathbf{w}^T \Pi_i^T \bar{\mathbf{y}} \\ &= \arg \max_{\mathbf{y} \in \mathcal{Y}(\mathbf{x}_i)} (2\mathbf{y}_i + \Pi_i \mathbf{w})^T \bar{\mathbf{y}}, \end{aligned} \tag{9}$$

by noting that $\|\mathbf{y}_i\|^2 = \|\mathbf{y}\|^2$ does not depend on the particular choice of $\mathbf{y}$ and the first term changed sign due to the factoring under $\bar{\mathbf{y}}$. Trough this shrewdness, the maximization oracle can be efficiently computed as in Eq. 6.

5. Experimental settings and results

The proposed algorithm for the identification of leaders in crowds has been tested on three publicly available datasets employed in crowd tracking and group detection tasks: `stu003`, `eth` and `GVEII`. Visual examples of the datasets and the achieved results are shown in Fig. 5. `stu003` and `GVEII` present mildly dense but highly group-structured crowds, characterized by the high variability of groups motion patterns. Conversely, `eth` is a low density crowd scenario where people tend to follow straight paths. In all scenarios most of the group are pairs (65% on average) but triplets and larger groups are present as well. All videos were reduced to 10 fps and the annotation was performed by a set of three independent observers that marked both the leaders and ordered the members by relevance inside each group. Majority of votes has been considered to establish the ground truth data. For testing purposes, we employed the ground truth trajectories and group annotation as the input of our algorithm. The training of the SSVM is performed independently on every video sequence on the first 20% of the groups. In all the experiments the maximum lag K of the time-lagged features of Sec. 3.1 is set to 5 seconds quantized in 6 intervals. The parameters c_d, c_s, c_r of the model based features (Sec. 3.2) have been obtained through grid search on the training data. Leader identification accuracy results are reported and compared in Tab. 1 with the works of Kjærgaard *et al.* [5] and Sanchez-Cortes *et al.* [13] and a SVM baseline, where the leader in a group is the member, among the properly labeled ones, with higher distance from the margin. We trained the SVM with the same PageRank scores used as features in our structured model. Details about leaders ranking in secondary positions are shown in Fig. 3, while Fig. 4 shows the different features importance as the crowd type varies.

6. Discussion

The comparison of the obtained results with other methods in Tab. 1 highlights the ability of the proposed learning framework to deal with the complexity of finding leaders in groups. In particular, sociological models of leadership

(a) `GVEII` (b) `stu003` (c) `eth`

Figure 5: Visual results on the employed datasets. Groups are identified by the color of the shape containing their members and leaders are marked with dots of pertinent colors.

exist, but are always context dependent and should be properly chosen. By learning to combine model-based features, time-lagged cues, individual and group peculiarities we outperform state of the art approaches and define a joint leadership model specifically tailored to the observed scenario, as highlighted by different dominant features in Fig. 4.

The positive results obtained against the SVM approach highlights the importance of the group structure in defining the leader itself. The SVM neglects the contribution of others group elements while, on the opposite, we evaluate the influence of all the members through the adoption of structured learning. We experimentally observed how the leadership is a shared concept among group members and thus all the members influence should be accounted in a unique framework to profitably obtain reliable results. Eventually, the adoption of a learning framework exhibited the capability to adapt to different scenarios where the leadership models may vary according to the people activity, their social identity and their purposes.

References

[1] M. Andersson, J. Gudmundsson, P. Laube, and T. Wolle. Reporting leaders and followers among trajectories of moving point objects. *GeoInformatica*, 12(4):497–528, Oct. 2007. 2

[2] A. Carmi, L. Mihaylova, F. Septier, S. K. Pang, P. Gurfil, and S. Godsill. MCMC-based tracking and identification of leaders in groups. In *International Conference on Computer Vision Workshops*, pages 112–119, Nov. 2011. 2

[3] R. Challenger, C. W. Clegg, and M. Robinson. Understanding Crowd Behaviours. Technical report, June 2009. 1

[4] S. Freud. *Massenpsychologie und Ich-Analyse. Die Zukunft einer Illusion*. Psychoanalytischer Verlag, Vienna, 1921. 1

[5] M. Kjargaard, H. Blunck, M. Wustenberg, K. Gronbask, M. Wirz, D. Roggen, and G. Troster. Time-lag method for detecting following and leadership behavior of pedestrians from mobile sensing data. In *International Conference on Pervasive Computing and Communications*, pages 56–64, Mar. 2013. 2, 3, 5

[6] E. Kometani and T. Sasaki. On the Stability of Traffic Flow. *Journal of Operations Research of Japan*, 2:11–26, 1958. 3

[7] G. Le Bon. *The Crowd: A Study of the Popular Mind*. Macmillan, 1896. 1

[8] J. J. Lee and J. H. Jones. Traffic Dynamics: Visual Angle Car Following Models. *Traffic Engineering & Control*, 8:348–350, 1967. 3

[9] S. Lemercier, A. Jelic, R. Kulpa, J. Hua, J. Fehrenbach, P. Degond, C. Appert-Rolland, S. Donikian, and J. Pettr. Realistic Following Behaviors for Crowd Simulation. *Computer Graphics Forum*, 31(2):489–498, May 2012. 3

[10] T.-Y. Li, Y.-J. Jeng, and S.-I. Chang. Simulating virtual human crowds with a leader-follower model. In *Conference on Computer Animation*, pages 93–102, 2001. 2

[11] L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank Citation Ranking: Bringing Order to the Web. Technical report, 1999. 2

[12] T. Robin, G. Antonini, M. Bierlaire, and J. Cruz. Specification, estimation and validation of a pedestrian walking behavior model. *Transportation Research Part B: Methodological*, 43(1):36–56, Jan. 2009. 2, 3, 4

[13] D. Sanchez-Cortes, O. Aran, M. Mast, and D. Gatica-Perez. A Nonverbal Behavior Approach to Identify Emergent Leaders in Small Groups. *IEEE Transactions on Multimedia*, 14(3):816–832, June 2012. 2, 5

[14] G. Tarde. *The laws of imitation*. Henry Holt and Company, New York, NY, USA, 1903. 1, 2

[15] I. Tsochantaridis, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research*, 6:1453–1484, 2005. 2, 4, 5

[16] J. C. Turner, M. A. Hogg, P. J. Oakes, S. D. Reicher, and M. S. Wetherell. *Rediscovering the social group: A self-categorization theory*. Basil Blackwell, Cambridge, MA, US, 1987. 1

[17] T. Yu, S.-N. Lim, K. Patwardhan, and N. Krahnstoever. Monitoring, recognizing and discovering social networks. In *Conference on Computer Vision and Pattern Recognition*, pages 1462–1469, June 2009. 2