

Social Odometry in Populations of Autonomous Robots

Álvaro Gutiérrez¹, Alexandre Campo²,
Francisco C. Santos², Carlo Pinciroli², and Marco Dorigo²

¹ ETSIT, Universidad Politécnica de Madrid, Madrid, Spain
aguti@etsit.upm.es

² IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium
{acampo, fsantos, cpinciro, mdorigo}@ulb.ac.be

Abstract. The improvement of odometry systems in collective robotics remains an important challenge for several applications. In this work, we propose a localisation strategy in which robots have no access to centralised information. Each robot has an estimate of its own location and an associated confidence level that decreases with distance travelled. Robots use estimates advertised by neighbouring robots to correct their own location estimates at each time-step. This simple *online* social form of odometry is shown to allow a group of robots to both increase the quality of individuals' estimates and efficiently improve their collective performance. Furthermore, social odometry produces a successful self-organised collective pattern.

1 Introduction

Many robotics applications require localisation methods to achieve different tasks. Many different solutions to the localisation problem have been implemented. Among these, odometry is probably the most used as it provides easy and cheap real time position information by the integration of incremental motion information over time. Unfortunately, this integration causes an accumulation of errors during the movement of the robot. Many different approaches have been implemented to deal with systematic and non-systematic localisation errors [1]. Some implementations have used Kalman filters [2]. Thrun and colleagues in [3] create a map of indoor environments combining the idea of posterior estimation with incremental map construction using maximum likelihood estimators.

Some applications in multirobot exploration are implemented without using odometry or dead-reckoning techniques. In [4], a group of robots remains stationary while the other team is in motion. In [5], only one robot is allowed to move while the others act as immobile landmarks. In [6], a chain between two specific areas is created, so that the rest of the group can follow it. In [7], the LOST method enables a team of robots to navigate in an indoor environment. Each robot updates the new optimal path to the goal communicating with a central computer.

All these implementations have a number of different limitations: *i)* they are power consuming in terms of computation because of the Kalman filters and use of maps, *ii)* some robots are not allowed to move while others are tracking distance between them, *iii)* robots must maintain visual contact at all times with the rest of the group, and *iv)* in some cases robots have to communicate with a central device to update or download maps, synchronise movements, or update positions.

The solution we propose in this paper exploits self-organised cooperation in a group of robots to reduce each individual location error. In a nutshell, each robot location knowledge consists of an estimate of its own location and an associated confidence level that decreases with the distance travelled. In order to maximise its confidence about its estimate, each individual tries to update it using the information available in its neighbourhood. Estimated locations, confidence levels and actual locations of the robots co-evolve in parallel in order to guide each robot to the correct objective. This simple *online* social dynamics is shown to allow the population of robots to both reduce individual's errors and efficiently reach a common objective.

In the rest of this paper, we present the details of our solution and analyse its performance in a particular task in which collective correction of odometry errors is beneficial.

2 The Task

To study cooperative odometry, we have devised a task of central place foraging [8] in which robots need to explore the environment to find resource sites and bring items back to a central place. The robots can perceive the central place and the resource site only when they are closer than a threshold distance, given by their sensorial capabilities. Initially, robots are scattered in the arena and they explore the environment to locate both areas. Robots rely on odometry to maintain an estimate of the location of each area (central place and resource site). As soon as a robot comes back to an area, the corresponding location estimate is reset and odometry errors affecting this estimate are discarded.

In an ideal case, robots would make no mistake in estimating the location of the two areas. They could travel endlessly from one place to the other without drifting away. As soon as errors are introduced in the odometry system, estimated locations differ from true locations. A robot may not manage to go back to a given area, and may end up lost. In that case, the robot is doomed to explore again the environment to find the area. To reduce the impact of odometry errors, robots keep in memory both estimated area locations and share this information among them.

3 Mobile Robot Positioning

The accuracy of odometry measurements depends on the kinematics of the robot (see [9]). Let the location of a robot at time $k - 1$ be $L_{k-1} = [x_{k-1} \ y_{k-1} \ \theta_{k-1}]^T$

where (x_{k-1}, y_{k-1}) are the Cartesian coordinates and θ_{k-1} is the orientation with respect to a global reference frame. A rotation $\Delta\theta_k$ and a translation $\Delta\rho_k$ move the robot to a new location L_k :

$$L_k = L_{k-1} + \begin{bmatrix} \Delta\rho_k \cos(\theta_{k-1} + \Delta\theta_k/2) \\ \Delta\rho_k \sin(\theta_{k-1} + \Delta\theta_k/2) \\ \Delta\theta_k \end{bmatrix} \quad (1)$$

Equation 1 is not taking into account problems such as slippage, unequal wheels diameters, wheels misalignments, etc. These errors can be classified as either *systematic* or *non-systematic errors* [10]. The first ones can be modelled and corrected [11], while the last ones can not be corrected and many classical techniques have been implemented to cope with them [12].

When modelling non-systematic errors, each computed robot position is surrounded by a characteristic *error ellipse* that represents a region of uncertainty in which the actual location lies. This region grows with the distance travelled, and it is reset to zero when the robot can localise itself exactly thanks to an environmental landmark (entering one of the two areas in our case). The error ellipse model is based on the covariance matrix of the robot's location defined as:

$$Cov(L_k) = J_{L_{k-1}} f \cdot Cov(L_{k-1}) \cdot J_{L_{k-1}} f + J_{U_k} f \cdot Cov(U_k) \cdot J_{U_k} f \quad (2)$$

where $J_{L_{k-1}} f$ and $J_{U_k} f$ are the Jacobians of f with respect to L_{k-1} and U_k , f is the $(x, y, \theta, \Delta U_r, \Delta U_l)$ vector, U_k is the $(\Delta U_l, \Delta U_r)$ vector and ΔU_l and ΔU_r are the displacements of the left and right wheels respectively. The covariance matrix $Cov(L_0)$ has an initial value of 0.

4 Methods

4.1 Experimental Setup

Our experiments are carried out using a simulation software based on ODE¹. The arena is a bounded square area of 3x3 m² and the robots are randomly scattered in the centre of the arena at the beginning of the experiment. The ground of the arena is white except for the two areas: central place is a black circle of 20 cm radius and the resource site is a grey circle of the same dimensions. Robots differentiate areas using an infrared sensor directed to the ground. A specific simulated range and bearing communication board, based on infrared sensors, allows robots to send messages to each other when their interdistance is less than 25 cm.

We introduce errors to simulate the imperfect response of the range and bearing sensor. Noise is added to the bearing ($\pm 20^\circ$) and range (± 2.5 cm) values. Moreover, each message emitted can be lost with a probability that varies linearly from 1% when the sender-receiver distance is less than 1 cm, to 50% when the two robots are at 25 cm from each other. Errors have also been introduced on the encoder sensors chosen uniformly random in $\pm 20\%$ of the maximum movement at each time step.

¹ <http://www.ode.org>

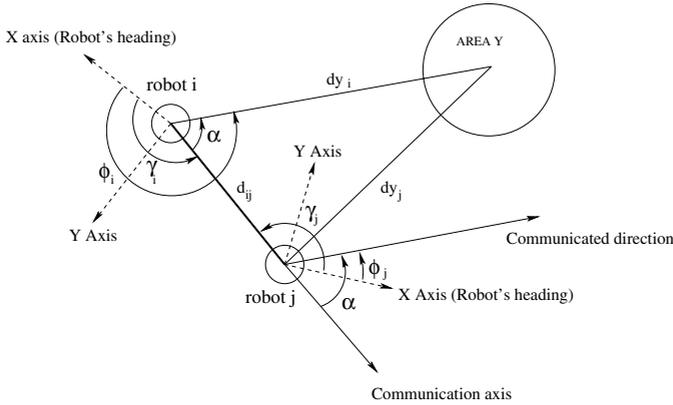


Fig. 1. Robots sharing information about the estimated location of area Y

4.2 Learning from Others

Figure 1 shows how information about the estimated location of area Y is transmitted from robot i to robot j . In a first step, robot i transmits its estimate of the distance dy_i and direction ϕ_i of area Y to robot j . For the direction, the value transmitted is the angle α , obtained from ϕ_i using the communication beam as reference axis: $\alpha = \phi_i - \gamma_i$. In a second step, robot j transforms the received data into its own coordinates system. First, it calculates the direction pointed by robot i as $\phi_j = \gamma_j + \alpha - \pi$. Then, robot j calculates the location $loc_j = (x, y)$ of area Y related to its own reference frame using robot i information and the simple trigonometric equations $x = d_{ij} \cdot \cos\gamma_j + dy_i \cdot \cos\phi_j$ and $y = d_{ij} \cdot \sin\gamma_j + dy_i \cdot \sin\phi_j$.

At this stage, robot j has the opportunity to adopt the estimate of the neighbour, to keep its own or to produce an updated location (loc_up_j) based on both. Given that estimates get worse with distance travelled, the robots use the inverse of the distance travelled as a confidence level of their estimated location. This confidence level, denoted by ϵ_i for robot i , respectively ϵ_j for robot j , is part of any communicated location and informs about the reliability, or quality, of the information. To calculate loc_up_j , we use the so-called pairwise comparison rule [13] for the social learning dynamics, which makes use of the Fermi distribution:

$$c = \frac{1}{1 + e^{-\beta(\epsilon_i - \epsilon_j)}}, \tag{3}$$

where β measures the importance of the relative confidence levels in the decision making. We use a weighted average to merge locations and confidence levels using the Fermi function: $loc_up_j = c \cdot loc_j + (1 - c) \cdot loc_i$ and $\epsilon_up_j = c \cdot \epsilon_j + (1 - c) \cdot \epsilon_i$

5 Results

The performance of the robots in the foraging task under study is measured as the number of total round trips completed from the central place to the resource site and back during one simulated hour.

Figure 2a shows box-plots of the performance with respect to different β values tested. We observe an optimal value $\beta = 10^{-3}$ showing that robots efficiently rely on imitation to increase their collective performance.

Next, we carry out a comparison of different behaviours:

- **no communication:** robots do not communicate and are affected by odometry errors.
- **no odometry error:** robots communicate and are not affected by odometry errors.
- **covariance knowledge:** robots communicate and robots update new location using their own covariance matrix value and the one offered by their neighbours.
- **global communications:** robots communicate with each other globally. Each robot updates its estimates by averaging the knowledge of the whole group.
- **local communications:** robots communicate with each other locally, β is set to its optimal value previously determined, that is, the behaviour detailed in section 4.

Figure 2b reports the outcome of the comparisons. In the **no communication** behaviour, robots rely solely on error prone odometry to find the areas. Once lost, they have to explore the environment and find the areas by chance, which explains the poor performance. In the **no odometry error** case, robots know

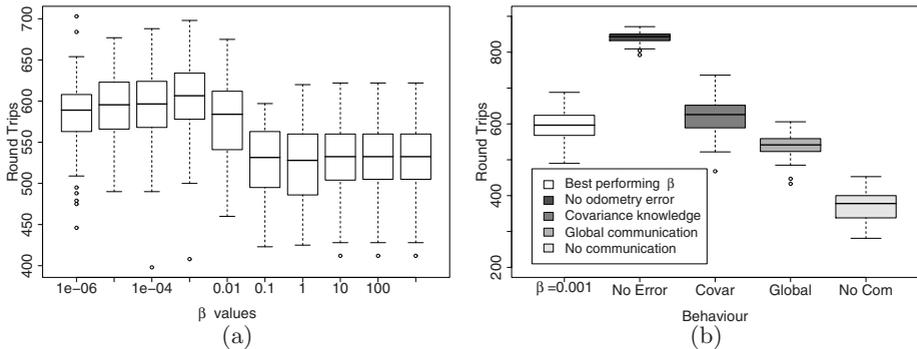


Fig. 2. (a) Task performance using local communications as a function of parameter β . (b) Task performance for the different individual behaviours tested (30 replications for each boxplot). Each box comprises observations ranging from the first to the third quartile. The median is indicated by a horizontal bar, dividing the box into the upper and lower part. The whiskers extend to the farthest data points that are within 1.5 times the interquartile range. Outliers are shown as dots.

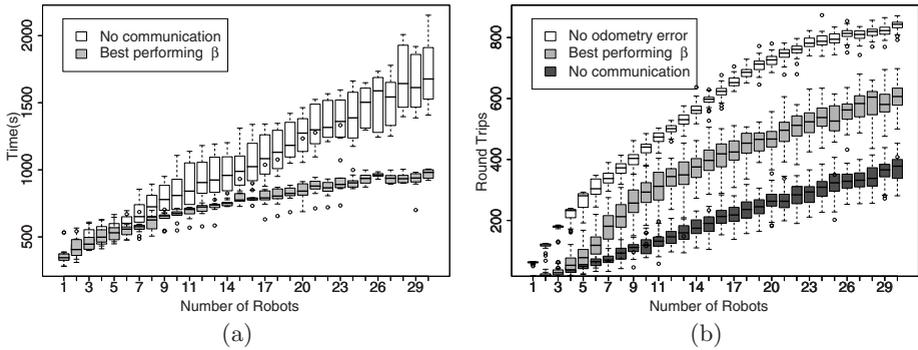


Fig. 3. (a) Time employed for a group of 30 robots to initially localise the central place and the resource site. (b) Task performance as function of group size for three representative cases (30 replications for each boxplot).

accurately the location of both areas. When robots employ **global communication**, they tend to perform worse than with **local communication**, mainly because of the negative influence of lost robots on the knowledge of the group. Interestingly, the **covariance knowledge** behaviour does not show statistically different performance with **local communications** with the best performing β . Not taking into account ideal behaviours that are not implementable in reality, our collective behaviour is therefore exhibiting maximal performance.

We focus now on the comparison between the **no communication** and the **local communication** behaviours. We study the time required to have each robot visit the central place and the resource site at least once. Figure 3a shows these times as a function of the number of robots. With few robots, the two behaviours perform equally well, which is explained by the infrequent encounters of the robots and the consequent low amount of communications. As the number of robots in the experiment increases, we clearly observe that using **local communication** allows the robots to find the areas faster. Using communication, robots are intrinsically carrying out a recruitment process which speeds up the initial exploration phase.

Figure 3b shows the performance of the robots with respect to increasing density, using three different behaviours. The performance of **local communication** is always in between the performances of **no odometry error** and **no communication**. When the number of robots is small (less than 10), the performance of **local communication** is increasing very fast. With more than 10 robots, the performance still increases but at a lower rate. This is due to two counterbalancing effects, namely the improvement of the knowledge of areas' locations through communication versus the disruption of the measure of odometry caused by a higher number of obstacle avoidance events. Results suggest that a high density of robots disrupts performance and there is most likely an optimal density to carry out the foraging task as reported in [14]. We also see that local communication allows the robots to cope with density to some extent and have performance that scales linearly in a wide range of situations.

6 Conclusions

In this paper we have described a social strategy in which robots use pairwise local communication to share knowledge about specific locations to improve their performance in a foraging task. By letting the robots use the estimates of the others, we engineer an efficient and decentralised knowledge sharing mechanism which allows the robots to achieve their goals, both from an individual and group perspective. This simple mechanism drives the system to a successful collective pattern that none of the individuals is able to achieve independently.

We show that local communications are more effective than global communications which would additionally require either more expensive devices or a central system taking care of routing robots' communications. Social odometry does not rely on any internal model provided to each robot, but exclusively on a general collective dynamics of knowledge sharing. This has obvious advantages, given that less effort is needed for environmental dependent parameter tuning.

Preliminary observations suggest that the β parameter depends on the distance between the areas, suggesting that the robots can tune parameter β knowing the distance between the central place and the resource site. The tuning of the parameter can be done off-line, where the designer introduces the value or on-line, where robots update the β parameter once they have located both areas. In the future, we intend to implement and test this strategy on real robots, emphasising the online tuning of the parameter β .

Finally, the performance of the social odometry allows an optimistic forecast concerning the use of online self-organised methodologies in the field of collective robotics.

Acknowledgments. A. Gutiérrez acknowledges support from the Consejo Social of the Universidad Politécnica de Madrid via the Fifth PhD Formation programme. A. Campo, F. C. Santos and M. Dorigo acknowledge support from the F.R.S.-FNRS. C. Pinciroli acknowledges support from *COMP2SYS*, a Marie Curie Early Stage Research Training Site funded by the European Community's Sixth Framework Programme. This work was partially supported by the *SWARMANOID* project, funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission and by the *ANTS* project, an *Action de Recherche Concertée* funded by the Scientific Research Directorate of the French Community of Belgium. The information provided is the sole responsibility of the authors and does not reflect the European Commission's opinion. The European Commission is not responsible for any use that might be made of data appearing in this publication.

References

1. Wang, C.M.: Location estimation and uncertainty analysis for mobile robots. *Autonomous Robot Vehicles* 1(1), 90–95 (1990)
2. Larsen, T., Bak, M., Andersen, N., Ravn, O.: Location estimation for autonomously guided vehicle using an augmented Kalman filter to autocalibrate the odometry. In: *FUSION 1998 Spie Conference*, pp. 33–39. CSREA Press, Las Vegas (1998)

3. Thrun, S., Burgard, W., Fox, D.: A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp. 321–328. Robotics and Automation Society, NJ (2000)
4. Kurazume, R., Hirose, S.: An experimental study of a cooperative positioning system. *Autonomous Robots* 8(1), 43–52 (2000)
5. Grabowski, R., Navarro-Serment, L., Paredis, C., Khosla, P.: Heterogeneous teams of modular robots for mapping and exploration. *Autonomous Robots* 8(2), 293–308 (2000)
6. Nouyan, S., Campo, A., Dorigo, M.: Path formation in a robot swarm: Self-organized strategies to find your way home. *Swarm Intelligence* 2(1), 1–23 (2008)
7. Vaughan, R., Stoy, K., Sukhatme, G., Matarić, M.: LOST: Localization-space trails for robot teams. *IEEE Transactions on Robotics and Automation* 18(5), 796–812 (2002)
8. Balch, T.: Reward and diversity in multirobot foraging. In: *IJCAI-1999 Workshop on Agents Learning About, From and With other Agents*, pp. 15–21. Morgan Kaufman, San Francisco (1997)
9. Klarer, P.: Simple 2-d navigation for wheeled vehicles. Technical report, Sandia Report SAND88-0540, Sandia National Laboratories, Livermore, CA (1988)
10. Feng, L., Borenstein, J., Everett, H.: *Where am I? Sensors and Methods for Autonomous Mobile Robot Positioning*. University of Michigan Press, Ann Arbor (1994)
11. Borenstein, J., Feng, L.: Measurement and correction of systematic odometry errors in mobile robots. *IEEE Transactions on Robotics and Automation* 12, 869–880 (1999)
12. Chong, K., Kleeman, L.: Accurate odometry and error modelling for a mobile robot. In: *IEEE International Conference on Robotics and Automation*, pp. 2783–2788. Robotics and Automation Society, NJ (1997)
13. Santos, F.C., Pacheco, J.M., Lenaerts, T.: Cooperation prevails when individuals adjust their social ties. *PLoS Computational Biology* 2(10), e140 (2006)
14. Goldberg, D., Matarić, M.J.: Interference as a tool for designing and evaluating multi-robot controllers. In: *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pp. 637–642. AAAI Press, Menlo Park (1997)