

The Impact of Informed Individuals on Collective Object Sorting

Andrew Vardy

Department of Computer Science / Faculty of Engineering & Applied Science
Memorial University of Newfoundland, St. John's, Canada, av@mun.ca

Abstract

In this presentation preliminary evidence will be shown to indicate that a minority of informed individuals can influence the collective behaviour of their peers in the distributed object sorting task. From the applications standpoint this suggests that the collective decisions as to where a robot swarm places its clusters of sorted materials can be governed by informing only a minority of robots. This is the first study that integrates *a priori* knowledge for a proportion of robots in this task. [Note: This work has not yet been submitted for publication. A paper describing the *cache consensus* algorithm that is modified for these experiments is under review.]

Introduction

In the distributed object sorting task a group of robots operates on a set of objects, transforming them from an arbitrary configuration to a set of segregated clusters. This is a paradigmatic example of a task with a global outcome achieved through the local interaction of relatively simple agents with limited capacity for computation and communication (we assume no communication). In their seminal paper Deneubourg et al. (1990) proposed a model whereby objects have a probability of being collected or deposited based on local object density. Various groups have implemented variations on this concept in physical robots (Beckers et al., 1994; Martinoli et al., 1999; Melhuish et al., 2001; Verret et al., 2004; Vardy, 2012). However, there has been no research so far on the impact of informed individuals on the sorting process. In the object sorting context, an informed individual is one with *a priori* knowledge as to where objects of a particular type are to be placed. The notion of individuals influencing the global decisions of their group is of great scientific interest, but also has practical implications. For distributed object sorting, informing a subset of individuals as to the desired location of clusters would enable users to influence a swarm that has already been deployed.

Couzin et al. (2005) found that the direction of movement for a group of agents could be dictated by a small proportion of informed individuals. Surprisingly, they also found that the informed proportion could be reduced as the population size increased and still maintain the same influence. We are

interested in whether this phenomenon applies for a group of robots sorting objects.

Cache Consensus

Our sorting algorithm is known as *cache consensus*¹. It requires each agent to maintain a set of cluster positions, known as *cache points*, in memory. Each cache point represents the visible centroid of a previously observed cluster, transformed from the robot's local coordinate frame to the coordinate frame the robot uses to localize. Localization is an extremely well-studied problem and many solutions are available. Examples include satellite-based positioning for outdoor systems (e.g. GPS), overhead tracking systems using cameras, map-based probabilistic localization using on-board sensors, and visual homing to previously visited places (Churchill and Vardy, 2012).

The algorithm works as follows: A robot that is not carrying a puck will wander randomly, examining visible clusters and considering each as a potential target for pick-up. A Deneubourg et al. style probabilistic rule determines whether a cluster is targeted and an individual puck is further selected for collection. Once a puck has been collected, the robot homes directly towards the cache point corresponding to the carried object type and deposits its puck at the cluster growing around that cache point. Meanwhile, the robot compares all visible clusters with remembered cache clusters and will adopt any suitable cluster that is larger than the remembered cache cluster of the same type.

Preliminary Results

To inform an individual robot about the desired position of clusters we simply set the positions of cache points and make them immutable. Preliminary results indicate that a minority of informed individuals is sufficient to guide the majority. Figure 1 shows snapshots from our simulation, which closely models the behaviour of a set of SRV-1 robots with forward facing fish-eye lens cameras and passive grippers (i.e. the robot has to back up to release its puck). The preset cache positions for informed individuals are equally

¹A paper describing this algorithm is under review.

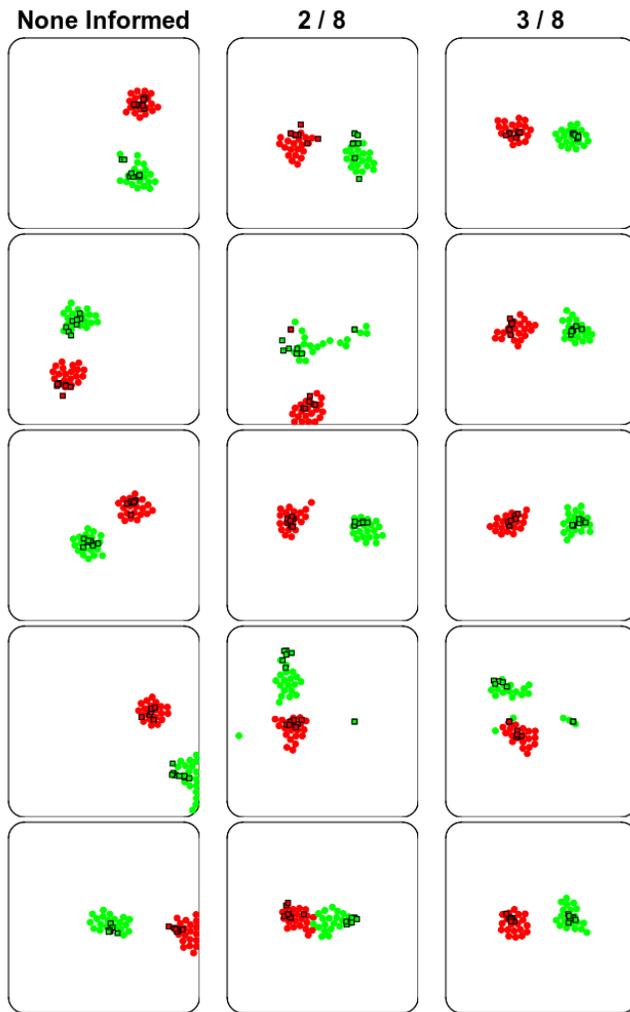


Figure 1: Snapshots of the final state of our simulation using 8 robots sorting objects of two types (red and green). Each column shows the results from five different trials lasting 10,000 time steps (a duration sufficient for convergence). The objects themselves are shown as coloured circles whereas the cache points are shown as squares, coloured to match their corresponding types.

positioned along a horizontal line that bisects the arena. The left-most column shows the result of the sorting process without informed individuals—the clusters form in arbitrary positions. In the right-most column 3 out of 8 individuals are informed and this is sufficient to achieve adherence to the preset cluster positions for four out of five trials. When the number of informed individuals increases beyond 3/8 the

converged clusters lie over their preset positions in all five trials. In the middle column of figure 1 the final clusters lie at the preset positions for 3 out of 5 trials.

Conclusions and Future Work

The results above indicate that a minority of informed individuals can guide the placement of clusters to desired positions. In applications of distributed sorting where direct communication with the entire swarm is difficult, we can still provide control by informing a small proportion of the robots. At the moment the required proportion of informed individuals has not been well-established (although our tests above indicate that it may be in the range $[\frac{3}{8}, \frac{1}{2}]$). Further experiments are clearly required using a greater number of robots. However, care must be taken when interpreting results for increased population sizes since the degree of spatial interference is also dependent on the population size.

References

- Beckers, R., Holland, O., and Deneubourg, J.-L. (1994). From local actions to global tasks: Stigmergy and collective robotics. In Brooks, R. and Maes, P., editors, *Artificial Life IV*, pages 181–189.
- Churchill, D. and Vardy, A. (2012). An orientation invariant visual homing algorithm. *Journal of Intelligent and Robotic Systems*. (accepted 13 July, 2012).
- Couzin, I., Krause, J., Franks, N., and Levin, S. (2005). Effective leadership and decision making in animal groups on the move. *Nature*, 433:513–516.
- Deneubourg, J. L., Goss, S., Franks, N., Sendova-Franks, A., Detrain, C., and Chrétien, L. (1990). The dynamics of collective sorting robot-like ants and ant-like robots. In *First Int. Conf. on the Simulation of Adaptive Behaviour*, pages 356–363.
- Martinoli, A., Ijspeert, A., and Gambardella, L. (1999). A probabilistic model for understanding and comparing collective aggregation mechanisms. In Floreano, D., Nicoud, J.-D., and Mondada, F., editors, *Advances in Artificial Life*, volume 1674 of *Lecture Notes in Computer Science*, pages 575–584. Springer Berlin / Heidelberg.
- Melhuish, C., Wilson, M., and Sendova-Franks, A. B. (2001). Path sorting: Multi-object clustering using minimalist robots. In *Advances in Artificial Life - Proceedings of the 6th European Conference on Artificial Life (ECAL)*.
- Vardy, A. (2012). Accelerated patch sorting by a robotic swarm. In *Canadian Conference on Computer and Robot Vision*. IEEE Xplore.
- Verret, S., Zhang, H., and Meng, M. Q.-H. (2004). Collective sorting with local communication. In *IEEE/RSJ International Conference on Robots and Systems (IROS)*.