Measuring Specialization in Collaborative Swarm Systems

Ling Li, Alcherio Martinoli, Yaser S. Abu-Mostafa

Learning Systems Group — http://www.work.caltech.edu

CORO Research Group — http://www.coro.caltech.edu

The Original Stick Pulling Experiment

Multiple robots search a circular arena and pull sticks out of holes. The length of sticks has been chosen so that collaboration between two robots is needed for the task. The learning issue is to find the optimal gripping time parameter (GTP) for each robot, in order to improve the team performance. Instead of experimenting with real robots, we use a microscopic probabilistic model in simulation.

Figures below show that in addition to the performance improvement with learning, the robots usually become specialized. This is quite interesting since we never explicitly reward diversity in our learning algorithm and there is no explicit communication between robots.

Some previous researches showed with a systematic study that when the number of robots was no more than the number of sticks, there was an advantage in being specialized. When there are more robots than sticks, specialization is not so important; We observe the dropping of diversity in the plot.

The diversity under global reinforcement changes more dramatically, implying a more faithful estimate of specialization.

Motivation and Aims

An initially homogeneous robot team may become heterogeneous with learning. We would like to measure to what degree the teammates become specialized as a function of the task constraints and environmental conditions. This quantitative metric would enable the investigation of issues like the impact of specialization on performance, and conversely, the impact of task constraints on specialization.

Definition of Specialization

Specialization is a mixed concept of both diversity and adaptation. We define it as the part of diversity that is incented by the need of performance improvement.

Diversity Measure

Diversity essentially means difference. We need a difference metric between individuals so that we can group them into clusters.

We use two metrics to measure the diversity, the number of different clusters, or the hierarchic social entropy based on some hierarchy clustering.

Results

The degree of specialization is affected strongly by the environmental conditions and the task constraints but little by the swarm size alone. Local and global reinforcement signals also have different effect on the obtained specialization.

Future Work

Future investigation of the specialization issue may include a more integrated and mathematical definition of specialization, and explicitly applying a priori specialization knowledge in learning to obtain more stable learned solutions.

Generalized Stick Pulling Experiments

We extend the original experiment so that, in order to pull a stick out, k robots are required to collaborate within a certain time window at a certain place.

Sequential Collaboration. We make the sticks longer; One robot can only pull a stick out by $1/k$ of its length.

Parallel Collaboration. We make the sticks heavier and shorter; Exactly k robots are needed simultaneously to lift a stick and pull it out.

We expected that the specialization required in parallel collaboration would be higher than that required in sequential collaboration.

However, the results are counterintuitive. The diversity in parallel collaboration does not surpass that in sequential collaboration until the number of robots is large enough. This reveals a point we did not know before. When there are too many needs for specialization and too few robots, each robot tends to having more roles and finally tends to being homogeneous.

In parallel collaboration, discrepancy between local and global reinforcement signals leads to large gap between diversity measurements under these two signals.

The diversity saturates in sequential collaboration but not in parallel collaboration due to different task constraints. We predict that with more robots, the diversity will also saturate in parallel collaboration.

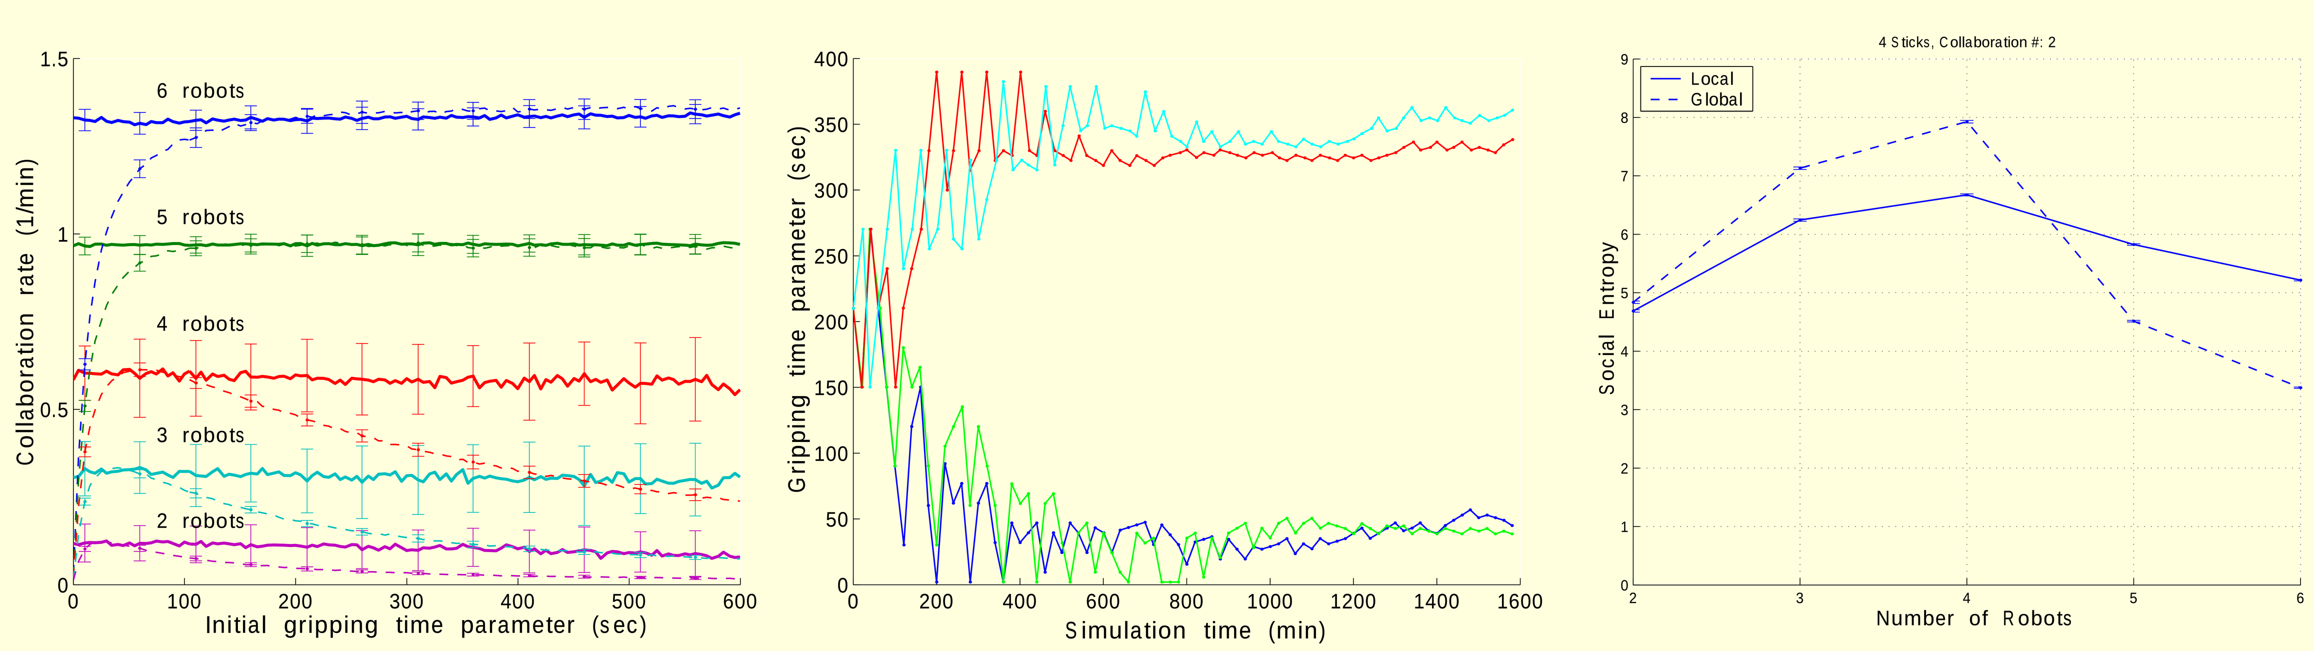

The original stick pulling experiment. Robots were initially given a GTP. With learning, they adjusted their GTPs and achieved higher performance. (Left) Team performance (stick-pulling rates) after learning. Dashed curves are performance without learning. (Middle) 4 robots had 210s as the initial GTP. At the end of the simulation, they formed two clusters, one with large GTPs and the other with small GTPs. (Right) Hierarchic social entropy with different reinforcement signals.

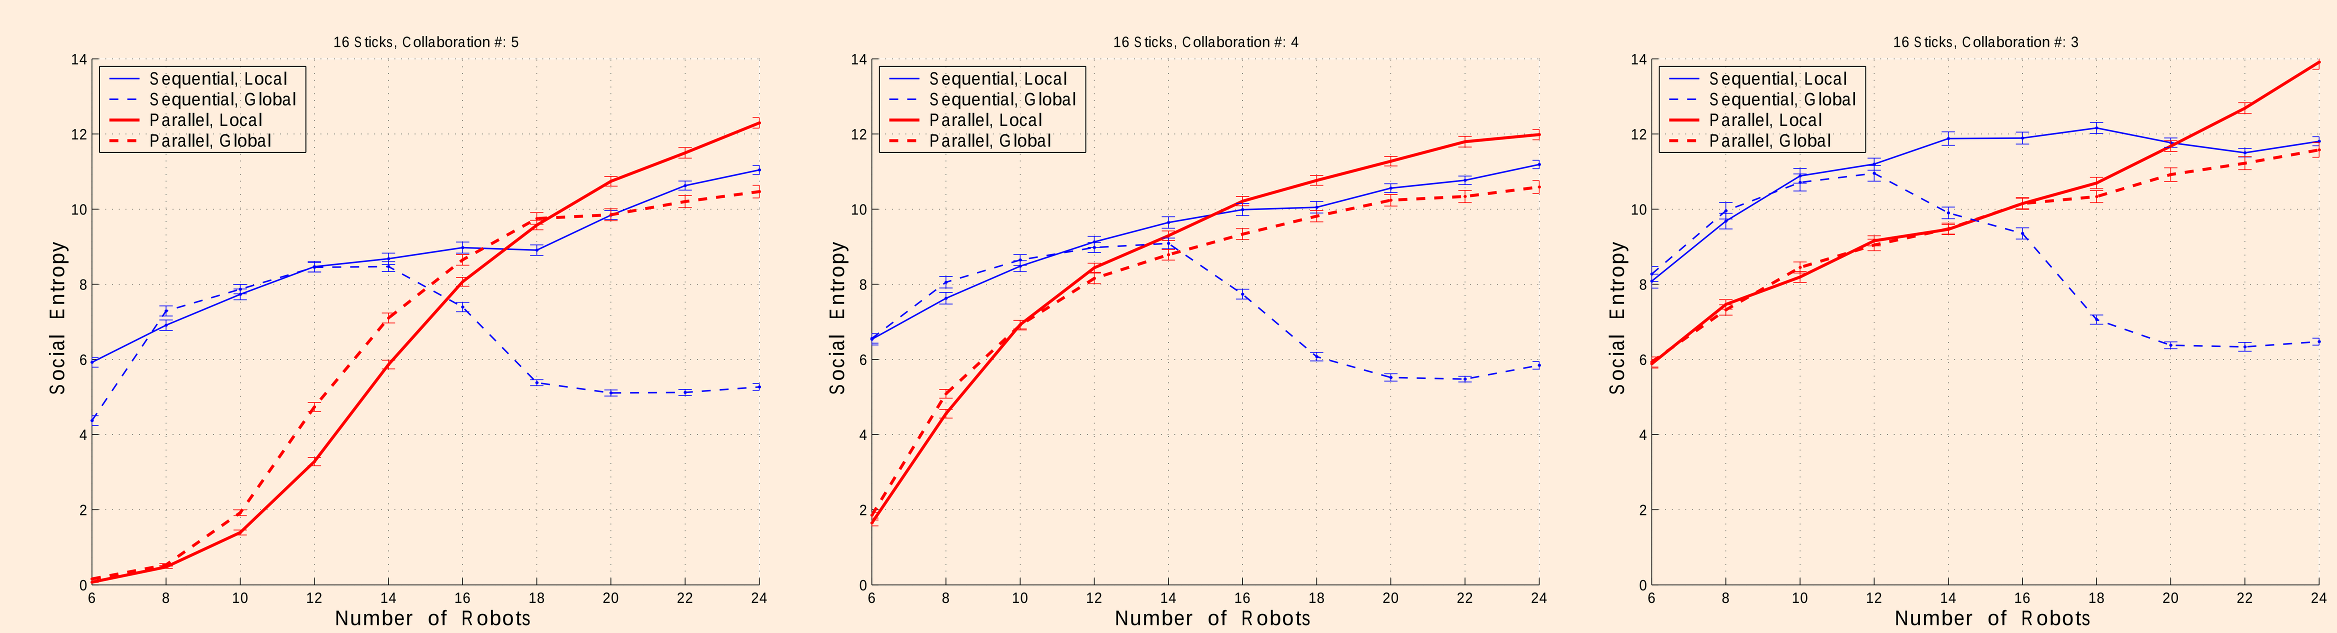

Social entropy as a function of swarm size and reinforcement type. The thin blue curves are social entropy in sequential collaboration, while the thick red curves are that in parallel collaboration. Dashed curves are those using global reinforcement. From left to right, the number of grips needed in collaboration is 5, 4, and 3, respectively. The curves seem continuous from left to right with some common trends.