

A Reading of
Philip R. Cohen and Hector J. Levesque's
Teamwork

"Teamwork" attempts to design a practical analysis (i.e. with implementation as its goal) of what joint action means, and more specifically, how it can be derived from their previous concepts of beliefs, goals and commitment. The authors try to reach a model that supports the robustness that is usually found in joint actions, especially in cases of dynamic conditions and changing beliefs in the individual agents.

In their model, Cohen and Levesque stress the fact that to jointly intend to do something, there needs to be mutual knowledge between the agents about their intentions and commitment to those intentions.

Furthermore, it is important to be able to communicate when the activity is over, whether this is due to success, unachievability or irrelevance of the goal the team is trying to attain.

Their line of thought is very similar to the earlier "On Acting Together" (which I incidentally read first - see: my critique of that paper), and revolves around the idea of a *weak goal of agent x with respect to a team*. This means that once agent x believes that the goal has been achieved, or will never be achieved, he now has a new goal: to create a mutual belief within the team about this new state of affairs.

The only extension over the previous paper is on their analysis of sequential joint action and stepwise sequential action, which I did not find to be radically revealing.

While "Teamwork" is in line with "On Acting Together" in trying to formulate a precise and generative approach to the problem of joint intention, I felt (again) that there is too much of a distance between these formulations and real implementation. Even though their ideas are valid, I cannot see how they extend in any way over folk-conceptions of what joint intention is, and also how in this particular case the rigid definition-theorem approach adds value to their claims.

It would have been interesting to see a more empirical analysis, in particular regarding the speech acts that enhance joint action and planning (these are hinted to in the conclusions of this paper), and how these have been (or can be) put into effect in the design of a collaborative machine. In general, I think it would be interesting to see Cohen and Levesque's theories work in real-world human machine collaboration scenarios.