Guy Hoffman - Fall 2003

A Reading of
Barbara J. Grosz's
Collaborative Systems

"Collaborative Systems" is an overview paper, and accordingly, this has its advantages and its shortcomings. Grosz covers many aspects of collaborative machines, offering a good survey of the various characteristics and requirements of systems that collaborate. On the other hand, since this is an address rather than a structured research paper, I often felt that her arguments drifted from topic to topic instead of supporting a firm thesis.

Grosz's main argument seems to be that collaboration needs to be built into artificially intelligent systems from the beginning, if they are ever to participate in meaningful cooperative activity. This is not clearly supported in he paper, but one can say that her main path to supporting this claim is by showing that collaboration is radically different from other types of joint activity, such as subcontracting, interaction and, yes, coercion. She also builds upon the claim that collaboration is not just the sum of the individuals' activities, but rather something that requires an inherent collaborative ingredient. This ingredient then is broken down Bratman's factors: *mutual responsiveness, commitment to the joint activity, to mutual support and the meshing of subplans*.

Even though many of her examples shed an interesting light on collaboration, I hardly find all of it to be too relevant to the design of collaborative systems. Asking psychological questions about commitment (in particular when contrasting collaboration with interaction and subcontracting) makes sense in a human-human scenario, but is not the right question to raise in the realm of artificial systems. This is partly due to the fact that in any scenario I can think of, machines will still be subordinate to humans, and therefore by definition committed to the human's goals and intents.

The paper makes some good points in support of collaborating machines, but misses a main one: that we are collaborative systems by our very nature, and are therefore trained to collaborate, a human trait which we would want to tap into. In particular I felt that *learning* was missing from Grosz's discourse on collaboration.

An important point to take from this address, though, is that of intent-conflict resolution and error-recovery. Error recovery in a collaborative system must always take into account the common goal, both in figuring out an alternative, and - not less important - in signaling failures and hurdles to the collaborating human, set in the larger context of the collaboration and the joint intention frame.