"Don't think of a Killer Robot" or Framing the public debate/perception of AI

Subbarao Kambhampati
Computer Science, Arizona State University
[Statement at IJCAI 2015 Panel on Who Speaks for AI]

AI is finally getting significant media exposure and generating public debate. It is hard to watch any media these days without coming across some mention of the progress and impact of AI. This is of course good news; it has certainly been a while since the community had this kind of popular reception. With a few exceptions, however, we seem to be mostly reacting rather than framing this coverage and debate.   One worrying consequence of this has been that a significant amount of public debate and coverage is being framed in terms of "stopping evil AI". The whiplash inducing speed with which the popular press went from "AI has been a disappointment" to "AI is going to take over" is a spectacle to behold in and of itself.

Of course, some of this dooms-day preoccupation is caused by the concerns expressed by people outside AI, as well as some over-the-top reporting.   But we in the AI community too seem to be complicit to a certain extent in lending prominence to this preoccupation. A recent example is the extent to which   the coverage of AI surrounding this IJCAI has been about autonomous weapons.

I think it is important that the AI community take a more active role in framing the public discourse on AI. We need to focus on emphasizing the societal benefits of AI rather than dwell merely on the possible safety and security issues. Organizations such as AAAI certainly have a role to play here, by educating the public as well as policy makers about the benefits of AI. As a step in this direction, for IJCAI 2016 next year, we will have "Human-Aware AI" as the theme of the conference. While some variation on stopping the risks of AI may well have rhymed better with the current media preoccupation, we deliberately decided on this more general theme, so as to highlight the many positive interactions AI systems will continue to have in cohabiting with humans be they in the form of   human-robot teams, human-in-the-loop decision support systems, human-machine collaborative scenarios.

The issue of "de-fragmenting" the Artificial Intelligence community is an important one. That the myriad energetic sub-communities of AI are thriving is a cause for celebration no-doubt, but we all believe that the enterprise of AI is bigger than the sum of its sub-communities. We can foster this by highlighting integrated systems, and papers that cross sub-community boundaries, and supporting vigorous tutorial exchange across sub-communities. Flagship conferences such as AAAI and IJCAI certainly have a role to play here, by providing a unified rather than fragmented/federated forums for the field. This is one reason why we are avoiding horizontal sub-tracks at IJCAI 2016.

Coming to the role of professional societies such as AAAI in regards the policy positions, in my view the primary role of a professional society is not so much to take positions on behalf of its membership, but rather to facilitate informed debate among its members, law makers, as well as the public about all sides of an issue of relevance to the field. As the old Fred Friendly saying goes, the mandate is "..not to make up your mind, but to open it. To make the agony of decision making so intense that you can escape only by thinking." This can, and should, be done through policy forums in the professional magazines (such as AI Magazine) and conferences.