Working In Contemporary Teams

Dr Chitra Dey

Department of Management Studies

Anna University

Week – 03

Lecture - 13

States and Outcomes in Human-AI teams

Welcome to lecture 4 of module number 3. In the previous lecture, we had looked at important inputs and processes in human AI teams. We had discussed how teams should be composed and structured so that they can function smoothly. In this lecture, we will be looking at the important emergent states and outcomes in human AI teams. We will complete the lecture by examining the challenges faced by team members when they work with AI and what is the future for such human AI teams. We will be using the input-process-output model to better understand team functioning.

In this lecture, we will be focusing on two emergent states and outputs which are critical for human AI teams. Emergent states develop or evolve as team members interact with each other. Researchers have identified many important emergent states which are critical for team functioning. One of the emergent states with which we all must be familiar is team cohesion, that is the extent to which team members like working with each other.

However, we are going to focus on two emergent states which are critical for human AI functioning and they are team cognition and team trust. Finally, we will be wrapping up by looking at team performance and the workload which comes on team members as a result of working with an AI counterpart. But before we go into that, let us look at two instances where robots find a lot of application in teamwork. Industries like the construction, military, firefighting all have a lot of instances where robots are integrated into the teams. Consider a construction team which is evaluating a building which has been damaged.

This team will consist of construction specialists, it will consist of technicians who will be controlling the robot and the robot itself. The construction specialist will give the instructions on what has to be monitored and the technicians will control the robots and its movements for taking these measures. Another instance where robots can be used

extensively in a dangerous or a dynamic setting is in firefighting. This team will consist of firefighters, again technicians who will control the robots and the robot itself. Robots can in fact be sent ahead of the human team members to go into buildings, survey and make an estimate of the damage and also identify survivors.

 Integrating robots into such kind of teams greatly reduces the danger or the threats faced by the human team members. Having seen this, let us see what are the team dynamics which are important for the functioning of human AI teams. Team cognition and team trust have been identified as two very important factors which will determine whether the human team members in the team actually end up using the AI teammate effectively or not. Team cognition refers to the manner in which knowledge which is important for the team is organized, represented and then distributed among the team members. The team cognitive processes include processes like information sharing and team cognitive states involves team mental models and team situation awareness.

 Team mental models refer to the knowledge which is held together or shared by all the team members. It can be about the task which the team is doing or it could be about the knowledges, skill and competency which are held by team members. For example, if you have an important technical doubt about the work that is being done by the team and you know who to approach in this particular case to clear your doubt, what you are relying on is the team mental model, that is, the knowledge of who knows what in the team. Team situation awareness reflects the extent to which the team members are aware of the dynamic   changes, are aware of the changes in the environment in which they are working and understanding what is the impact of that change on the work that they are doing. It has been found   out in dynamic and dangerous situations that situation awareness is critical for the survival of the team itself.

 So, given the importance of team cognition, let us see what research tells us about team cognition in human AI teams.   The first aspect that we have to consider here is information sharing. It is very important that the team members learn to share relevant information in a timely manner with each other. If this does not happen, it can greatly increase the workload which is taken by the team members. If we go back to the example of the fire fighting squad which has entered the building which is on fire, then we will see that the team members continuously need to be communicating with each other about the information they are getting.

 This will include their locations, imminent threats and location of survivors. When this communication is being relayed with each other, it has to be shared in such a way that the person who is receiving this   information can quickly understand what is the implication for himself or herself. If this does not happen properly, in that kind of a

situation, it can cause a lot of confusion in the mind of the person who is listening to that information. Therefore, the robots which are sent in as team mates should be explicitly programmed to capture the information and convey it in a format which the team members can immediately understand and act upon. The second important aspect when it comes to team cognition are the shared mental models which is the understanding of the task and the team.

Now, if you were to consider a human-human team, a shared mental model evolves as the team members interact with each other. In fact, consider your first day on the project. You wouldn't even know who the other team members are or what they are good at doing. But as the work progresses and as you work with the team members, you start to understand the strengths and weaknesses of your fellow team members and you also understand their preferences in terms of how the work should be done. Now, all of this happens between human-human team members as they work with each other.

But when it comes to the AI counterpart, the shared mental model development happens only to the extent to which the AI has been programmed at the beginning of the team's life cycle. If the AI has a mechanism of deep learning or machine learning, it can be programmed to learn more about the task as the team works on the task. However, it is very difficult to program AI to understand about the feelings and emotions of the team members with whom they are working. Situational awareness is another important aspect in team cognition. It refers to the extent to which the team members are aware of the changes in the environment and its impact on the team's task.

In dangerous and dynamic situations, situational awareness can mean life or death for team members. The AI has to be explicitly programmed to measure critical factors and convey this information to team members in a manner in which they can quickly understand and act upon. So, how can we improve team cognition in human AI teams? Well, first of all, we have to understand the limitations for artificial intelligence. Most of the AI that we have today is narrow AI. That is, it is very good at doing what it has been programmed to do.

Therefore, when the team is designing the AI in its initial phases, it has to take into consideration what all the AI will be expected to do and the team members have to be trained in terms of what are the limitations on what the AI can deliver. Another important aspect which comes in human AI teams is programming the team model versus the task model. If the work of the team can be broken down into various activities and sequenced properly, then the AI can be programmed to take care of the task model. But it is very difficult to program AI to be able to understand the emotions and feelings of team members unlike other human beings. Therefore, one of the challenges which

researchers are grappling with is understanding how AI can be made to understand or detect changes in human emotions and understand human preferences   when working with humans.

 Team cognition in human AI teams can also be improved by improving the interface. Humans are more likely to use an AI if the interface with the AI is easy to use, easy to access and easy to interpret. A very simple touch screen can relay a wealth of information, but if the graphics on the touch screen is not easy to understand or not visible, then it is very likely that the human teammates will stop using the AI. Finally, like I said, AI is being programmed to capture the team model along with the task model to improve team cognition.   The next important emergent state that we need to understand in human AI teams is trust.

  Trust is the ability of team members to become vulnerable to the actions of other team members. It is the situation where one team member relies upon his other teammate. Trust is very crucial because it has been found out that if humans do not trust the AI teammate, then it is very unlikely that they will use the AI teammate. So let us see how we can build trust in teams. Perceived trustworthiness of perceived trustworthiness in AI depends on various factors.

  The first and most important being the characteristics of the agent being trusted. An agent which is tangible is more likely to be trusted than an agent which is not. That is, a robot which has a physical body is more likely to be trusted than a virtual agent which the humans cannot see. Also, if there is transparency in the way the AI processes the information and comes to decisions, then it is more likely that the human team member will trust the AI. Another important factor is the reliability of the AI.

  If the AI is consistent in terms of being correct or giving the correct advice or information to the human members, then it is likely that trust will develop. However, if the AI makes a mistake then it can cause a deep dent in the trust levels which the human team members have on the AI. Autonomy is another important aspect upon which the human will decide whether to allocate work to the agent or not. It has been found that completely autonomous agents are more efficient than semi-autonomous AI. However, if humans have the decision making authority and the ability to allocate tasks to agents, then when they allot tasks to semi-autonomous agents, they are less likely to allot more work to the semi-autonomous agent.

  The next characteristic which is important for trustworthiness or the development of trust is the   characteristic of the agent doing the trusting. The agent should have the human team member should have the willingness to work with AI and this to a great

extent is determined by the previous experience which the team member has had with AI. Perceived trustworthiness is also determined by the characteristics of the work and the characteristics of the environment. It has been found out that when the purpose of the team's task is of an interpersonal nature, humans have less tendency to rely upon the AI. However, if the task is of a highly technical nature, then the humans have a higher tendency to trust and rely upon the AI.

So when it comes to trust in human AI teams, two important things have to be considered. One is the presence of the trust and second is the appropriateness of trust. The first question that we need to ask is, is it possible for the human to trust the AI in the team and vice versa? The second is, what should be the level of trust between the team members? Now, here it is important to understand that trust operates on a continuum. At one extreme end, we have over trust which can cause complacency and neglect on the part of team members. While at the other end, we have under trust where the team members disregard the suggestions which are given by the AI counterpart.

You would have seen examples of airline crews where the human members of the team trusted too much on the AI and neglected certain important warnings which came from other systems. It also has examples where the human team members have disregarded the suggestions given by the AI because of the lack of trust. We complete the discussion on trust in human AI teams by looking at an example where two different robots were integrated into a SWAT team. One robot was called as the Husky A200, a heavier and less maneuverable robot and the MMP 300 which was a much more lighter, smaller and easily maneuverable robot. So, if we are looking at how SWAT team functions, we would think that these teams would prefer the smaller more maneuverable robot over the slower more steadier robot.

However, the research found out that the team members preferred the Husky A200 which was the slower robot. Why? Three important factors determined the preference for the Husky A200 and they are the predictability, reliability and identity. SWAT team preferred to have a robot which moved slowly and therefore required less monitoring and which gave a steady stream of information. They also wanted to identify with the Husky A200 which was a more rugged version of the robot. This example shows us that when we are integrating AI or robots into any kind of team, we have to take into consideration the perceived identity of the team members and also we have to take into consideration realistic assessments or requirements of tasks because it will have a great bearing on how the AI is finally utilized by the human team members.

Next, we look at the contextual factors in human AI teams. Contextual factors are factors which change the level or direction of relationship between characteristics and

outcomes. Like we just saw, team identification is one very important factor which can change the relationship between characteristics and outcomes. The SWAT team identified with a robot which they felt aligned with the identity of the team and therefore they were more likely to use it and they felt like they could rely on that robot more than they could rely on the smaller more maneuverable robot. Therefore, we have to take into consideration the identity of the team and how the robot gels with that identity to be able to predict whether the AI will eventually be used by the team members.

At the organizational level, the training which is given to the team members and the acceptance of AI in teams by important role models will go to a great extent in increasing the acceptance of AI in teams. Next, let us look at three important team outcomes. We look at performance which is usually measured by effectiveness and efficiency. We look at workload and finally at team member satisfaction. These are some of the important outcomes which have been identified in team research for human AI teams.

When it comes to research on performance in human AI teams, we have used a range of measures which have varied from objective measures to attitude measures. And the important and interesting thing here is that results of performance have varied depending on the measure. Team task performance measures for instance are usually measured using the processing of targets and the time measures. That is the time it takes to complete a particular task, the amount of time in which the team members sit idle and so on. Subjective measures of performance on the other hand are the extent of satisfaction which team members feel in working in such a team.

It has been found out that subjective measures of performance were better for human team mates than for autonomous team mates. That is, humans prefer working with other humans more than working with autonomous team mates. Another important insight which emerges from research on performance is that when we give training to the human team members for goal setting and role clarifying in a human AI team, then it improves performance. And it is important to notice that the improved performance is not only in terms of absolute measures but also in terms of the affect. That is, team members were more satisfied in working with those teams when they set goals and had their roles clarified before working, before starting to work in the human AI team.

However, training to treat AI as a team member did not improve absolute measures of performance. However, it had a significant impact on the effect. That is, the rate of satisfaction which team members felt in working with an AI counterpart greatly improved when there was training to treat the AI as a team member. However, it did not increase the absolute measures of performance. Therefore, research on performance tells

us that there is a possibility of a trade-off.

 That is, when we look at subjective and objective measures of performance, subjective measures of performance tell us that the performance of the team is improving as team members see AI as a team mate. However, there is no significant improvement in absolute measures. The next important aspect when it comes to performance in human AI teams is the workload. Workload refers to the amount of time and dedication which the team members have to consciously put in to complete a task. Human team members experienced higher levels of workload when there was higher transparency and low reliability in the AI counterpart.

 That is, when the instructions which were given by the AI counterpart were not reliable, it greatly increased the workload of the human team member. Also, when the AI shared information in a format that was not easy to understand, that also increased the workload of the human team member. The last aspect which can greatly distress the human team members is when there is an AI failure. If we go to the example of autonomous cars, then at what point should the AI give control of the car to the human passenger in the car. Having examined the emerging states and the outcomes which are important for human AI teams, let us look at some of the challenges which researchers and practitioners are facing when considering human AI teams.

 The first challenge is the bias systems.   Human beings have a tendency to anthropomorphize AI. That is, they start considering the AI as another human team mate. If they do not see the AI as another team mate, the likelihood of them trusting and using the AI decreases. One way to get around this challenge is by introducing certain features into the AI which mimics the operations of other humans. Another important challenge which we face in human AI teams is algorithm aversion.

 Algorithm aversion refers to a bias against AI   which can reflect either in a negative behavior or attitude towards the AI on the part of the human team members. Considering that the representation of AI in our general literature, history and even comics for that matter is negative, usually people have a negative attitude towards AI. This can be remedied by giving training to the people in terms of understanding how AI can contribute to teamwork and also understanding what are the limitations of the AI. Another aspect which is challenging in human AI teams is system integration. System integration is the process of bringing together different subsystems and ensure that they function as a whole.

 The problem here is team members in a human AI team consist of human team members, autonomous agents and other factors which help the team function. There

could be a lot of differences between each of these team members in terms of competencies, skills and knowledge and therefore integrating all of this can be a big challenge. Also, a lot of pressure adds up on the human team members when the AI stops functioning in the midst of a task. Therefore, given all this, system integration is a huge challenge for human AI teams. There are a lot of ethical concerns also when it comes to humans and AI working together.

While it is permissible to treat AI in a certain manner, especially when it comes to algorithmic optimization, we cannot treat humans in the same way. Any optimization which is done when there is a human team member involved has to take into consideration the well-being of the human team member first and foremost. Also, humans in human AI teams can consider their AI counterparts as dehumanizing their experiences. One of the industries where AI finds a lot of application is healthcare, particularly in rehabilitation where repeated instructions need to be given to the patients.

AI and robots can find a lot of usage. But the patients sometimes find that the use of a robot is dehumanizing their experience of rehabilitation. So, finding a way around this will be very important for integrating AI into healthcare teams. Another huge challenge which is faced in human AI teams seems to be the trade-off between different measures of performance. The question that arises here is what are we going to prioritize? We saw that subjective measures of performance favored humans over autonomous agents. We also found that training human team members to treat AI as a team member improve the satisfaction level but not absolute measures of performance.

So, there seems to be a trade-off and a lot more research has to be carried out to understand which measures of performance need to be given importance and in what context. Sometimes, it might be more important for the AI to gain trust and acceptance than have performance in terms of absolute measures. Research has to identify what those contexts are and we have to give priority to those measures of performance. The last challenge which we have to deal with in human AI teams is dealing with workers' fears in working with AI.

There are two important aspects to it. First one is, humans fear that AI and robots will replace them in work. Another one is, humans will doubt whether they have the ability to keep the pace, keep up with the pace of change that is happening in the AI. We have to constantly train and educate the team members about the limitations and capabilities of the AI so that the human team members understand that AI is going to complement in their work and is actually going to help them work better. So, what does the future hold for human AI teams? Well, research and practitioners are focusing on design teams that are better equipped to work with AI. This includes educating human team members

about the limitations of AI and what they can reasonably expect from an AI counterpart.

It is also important to build trust between human team members and autonomous agents by giving them chance to work in simulated environments before sending them out to work on the field. One of the challenges which are faced by human AI team researchers and practitioners is designing performance evaluations for human autonomous teams. The future for human AI teams also includes the design of interfaces and models that support social dynamics as much as task work. Two ways in which this can be done, one is to imbibe human characteristics in the AI design.

I had given the example of backchannelling. Backchannelling happens when one person who is listening to another indicates his or her commitment to the conversation by giving verbal and non-verbal feedback. When the listener backchannels, the person who is talking feels that they are paying more attention to what is being said. Imbibing elements of backchannelling in the AI or the robot can go a long way in making the human team member feel that the AI is invested in the conversation that they are having. Another way in which we can improve interface is by designing a social artificial intelligence layer over the task layer which is programmed in the AI. To summarize, in this lecture, we have looked at two important emergent states, team cognition and team trust.

Then we have looked at the contextual factors which are important for team functioning. Finally, we have looked at what are the important team outcomes for human AI teams. Then we have seen what are the challenges which people face when working in human AI teams and how these challenges can be met so that we have a good future for human AI teams. I leave you with some interesting topics that you may want to look up. One important area which is challenging both researchers and practitioners is how to develop performance evaluation systems in human AI teams.

Another challenge is developing shared mental models. The last challenge which is perplexing researchers is how to identify with human AI teams. With this, I say bye to you till we meet again.