Collectively Intelligent Teams: Integrating Team Cognition, Collective Intelligence, and AI for Future Teaming

Lorenzo Barberis Canonico
Clemson University, SC
lorenzb@g.clemson.edu

Christopher Flathmann
Clemson University, SC
cflathm@g.clemson.edu

Dr. Nathan McNeese
Clemson University, SC
mcneese@clemson.edu

In this paper we propose a new model for teamwork that integrates team cognition, collective intelligence, and artificial intelligence. We do this by first characterizing what sets team cognition and collectively intelligence apart, and then reviewing the literature on "superforecasting" and the ability for effectively coordinated teams to outperform predictions by large groups. Lastly, we delve into the ways in which teamwork can be enhanced by artificial intelligence through our model, finally highlighting the many areas of research worth exploring through interdisciplinary efforts.

INTRODUCTION

Cognitive science is not just a study about the human mind: it's also a study in information processing and neurological behavior. Through those lenses, it becomes possible not only to examine individual behavior, but also to investigate social behavior. Team science lies at the intersection between researching individual and group psychology, and actively deals with how collections of individual people coordinate most effectively.

This kind of phenomenon however is replicated at a larger scale with intellectually productive crowd efforts. A primary example is Foldit, a crowdsourcing effort for biochemistry and protein folding that uncovered in just three weeks the structure of an enzyme related to AIDS that had eluded scientists for 15 years (Malone, 2018). This type of phenomenon taps into resources and skills needed to perform an activity that are distributed widely or reside in places that are not known in advance (Malone et al., 2009).

So far, the team cognition literature has mostly focused on the ability for teams to develop shared situational awareness or a shared understanding of a complex task in many practical contexts (Cooke et al., 2003). On the other hand, the collective intelligence literature has mostly focused on the ability of large decentralized groups to aggregate and process large amounts of information to produce insights far superior to those of individuals and teams, especially in scenarios marked by deep uncertainty (Atanasov et al., 2016).

In this paper, we analyze the results of "superforecasting" (teams of forecasters outperforming prediction markets) as the starting point for a new way to think about teamwork that merges the insights from team cognition and collective intelligence to challenge many assumptions held by each school of thought. Thus, one of the contributions of this paper is a comparative analysis of the team cognition and collective intelligence literature to identify the connections between the concept of a shared mental model and that of the "wisdom of the crowd". Furthermore, we analyze the results from the superforecasting research to develop a model that enables the coordination individuals whose information processing enables both team cognition as well as collective intelligence to emerge. Lastly, we demonstrate how to incorporate artificial intelligence into our model in order to enable team's information processing to become the input of a neural network to enhance the team's collectively intelligence to outperform smart crowds.

REVIEW

Team Cognition

The concept of a team naturally lies at the core of team cognition. As opposed to groups or aggregations of individuals, teams are composed of members with specific roles and responsibilities related to a complex task. One of the best ways to understand teamwork is to think of teams as a form of highly specialized groups who interact dynamically, interdependently, and adaptively toward a common and valued goal/objective/mission (McNeese et al., 2014; Salas et al., 1992). At the core of team effectiveness is communication. By acquiring, sharing and processing information, teams effectively communicate and coordinate towards a common objective (Hinsz et al., 1997). From this process emerges team cognition, a larger form of cognition that is very akin to that of individuals and yet happens at a larger scale.

Team cognition is thus the glue that can hold a team together by improving communication, coordination, and awareness of the associated teammates. There are two main perspectives that conceptualize team cognition: the ecological interaction approach and the shared knowledge approach (McNeese et al., 2016). The ecological interaction approach states that team cognition is a form of team interaction involved in communication and coordination (Cooke et al., 2013). The shared knowledge approach is defined by an input-process-output paradigm where the input is individual team members' knowledge, the process is the sharing of such knowledge, and the output is shared cognition typically represented in the form of a shared mental model (Mohammed et al., 2010). Research under both paradigms identifies team cognition as a major predictor of team performance, highlighting how breakdowns in team cognition result in performance losses and how improvements in team cognition result in performance gains (Wilson et al., 2007; Cooke et al., 2003, 2013).

A critical aspect to emphasize is how team cognition is not simply the sum of the individual cognitions of the team members, but is rather the team-level cognition emerging from the interactions of the team members (Cooke et al., 2007). In its most simple form, team cognition is a process (communication and coordination) that produces a shared understanding of both teamwork and taskwork (Mohammed et al., 2010; Fiore & Salas, 2004). Under the shared knowledge perspective, team cognition results in the production of shared mental models.

Mental models are the mechanisms that enable humans to describe and explain a system's purpose, form, function, and predictable future states (Rouse & Morris, 1986). In the context of teams, shared mental models are the emergent property of team cognition as the team members developed a shared understanding and mental representations of knowledge about the team's environment (Cannon-Bowers et al., 1990; McNeese et al., 2016).

Prior research highlights how compatible and shared mental models lie at the foundation of the ability for experienced teams to coordinate, anticipate, predict, and adapt to both tasks as well as to each other's needs (Fiore et al.,2010). To put it in another way, the sharing of mental models among teammates enables each team member to describe, explain, and predict future events at the team level (Graham et al., 2004; Mathieu, 2000).

However, despite the heightened awareness, knowledge and understanding that comes with shared mental models, teams often still fall prey to the biases that affect individuals. A major example is that groups can be primed to over-emphasize solutions from one problem to subsequent ones. Priming in groups can inhibit creativity to solve complicated problems and cause groups to resemble individuals in terms of mental set or habitual routine (Hinsz et al., 1997). Furthermore, large teams are often inefficient at storing information. It is estimated that groups use only about 70% of their storage capacity because of the losses incurred from the collaboration required to remember at the group level (Hinsz et al., 1997). Beyond that, in scenarios marked by deep uncertainty, where probabilistic thinking is key to navigating the situation, teams tend to exaggerate the individual tendency to neglect base-rate information when making a judgement (Hinsz et al., 1997). This research suggests that teams can exaggerate the biases, errors, or tendencies of information processing that occur among individuals.

In essence, teams not only display magnified cognitive capacity but they also display unique cognitive abilities that emerge from the interaction between the team members. Those abilities however are not unlimited, and are often constrained by the very biases that plague individual decision making. However, much in the same way team cognition emerges from individuals to produce intelligence and behaviors beyond the individual, a different type of intelligence emerges from large groups and crowds that cannot be reduced to behaviors of the individual -- collective intelligence.

Collective Intelligence

The best way to understand collective intelligence it to investigate prediction markets: platforms were a decentralized users can place bets on outcomes into the future (Surowiecki, 2005; Malone, 2009). The ability of prediction markets to outperform experts and pools alike is a primary display of "the wisdom of the crowd" -- the kind of collective intelligence emerging from groups of decentralized information seekers.

DARPA's experience with prediction markets serves as a perfect illustration. The DAGGRE geopolitical forecasting market was a combinatorial prediction market sponsored by IARPA that recruited participants to place bets on the geopolitical events. Over the 20 months the DAGGRE market was active, more than 3000 forecasters participated, with an average of about 150 forecasts per week. The overall market accuracy, as reflected the prices associated with each event, was about 38% greater than the baseline system at over 400 geopolitical questions (Laskey et al., 2015). These results extend beyond geopolitics. For example, when compared to concurrent major opinion polls on U.S. presidential elections, the Iowa Electronic Market forecasts were more accurate 451 out of 596 times (Hanson, 2005).

Markets by their very nature provide strong economic incentives for individuals to correct systematic biases, such as overconfidence or underconfidence. A rational trader would place bets that are profitable in expectation, realigning prices with historical base rates (Atanasov et al., 2016). Furthermore, they are able to handle more complexity than an individual or centralized body could grasp because "knowledge that is implicit, dispersed, and inaccessible by traditional, conscious methods can be organized through markets to create more rational calculation than can elite experts" (Watkins, 2007).

Prior research has identified four conditions that enable the emergence of collective intelligence in a crowd (Surowiecki, 2005):

1. Diversity of opinion: each person should have some private information, even if it's just an eccentric interpretation of the known facts
2. Independence: people's opinions are not determined by the opinions of those around them
3. Decentralization: people are able to specialize and draw on local knowledge
4. Aggregation: some mechanism exists for turning private judgments into a collective decision

Such valuable, productive and intelligent behavior emerges from decentralized groups of people that explore and aggregate local information into collectively useful knowledge (Surowiecki, 2005). Even though the decentralized nature of collective intelligence stands in sharp contrast to the centralized and interdependent nature of team cognition, the two phenomena can be understood as manifestations of the same emergent properties. As long as the four properties are maintained, collective intelligence emerges to overcome many of the systemic biases that plague teams.

However, collective intelligence comes with its own set of challenges and constraints. Prior research on large datasets from prediction market transaction prices display long-shot bias: high-likelihood events are underpriced, and low-likelihood events are overpriced (Page & Clemen, 2012). Beyond prediction markets, crowds often display a narrowed focus of attention, redundant memories, accentuation of processing strategies, and shared distribution of information (Hinsz et al., 1997). This phenomenon is often attributed to the reaction of individuals exposed to majority influence as they abandon divergent thinking (i.e., with more thoughts covering a wider range of categories or perspectives) for convergent thinking (i.e., with a narrow range of focus and less cognitive effort) (Hinsz et al.,1997).

In essence, once an individual stops being part of a team or small group and becomes part of a crowd, a unique set of social pressures can erode the basis for collective intelligence.

Superforecasting

Recent results from IARPA's geopolitical forecasting tournaments set the foundation for a way to possibly rethink team cognition and collective intelligence. Specifically, results from prediction polls enabled researchers to identify the characteristics of optimal forecasting as well as a model to create teams of forecasters (Atanasov et al., 2016).

There are three important distinctions between prediction polls and other polls or surveys (Atanasov et al., 2016). .

1. In prediction polls, participants are asked for probabilistic forecasts, rather than preferences or voting intentions. Forecasts are elicited in a dynamic context. Forecasters update their predictions whenever they wish, and feedback is provided when events are resolved.
2. Forecasters compete against other forecasters -- competitive features encourage better search processes and more accurate inferences.
3. Prediction polls rely on crowds with dozens, hundreds, or thousands of individuals who may be knowledgeable but are not necessarily subject matter experts, which distinguishes polls from expert elicitation techniques

In IARPA's geopolitical prediction tournament, more than 2,400 participants made forecasts on 261 events. Forecasters were randomly assigned to either prediction markets (continuous double auction markets where participants bet on outcomes), or prediction polls in which they submitted probability judgments, independently or in teams, and were ranked based on Brier scores. In both seasons of the tournament, team prediction polls outperformed prediction markets when forecasts were statistically aggregated using temporal decay, differential weighting based on past performance, and recalibration. Atanasov et al. (2016) conclude that "team prediction polls created a mix of intrateam cooperation and inter-team competition. This mixed environment may be better than a purely competitive one if individuals share information with each other and pool their knowledge".

These results suggest that prediction polls with proper scoring feedback, collaboration features, and statistical aggregation can distill the wisdom of crowds in teams as opposed to large groups.

NEW TEAMWORK MODEL

On the surface, team cognition and collective intelligence appear to be irreconcilable. Teams cognition is heavily dependent on information sharing and verbal communication, whereas collective intelligence relies on the opposite: local information gathering and independent analysis (Graham et al., 2004) . Team cognition benefits from centralized leadership, whereas collective intelligence is predicated upon decentralization (Fiore et al., 2010). Teams exaggerate the cognitive biases of individuals, whereas wise crowds compensate for them (Atanasov et al., 2016).

However, by delving deeply into the phenomenon of "superforecasting", a new type of team cognition can be identified as the properties of the wisdom of the crowd can emerge from small teams operating in the right structure. In turn, this insight sets the foundation for methods to combine artificial intelligence with collective intelligence. We thus develop a new model (Figure 1) for teamwork that unifies team cognition, collective intelligence, and artificial intelligence.

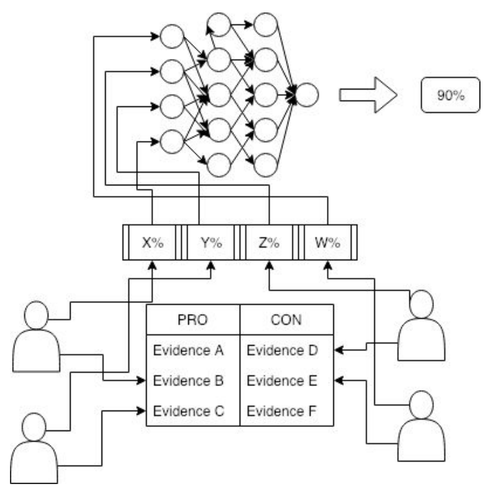

Figure 1. Each teammate shares evidence for a team discussion, and expresses a probability estimate that is then aggregated by the neural network

Collectively Intelligent Teams

A collectively intelligent team is a team whose shared mental model enables higher degrees of collective intelligence thloyed at the individual level. Specifically, each individual is trained on specific cognitive strategies to improve their decision-making. Just like superforecasters, collectively intelligent teammates are trained in probabilistic thinking, evidence-based discussions, depersonalized intellectual conflict, diversity of viewpoints, commitment to truth, and data-driven decision-making. All of these practices can dramatically enhance the teammate's team cognition.

The second major component is a set of guidelines to the team for information processing. Elite superforecasters in the literature are shown to engage in many distinct behaviors that were predictors of accuracy: they cognitively triaged (deciding how to allocate effort across questions). They asked five times more specific questions than average and were answered six times more frequently. They made comments roughly one third longer, and made between nine and thirteen times more general comments than average. Beyond merely being more likely to gather news and opinion pieces related to the forecast, superforecasters were between six and ten times more likely to share news links with their teammates) (Mellers et al., 2015).

The last component is a layer of artificial intelligence to aggregate the team's insights. An often overlooked result of the superforecasting research is how applying an extremizing algorithm that aggregates the forecasters' forecasts and weighs them based on track record and diverse point of view outperformed 99% of the individual super-forecasters (Tetlock & Gardner, 2016). Already prediction-polled superforecasters outperform prediction markets, the best collectively intelligent mechanism to date, so the fact that shifting from a statistical aggregation rule to an algorithm leads to even better predictive power is remarkable. Prior research in both prediction markets and prediction polling has identified the choice of an aggregation function materially changes the value of the probability estimates of the system (Atanasov et al., 2013; Atanasov et al., 2016).

The advances brought about a simple algorithm open up the opportunity to incorporate AI in the forecasting process. Prior research has shown how collectively intelligent teams can be created through cognitively optimized software applications. Specifically, a recent radiology study at Stanford has shown that a team radiologist coordinated with a probabilistic interface gained a 22% margin on the state-of-the-art deep-learning solution and a 33% margin on individual radiologists as a whole (Rosenberg et al., 2018). This result is not isolated to healthcare but also works in financial forecasting. A recent Oxford study on using the same probabilistic interface to coordinate traders to predict four economic indicators showed a 26% increase in prediction accuracy as compared to individual forecasts (Rosenberg et al., 2018).

However these results occur with the use of simple algorithms that aggregate the individual forecasts to produce an optimal estimate. Neural networks can process much deeper correlations between individual forecasts over time, therefore machine learning can yield to significant gains for collectively intelligent teams that are orders of magnitude better than base performance.

In essence, our model for collectively intelligent teamwork relies upon each teammate being trained by probabilistic cognitive strategies, the team sharing information through a protocols that preserve independence, and most importantly the use of neural networks to optimize the collective intelligence of the team.

DISCUSSION

The superforecasting research could not be understated: those results robustly demonstrate how small teams can outperform the most collectively intelligent mechanism known to date -- prediction markets -- in a significant way. This evidence suggests that a new type of intelligence can be unlocked through teamwork that goes beyond what the team cognition literature has been demonstrating. Team cognition expands from perception into prediction as a team of superforecasters displays both team cognition and collective intelligence, two emergent phenomena thought to be separate up until that point.

Researching superforecasters expands how we think about collectively intelligence as well as team cognition. For the former, it means that you no longer necessarily need a crowd to have collective intelligence as long as teams are trained through particular cognitive strategies when engaging in forecasting. For team cognition, it means that the possibility landscape for teamwork is vaster than previously thought, and is relevant beyond the military setting and responding to physical situations and can move into a higher-order form of information processing like forecasting.

Furthermore, our model bridges the literature gap between team cognition and collective intelligence. Specifically, we note that shared mental models are to team cognition, what the wisdom of the crowd is to collective intelligence: they are emergent phenomena of effectively coordinated groups of people. Through the collectively intelligent team model, these two perspectives can be integrated. Not only can teams rapidly developed a sophisticated understanding of their situation and environment, but can also work together to produce remarkable insights under deep uncertainty by thinking about the future in a probabilistic manner.

The collectively intelligent team model is also useful as the foundation for new technology. A collectively intelligent team emerges when a team's shared mental model enables each teammate to enhance their forecasting process so that they can each produce better estimates that can be fed as input to a neural network that learns over time how to calibrate the weight it assigns to each team member's opinion in particular contexts. The use of AI in this case creates a highly adaptive coordination mechanism that enables the team to retain independent thought while still collaborating. Rosenberg et al's (2018) work on coordinating teams through probabilistic decision making is merely an initial attempt at developing human-centered interfaces that can transform team cognition into collective intelligence.

Interestingly enough, the literature on team cognition provides many results that are consistent with collectively intelligent teamwork. For example, Google's Project Aristotle illustrates how the stronger predictors of team success in knowledge-work shows are equality of conversational turn-taking and higher levels of social sensitivity (Duhigg, 2016). This is consistent with the observations that superforecasters deeply value hearing everyone else's opinion before expressing their own (Tetlock & Gardner, 2016). Furthermore, prior research has shown that asking team members to defend their position induces a cognitive strategy relying on generative confirmatory evidence, whereas asking not expecting them to leads for greater freedom to explore arguments counter to initial positions (Hinsz et al., 1997). Teams of superforecasters do this by keeping track of arguments from different viewpoints in favor or against a particular forecast (Tetlock & Gardner, 2016).

CONCLUSION

Through a deep analysis of team cognition and collective intelligence, we developed a new model that integrates the two phenomena. This model sets the basis for a new type of teamwork that exhibits both properties, as evidenced by the results on teams of superforecasters. Lastly however, we go one step further to show how artificial intelligence can be used to cognitively enhance collectively intelligent teams to reach a new level of intelligence that has only begun to be fully explored.

REFERENCES

Atanasov, Pavel, et al. "The marketcast method for aggregating prediction market forecasts." International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction. Springer, Berlin, Heidelberg, 2013.

Atanasov, P., Rescober, P., Stone, E., Swift, S. A., Servan-Schreiber, E., Tetlock, P., ... & Mellers, B. (2016). Distilling the wisdom of crowds: Prediction markets vs. prediction polls. Management science, 63(3), 691-706.

Cannon-Bowers, J. A., Salas, E., & Converse, S. A. (1990). Cognitive psychology and team training: Training shared mental models and complex systems. Human factors society bulletin, 33(12), 1-4.

Cooke, N. J., Gorman, J. C., Winner, J. L., & Durso, F. T. (2007). Team cognition. Handbook of applied cognition, 2, 239-268.

Cooke, N. J., Gorman, J. C., Myers, C. W., & Duran, J. L. (2012). Theoretical underpinnings of interactive team cognition. *Theories of team cognition: Cross-disciplinary perspectives*, 187-207.

Duhigg, C. (2016). What Google learned from its quest to build the perfect team. The New York Times Magazine, 26, 2016.

Fiore, S. M., Rosen, M. A., Smith-Jentsch, K. A., Salas, E., Letsky, M., & Warner, N. (2010). Toward an understanding of macrocognition in teams: predicting processes in complex collaborative contexts. Human Factors, 52(2), 203-224.

Graham, J., Schneider, M., Bauer, A., Bessiere, K., & Gonzalez, C. (2004, September). Shared mental models in military command and control organizations: Effect of social network distance. In Proceedings of the Human Factors and Ergonomics Society Annual Meeting (Vol. 48, No. 3, pp. 509-512). Sage CA: Los Angeles, CA: SAGE Publications.

Hinsz, V. B., Tindale, R. S., & Vollrath, D. A. (1997). The emerging conceptualization of groups as information processors. Psychological bulletin, 121(1), 43.

Hollenbeck, J. R., Ilgen, D. R., LePine, J. A., Colquitt, J. A., & Hedlund, J. (1998). Extending the multilevel theory of team decision making: Effects of feedback and experience in hierarchical teams. Academy of Management Journal, 41(3), 269-282.

Laskey, K. B., Hanson, R., & Twardy, C. (2015, July). Combinatorial prediction markets for fusing information from distributed experts and models. In 2015 18th International Conference on Information Fusion (Fusion) (pp. 1892-1898). IEEE.

Letsky, M., Warner, N., Fiore, S. M., Rosen, M., & Salas, E. (2007). Macrocognition in complex team problem solving. OFFICE OF NAVAL RESEARCH ARLINGTON VA.

Malone, T. W. (2018). Superminds: The Surprising Power of People and Computers Thinking Together. Little, Brown.

Malone, T. W., Laubacher, R., & Dellarocas, C. (2009). Harnessing crowds: Mapping the genome of collective intelligence.

Mathieu, J. E., Heffner, T. S., Goodwin, G. F., Salas, E., & Cannon-Bowers, J. A. (2000). The influence of shared mental models on team process and performance. Journal of applied psychology, 85(2), 273.

McNeese, N. J., Reddy, M. C., & Friedenberg, E. M. (2014, September). Towards a team mental model of collaborative information seeking during team decision-making. In Proceedings of the Human Factors and Ergonomics Society Annual Meeting (Vol. 58, No. 1, pp. 335-339). Sage CA: Los Angeles, CA: SAGE Publications.

McNeese, N. J., & Cooke, N. J. (2016, July). Team Cognition as a Mechanism for Developing Collaborative and Proactive Decision Support in Remotely Piloted Aircraft Systems. In International Conference on Augmented Cognition (pp. 198-209). Springer, Cham.

Mellers, B., Stone, E., Murray, T., Minster, A., Rohrbaugh, N., Bishop, M., ... & Ungar, L. (2015). Identifying and cultivating superforecasters as a method of improving probabilistic predictions. Perspectives on Psychological Science, 10(3), 267-281.

Orlitzky, M., & Hirokawa, R. Y. (2001). To err is human, to correct for it divine: A meta-analysis of research testing the functional theory of group decision-making effectiveness. Small Group Research, 32(3), 313-341.

Page, L., & Clemen, R. T. (2012). Do Prediction Markets Produce Well‐Calibrated Probability Forecasts?. The Economic Journal, 123(568), 491-513.

Rosenberg, L., Lungren, M., Halabi, S., Willcox, G., Baltaxe, D., & Lyons, M. (2018, November). Artificial swarm intelligence employed to amplify diagnostic accuracy in radiology. In 2018 IEEE 9th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)(pp. 1186-1191). IEEE.

Rouse, W. B., & Morris, N. M. (1986). On looking into the black box: Prospects and limits in the search for mental models. Psychological bulletin, 100(3), 349.

Salas, E., Dickinson, T. L., Converse, S. A., & Tannenbaum, S. I. (1992). Toward an understanding of team performance and training.

Salas, E., Nichols, D. R., & Driskell, J. E. (2007). Testing three team training strategies in intact teams: A meta-analysis. Small Group Research, 38(4), 471-488.

Salas, E., DiazGranados, D., Klein, C., Burke, C. S., Stagl, K. C., Goodwin, G. F., & Halpin, S. M. (2008). Does team training improve team performance? A meta-analysis. Human factors, 50(6), 903-933.

Surowiecki, J. (2005). The wisdom of crowds. Anchor.

Tetlock, P. E., & Gardner, D. (2016). Superforecasting: The art and science of prediction. Random House.

Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In Theories of group behavior (pp. 185-208). Springer, New York, NY.