Directed Epistemic Network Analysis

Ariel Fogel[1(✉)], Zachari Swiecki[2], Cody Marquart[1], Zhiqiang Cai[1], Yeyu Wang[1], Jais Brohinsky[1], Amanda Siebert-Evenstone[1], Brendan Eagan[1], A. R. Ruis[1], and David Williamson Shaffer[1]

[1] University of Wisconsin–Madison, Madison, WI, USA
ajfogel@wisc.edu
[2] Monash University, Melbourne, VIC, Australia

Abstract. Quantitative ethnographers across a range of domains study complex collaborative thinking (CCT): the processes by which members of a group or team develop shared understanding by making cognitive connections from the statements and actions of the group. CCT is difficult to model because the actions of group members are interdependent—the activity of any individual is influenced by the actions of other members of the group. Moreover, the actions of group members engaged in some collaborative tasks may need to follow a particular *order*. However, current techniques can account for either interdependence or order, but not both. In this paper, we present directed epistemic network analysis (dENA), an extension of epistemic network analysis (ENA), as a method that can simultaneously account for the interdependent and ordered aspects of CCT. To illustrate the method, we compare a qualitative analysis of two U.S. Navy commanders working in a simulation to ENA and dENA analyses of their performance. We find that by accounting for interdependence but not order, ENA was not able to model differences between the commanders seen in the qualitative analysis, but by accounting for both interdependence and order, dENA was able to do so.

Keywords: Complex collaborative thinking · Directed epistemic network analysis (dENA) · Epistemic network analysis (ENA)

1 Introduction

Many quantitative ethnographers study *complex collaborative thinking* (CCT) [12, 8, 17], often conceptualized as the process by which members of a group develop shared understanding. Researchers attempting to model CCT face the challenge of accounting for *interdependence*, or the direct influence of group activity on the actions of individual members. Interdependence implies that the relationship between events are inherently *temporal*—that is, individual actions are more strongly influenced by immediately preceding actions than by more distant ones [21]. The concurrent interdependent and temporal aspects of CCT make assessing individuals difficult, as models must account for individual contributions in the context of group contributions within recent temporal context. Additionally, some collaborative tasks require group members to perform a series of actions in a specific *order* to accomplish the shared goal. Current modeling

A. R. Ruis and S. B. Lee (Eds.): ICQE 2021, CCIS 1312, pp. 122–136, 2021.
https://doi.org/10.1007/978-3-030-67788-6_9

approaches can account for some combinations of interdependence, temporality, and order, but no extant techniques can account for all three simultaneously. In this paper, we propose a new technique, *directed epistemic network analysis* (dENA), an extension of epistemic network analysis (ENA) [18], that can construct models of CCT that account for interdependence, temporality, and order. We tested this novel technique on a well- dataset for which there are published findings on CCT, and we compared dENA models of CCT with ENA models. We found that by accounting for interdependence but not order, ENA was not able to model differences between individuals seen in the qualitative analysis, but by accounting for both interdependence and order, dENA was able to do so.

2 Theory

Quantitative ethnographers study tasks involving CCT across a wide range of domains, from collaborative problem solving by students learning social science research methods [12] to interactions between children and robots [8]. This prior works suggests that CCT involves three key elements: interdependence, temporality, and order.

In collaborative contexts, *interdependence* has been defined in terms of the coordination of behavior and information between group members [3]. In particular, as individuals participate in collaborative activities, they add information to the *common ground*, or the set of shared knowledge and experiences resulting from interactions with other members [1]. Group members then respond to information in the common ground, which influences subsequent actions and interpretations [4]. Thus, interdependence can be more generally defined as the influence of the group's actions on those of a given member of the group.

The interdependence of group processes also means that CCT has an important *temporal* aspect: events at any given point in time are influenced by events that occurred previously. Suthers and Desiato [21] argue that the *recent temporal context*, or immediately preceding events, have the greatest influence on the interpretation of subsequent actions and interactions. That is, actions taken by some group members affect the likelihood of actions taken by others in the near future [5]. For example, if one individual asks a question, another will likely respond in short order.

In addition, the *order* in which events unfold in the recent temporal context affects group members' future actions and interactions [15]. Some collaborative tasks require groups to follow a particular sequence of events, such as the specific order of actions naval navigation teams need to take to accurately track the position of their ships [6].

Thus, it may be important to simultaneously account for all three of these elements in models of CCT. However, existing approaches account for some, but not all, of these characteristics. For example, Sequential Pattern Mining and Lag Sequential Analysis account for order and temporality by using sliding windows to identify sequences of discourse moves occurring in the same recent temporal context [7, 14]. A key limitation of both methods is that while they can model the interdependence of a team as a whole, neither technique models *individual* CCT *in relation* to the contributions of collaborators; that is, they do not account for the how one specific individual's actions depend on others in a team or group.

Epistemic network analysis (ENA), on the other hand, models temporality and interdependence in collaborative activity [23] but does not explicitly represent the order of events in the model. That is, it can account for how he how one specific individual's actions depend on recent actions of others in a team or group, but not the order in which events occurred within a window of time. The algorithm accumulates *codes*, or indicators of meaningful discourse moves made by group members that occur *cotemporally*—that is, within the same recent temporal context. The algorithm visualizes the relative co-occurrence of codes using weighted network diagrams, with each node in the network corresponding to a code. The network diagram for each unit of analysis, or group member, is summarized using an ENA score and plotted in the same metric space. This enables researchers to directly compare the actions of specific individuals in the context of the group.

Recent work comparing ENA to Sequential Pattern Mining found that, in at least one context, ENA outperformed Sequential Pattern Mining as a measure of CCT [22]. This suggests that accounting for order may be less important than accounting for cotemporality, but there are contexts where accounting for order in CCT has been shown to be important, as in Hutchin's study of naval navigation teams [6].

Some studies have sought to use ENA to accomplish this goal. D'Angelo and colleagues [2] used ENA to model order by using two nodes for each code instead of one. One node represented instances when an individual was responding *to* the code; the other represented instances when an individual responded *with* the code. While D'Angelo et al. [2] were able to use this method to compare the operative discourse of surgical residents as they performed simulated procedures, there are drawbacks to this approach. In particular, the resulting ENA visualizations are difficult to interpret: it is hard to keep track of the difference between responding *to* and responding *with* the same code. Also, this approach doubles the number of codes, which can cause overfitting problems in models with fewer units of analysis.

Others have accounted for order by including only those connections that occurred in a particular sequence [16]. However, this approach only works if there is a strong justification for defining one or more connections as meaningful only in one order; and doing so makes the interpretation of resulting visualizations more difficult, as the network graphs do not indicate the directionality of the connections.

In what follows, we propose an approach to account for the ordering of events within an ENA framework. The technique, directed ENA (dENA), tracks both what individuals respond *with* and what they respond *to* as they act within the context of a group. Rather than accomplishing this by representing each code with two nodes, dENA represents the directionality of a response using triangles between each set of codes. An individual's network is then summarized with two ENA scores, drawn as a vector, to represent their overall responses as well as the common ground to which they responded. Thus, dENA provides a method to compare the discourse of individuals within a group, accounting simultaneously for interdependence, temporality, and order.

To test this technique, we use data on Navy air defense warfare teams that has previously been analyzed by quantitative ethnographic researchers using a variety of methodologies, including ENA [22, 23]. We compared ENA and dENA to address the following research questions:

1. What differences does *ENA* show between two units of analysis that are qualitatively different?
2. What differences does *dENA* show between two units of analysis that are qualitatively different?

3 Study 1

3.1 Methods

Data. We analyzed discourse data collected from U.S. Navy air defense warfare teams engaging in training scenarios. Each team's goal was to perform the *detect-to-engage* sequence, in which they must detect and identify nearby vessels, or *tracks*, and assess whether these tracks pose a threat to the Navy ship.

The detect-to-engage sequence typically begins with the detection and identification of a track. When a track's identity is uncertain, team members continue monitoring its behavior and make an assessment as to whether its behavior is threatening. Based on these assessments, teams decide to warn tracks of imminent attach or engage them in combat.

Every team participated in the same four training scenarios, with each scenario testing how the team handled a different set of tracks. Each team consisted of six participants: two commanders and four support roles. In this analysis, we focused on interactions between one of the command roles, the Tactical Action Officer (TAO), and other members of the team.

Coding. We analyzed the transcripts using the codes in Table 1, which were developed by Swiecki et al. [23] using a grounded approach informed by prior qualitative analyses on similar data [11] as well as existing air defense warfare literature [13]. The nCodeR package for the R statistical programming language was used to develop an automated classifier for each code in Table 1 [10]. All codes were validated at a kappa threshold of 0.65 and a rho threshold of 0.05.

Epistemic Network Analysis. We used ENA [18] to visualize and test differences in the discourse of commanders in both conditions. We conducted the analysis using the rENA package [9] for the R programming language.

The data were segmented by team and scenario. This segmentation defined the ENA *conversations*, or the set of utterances made by the team members over the course of the scenario. The ENA algorithm slid a moving window [20] of 5 lines over the conversations to identify co-occurrences between codes for each unit of analysis (i.e., for each commander). Once the co-occurrences were accumulated, they were transformed into high-dimensional vectors for each unit of analysis and normalized and centered. Next, the algorithm performed a dimensional reduction on the unit vectors via singular value decomposition. This process resulted in an ENA score for each unit of analysis on each dimension.

Each unit's score was visualized by projecting the resulting values from the first two dimensions of the dimensional reduction into a lower dimensional space. In addition

Table 1. Qualitative codes, definitions and examples.

Code	Definition	Example
DETECT/IDENTIFY	Talk about radar detection of a track or the identification of a track, (e.g., vessel type)	IR/EW NEW BEARING, BEARING 078 APQ120 CORRELATES TRACK 7036 POSSIBLE F-4
TRACK BEHAVIOR	Talk about kinematic data about a track or a track's location	AIR/IDS TRACK NUMBER 7021 DROP IN ALTITUDE TO 18 THOUSAND FEET
SEEKING INFORMATION	Asking questions regarding track behavior, identification, or status	TAO CO, WE'VE UPGRADED THEM TO LEVEL 7 RIGHT?
DETERRENT ORDERS	Giving orders meant to warn or deter tracks	TIC AIR, CONDUCT LEVEL 2 WARNING ON 7037
DEFENSIVE ORDERS	Giving orders to prepare defenses or engage hostile tracks	TAO/CO COVER 7016 WITH BIRDS

to scores, the method produces a weighted network graph for each unit of analysis whose nodes represent the codes and whose edges represent the relative frequency of co-occurrence between those codes. Thicker and more saturated lines indicate more frequent co-occurrence.

ENA uses an optimization routine to *co-register* each ENA score with its associated weighted network graph. By projecting each unit's ENA score and weighted network graph into the same metric space, ENA enabled us to compare which connections were stronger between different units of analysis, as well as define the dimensions along which units of analysis differed. The algorithm positions the nodes in the metric space by minimizing the distance between the plotted points and the centroids of each network. ENA also computes a representation to visually compare two networks in the same space by subtracting the connection weights of one unit's network from the other network and plotting the network difference.

3.2 Results

Qualitative Results. Previous analyses [22] suggested that a key issue for commanders in these scenarios is maintaining awareness of tactical information.

For example, the following excerpt shows part of one team's activity (Team 1) as they dealt with a hostile combat helicopter (which they refer to as "TRACK 7023"). The excerpt begins when the TAO orders the team to issue a warning to the helicopter ("LEVEL 3 WARNING") and to track it with onboard weapons.

Line	Team member	Utterance
6447	TAO	AIR TAO, COVER TRACK 7023 LEVEL 3 WARNING, SAY AGAIN, LEVEL 3 WARNINGS
6448	ADWC	AIR, COVERING INSIDE TRACK 7023
6449	TAO	ASK HIM TO VECTOR TO 000
6450	IDS	WHERE IS 23?
6451	TAO	THEY BEAR 047, 12 MILES

The action unfolds as follows:

6447. The TAO issues two orders. The first is a DEFENSIVE ORDER to the ADWC to "COVER TRACK 7023" the second is a DETERRENT ORDER to the IDS to warn the track to move away from the area.

6448. In response, the Air Defense Warfare Coordinator (ADWC) begins tracking the hostile helicopter with the onboard ship weapons.

6449. The TAO tells the Identification Supervisor (IDS) what the content of the warning should be: specifically, to tell the helicopter to turn to 0°, or due north, away from the ship.

6450. The IDS responds by SEEKING INFORMATION on the position of the helicopter, asking "WHERE IS 23?"

6451. The TAO, in turn, responds by describing the relevant part of the TRACK BEHAVIOR, which in this case is the helicopter's most recent position: "THEY BEAR 047, 12 MILES," meaning they are 12 miles north-east of the ship—and thus, turning north will lead them away.

In other words, Team 1 maintained tactical awareness because when a supporting member of the team (the IDS) was SEEKING INFORMATION about TRACK 7023's recent behavior, the TAO responded with the relevant information about the TRACK BEHAVIOR.

In the next excerpt, a second team (Team 2) also maintained awareness of tactical information about the helicopter designated as TRACK 7023, albeit in a different way than Team 1. The excerpt begins after the TAO orders the ADWC to attack the hostile helicopter, and a news helicopter is shot down by mistake:

Line	Team member	Utterance
4857	TAO	AWC/TAO HOW DO WE KNOW THAT WE SHOT DOWN THE NEWS HELO?
4858	ADWC	BECAUSE COUPLE OF SECONDS AFTER I FIRED A SHOT IT FLEW RIGHT ACROSS OUR BOW AND TOOK THE HIT
4859	TAO	AND THEN DISAPPEARED OFF THE SCREEN?
4860	ADWC	THAT'S AFFIRMATIVE YOU MIGHT WANT TO GO UP TO BRIDGE AND ASK

The action unfolds as follows:

4857. The TAO is SEEKING INFORMATION about how the ADWC knows the team had hit a news helicopter: "HOW DO WE KNOW THAT WE SHOT DOWN THE NEWS HELO?"

4858. The ADWC responds by describing the TRACK BEHAVIOR of NEWS HELO: shortly after the ADWC fired on Track 7023 (“COUPLE OF SECONDS AFTER I FIRED A SHOT”) the news helicopter flew into the line of fire (“FLEW RIGHT ACROSS OUR BOW AND TOOK THE HIT”).

4859. The TAO, now SEEKING INFORMATION to confirm that the news helicopter was actually hit, asks whether the helicopter disappeared off the screen.

4860. The ADWC responds that the helicopter did disappear (TRACK BEHAVIOR).

In other words, unlike the TAO on Team 1, who *responded* to SEEKING INFORMATION by *providing* TRACK BEHAVIOR, the TAO from Team 2 was SEEKING INFORMATION *from others* in order to *discover more* about the TRACK BEHAVIOR that led to the downing of a non-combatant.

Epistemic Network Analysis. In the ENA model shown in Fig. 1, the top of the space contains connections to SEEKING INFORMATION. The bottom of the space contains DETECT/IDENTIFY and TRACK BEHAVIOR, codes most associated with generating information about the tracks. This suggests that the commanders’ discourse was most distinguished on the vertical dimension in terms of *Providing Information* and *Seeking Information.*

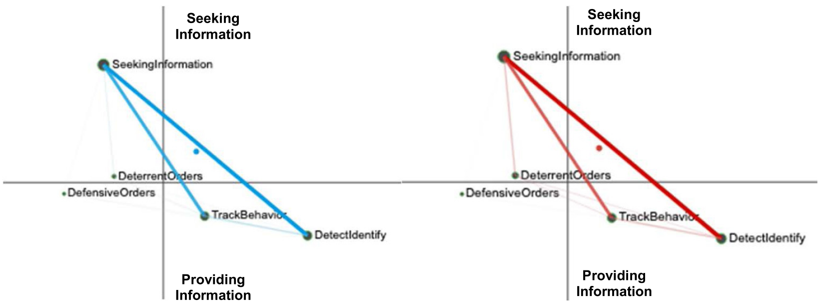

Fig. 1. ENA plots for the TAOs from Team 1 (blue) and Team 2 (red).

The ENA scores for both Team 1’s TAO (Fig. 1, left) and Team 2’s TAO (Fig. 1, right) are positioned toward to the top of the space. This means their discourse focused on connections to Seeking Information. For both TAO’s, the strongest connections were between SEEKING INFORMATION and DETECT/IDENTIFY as well as SEEKING INFORMATION and TRACK BEHAVIOR.

Comparing the networks of the two TAOs using a network difference graph shows that the TAOs’ networks are very similar (see Fig. 2). The ENA scores for both TAOs overlap, meaning there is little difference in the summary of these networks’ connections. Additionally, the strong connections in the individual plots between SEEKING INFORMATION and DETECT/IDENTIFY as well as SEEKING INFORMATION and TRACK BEHAVIOR are thin and desaturated, suggesting that there is little difference in the weight or strength of those connections between the networks. The network difference graph thus suggests little difference in the discourse of the two teams’ TAOs.

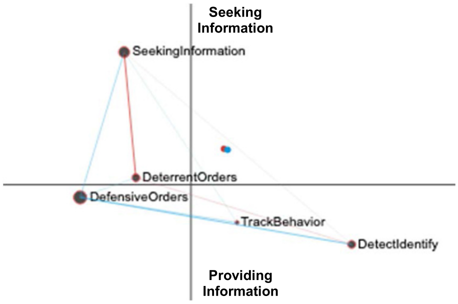

Fig. 2. Network difference graph for the TAOs from Team 1 (blue) and Team 2 (red).

These results suggest that both commanders were involved in exchanges in which DETECT/IDENTIFY and TRACK BEHAVIOR information was missed or incomplete and had to be clarified. These findings are consistent with those of Swiecki et al. [22], who found that teams often had to ask explicitly for Tactical Information to be repeated.

3.3 Discussion

The qualitative results show that both TAOs were involved in situations in which connections were made between SEEKING INFORMATION and TRACK BEHAVIOR. However, the TAO in Team 1 *responded* to SEEKING INFORMATION with TRACK BEHAVIOR while the TAO in Team 2 was SEEKING INFORMATION *in order to understand* TRACK BEHAVIOR. That is, the Team 1 TAO was providing information about TRACK BEHAVIOR *after* another team member was SEEKING INFORMATION. The Team 2 TAO did the opposite: they were responding when another team member described TRACK BEHAVIOR by SEEKING INFORMATION (that is, *additional* information) about the tactical situation.

In other words, although the ENA plots show that both TAO's made connections between SEEKING INFORMATION and TRACK BEHAVIOR:

1. What the two TAOs were doing was actually quite different;
2. These differences can be seen in the order in which the codes occur in the discourse; and
3. These differences are not apparent because the unordered ENA model treats as identical: (a) the TAO *providing* information to other team members who are asking about the tactical situation, and (b) the TAO *seeking* additional information from other team members to clarify the tactical situation.

4 Study 2

4.1 Methods

Directed Epistemic Network Analysis. We conducted the second study using the same data and coding scheme as the first. For this study, we developed and implemented *directed epistemic network analysis* (dENA), an extension of ENA that accounts for the order in which a connection between two codes occurred.

The ENA algorithm uses a moving window to identify connections formed from a current line of data (e.g., turn of talk), or *response*, to the preceding lines within the window, or *common ground*. These connections counts are then accumulated into a *symmetric adjacency matrix* for each unit of analysis: the number of connections from any code A to code B are the same as the number of connections from B to A.

Rather than produce a symmetric adjacency matrix for each, unit of analysis, dENA accounts for the order in which the connection occur by constructing an *asymmetric adjacency matrix* for each unit of analysis: the number of connections from any code A to code B may be different than the number of connections from B to A. Figure 3 highlights how this matrix is created using coded data from the first qualitative example in Study 1. With a moving window of 5 lines, dENA collapses code occurrences from the common ground (lines 2–5) and response (line 6) using a binary summation. The binarized summation of the ground and response code occurrences are then represented as vectors and multiplied to construct an asymmetric adjacency matrix.

Line	Seeking Information	Detect/ Identify	Track Behavior	Deterrent Orders	Defensive Orders
1	0	1	0	0	0
2	0	0	0	1	1
3	0	0	0	0	0
4	0	0	0	0	0
5	1	0	0	0	0
6	0	0	1	0	0

Ground Summation	1	0	0	1	1
Response Summation	0	0	1	0	0

Fig. 3. The top section shows a moving window (indicated by the red dashed line) on coded data. Line 6 represents the current turn of talk, or the *response* line. The portion of the window shaded gray represents the recent temporal context, or the *ground* lines. The bottom section shows the ground and response summation vectors for the indicated window.

The asymmetric adjacency matrices are then transformed to create two high-dimensional vectors for each unit of analysis, the ground vector and the response vector. The ground vector represents the connections formed from the codes in the common ground to the codes in the unit's responses. The response vector represents the connections formed from the unit's contributions back to the contributions in the common ground. Put another way, the ground vector summarizes what a given unit *responds to*, while the response vector summarizes what a given unit *responds with.*

The ground and response vectors for all units are normalized and centered and the algorithm performs a dimensional reduction via a singular value decomposition of the matrix of either the ground or response vectors. This process, which involves the same mathematics used in ENA, results in a pair of dENA scores for each unit of analysis in the lower dimensional space: a *ground score* and a *response score*.

The scores are visualized by plotting them in the lower dimensional space resulting from the dimensional reduction (see arrow in Fig. 4). For each unit, its scores are represented by a vector with its head at the response score and tail at the ground score. Subpopulations within the data are summarized by independently calculating the mean response and ground scores. The scores are then connected by a vector from ground mean to response mean.

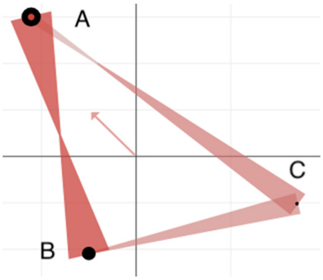

Fig. 4. Sample unit's dENA network. The overall size of the nodes represents the relative response strength with each code. The red dot in the middle of CODE A represents self-connections. The triangles represent directed connections: more saturated and opaque connections mean stronger connections. The network is summarized by two scores, connected as a vector. The tail of the vector is a ground score summarizing the ground connections and the head is the response score representing the response connections.

To help interpret these vectors, the algorithm co-registers *directed and weighted network graphs* in the same low dimensional space. For each unit, its graph shows the strength and directionality of the connections it made.

The nodes of the network correspond to the codes, and node size is proportional to the number of occurrences of that code *as a response* in the data, with larger nodes indicating more responses. The color and saturation of the circle within each node is proportional to the number of self-connections for that code: that is, when a code appears in both the response and ground of a given window. Colored circles that are larger and more saturated reflect codes with more frequent self-connections. For example, Fig. 4 suggests that roughly a quarter of responses made with code A were responding to code A.

The relative frequency of connections between any two codes are represented by two triangles. The connections between code A and code C, for example, are represented by one triangle with its base at code A pointing towards code C and the other with its base at code C pointing towards code A. Thicker and more saturated triangles indicate more frequent connections. Each triangle represents an ordered connection between two codes. The base of the triangle represents the common ground that a code was responding to and the opposite vertex of the triangle represents the response. Thus, in a given unit's network, the triangle with a base at code C and vertex directed towards code A would be interpreted as that unit's relative response to code C with code A.

The location of the tapered point where two triangles meet indicates the relative proportion of responses of one code to the other. For example, if the two triangles meet closer to code A's node, connections of A in response to C were more frequent than the reverse.

Network nodes in dENA are positioned in the space using the same optimization routine used in ENA, except that in dENA, the algorithm minimizes the distance between the *means of the ground and response scores* and the centroids of the corresponding networks. As a result, the dENA metric space can be interpreted based on the location of the nodes. Units with vectors on the right side of the space have more frequent connections between the codes on the right side of the space. Similarly, units with vectors on the left have more frequent connections between the codes on the left side of the space. The vector that represents a unit of analysis shows the directionality of the network: which nodes the unit is responding to (ground score, or tail of the vector) and what responses are made (response score, or head of the vector). Like the unit vectors, the network graphs can be averaged for subpopulations within the data to view their overall patterns of directed connections. We visually compared the dENA network graphs of two units of analysis in the same space by subtracting the connection weights of one unit's network from the other and plotting the network difference.

4.2 Results

Qualitative Results. Recall that the qualitative results from Study 1 show that the TAO on Team 1 responded to SEEKING INFORMATION with TRACK BEHAVIOR to help the team maintain tactical awareness of the situation. In contrast, the TAO on Team 2 responded to TRACK BEHAVIOR by SEEKING INFORMATION to understand the tactical situation.

Directed Epistemic Network Analysis. We examined the directed epistemic networks of the TAOs from Team 1 and Team 2. The left side of the space contains DETECT/IDENTIFY and TRACK BEHAVIOR, codes associated Providing Information. The right side of the space contains connections to SEEKING INFORMATION. This means the TAOs' discourse was most distinguished in terms of Providing Information and Seeking Information (Fig. 5).

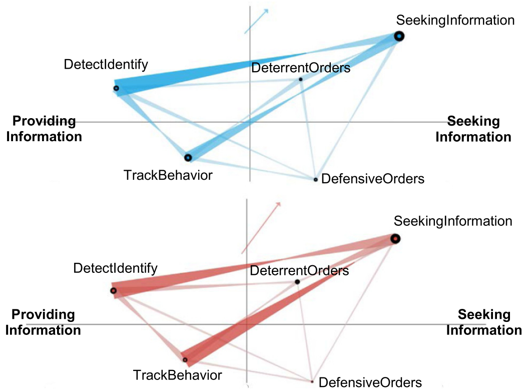

Fig. 5. dENA plots for the TAOs from Team 1 (blue) and Team 2 (red).

The strongest directed connections in each TAO's network were between SEEKING INFORMATION and TRACK BEHAVIOR and SEEKING INFORMATION and DETECT/IDENTIFY: the triangles with the widest base, darkest saturation, and heaviest opacity are based at TRACK BEHAVIOR and DETECT/IDENTIFY and point towards SEEKING INFORMATION. This means that relative to other directed connections, each TAO responded more to TRACK BEHAVIOR and DETECT/IDENTIFY with SEEKING INFORMATION.

Each TAO's vector summarizes the information in their respective networks. Overall, both TAOs were responding to Providing Information with Seeking Information. The ground scores, or tails of the summary vectors, are positioned closer to the side associated with Providing Information. The response scores, or heads of the summary vectors, are closer to the side associated with Seeking Information.

The network subtraction shows that the greatest difference in their networks was the connection between SEEKING INFORMATION and TRACK BEHAVIOR (Fig. 6). The TAO on Team 1 (blue) responded strongly *to* SEEKING INFORMATION *with* TRACK BEHAVIOR relative to the TAO on Team 2. In contrast, the TAO on Team 2 responded strongly *with* SEEKING INFORMATION *to* TRACK BEHAVIOR compared to the other TAO.

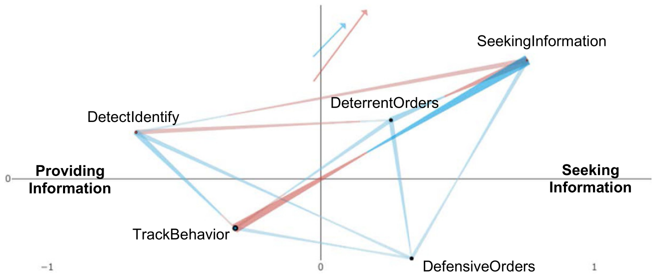

Fig. 6. dENA network difference plot for the TAOs from Team 1 (blue) and Team 2 (red).

The two triangles meet at a point approximately midway between SEEKING INFORMATION and TRACK BEHAVIOR. This means that the TAO on Team 1 responded to SEEKING INFORMATION with TRACK BEHAVIOR with approximately the same relative frequency as the TAO on Team 2 (red) responded to TRACK BEHAVIOR with SEEKING INFORMATION.

Taken together, this suggests that the key difference in how the two TAOs made connections between TRACK BEHAVIOR and SEEKING INFORMATION was the order of the connections rather than their relative frequency.

4.3 Discussion

The dENA results show that both TAOs were involved in situations where tactical information was missed and needed to be clarified. Each TAO's dENA plots shows strong responses between TRACK BEHAVIOR and DETECT/IDENTIFY, on one hand, and SEEKING INFORMATION on the other. Comparing the differences between the TAOs shows that the

Team 1 TAO responded to SEEKING INFORMATION *with* TRACK BEHAVIOR while the TAO on Team 2 was SEEKING INFORMATION *in response to* TRACK BEHAVIOR. These results align with the qualitative results of Study 1—the Team 1 TAO was providing tactical information to other team members while the TAO on Team 2 was SEEKING INFORMATION about TRACK BEHAVIOR to better understand the tactical situation.

5 General Discussion

In this paper, we compared the results of a qualitative analysis with those of two quantitative analyses leveraging ENA. Using data on U.S. Navy air defense warfare teams, we focused on the discourse of two tactical officers (TAOs) on two different teams that participated in the same training simulations. Our qualitative analysis showed that both teams were involved in situations where information was reiterated to maintain awareness of the tactical situation. However, the TAO on Team 1 *provided* tactical information *to* team members about hostile track behavior, while the TAO on Team 2 *requested* information *from* team members to gain a better understanding of the situation. This suggests a need to consider the order in which the two TAOs made connections between SEEKING INFORMATION and TACTICAL INFORMATION to evaluate their performance and give relevant feedback.

The ENA model (Study 1) showed that both TAOs were *involved* in situations requests were made for tactical information. However, the model showed little difference between the two TAOs. In contrast, the dENA model (Study 2) was aligned with the qualitative findings: relative to the TAO on team 2, the TAO on Team 1 responded more often *to* SEEKING INFORMATION by describing TRACK BEHAVIOR; relative to the TAO on Team 1, the TAO on Team 2 responded more often to TRACK BEHAVIOR by SEEKING INFORMATION.

These results demonstrate dENA's ability to model the interdependent, temporal, *and* ordered aspects of Complex Collaborative Thinking (CCT). Like ENA, dENA uses a moving window to account for the ways group members respond to events in the recent temporal context [20], and accounts for interdependence by measuring connections within that window to model individuals' responses to information in the common ground [1]. But unlike ENA, dENA preserves the order of events, enabling the algorithm to model the influence of information *from* the common ground *on* individuals' responses [4].

This approach builds on, but also extends, the visual affordances of ENA. ENA represents networks in two ways: (1) as network graphs, where the strength of connections between nodes is represented by line thickness and saturation; and (2) as a set of points in a metric space that allows networks to be compared statistically in terms of their content. Crucially, ENA co-registers a statistical model of networks with their individual network graphs such that the dimensions of the statistical model can be interpreted using the positions of nodes in the network. dENA retains these key properties but adapts them to account for the order in which the events that generate the networks occur. First, the network graphs in dENA account for directionality between nodes by representing each pairwise connection with a pair of triangles, such that the thickness and saturation of the triangles represents the total strength of the pairwise connection, and the relative heights of the triangles represent the proportion of connections going in each direction. The number of times each node is part of a response is modeled by the size of the node, and self-references are represented as a proportion of the node size. Second, each network is represented a *vector* in the metric space pointing from the position of its accumulated ground to the position of its accumulated response. dENA co-registers the network graphs with the vector corresponding to each network such that the dimensions of the metric space can be interpreted in terms of the flow of information in the networks.

This study has several limitations. The analysis was conducted using a single dataset. However, the data we used—and the analysis we provided—was only meant to provide an example of how dENA can model interdependent, temporal, and ordered elements common in a CCT context. Future work will, of course, explore using this approach to model other domains of CCT. Second, we have yet to develop a method for testing whether differences between network vectors are significant. We intend to incorporate significance testing in future work. Additionally, the current implementation of dENA does not yet account for sequences of more than two events. Our future work will explore extensions of dENA to account for and represent longer sequences.

These limitations notwithstanding, this work shows dENA is a method that can analyze the CCT of individuals in the context of a group, representing ordered connections in a way that visually captures qualitative differences that unordered models cannot. Unlike other approaches, dENA is thus able to present an interpretable visualization capable of simultaneously accounting for interdependence, temporality, *and* order.

Acknowledgements. This work was funded in part by the National Science Foundation (DRL-1661036, DRL-1713110), the Wisconsin Alumni Research Foundation, and the Office of the Vice Chancellor for Research and Graduate Education at the University of Wisconsin-Madison. The opinions, findings, and conclusions do not reflect the views of the funding agencies, cooperating institutions, or other individuals.

References

1. Clark, H.H.: Using Language. Cambridge University Press, Cambridge (1996)
2. D'Angelo, A.-L.D., Ruis, A.R., Collier, W., Shaffer, D.W., Pugh, C.M.: Evaluating how residents talk and what it means for surgical performance in the simulation lab. American Journal of Surgery. In press (2020)
3. DeChurch, L.A., Mesmer-Magnus, J.R.: The cognitive underpinnings of effective teamwork: a meta-analysis. J. Appl. Psychol. **95**(1), 32–53 (2010)
4. Dillenbourg, P.: Collaborative Learning: Cognitive and Computational Approaches. Advances in Learning and Instruction Series. Elsevier Science Inc, New York (1999)
5. Halpin, P.F., von Davier, A.A.: Modeling collaboration using point processes. In: von Davier, A.A., Zhu, M., Kyllonen, P.C. (eds.) Innovative Assessment of Collaboration, pp. 233–247. Springer, Cham (2017)
6. Hutchins, E.: Cognition in the Wild. MIT Press, Cambridge (1995)
7. Kapur, M.: Temporality matters: advancing a method for analyzing problem- solving processes in a computer-supported collaborative environment. Int. J. Comput. Supported Collaborative Learning **6**(1), 39–56 (2011)

8. Kim, Y., Tscholl, M.: The dynamic interaction between engagement, friendship, and collaboration in robot children triads. In: Eagan, B., Misfeldt, M., Siebert-Evenstone, A. (eds.) ICQE 2019. CCIS, vol. 1112, pp. 307–314. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-33232-7_27
9. Marquart, C.L., Swiecki, Z., Collier, W., Eagan, B., Woodward, R., Shaffer, D. W. rENA: Epistemic Network Analysis (Version 0.1.3) (2018)
10. Marquart, C.L., Swiecki, Z., Eagan, B., Shaffer, D. W.: ncodeR (Version 0.1.2) (2018)
11. Morrison, J.G., Kelly, R.T., Moore, R.A., Hutchins, S.G.: Tactical decision making under stress (TADMUS) decision support system, vol. 13 (1996)
12. Nachtigall, V., Sung, H.: Students' collaboration patterns in a productive failure setting: an epistemic network analysis of contrasting cases. In: Eagan, B., Misfeldt, M., Siebert-Evenstone, A. (eds.) ICQE 2019. CCIS, vol. 1112, pp. 165–176. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-33232-7_14
13. Paris, C., Johnston, J.H., Reeves, D.: A schema-based approach to measuring team decision making in a Navy combat information center. The Human in Command, pp. 263–278 (2000)
14. Perera, D., Kay, J., Koprinska, I., Yacef, K., Zaïane, O.R.: Clustering and sequential pattern mining of online collaborative learning data. IEEE Trans. Knowl. Data Eng. **21**(6), 759–772 (2009)
15. Reimann, P.: Time is precious: variable- and event-centred approaches to process analysis in CSCL research. Comput. Supported Learn. **4**, 239–257 (2009)
16. Ruis, A.R., Rosser, A.A., Quandt-Walle, C., Nathwani, J.N., Shaffer, D.W., Pugh, C.M.: The hands and head of a surgeon: modeling operative competency with multimodal epistemic network analysis. Am. J. Surgery **216**(5), 835–840 (2018)
17. Ruis, A.R., Siebert-Evenstone, A.L., Pozen, R., Eagan, B., Shaffer, D.W.: Finding common ground: a method for measuring recent temporal context in analyses of complex, collaborative thinking. In: Lund, K., Niccolai, G., Lavoué, E., Hmelo-Silver, C., Gwon, G., Baker, M. (eds.) 13th International Conference on Computer Supported Collaborative Learning (CSCL), vol. 1, pp. 136–143. International Society of the Learning Sciences, Lyon, France (2019)
18. Shaffer, D.W., Collier, W., Ruis, A.R.: A tutorial on epistemic network analysis: analyzing the structure of connections in cognitive, social, and interaction data. J. Learn. Analytics **3**(3), 9–45 (2016)
19. Shaffer, D.W.: Quantitative Ethnography. Cathcart Press, Madison (2017)
20. Siebert-Evenstone, A., Arastoopour Irgens, G., Collier, W., Swiecki, Z., Ruis, A.R., Shaffer, D.W.: In search of conversational grain size: Modelling semantic structure using moving stanza windows. J. Learn. Anal. **4**(3), 123–139 (2017)
21. Suthers, D.D., Desiato, C.: Exposing chat features through analysis of uptake between contributions. In: 2012 45th Hawaii International Conference on System Sciences, pp. 3368–3377. IEEE, Hawaii (2012)
22. Swiecki, Z., Lian, Z., Ruis, A., Shaffer, D.W.: Does order matter? Investigating sequential and cotemporal models of collaboration. In: Lund, K., Niccolai, G.P., Lavoué, E., Gweon, C.H., Baker, M. (eds.) 13th International Conference on Computer Supported Collaborative Learning (CSCL), vol. 1, pp. 112–119. International Society of the Learning Sciences, Lyon, France (2019)
23. Swiecki, Z., Ruis, A.R., Farrell, C., Shaffer, D.W.: Assessing individual contributions to collaborative problem solving: a network analysis approach. Comput. Hum. Behav. **104** (2020)