Epistemic Association Rule Networks: Incorporating Association Rule Mining into the Quantitative Ethnography Toolbox

Andres Felipe Zambrano[1[0000-0003-0692-1209]], Ryan S. Baker[1[0000-0002-3051-3232]], Shruti Mehta[1[0009-0004-0928-8677]], and Amanda Barany[1[0000-0003-2239-2271]]

[1] University of Pennsylvania, Philadelphia, PA, United States
azamb13@upenn.edu

Abstract. Since the formalization of Quantitative Ethnography (QE) as a methodology, Epistemic Network Analysis (ENA) has been the most widely used analytical tool in the community. ENA has proven itself highly useful for QE research, particularly in modeling temporal associations between code pairs. However, integrating additional techniques that systematically reveal more complex patterns in data can not only supplement the insights derived from ENA but also broaden the range of research that can be conducted through a QE lens. To that end, this paper proposes Association Rule Mining (ARM) as an additional technique for QE. We introduce a new visualization of the results of ARM, which we term Epistemic Association Rule Networks (EARN), which combines ENA's visualization strengths with ARM's ability to identify more complex patterns. Using human-human tutoring transcripts, we illustrate how ARM and EARN can complement ENA by offering insights on directional conditional relationships between groups and pairs of constructs, offering a more nuanced understanding of complex phenomena.

Keywords: Association Rule Mining, Epistemic Network Analysis, Epistemic Association Rule Networks.

1 Introduction

Since the formalization of Quantitative Ethnography as a methodology, Epistemic Network Analysis (ENA) has been the most popular analytical tool in the community [1]. However, members of the QE community have argued that it is important to differentiate QE as a methodology from ENA as an analysis tool [2]. QE is a multidisciplinary field encompassing a variety of methodological approaches and tools where ethnography and quantification are both foundational [3]. Given this framing, there are a range of additional quantification techniques that can complement the insights provided by standard ENA models, so long as these techniques aid in the systematic and fair interpretation of the contextual meaning represented in the data. For example, models such as Process Mining [4] and CORDTRA [5] have been used to recognize chronological sequences and patterns that the standard ENA model could not show. Inspired by these ideas for complementing standard ENA models, we propose Association Rule Mining (ARM; [6]) as another technique that might be useful for the QE toolbox.

ARM is a data mining technique extensively utilized by data scientists to identify relationships, patterns, and associations among items in large databases. Like Epistemic Network Analysis, ARM attempts to recognize associations through the occurrence and co-occurrence of constructs within specified timeframes, instances, or stanzas. However, there are notable differences between the two methods. In specific, ARM is capable of discovering rules or associations involving groups of more than two constructs and provides a variety of metrics to quantify these associations, including conditional relationships between constructs [7]. These features of ARM suggest it has the potential to offer insights about complex phenomena that are both complementary and supplementary to ENA. However, while some visualization techniques like Association Rule Networks (ARN; [8, 9]) have been developed to present ARM findings, they generally do not match the interpretability of ENA models.

In this paper, we explore how ENA and ARM can be integrated to enhance data analysis. Specifically, we introduce a novel network visualization called Epistemic Association Rule Networks (EARN), which adapts the traditional ENA visualization to display association rule patterns based on ARM metrics. We use transcripts from human-human math tutor lessons as a case study to illustrate the utility of ARM and EARN as a complement to standard ENA. Our goal is to show how the tools, metrics, and techniques commonly used by the ARM community may be used to enrich QE research.

2 Two Approaches

2.1 Epistemic Network Analysis

Epistemic Network Analysis (ENA) is a quantitative ethnographic technique designed to model the structure of connections between a set of constructs. A common assumption of the method is that the selected constructs represent different aspects of cognition or behavior relevant to the study and their structure of connections is meaningful to analyze the discourse and culture embedded in that data [10, 11]. ENA has been used for a wide range of analyses, including learning processes in digital learning environments [12]; task performance [13]; gaze patterns [14]; team communication [15]; social media [16]; and video game players' behaviors [17].

The connections identified by ENA are modeled by quantifying the co-occurrence of constructs within and across lines of coded data, which are segmented into conversations of discrete data and organized into units of analysis that can be visually and statistically compared through weighted networks of co-occurrences. Typically, the ENA algorithm uses a moving window to accumulate these connections between codes in the current line and prior lines in a conversation [11, 18]. These matrices of code connections are then cumulatively aggregated across all lines for each analysis unit, and normalized to adjust for variations in the number of coded lines per unit. Subsequently, ENA employs singular value decomposition (and additional rotations in cases where the main goal is to compare 2 or multiple groups) to derive two orthogonal dimensions that maximize the variance explained by each dimension, facilitating the visualization and interpretation of these networks. For a more detailed explanation of the mathematical processes involved, refer to [11, 19].

Networks in ENA are visualized using network graphs, where nodes represent the codes, and edges indicate the relative frequency of co-occurrence between two codes. Each unit of analysis is represented by two coordinated visualizations: (1) a plotted point, which marks the location of that unit's network in the low-dimensional projected space, and (2) a weighted network graph, which represents the strength of connections between each pair of codes. The positions of the nodes in the network graph are calculated through an optimization routine that aims to minimize the discrepancy between the plotted points and their corresponding network centroids. This alignment of network graphs with the projected space allows the positions of the nodes (and the connections they establish) to interpret the dimensions of the projected space, thereby explaining the positions of plotted points within that space, although with the limitations inherent to any dimensionality reduction.

ENA enables the comparison of units of analysis through examination of the positions of plotted points, individual networks, mean plotted point positions, and mean networks. Additionally, comparisons can be made using network difference graphs generated by subtracting the weight of each connection in one network from the corresponding connections in another network and rotating the networks to maximize the variance of the groups being compared towards one single axis, highlighting visually the differences between them.

Ultimately, the strengths of this technique lie in the capacity for epistemic networks to account for associations across multiple themes in complex temporal data. Epistemic networks also create a common comparison space to understand relative patterns within and across units of analysis. However, like any approach, ENA alone cannot address all types of inquiries. In particular, ENA does not look at relationships between more than two constructs at a time, and traditional ENA also does not give ordered (chronological), or conditional (one-directional – where A implies B, but B may not imply A -- but not necessarily causal or chronological) directionality of associations. Although chronological directionality has been developed in ordered networks [20], conditional directionality has not been addressed yet. As ARM can consider more than 2 constructs at a time and offers conditional but not chronological directions for associations, it offers information that is complementary to ENA.

2.2 Association Rule Mining

Association Rule Mining (ARM) is a data mining technique that identifies associations among items or variables within large datasets [6]. ARM primarily seeks to discover rules framed as "If → Then" statements. The "If" portion of the rule, known as the antecedent, comprises one or more conditions whose fulfillment suggests the likelihood of the "Then" outcome or the consequent. For instance, a rule derived through ARM might state: If a teacher provides instruction and encouragement, then it is likely that the teacher also offers feedback. Though the two methods have many similarities, the primary distinction between ARM and ENA lies in the scope of connections or rules they can identify. Standard ENA primarily focuses on the co-occurrence of pairs of constructs. In contrast, ARM is not limited to dyadic associations; it can identify rules involving multiple constructs. For instance, in the above teacher mentoring example, there were two constructs in the antecedent (providing instruction and encouragement) and one in the consequent (giving feedback). Thus, ARM can potentially complement standard ENA by uncovering associations among groups of three or more constructs that co-occur within a defined conversation.

Additionally, the rules identified by ARM are not necessarily symmetrical, offering further insight into the relationships between the two or more constructs being considered by the rule. For example, the rule "If (teacher provided encouragement) → Then (teacher offered feedback)" might be more prevalent than its inverse, where the antecedent and consequent are swapped. In this example, ARM does not necessarily identify those cases where *Instruction* occurred chronologically before *Offering Feedback*, as ONA [20] would recognize. This association rule specifically states that the teacher consistently *Offered Feedback* when *Providing Encouragement*. But note that this is different from discovering that the teacher consistently *Provided Encouragement* when *Offering Feedback*. Furthermore, none of these patterns depends on the chronological order. It is possible for some constructs (consequents) to consistently be associated with others (antecedents) without the reverse being true, or without having a unique chronological/temporal association between them. In this way, the nuanced understanding of conditional rules provided by ARM expands upon what can be learned from the paired associations displayed in standard ENA and the chronological/order relations displayed in ONA.

An alternative method, Sequential Pattern Mining (SPM, [21]), was proposed to analyze antecedent-consequent relationships while also accounting for their sequential order. Like ONA, SPM identifies only items, events, or constructs as antecedents if they temporally precede the consequent within a specified timeframe or stanza. The choice of ARM versus SPM, like the choice between ENA and ONA, depends on the research goals and context. In this article, we focus on comparing ARM to standard ENA more broadly rather than focusing specifically on their SPM and ONA variants.

2.3 Association Rule Mining Metrics

Building on the conditional and not necessarily symmetric associations identified by ARM, a variety of metrics for quantifying the goodness of specific rules have been proposed. The two most widely employed metrics are support and confidence. Support measures the probability of the co-occurrence of two or more constructs across the entire dataset. For example, when coding different conversations, the support of a rule can be calculated as the number of conversations where the 2 constructs appeared, divided by the total number of conversations. In the context of a moving stanza (the most commonly used aggregation method in ENA for building connections), support might be calculated as the ratio of stanzas where the constructs co-occur to the total number of stanzas:

$$Support\ (A, B) = \frac{\#Stanzas\ containing\ A\ AND\ B}{\#\ Stanzas}$$

As it relies solely on co-occurrences and not on conditional probabilities based on the presence of either the antecedent or the consequent, support is a symmetric metric that reflects the prevalence of an association, preferring common patterns to rare combinations. This metric is conceptually similar (though not mathematically equivalent) to the connection weights in ENA after normalizing the adjacency vectors to calculate relative frequencies of co-occurrence (which divides each component of the vector by the length of the vector, scaling the length of the vector to 1), before applying singular value decomposition or other dimensionality reduction techniques [10, 11, 19].

A second metric, confidence, measures the likelihood that the consequent appears when the antecedent is present. It is mathematically calculated as the ratio of the support for the co-occurrence of both antecedent and consequent to the support for the occurrence of the antecedent alone:

$$Confidence(A \rightarrow B) = \frac{\#Stanzas\ containing\ A\ AND\ B}{\#\ Stanzas\ containing\ A} = \frac{Support(A, B)}{Support(A)}$$

This metric is particularly useful for complementing standard ENA as it adds a measure of conditional relationship, aiding in the identification of potential asymmetries in the associations revealed by both standard ENA and the support metric.

Beyond confidence and support, additional metrics are used to evaluate association rules, such as Cosine, Lift, Jaccard, and Cohen's Kappa [7]. These metrics are collectively referred to as interestingness metrics. Merceron and Yacef [22] suggest that these metrics, particularly Cosine and Lift, may be more appropriate for analyzing educational data than the more commonly used measures of confidence and support. Bazaldua and colleagues [23] argue for using Jaccard after taking confidence and support into consideration. Given the differences in the sets of rules that interestingness metrics can reveal [23], integrating ARM could allow researchers to include more varied metrics, uncovering sets of rules that neither support nor standard ENA would detect. Although this paper primarily focuses on the traditional ARM metrics of support and confidence to draw on the parallels and differences between ARM and standard ENA, exploring a variety of metrics could also significantly enhance analysis within ENA. This possibility will be further explored in the discussion section.

3 Methods

3.1 Dataset & Codebook

For this research, we utilized transcripts from three tutoring sessions orchestrated by a non-profit organization to support mathematics learning for students from high-poverty schools in an urban area of the northeastern United States. These sessions were conducted virtually in small groups, with one tutor and one or two students. All participants were 9th-grade students enrolled in Algebra I.

Table 1 presents the codebook used in this study, originally proposed in [24]. Using the same dataset we employ in this study (plus an additional small transcript used for inter-rater reliability checking), this codebook was created in the process of evaluating various methods of incorporating ChatGPT into the codebook-development process. It was developed by prompting ChatGPT for suggestions for constructs to explore teaching methodologies in online tutoring sessions, supplying ChatGPT with the dataset, and then by one expert qualitative coder refining ChatGPT suggestions based on their expertise.

Table 1. Codebook originally proposed in [24].

Construct	Definition	Examples
Greeting (G)	The initial interaction between the tutor and student, often at the beginning or end of the session. Anytime a salutation or farewell is exchanged.	"Hello." "Cheers." "I'll see you in a couple of days." "Enjoy the rest of your day." "Bye."
Instruction (I)	Specific instructions or directions posed by the tutor throughout the lesson.	"So, we're going to test it out. I'm going to have you guys work on this do now right here." "Go ahead and fill that out."
Guiding Feedback (GF)	Guided practice through a math problem by the tutor. Feedback on the student's work or response and clarification or explanation of a concept or instruction.	"Not quite. I'm not sure why you have these X's." "No, not quite one x because you divided the negative three by three but did you divide the x by x?"
Aligning to Prior Knowledge (PK)	Instances when the tutor brings attention to a previous math concept that a student knows or has discussed in a session.	"Remember, what does factor mean?" "But remember, what's in your parentheses should be the same if you did it right."
Check for Understanding and Engagement (UE)	The tutor presents checks for understanding as questions to students. Students answer questions or provide input to tutor's questions.	Tutor: "How do you figure out what's halfway?" "Why do you think we might have done that?" Student: "We should have negative three minus equals four."
Technical or Logistics (TL)	Tutor comments related to the technical aspects or logistics of the lesson.	"Your camera's looking at the ceiling." "Did you lose connection?"
Encouragement (E)	Affirmative statements from the tutor recognizing student's efforts, answers, or performance. Anytime the tutor provides a positive acknowledgment or praise.	"Perfect." "Good job." "Great." "You're getting it, man." "We are going to get it"
Time Management (TM)	Statements regarding the duration left, the need to move on, or how much has been covered. Any mention of time, pacing, or the order of topics.	"We have about 5 minutes." "Class is almost halfway over." "This should go fairly quickly so we can finish the lesson today."

Barany et al. [24] compared this codebook with three other codebooks derived from the same data set, including a codebook developed solely by a human coder, comparing codebooks in terms of clarity, ease of use, mutual exclusivity, inter-rater reliability, and conceptual overlap. Clarity, ease of use, and mutual exclusivity were measured using 5-point Likert-scale coders' rankings after two rounds of human coding, involving a new set of coders who had not previously seen the codebook or data. The codebook selected for this study showed a higher average rating across these three measures. It also achieved higher inter-rater reliability between the human coders. Additionally, we engaged in social moderation of the entire dataset [25] by two coders, including one of the original coders in the discussion process, to achieve complete agreement.

3.2 Comparison between ENA and ARM

To compare ARM and standard ENA, we initially focused on one of the three sessions available in this dataset. This particular session was selected because it featured the most active interaction between the students and the tutor, providing a rich sample of instances for investigating the potential differences in what findings ARM and ENA could reveal.

In our dataset, each line corresponds to a tutor or student sentence. We opted for a moving stanza of size 8, based on the typical number of sentences a tutor speaks before being interrupted or responded to by a student. Both ARM and standard ENA utilized this stanza size to ensure comparability of results. We also experimented with other stanza sizes ranging from 6 to 10 (based on the common range of tutor utterances before being interrupted or responded to) but did not observe significant variations in the outcomes. Given that each session was a continuous one-hour block without breaks or abrupt interruptions, each session was our conversation variable.

To facilitate comparison between ARM and standard ENA, and to make the results of ARM easier to understand, we developed a new visualization of ARM, Epistemic Association Rule Network (EARN), adapting ENA diagrams created using the ENA Web Tool (version 1.7.0) [26]. Although past approaches have proposed diagrams for visualizing rules found using ARM (e.g. [8, 9]), these visualizations usually focus on one consequent as the objective or final node of the diagram or focus only on rules of size 2, and mainly consider confidence for displaying the edges of the network. Although these approaches have been useful for finding and visualizing the associations with a specific target construct, they cannot show a general view of the associations present in the dataset, unlike ENA. Additionally, the location of the nodes in these approaches does not offer the same type of interpretability provided by ENA models. For this reason, we preferred to adapt ENA diagrams to incorporate ARM findings rather than using these visualizations of ARM.

In this new visualization, we show the directional associations found by ARM as arrows in an ENA diagram. In this adaptation, nodes are located in the same position as in the original ENA diagrams to facilitate the interpretation of differences between diagrams. The node sizes represent the occurrence probability of each construct, while the line width indicates the support of the rules involving these constructs. The direction of the edges points from the antecedent to the consequent, based on the rule that maximizes confidence for the specific constructs connected by each edge, thus defining the conditional direction of the associations. In cases where the association involves more than two constructs, such associations are depicted with dashed curved lines, with directionality determined by the rule with the highest confidence. Dashed lines pointing in both directions signify that both constructs are part of the antecedent or the consequent of the association. For instance, in the rule involving the constructs Guiding Feedback, Checking Understanding, and Instruction (Fig. 1), the highest confidence was found for the rule: If (Instruction AND Guiding Feedback) → Then (Checking Understanding). This rule is therefore represented by a bidirectional dashed line between Instruction and Guiding Feedback (antecedent), with additional dashed lines pointing towards Checking Understanding (consequent). In this network (and also in the corresponding ENA), only line weights (lw) or supports greater than 0.2 are included to simplify the visual comparison.

4 Results

Fig. 1 shows a comparison between the standard ENA model (left) and the EARN model (right). Table 2 compares the connections (line weights) observed using ENA and the associations identified using ARM. In this initial analysis, standard ENA and ARM offer similar insights into the data. In both analyses, the strongest connection was found between Guiding Feedback and Checking Understanding (lw=0.917, support=0.459). This association frequently occurred when the instructor provided feedback following a question asked to check students' understanding or when the tutor initially offered guidance or feedback and then posed a question to assess understanding. The second strongest connection was found between Checking Understanding and Instruction (lw=0.731, support=0.357), suggesting that these constructs also frequently co-occurred.

This is seen in the following fragment where the tutor aimed to teach students how to solve an equation with the variable in the denominator:

> *Tutor: We always multiply what's in the denominator (Guiding Feedback). So we're going to multiply that on both sides, equals two (Instruction). And now, what should we do from here now? (Checking Understanding).*

In this fragment, the tutor guided the students to multiply both sides of the equation by the denominator, then provided instructions on how to multiply, and subsequently asked for the next steps to ascertain whether the students understood the process of solving the equation once the denominator was eliminated.

However, this fragment illustrates a further point: not only does Checking Understanding co-occur with both Instruction and Guiding Feedback, but the three constructs also frequently appear together. This is captured by the association rule IF Guiding AND Instruction THEN Checking, which (due to its much higher support and confidence) indicates that the specific combination of Guiding and Instruction is more indicative of Checking than other combinations of these three constructs. This type of connection involving more than two constructs cannot be directly observed in standard

ENA. Although ENA can demonstrate that three constructs have strong pairwise connections, as shown in Fig. 1 and Table 2, it does not imply a collective strong connection among all three constructs–all it can do is show which pairwise combinations are strongest. By contrast, the support metric computed on each possible If→Then combination of these three constructs indicates a more nuanced relationship between the constructs. This result suggests that EARN can offer additional/different insights into the co-occurrences of groups of constructs.

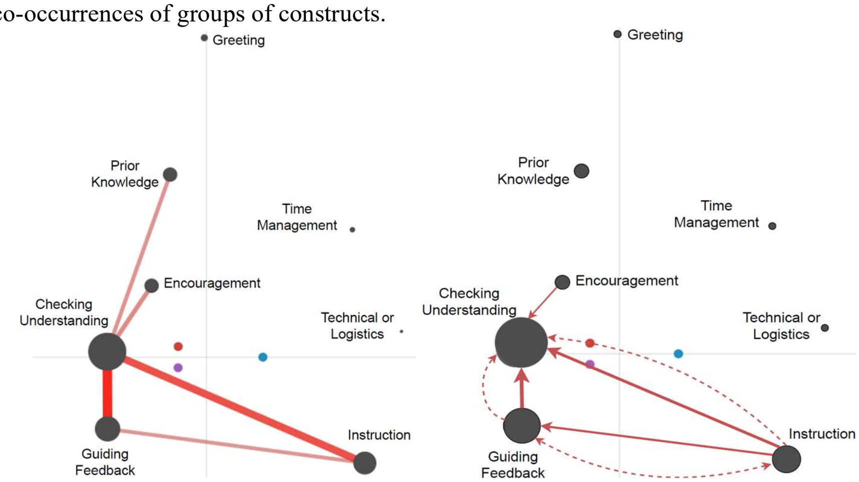

Fig. 1. Visual comparison between connections observed using standard ENA (left) and associations found using ARM (right). Only line weights or supports higher than 0.2 are shown. Dashed curved lines indicate associations that involve more than two constructs. Directions indicate the rule with the highest confidence between each group of constructs.

This ability of EARN to provide If→Then conditional associations is an important difference from standard ENA. Although the current version of ENA can display temporal order [20], an aspect also captured by SPM [21], it does not consider conditional relationships. In many scenarios, the conditional relationship may be more significant than the temporal one. For instance, in the following fragment, the tutor was providing feedback on a problem that asked the student to identify fractions that would be undefined under certain conditions:

Student: I was focused on the numerator. That's why I kept getting them wrong (Checking Understanding).
Tutor: All right. Good job. All right, fellas (Encouragement). I want y'all to try out (Instruction). Y'all doing good (Encouragement). [Student Name], the very last one is the second one (Guiding Feedback, referring that the last expression was the second one of the options that could be undefined).

Recognizing the importance of encouragement in educational settings, particularly during feedback, it is crucial to identify if the tutor is providing encouragement simultaneously with feedback, rather than merely determining whether the encouragement precedes or follows the feedback. Even if the overall frequency of encouragement and feedback within a session is low (which would result in an ENA visualization that does not highlight this connection), it may be useful to understand if tutors are particularly likely to encourage students when giving feedback (i.e. the proportion of encouragement is particularly high during feedback instances). Thus, the confidence metric used by ARM, which accounts for this conditional relationship rather than focusing on the temporal order, can provide valuable nuance to the analysis.

Table 2. Comparison between connections observed using standard ENA and associations found using ARM. Only line weights or supports higher than 0.2 are shown. The rule with the highest confidence for each group of constructs is shown in bold.

Association (IF)	Association (THEN)	Confidence (ARM)	Support (ARM)	Line Weight (ENA)
Guiding	**Checking**	**0.843**	0.459	0.917
Checking	Guiding	0.578		
Instruction	**Checking**	**0.822**	0.357	0.731
Checking	Instruction	0.449		
Encouragement	**Checking**	**0.670**	0.201	0.445
Checking	Encouragement	0.253		
Aligning	**Checking**	**0.805**	0.197	0.392
Checking	Aligning	0.248		
Instruction	**Guiding**	**0.625**	0.271	0.379
Guiding	Instruction	0.498		
Guiding AND Instruction	**Checking**	**0.848**	0.230	–
Checking	Guiding AND Instruction	0.290		
Instruction AND Checking	Guiding	0.645		
Guiding	Instruction AND Checking	0.423		
Guiding AND Checking	Instruction	0.502		
Instruction	Guiding AND Checking	0.530		

Determining these conditional relationships is crucial for identifying potential asymmetries in the co-occurrence of constructs. For example, the confidence levels in the rules identified through ARM indicate that if the tutor employed instruction or guiding feedback, there was a high likelihood that they were also checking for understanding simultaneously (confidence of 0.843 and 0.822, respectively). For instance, in the following fragment, the tutor began by checking students' understanding and, upon not receiving a direct response from the student, opted to provide more guiding feedback and instructions:

Tutor: So, we're going to divide twelve on both sides, right? (Checking Understanding). That's all we are doing (Guiding Feedback). That's going to help cancel out (Guiding Feedback). So, we divide twelve on both sides (Instruction). So now, our final answer is going to be x equals sixty-one over twelve (Guiding Feedback).

In contrast, the low confidence for the reverse rules (swapping the antecedent and the consequent) suggests that checking for understanding does not necessarily mean that the tutor was also providing guiding feedback (confidence=0.578) or instruction (confidence=0.449). An example of this can be seen in the following example (which precedes the previous fragment):

Tutor: So now, how will we get x by itself? (Checking Understanding) How will we solve for x right now? (Checking Understanding).
Student: So, we uh…
Tutor: Twelve x is the same as what operation? (Checking Understanding) Addition, multiplication, subtraction, multiplication, right? (Checking Understanding).
Student: Multiplication
Tutor: So what's the opposite of that? (Checking Understanding).

In this case, the tutor posed several questions to the student but did not offer any additional explanations. Although the confidence measure in both conditional directions for these rules is still relatively high (all above 0.4), there might be scenarios where the If→Then relationship applies distinctly in one direction only. In such cases, failing to acknowledge the specific direction of the association could be misleading, further illustrating the utility of EARN as an additional approach with potential value for QE research.

5 Discussion & Conclusion

In this article, we compare association rule mining (ARM) and epistemic network analysis (ENA) to better understand the additional contribution that association rule mining networks could provide to quantitative ethnography (QE) research. Our findings reveal that both ARM and ENA produce many overlapping findings, but that the nuanced differences between the two techniques can enhance our understanding and analysis of the data.

The first key difference between ARM and standard ENA lies in how they handle groups of constructs. Take the situation where multiple pairs of constructs all co-occur. In some cases, these pairs may actually represent bigger groups of 3 or 4 constructs that regularly co-occur; in others, single pairs can often be seen without the other constructs. This is a distinction that standard ENA does not elicit. ARM addresses this gap by identifying cases where groups of more than two constructs are highly common, thus complementing the analysis that can be done with traditional ENA.

Additionally, ARM can reveal conditional relationships between constructs that standard ENA does not distinguish. For instance, some tutors might frequently ask questions to check student understanding but provide little guidance or instructions, while other tutors might offer extensive instructions and guidance without directly checking student understanding. Although both scenarios might show similar patterns of co-occurrence in ENA, ARM can differentiate them by evaluating the conditional probabilities of these interactions. In this case, complementing standard ENA with ARM enables the development of a more nuanced and effective understanding of the relationships between codes, when relationships are asymmetric.

Note that this is distinct from the ability of Ordered Network Analysis (ONA) to identify antecedent/consequent temporal relationships; ARM identifies conditional relationships rather than temporal relationships. In other words, there are cases where the order of events is less important than simply recognizing that if one construct appears, another does too (but not vice-versa). In addition, a variant of ARM, named Sequential Pattern Mining (SPM; [21]), is widely used to identify temporal relationships while also offering the possibility of recognizing groups of more than 2 constructs. Therefore, even in those contexts where the order of constructs is more important than their conditional relations, SPM can also offer valuable additional insights.

In this study, we mainly evaluated association rules using the most commonly used metrics within that methodological community: Confidence and Support. However, these measures can be disproportionately influenced by constructs that appear very frequently in the dataset. As a result, rules or associations involving these frequently occurring constructs may monopolize the findings, potentially overshadowing other results that could be more insightful for researchers.

This risk also exists in standard ENA. As with the Support metric, standard ENA models count the occurrences of each connection, which are then normalized across each adjacency vector by dividing each number of co-occurrences by the magnitude or norm of the adjacency vector. Consequently, a construct that appears frequently overall is likely to have more connections with other constructs, and this prevalence remains even after normalization. This is the main reason for the strong similarities between the line weights of standard ENA models and the Support metric (in fact, Pearson's R is 0.959), both of which are heavily influenced by a construct's overall rate of occurrence, making it challenging to determine whether there is an actual association between constructs and sometimes potentially leading to trivial findings (e.g. overemphasizing connections between the most common constructs).

Although the QE community has tried to mitigate this issue by considering different projections, rotations, and visualizations, and particularly by continually referring back to the actual data to close the interpretative loop when using ENA, the ARM community has addressed the issue by proposing alternative metrics, typically referred to as interestingness metrics [7]. For instance, Merceron and Yacef [22] recommend calculating interestingness metrics that go beyond the commonly used Confidence and Support, specifically suggesting Cosine and Added Value (similar to Lift) for each association rule as a way to identify the most relevant rules. Cosine is similar to Support, but differs because it checks if an association exists due to the frequent occurrence of certain items together rather than just due to a high individual frequency of any of the constructs [22]. As a quick note, the metric referred to as Cosine by association rule mining users should not be confused with the Cosine Normalization used in epistemic networks [10], which, as mentioned before, has a mathematical definition closer to Support and shares similar limitations. Similarly, Added Value and Lift go beyond Confidence by assessing whether a strong association exists between two items or is merely a result of the high occurrence rate of the consequent or Then part of the association rule [22]. Bazaldua et al. [23] have also argued for other metrics, such as Jaccard, by identifying which interestingness metrics correlated with the rules that human researchers actually report as being most interesting among rules with high Support and Confidence, suggesting again the need to go beyond these two metrics. Interestingness metrics can therefore provide further valuable quantification of the association between multiple constructs and highlight different relations.

As demonstrated here, it is feasible to create representations (Epistemic Association Rule Networks; EARN) by utilizing the initial rotations and visualizations from standard ENA and then drawing edges based on ARM. Additionally, these metrics can be considered not only for modifying existing ENA model visualizations but also for generating new versions of them. For example, it is possible to replace ENA's adjacency vector with a vector defined by a chosen interestingness metric, instead of merely counting co-occurrences. This approach allows the resulting ENA visualizations to position nodes and assign line weights in a manner that can highlight connections that are not mainly explained by a high frequency of a single construct and might be more relevant for researchers.

Furthermore, EARN can be used to develop difference models, as is commonly done in standard ENA [11]. In this approach, the researcher can develop individual models using any ARM metric and then subtract them to create difference models. This is already commonly done (without the visualizations used in the QE community) by the Differential Sequence Mining algorithm [27]. In the models we propose here, nodes can be positioned according to standard ENA, while the edges could highlight the differences in two models based on the chosen interestingness metric.

In conclusion, although standard ENA offers a wide range of possibilities to analyze patterns embedded in discourse data, ARM can provide complementary analysis by providing additional nuance about the associations between constructs within the data that might be helpful for researchers. Tools like WebENA and rENA [26, 28], widely used in the QE community, could also incorporate some of the ideas discussed in this paper to offer more types of functionality to QE researchers to help them develop an even deeper understanding of their data. Moreover, researchers who usually use ENA can also explore some of the tools, metrics, and techniques commonly used by the ARM community to complement their QE toolbox. Considering both approaches can therefore help to enrich QE research and discussions.

Acknowledgments. Andres Felipe Zambrano thanks the Ministerio de Ciencia, Tecnología e Innovación and the Fulbright-Colombia commission for supporting his doctoral studies through the Fulbright-MinCiencias 2022 scholarship. We extend thanks to the Learning Engineering Virtual Institute Engagement Hub project for facilitating access to datasets from Saga Education.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Kaliisa, R., Misiejuk, K., Irgens, G.A., Misfeldt, M.: Scoping the Emerging Field of Quantitative Ethnography: Opportunities, Challenges and Future Directions. In: Advances in Quantitative Ethnography. pp. 3–17. Springer, (2021).

2. Shaffer, D.W., Ruis, A.R.: Is QE Just ENA? In: Advances in Quantitative Ethnography. pp. 71–86. Springer, (2023).
3. Arastoopour Irgens, G., Eagan, B.: The Foundations and Fundamentals of Quantitative Ethnography. In: Advances in Quantitative Ethnography. pp. 3–16. Springer, (2023).
4. Melzner, N., Greisel, M., Dresel, M., Kollar, I.: Using Process Mining (PM) and Epistemic Network Analysis (ENA) for Comparing Processes of Collaborative Problem Regulation. In: Advances in Quantitative Ethnography. pp. 154–164. Springer, (2019).
5. Zhang, S., Liu, Q., Cai, Z.: Exploring primary school teachers' technological pedagogical content knowledge (TPACK) in online collaborative discourse: An epistemic network analysis. British Journal of Educational Technology. 50, 3437–3455 (2019).
6. Agrawal, R., Imieliński, T., Swami, A.: Mining association rules between sets of items in large databases. In: International Conference on Management of Data. pp. 207–216. Association for Computing Machinery, New York, NY, USA (1993).
7. Tan, P.-N., Kumar, V., Srivastava, J.: Selecting the right interestingness measure for association patterns. In: Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 32–41. Association for Computing Machinery, New York, NY, USA (2002).
8. Chawla, S., Arunasalam, B., Davis, J.: Mining Open Source Software (OSS) Data Using Association Rules Network. In: Whang, K.-Y., Jeon, J., Shim, K., and Srivastava, J. (eds.) Advances in Knowledge Discovery and Data Mining. pp. 461–466. Springer Berlin Heidelberg, Berlin, Heidelberg (2003).
9. Pandey, G., Chawla, S., Poon, S., Arunasalam, B., Davis, J.G.: Association Rules Network: Definition and Applications. Statistical Analysis and Data Mining: The ASA Data Science Journal. 1, 260–279 (2009).
10. Shaffer, D.W., University of Wisconsin-Madison, USA, Ruis, A.R., University of Wisconsin-Madison, USA: Epistemic network analysis: A worked example of theory-based learning analytics. In: Handbook of Learning Analytics. pp. 175–187. Society for Learning Analytics Research (SoLAR) (2017).
11. 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. Journal of Learning Analytics. 3, 9–45 (2016).
12. Shaffer, D.W., Hatfield, D., Svarovsky, G.N., Nash, P., Nulty, A., Bagley, E., Frank, K., Rupp, A.A., Mislevy, R.: Epistemic Network Analysis: A Prototype for 21st-Century Assessment of Learning. International Journal of Learning and Media. 1, 33–53 (2009).
13. 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. The American Journal of Surgery. 216, 835–840 (2018).
14. Andrist, S., Collier, W., Gleicher, M., Mutlu, B., Shaffer, D.: Look together: analyzing gaze coordination with epistemic network analysis. Frontiers in Psychology. 6, (2015).

15. Sullivan, S., Warner-Hillard, C., Eagan, B., Thompson, R.J., Ruis, A.R., Haines, K., Pugh, C.M., Shaffer, D.W., Jung, H.S.: Using epistemic network analysis to identify targets for educational interventions in trauma team communication. Surgery. 163, 938–943 (2018).
16. Dubovi, I., Tabak, I.: Interactions between emotional and cognitive engagement with science on YouTube. Public Understanding of Science. 30, 759–776 (2021).
17. Zambrano, A.F., Barany, A., Ocumpaugh, J., Nasiar, N., Hutt, S., Goslen, A., Rowe, J., Lester, J., Wiebe, E., Mott, B.: Cracking the code of learning gains: Using Ordered Network Analysis to Understand the Influence of Prior Knowledge. In: International Conference on Quantitative Ethnography. pp. 18–33. Springer (2023).
18. Siebert-Evenstone, A.L., Arastoopour Irgens, G., Collier, W., Swiecki, Z., Ruis, A.R., Williamson Shaffer, D.: In Search of Conversational Grain Size: Modeling Semantic Structure using Moving Stanza Windows. Journal of Learning Analytics. 4, 123–139 (2017).
19. Bowman, D., Swiecki, Z., Cai, Z., Wang, Y., Eagan, B., Linderoth, J., Shaffer, D.W.: The mathematical foundations of epistemic network analysis. In: Advances in Quantitative Ethnography: Second International Conference, ICQE 2020, Malibu, CA, USA, February 1-3, 2021, Proceedings 2. pp. 91–105. Springer (2021).
20. Tan, Y., Ruis, A.R., Marquart, C., Cai, Z., Knowles, M.A., Shaffer, D.W.: Ordered network analysis. In: International Conference on Quantitative Ethnography. pp. 101–116. Springer, (2022).
21. Srikant, R., Agrawal, R.: Mining sequential patterns: Generalizations and performance improvements. In: Advances in Database Technology — EDBT '96. pp. 1–17. Springer Berlin Heidelberg, Berlin, Heidelberg (1996).
22. Merceron, A., Yacef, K.: Interestingness Measures for Association Rules in Educational Data. In: Educational Data Mining. pp. 57–66 (2008).
23. Bazaldua, D.A.L., Baker, R., Pedro, M.O.S.: Comparing Expert and Metric-Based Assessments of Association Rule Interestingness. In: Educational Data Mining (2014).
24. Barany, A., Nasiar, N., Porter, C., Zambrano, A.F., Andres, A., Bright, D., Choi, J., Gao, S., Giordano, C., Liu, X., Mehta, S., Shah, M., Zhang, J., Baker, R.S.: ChatGPT for Education Research: Exploring the Potential of Large Language Models for Qualitative Codebook Development. In: International Conference on Artificial Intelligence in Education. Springer, (2024).
25. Shaffer, D.W., Ruis, A.R.: How We Code. In: Advances in Quantitative Ethnography. pp. 62–77. Springer, (2021).
26. Marquart, C.L., Hinojosa, C., Swiecki, Z., Eagan, B., Shaffer, D.W.: Epistemic network analysis (Version 1.7.0), http://app.epistemicnetwork.org, (2021).
27. Kinnebrew, J.S., Loretz, K.M., Biswas, G.: A contextualized, differential sequence mining method to de-rive students' learning behavior patterns. Journal of Educational Data Mining. 5(1), 190–219 (2013).
28. Marquart, C.L., Swiecki, Z., Collier, W., Eagan, B., Woodward, R., Shaffer, D.W.: RENA: Epistemic network analysis [Software], https://CRAN.R-project.org/package=rENA, (2021).