Studying Affect Dynamics using Epistemic Networks

Shamya Karumbaiah and Ryan S Baker

University of Pennsylvania, USA
shamya@upenn.edu, ryanshaunbaker@gmail.com

Abstract. The study of how affect develops and manifests over time during learning is a popular area of research called affect dynamics. Students' affective states are recorded in authentic settings like classrooms using direct observations by culturally sensitive, trained, and certified coders. A popular approach to studying affect dynamics in the last decade involved a transition metric called the L statistic. However, recent studies have reported statistical errors and other discrepancies with L statistic leading to questions about its reliability. Thus, we turn to epistemic network analysis (ENA), an emerging technique that is gaining popularity in studying the structure of temporal interconnections between codes. In this paper, we present an alternative approach to study affect dynamics by extending ENA to include directionality in the network edges to capture transitions. We also propose a new approach to running significance tests on network edges to identify significantly likely transitions. Then, we apply the two techniques – L statistic and ENA - to a previously collected affect dataset from a middle school math class, in order to better understand the trade-offs between these methods. Our analysis revealed that ENA could be a promising new approach to conduct affect dynamics analysis. In addition to avoiding statistical errors seen in L statistic, ENA offers better visualization which better emphasizes the magnitude of a transition's strength. We discuss the assumptions in ENA that need to be vetted further and the possibility for new kinds of analysis in the future for affect dynamics research using ENA.

Keywords: Affect Dynamics, Transition Analysis, Epistemic Network Analysis, L Statistic, Temporal Sequences.

1 Introduction and Motivation

Affect within intelligent tutors and other types of adaptive and artificially intelligent educational systems has been shown to correlate with a range of other important constructs including self-efficacy [17], analytical reasoning [7], motivation [21], and learning [4, 8]. Affect-sensitive interventions have been designed to improve student engagement [22], learning gains [8, 9] and overall experience [10].

Developing effective real-time interventions depends on understanding how affect develops and manifests over time, an area of research termed *affect dynamics* (i.e. [14]), with a large body of research examining how students transition from one affective state to the next during learning activities (see [11] for review). These studies have been conducted in a wide range of contexts and demographics, including students in middle school (private), high school (public and private), undergraduate programs and graduate schools, with a particular focus on contexts in the United States and in the Philippiness.

Identifying student affect for research and to develop automated models is complex and nuanced. Affect data collection in authentic settings like classrooms typically involves direct observations by culturally sensitive, trained, and certified coders, or intensive video data collection and coding procedures [20]. Students are observed using widely-used protocols like BROMP [3] to reduce rater bias and observer effects and enable rigorous quantitative analysis. The output of such a method is temporally sequenced codes representing the field observations of student affect.

A popular transition metric used in affect dynamics research is the L statistic. It calculates the likelihood that a student in a given affective state will transition to a next state, given the base rate of the next state. In 2019, Karumbaiah and colleagues [12] provided mathematical evidence that several past studies applied the transition metric (L) incorrectly - leading to invalid conclusions of statistical significance. They proposed a corrected method which shifts the chance value of L from 0 to a positive value dependent on the number of affective states studied. Although this solution attends to the primary statistical error in past work, it makes the statistic difficult and non-intuitive to interpret. For example, it is possible for an L value to be above 0, but statistically significantly below the chance value, a situation likely to confuse researchers and readers and lead to incorrect conclusions. In addition, a recent study has reported that there are further issues with the L statistic involving states with high base rates [16], frequently seen for the affective state of engaged concentration [2]. Another simulation study with L produced results above chance levels for randomly generated affect sequences, if the sequences were short [5]. These continuing issues with the L statistic suggest an alternate approach may be warranted. We turn to the emerging field of quantitative ethnography for alternative approaches to conduct affect dynamics analysis.

A technique that is gaining popularity to study the structure of temporal interconnections between codes is epistemic network analysis (ENA). ENA models the pattern of association in coded data by building a network of relationships among the codes [23]. For affect dynamics, this means exploring associations between students' affective states during learning. There are five preliminary reasons why ENA could be a useful approach for affect dynamics research. First, the ENA network offers an intuitive way of visualizing probabilistically likely connections between affective states. Second, the edge thickness in the network (representing the strength of the association) offers a straightforward approach to interpreting the magnitudes of the transition strength – an indicator often overlooked in previous work focusing on the statistical significance of transitions. Third, unlike the L statistic, ENA can represent the case where there are multiple affective states occurring during the same time interval (seen more in the use of affect detectors as data sources than in field observation data – i.e. [19]). Fourth, ENA could enable identification of the most salient differences between the transitions observed in different student subgroups or learning activities and visualize them clearly as difference networks. Fifth, ENA could also help researchers study changes in the strength of connections over time.

There are two main limitations of using ENA as-is for affect dynamics analysis. First, we need to go beyond simple associations between the codes to also capturing the directionality of the co-occurrence. For instance, it is not enough to know that confusion and frustration has a strong association. We also need to know whether confusion transitions to frustrations or vice-versa or both. Second, we need to establish ways to conduct statistical tests on the strengths of these associations to identify significant transitions. In this paper, we present an alternative approach to study affect dynamics using ENA, also proposing ways to overcome the current limitations of the ENA tool for this type of analysis. We then compare this method of analyzing affect dynamics to the currently popular method of using the L statistic and discuss the strengths and weaknesses of the two approaches.

2 Affect Dynamics Analysis with L Statistic – Prior Work

Given an affect coding sequence, the ***L*** statistic [17] calculates the likelihood that an affective state (*prev*) will transition to a subsequent (*next*) state, given the base rate of the next state occurring.

$$L(prev \rightarrow next) = \frac{P(next|prev) - P(next)}{1 - P(next)} \quad (1)$$

The expected probability for an affective state, *P(next),* is the percentage of times that the state occurred as a next state. The conditional probability, *P(next|prev)* is given by:

$$P(next \mid prev) = \frac{Count(prev \rightarrow next)}{Count(prev)} \quad (2)$$

where Count(prev → next) is the number of times the prev state transitioned to the next state, and Count(prev) is the number of times the state in prev occurred as the previous state.

The value of ***L*** varies from -∞ to 1. The sign and the magnitude of ***L*** has been thought to be intuitively understandable as the direction and size of the association (see [8]). ***L*** = 0 has generally been treated as chance association, while ***L*** > 0 and ***L*** < 0 are treated as transitions that are more likely or less likely (respectively) than chance. To perform affect dynamics analysis across all students in an experiment, first the L value for each affect combination is calculated individually per student. Next, for each transition, a two-tailed one-sample t-test is conducted to test whether the likelihood is significantly greater than or equivalent to zero, across students [1]. More recently, researchers have added a step where a Benjamini-Hochberg post-hoc correction procedure is used to control for false positive results since the set of hypotheses involves multiple comparisons [11].

One special case that is not fully discussed in most of the literature is self-transitions where the student remains in the same affective state for more than one step in a sequence. Close to half of the previous studies have removed self-transitions during the data preparation stage (see discussion in [11]). This straightforward procedure seems quite logical, but violates the statistical assumption of independence between *prev* and *next* states as *next* state can now only take values other than that of *prev* state. Hence, when self-transitions are excluded, $\mathrm{P}(next|prev) \neq P(next)$ for transitions at chance, and for a state space with n affective states (n $>$ 2), the value of ***L*** at chance is [11]:

$$L = \frac{1}{(n-1)^2} \qquad \text{if self-transitions are excluded}$$

This finding showed that the L statistic must be interpreted differently depending on how many affective states are being observed; several past published studies which treat chance ***L*** as 0 therefore treat relationships significantly less likely than chance as significantly more likely than chance (see [12]). Although adjusting the chance value for ***L*** offers a remedy for this error, it complicates the interpretation of the statistic. While running statistical tests or making sense of the transition patterns, researchers have to be cautious about choosing the correct chance ***L*** value and interpreting values accordingly.

Additional problems have been revealed in terms of the reliability of the ***L*** statistic, based on analysis being conducted at the student-level and ignoring the within-student sample size. A recent simulation study [5] reported that ***L*** frequently produced results above chance levels for randomly generated sequences, if the sequences were short. Their study recommends a minimum sequence length of 20 per student for 4 states to avoid invalid values and much longer sequences (in excess of 50) to avoid spurious results. A sequence length of 50 would translate to a minimum observation session of 2.5 hours in a typical data collection setting with 10 students and a 20-second observation grain size. This is impractical to achieve in traditional classrooms where a class period is often under an hour long. In other words, ***L*** may be unreliable for use with the BROMP data collection where it is currently commonly used [11].

3 Affect Dynamics Analysis with ENA – An Alternate Approach

Epistemic network analysis (ENA) is a method used to identify and quantify connections among elements in a coded data [23]. An epistemic network - originally developed to model cognitive networks [23] - represents the structure of connections and the strength of association among the codes. To interpret situated events, *codes* are used as the socially and culturally organized ways of seeing these recorded actions. In authentic classroom settings, a protocol like BROMP [3] produces a sequence of data coded in terms of a validated list of codes of student affect and behavior. BROMP has been used by over 160 researchers and practitioners in seven countries for field observations with culturally sensitive coding schemes revised for and adapted to different cultural contexts [3]. In this section, we explore the use of ENA to model affect dynamics by capturing the temporal interconnections between the recorded affective codes.

3.1 Capturing Directionality

Traditionally, most ENA research has been interested in the co-occurrences of codes over time and thus focuses on symmetric data where an edge between two codes imply that they are connected to each other, with no explicit consideration of the direction of the connection. When adopting ENA for transition analysis like in the case of affect dynamics, the direction of the connection needs to be established clearly. This can be achieved by extending the approach suggested by Shaffer [23] wherein each code assumes two functions – sending and receiving. These two functions can be represented as two separate nodes in the network. For instance, the affective state of confusion (CON) can be denoted as confusion_sender (CON_S) if it is the current state in the transition and confusion_reciever (CON_R) if it is the next state in the transition. Thus, a transition from confusion to frustration (FRU) will be represented as the codes CON_S (time N) and FRU_R (time N+1). This change can be captured at the data preparation stage. For an example affect sequence of ENG (engaged concentration)-CON-CON-BOR (bored)-BOR-FRU, Table 2 shows the transformation into the vector encodings representing the 5 transitions in the sequence. Each row has 2 columns that are set to 1 – one for the sender (_S) and one for the receiver (_R). Note that the first state in the sequence is ignored as it does not represent a transition.

Table 1. Vector Encodings Preserving the Direction of the Transitions ENG-CON-CON-BOR-BOR-FRU

Affect	Transition	ENG_S	CON_S	FRU_S	BOR_S	ENG_R	CON_R	FRU_R	BOR_R
ENG	-	-	-	-	-	-	-	-	-
CON	ENG-CON	1	0	0	0	0	1	0	0
CON	CON-CON	0	1	0	0	0	1	0	0
BOR	CON-BOR	0	1	0	0	0	0	0	1
BOR	BOR-BOR	0	0	0	1	0	0	0	1
FRU	BOR-FRU	0	0	0	1	0	0	1	0

3.2 Removing Self-Transitions

The decision of whether to include or exclude self-transitions depends on the research goals. If some affective states are particularly persistent - for instance engaged concentration [2] - including self-transitions could lower the transition probabilities for transitions to new affective states, and/or in some cases, cause them to become non-significant. In contrast, excluding self-transitions may inflate the frequency of seeing transitions between affective states. Recent research on affect dynamics has focused on between states rather than their persistence [4, 6, 11]. In keeping with this, our analysis within this paper also excludes self-transitions. Self-transitions can be removed from the vector encodings by eliminating all the transitions where the sender (_S) and receiver (_R) states match, such as the third row in Table 2 (CON-CON) where

CON_S=1 and CON_R=1. The Exclusion of self-transitions could also be done at the beginning of data preparation by collapsing self-transitions in the original sequence into a single state.

3.3 Choosing ENA Parameters

Our unit of analysis is a student. All the affect codes for a student from a single observation session will constitute a conversation, in terms of ENA analysis. Since we are interested in analyzing the transitions between two states, we limit the temporal context for ENA analysis to a single vector encoding a single transition. Thus, the moving window size (grain-size of time for co-occurrence) is set to 1. The resulting epistemic network will not have any edge between the sender nodes (say CON_S and ENG_S) or between the receiver nodes (say CON_R and ENG_R). As we remove self-transitions, there also will not be any edge between the sender and receiver nodes for the same state (say ENG_S and ENG_R). An example epistemic network with three affective states (*Engaged Concentration*, *Frustration*, *Confusion*) and without self-transitions is given in Figure 1. In this example, stronger connections (thicker edges) are seen in these four (of six possible) transitions:

- ENG_S and CON_R - Engaged Concentration -> Confusion
- ENG_S and FRU_R - Engaged Concentration -> Frustration
- FRU_S and ENG_R – Frustration -> Engaged Concentration
- CON_S and ENG_R – Confusion -> Engaged Concentration

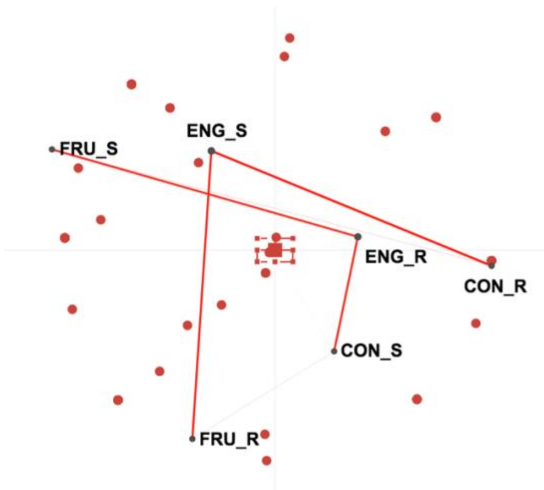

Fig. 1. An example epistemic network

3.4 Significance Test for Transition Strength

Most ENA research focuses on network level analysis with statistical tests comparing one network to another. In case of transition analysis like affect dynamics, our focus shifts from the whole network to individual transitions. Thus, it is necessary to quantify the transition strength and establish significance tests on the edges to determine which of the transitions are significant. We do so by extracting the edge weights of the resulting network for each individual student. While the output network in the web version of ENA tool presents the mean strength of the edges (Figure 1), we can extract the line weights for individual students by clicking each data point and hovering the cursor over the edges of the individual network. This could be time consuming when there is a bigger data set or a higher number of codes. Thus, we recommend extracting the network line weights for individual student networks from the *$line.weights* variable in rENA - its R implementation [15]. The output matrix will have ${}^{n}_{2}C$ columns representing all possible edges, where n is the number of affective states. For instance, if a data set has 4 affective states, there will be weights for 28 possible edges for each student. These include self-transitions (4), transitions between senders (6) and transitions between receivers (6) – all of which are invalid in our network configuration. Thus, we will run a one-sample [two-tailed] t-tests on the normalized transition strengths (line weights) only on the weights of the twelve edges that remain after removing the sixteen invalid edges. Lastly, a Benjamini-Hochberg post-hoc correction procedure is used to control for false positive results since the set of hypotheses involves multiple comparisons.

4 Example Analysis

In this section, we apply the two techniques – L statistic and ENA - to a previously collected affect dataset, in order to better understand the trade-offs between these methods.

4.1 Data

The data used in this analysis was collected through field observations of 838 students using ASSISTments, a computer based learning system for middle school math [6]. The coders used BROMP (Baker Rodrigo Ocumpaugh Monitoring Protocol; [3]) to code 3,127 observations of student affect and behavior. The observation data is highly skewed, with approximately 82% of observations coded as engaged concentration, 10% coded as boredom, 4% coded as confused, and 4% coded as frustration. This affect distribution is consistent with past research on affect prevalence in systems such as ASSISTments.

4.2 L Statistic Result

Table 3 lists the result of the affect dynamics analysis with L statistic. In this case, as there are 4 affective states, the L value at chance is 0.11 (see Table 1), and this value is used within the one-sample t-tests. The transitions that are statistically significantly likely than chance (after post-hoc correction) are *Engaged Concentration -> Boredom*, *Confusion -> Engaged Concentration*, *Frustration -> Engaged Concentration*, and

Boredom -> *Engaged Concentration*. Note that the transitions *Boredom* -> *Confusion* and *Boredom* -> *Frustration* are both significant and have an L value greater than zero but are not significantly positive because they both have L values less than the chance value.

Table 2. L Statistic Results for ASSISTment BROMP Data

Sender State	Receiver State	L	p-value	Adjusted Alpha
ENG	CON	0.171	0.091	0.0333
ENG	FRU	0.153	0.243	0.0375
ENG	**BOR**	**0.349***	**<0.001**	0.0042
CON	**ENG**	**0.539***	**<0.001**	0.0083
CON	FRU	0.049	0.286	0.0416
CON	BOR	0.085	0.683	0.0458
FRU	**ENG**	**0.552***	**<0.001**	0.0125
FRU	CON	-0.003*	0.015	0.0250
FRU	BOR	0.120	0.867	0.0500
BOR	**ENG**	**0.669***	**<0.001**	0.0167
BOR	CON	0.033*	0.022	0.0292
BOR	FRU	0.016*	0.012	0.0208

*All significant transitions
Significantly positive transitions are in bold

4.3 ENA Result

Figure 2 and Table 4 present the result of the epistemic network analysis. From the visual inspection of the resulting network (Figure 2), we can observe that the two strongest connections translate to strong transitions between *Engaged Concentration* -> *Boredom* and *Boredom* -> *Engaged Concentration*. There are four other connections that have relatively medium strength - *Engaged Concentration* -> *Confusion*, *Engaged Concentration* -> *Frustration*, *Confusion* -> *Engaged Concentration*, and *Frustration* -> *Engaged Concentration*. Other transitions do not have visibly strong edges.

While the network provides a visualization to emphasize the relative strengths of the transitions, it is also useful to quantify the transition strength to compare the degree of difference and run significance tests. Table 4 summarizes the method described in section 3.3 to achieve this with rENA. All the visible connections in the network (Figure 2) have a significantly positive transition strength. The two strongest transitions (*Engaged Concentration* -> *Boredom* and *Boredom* -> *Engaged Concentration*) have close to three times the transition strength as the other four statistically significantly positive connections.

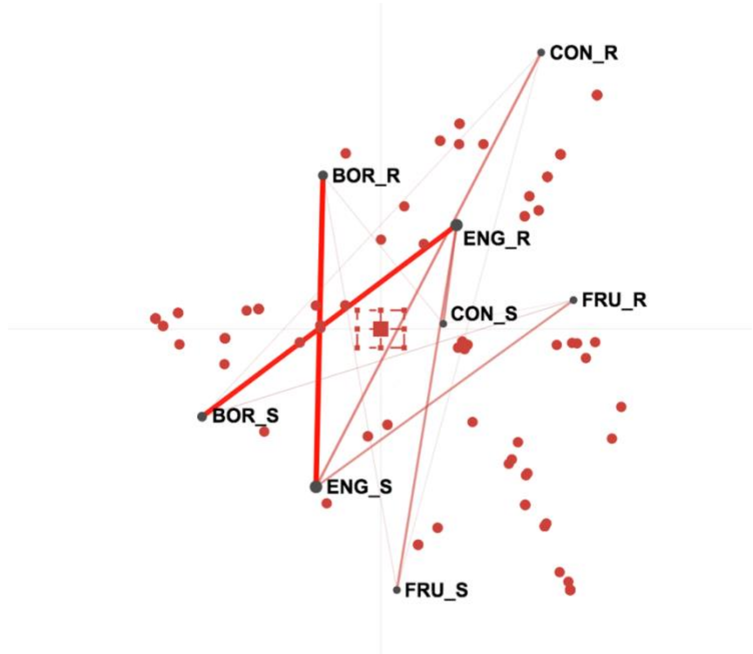

Fig. 2. Epistemic network capturing transitions in ASSISTment BROMP Data

Table 3. Significance of Transition Strengths in the Epistemic Network

Sender State	Receiver State	Transition Strength	p-value	Adjusted Alpha
ENG	**CON**	**0.099**	**<0.001***	0.0042
ENG	**FRU**	**0.083**	**<0.001***	0.0083
ENG	**BOR**	**0.256**	**<0.001***	0.0125
CON	**ENG**	**0.088**	**<0.001***	0.0167
CON	FRU	-0.027	0.001*	0.0292
CON	BOR	-0.018	0.057	0.0416
FRU	**ENG**	**0.090**	**<0.001***	0.0208
FRU	CON	-0.021	0.053	0.0375
FRU	BOR	-0.015	0.145	0.0458
BOR	**ENG**	**0.240**	**<0.001***	0.0250
BOR	CON	-0.019	0.046	0.0333
BOR	FRU	-0.002	0.907	0.0500

*All significant transitions
Significantly positive transitions are in bold

There are some similarities and differences between the results of L statistics and ENA-based analyses of significance. In both the results, all the significant transitions are either into or out of *Engaged Concentration* – the state with the highest base rate (80%). Four transitions are more likely than chance within each paradigm. In comparison to the results with L statistics, ENA finds two additional transitions to be more likely than chance: *Engaged Concentration* -> *Confusion* and *Engaged Concentration* -> *Frustration.* Both of these transitions have positive L values that are above chance (0.17 and 0.15 respectively) but are not significant. They are not among the strongest ENA results visualized in the network (Fig 2). Their quantified transition strength is less than 0.1. Unlike the L results, there is a visually clear distinction between the two strongest transitions involving *Engaged Concentration* and *Boredom* and all the other four significant transitions in the ENA results. With L statistic, there is no such obvious distinction - the significantly positive L values range from 0.35 to 0.67.

5 Comparing L Statistic and ENA Approaches to Affect Dynamics Analysis

Validity. One of the primary concerns with the L statistic is that it breaks down when self–transitions are removed from the affect sequence. In contrast, ENA does not make independence assumptions between two states, relying solely on their co-occurrences. Although ENA has not been traditionally used for transition analysis, this paper has demonstrated an extension that can serve this purpose. As such, it appears that using ENA for affect dynamics analysis is a reasonable choice, whether self-transitions are included or excluded.

Interpretability. L statistic offers a simple approach to quantify the transition likelihood, which can be used to run significance tests. However, in the case when self-transitions are excluded, the results need to be interpreted with a non-zero value for chance, requiring counterintuitive interpretation. This property has led to incorrect statistical tests being run and incorrect interpretations in past work. In contrast, ENA does not have an established approach to run significance tests on the transition strength as such. In this paper, we propose a new approach to do this – the properties of which we offer initial analysis of but which is yet to be vetted thoroughly.

Visualization. ENA offers a straightforward approach to visualize the results of an affect dynamics analysis as a network with nodes representing the affective states and the thickness of the edges representing the strength of the transition between two affective states. Previously affect dynamics analyses using the L statistic have focused more on the significance of the transitions than on the magnitude of the likelihood itself. With ENA's network visualization, it is much easier to see which transitions are most prevalent.

Co-occurring affective states. L statistic requires an affect sequence with a single state active at a given interval of time. This restricts the possibility that multiple affective states may occur during the observation interval (e.g., *Confusion* and *Frustration*), and makes it challenging to handle cases where there is a disagreement between the two coders who are observing the same student at the same time, or where an automated detector of affect is unable to distinguish which of two affective states is occurring. With ENA, it possible to work with data where multiple affective states co-occur as ENA uses vector encodings for representing transitions instead of the single state representation.

6 Discussion and Conclusions

Affect dynamics is the study of how students' affect develops and manifests over time during learning. Past affect dynamics research has analyzed the likelihood of transition between states using the ***L*** statistic. Researchers have identified many limitations to the L statistic. We explore ENA as an alternative approach to model the temporal interconnections between affective codes. In this paper, we extend ENA to include directionality in the network edges to capture transitions. The resulting network represents affect as nodes and the strength of transition as the weight of the edge between them. Our analysis with the affect observation data from a middle school math class reveals that ENA could be a promising new approach to conduct affect dynamics analysis. First, ENA appears to avoid key limitations of using ***L*** when self-transitions are removed from the affect data. Second, ENA offers better visualization which better emphasizes the magnitude of a transition's strength. Third, ENA could be used when more than one affective state is active at once.

ENA is an emerging technique, and it needs to be vetted further for its assumptions and implications for different research contexts and practices. Take, for example, the implications of highly imbalanced codes. It is common in affect datasets to have a high base rate for certain states like Engaged Concentration (82% in the data analyzed in this paper). What is striking with the use of ENA is that all the transitions into or out of Engaged Concentration are significantly likely. Further research is needed to analyze the impact of dominant codes in ENA. Mello and Gasevic [18] did some preliminary analysis on this topic and found that excluding dominant codes had drastic impacts on the resulting networks. Unfortunately, removing important codes from the analysis may not be a viable option for researchers but may present a challenge for the use of ENA for some data sets.

One of the contributions of this paper is that it demonstrates the use of ENA for transition analysis with qualitative codes. We also propose a new approach to running significance tests on network edges to identify significantly likely transitions. This has implications beyond affect dynamics. For instance, we could identify the dialog moves that are more likely to precede or follow other dialog moves in a collaborative discourse. It could also be used when the directionality between the epistemic network nodes is important. For example, in a network representing citations between authors –

author A's citation of author B is not necessarily the same as author B's citation of author A.

In our current analysis, we see that doing transition analysis with ENA produced significant results for all transitions with a positive transition strength. Further analysis is needed to confirm that the chance value of transition strength in ENA is indeed zero. Along similar lines, further research is needed to investigate if shorter affect sequences could lead to spurious results with ENA, as has been seen for the L statistic (i.e. [5. 16]).

Beyond these potential benefits, ENA opens up the possibility for new kinds of analysis in the future. Traditionally, affect dynamics research has looked at the transitions between two temporally immediately adjacent states. With ENA, we have the provision to examine co-occurrence of states at a coarser level by experimenting with moving window sizes greater than one.

Likewise, ENA offers difference networks (or subtracted networks) to enable identification of the most salient differences between two or more networks. This could be used to identify differences in transition patterns in student subgroups, such as whether there are differences between students in the US vs. the Philippines or between classroom studies and laboratory studies of affect. It could also be used to recognize students' affective trajectories by visualizing difference networks in time intervals or during important moments in learning.

Acknowledgments

We would like to thank the Penn Center for Learning Analytics (PCLA) for supporting this work and the discussions at the 2019 workshop on advanced ENA in the first international conference on quantitative ethnography (ICQE).

References

1. Baker, R.S.J.d., Rodrigo, M.M.T., Xolocotzin, U.E. (2007). The Dynamics of Affective Transitions in Simulation Problem-Solving Environments. *Proceedings of the Second International Conference on Affective Computing and Intelligent Interaction*.
2. Baker, R.S., D'Mello, S., Rodrigo, M.M.T., Graesser, A. (2010). Better to be frustrated than bored: The incidence, persistence, & impact of learners' cognitive-affective states during interactions with 3 different computer-based learning environments. *Int'l J. Hum-Comp. Stu., 68* (4),223-41.
3. Baker, R. S., Ocumpaugh, J. L., & Andres, J. M. A. L. (2020). BROMP quantitative field observations: A review. *Learning Science: Theory, Research, and Practice. New York, NY: McGraw-Hill*.
4. Bosch, N., & D'Mello, S. (2017). The Affective Experience of Novice Computer Programmers. *International Journal of Artificial Intelligence in Education*, 1-26.
5. Bosch, N. & Paquette, L. (under review). What's next? Edge cases in measuring transitions between sequential states. Submitted.
6. Botelho, A.F., Baker, R., Ocumpaugh, J., Heffernan, N. (2018) Studying Affect Dynamics and Chronometry Using Sensor-Free Detectors. *Proceedings of the 11th International Conference on Educational Data Mining*, 157-166.

7. D'Mello, S., Person, N., Lehman, B. (2009). Antecedent-Consequent Relationships and Cyclical Patterns between Affective States and Problem Solving Outcomes. *In AIED*, 57-64.
8. D'Mello, S. Graesser, A. (2012). Dynamics of Affective States during Complex Learning. *Learning and Instruction, 22*, 145-157.
9. DeFalco, J. A., Rowe, J. P., Paquette, L., Georgoulas-Sherry, V., Brawner, K., Mott, B. W., Baker, R. S., & Lester, J. C. (2018). Detecting and addressing frustration in a serious game for military training. *International Journal of Artificial Intelligence in Education*, *28*(2), 152-193.
10. Karumbaiah, S., Lizarralde, R., Allessio, D., Woolf, B. P., Arroyo, I., & Wixon, N. (2017). Addressing Student Behavior and Affect with Empathy and Growth Mindset. In *Educational Data Mining*.
11. Karumbaiah, S., Andres, J. M. A. L., Botelho, A. F., Baker, R. S., & Ocumpaugh, J. S. (2018). The Implications of a Subtle Difference in the Calculation of Affect Dynamics. *In 26th International Conference for Computers in Education.*
12. Karumbaiah, S., Baker, R. S., & Ocumpaugh, J. (2019). The Case of Self-transitions in Affective Dynamics. *In International Conference on Artificial Intelligence in Education* (pp. 172-181). Springer, Cham.
13. Karumbaiah, S., Baker, R. S., Ocumpaugh, J. & Andres, J. M. A. L. (2020). A Re-Analysis and Synthesis of Data on Affect Dynamics in Learning. Submitted.
14. Kuppens, P. (2015).It's about time: A special section on affect dynamics. *Emotion Rev., 7*(4), 297-300.
15. Marquart, C. L., Swiecki, Z., Collier, W., Eagan, B., Woodward, R., & Shaffer, D. W. (2019). rENA: epistemic network analysis.
16. Matayoshi, J. & Karumbaiah, S. (under review). Adjusting the L Statistic when Self-Transitions are Excluded in Affect Dynamics. Submitted.
17. McQuiggan, S. W. & Lester, J. (2009). Modeling affect expression and recognition in an interactive learning environment. *International Journal of Learning Technology, 4*(3-4), 216-233.
18. Mello, R. F., & Gašević, D. (2019). What is the Effect of a Dominant Code in an Epistemic Network Analysis? In *International Conference on Quantitative Ethnography* (pp. 66-76). Springer, Cham.
19. Nye, B. D., Karumbaiah, S., Tokel, S. T., Core, M. G., Stratou, G., Auerbach, D., & Georgila, K. (2018). Engaging with the scenario: Affect and facial patterns from a scenario-based intelligent tutoring system. In *International Conference on Artificial Intelligence in Education* (pp. 352-366). Springer, Cham.
20. Okur, E., Aslan, S., Alyuz, N., Esme, A. A., & Baker, R. S. (2018). Role of socio-cultural differences in labeling students' affective states. In *International Conference on Artificial Intelligence in Education* (pp. 367-380). Springer, Cham.
21. Rodrigo, M. M. T., Anglo, E., Sugay, J., Baker, R. (2008). Useofun-super- vised clustering to characterize learner behaviors and affective states while using an intelligent tutoring system. In *International Con-ference on Computers in Education*, 57-64.
22. Sanghvi, J., Castellano, G., Leite, I., Pereira, A., McOwan, P. W., & Paiva, A. (2011). Automatic analysis of affective postures and body motion to detect engagement with a game companion. *In Proceedings of the 6th International Conference on Human-robot Interaction* (pp. 305-312). ACM.
23. Shaffer, D. W., Collier, W., & Ruis, A. R. (2016). A tutorial on epistemic network analysis: Analyzing the structure of connections in cognitive, social, and interaction data. *Journal of Learning Analytics*, *3*(3), 9-45.