Identifying and Modeling Epistemic Emotions in Collaborative Problem-Solving

Sifatul Anindho
Colorado State University
sifatul.anindho@colostate.edu

Nathaniel Blanchard
Colorado State University
nathaniel.blanchard@colostate.edu

ABSTRACT

Epistemic emotions such as confusion and curiosity play a central role in learning. However, they are difficult to identify in collaborative problem-solving. Existing approaches usually rely on behavioral coding or self-reports, both of which results in contradictory empirical results in different populations and environments, raising concerns about generalization, construct validity and what constitutes the gold-standard ground truth. Here, I propose and outline my ongoing investigation of measurement and analysis of epistemic emotions in collaborative problem solving. I adopt a retrospective cued-recall framework to capture self-reported cognitive-affective states, providing a minimally disruptive and ecologically valid method of annotating internal experience. I then characterize the temporal and relational structure of these states using ordered network analysis, demonstrating how epistemic emotions exhibit cooccurrence and transition dynamics in collaborative problem-solving. Finally, I propose to systematically compare multiple representations of affect to examine how annotation methods, temporal alignment, and representational choices influence what is measured, observed, and interpreted as epistemic emotional dynamics.

Keywords

Epistemic Emotions, Collaborative Problem-Solving, Affective Computing, Ordered Network Analysis, Multimodal Modeling

1. INTRODUCTION

Epistemic emotions, such as confusion and boredom, can play a central role in learning, influencing how learner's interpret information and construct knowledge [1], and this influence can sometimes be quite complex. Baker et al. [2] showed that some of these emotions can have completely different effects on learning depending on their incidence and persistence. Confusion, for example, is an epistemic emotion often associated with impasse or difficulty. Researchers find that confusion can actually be beneficial for learning when appropriately resolved, leading to deeper understanding. But when it persists without resolution, it's shown to be detrimental in the same population groups [3]. This highlights that the impact of epistemic emotions depends not only on their presence, but also on their temporal dynamics. Therefore, developing computational models that can automatically detect and track these states in real-world learning environments is of growing interest in the EDM commmunity. Such models have the potential to power intelligent tutoring systems that provide pedagogical interventions which support productive phases of such emotions and break disruptive cycles.

Epistemic emotions often manifest differently across learning environments. Prior work has shown that both the incidence and temporal dynamics of affective states vary across learning contexts and interaction conditions, with the same emotion exhibiting different patterns and consequences depending on the task and environment [2, 3]. A relatively less explored learning environment is collaborative-problem solving (CPS). CPS is a learning paradigm when two or more individuals come together to solve problems with a shared goal [4]. In such settings, identification of epistemic emotions becomes challenging, since learning emerges through interaction, coordination, and shared reasoning among multiple individuals. In CPS, emotions are shaped by social dynamics, including negotiation, disagreement, and joint attention. As a result, these states are often subtle, context dependent, and regulated in ways that differ from individual learning scenarios. This introduces additional challenges for measurement. There is no clear gold-standard ground truth for cognitive-affective states in CPS, and internal experiences must be inferred through self-reports or external observations, both of which are inherently imperfect. In particular, social masking and regulation—where individuals suppress, exaggerate, or modulate their expressions in response to group dynamics—can lead to misalignment between observable signals and true internal states [5]. Conversely, the temporal imprecision and recall bias in self-reports complicates the alignment between subjective labels and multimodal behavioral data [6].

Here, I briefly discuss my ongoing work and outline a research roadmap for investigating how epistemic emotions in collaborative problem-solving can be more reliably measured. The central premise is that cognitive-affective states must be conceptualized as temporally dynamic, context dependent, and methodologically constructed phenomena. To

Sifatul Anindho, and Nathaniel Blanchard. Identifying and Modeling Epistemic Emotions in Collaborative Problem-Solving. In Anthony Botelho, Maria Mercedes T. Rodrigo, Adish Singla, Hiroaki Ogata, Hyojeong So, and Young Hoan Cho (eds.) Proceedings of the 19th International Conference on Educational Data Mining, Seoul, Republic of Korea, June, 2026, pp. 860–864. International Educational Data Mining Society (2026).

https://doi.org/10.5281/zenodo.21040002

this end, this work is organized into three components. First, I employ retrospective cued recall to capture self-reported cognitive-affective states, providing a non-intrusive and minimally disruptive measure of internal experience. Second, I analyze the temporal and relational structure of these states using ordered network analysis, revealing patterns of co-occurrence and transition that challenge static labeling assumptions. Third, I examine how different temporal resolutions and aggregation strategies influence the interpretation of these states, drawing on recent work on affective chronometry and dynamics. Together, these components aim to characterize what different forms of "ground truth" capture, and how methodological choices shape our understanding of cognitive-affective processes in collaborative learning.

2. THE PROBLEM AND STATE OF THE ART

Research in educational psychology has long argued that emotions are not peripheral to learning, but are tightly coupled with attention, motivation, self-regulation, and knowledge construction. Control-value theory, for example, frames achievement emotions as arising from learners' appraisals of control and value [7]. In parallel, research on epistemic emotions has emphasized that states such as confusion are especially relevant during complex sense-making because they arise from cognitive disequilibrium, uncertainty, and the effort to resolve gaps in understanding [8, 3]. Researchers have also shown that affect in learning could be modeled as dynamic, contextual, and tied to the learner's meaning-making process rather than as a static display of basic emotion [1].

Within EDM, this perspective motivated a substantial body of work on detecting learner affect from observations, interaction traces, and multimodal signals. A particularly influential line of research focused on building reliable observational infrastructures for studying affect in authentic educational settings. The Baker Rodrigo Ocumpaugh Monitoring Protocol (BROMP) was developed to support minimally disruptive field observation of student affect and behavior, and became one of the field's most widely used protocols for collecting affect labels in classrooms and digital learning environments [9]. This foregrounded a measurement problem that is easy to underestimate in machine learning: affect detectors are only as meaningful as the labels and populations on which they are built. Ocumpaugh et al. [10] demonstrated that affect detectors trained in one population or setting may not generalize well to another, even when the nominal task appears similar. Zambrano et al. [11] compared different ground-truth measures of emotion and showed that the choice of measurement approach can materially shape the conclusions drawn from affective modeling studies.

Previous research also suggests that states such as confusion, frustration, and surprise are meaningful in how they unfold and interact over time [12]. For instance, confusion may support learning when it transitions into productive inquiry, yet hinder progress when it persists or co-occurs with disengagement. However, empirical approaches that aggregate affective experiences across entire activities or treat them as independent labels have struggled to recover consistent evidence for these theoretical dynamics [13, 14]. This disconnect points to a key limitation in current modeling approaches: by collapsing temporal sequences into static representations, they discard the relational and transitional structure that defines epistemic emotions.

These concerns become sharper in collaborative problem-solving. Compared to individual tutoring or solo problem-solving, collaborative settings introduce social regulation and coordination. Learners may mask confusion to preserve competence in front of peers, suppress frustration to maintain group harmony, or display attentiveness strategically while being disengaged internally [5]. This creates a particularly difficult setting for automatic affect detection because the mapping from observable signal to felt state is neither direct nor stable [15, 7]. For this reason, collaborative learning is not simply another application domain for affect models developed elsewhere. It is a setting that stresses the assumptions built into those models.

Taken together, this literature points to three key design requirements for the present work. First, measurement must be treated as a primary research problem rather than a preprocessing step, echoing prior work on observational rigor, population validity, and the instability of ground truth in affect modeling [**?**, 10, 11]. Second, epistemic emotions should be conceptualized as dynamic and relational processes rather than independent, static categories, consistent with work on affective dynamics, confusion resolution, and state transitions [3]. Third, collaborative problem-solving should be treated as a distinct affective context, where social regulation, masking, and interactional dependencies complicate the mapping between observable signals and internal states. Building on these insights, I investigate the problem of ground truth for epistemic emotions by combining construct-aligned annotation, analysis of temporal and relational structure, and systematic examination of how different measurements of affect shape our understanding of cognitive-affective processes in collaborative learning.

3. THEORETICAL FRAMING

The theoretical framing of my doctoral research is grounded in the premise that epistemic emotions in collaborative problem solving are mutually inclusive and context-dependent processes, rather than static and independently observable states. Drawing from educational psychology, epistemic emotions arise from learners' interactions with uncertainty, task demands, and social context, and their role in learning depends on how they evolve over time and in relation to other states. From this perspective, the notion of a single, fixed ground truth for affect is inherently problematic. Instead, what is commonly treated as "ground truth" should be viewed as a methodological construct shaped by choices in measurement, temporal aggregation, and representation. To adress this, I ground the investigation in a sequence of studies that progressively examine how epistemic emotions can be measured and represented in collaborative problem-solving in a meaningful and valid way. Some of this work has already been carried, and the remaining is in plan.

3.1 Foundations (Completed Work)

Initially, I explored the space of internal states that arise in collaborative contexts. Using a retrospective cued-recall paradigm, I collected data of collaborative problem-solving sessions where participants were asked to verbalize their internal states while reviewing recordings of their group in-

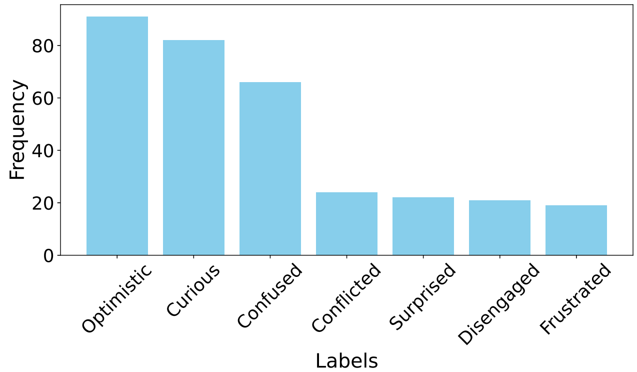

(a) Overall frequency distribution of reported affective states

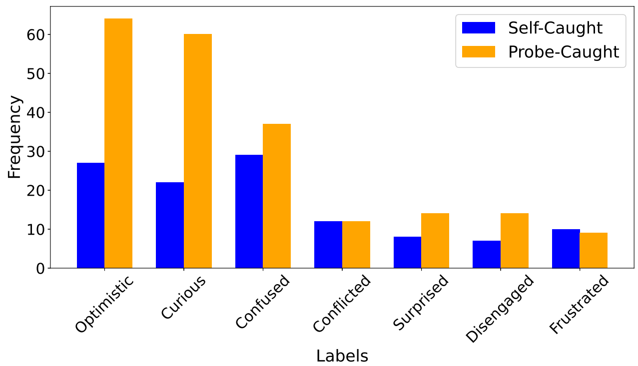

(b) Distribution across reporting mechanisms

Figure 1: Distribution of reported cognitive-affective states in our dataset

teractions. These internal monologues were transcribed and thematically analyzed, revealing a set of prevalent epistemic states—including labels such as *conflicted* and *reserved*—that are not commonly emphasized in prior work on individual learning. Analysis of n-grams and manually extracted keywords further supported the consistency of these states, and semantic similarity analysis confirmed coherence within and across identified categories. This study established both the diversity and context-specificity of epistemic emotions in collaborative settings [16].

Building on this, I developed a structured self-report framework in which participants labeled their cognitive-affective states using a more structured retrospective protocol that provided a subset of the initial labels, either spontaneously (self-caught) or in response to periodic probes (probe-caught). This enabled the collection of a more formalized dataset of retrospectively reported affective states that are temporally anchored across multiple groups. along with rich metadata including timestamps, reporting type, and participant roles. Analysis of the resulting data showed that frequency distributions of reported states were largely consistent across reporting mechanisms, with subtle but meaningful differences (Figure 1). For example, states such as *confused* were reported more frequently in self-caught conditions, while probe-caught reports tended to capture lower-arousal states, suggesting that participants may use available labels as approximations for a broader "neutral-engaged" condition. These findings highlight how measurement choices influence not only what is captured, but how affective states are interpreted. I further examined the temporal characteristics of these states by analyzing their distribution over normalized task time. Figure 2 shows a scatterplot of affective state reports over normalized collaborative task time. This analysis revealed systematic patterns: disengagement increased over time, likely reflecting fatigue or waning attention, while confusion persisted over longer durations, consistent with its characterization as a sustained epistemic state [3]. Frustration exhibited mid-task peaks, aligning with its role as an intermediate response to prolonged difficulty [1], and surprise appeared as a transient, intermittent state [2]. These results reinforce that epistemic emotions are not static events, but temporally structured processes with distinct dynamics [17].

Next, I investigated the relational structure of these epistemic emotions using ordered network analysis (ONA) (see Figure 3) [18]. This work showed that epistemic emotions form a structured system of co-occurrences and transitions that cannot be recovered from frequency-based summaries alone. In particular, I identified a stable "epistemic core" linking curiosity, optimism, and confusion, while also observing that reporting mechanisms (self-caught vs. probe-caught) and group performance (faster vs. slower groups) emphasize different aspects of this structure. For example, the roles of confusion and disengagement shift in their relationship to conflict depending on group dynamics, suggesting that the meaning of a state is contingent on its position within an interactional sequence [19].

3.2 Proposed Analysis (Future Work)

The next stage of this doctoral research centers on a systematic investigation of ground truth by comparing different annotation methods, temporal resolutions, and data modalities. Specifically, I propose a set of studies organized around three questions. First, how do different annotation methods capture epistemic emotions? I will compare self-reported states obtained through retrospective cued recall with externally observable indicators derived from multimodal behavioral data, including facial action units, vocal prosody, and conversational features. Agreement will be assessed using both categorical metrics (e.g., Cohen's κ, F1) and distributional comparisons, while also examining where and why divergence occurs. Rather than treating disagreement as error, I will analyze systematic patterns of misalignment, such as cases where internal reports indicate confusion but observable signals remain neutral, as potential evidence of social masking or regulation.

Second, how does temporal alignment influence interpretations of affect? I will investigate the temporal fidelity of different signals by comparing self-reported events with behavioral data across varying temporal windows (e.g., $\pm$1s, $\pm$5s, $\pm$10s) and aggregation strategies. This includes evaluating whether certain modalities exhibit leading or lagging relationships with reported states, and whether affective episodes are better represented as point events or temporally diffuse intervals. Building on prior work in affective chronometry, I will quantify differences in persistence, onset, and decay across representations, examining how these properties shift under different temporal assumptions.

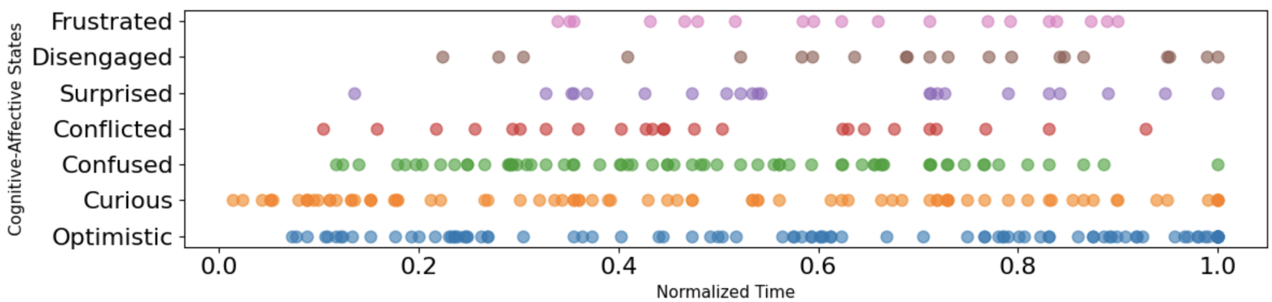

Figure 2: Distribution of affective state reports over normalized collaborative session time

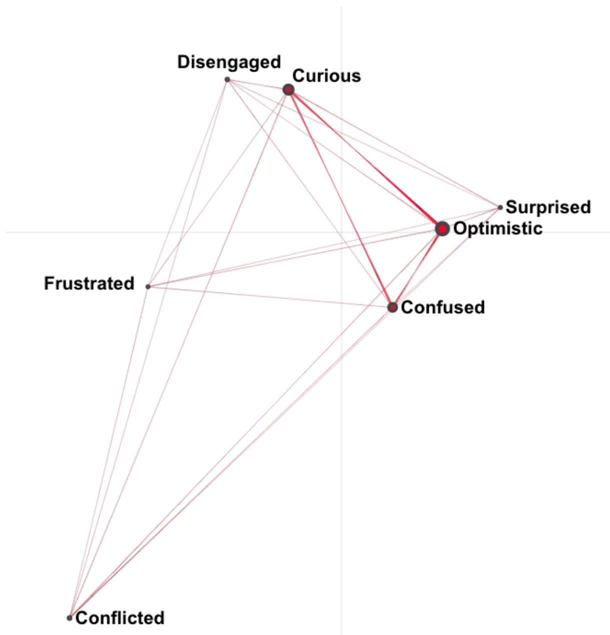

Figure 3: ONA network for affective states across all CPS groups. The first two dimensions explain 34.3% and 21.3% of the variance

Third, how do representational choices shape conclusions about affective dynamics? Using ordered network analysis and related sequence-based methods, I will construct relational representations of affect from each operationalization (e.g., raw self-reports, temporally aggregated labels, multimodal-derived signals) and compare their structural properties. This includes examining differences in cooccurrence patterns, transition probabilities, and higher-order structures such as stable "cores" or hubs of interaction. By comparing these structures across representations, I aim to identify which aspects of affective dynamics are robust to measurement choices and which are artifacts of specific annotation or aggregation strategies.

Across these studies, the goal is not to identify a single "correct" ground truth, but to characterize the space of plausible representations and the conditions under which they converge or diverge. In doing so, this work treats ground truth as a multi-faceted construct that must be interrogated empirically. The outcome is a framework for understanding how methodological decisions—ranging from annotation protocols to temporal segmentation—shape both the analysis and interpretation of cognitive-affective states in collaborative learning.

4. CONTRIBUTION AND IMPACT

This research aims to contribute to both the learning sciences and computer science communities by reframing the problem of affect modeling as a question of measurement, representation, and validity.

From a learning sciences perspective, this work advances the study of epistemic emotions in collaborative settings by emphasizing construct validity and temporal dynamics. By leveraging retrospective cued recall, it provides a minimally disruptive yet theoretically grounded approach to capturing internal states in naturalistic environments. Further, by analyzing the temporal and relational structure of these states, this work contributes to a deeper understanding of how epistemic emotions unfold during collaborative problem-solving, moving beyond frequency-based accounts toward process-oriented representations. More broadly, it contributes to ongoing discussions in EDM regarding the nature of ground truth, the role of measurement in learning analytics, and the interpretation of affective data.

From a computer science perspective, this work frames affect modeling as a data-centric machine learning problem, where the primary challenge lies in the definition and reliability of supervision signals rather than model capacity alone. By systematically examining how different annotation strategies, temporal resolutions, and representations influence downstream analysis, this research provides insights into learning under weak, subjective, and temporally misaligned supervision. These challenges are not unique to affect modeling, but arise in many real-world machine learning settings involving human-centered data. As such, this work contributes to broader efforts in machine learning to understand label uncertainty, temporal misalignment, and the relationship between data collection methods and model behavior.

Finally, this research has implications for the design of affect-aware systems in education. By clarifying what different forms of "ground truth" capture and where they fall short, this work provides a foundation for developing models and systems that are more interpretable, robust, and aligned with the constructs they aim to represent. In doing so, it can support the development of educational technologies that respond not only to observable behavior, but to the underlying cognitive-affective processes that drive learning.

5. REFERENCES

[1] Sidney D'Mello and Art Graesser. Dynamics of affective states during complex learning. *Learning and Instruction*, 22(2):145–157, 2012.

[2] Ryan SJd Baker, Sidney K D'Mello, Ma Mercedes T Rodrigo, and Arthur C Graesser. Better to be frustrated than bored: The incidence, persistence, and impact of learners' cognitive–affective states during interactions with three different computer-based learning environments. *International Journal of Human-Computer Studies*, 68(4):223–241, 2010.

[3] Sidney D'Mello, Blair Lehman, Reinhard Pekrun, and Art Graesser. Confusion can be beneficial for learning. *Learning and Instruction*, 29:153–170, 2014.

[4] Friedrich Hesse, Esther Care, Juergen Buder, Kai Sassenberg, and Patrick Griffin. A framework for teachable collaborative problem solving skills. In *Assessment and teaching of 21st century skills: Methods and approach*, pages 37–56. Springer, 2014.

[5] Marion K Underwood. Peer social status and children's understanding of the expression and control of positive and negative emotions. *Merrill-Palmer Quarterly (1982-)*, pages 610–634, 1997.

[6] Ulrich W Ebner-Priemer, Janice Kuo, Stacy Shaw Welch, Tanja Thielgen, Steffen Witte, Martin Bohus, and Marsha M Linehan. A valence-dependent group-specific recall bias of retrospective self-reports: A study of borderline personality disorder in everyday life. *The Journal of nervous and mental disease*, 194(10):774–779, 2006.

[7] Reinhard Pekrun. The control-value theory of achievement emotions: Assumptions, corollaries, and implications for educational research and practice. *Educational psychology review*, 18(4):315–341, 2006.

[8] Arthur C Graesser. Emotions are the experiential glue of learning environments in the 21st century. *Learning and Instruction*, 70:101212, 2020.

[9] Jaclyn Ocumpaugh, Ryan S. Baker, and Ma. Mercedes T. Rodrigo. Baker rodrigo ocumpaugh monitoring protocol (bromp) 2.0 technical and training manual. Technical report, Teachers College, Columbia University and Ateneo Laboratory for the Learning Sciences, New York, NY, USA and Manila, Philippines, 2015.

[10] Jaclyn Ocumpaugh, Ryan Baker, Sujith Gowda, Neil Heffernan, and Cristina Heffernan. Population validity for educational data mining models: A case study in affect detection. *British Journal of Educational Technology*, 45(3):487–501, 2014.

[11] Andres Felipe Zambrano, Nidhi Nasiar, Jaclyn Ocumpaugh, Alex Goslen, Jiayi Zhang, Jonathan Rowe, Jordan Esiason, Jessica Vandenberg, and Stephen Hutt. Says who? how different ground truth measures of emotion impact student affective modeling. *International Educational Data Mining Society*, 2024.

[12] Sidney D'Mello and Art Graesser. The half-life of cognitive-affective states during complex learning. *Cognition & Emotion*, 25(7):1299–1308, 2011.

[13] Shamya Karumbaiah, Ryan S Baker, and Jaclyn Ocumpaugh. The case of self-transitions in affective dynamics. In *International conference on artificial intelligence in education*, pages 172–181. Springer, 2019.

[14] Jaclyn Ocumpaugh, Nidhi Nasiar, Andres Felipe Zambrano, Alex Goslen, Jessica Vandenberg, Jordan Esiason, Jonathan Rowe, and Stephen Hutt. Refocusing the lens through which we view affect dynamics: The skills, difficulty, value, efficacy and time model. In *Proceedings of the 15th International Learning Analytics and Knowledge Conference*, pages 192–203, 2025.

[15] James J Gross. Emotion regulation: Current status and future prospects. *Psychological inquiry*, 26(1):1–26, 2015.

[16] Sifatul Anindho, Videep Venkatesha, Mariah Bradford, Anne M Cleary, and Nathaniel Blanchard. An exploration of internal states in collaborative problem solving. In *International Conference on Human-Computer Interaction*, pages 135–150. Springer, 2025.

[17] Sifatul Anindho, Videep Venkatesha, and Nathaniel Blanchard. A methodological framework for capturing cognitive-affective states in collaborative learning. *arXiv preprint arXiv:2507.01166*, 2025.

[18] Yuanru Tan, Andrew R Ruis, Cody Marquart, Zhiqiang Cai, Mariah A Knowles, and David Williamson Shaffer. Ordered network analysis. In *International conference on quantitative ethnography*, pages 101–116. Springer, 2022.

[19] Sifat Anindho, Videep Venkatesha, Jaclyn Ocumpaugh, and Nathaniel Blanchard. Ordered network analysis of epistemic emotions during collaborative problem solving. In *Artificial Intelligence in Education*, Lecture Notes in Computer Science, 2026. In press.