The Case of Self-Transitions in Affective Dynamics

Shamya Karumbaiah, Ryan S Baker and Jaclyn Ocumpaugh

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

Abstract. Affect dynamics, the study of how affect develops and manifests over the course of learning, has become a popular area of research in learning analytics. Despite some shared metrics and research questions, researchers in this area have some differences in how they pre-process the data for analysis [17]. Specifically, researchers differ in how they treat cases where a student remains in the same affective state in two successive observations, referred to as *self-transitions*. While most researchers include these cases in their data, D'Mello and others have argued over the last few years that these cases should be removed prior to analysis. While this choice reflects the intended focus in their research paradigm on the transitions out of an affective state, this difference in data preprocessing changes the meaning of the metric used. For around a decade, the community has used the metric *L* to evaluate the probability of transitions in affect. *L* is largely believed to have a value of 0 when a transition is at chance, and this is true for the original use of the metric. However, this paper provides mathematical evidence that this metric does not have a value of 0 at chance if self-transitions are removed. This shift is problematic because previously published statistical analyses comparing *L* values to the value at chance have used the wrong value, incorrectly producing lowered p values and in many cases reporting transitions as significantly more likely than chance when they are actually less frequent.

Keywords: Affect dynamics, L statistics, Student affect, Engagement, Self-transitions, Data preprocessing.

1 Introduction

In a data mining pipeline, data preprocessing is often considered a step separate from analysis. Data preprocessing steps like cleaning, sampling, normalization/standardization, and imputation are performed to clean and consolidate the collected data into a format ready for input into an analytical technique. The choices made during pre-processing may, in many cases, have relatively minor implications on the analysis to follow. But in some cases, a seemingly small, theoretically-justified preprocessing step can change the meaning of the metric used in the analysis. In this paper, we present one such example of a misinterpreted metric, D'Mello, Taylor, and Graesser's [2007] *L*, that was used in affect dynamics research for over ten years, in over a dozen published studies [1-5,8-11,13,14,17,19-21,23-25]. However, a closer look at the way data is preprocessed in some of these studies reveals how it changes the meaning of the metric.

We discuss the implication of this new finding on the results of these studies, some of which appear to have reported results in the wrong direction due to this shift.

As mentioned, this work occurs in the domain of affect dynamics [18] - an area of research that studies how students transition between different emotional states, in this case in a learning setting. Based on increasing evidence that student affect is associated with learning and long-term outcomes [7,22], affect has been used to understand the design of a learning environment [16] and affect-sensitive interventions have been designed and tested in some systems [12,15]. Understanding how affect manifests over time is useful when designing real-time educational interventions that work with natural patterns and transitions in affect.

Perhaps the mostly widely used metric in research on affect dynamics is D'Mello, Taylor, and Graesser's [2007] L statistic. It measures whether a transition from one affective state to another is more likely than the second state's base rate. Approximately 20 studies have used this statistic to study the transitions between different emotional states of interest [17].

During data preprocessing, one key methodological question is whether self-transitions (when a student remains in the same affective state both before and after) should be considered or excluded from calculations, with most of the studies by D'Mello and his colleagues excluding self-transitions [3-5,8-11,17] and most of the work by other research groups including them [1,2,13,14,19-21,23-25]. A recent review found that the exclusion of self-transitions leads to a higher proportion of transitions being found to be more likely than chance [17]. If valid, this result would suggest that it is beneficial to exclude self-transitions to increase statistical power. However, in one recent paper that excluded self-transitions, the researchers reported all the transitions into engaged concentration were more likely than chance [5], a mathematically impossible result. Further investigation with the original authors of this paper indicated that this result was not a typing error, raising questions about the validity and interpretation of this widely-used metric. In this paper, we extend prior work by explicitly investigating the mathematical basis of the L statistic, both when self-transitions are included and when they are excluded, to see how this impossible result was obtained and what its implications are for the use of this statistic.

2 *L* statistics and Affect Dynamics Analysis

Given an affect sequence, the L statistic [10] calculates the likelihood that an affective state (*prev*) will transition to a subsequent (*next*) state, given the base rate of the next state occurring.

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

The expected probability, *P(next)* for an affective state is the percentage of times that the state occurred as a next state. Thus, the first affective state in the sequence of a student will be excluded from this calculation since this state cannot take the role of a next state. Similarly, the calculation of the *prev* state excludes the last state in the sequence. The conditional probability, *P(next|prev)* is given by:

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

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.

There are several special cases in the calculation of *L* where there is no consensus in the literature on how to perform the calculation, and [17] has recommended the following treatment:

1. When any affective state (A_n) being considered in a given study is not present for a given student's observation period:
   a. Transitions to A_n do not occur for that student. In this case, *P(next)* = 0 and *P(next | prev)* = 0, and thus, *L* = 0.
   b. Transitions from A_n also do not occur. In this case, we do not know what affective state would have followed A_n, and thus, *L* = undefined.
2. Following from case 1, if a student remains in a single affective state (A_s) throughout an observation period, all other affective states being considered in the study behave as A_n. However, the calculations differ based on whether or not the self-transitions are included.
   a. If self-transitions are included in the analyses:
      (1) Transitions from A_s to any other affective state (e.g., A_n) do not occur, and therefore, as in 1a, *L* = 0 for any transition out of A_s.
      (2) Transitions to A_s from any other affective state (e.g., A_n) do not occur, and therefore, as in 1b, *L* = undefined.
   b. If self-transitions are discarded in the analyses, an affect sequence consisting of repeated observations of the same affective category is reduced to a single observation of that affective state. In this case, no transitions occur, and therefore *L* = undefined for all possible sequences being studied.

It is not always clear how these special cases are treated in past research. In this study, we follow [17]'s definition of *L* as outlined above.

The value of *L* varies from -∞ to 1. D'Mello and Graesser [8] state in page 7 that "the sign and the magnitude of *L* is intuitively understandable as the direction and size of the association". As has been expanded in subsequent papers [1,3-5,8,9,11,13,14,17,19-21,23-25], *L* = 0 is treated as chance, 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, as [8, pg. 7] recommends, the researcher runs "one-sample [two-tailed] t-tests to test whether likelihoods were significantly greater than or equivalent to zero (no relationship between immediate and next state)", on the sample of individual student L values for each transition. Lastly, a Benjamini-Hochberg post-hoc correction procedure is often used [1,5,17,21,23-25] to control for false positive results since the set of hypotheses involves multiple comparisons.

3 Analysis

This straightforward procedure seems quite logical, but the result seen in [5], where, after removing self-transitions, all transitions into the affective state of engaged concentration were more likely than chance, suggests that something may be wrong. As such, it may be worth examining the mathematical assumptions of this procedure. Specifically, while calculating the transition likelihood from the affective state of M_t ($prev$) to M_{t+1} (*next*), D'Mello explains that, "...if M_{t+1} and M_t are *independent* [emphasis added], then $Pr(M_{t+1}|M_t) = Pr(M_{t+1})$" [8]. However, removing self-transitions breaks the independence between M_{t+1} and $M_t$as M_{t+1}can now only take values other than Mt. Hence, when self-transitions are excluded, $Pr(M_{t+1}|M_t) \neq Pr(M_{t+1})$.

Another sign of potential problems is found in [8], when that paper draws an analogy between L statistics and Cohen's kappa, saying, "The reader may note significant similarity to Cohen's kappa for agreement between raters and indeed the likelihood metric can be justified in a similar fashion." Although this analogy seems compelling, it is worth noting that there is a striking difference between the range of values the two statistics can take. While the value of L varies from -∞ to 1 [2], the value of Cohen's kappa varies from -1 to 1 [6].

These raise the question: if a transition occurs at chance, and self-transitions are excluded, is the value of L still 0?

3.1 Understanding how removing self-transitions affect *L* Values

Differences between a calculation based on a transition pattern (L) and a calculation based on a confusion matrix (e.g., Cohen's k) mean that the chance value takes a different value for L than for Cohen's k when transitions are altered. To illustrate, let's take an example with three states, A, B, and C, which allows for a total of nine unique transitions (AA, AB, AC, BB, BA, BC, CC, CA, and CB). We will consider the hypothetical sequence, ABBCAACCBA.

First, let us consider the case where we keep self-transitions within our calculations. Our hypothetical sequences contain all the 9 possible transitions occurring each occurring exactly once. As Table 1 shows, this makes all the possible transition types equally likely (as each occurs at the frequency expected given the base rate of the next state).

Table 1. *L* statistics calculation for an example sequence of *ABBCAACCBA* when *self-transitions* are included

Transition	Count	*P(next|prev)*	*P(next)*	*L*
A->A	1	0.33	0.33	0
A->B	1	0.33	0.33	0

A->C	1	0.33	0.33	0
B->A	1	0.33	0.33	0
B->B	1	0.33	0.33	0
B->C	1	0.33	0.33	0
C->A	1	0.33	0.33	0
C->B	1	0.33	0.33	0
C->C	1	0.33	0.33	0

Now, consider the transition AB, where A is the *prev* state and B is the next state. The expected probability, *P(next)*, is *P(B_next)* i.e., the probability of occurrence of B in the next state.

$$P(next) = P(B\_next) = \frac{2}{6} = 0.33$$

The conditional probability, *P(next | prev)* , is *P(B_next | A_prev).* Note that we are not including the last instance of A as it cannot take the *prev* state in any transition. Using equation (2), we have

$$P(next \mid prev) = P(B\_next \mid A\_ prev) = \frac{1}{3} = 0.33$$

Substituting in equation (1), we get,

$$L(A \to B) = \frac{0.33 - 0.33}{1 - 0.33} = 0$$

This holds true for all the transitions. Recall that in Table 1, the conditional probability, *P(next|prev)*, is equal to the expected probability, *P(next).* Thus, when self-transitions are included, all the transition likelihoods in this example take a value of zero, in line with the claim made in [D'Mello, p.7].

Next, we consider what happens to the *L* value at chance when we omit self-transitions. If we consider the same hypothetical sequence (ABBCAACCBA), only six unique transitions remain ABCACBA. Though this sequence is different, each affective state is equally followed by all affective states. Again, consider the transition AB, where A is the *prev* state and B is the next state. The probability that B is the next state remains the same as it did when self-transitions were included.

$$P(next) = P(B\_next) = \frac{2}{6} = 0.33$$

However, the removal of A->A sequences results in value of *P(next | prev)* that is different than in the original sequence.

$$P(next \mid prev) = P(B\_next \mid A\_ prev) = \frac{1}{2} = 0.5$$

Finally, we obtain.

$$L(A \rightarrow B) = \frac{0.5 - 0.33}{1 - 0.33} = 0.25$$

Table 2. *L* statistics calculation for an example sequence of *ABBCAACCBA when self-transitions are excluded*

Transition	Count	*P(next|prev)*	*P(next)*	*L*
A->B	1	0.5	0.33	0.25
A->C	1	0.5	0.33	0.25
B->A	1	0.5	0.33	0.25
B->C	1	0.5	0.33	0.25
C->A	1	0.5	0.33	0.25
C->B	1	0.5	0.33	0.25

This value is obtained for all six possible transitions. As we see in Table 2, when all affective states allowed are equally likely as the next state, *L* = 0.25, not 0. Since self-transitions are excluded, a given state can only transition to the other two states as opposed to the three states in total. This contrasts with the claim that *P(next|prev)* = *P(next)* [D'Mello, p.7] and increases the conditional probability (i.e, *P(next|prev)*) to one out of two while the expected probability (i.e, *P(next)*) remains at two out of three. Thus, for a state space with three states, the chance value of *L* is at 0.25 instead of 0.

3.2 Redefining Chance *L* Value

We now generalize our observations above for a state space with n affective states ($n > 2$) and determine what *L* value would be expected at chance. Such a state space would have n^2 unique transitions if we include self-transitions, but only has $n^2 - n$ unique transitions if we exclude self-transitions. Thus, at chance, the expected probability is

$$P(next) = \frac{n}{n^2} = \frac{1}{n}$$ if self-transitions are included

$$P(next) = \frac{n-1}{n^2 - n} = \frac{1}{n}$$ if self-transitions are excluded

However, at chance, the conditional probability is

$$P(next \mid prev) = \frac{1}{n}$$ if self-transitions are included

$$P(next \mid prev) = \frac{1}{n-1}$$ if self-transitions are excluded

Plugging these into the original equation of L (equation 1), the value of L at chance is

$L = 0$	if self-transitions are included
$L = \frac{1}{(n-1)^2}$	if self-transitions are excluded

Generally, affect dynamics is studied in terms of the four academic emotions of confusion, frustration, boredom and engaged concentration (emotions like delight and surprise are also sometimes considered, somewhat more rarely). The otherwise unlabeled data segment in the timeline, which occurs when the primary states being investigated are not found, are sometimes given the label NA and considered in the analyses. In such a setup (n = 5), the L value at chance is L=0.0625. For the smallest reasonable state space with n = 3, the L value at chance is at its maximum, 0.25. As the number of affective states observed increases, the impact of the difference between including and excluding self-transitions decreases (Table 3).

Table 3. The value of L that represents chance, for varying state space

n	3	4	5	6	7	8
chance *L*	0.25	0.11	0.0625	0.04	0.0277	0.0204

4 Implications

The primary implication of this new finding is on the interpretation of the L value. If an affective dynamics study excludes self-transitions, the threshold to understand the direction of the transition must be set based on the number of affective states studied (see Table 3). For instance, for a study with four affective states, the transitions with L value less than 0.11 should be interpreted as being less likely than chance. Importantly, the test for significance of these transitions must set the null hypothesis at the appropriate chance levels and not zero.

This finding, thus, has implications on past published studies as well. In past studies which excluded self-transitions [3-5,8-11,17], we need to reconsider the results in terms of what the correct chance value was. Since these papers conducted hypothesis tests with $L = 0$ as the null hypothesis, they are likely to have overstated their possible effects, possibly finding positive results where negative results would have been more accurate. As such, these results need to be reanalyzed with the appropriate chance values for L (given in Table 3) to get the new significance values. For instance, in [5], the transition from boredom to frustration is reported to have an $L = 0.036$ and is significant with $p < 0.001$ – indicating that the transition from boredom to frustration is more likely than chance. But, with n = 5, the reported L value actually denotes a negative transition as the reported L value is less than the L value at chance (0.0625, as shown in Table 3).

As such, it becomes essential to rerun the t-test on the original data with the null hypothesis of L = 0.0625 to confirm if this transition is actually significantly less likely than chance.

It is important to once again note that not all past publications using L are affected by this finding. Over half of the past studies using this metric included self-transitions [1,2,13,14,19-21,23-25] and are therefore unchanged by this finding. The choice of whether or not one ought to include self-transitions in an affect dynamics analysis depends on the research goals and questions of the study. As [17] suggests, excluding self-transitions reveals a larger number of affective patterns that might otherwise be suppressed by the presence of persistent affective states. Including self-transitions in analysis helps us to better understand each state's persistence, but dilutes the transitions between different affective states. Better understanding transitions is likely important in theoretical models, but understanding true persistence might be particularly useful for algorithms being used to trigger interventions, for example.

5 Conclusion

In this paper, we demonstrate that a commonly-used metric in affect dynamics research has been incorrectly interpreted when a common pre-processing step is also taken. The past 18 studies in this area can be divided into two groups - 10 studies that includes self-transitions [1,2,13,14,19-21,23-25] and 8 that excludes self-transitions [3-5,8-11,17]. The studies that excluded self-transitions did so in order to concentrate on the transitions between states rather than on the persistence of each state [4]. While this focus can be justified, this paper demonstrates that doing so changes the interpretation of a key metric, and that the previous papers that excluded self-transitions did not account for this, invalidating many of their results.

Specifically, we find that when self-transitions are excluded, the value for L that represents chance shifts from 0 to $1/(n-1)^2$, where n is the number of affective states studied. This is because the exclusion of self-transitions leads to a violation of the assumption of independence in the equations used to calculate L. This new finding has a direct impact on the validity of the claims made by the 8 studies that excluded self-transitions as all the t-tests conducted in these studies have used $L = 0$ in their null hypothesis. As illustrated in section 4, the t-tests in these studies should be re-run and re-examined for effects that switch from significantly more likely than chance to null effects or even effects that are significantly less likely than chance.

In conclusion, this paper illustrates the impact of a seemingly subtle data preprocessing step in the interpretation of the results of an analysis. As the use of data mining and automation becomes widespread in areas like education, we need to be more cautious about the impact of all the changes we do to the data processing pipeline - however independent the stages of the pipeline may look like. In some cases, as illustrated in this paper, a simple preprocessing step could potentially imply that you are attempting to answer a different research question. It is also necessary to be mindful of the underlying reasons and assumptions behind each step in the data mining pipeline. Only by carefully considering the validity of our complete processes can we ensure that our findings are valid, and that the adaptive systems we develop using those findings are optimally effective for learners.

Acknowledgements

We would like to thank Penn Center for Learning Analytics and Nigel Bosch and Luc Paquette for comments and discussions that helped to motivate this research and our thinking about the issues in this paper.

References

1. Andres, J.M.L., & Rodrigo, M.M.T. (2014). The Incidence and persistence of affective states while playing Newton's playground. 7th IEEE International Conference on Humanoid, Nanotechnology, Information Tech., Communication and Control, Environment, and Management.
2. Baker, R.S., Rodrigo, M.M.T., Xolocotzin, U. (2007). The dynamics of affective transitions in simulation problem-solving environments. *International Conf. on Affective Computing and Intelligent Interaction*. Springer Berlin Heidelberg, 666-67.
3. Bosch, N., & D'Mello, S. (2013). Sequential patterns of affective states of novice programmers. In *The 1st Workshop on AI-supported Education for Computer Science (AIEDCS 2013)*, 1-10.
4. Bosch, N., & D'Mello, S. (2017). The Affective Experience of Novice Computer Programmers. *International Journal of Artificial Intelligence in Education*, 1-26.
5. 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.
6. Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and psychological measurement*, *20*(1), 37-46.
7. Craig, S., Graesser, A., Sullins, J., & Gholson, B. (2004). Affect and learning: an exploratory look into the role of affect in learning with AutoTutor. *Journal of educational media*, *29*(3), 241-250.
8. D'Mello, S. Graesser, A. (2012). Dynamics of Affective States during Complex Learning. *Learning and Instruction, 22*, 145-157.
9. 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.
10. D'Mello, S., Taylor, R., & Graesser, A. (2007). Monitoring Affective Trajectories during Complex Learning. D. McNamara & J. Trafton (Eds.), *Proc. 29th Annual Cognitive Science Soc.,* 203-8.
11. D'Mello, S., Graesser, A. (2010). Modeling cognitive-affective dynamics with Hidden Markov Models. *Proceedings of the 32nd Annual Cognitive Science Society*, 2721-2726.
12. 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.
13. Guia, T.F.G., Rodrigo, M.M.T., Dagami, M., Sugay, J., Macam, F., Mitrovic, A. (2013) An exploratory study of factors indicative of affective states of students using SQL-Tutor. *Research & Practice in Technology Enhanced Learning, 8*(3), 411-430.

14. Guia, T.F.G., Sugay, J., Rodrigo, M.M.T., Macam, F., Dagami, M., Mitrovic, A. (2011). Transitions of Affective States in an Intelligent Tutoring System. *Proc. Philippine Computing Soc.* 31-5.
15. 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 *EDM*.
16. Karumbaiah, S., Rahimi, S., Baker, R. S., Shute, V. J., & D'Mello, S. (2018). Is Student Frustration in Learning Games More Associated with Game Mechanics or Conceptual Understanding. In *International Conference of Learning Sciences*.
17. 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*
18. Kuppens, P. (2015). It's about time: A special section on affect dynamics. *Emotion Rev., 7*(4), 297-300
19. McQuiggan, S. W., Robison, J. L., & Lester, J. C. (2010). Affective transitions in narrative-centered learning environments. *Educational Technology & Society, 13*(1), 40-53.
20. McQuiggan, S., Robison, J., & Lester, J. (2008). Affective transitions in narrative-centered learning environments. In *Intelligent Tutoring Systems,* 490-499). Springer Berlin/Heidelberg.
21. Ocumpaugh, J., Andres, J.M., Baker, R., DeFalco, J., Paquette, L., Rowe, J., et al. (2017). Affect Dynamics in Military Trainees using vMedic: From Engaged Concentration to Boredom to Confusion. In *International Conf. on Artificial Intelligence in Ed.*, 238-249. Springer, Cham.
22. Pardos, Z. A., Baker, R. S., San Pedro, M. O., Gowda, S. M., & Gowda, S. M. (2013, April). Affective states and state tests: Investigating how affect throughout the school year predicts end of year learning outcomes. In *Proceedings of the Third International Conference on Learning Analytics and Knowledge* (pp. 117-124). ACM.
23. Rodrigo, M.M.T., Baker, R., Agapito, J., Nabos, J., Repalam, M., Reyes Jr, S., San Pedro, M.O.C. (2011). The effects of an embodied conversational agent on student affective dynamics while using an intelligent tutoring system. *IEEE Transactions on Affective Computing*, 2(4), 18-37.
24. Rodrigo, M.M.T., Baker, R., Agapito, J., Nabos, J., Repalam, M., Reyes, S., San Pedro, M.O.C. (2012). The effects of an interactive software agent on student affective dynamics while using; an intelligent tutoring system. *IEEE Transactions on Affective Computing, 3*(2), 224-36.
25. Rodrigo, M.M.T., Rebolledo-Mendez, G., Baker, R., du Boulay, B., Sugay, J.., Lim, S., Luckin, R. (2008) The effects of motivational modeling on affect in an intelligent tutoring system. *Proc. of International Conference on Computers in Education*, 57, 64.