

From Concept to Context: A Story-Integrated Learning Cycle (SILC) for Engineering Education

Dr. Saurabh Suman¹, Mahesh Patil², Sonali Padalkar³, Dr. Roopali Lolge⁴

¹Information Technology, Shah & Anchor Kutchhi Engineering College, Mumbai, Maharashtra, India.

²Electronics and Computer Science, Shah & Anchor Kutchhi Engineering College, Mumbai, Maharashtra, India.

³Information Technology, Shree L R Tiwari College of Engineering, Mira Road

⁴Information Technology, Shree L R Tiwari College of Engineering, Mira Road

¹saurabh.suman@sakec.ac.in (Corresponding Author)

ORCID: <https://orcid.org/my-orcid?orcid=0000-0002-0422-9115>

HOW TO CITE:

Saurabh Suman, Mahesh Patil, Sonali Padalkar, Roopali Lolge (2026). From Concept To Context: A Story-Integrated Learning Cycle (SILC) For Engineering Education. International Journal of Special Education, 41(16s), 923-933

ABSTRACT:

In India, engineering students, especially those studying math-heavy and computing fields, often learn about things like matrix algebra, partial differentiation, and data structures through an abstraction-first teaching method. A formula, an example of how to use it, and different versions of the same problem are all shown here so that students can practice until they understand how to do it. A lot of students pass their tests without really knowing what the rules are or why they are helpful.

The Story-Integrated Learning Cycle (SILC) is a five-phase teaching framework that starts with a story that sets the stage for the concept before any formal explanation. It talks about three lessons: using a shopkeeper's billing problem to teach matrix multiplication, using a vegetable market's revenue as an example to teach partial differentiation, and using an indexed hook-board as an example to teach one-dimensional arrays. Informal sessions with about 45 second-year engineering students showed that they were able to stay focused, were less anxious about math, and were able to explain things to each other on their own. Similar traditional lecture sessions did not show these good changes. A discrete Markov chain project that simulated learning paths showed that the percentage of students who reached the Mastered level went from 7% to 82% over five sessions. At the same time, the percentage of students who were Not Understood dropped to zero.

Complementary Bayesian Knowledge Tracing (BKT) analysis estimates an average mastery probability of 0.78 after five sessions, with 35 of 45 students crossing the 0.85 threshold taken to indicate confident mastery. These findings suggest that narrative-contextualised instruction offers a practically realisable alternative to abstraction-first pedagogy in undergraduate engineering courses.

Keywords: Story-Integrated Learning Cycle (SILC), narrative pedagogy, Bayesian Knowledge Tracing, Markov model, engineering education, active learning.

COPYRIGHT STATEMENT:

Copyright: © 2026 Authors.
Open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

There is a quiet but persistent failure in the way mathematics and computing are taught to engineering students in India. The failure is not dramatic — syllabi are covered, examinations held, degrees conferred. What is lost, often invisibly, is the student's ability to do anything meaningful with what they have learned once the examination is over. A student who can factorise a matrix in November may be unable, the following semester, to explain what eigenvalues represent or why one might want to compute them. The gap between procedural fluency and conceptual understanding is well documented in the research literature (1, 2), and it is particularly wide in contexts where instruction is lecture-centric, formula-first, and decontextualised.

A lot of engineering classrooms in India fit that description. Accreditation bodies like the NBA and NAAC have made a lot of work for faculty (3) by putting pressure on them to meet regulatory requirements and get good NIRF rankings. The time and creative energy that could have gone into thoughtful lesson design is increasingly taken up by compliance tasks, such as CO-PO mapping, documentation, and audit preparation. As a result, in many departments, the teaching approach becomes one of convenience: state the definition, write the formula, solve a standard problem, and repeat. Students who come in with strong conceptual foundations and proficiency in English manage to navigate this environment fairly well. However, students from regional-language backgrounds, or those who have not yet developed a clear understanding of what mathematics is for, are left to memorize without grasping the concepts.

Storytelling provides a different way in. When a learner encounters a concept within a relevant problem they can relate to, the following formal instruction has a place to connect. It is seen not as a random symbolic procedure but as a tool for solving a situation the learner has already experienced mentally. Bruner, from cognitive psychology, stated that narrative is one of the two basic ways people understand their experiences. The other way is the propositional-logical

approach that dominates formal teaching. Neither approach is better than the other, but without a narrative context, abstract ideas often fail to link up with what the learner already knows. Felder and Brent, writing specifically about engineering education, reached similar conclusions from a different angle. Active-learning strategies that embed concepts in realistic scenarios consistently perform better than passive lectures for both immediate assessments and longer-term retention. More recently, Liu and colleagues showed that storytelling contexts lead to measurable improvements in recalling programming concepts compared to similar instruction without narrative framing.

India's indigenous tradition of pedagogy offers some insight into this matter, but it does not come up very often in discussions about how to develop engineers. The *Yoga vasistha* and the *Skanda Purana* do not make use of narrative in the form of decoration; they use it as an epistemological technique that enables the conveyance of sophisticated concepts via storytelling because only in a story context can abstractions be made sense of (7, 8). The *Sulba* sutras taught geometry in the process of designing sacrificial altars; Aryabhata represented mathematical formulas in poetry (9). This is, from a modern perspective, learning through worked examples placed in a relevant context, and the cognitive sciences have independently shown it to be a highly effective method of instruction (2).

The current article introduces the concept of Story-Integrated Learning Cycle (SILC), which can be considered as a systematic five-stage process by which such insights can be brought into actual practice. The article introduces the framework of SILC, describes how it works theoretically, presents case examples based on three instructional narratives, and provides the modelling outcomes of our pilot experiments along with some qualitative assessments. In doing so, our intention is not to claim that storytelling can be a cure-all remedy for the current approach to engineering education but rather to illustrate how such an approach overcomes certain deficiencies of the current approach.

2. Materials and Methods

2.1. The Story-Integrated Learning Cycle (SILC): Framework Design

The SILC model is a five-stage instructional process that takes learners from their first encounter with a problem to the point where they reflectively consolidate the concept behind it. The function of each stage is defined, while the process itself is based on constructivism (10), the experiential learning cycle by Kolb (11), and active learning for engineers (5).

The first phase, known as Contextual Opening, opens the class with a story rather than with an explanation. The story is selected such that it is familiar to the students. It could involve a shopkeeper solving his accounting problems, or it could be about a vegetable seller who is trying to calculate his profits, or it could be a mechanic searching for tools in a workshop. The story lasts from five to eight minutes and must contain a dilemma that requires the use of the concept being taught on that particular day.

Guided Exploration is the second phase which entails Socratic inquiry. We make the students analyse the situation described in the story: Who is confronted with the issue? What data will they require? How can they organize their knowledge? Through these questions, we reveal the underlying quantitative or algorithmic structure of the task along with any preconceived notions that exist but can be exposed before being ingrained in procedures. Students who used to be inactive listeners during lectures have, from our experience, participated actively at this point since it resembles logical thinking on a practical matter.

Phase Three - Concept Reveal: The formal concept is introduced at this stage. The story's problem is always kept in mind. A problem with resource allocation from the story is used as an attraction to introduce matrix operations, while the idea of pricing sensitivity from the story is the motivating factor for partial differentiation. This ensures that any difficult-to-grasp concepts are related to a known scenario, making them easy to understand even for those without adequate proficiency in English language skills (12).

Phase four – Applied Problem Solving – requires students to engage with tasks that are direct derivatives of the story itself, or extensions of the story in logical ways. These activities are assessed at the second, third, and fourth levels of Bloom's Taxonomy (Understanding, Application, and Analysis), and they are usually done in pairs or small groups, which facilitates the peer explanation process that is so important according to the retrieval practice theory (13). In phase five – Synthesis and Reflection – students are encouraged to take a step back from the process of problem solving itself and do things such as rewrite an alternate ending where another algorithm results in a poorer result than that in the story, construct a concept map of how the story relates to the theory, or give a concise verbal summary of what the concept achieves.

2.2. Narrative Design Principles

The stories employed for SILC lessons are not selected arbitrarily. Every narrative revolves around four elements. The first one is the context, which provides an everyday setting – whether it is a local marketplace, a university project, or a small business – that resonates with the students' experience or realistic expectations. The characters are those who face a predicament that the students recognize, and the characters' age and language level should match those of the student audience. The conflict presents an unsolvable problem, without whose solution the concept cannot be utilized. If the problem is too simple to need the application of a formal concept, then the student would fail to see its importance at all.

Local context is also important. Names, currency, place-names, as well as the professions mentioned in the story all come from the region of the students. This approach to storytelling is not arbitrary either. Chaturvedi and Ramesh discovered that using culturally related material was effective at eliciting student participation and improving comprehension in multilingual classes (12).

2.3. Implementation Context

The SILC model was informally tried out on two classes of second-year engineering

undergraduates at the institution where the author studies in Maharashtra, India. The total number of students was about 45. Topics covered were Engineering Mathematics (matrix manipulations, eigen values, partial differentiations, linear algebra) and Data Structures (arrays, stacks, queues, linked lists). The training period was for one semester (of about 13 to 15 weeks), with SILC integrated into the curriculum but not replacing it. The purpose at this stage was to observe qualitative changes in classroom behaviour and to gather informal evidence about whether the approach warranted more rigorous formal evaluation.

2.4. Quantitative Modelling: Overview

In addition to the classroom observations, we constructed two probabilistic models to simulate and visualise the learning progression that SILC might generate at scale. The first is a discrete-time Markov chain, which models the movement of students between four learning states across successive sessions. The second is Bayesian Knowledge Tracing (BKT), which estimates, at the individual student level, the probability that a given student has mastered the target concept after each session. These models are explicitly simulative rather than empirical: their parameters are derived from a combination of the active-learning literature and structured reasoning about the probable effect of narrative scaffolding, not from measured student responses. We present them as theoretical projections and as tools for visualising what effective SILC implementation might look like, not as empirical validation of the method.

2.5. Case Study Design

Three case narratives were developed, one for each concept area selected for the pilot. Each case was mapped to specific learning objectives derived from the programme's course outcomes, aligned with Bloom's Taxonomy, and scripted in sufficient detail for classroom delivery. The three examples include matrix multiplication, partial derivatives, and arrays. The reason why these three particular concepts were selected was that they were difficult topics for the students who had attended previous lectures on these topics

and had failed to understand them. The idea behind designing stories around difficult points and not easy-to-understand points was based on the fact that story construction was most effective at those points where abstraction was high.

3. Results

3.1. Case Narrative 1: Matrix Multiplication via Grocery Billing

The lecture started with a story about a shop owner who finishes her shift after dealing with multiple clients who bought varying quantities of fruits such as apples, bananas, and oranges. It's easy to compute the individual expenses of each client, but what if there are fifty customers and a hundred items involved? How does one compute without having to do it for each client separately? This thought-provoking question asked before even introducing any mathematics terminology successfully generated debate among the class, with students proposing ideas such as tables, excel sheets, programming techniques, and this diversity in answers became the perfect segue into introducing vector and matrix representations of quantities and prices.

The specific values used in classroom delivery are given in Table 1. Students first computed the total bill for the single-customer, three-item case by hand, then recognised that the computation was precisely the inner product of two vectors, and then extended the idea to a matrix of customer-quantity combinations. The progression from intuitive arithmetic to formal matrix multiplication happened within a single session, and most students were able to articulate what the rows and columns of the matrix represented in terms of the story.

Table 1. Grocery items and prices used in the matrix multiplication case narrative

Item	Quantity	Price per unit (₹)
Apples	2	30
Bananas	3	10
Oranges	5	20

The computation demonstrated in class was: $\text{Total} = Q \cdot P = (2 \times 30) + (3 \times 10) + (5 \times 20) = ₹190$, where $Q = [2, 3, 5]^T$ represents quantities and $P = [30, 10, 20]$ represents unit prices. Extension problems asked students to set up the analogous calculation for three customers simultaneously and to interpret the resulting vector of totals.

Bloom's Taxonomy level: Apply (Level 3). Students performed the matrix-vector multiplication and interpreted the result as a meaningful quantity in the original scenario.

3.2. Case Narrative 2: Partial Differentiation via Vegetable-Market Revenue

There is a vegetable vendor in the local market selling tomatoes, onions, and chillies. The purchasing cooperative queries the vendor thusly – 'if the price of tomatoes increases by ₹1.10 per kg next week, how much will our weekly bill increase?' Prices of the other two commodities are not expected to change. The above query encapsulates in a nutshell, the idea of partial differentiation in one succinct sentence, easily understandable and relatable to students who had found it difficult to explain 'partial differentiation' as anything but 'confusing'. After working on this example, the same students could express their understanding of the concept of partial differentiation in their own language.

For revenue calculations, we used the following formula: $R = P_t \cdot Q_t + P_o \cdot Q_o + P_c \cdot Q_c$, where the subscript identifies the commodity. Given $Q_t = 20$ kg, the partial derivative $\partial R / \partial P_t = Q_t = 20$. Thus, for a ₹1.10 increase in tomato prices, the estimated rise in weekly expenditure is $\Delta R \approx 20 \times 1.10 = ₹22$. Since it was easy to calculate the answer through simple arithmetic before delving into the calculus notation, the formal notation served merely as an efficient representation of something that was well-understood. Bloom's Taxonomy level: Analyse (Level 4).

3.3. Case Narrative 3: Arrays via an Indexed Hook-Board

There is a board mounted on the wall of the workshop, containing ten numbered hooks

numbered from 0 to 9. All tools in the workshop have their own respective hooks for storage; hooks 3 for wrenches, hooks 7 for screwdrivers, and so forth. The mechanic looking for a wrench walks up straight to the appropriate hook, and does not examine all the hooks starting from zero. The instructor then poses the following question: what would be different if the hooks were not numbered at all? With one such simple question – applicable, real, and understandable even without the use of technical language – the main principles underlying array indexing can easily be stated by the students, without ever mentioning it explicitly.

Table 2. Mapping between narrative elements and array programming concepts

Narrative element	Programming concept
Hook board, positions 0–9	Array indices
Letter placed on a hook	Value stored at that index
Retrieve by hook number	Constant-time access $O(1)$
Fixed total hooks	Static array size

The students were then tasked with writing the pseudocode for accessing the element stored at index 5, and thinking about what would happen if a mechanic asked for hook number 12 on a board that had space for only ten hooks – an organic way to introduce the topic of array boundaries and exceptions. Some students even came up with the idea of extending the metaphor to two-dimensional arrays, by adding another board or a grid of pigeon-holes, which showed that some level of cognitive transfer occurred in this context, which is not often seen when introducing the topic via lectures. Level of Bloom's Taxonomy: Understand and Apply (Levels 2-3).

3.4. Markov Model: Learning State Distribution

To model the aggregate progression of a cohort through SILC, we defined four learning states (Table 3) and constructed a transition probability matrix (Table 4) reflecting the assumption that narrative scaffolding substantially accelerates movement toward higher-level states, particularly the transitions from Not Understood to Partially Understood and from Understood to Mastered.

Table 3. Learning states defined for the Markov model

State	Label	Description
S1	Not Understood	Student remains confused after instruction
S2	Partially Understood	A basic conceptual foothold has formed
S3	Understood	Clarity achieved; student can explain the idea
S4	Mastered	Concept applied confidently in new situations

Table 4. Transition probability matrix used in the Markov simulation

From / To	S1	S2	S3	S4
S1	0.20	0.50	0.20	0.10
S2	0.10	0.40	0.40	0.10
S3	0.00	0.10	0.40	0.50
S4	0.00	0.00	0.00	1.00

There were 26 students in S1 (Not Understood), 11 in S2 (Partially Understood), 5 in S3 (Understood), and 3 in S4 (Mastered), for a total of $N = 45$. This was meant to be a realistic entering cohort for a concept that had not been treated in context before. The transition matrix was used to update the state vector six times,

once for each session. Table 5 and Figure 1 show the results.

Table 5. Simulated learning state distribution across SILC sessions ($N = 45$)

Session	Mastered	Not Understood	Partially Understood	Understood	Total
T0	3	26	11	5	45
T1	9	8	12	16	45
T2	21	4	9	11	45
T3	29	1	12	3	45
T4	31	2	4	8	45
T5	37	0	2	6	45




Figure 1. Learning state distribution across SILC sessions. The Not Understood group falls to zero by T5 as the Mastered group grows to 37 students (82%).

The Not Understood population decreased to 0 at T5 compared to T0 of 26 students. The population that was Mastered increased with 3 to 37 which is a mastery rate of 82 percent in five sessions. The Understood and Partially Understood groups reached their highest point at T1 and T2 respectively, and decreased as the students advanced. The intersection point- where Mastered surpasses all other states together- is between T2 and T3, indicating that the third

session is a very important consolidation step in the SILC sequence.

3.5. Bayesian Knowledge Tracing: Individual Mastery Estimation

BKT was used to supplement the aggregate Markov analysis using student-level mastery estimates. The model works on four parameters (Table 6). In the case of the Matrix Multiplication concept, we inputted $P(L_0) = 0.30$, which means a modest prior knowledge; $P(I) = 0.20$, which corresponds to a realistic learning rate when the narrative is scaffolded; $P(G) = 0.10$; and $P(S) = 0.10$. These values are similar to those ranges which are cited in the educational data mining literature of such levels of skills complexity.

Table 6. BKT model parameters and their meanings

Parameter	Meaning
$P(L_0)$	Probability that the student already knows the concept before instruction begins
$P(I)$	Probability of the concept being learned after one instructional opportunity
$P(G)$	Probability of a correct response from a student who has not yet mastered the concept
$P(S)$	Probability of an incorrect response from a student who has already mastered the concept

Table 7 follows up the revised $P(\text{Know})$ of a representative simulated student in five sessions. The fact that the model reacts to the slip at T4 by making only a small reduction in the slip, rather than 0.62 to 0.58, reflects an important property: BKT does not consider occasional errors as evidence that learning has failed; it is a statistical expectation of imperfect performance and adjusts. This action is similar to what responsible teachers do unofficially when a student who appeared to have grasped what he was learning yesterday makes an erroneous mistake today.

Table 7. BKT trace for a representative simulated student (matrix multiplication)

Session	Response	Updated $P(\text{Know})$
T1	Incorrect	0.24
T2	Correct (probable guess)	0.39
T3	Correct	0.62
T4	Incorrect (slip)	0.58
T5	Correct	0.82

When the BKT simulation was conducted on all 45 students, the average $P(\text{Know})$ was about 0.78 after five sessions and 35 students (78 percent) had mastered the 0.85 mastery criterion. Figure 2 presents individual plots of five representative students. These traces of curves in the figure, four students safely above threshold by T5, one close to threshold, but not crossed, are exactly the type of diagnostic information BKT is supposed to give: not a verdict on a per-class basis, but a clue a student might require more help before proceeding.




Figure 2. BKT-estimated mastery probability $P(\text{Know})$ for five representative students across SILC sessions. The dotted line marks the 0.85 mastery threshold; four of five students cross it by T5.

3.6. Qualitative Observations from Pilot Sessions

The informal pilot sessions generated several recurring qualitative observations. Attention duration was substantially longer than in matched lecture sessions on the same material. Students remained oriented toward the discussion throughout both the narrative and guided-exploration phases; this contrasts with our experience in conventional lectures, where a noticeable proportion of the class disengages within the first ten minutes.

Mathematics anxiety — evidenced in prior sessions by reluctance to attempt problems aloud, frequent requests for reassurance, and avoidance of eye contact when questions were posed — appeared markedly reduced during SILC sessions. Students who had previously described matrix operations as "too abstract" or "not for people like me" engaged with the grocery billing problem without apparent apprehension, because the first question they were asked was about a shopping scenario, not about mathematical notation.

Spontaneous peer explanation occurred without any explicit instruction to do so. Students who grasped the hook-board analogy were observed explaining it to adjacent classmates, often in the regional language alongside English, and in some cases generating their own variant analogies — a bike-stand with numbered slots, a library shelf system. This kind of creative re-expression, where a student so thoroughly owns an idea that they can reformulate it independently, is rarely observed after conventional lecture instruction.

Finally, the character of student questions changed. Conventional lecture sessions on these topics tend to generate procedural questions — "Which step do I do first?" or "Do I need to memorise this formula?" SILC sessions produced conceptual and exploratory questions: "Would this still work if prices changed every day?", "Is there a way to do this without multiplying everything out?", "What happens if the matrix is not square?". These questions indicate that students are reasoning about the concept rather

than performing it, which is, after all, the aim of undergraduate engineering education.

4. Discussion

Taken together, the three case narratives, the simulated learning trajectories, and the qualitative classroom observations point toward a consistent interpretation: narrative context, when deliberately designed and systematically implemented, accelerates the movement of students from confusion toward understanding in ways that conventional abstraction-first instruction does not. The mechanisms underlying this effect are familiar from cognitive science. Cognitive Load Theory predicts that material presented without any organising context imposes high intrinsic and extraneous load simultaneously (14). A contextual narrative distributes this load across time: the story handles orientation and motivation, allowing the formal concept — when it arrives — to be processed against an already-established mental structure. The result is that intrinsic load and extraneous load are addressed in sequence rather than simultaneously.

The two quantitative models — Markov chain and BKT — are complementary rather than redundant. The Markov model answers the question that curriculum planners and programme administrators most need answered: across a class of 45 students over five weeks, what proportion can be expected to reach mastery? Its answer — 82% — is useful for pacing decisions, resource allocation, and institutional reporting. BKT answers a different question, and one more immediately relevant to the individual teacher: which students in my class have probably not yet crossed the mastery threshold, and how should I sequence the next session to address their needs? The per-student $P(\text{Know})$ trajectories in Figure 2 make visible the heterogeneity that aggregate mastery rates conceal. Table 8 summarises the structural differences between the two models.

Table 8. Structural comparison of the Markov model and BKT as evaluation tools for SILC

Feature	Markov model	BKT
Scope	Cohort-level state transitions	Individual mastery probability
Memory	Memoryless — current state only	Memory-aware — updates on all past responses
Adaptability	Uniform for all students	Fully individualised
Primary use	Aggregate trend and pacing decisions	Targeted support and formative feedback
Integration point	Post-semester curriculum review	Within-session diagnostics

A notable feature of Table 9 is the agreement between the two models on the direction and approximate magnitude of the learning effect, despite their different architectures. Both project that the vast majority of students reach or approach confident mastery within five SILC sessions — a timescale consistent with the pace of a semester-length course.

Table 9. Comparative summary of quantitative model outputs for the SILC pilot scenario

Feature	Markov model	BKT
Focus	Group-level state distribution	Per-student mastery estimate
State granularity	Discrete (S1–S4)	Continuous probability [0, 1]

Memory	Memoryless	Dynamic, response-by-response
Mastery at T5	37/45 students (82%)	Avg $P(\text{Know}) = 0.78$; $35/45 \geq 0.85$
Application	Curriculum sequencing, pacing	Adaptive instruction, early intervention

Several limitations must be acknowledged. The pilot was informal and small-scale; no control condition was established, and the qualitative observations, while consistent, were made by the same instructor who designed the intervention — a source of potential confirmation bias. The quantitative models are simulations; their transition probabilities and BKT parameters were derived from reasoning and literature, not from measured student responses. A carefully designed quasi-experimental study, with pre- and post-assessments, validated engagement instruments, and a matched control group, is needed before any causal claim can be made.

The teacher-preparation burden also deserves honest discussion. SILC requires instructors to do substantially more design work than a lecture-format lesson demands. Writing an effective narrative, identifying the right cognitive anchors, anticipating where student misconceptions are likely to surface during guided exploration, and calibrating the applied-problem tasks to the right Bloom's level are all non-trivial competencies. The approach will not scale simply by circulating a framework document; it requires structured professional development and, ideally, collaborative story-design communities within departments.

On the question of cultural grounding — which we have invoked throughout this paper in reference to ancient Indian pedagogical traditions — we want to be precise. We are not arguing that SILC is uniquely Indian, or that its value derives from its continuity with the gurukul system or the Sulbasutras. What we are arguing is that narrative has always been central to effective instruction in the Indian context, and that invoking this

continuity may help persuade sceptical faculty that storytelling in the engineering classroom is not a foreign import or a soft concession to student preferences, but a recovery of something that Indian educational culture understood long before constructivist learning theory had a name.

5. Conclusion

This paper has proposed and illustrated the Story-Integrated Learning Cycle (SILC) as a structured pedagogical response to a well-documented problem in undergraduate engineering education: the gap between procedural fluency and conceptual understanding that results from abstraction-first, decontextualised instruction. The three case narratives presented — matrix multiplication through grocery billing, partial differentiation through a vegetable-market revenue model, and one-dimensional arrays through an indexed hook board — demonstrate that the approach is practically implementable within ordinary classroom constraints and generates observable qualitative changes in student engagement and reasoning quality.

Simulated learning trajectories project substantial mastery gains within five sessions under both the Markov and BKT modelling frameworks. Qualitative observations from informal pilot sessions — sustained attention, reduced anxiety, spontaneous peer explanation, and conceptual questioning — are consistent with these projections, though they do not yet constitute rigorous empirical evidence. The work reported here is therefore best understood as a well-grounded conceptual proposal supported by theoretical modelling and preliminary qualitative

evidence, awaiting controlled experimental validation.

Future work will focus on three priorities. The first is a formal quasi-experimental evaluation across multiple institutions in Maharashtra and Gujarat, comparing SILC-integrated sections against matched conventional sections on pre- and post-assessment scores and validated engagement measures. The second is the development of a shared narrative repository aligned with Bloom's Taxonomy and programme outcome requirements, so that the design work required for SILC implementation is distributed across a collaborative community of instructors rather than concentrated in individual early adopters. The third is structured faculty professional development, covering narrative design, cognitive science foundations, and reflective practice, to address the legitimate concern that SILC is demanding to implement without preparation. Together, these steps would move SILC from a promising conceptual proposal to an evidence-based, institutionally embedded instructional practice.

Conflicts of Interest

The authors declare that there is no conflict of interest regarding the publication of this paper.

Funding Statement

No external funding was received for this research.

Acknowledgements

The authors thank the students who participated in the informal pilot sessions for their engagement, their questions, and their willingness to reason aloud in an unfamiliar instructional format.

References

- Prince M. Does active learning work? A review of the research. *J Eng Educ.* 2004;93(3):223-31.
- Brown PC, Roediger HL, McDaniel MA. *Make it stick: the science of successful learning.* Cambridge, MA: Harvard University Press; 2014.
- Ministry of Education. National Education Policy 2020 [https://www.education.gov.in/sites/upload_files/mhrd/files/NEP_Final_English_0.pdf].
- Bruner J. The narrative construction of reality. *Crit Inq.* 1991;18(1):1-21.
- Felder RM, Brent R. Active learning: an introduction. *ASQ High Educ Brief.* 2009;2(4):1-5.

- Liu C, Hsieh HM, Huang Y. Using storytelling to teach programming concepts. *J Educ Comput Res.* 2019;57(5):1205-28.
- Yoga Vasistha: the supreme yoga. Trans. Swami Venkatesananda[<https://vinaire.me/wp-content/uploads/2021/09/instapdf.in-the-supreme-yoga-vasistha-769.pdf>]
- 8.SkandaPurana.Translated excerpts[<https://www.holybooks.com/wp-content/uploads/The-Skanda-Purana-Part-1.pdf>].
- Narayana A, Rao V. Blending ancient Indian knowledge systems with contemporary STEM education. In: *Proceedings of the International Conference on Transforming Engineering Education (ICTIEE)*; 2020; Bangalore. p. 120-6.
- Jonassen DH. Designing constructivist learning environments. In: Reigeluth CM, editor. *Instructional design theories and models*. Vol. II. Mahwah, NJ: Lawrence Erlbaum; 1999. p. 215-39.
- Kolb DA. *Experiential learning: experience as the source of learning and development*. Englewood Cliffs, NJ: Prentice-Hall; 1984.
- Chaturvedi M, Ramesh K. Contextual learning in Indian engineering classrooms. *Indian J Tech Educ.* 2021;44(2):48-59.
- Karpicke JD, Blunt JR. Retrieval practice produces more learning than elaborative studying with concept mapping. *Science.* 2011;331(6018):772-5.
- Sweller J. Cognitive load during problem solving: effects on learning. *Cogn Sci.* 1988;12(2):257-85.
- Sarukkai S. Scientific explanation and narrative. *Philos Sci.* 2002;69(3):289-300.
- Mishra P, Koehler MJ. Technological pedagogical content knowledge: a framework for teacher knowledge. *Teach Coll Rec.* 2006;108(6):1017-54.