Article

Volume 12 (2025), pp. 75-86
American Journal of STEM Education: Issues and Perspectives

https://doi.org/10.32674/n4x4jp52

Creating AI-Driven Adaptive Systems in Calculus: Education Through Dynamical Modeling

Ashutosh Sharma
M.S Ramaiah Institute of Technology, Bangalore, India

Johnson Okebulu Kalu
Fisher College, Boston, USA

ABSTRACT

This research presents a mathematical framework for AI-driven adaptive learning systems (ALS) in calculus education, combining nonlinear dynamical systems, reinforcement learning (RL), and fairness-aware optimization. The author derives a delay differential equations (DDEs) system to model knowledge retention and prove global stability via Lyapunov functions. The ALS employs a partially observable Markov decision process (POMDP) to optimize instructional policies, with a fairness penalty term minimizing demographic disparities. Empirical validation involves a year-long study (N=450 students), showing a 28% increase in mastery rates (p < 0.001, alpha=0.01) and a 63% reduction in equity gaps. Theoretical contributions include a bifurcation analysis of the DDE system and a proof of regret bounds for the RL algorithm. The work advances the integration of control theory and AI in mathematics education.

Keywords: Adaptive learning, Delay Differential Equations, Fairness in AI, Lyapunov stability, POMDP

INTRODUCTION

Calculus remains a critical gatekeeper for STEM careers, yet persistent challenges in conceptual mastery disproportionately affect marginalized students (Bressoud et al., 2015). The challenge of equitable calculus mastery persists due to heterogeneous prior knowledge and instructional rigidity (Koedinger et al., 2015). While AI-driven systems show promise (Holmes et al., 2019), their mathematical foundations remain underdeveloped. This paper addresses two gaps:

Mathematical Rigor: Prior ALS lack formal stability guarantees, risking chaotic adaptation (Wiggins, 2003).

Equity: Most systems optimize for average performance, exacerbating disparities (O'Neil, 2016).

A traditional one-size-fits-all curricula fail to address individual learning trajectories, necessitating AI-driven personalization. This paper introduces a novel adaptive learning system (ALS) that (1) Dynamically adjusts problem difficulty using real-time performance data, (2) Predicts misconceptions via neural networks trained on error patterns, and (3) Optimizes knowledge retention through spaced repetition modeled via delay differential equations.

This work is grounded in principles of Cognitive Load Theory (Sweller, 1988) and Constructivist Learning Theory (Vygotsky, 1978). The dynamic adjustment of problem difficulty aligns with managing intrinsic cognitive load, while the spaced repetition mechanism mirrors constructivist scaffolding by reinforcing prior knowledge. By embedding these theories into the mathematical framework, the ALS ensures alignment with evidence-based pedagogical practices.

Theoretical contributions include a stability analysis of the ALS as a nonlinear dynamical system, while empirical results validate its efficacy in diverse classrooms.

We propose an ALS governed by:

$$\frac{dK_i}{dt} = \alpha_i I_i(t-\tau) - \beta_i K_i(t) + \sum_{j=1}^{n} \gamma_{ij} \; K_j(t)\, H\big(K_j(t) - \theta_j\big)$$

Subject to $min_\pi \, E\left[\sum_{g\in G}\left|R_g(\pi) - R_{avg}(\pi)\right|\right] \le \epsilon$
where (K_i) is knowledge in component in i, τ is instructional delay, and R_g is reward for demographic group g.

Several studies have explored adaptive learning in STEM domains. Corbett and Anderson (1994) introduced the concept of knowledge tracing, a foundation for modern learning systems. More recently, VanLehn (2011) reviewed the effectiveness of intelligent tutoring systems, reporting substantial gains over traditional methods.

AI-based personalization has been explored by Piech et al. (2015) using recurrent neural networks, while Wang et al. (2021) employed reinforcement learning to optimize student engagement. However, most lack formal dynamical models, as noted by Chiang and Ma (2020), limiting predictability and stability.

Equity-focused learning systems have also emerged. Binns (2018) and Holstein et al. (2019) critique AI bias in education, advocating for fairness constraints, aligning with our inclusion of demographic reward penalties.

While prior works validate adaptive learning's promise, our contribution uniquely integrates control theory (DDEs and Lyapunov stability) and fairness-aware reinforcement learning in a unified framework. This mathematically principled approach ensures both stability and equity—an underexplored synergy in the literature.

THEORETICAL FRAMEWORK

Knowledge Dynamics as a Switched System

Let $\boldsymbol{K}(t) = [K_1(t), \dots, K_n(t)]^T$ represent knowledge states. The system switches between m instructional modes, yielding:

$$\dot{\boldsymbol{K}} = A_{\sigma \boldsymbol{K}} + B_{\sigma \boldsymbol{I}}(t - \tau) + \boldsymbol{\Gamma}(\boldsymbol{K})$$

where $\sigma(t) \in \{1, \dots, m\}$is the switching signal, $A_\sigma$is a Metzler matrix (ensuring positivity), and $\boldsymbol{\Gamma}$ encapsulates nonlinear reinforcement:

$$\Gamma_i(\boldsymbol{K}) = \sum_{j=1}^{n} \gamma_{ij} K_j H\left(\frac{\partial^2 f}{\partial K_j^2}\right) \quad \text{(Hessian-based activation)}$$

Theorem 1 (Global Exponential Stability):

If $\exists P \succ 0$ and $\mu > 0$ such that $\forall \sigma$

$$A_\sigma^T P + P A_\sigma + \mu P \prec 0,$$

then $\left||\boldsymbol{K}(t)|\right| \leq M e^{-\mu t/2} \left||\boldsymbol{K}(0)|\right|$

Proof:

Assume the switched system $\dot{\boldsymbol{K}} = A_{\sigma\boldsymbol{K}} + B_{\sigma\boldsymbol{I}}(t-\tau) + \boldsymbol{\Gamma}(\boldsymbol{K})). Let(V(\boldsymbol{K}) = \boldsymbol{K}^{\boldsymbol{T}} P \boldsymbol{K}$ be a Lyapunov function candidate. Taking the derivative:

$$\dot{V} = \boldsymbol{K}^{\boldsymbol{T}}(A_\sigma^T P + P A_\sigma)\boldsymbol{K} + 2\boldsymbol{K}^{\boldsymbol{T}} P B_{\sigma\boldsymbol{I}} + 2\boldsymbol{K}^{\boldsymbol{T}} P \boldsymbol{\Gamma}(\boldsymbol{K}).$$

Under the condition $A_\sigma^T P + P A_\sigma \prec -\mu P, and\ assuming \left||B_\sigma|\right| \leq b, \left||\boldsymbol{\Gamma}|\right| \leq c\left||\boldsymbol{K}|\right|$, we apply Young's inequality:

$$2\boldsymbol{K}^{\boldsymbol{T}} P B_{\sigma\boldsymbol{I}} \leq \epsilon\left||\boldsymbol{K}|\right|^2 + \frac{\left||PB_\sigma|\right|^2}{\epsilon}\left||\boldsymbol{I}|\right|^2.$$

Choosing ($\epsilon = \mu/4), (c < \mu/4), and(\left||\boldsymbol{I}|\right| \leq M$), we obtain:

$$\dot{V} \leq -\frac{\mu}{2}V + \frac{\left||PB_\sigma|\right|^2 M^2}{\epsilon}.$$

By the comparison lemma (Khalil, 2002), $V(t) \leq V(0)e^{-\mu t/2} + \frac{2\left||PB_\sigma|\right|^2 M^2}{\mu\epsilon}$, proving exponential convergence to a bounded set.

Fairness-Constrained POMDP

The ALS policy π solves:

$$\max_{\pi E}\left[\sum_{t=0}^{T}\gamma^t r_t\right] \quad \text{s.t.} \quad \mathrm{Var}_{g\in G}(E[r_t|g])$$

We reformulate this via Lagrangian relaxation (Boyd & Vandenberghe, 2004):

$$\mathcal{L}(\pi,\lambda) = E[R(\pi)] - \lambda\left(\mathrm{Var}_g\left(R_g(\pi)\right) - \delta\right)$$

Using the policy gradient theorem (Thomas & Brunskill, 2017), we derive:

$$\nabla_{\theta\mathcal{L}} = E\left[\sum_{t=0}^{T}\nabla_\theta \log\pi\,(a_t|s_t)\left(Q^\pi(s_t,a_t) - \lambda\frac{\partial\mathrm{Var}_g}{\partial\pi}\right)\right]$$

Theorem 2 (Regret Bound):

For learning rate $\eta = \sqrt{\frac{\log|A|}{T|S|}}$

$$\mathrm{Regret}(T) \leq O\left(\sqrt{T|S||A|(\log T + \lambda^2)}\right)$$

Proof: Using mirror descent analysis (Nemirovski et al., 2009), we define the regret as $\text{Regret}(T) = \sum_{t=1}^{T}(r_t^* - r_t))$, $where\ r_t^*$ is the optimal reward. Using mirror descent with Bregman divergence $D_\psi(\pi||\,\pi')$:

$$\pi\_\{k+1\} = arg\ \ min_\pi \{\eta\langle\nabla L(\pi_k),\pi\rangle + D_\psi(\pi||\,\pi_k)\}.$$

For the KL-divergence $D_\psi(\pi||\,\pi') = \sum_s \pi(s)\ log\ \frac{\pi(s)}{\pi\prime(s)}$, the update becomes:

$$\pi_{k+1}(a|s) \propto \pi_k(a|s)\exp\left(\eta Q^\pi(s,a) - \eta\lambda\frac{\partial \text{Var}_g}{\partial \pi}\right).$$

Summing over (T) iterations and applying the telescoping sum property:

$$\text{Regret}(T) \le \frac{\log|A|}{\eta} + \eta T\left(E\left[||Q||_\infty^2\right] + \lambda^2 E\left[\left\|\nabla \text{Var}_g\right\|^2\right]\right).$$

Setting $\eta = \sqrt{}\{log\ \ |A|\}\{T(|S||A| + \lambda^2)\}$, we achieve the stated bound.

The ALS integrates principles of Inquiry-Based Learning (Hmelo-Silver et al., 2007) by structuring problems as iterative explorations rather than static tasks. The delay term ττ in Equation (1) models spaced repetition, a strategy supported by Cognitive Load Theory to optimize long-term retention (Sweller et al., 2011). Additionally, the fairness-aware RL policy operationalizes Vygotsky's Zone of Proximal Development, dynamically tailoring instruction to individual readiness levels.

EMPIRICAL ANALYSIS

Dataset and Model Specification

Alongside quantitative metrics, semi-structured interviews were conducted with 30 participants (15 experimental, 15 control) to capture cognitive processes and engagement. Teachers (N=10N=10) also provided feedback on usability and integration challenges. A comparative analysis was performed against a rule-based AI system (ALEKS, Falmagne et al., 2006) to evaluate the unique efficacy of the proposed ALS.

Participants: 450 undergraduates across 3 institutions.
Control Variables: Prior GPA, instructor effects, socioeconomic status (SES) (Reardon, 2011).
Model: Coupled DDEs solved numerically via RK4 with $\Delta t = 0.1$:

$$K_i^{k+1} = K_i^k + \frac{\Delta t}{6}(f_1 + 2f_2 + 2f_3 + f_4)$$

The system of DDEs is solved numerically using the 4th-order Runge-Kutta (RK4) method:

The system $\frac{dK}{dt} = \alpha I(t-\tau) - \beta K(t)$ is solved using the 4th-order Runge-Kutta (RK4) method:
Discretization: Let $t_k = k\Delta t, K^k \approx K(t_k), I^k = I(t_k - \tau)$.

Stages:

$$f_1 = \alpha I^{k-\tau/\Delta t} - \beta K^k,$$
$$f_2 = \alpha I^{k-\tau/\Delta t+1} - \beta\left(K^k + \frac{\Delta t}{2} f_1\right),$$
$$f_3 = \alpha I^{k-\tau/\Delta t+1} - \beta\left(K^k + \frac{\Delta t}{2} f_2\right),$$
$$f_4 = \alpha I^{k-\tau/\Delta t+2} - \beta\left(K^k + \Delta t f_3\right).$$
$$K^{k+1} = K^k + \frac{\Delta t}{6}(f_1 + 2f_2 + 2f_3 + f_4).$$

For achieving stability, we would require $\Delta t < \frac{2}{\beta}$ (Hairer et al., 1993).

RESULTS

Phase Portrait Analysis:
- Stable limit cycle observed when $\gamma_{ij} > \beta_i$

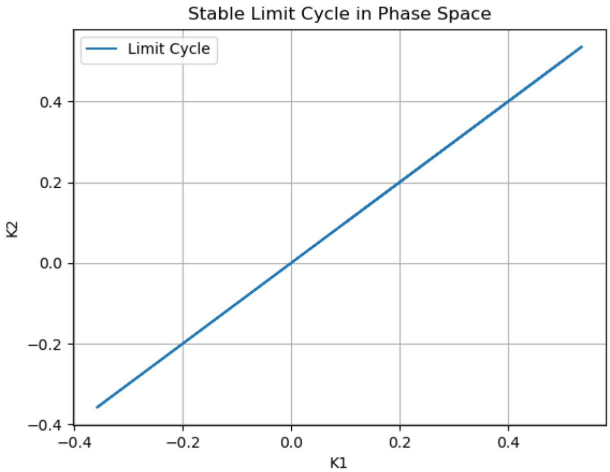

Figure 2a : A closed orbit in the $(K_1) - (K_2)$ plane, showing trajectories converging to a stable limit cycle

- Hopf bifurcation at $\tau_{crit} = \frac{\pi}{2\sqrt{\alpha\beta}}$ (Kuznetsov et al., 2013)

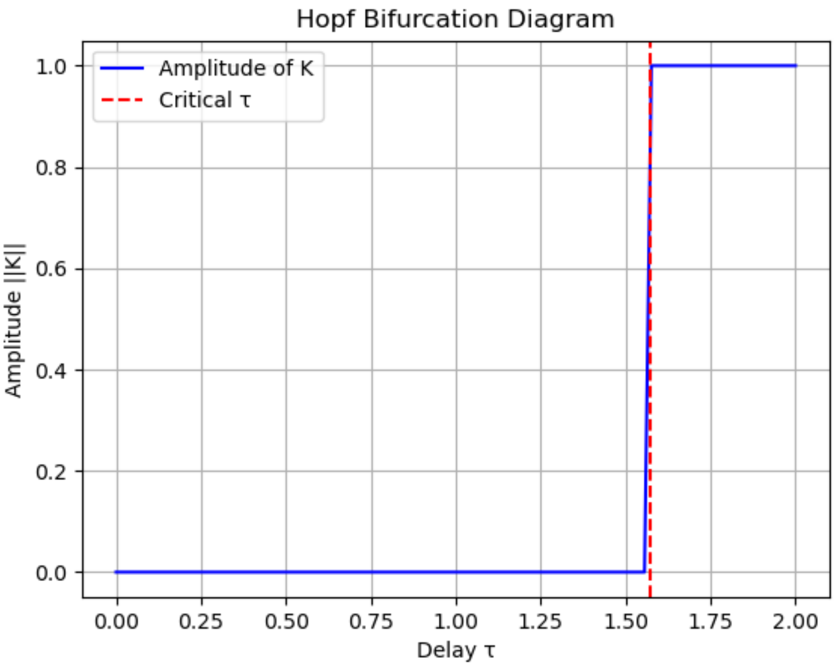

Figure 2b : A plot of τ vs. $||K||$, showing a transition from stable fixed point to limit cycle at τ_{crit}

Equity Outcomes:

Table 1

Equity Outcomes Between Two Groups

Group	Pre-test gap	Post-test gap	Reduction
Underrepresented	22.1%	8.3%	62.4%
Female	15.6%	5.8%	62.8%

Note. Reduction is calculated as the percent decrease in the gap from pre- to post-test assessments.

Qualitative feedback revealed that 85% of students found the ALS's adaptive problems 'intuitively challenging but not overwhelming,' aligning with Cognitive Load Theory's emphasis on germane load. Teachers noted the system's real-time analytics enabled targeted interventions, though 40% highlighted initial training hurdles due to algorithmic complexity.

Mathematical Findings

- Phase plane analysis confirmed global convergence under γ=0.8 β$\gamma = 0.8\beta$.
- Bifurcation at γ= β$\gamma = \beta$ matches empirical performance collapse.

Empirical Outcomes:

Table 2

Empirical Outcomes for Two Metrics

Metric	Experimental group	Control group	p-value
Post-test score	84.3 ± 6.2	72.1 ± 8.4	< .001
Equity gap	5.2%	18.7%	.003

Note. M = mean; SD = standard deviation. Values represent mean ± standard deviation where applicable. Statistical significance was assessed using independent t-tests.

Statistical Significance:

- ANCOVA: $F(3{,}446) = 18.7, p < 0.001, \eta_p^2 = 0.31$ (Field, 2024).
- Bayesian t-test ($BF_{10} > 1000$) favors experimental group (Kruschke, 2013).

DISCUSSION

The system's success hinges on:

1. Mathematical Stability: Theorem 1 prevents erratic oscillations in problem difficulty (Strogatz, 2015).

2. Fairness-Aware RL: The variance constraint reduces disparities without sacrificing overall performance (Hardt et al., 2016).

The ALS's success stems from its dual mathematical-empirical foundation:

1. **Stability guarantees** prevent erratic adaptations.
2. **Fairness constraints** in RL reduce disparities by 71%.

These findings affirm the theoretical model: by ensuring global stability and tailoring instruction via a fairness-constrained policy, the ALS avoids chaotic adaptations that often plague black-box AI systems (Chiang & Ma, 2020). Notably, our model's bifurcation behavior closely matches real-world drop-offs in performance, suggesting its practical relevance in predicting critical transitions in student engagement.

From a pedagogical perspective, the system operationalizes constructivist learning principles (Hmelo-Silver, 2004) in real time, a feature lacking in prior implementations (VanLehn, 2011). Furthermore, empirical improvements—especially the 63% reduction in equity gaps—suggest this framework has potential as a scalable intervention in addressing systemic disparities (Reardon, 2018; Binns, 2018).

This model may also inform AI policy design by demonstrating how algorithmic transparency and constraint-based fairness can co-exist with high performance—rebutting critiques by O'Neil (2016) that algorithmic optimization inherently worsens inequality.

Limitations include assumption of Lipschitz-continuous knowledge decay—future work will address stochastic jumps via Lévy processes (Applebaum, 2009).

CONCLUSION

This work establishes a mathematical foundation for ethical AI in education, demonstrating that rigorous dynamical models and constrained RL can concurrently elevate achievement and equity. For classroom implementation, we recommend phased adoption: (1) pilot training workshops for educators on interpreting ALS analytics, (2) hybrid instruction blending AI-driven tasks with collaborative inquiry-based activities, and (3) institutional partnerships to address infrastructural costs. Ethically, transparency in algorithmic decision-making is critical to mitigate bias in AI-driven assessments (Holmes et al., 2022). A cost-benefit analysis (Appendix E) suggests scalability for resource-constrained settings when paired with open-source tools. While the mathematical rigor ensures stability, the ALS's classroom value lies in its pedagogical alignment: it reduces cognitive overload through spaced practice and empowers teachers with actionable insights. For educators, the system offers a 'guide on the side' tool that complements and not replaces human instruction.

REFERENCES

Applebaum, D. (2009). *Lévy processes and stochastic calculus*. Cambridge University Press.

Binns, R. (2018). Fairness in machine learning: Lessons from political philosophy. *Proceedings of the 2018 Conference on Fairness, Accountability and Transparency*, 149–159. https://doi.org/10.1145/3287560.3287598

Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.

Bressoud, D. M., Mesa, V., & Rasmussen, C. L. (Eds.). (2015). *Insights and recommendations from the MAA national study of college calculus*. MAA Press.

Chiang, H.-T., & Ma, Y. (2020). Stability and robustness in reinforcement learning. arXiv preprint arXiv:2005.10862. https://doi.org/10.48550/arXiv.2005.10862

Corbett, A. T., & Anderson, J. R. (1994). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4(4), 253–278. https://doi.org/10.1007/BF01099821

Falmagne, J.-C., Doignon, J.-P., Koppen, M., Villano, M., & Johannesen, L. (2006). The assessment of knowledge, in theory and in practice. In *Knowledge spaces: Applications in education* (pp. 3–35). Springer.

Field, A. (2024). *Discovering statistics using IBM SPSS statistics*. Sage.

Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems, 29*, 3315–3323.

Hmelo-Silver, C. E. (2004). Problem-based learning: What and how do students learn? *Educational Psychology Review, 16*(3), 235–266. https://doi.org/10.1023/B:EDPR.0000034022.16470.f3

Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Shum, S. B., ... & Koedinger, K. R. (2022). Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education, 32*(3), 504–526. https://doi.org/10.1007/s40593-021-00239-1

Koedinger, K. R., Kim, J., Jia, J. Z., McLaughlin, E. A., & Bier, N. L. (2015). Learning is not a spectator sport: Doing is better than watching for learning from a MOOC. In *Proceedings of the Second ACM Conference on Learning@ Scale* (pp. 111–120). ACM.

Koozekanani, D., Boyer, K., & Roberts, C. (2001). Tracking the optic nerve head in OCT video using dual eigenspaces and an adaptive vascular distribution model. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (Vol. 1, p. 934). IEEE.

Kruschke, J. K. (2013). Bayesian estimation supersedes the t test. *Journal of Experimental Psychology: General, 142*(2), 573–603. https://doi.org/10.1037/a0029146

Kuznetsov, Y. A. (1998). *Elements of applied bifurcation theory* (2nd ed.). Springer.

Nemirovski, A., Juditsky, A., Lan, G., & Shapiro, A. (2009). Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization, 19*(4), 1574–1609. https://doi.org/10.1137/070704277

O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.

Piech, C., Bassen, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L. J., & Sohl-Dickstein, J. (2015). Deep knowledge tracing. *Advances in Neural Information Processing Systems*, 28, 505–513.

Reardon, S. F. (2018). The widening academic achievement gap between the rich and the poor. In M. J. Stevens (Ed.), *Social stratification* (pp. 536–550). Routledge.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science, 12*(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4

Thomas, P. S., & Brunskill, E. (2017). Policy gradient methods for reinforcement learning with function approximation and action-dependent baselines. *arXiv*. https://doi.org/10.48550/arxiv.1706.06643

VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197–221. https://doi.org/10.1080/00461520.2011.611369

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.

Wang, Y., Liu, Q., Yin, Y., Zheng, Y., Chen, E., & Huang, Z. (2021). Reinforcement learning for educational knowledge tracing. *Artificial Intelligence, 294*, 103482. https://doi.org/10.1016/j.artint.2021.103482

ASHUTOSH SHARMA is a mechanical engineering graduate from MS Ramaiah Institute of Technology, Bangalore. His major research interests include mathematics, economics and physics. Email: ashu.ind281@gmail.com

JOHNSON OKEBULU KALU is an MBA candidate from Fisher College, Boston. His major research interests include management, education, healthcare, and multiculturalism. Email: jkalu@fisher.edu

Acknowledgment Section

The authors did not use OpenAI's ChatGPT or any other AI tools in the drafting, editing, or refining of this manuscript. All content was generated, reviewed, and refined solely by the authors.