

ASPIRE

AI-Powered Personalized Learning

Sesh Murthy
Virginia de Sa

UC San Diego

SCHOOL OF COMPUTING, INFORMATION
AND DATA SCIENCES



Acknowledgements



Andrea Chiba
Cognitive Science



Amy Eguchi
Education Studies



Gordon McCord
Global Policy and Strategy

UC San Diego Contributors

Frances Hammock
Ko Woon Ohm
Leesa Anzaldo
Hsiao-Bing Cheng
Mike Holst
Adam Bowers

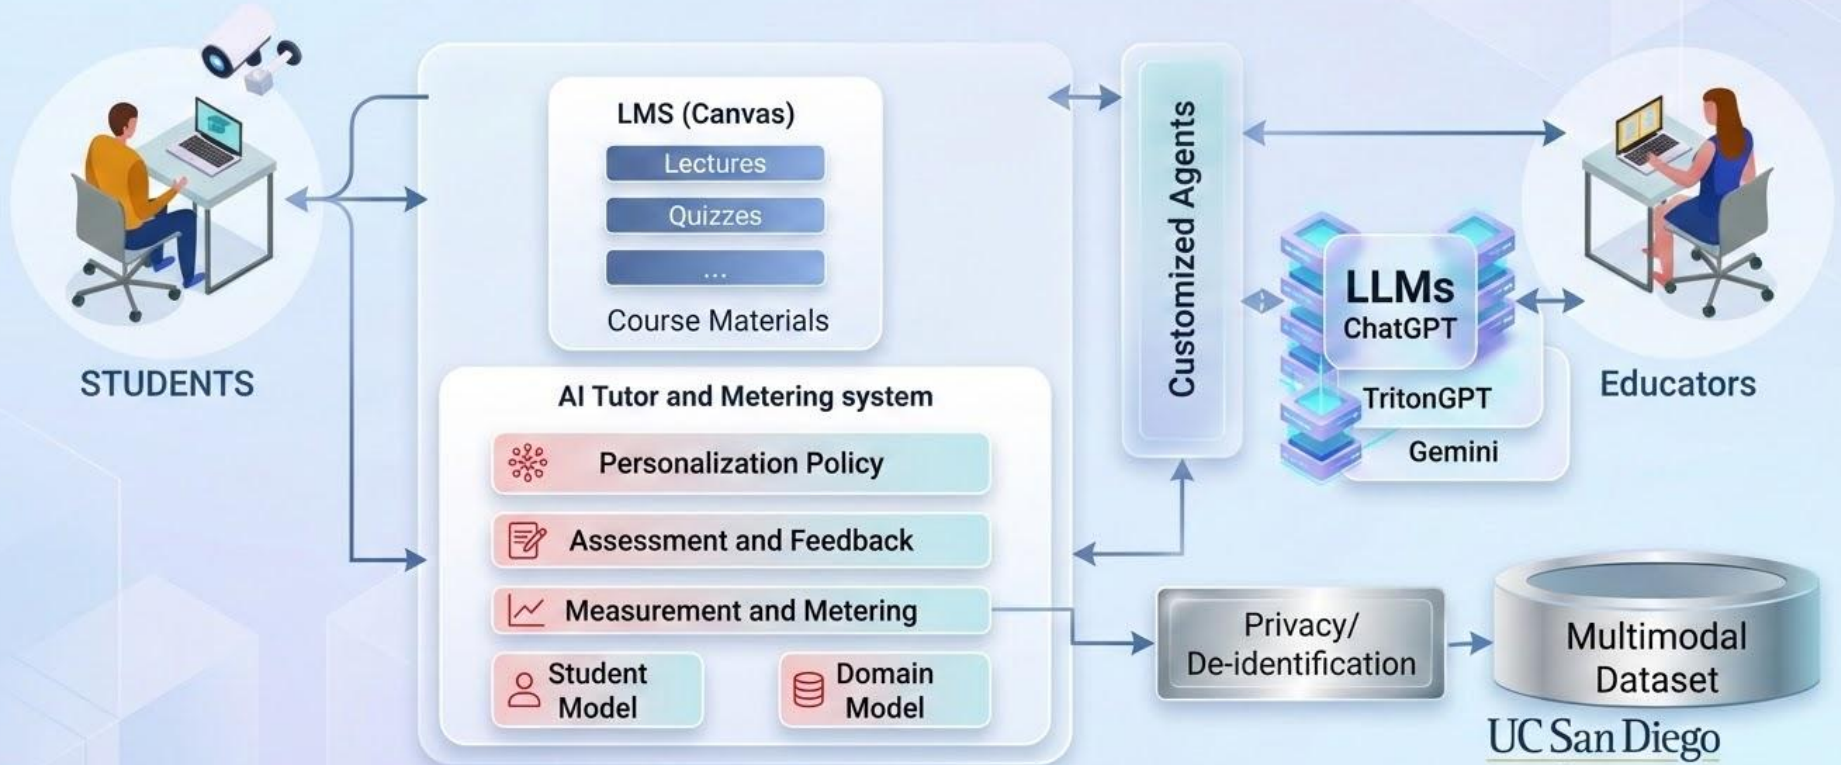
Umberto Mignozzetti
Rayan Saab
Anushri Eswaran
Shuangquan Feng
Philip Chi
Daniel Ritchie

Partners & Collaborators

Rob Rubalcaba, SDCC

Vocareum Team
Sanjay Srivastava
Jagmeet Singh
Portia Restuccia

Architecture of ASPIRE



Automating Workflows from Science of Learning

- Implemented on a highly scalable platform (Scales to 4 million users)
- Deployed in 12 courses with 2236 students
- Collect fine-grained data, analyze outcomes, test candidate mechanisms, and iteratively refine workflows



ASPIRE for Math Prep in GPS

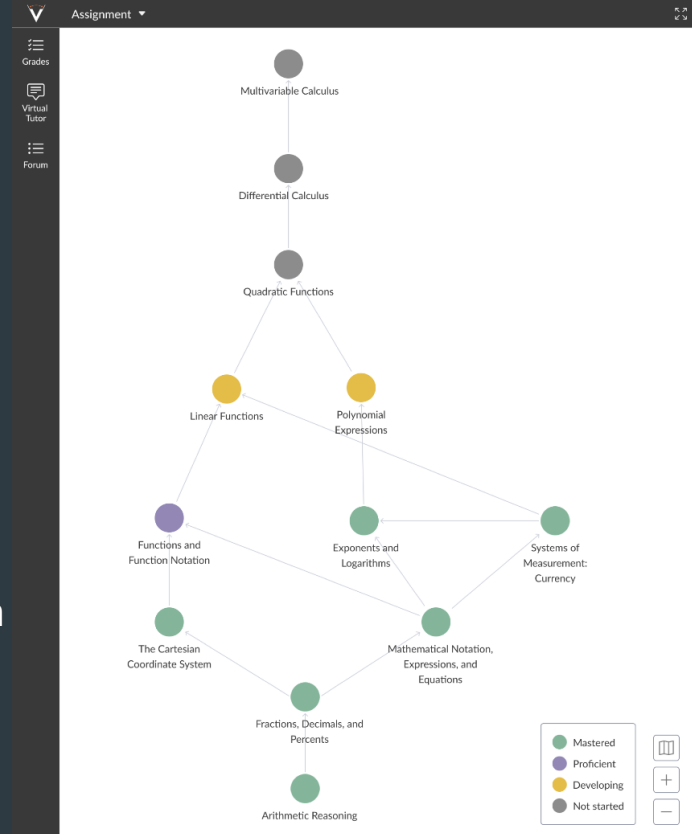
The GPS Math Problem

- Incoming GPS students arrive with highly uneven math backgrounds
- Math Camp and Econ/QM Prep must get everyone to a common floor very fast
- Faculty and TAs cannot realistically provide personalized remediation at scale

ASPIRE: A faculty controlled AI tutoring Platform

- Provides 24/7 personalized tutoring tied directly to GPS math and econ content
- Gives step-by-step feedback on every answer
- Lets instructors author, curate, and monitor both content and AI behavior
- This is not ChatGPT — it is a curated, course-embedded AI tutor aligned with our curriculum

[content from Gordon McCord (Associate Dean of Student Affairs and Associate Teaching Professor, GPS)]





Grades



Assistants



Forum

I need a hint to get started.

Virtual Tutor 1:40 PM Verified

To find the consumer and producer surplus, first determine the demand and supply equations using the given price and quantity pairs.

[Create forum thread](#) **You** 1:42 PM

what's the difference between consumer and producer surplus

Virtual Tutor 1:40 PM

Here's a brief explanation:

- **Consumer Surplus:** The extra benefit consumers get when they pay less than the maximum price they're willing to pay. It's the area between the demand curve and the equilibrium price.
- **Producer Surplus:** The extra benefit producers earn when they sell at a market price higher than their minimum acceptable price. It's the area between the supply curve and the equilibrium price.

Keep going—you've got this!

[Create forum thread](#) **You** 1:42 PM

How do I find the area of a triangle?



Ask a question



Concept

Consumer and Producer Surplus

Question

Question list

Progress map

Unlock next concept: 90%



Question 5

You are managing a car dealership. Your market research team tells you that at a price of \$20,000, customers demand 40 cars, and at a price of \$30,000, the demand drops to 20 cars. They also inform you that at a price of \$20,000, you can supply 10 cars, and at a price of \$40,000, you can supply 50 cars. What is the consumer and producer surplus?



Submit answer

Previous

Next

How ASPIRE Helps Students, Faculty, and TAs

- Visual *Student Model* lets students see what they know & where they are stuck
- Hundreds of vetted problems (800+) generated and refined for GPS content
- AI Tutor:
 - Explains mistakes step-by-step
 - Answers follow-up “why” questions
 - Works in any language
 - Reduces math anxiety and fear of asking questions
- “The AI Tutor was able to point out the specific step I missed and explain why... I found it very helpful to ask guiding questions without feeling silly”
- Daily AI digests summarize:
 - Which questions students struggled with
 - Common misconceptions
 - Where time should be spent in class and section
- TAs reported this made help sessions more targeted and efficient

GPS Summer Math Camp 2025 Results

- 110 students; 18,000+ graded problems completed
- 846 unique AI-generated questions
- 4,175 student questions to the AI tutor
- Learning signal
 - Engagement with ASPIRE statistically associated with test score improvement
 - Effects strongest among students starting in the bottom 75% of the class on the pre-test
 - Not causal —but promising for targeting our most vulnerable students.
- What's next
 - Math Camp content is now built and reusable
 - Natural expansion to Econ Prep & Fall Micro
 - Explore proper encouragement-design experiment to measure causal effects

The UCSD Foundational Math Problem

The Math Proficiency Gap

- Around 80% of the UCSD student body requires Calculus for degree completion
- Between 2020 and 2025, the number of students whose math skills fall below high school level increased nearly thirtyfold
- Roughly, one-eighth of entering students are not meeting high school standards; of these, 70% are below a middle school level

Math Preparatory Sequence and Performance

- The math department designed Math 2 and 3B to remediate middle school and high school math respectively
- Low promotion rates from Math 2 to 3B
- Failure rates in math 3B classes
 - F24 44.3%
 - W25 52.9%
- Pass rates in downstream Calculus courses remain **low**



Comprehensive Instructional Redesign in Math 3B

Weekly Assessment Schedule

- **Monday:** Online homework using **ASPIRE** providing personalized assistance to the first exposure to the material (NEW FA25)
- **Wednesday: Written Homework** focused on critical thinking and the formal mathematical writing required for Calculus courses (**SAME**)
- **Friday: Proctored Online Quizzes** with practice in PrairieLearn to ensure mastery (NEW FA25)

Other Resources (All NEW FA25)

- ASC@Math tutor (M-F 10A-8P)
- Practice PrairieLearn Questions 24/7
- Practice ASPIRE Questions/Help 24/7

[content from Frances Hammock, Assistant Teaching Professor, Mathematics]

Test Preparation and Foundation

- In person **Diagnostic Test** (NEW FA25)
 - Students received individualized feedback and assigned foundational units through **ASPIRE** to target areas of need (NEW FA25)
- Explicit teaching of study skills (NEW FA25)
- Written problem bank (**SAME** but improved)
- Problem bank required submission before each exam (NEW FA25)
- **Online practice exams** through **ASPIRE** with instant feedback (NEW FA25)



Other Class Differences

- FA25 Cohort had higher placement score requirement (more consistent and higher background level affects class atmosphere)
- Some (5-10%) harder material replaced with easier material
- Time per topic adjusted
- Carefully selected TAs



Student Needs and ASPIRE Solutions

Student Needs

- Students have learned mathematical misconceptions that are best addressed with individualized support.
- All students need support on developing college level study skills
- Varying knowledge gaps that require individualized support to close them effectively

Institutional Constraints/Challenges

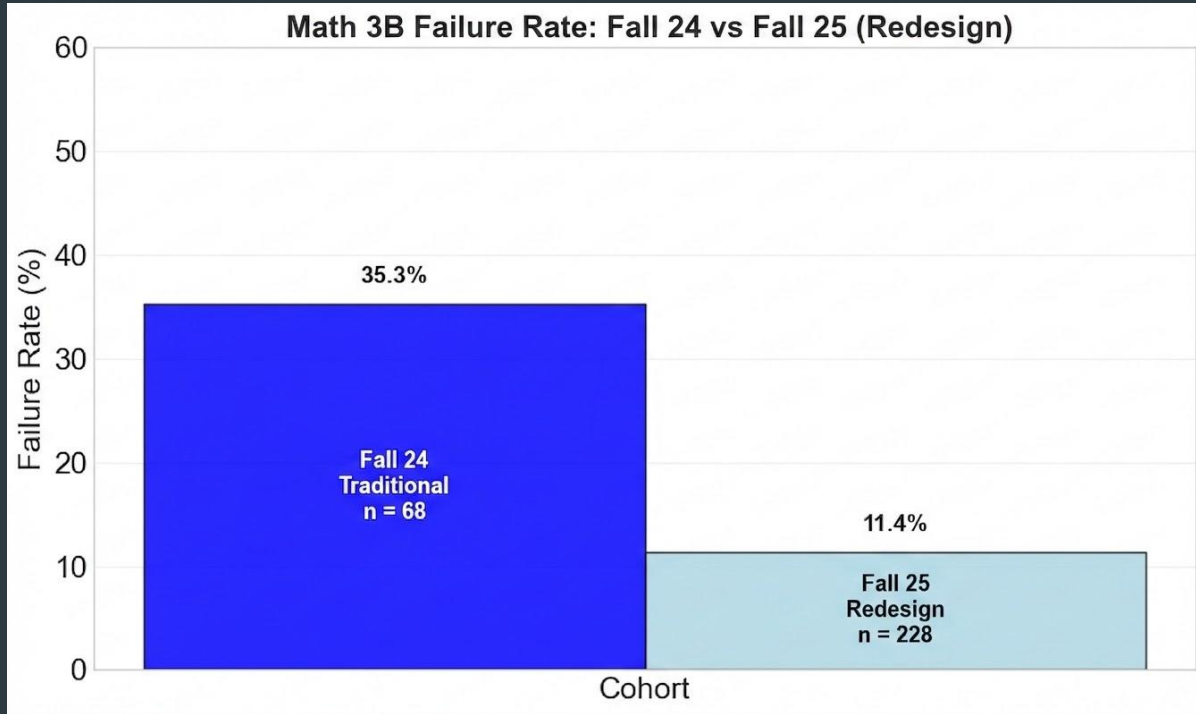
- Large class sizes
- Limited tutoring and office hours
- Not enough graders (and delays)
- Discomfort of students to accessing in person help
- Easy access to AI resources that solve homework entirely

ASPIRE AI Benefits

- Available 24 x 7
- Safe space to ask basic questions
- Grades all questions, not just a subset
- Catches mistakes in real time like a human tutor
- Generates instant individualized feedback; prevents reinforcement of misconceptions
- Provides individualized foundational review units with Aspire based on diagnostic test performances
- Scaffolded help to get started. Preserves productive struggle.



Impact on Student Outcomes in Math 3B

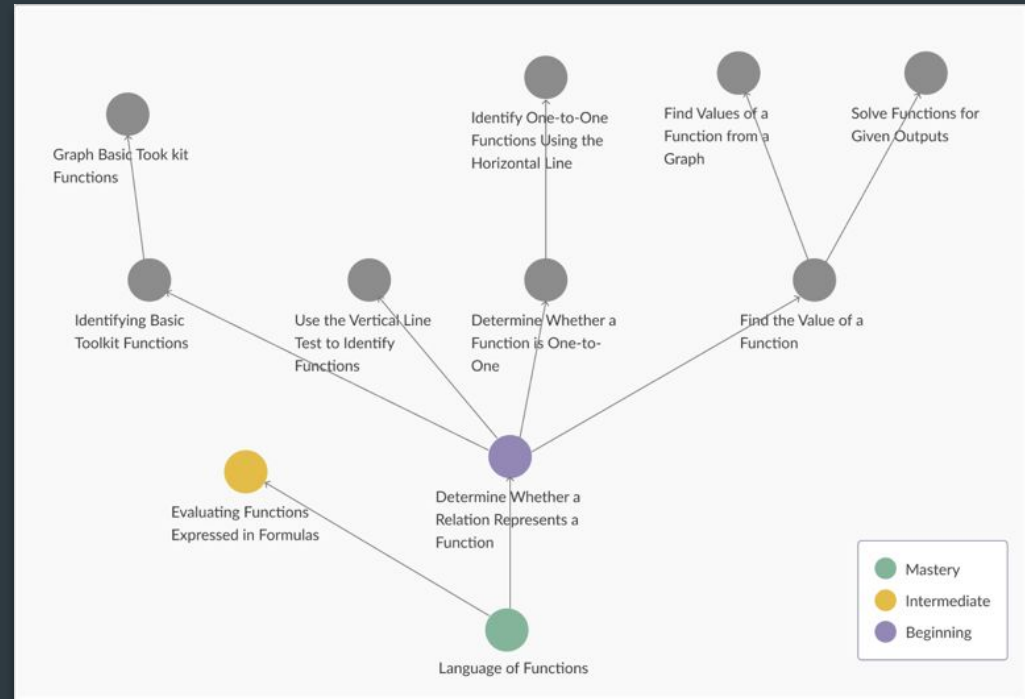


- Controlled for first attempt, quarter, placement score, and instructor
- ASPIRE added and other Significant Structural Interventions made in FA 25 (as described on slides 10, 11)



Student Model: A Dynamic Representation of Learning

- Digital Twin: Dynamic model of concept level mastery, knowledge gaps, and learner state
- Enables adaptive personalized instruction by AI Tutors
- Supports instructional decision-making
- Guides student to focus on required concepts



How the Tutor was Used

The ASPIRE Tutor is configured to support learning

- System guardrails
- Prompts that can be modified by the instructor
- Relevant context: the problem, any instructor materials, domain model
- Personalized based on the student model

10 week quarter (FA25), 280 students:

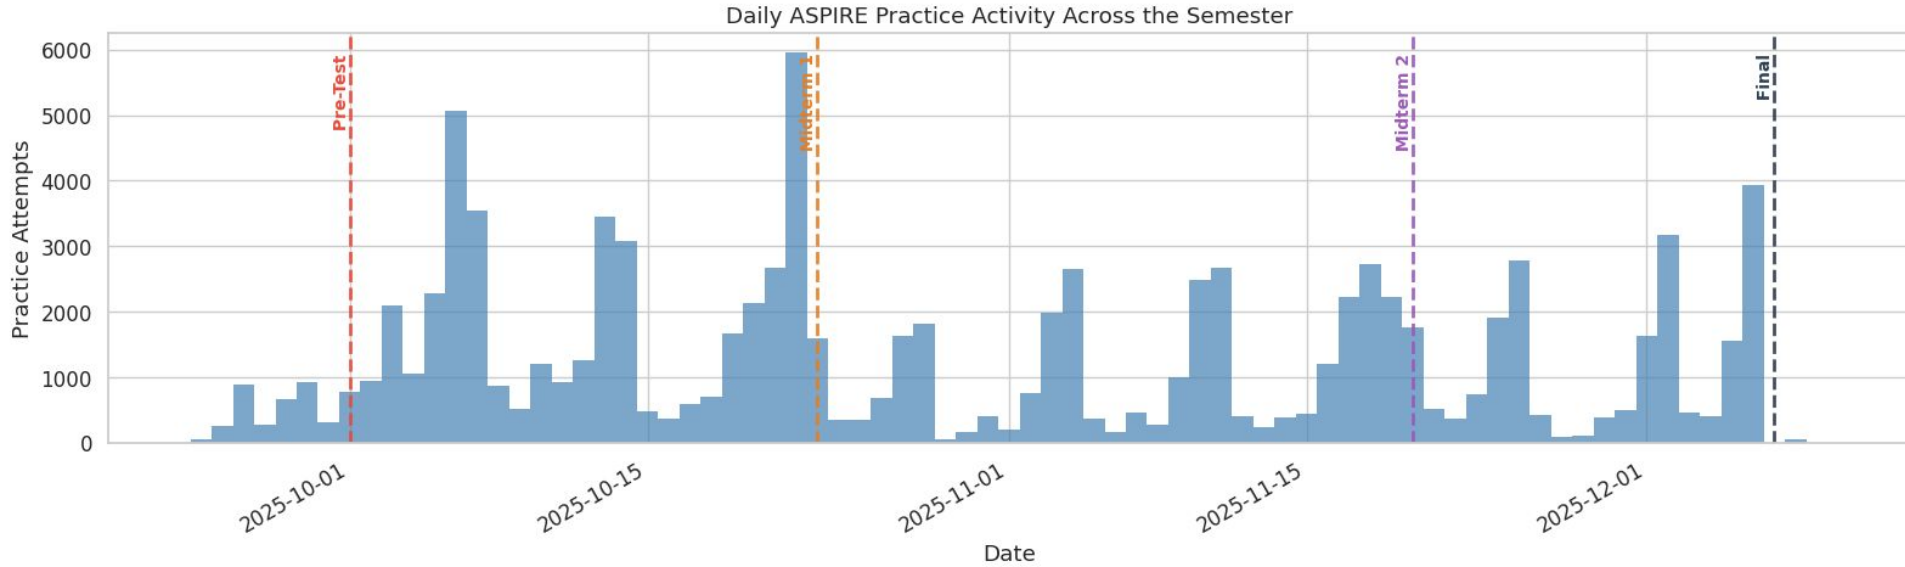
- 16,924 questions asked of the tutor
- 80,047 free response questions graded and feedback given

Constant iteration of the platform to implement best practices

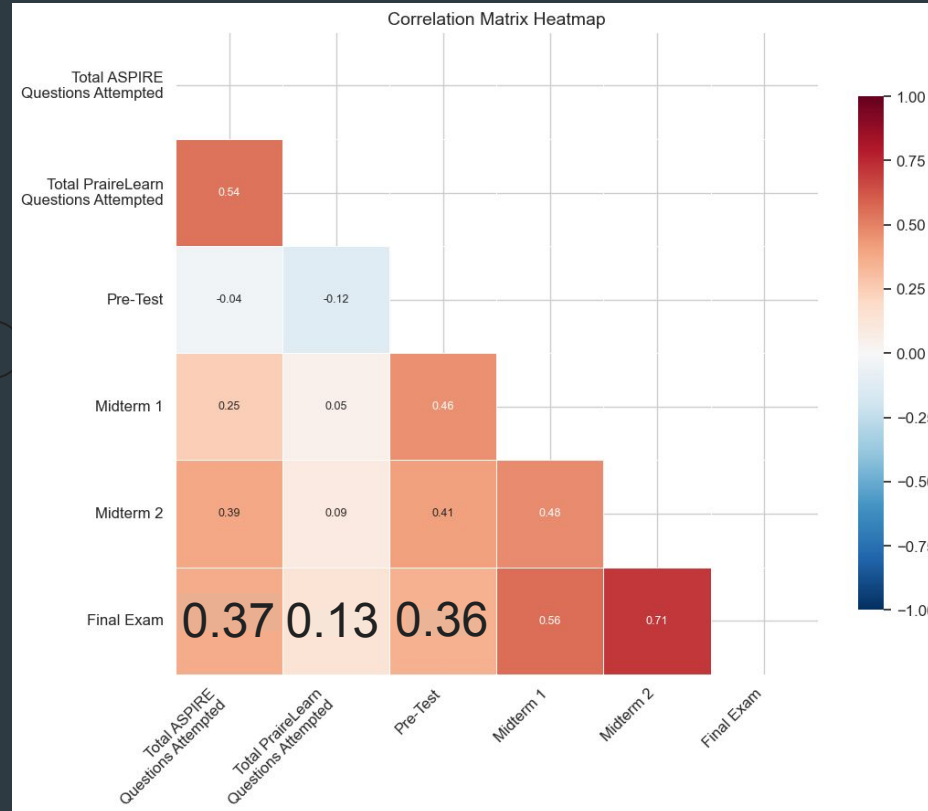
- Educators and students in the loop



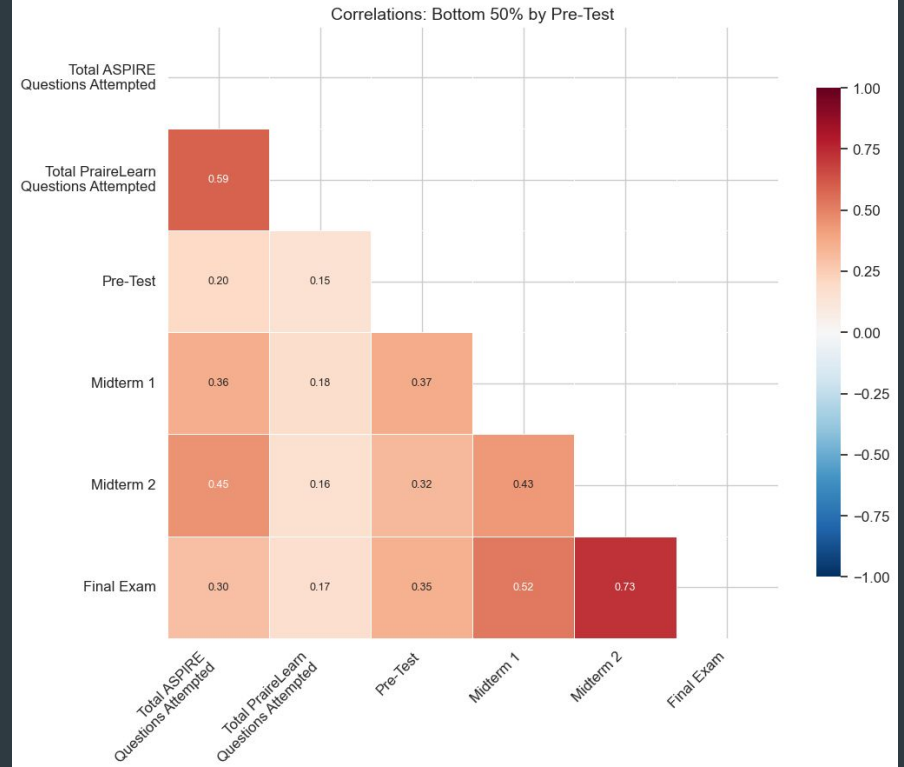
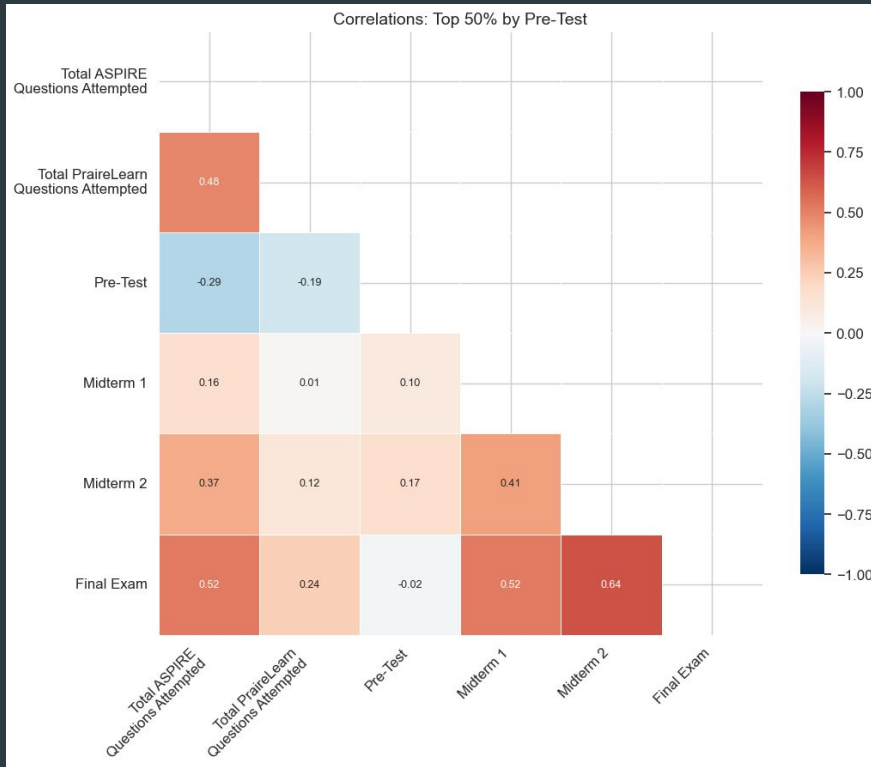
ASPIRE Use



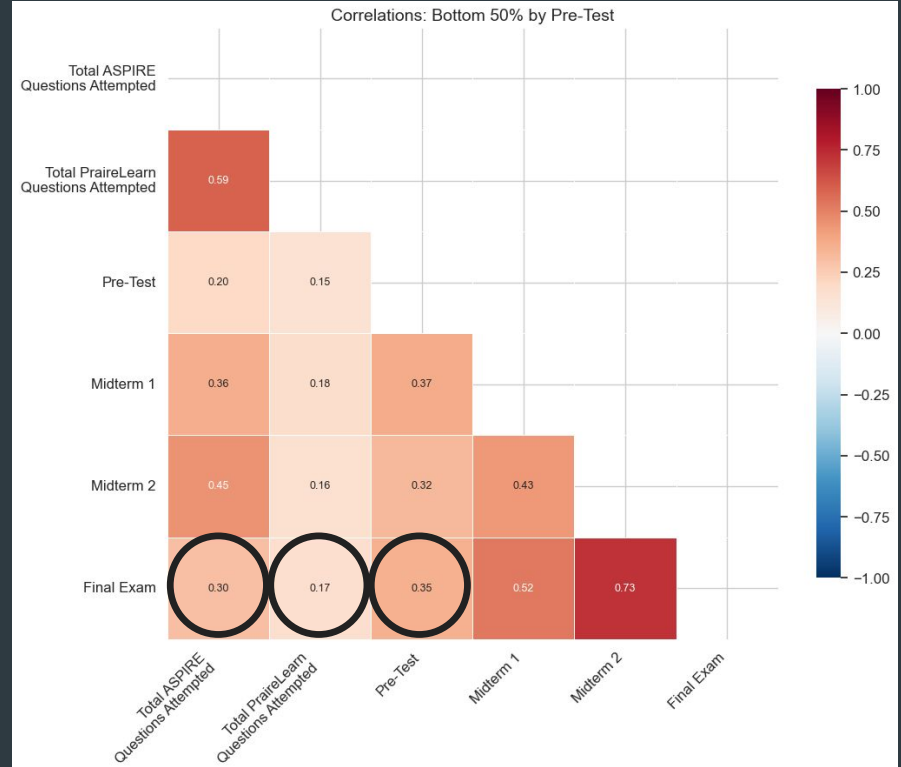
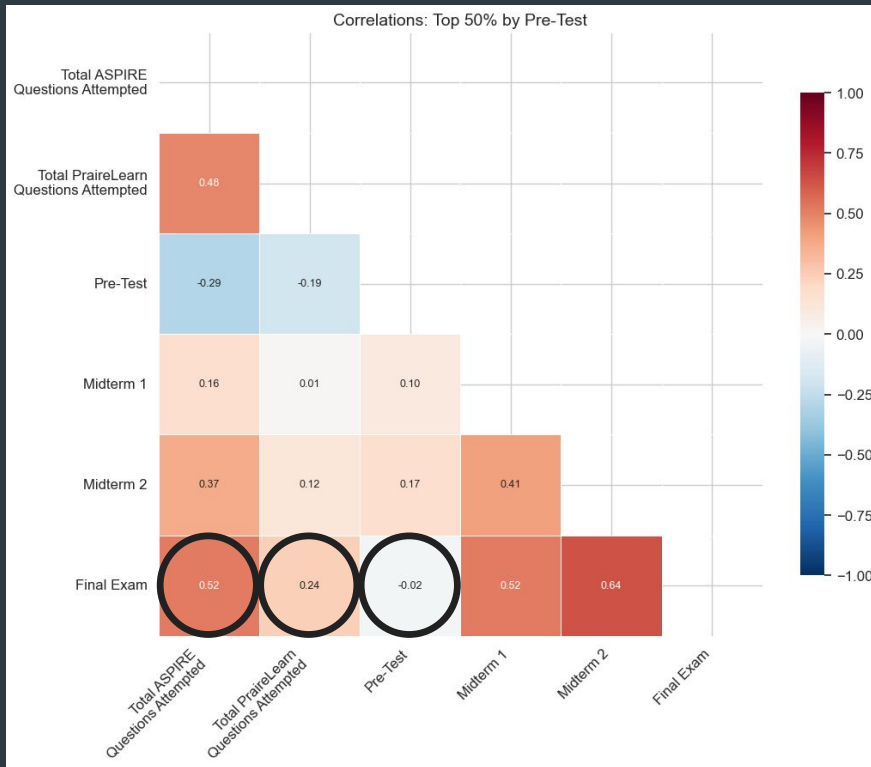
Correlation Matrix



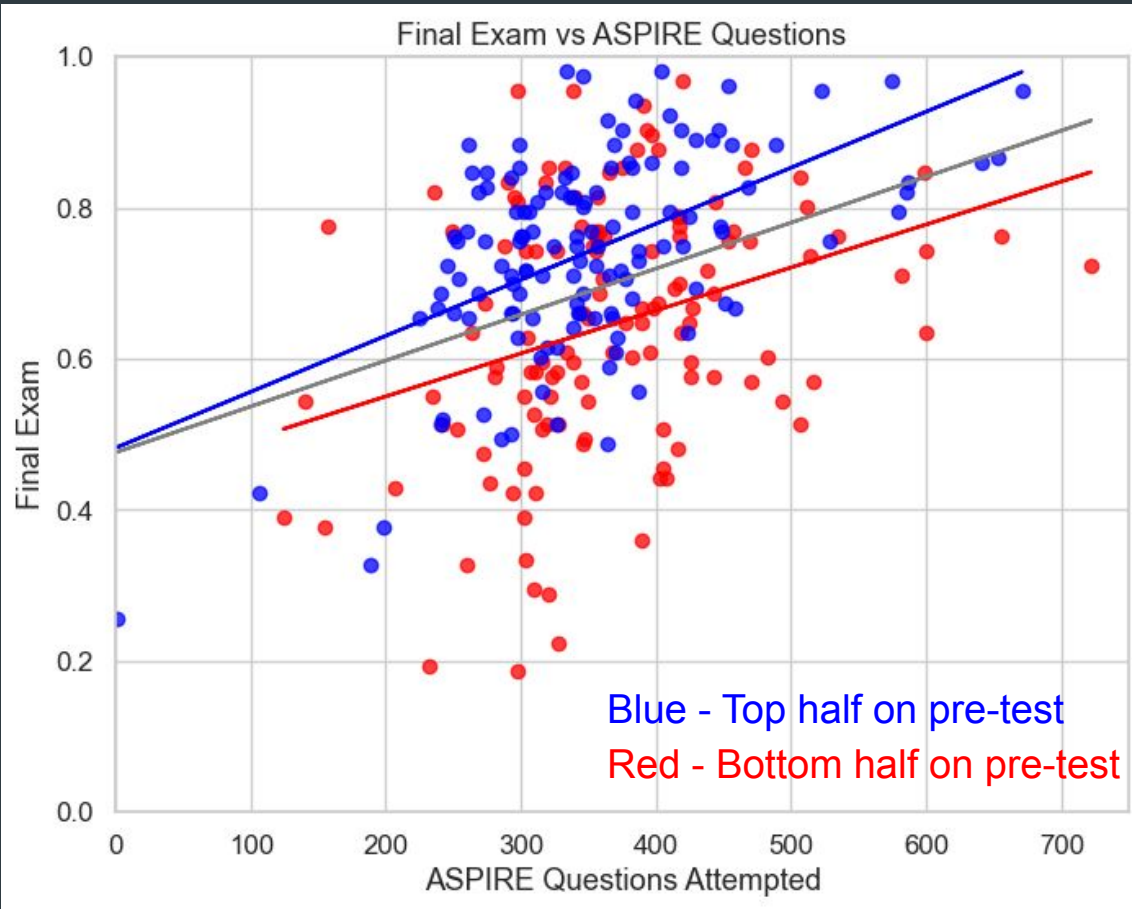
Correlation Matrices



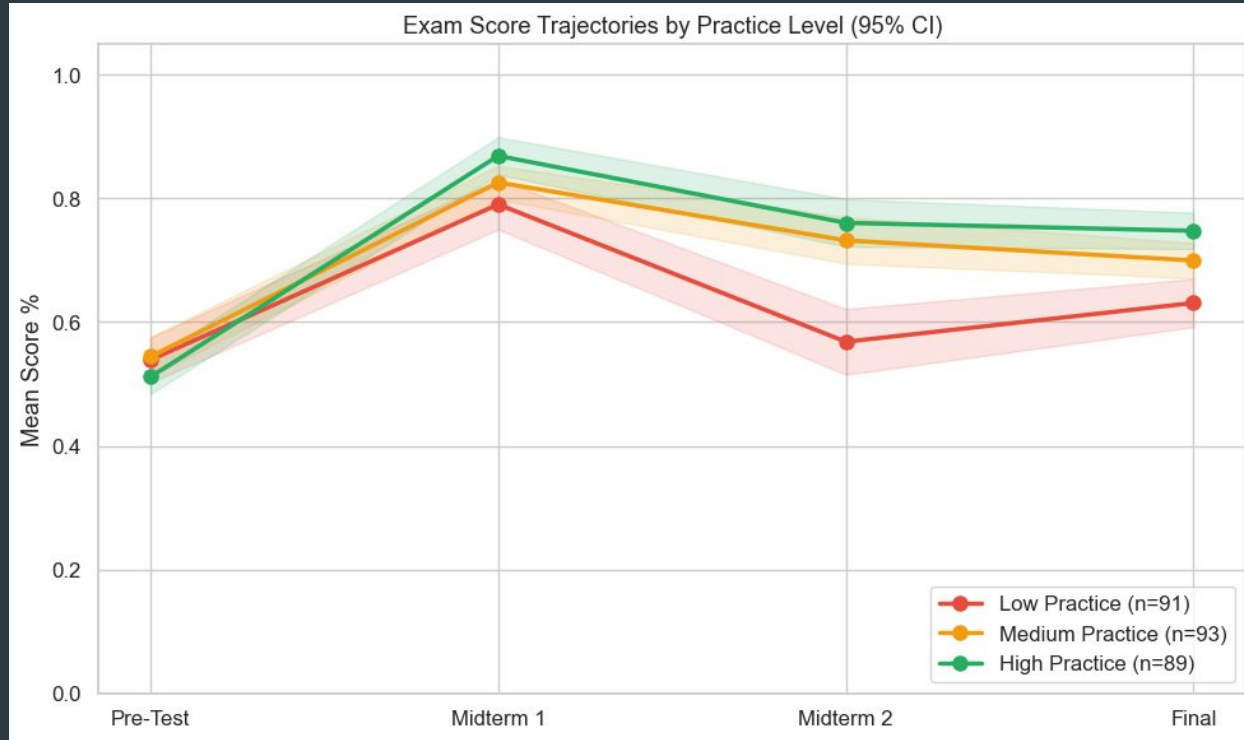
Correlation Matrices



More Practice → Improved Scores

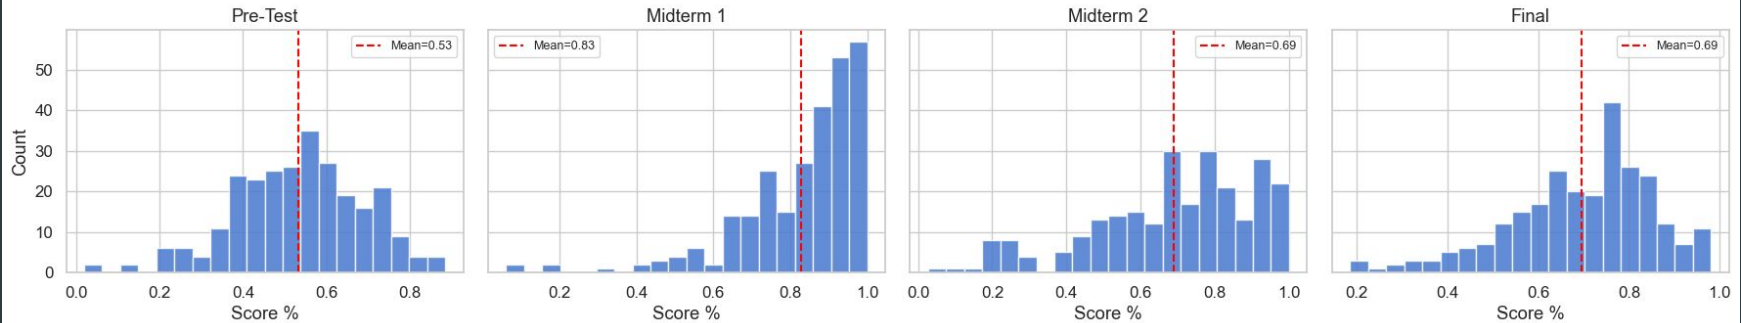


Higher Practice → Higher Scores

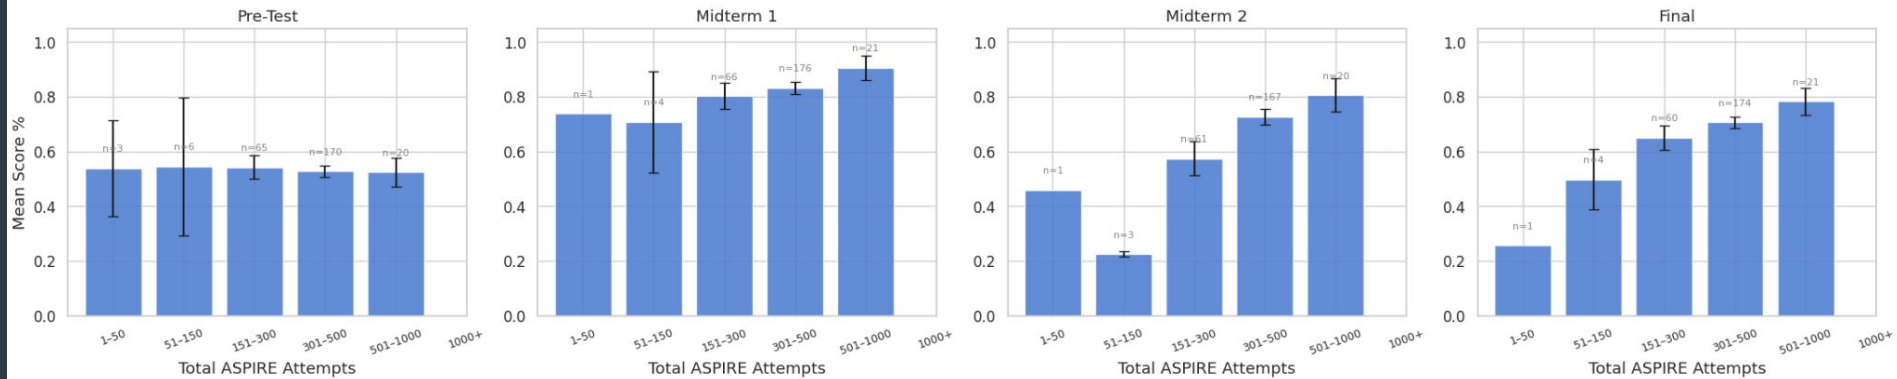


More Practice → Improved Scores

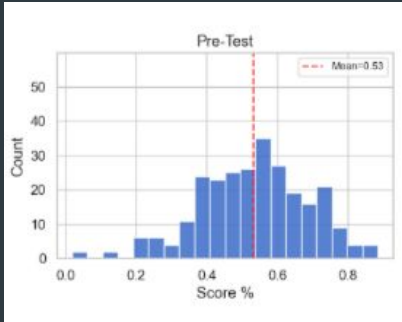
Exam Score Distributions



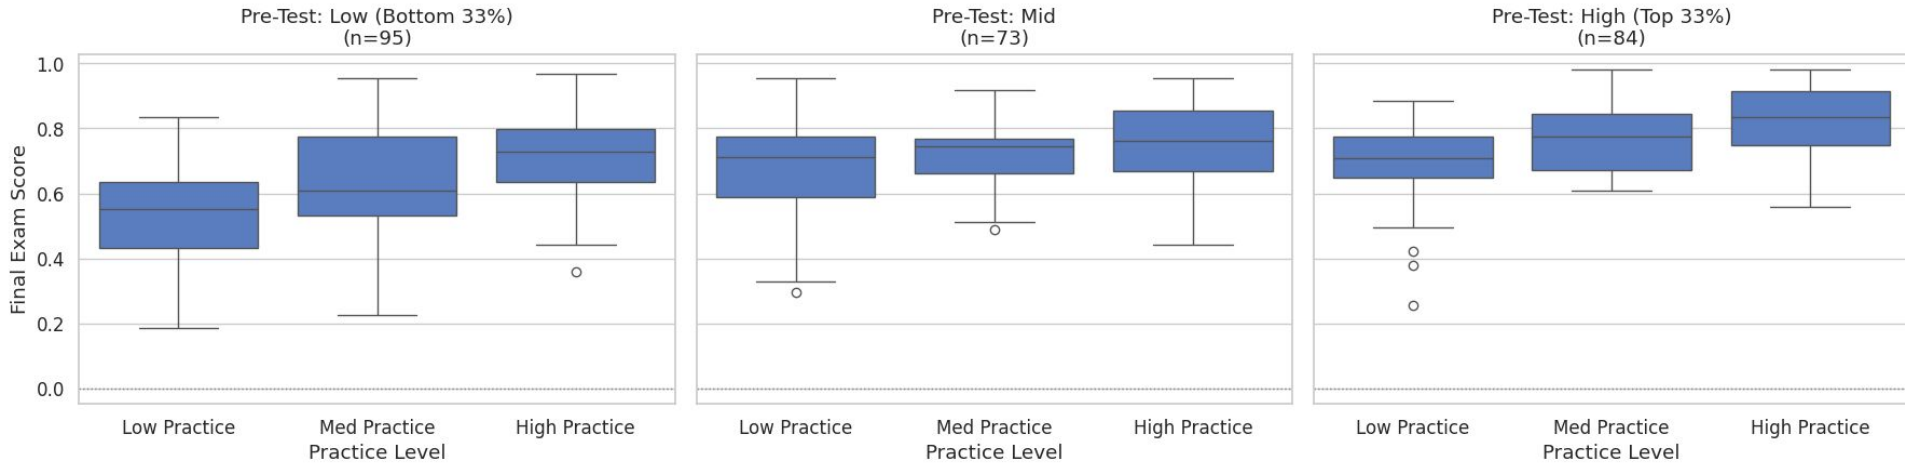
Dose-Response: More Practice → Better Exam Scores?



Consistent Effects Regardless of Pre-Test Grade



Practice Impact by Pre-Test Performance Level



Findings and Preliminary Mechanisms of Impact

- Failure rate reduced by $\frac{2}{3}$ in high failure gateway course
- Candidate mechanisms of impact (hypotheses)
 - Engagement drives outcomes
 - Course workflow design drives engagement
 - AI-tutor promotes productive struggle but instant feedback gives rewards
 - Cadence supports distributed practice
- Next:
 - Test mechanisms and isolate drivers
 - Expand upon personalization capabilities
 - Replicate and scale



Summary of Findings

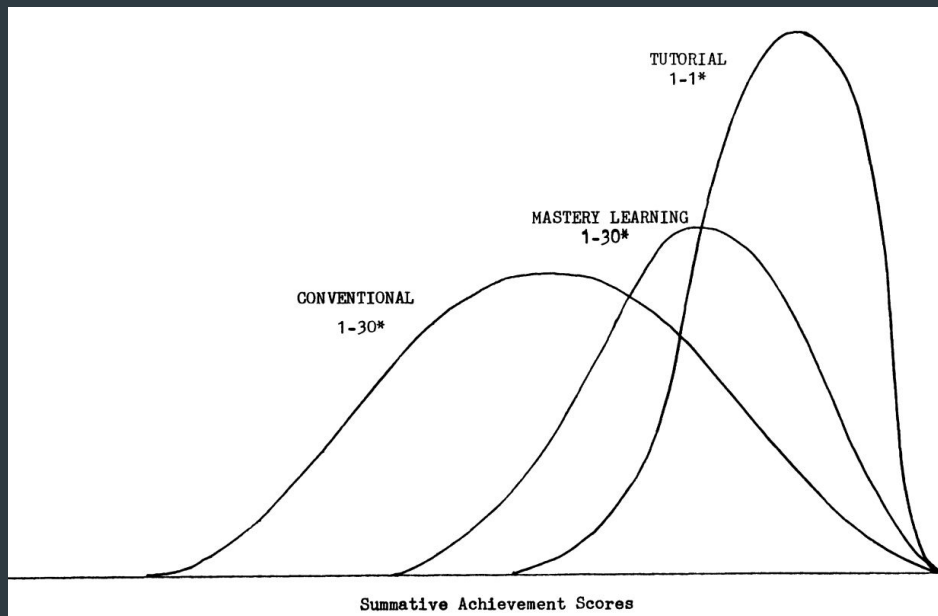
- ASPIRE is a scalable, instructor-controlled platform for personalized learning, integrated with Canvas.
- It uses a domain model and student model to provide personalized support at scale.
- In Math 3B, within a comprehensive course redesign, the failure rate fell by about two-thirds.
- Student outcomes were positively correlated with ASPIRE practice.



Questions?



Bloom's 2 σ Gap

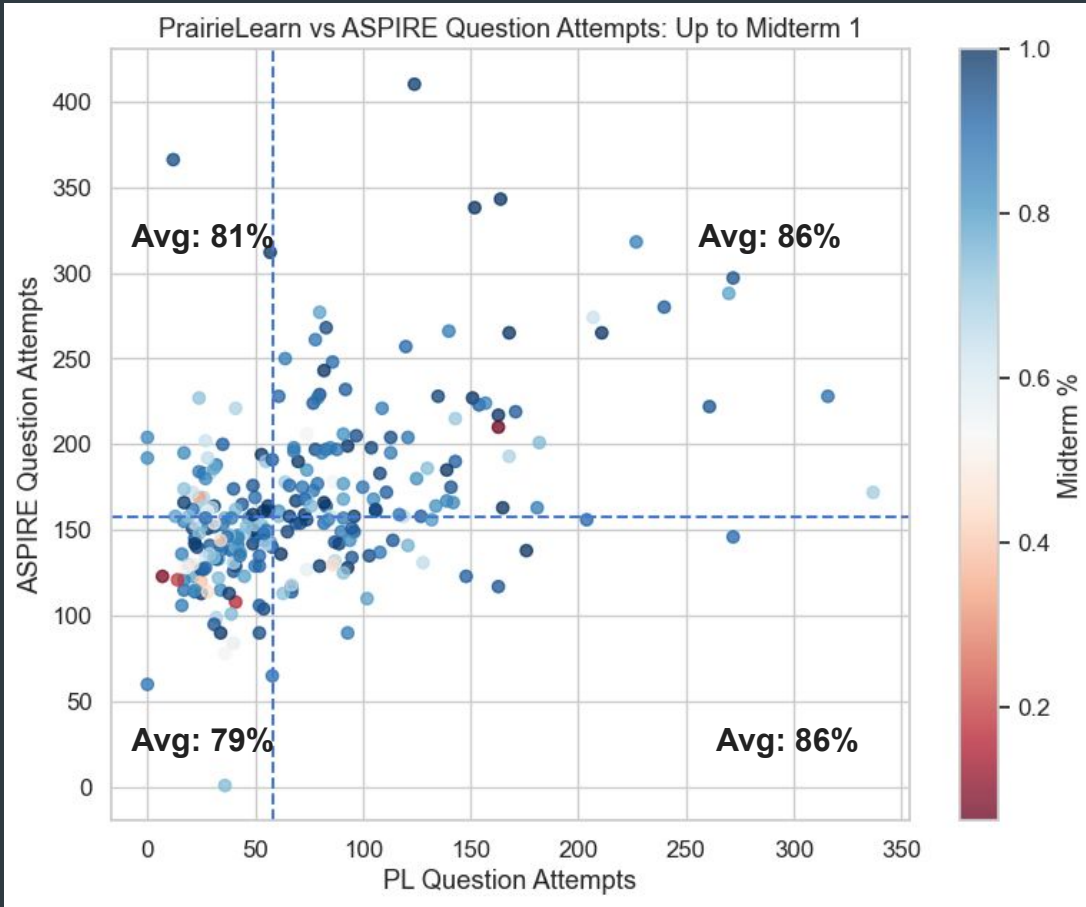


- Average tutored student was above 98% of the students in control
- Time on task: 90+%(tutoring) vs 65% (control)
 - One-on-one interaction with the student makes the student more attentive
- Mastery learning
- Personalized instruction with feedback.
 - Rapid assessment of student knowledge
 - Focus on knowledge gaps

The 2 Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-one Tutoring, Educational Researcher, Vol 13, No 6, (Jun - Jul 1984), pp. 4-16.

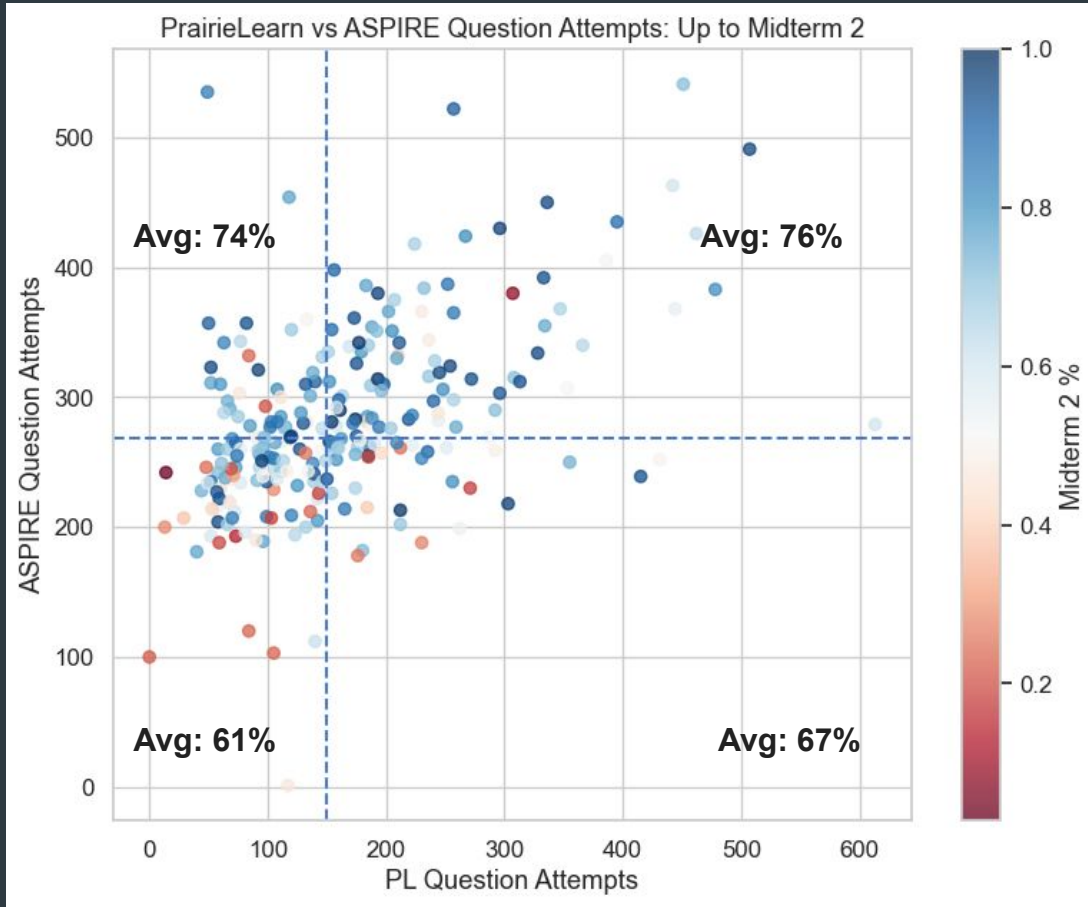


More Practice → Improved Scores



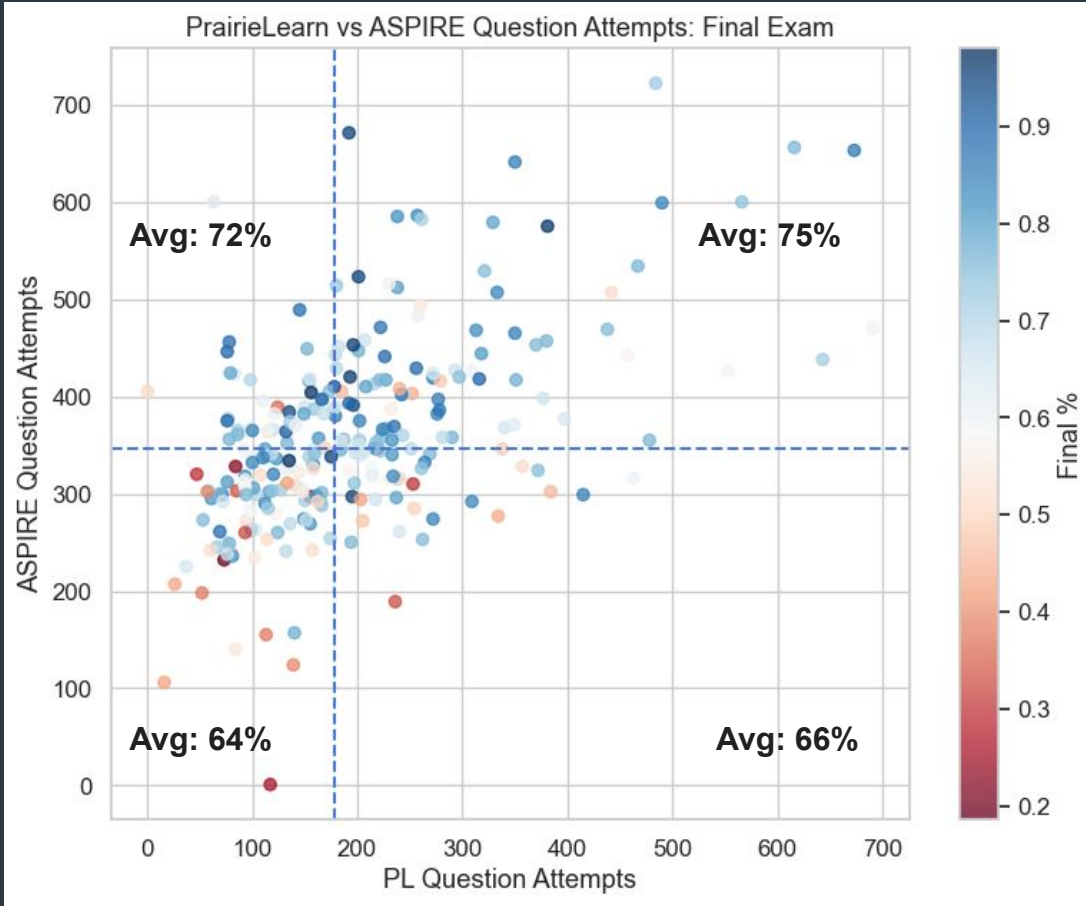
Dotted lines indicate median # of attempts for ASPIRE and PrairieLearn

More Practice → Improved Scores



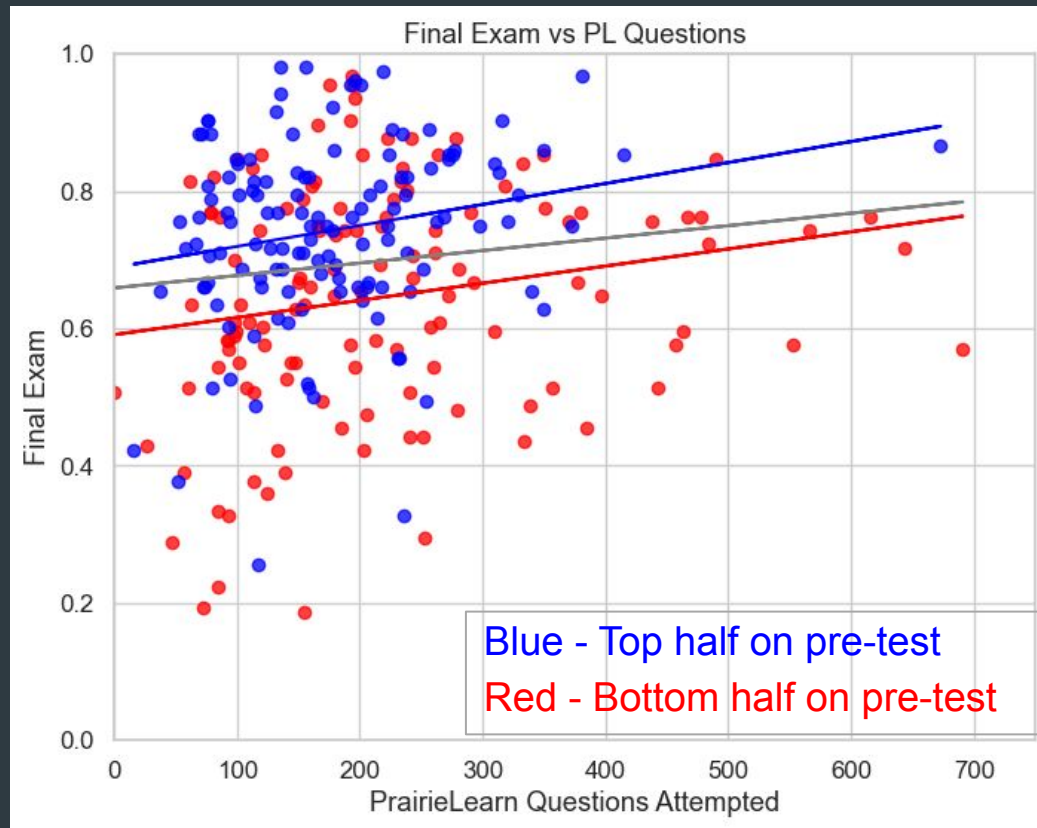
Dotted lines indicate median # of attempts for ASPIRE and PrairieLearn

More Practice → Improved Scores



Dotted lines indicate median # of attempts for ASPIRE and PrairieLearn

More Practice → Improved Scores



Course Resources & Exam Performance

