Discovering and removing barriers to learning

Ken Koedinger
Human-Computer Interaction & Psychology
Director of LearnLab
Carnegie Mellon University

NIPS Workshop on
Data Driven Education
Dec 10, 2013

General Argument

- Most of what we learn is outside our conscious awareness
- Thus, instructor intuition produces *flawed course materials*
- Instead: *Use data-driven models of learners to design instruction*
  - Machine discoveries predict data better than human-built models
- Instructional experiments demonstrate improved learning

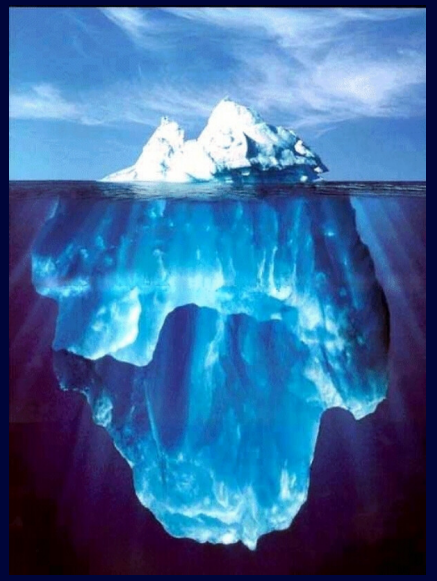

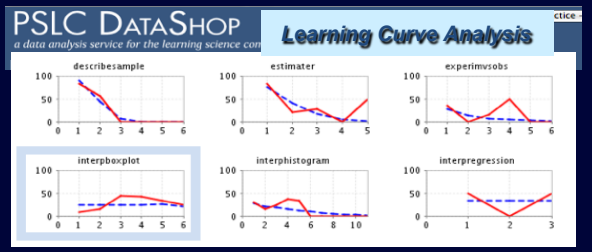

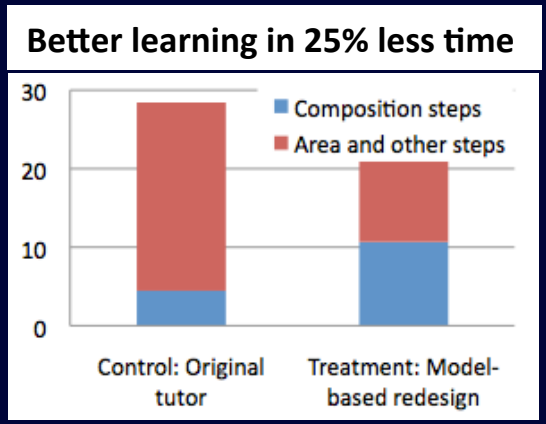

Cognitive Tutors:
Interactive Support for Learning by Doing

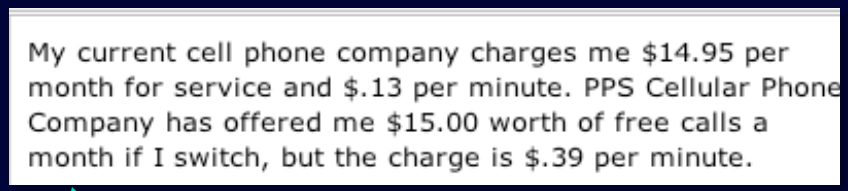

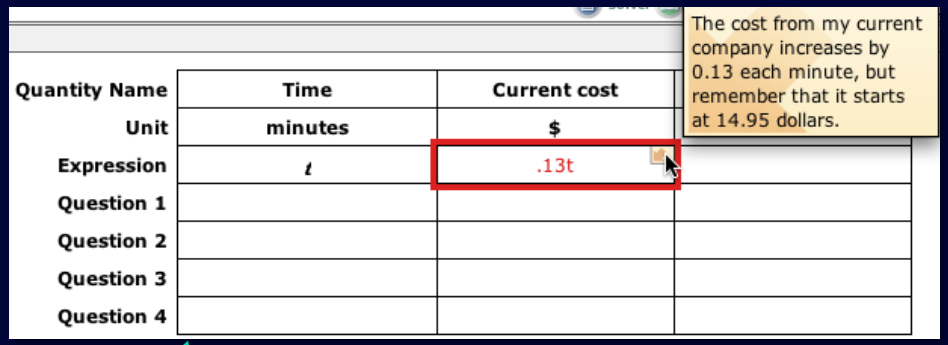

Authentic problems

Feedback *within* complex solutions

Progress…

Personalized instruction

Challenging questions

… individualization

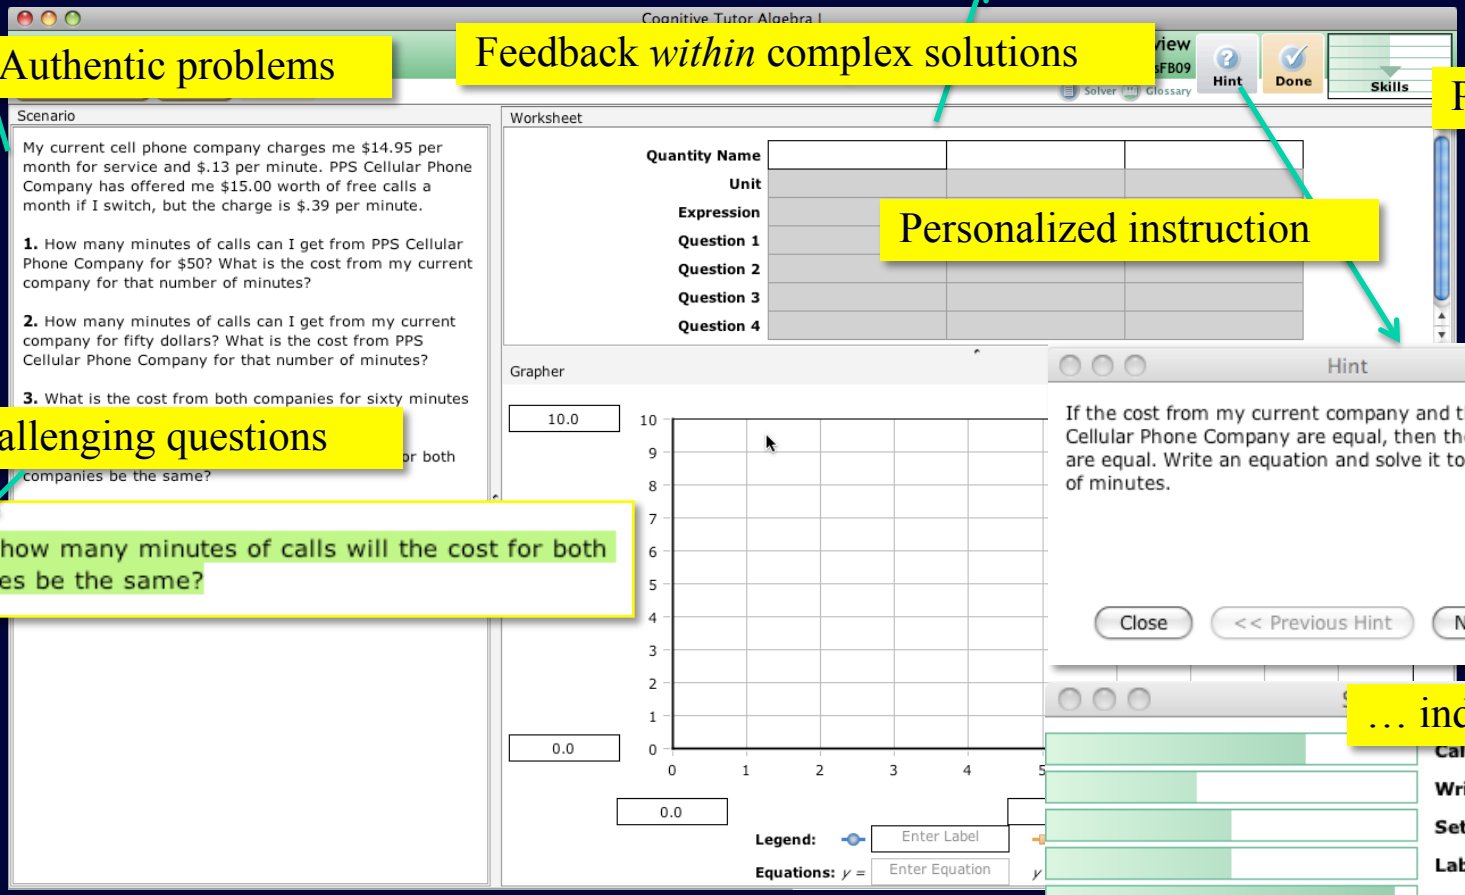

Cognitive model drives instruction

- Cognitive Model: A system that can solve problems in the various ways students can

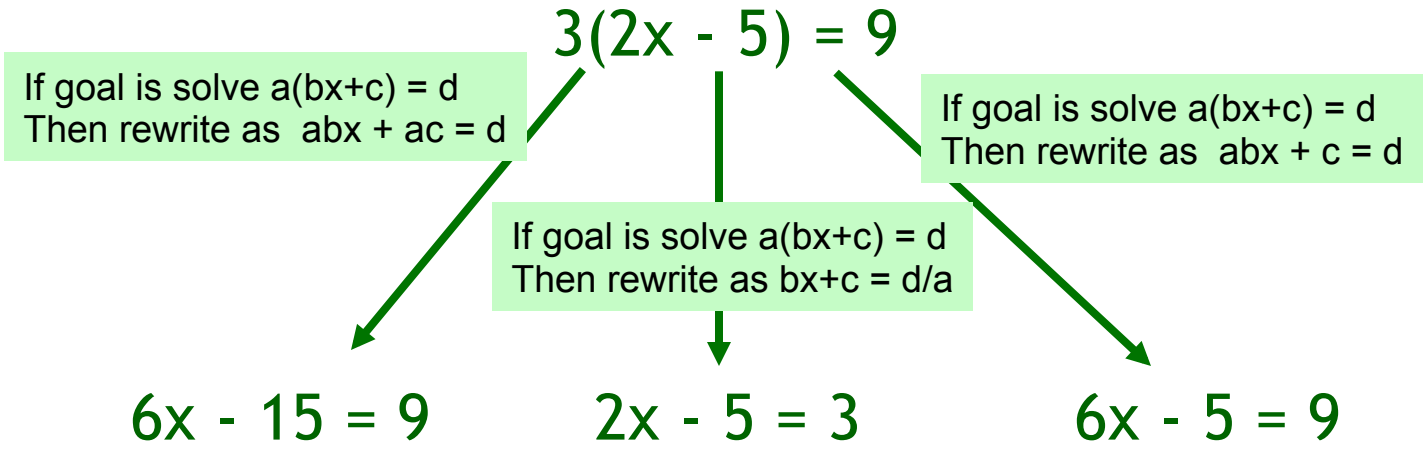

- Model Tracing: Follows student through their individual approach to a problem -> context-sensitive instruction

Cognitive model drives instruction

- Cognitive Model: A system that can solve problems in the various ways students can

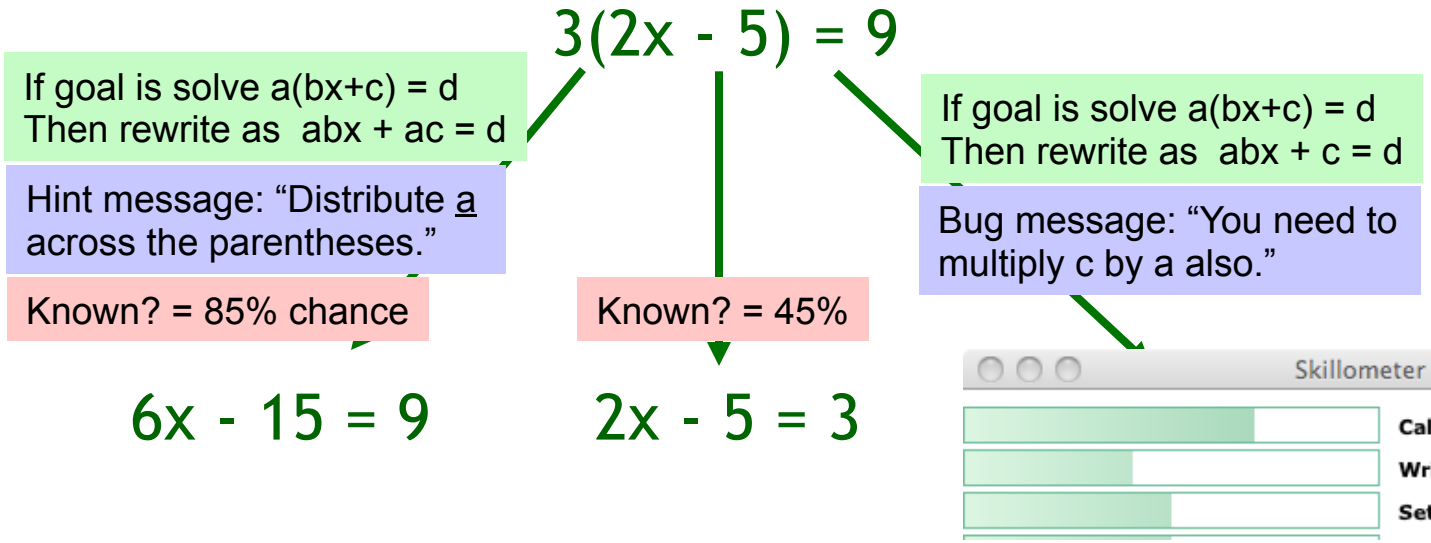

- Model Tracing: Follows student through their individual approach to a problem -> context-sensitive instruction
- Knowledge Tracing: Assesses student's knowledge growth -> individualized activity selection and pacing

You've had lots of experience with the English language.

You might say you know English.

But, do you know what you know?

Do educators know what they know?

Which is harder for algebra students?

Story Problem

As a waiter, Ted gets $6 per hour. One night he made $66 in tips and earned a total of $81.90. How many hours did Ted work?

Word Problem

Starting with some number, if I multiply it by 6 and then add 66, I get 81.90. What number did I start with?

Equation

x * 6 + 66 = 81.90

Math educators say:
story or word is hardest

Students:
equations
are hardest

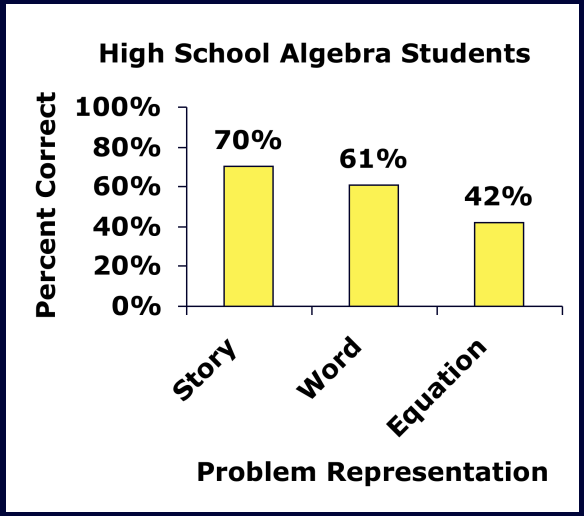

Expert blind spot!
Algebra teachers, especially, incorrectly think equations are easy

Experts can describe <30% of what they know!
(Clark et al)

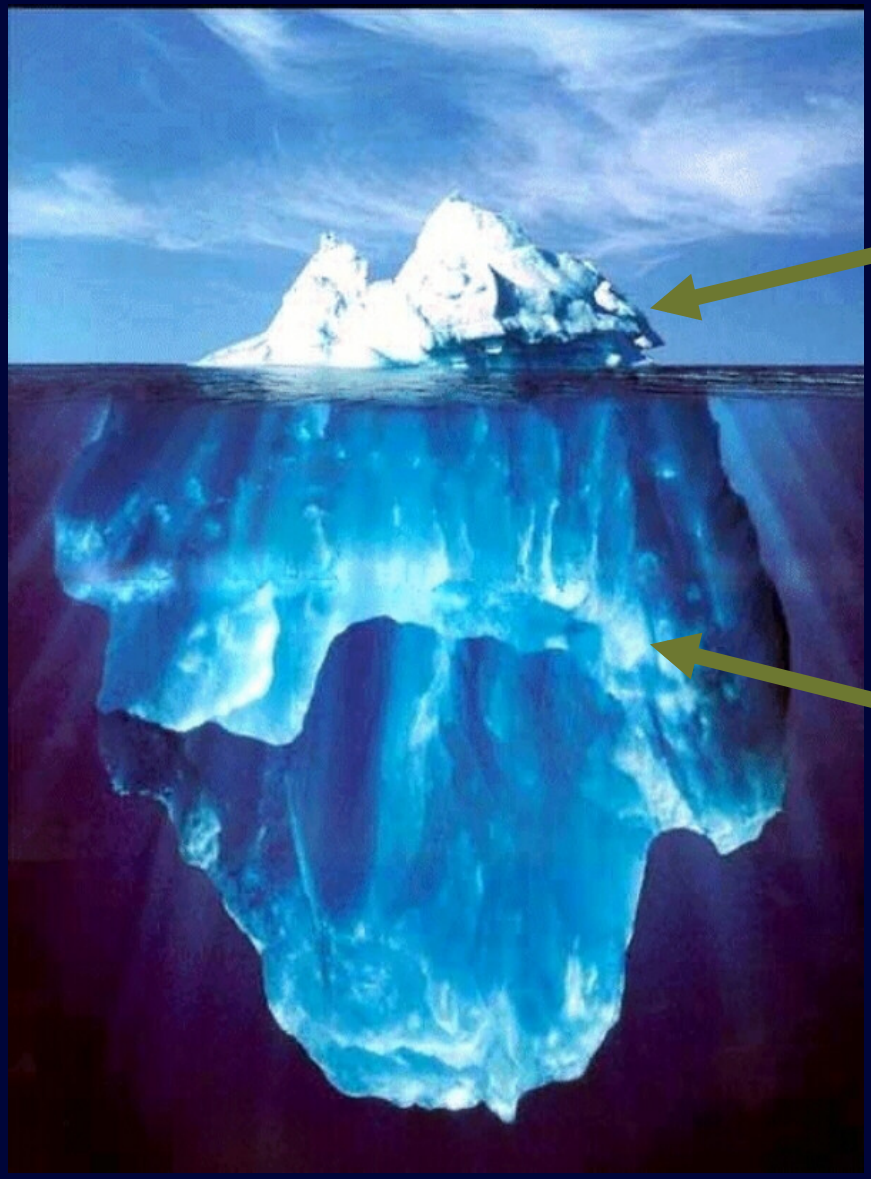

What we know about our own learning

What we do *not* know we know

You can't design well for what you don't know!

Repeat the mantra with me:

"I don't know what I know!"

"I don't know what I know!"

"I don't know what I know!"

Real World Impact of Data-Driven Cognitive Science

Algebra Cognitive Tutor

- Widespread intensive use
  ~600K students per year
  ~80 minutes per week
- Enhances student learning
- Still:
  Too many decisions driven by intuition

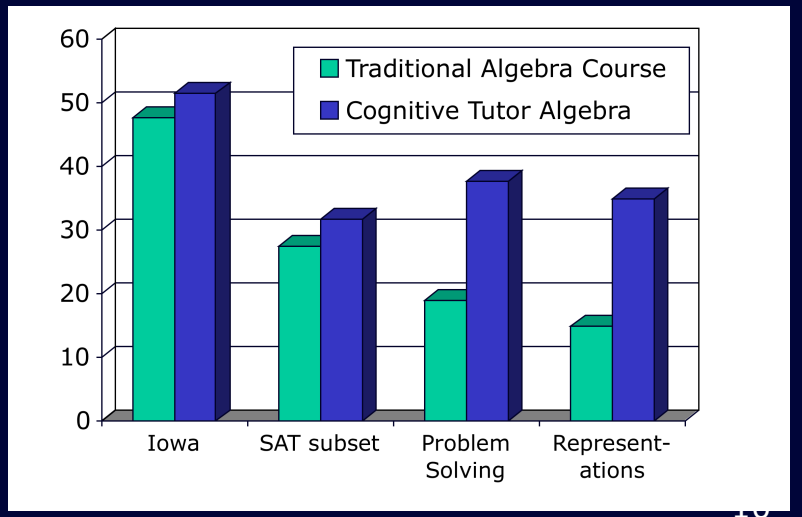

Koedinger, Anderson, Hadley, & Mark (1997). Intelligent tutoring goes to school in the big city.

Pane et al. (2013). Effectiveness of Cognitive Tutor Algebra I at Scale. Santa Monica, CA: RAND Corporation.

An Infrastructure for Course-Based *Data Collection & Improvement*

Ed tech + wide use = "Basic research *at scale*"

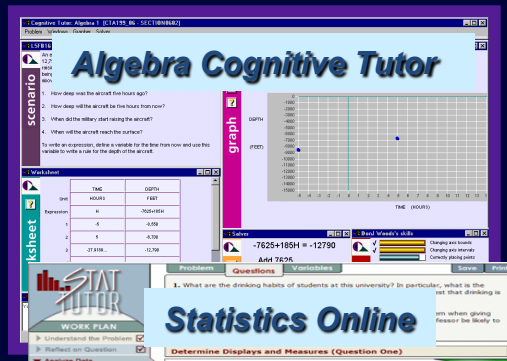

+

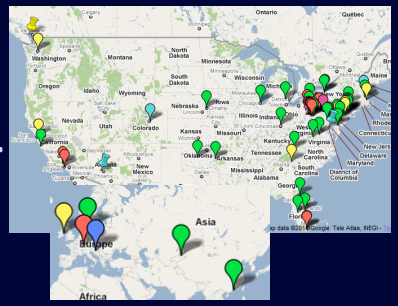

=

LearnLab
Pittsburgh Science of Learning Center

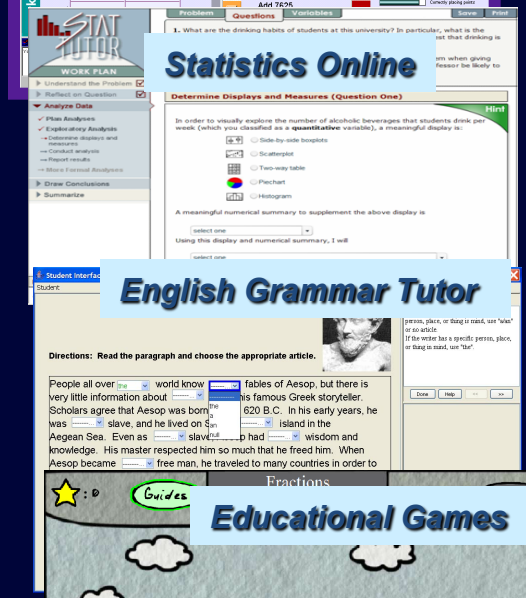

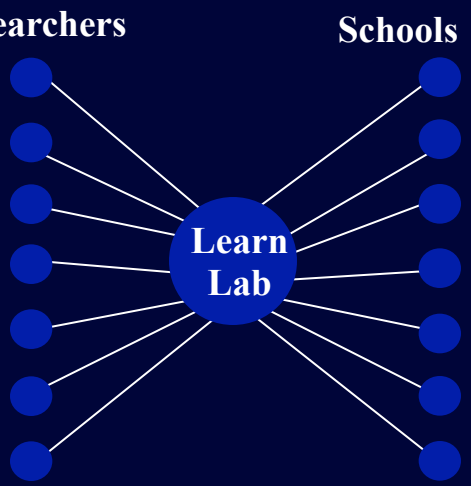

Since 2004
> 450 ed tech data sets in DataShop

> 320 *in vivo* experiments

Koedinger, Booth & Klahr (2013). Instructional complexity and the science to constrain it. *Science*.

Koedinger, Corbett & Perfetti (2012). The Knowledge-Learning-Instruction (KLI) framework. *Cognitive Science*.

Data-driven learner models of cognition+

- LearnLab collaborators have also modeled
  - Metacognition (Aleven, Roll, …)
  - Motivation (Baker, Nokes-Malach, …)
  - Student discussions (Rose, Resnick, …)

- Focus today: Discovering better *models of what is hard for students to learn*

Closing the data-improvement loop

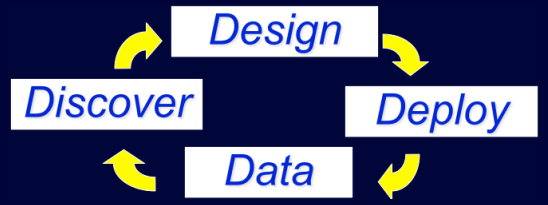

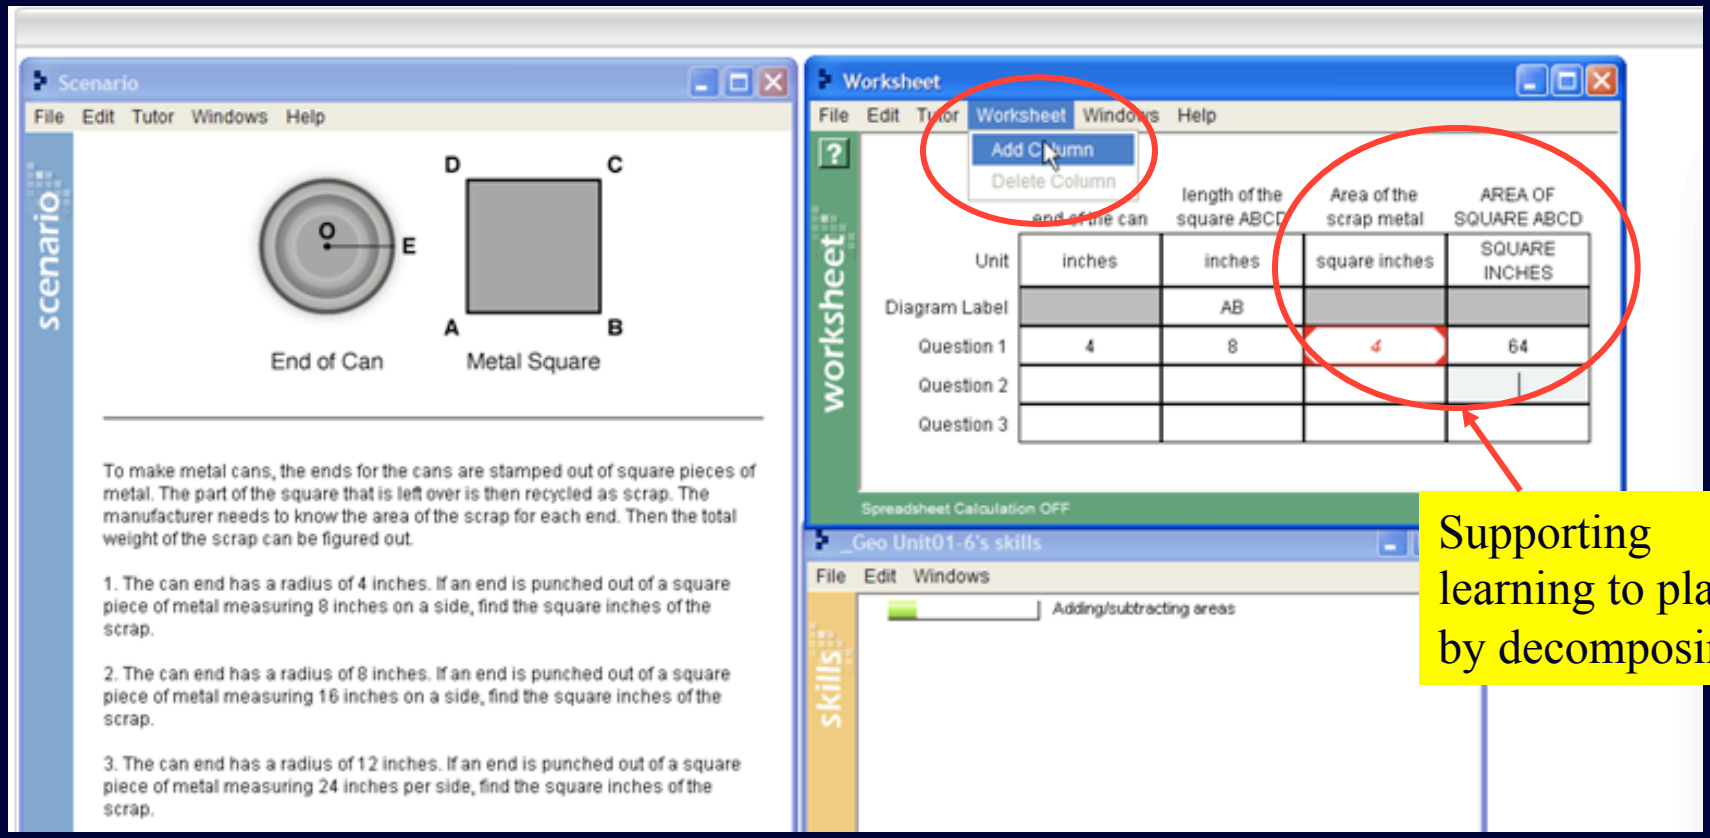

- Performance on many steps in many problems is logged & labeled …

Log data: Grade student steps & apply *label-based cognitive model*

Steps within problems that are assessed

Map of steps to knowledge comp's (Q-matrix)

Opportunities student has had to learn KC

Was student's first attempt on this step a correct one?

Student	Step (Item)	Textbook KCs	Opportunity	Success
A	prob1step1	Circle-area	1	0
A	prob2step1	Circle-area	2	1
A	prob2step2	Rectangle-area	1	1
A	prob2step3	Compose-by-addition	1	0
A	prob3step1	Circle-area	3	0

Single skill	Opportunity
Geometry	1
Geometry	2
Geometry	3
Geometry	4
Geometry	5

Different KC hypothesis changes prediction of student success

http://pslcdatashop.org

PSLC DataShop

a data analysis service for the learning science community

Geometry Area (1996-97)

logged in as kk1u@ANDREW.CMU.EDU logout
help | news | contact us

Dataset Info | Learning Curve | Error Report | Performance Profiler | Export

Line Graph | Step Rollup Table | LFA Values

Samples

deselect all

My Samples

Shared Samples

☑ **All Data**

Learning Curve

View By

- (•) Knowledge Component
- ( ) Student

Type

- ( ) Assistance Score
- (•) Error Rate
- [ ] View Predicted

Opportunity Cutoff

Min [ ] Clear

Max [ ] Clear

Refresh Graph

Knowledge Component Models

Primary: Geometry

Knowledge Components

deselect all | select all

☑ **Geometry**

Dataset: Geometry Area (1996-97)
Sample(s): All Data

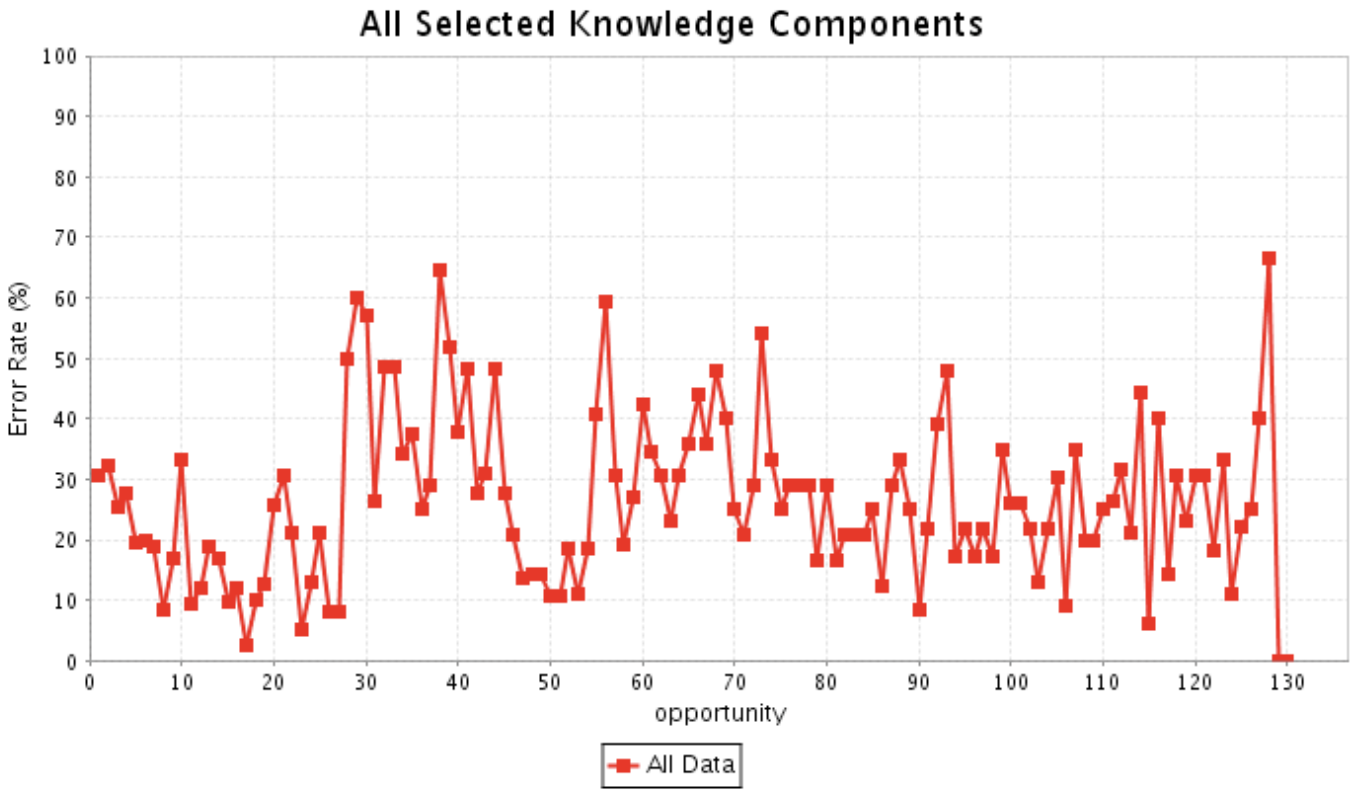

All Data

Opportunity Number	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Number of Observations	59	59	59	58	56	55	48	48	47	45	42	42	42	41	41	41	41	40	39	39
Opportunity Number	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40

Data mining using DataShop's learning curve analytics

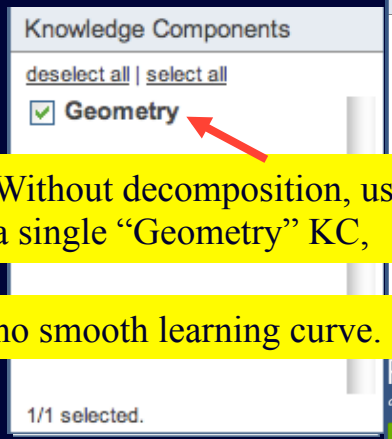

Without decomposition, using just a single "Geometry" KC,

no smooth learning curve.

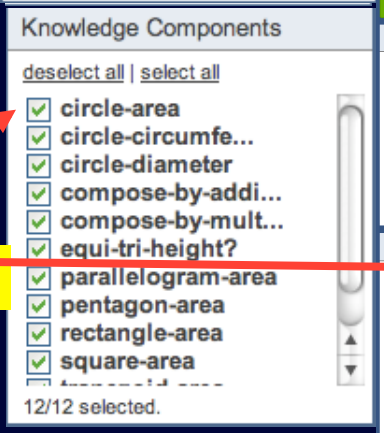

But with decomposition, 12 KCs for area concepts,

a smoother learning curve.

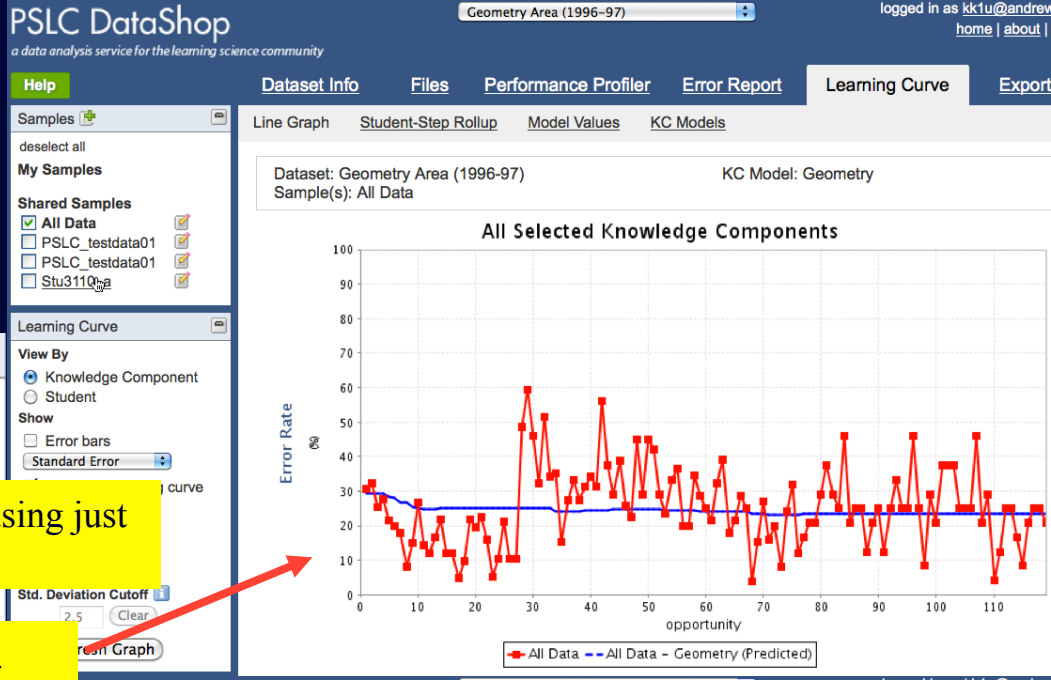

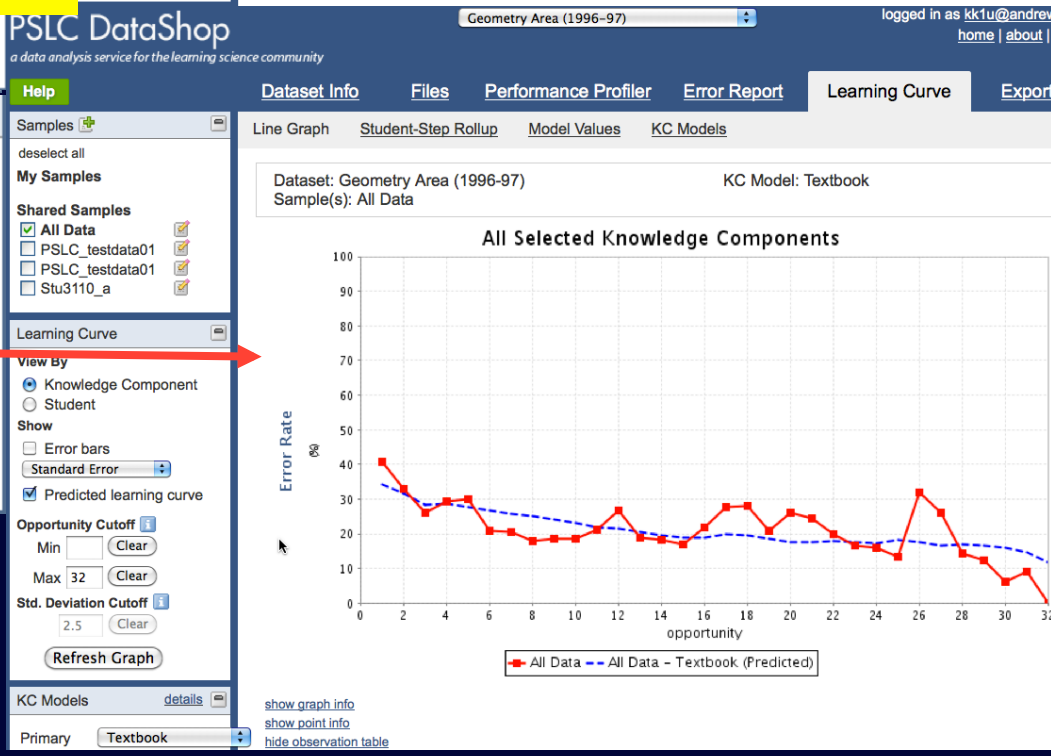

Visualizing learning curves to find opportunities for improvement

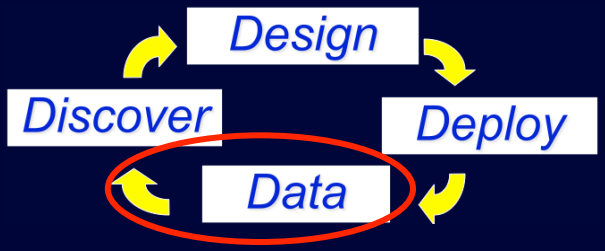

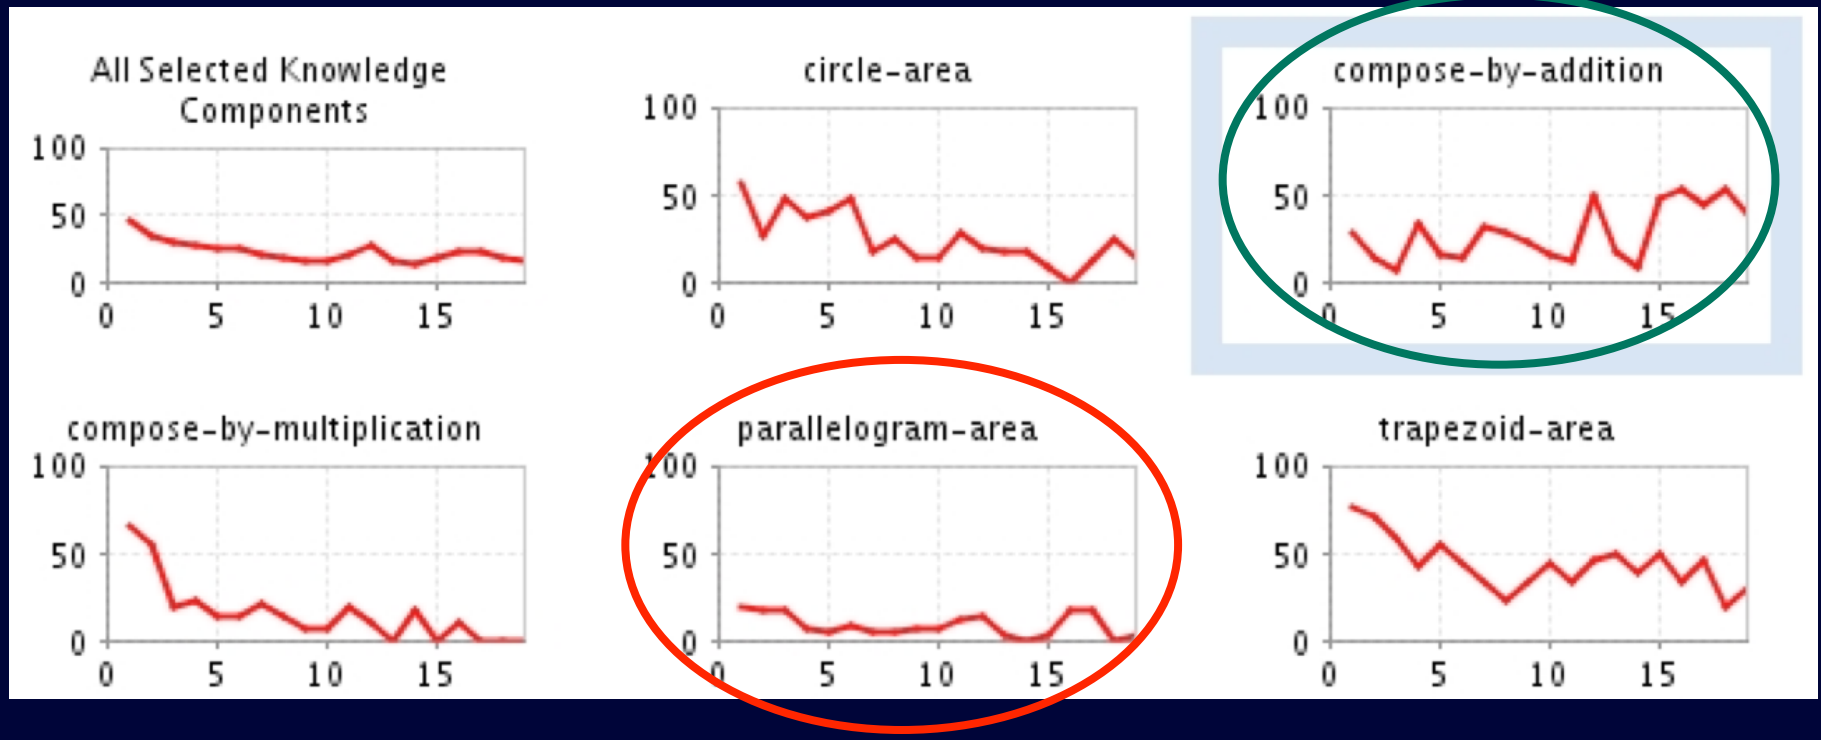

Low, long curve => remove busy work

Cen et al (2007)

High rough curve => concept/skill is more complex

Statistical version of cognitive model

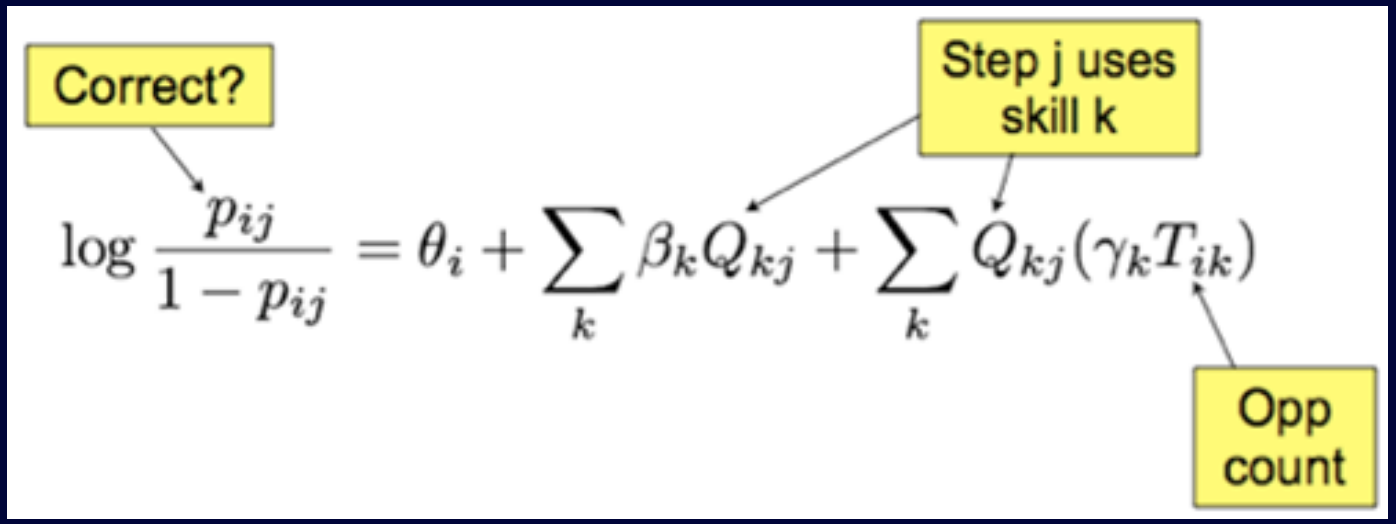

GIVEN:

- p_{ij} = probability student i gets step j correct
- Q_{kj} = each knowledge component k needed for this step j
- T_{ik} = opportunities student i has had to practice k

ESTIMATED:

- θ_i = proficiency of student i
- β_k = difficulty of KC k
- γ_k = gain for each practice opportunity on KC k

Cen, Koedinger, & Junker (2006)
Spada & McGaw (1985)

Machine learning to improve cognitive models, lower prediction error (RMSE)

Learning Factors Analysis algorithm uses *beam search*
- feature generation ops
- stat model as evaluation function

Like a genetic algorithm

Koedinger, McLaughlin, & Stamper (2012). Automated student model improvement. In *Proceedings of the Fifth International Conference on Educational Data Mining*. [Conference best paper.]

Dataset	Content area	RMSE		
		Orig in-use	Best-hand	Best-LFA
Geometry 9697	Geometry area	0.4129	0.4033	0.4011
Hampton 0506	Geometry area	NA	0.4022	0.4012
Cog Discovery	Geometry area	NA	0.3250	0.3244
DFA-318	Story problems	0.4461	0.4407	0.4405
DFA-318-main	Story problems	0.4376	0.4287	0.4266
Digital game	Fractions	0.4442	0.4396	0.4346
Self-explanation	Equation solving	NA	0.4014	0.3927
IWT 1	English articles	0.4262	0.4110	0.4068
IWT 2	English articles	0.3854	0.3854*	0.3806
IWT 3	English articles	0.3970	0.3965	0.3903
Statistics-Fall09	Statistics	0.3648	0.3527	0.3353

Variety of domains & technologies

11 of 11 improved models

Where do meaningful task features come from?

- Combine psychology & domain expertise

  Stick with me here …
  particularly if you're in CS or stats

  Say it with with me:

  "I don't know what I know"!

Same math but different difficulty!?

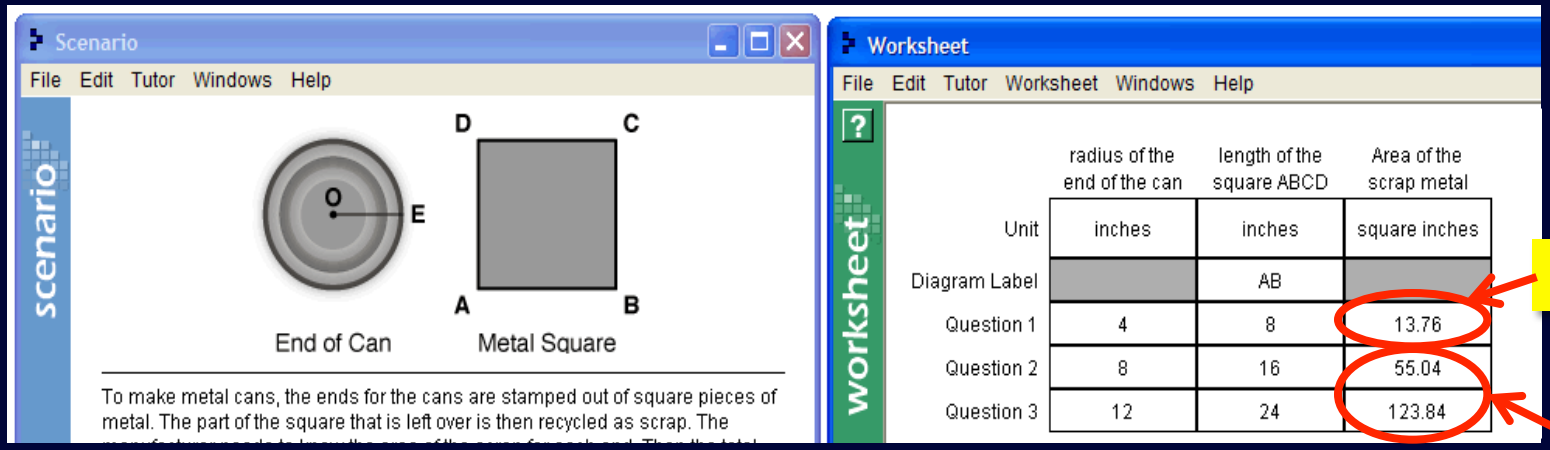

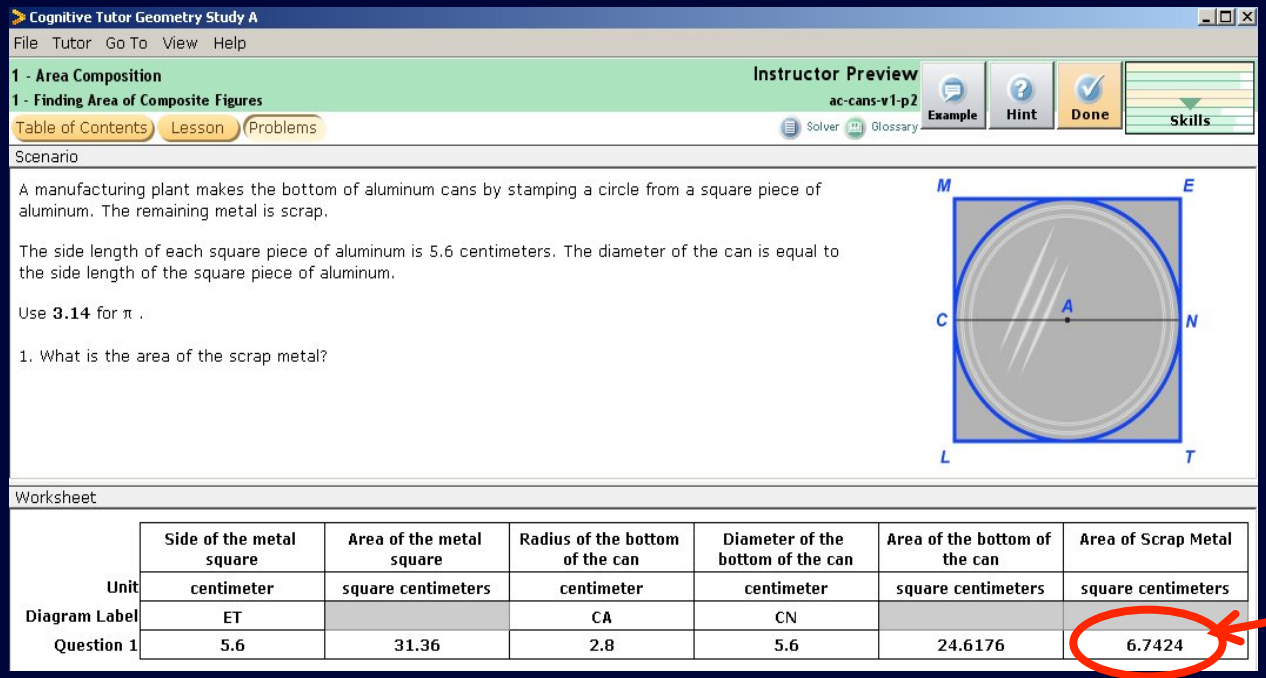

Answer: Different planning needs

Design
Discover
Deploy
Data

Scenario
File Edit Tutor Windows Help

scenario

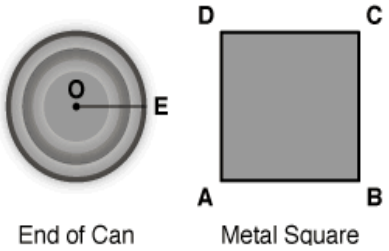

End of Can Metal Square

To make metal cans, the ends for the cans are stamped out of square pieces of metal. The part of the square that is left over is then recycled as scrap. The

Worksheet
File Edit Tutor Worksheet Windows Help

worksheet

	radius of the end of the can	length of the square ABCD	Area of the scrap metal
Unit	inches	inches	square inches
Diagram Label		AB	
Question 1	4	8	13.76
Question 2	8	16	55.04
Question 3	12	24	123.84

Plan to combine areas

Subtract

Cognitive Tutor Geometry Study A
File Tutor Go To View Help

1 - Area Composition
1 - Finding Area of Composite Figures
Instructor Preview
ac-cans-v1-p2
Example Hint Done Skills
Table of Contents Lesson Problems
Solver Glossary

Scenario

A manufacturing plant makes the bottom of aluminum cans by stamping a circle from a square piece of aluminum. The remaining metal is scrap.

The side length of each square piece of aluminum is 5.6 centimeters. The diameter of the can is equal to the side length of the square piece of aluminum.

Use **3.14** for π .

1. What is the area of the scrap metal?

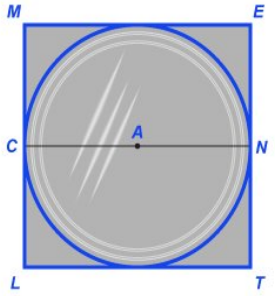

Scaffolding

Worksheet

	Side of the metal square	Area of the metal square	Radius of the bottom of the can	Diameter of the bottom of the can	Area of the bottom of the can	Area of Scrap Metal
Unit	centimeter	square centimeters	centimeter	centimeter	square centimeters	square centimeters
Diagram Label	ET		CA	CN		
Question 1	5.6	31.36	2.8	5.6	24.6176	6.7424

Combine Areas

Use data-driven model to redesign tutor

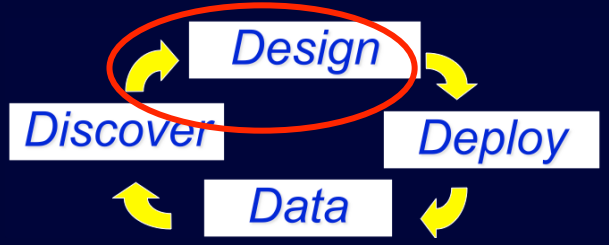

1. Change adaptive prob selection
   - New bars for discovered skills
   - Adjust optimization parameters

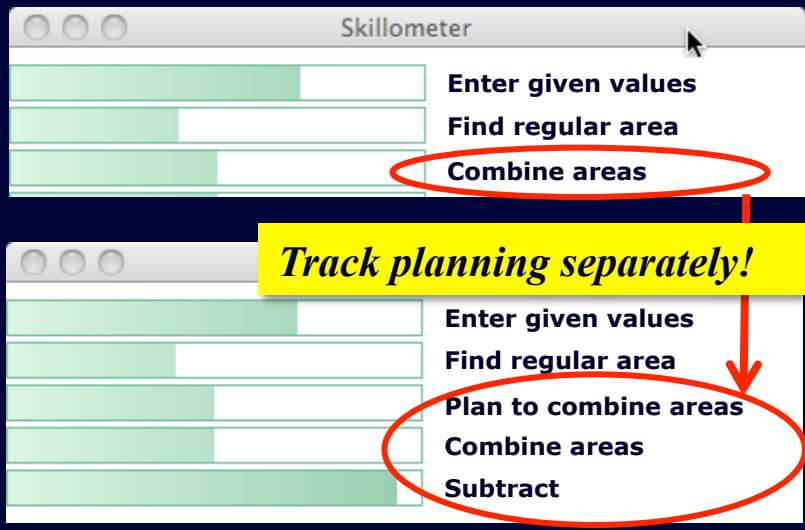

2. Sequence for gentle slope
   - Simple problems first

Worksheet

	Side of the metal square	Area of the metal square	Radius of the bottom of the can	Diameter of the bottom of the can	Area of the bottom of the can	Area of Scrap Metal
Unit	centimeter	square centimeters	centimeter	centimeter	square centimeters	square centimeters
Diagram Label	ET		CA	CN		
Question 1	5.6	31.36	2.8	5.6	24.6176	6.7424

Fade scaffolding!

3. Create new problems to focus on planning KCs
   - Next slide ..

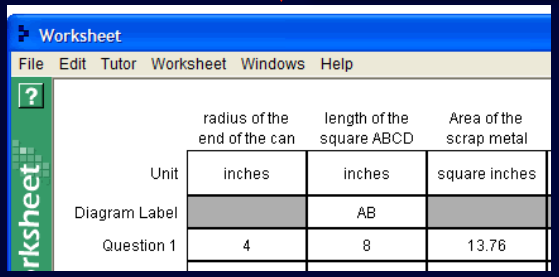

New problem type to *focus instruction* on greatest need

- Practice planning step *only*

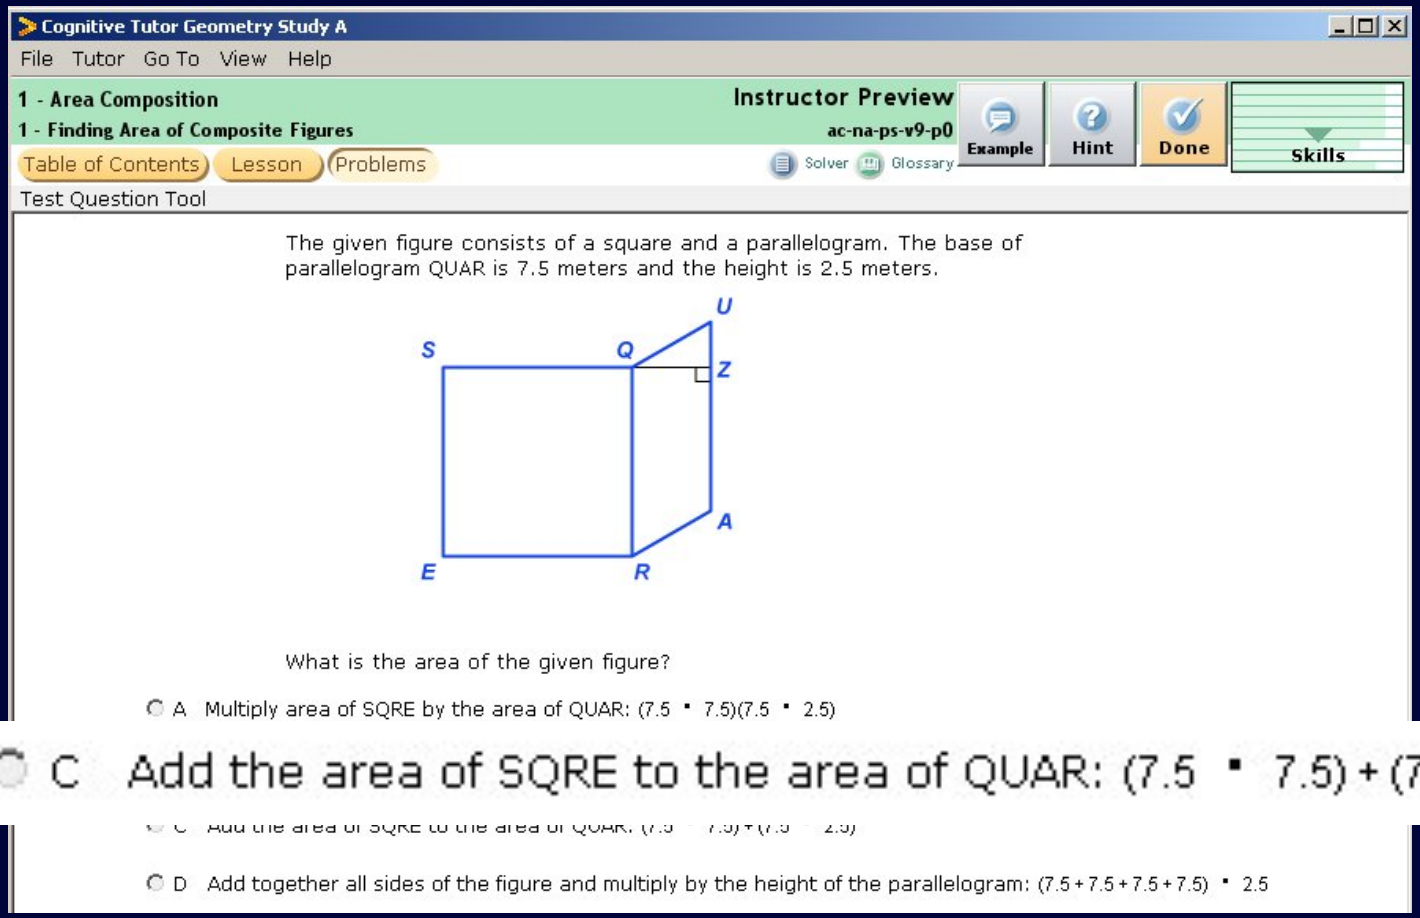

In vivo experiment: Random assignment A/B test embedded in ed tech used in schools

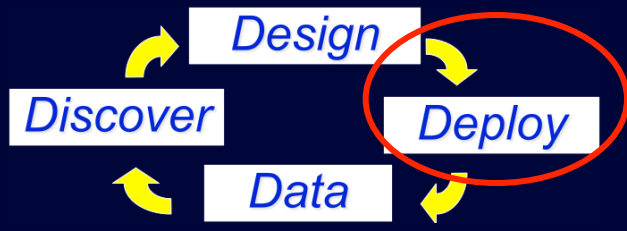

- Compare new vs. old tutors part of normal classroom use
- 103 high school students
- Measures:
  Pre-test, post-test,
  time to mastery

Model-based instruction is a better student experience

- More efficient: 25% less time

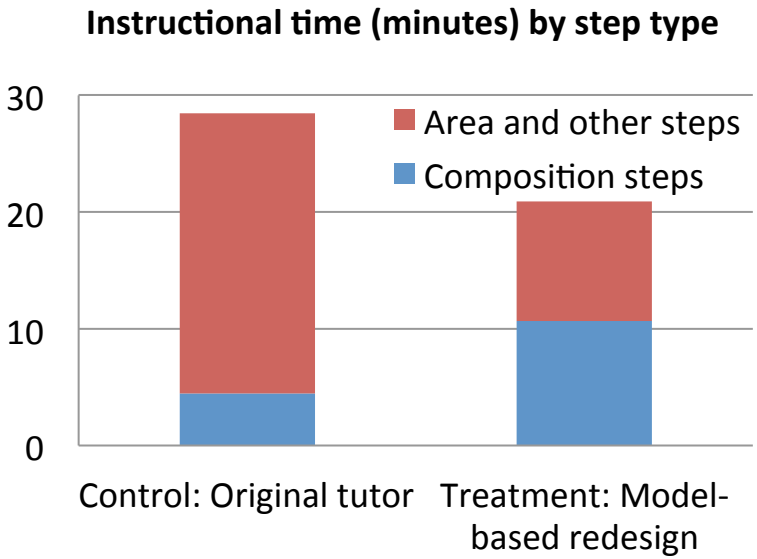

- And better learning of planning skills

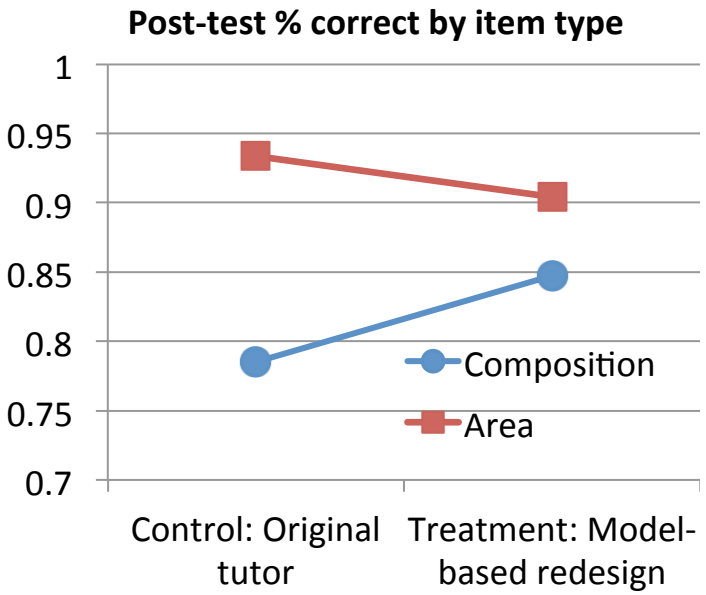

Where do meaningful task features come from?

- Past successes combine psychology & domain expertise
- Matrix factorization, deep belief networks *do not work*
  - Better predictions, no actionable insight
- Crowdsource task labeling
  - What makes tasks A & B harder than C & D
- Model *actual student learning!*

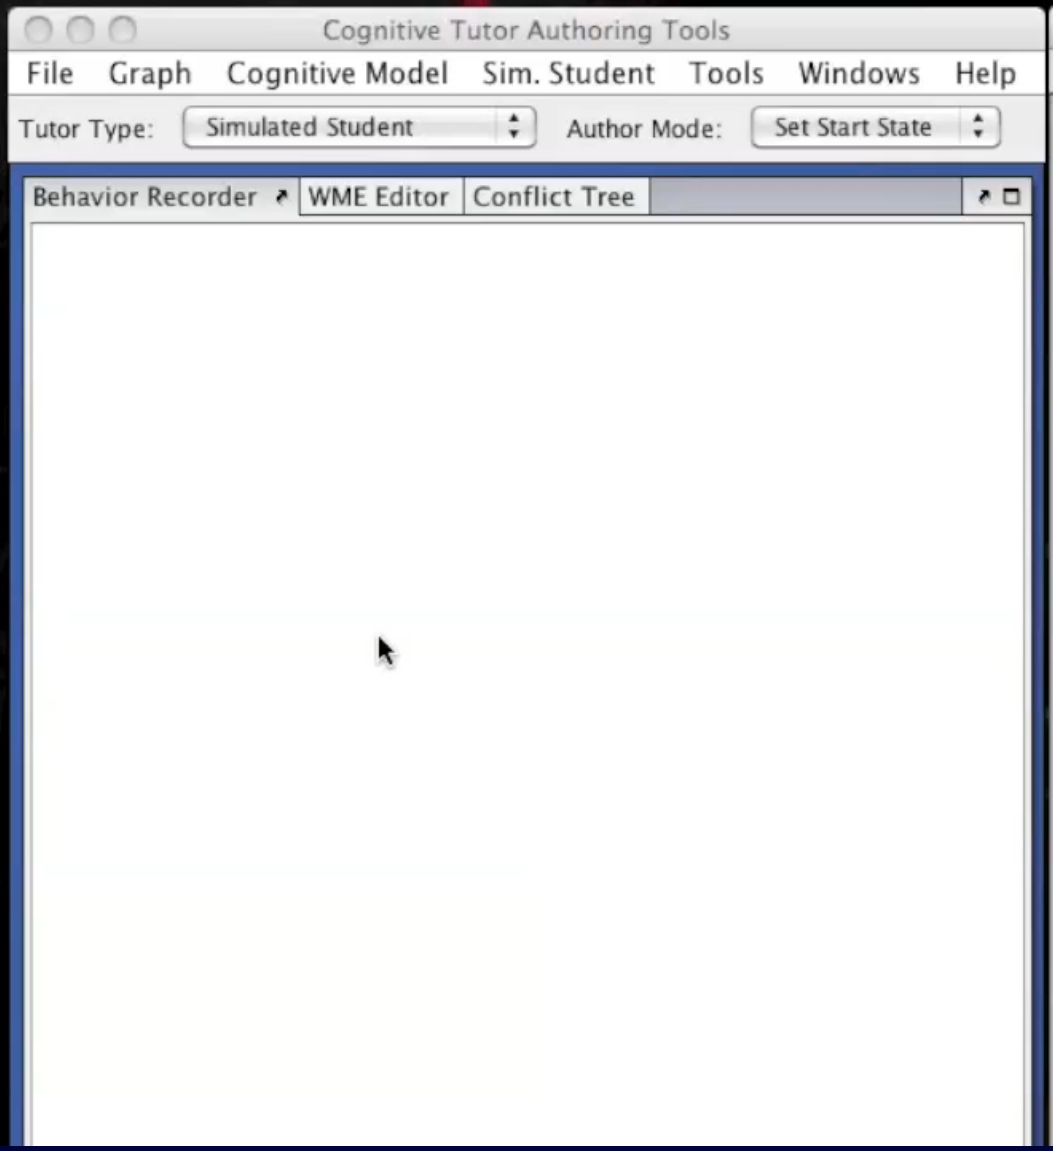

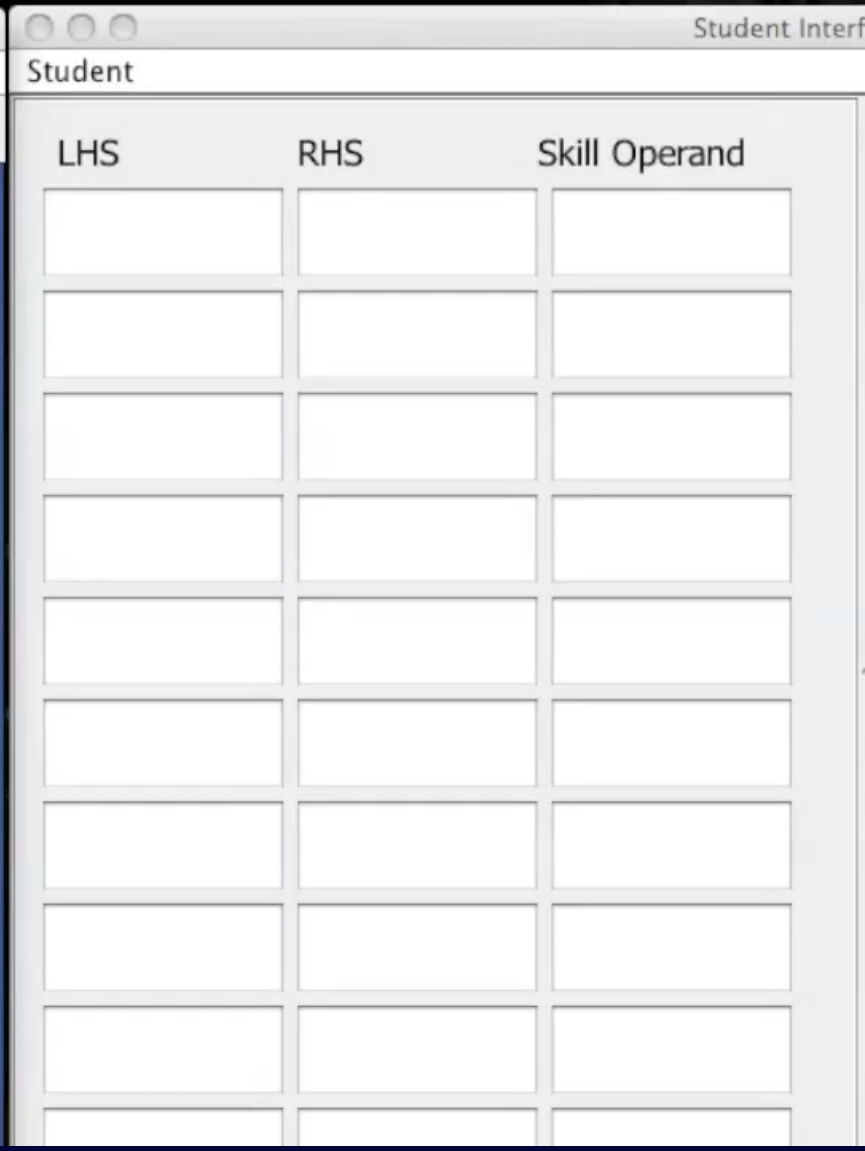

SimStudent learns like students do

- Representations from co-occurrence frequency via grammar induction
- Skills from imitation & feedback via hybrid of logic induction mechanisms

SimStudent Discovers Hidden Skills

- Tutor SimStudent to create cognitive models
- Evaluate model
  - Productions => features
  - Predict learning curves
- Result: More accurate cognitive models
- Interpretable & actionable

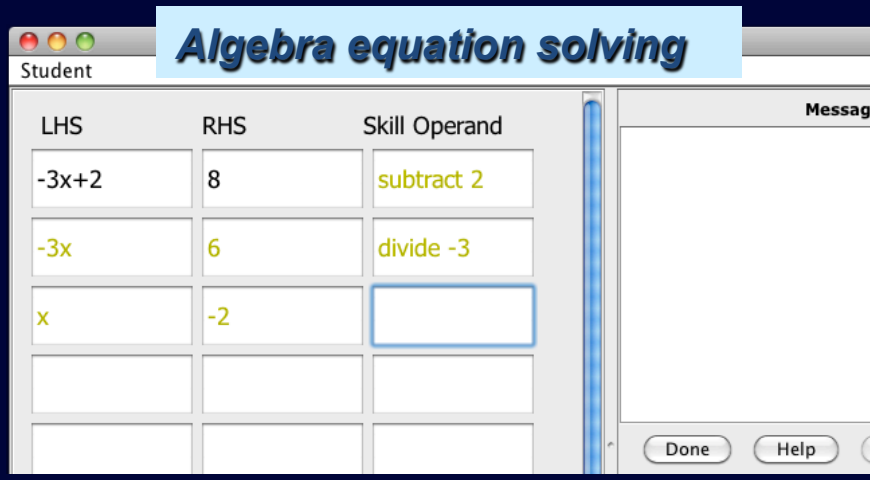

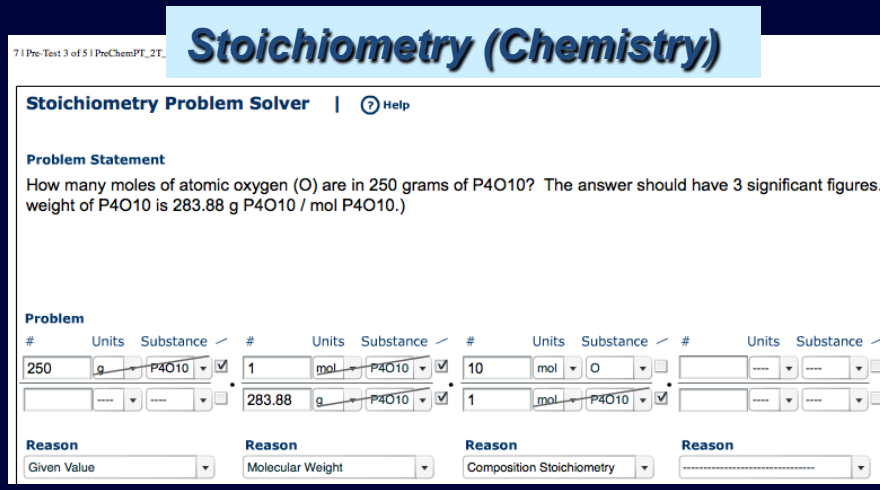

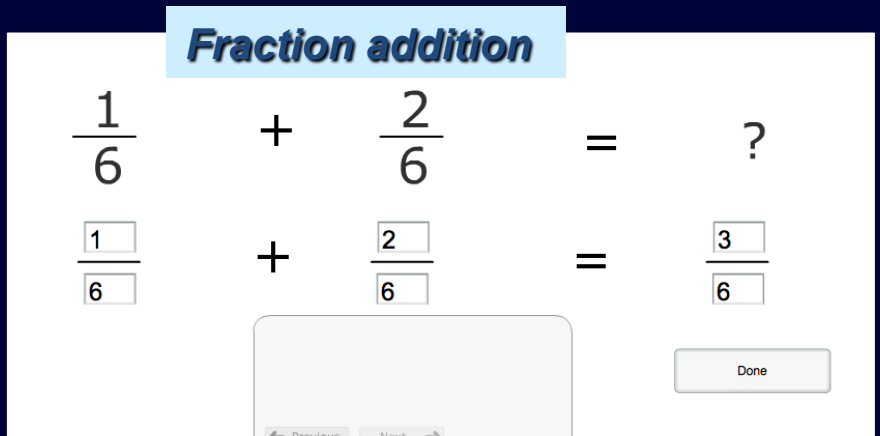

SimStudent discovered this tacit student difficulty in learning the algebraic representation!

Step
2x=12
6=3x
-x=5
-24=-4x

Data to model to better education

- You don't know what you know!
  - Intuitive design is not reliable
- Need data-driven *explanatory models* of learners
  - SimStudent potential: create intelligent tutors by teaching, pedagogically more effective than human-built tutors
- Improve science & practice of learning

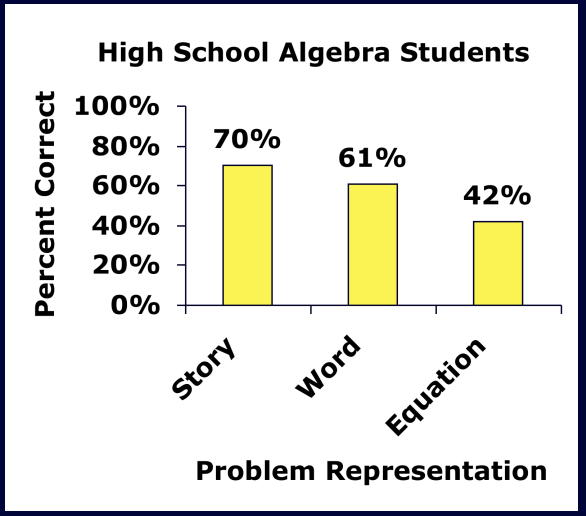

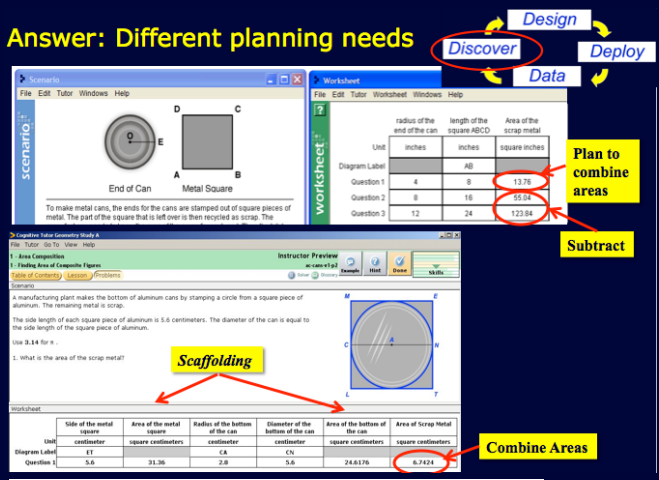

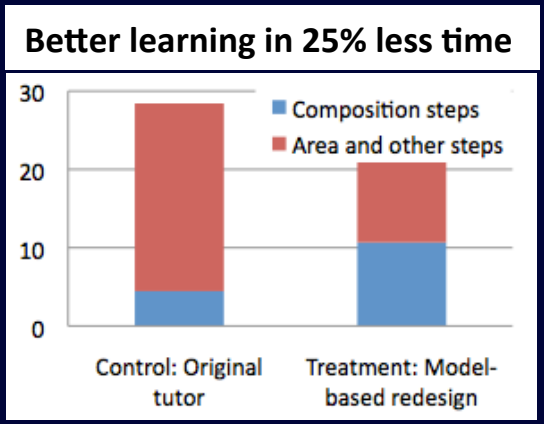

LearnLab
Pittsburgh Science of Learning Center

National Science Foundation
WHERE DISCOVERIES BEGIN

Thank you!

And thanks to many funders & colleagues:

- Funders

  ONR, Heinz/Grable/Mellon/Buhl/Pittsburgh foundations, Dept of Ed IES, National Science Foundation

- Cognitive Tutors

  John R. Anderson, Albert Corbett, Steve Ritter, …

- Algebra & Geometry Studies

  Mitchell Nathan, Mimi McLaughlin, John Stamper, Tristan Nixon …

- LearnLab
  - Charles Perfetti, David Klahr, Lauren Resnick, Vincent Aleven, Kurt VanLehn, Carolyn Rose …
  - All 200+ past & current members!

Different learned representation => different learned skills

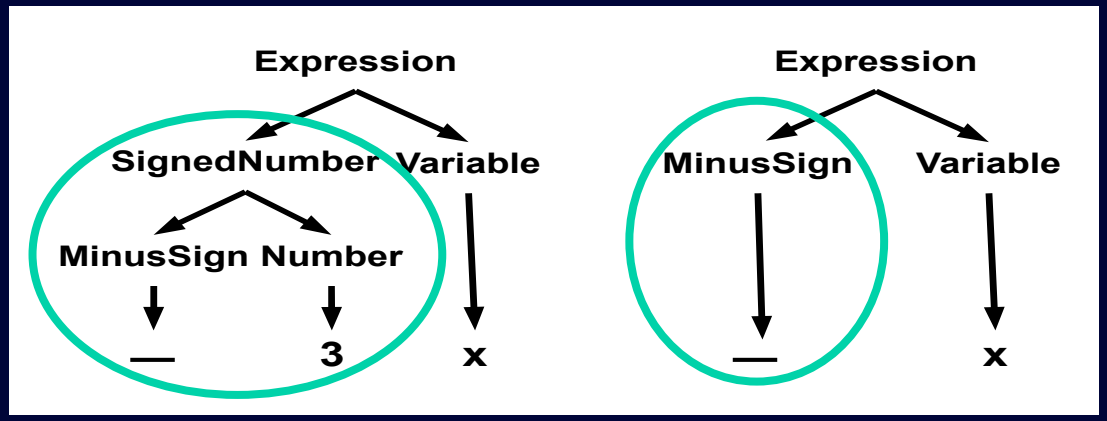

Skill simSt-divide (e.g. -3x = 6)

- Retrieval path:
  - Left side (-3, -3x)
  - Right side (6)
- Precondition:
  - Left side (-3x) does not have constant term
- Function sequence:
  - Divide both sides with the coefficient (-3)

Skill simSt-divide-1 (e.g. -x = 3)

- Retrieval path:
  - Left side (-x)
  - Right side (3)
- Precondition:
  - Left side (-x) is of the form -v
- Function sequence:
  - Generate one (1)
  - Divide both sides with -1

Request: Help build learning data infrastructure

- Recommendation #4 of 5 in Science:
- Worth doing?
- Looking for "DataLab partners"
  - Contribute data, methods, or infrastructure
  - Share same
  - LearnLab DataShop is seed

Abstract

We have developed analytic methods to discover barriers to student learning from data for educational technology use (see learnlab.org). Such discoveries can guide the redesign of instruction and our online experiments demonstrate enhanced learning outcomes. Our analytic methods span issues of student skill acquisition, metacognition, and motivation. Focusing on the first, I will illustrate how alternative cognitive models of learning can be evaluated by translating them to statistical models and predicting learning curve data for model comparison. We have used machine learning in a couple of ways to generate alternative cognitive models, one more practical and the other more cutting edge. The second involves a computational model of student learning, SimStudent, that learns as students do by using an intelligent tutoring system. The cognitive models SimStudent acquires have been demonstrated to yield empirically-verified discoveries not present in the human-designed cognitive models behind the intelligent tutoring systems. In other words, with SimStudent is the potential to not only create intelligent tutoring systems without AI programming, but to also produce systems that are pedagogically more effective than human-built tutoring systems.

Data-driven learner models

- Accurate assessment during learning
- Models of cognition, metacognition, motivation, discussion
- *Discover better models of what is hard to learn*

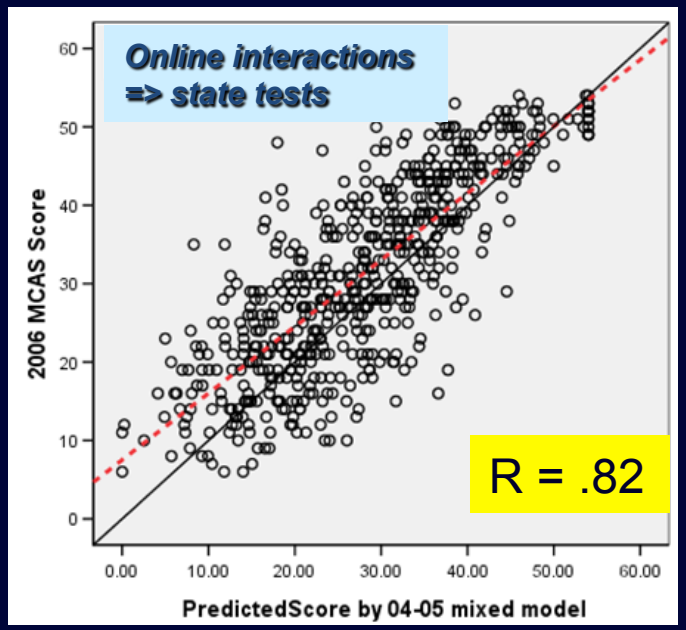

Data => discovery => outcomes

- Discover barriers to student learning
  - Change instruction & produce better learning
  - Domain skills, learning strategies, dispositions
- Use data from graded online activities
  - Translate cognitive to statistical models
  - Predict student performance over time
- Machine learning generates cognitive models
  - Feature gen & search (“genetic algorithm”)
  - SimStudent learns as students do
- Proving SimStudent works
  - Creates intelligent tutors without programming
  - more effective than human-built tutoring systems

Expertise is represented as display-based production rules

- Skill divide (e.g. -3x = 6)
  - What:
    - Left side (-3x)
    - Right side (6)
  - When:
    - Left side (-3x) does not have constant term

  =>
  - How:
    - Get-coefficient (-3) of left side (-3x)
    - Divide both sides with the coefficient

Learning mechanisms to generalize

- Perceptual information retrieval
- Preconditions on actions
- Action plans

- Original model required strong domain-specific functions, like Get-coefficient
  - More knowledge engineering required
  - Did not capture distinctions in student learning errors
    Like 3x=6 vs. -3x = 6