Agenda Items

- Big questions for CMDM domain modeling
- How to pose domain (assessment) model discovery problem
  - What is the general question that LFA, POKS, ePCA, Q-matrix, Rule Space is trying to answer?

Developing Cognitive Models of Academic Domains: How Can Educational Data Mining Help?

Ken Koedinger

Human-Computer Interaction & Psychology

Carnegie Mellon University

Focal Questions of CMDM Thrust

Today

1. How can we generate accurate cognitive models of students' *domain-specific* knowledge?
2. What models of *domain-general* processes best capture student learning?
   - learning & metacognition
   - motivation & affect
   - social aspects and instructional talk
3. By integrating domain-specific & -general models into *predictive models,* how can we *engineer* instructional interventions with big impact?

Domain-Specific Cognitive Models

- Question: How do students represent knowledge in a given domain?
- Answering this question involves deep domain analysis
- The product is a cognitive model of students' knowledge

Previously Stated Goals in Discovering Domain Models

- Improve model-discovery methods
  - Learning Factors Analysis
  - Partial Order Knowledge Structures (POKS)
  - Exponential-family Principle Component Analysis
- Improve human-machine interaction
  - Better process for task difficulty factor labeling
- Show models yield improved student learning

Want to expand that discussion

- What are the big, high impact scientific and practical questions?

What are important or interesting research questions for domain modeling?

Scientific questions:

- How do theories of learning constrain domain model discovery?
  - Is frequency of practice enough? What do we miss just using it?
- How can we infer students' domain rep's from performance & learning data?
- Can we model transfer (generalization of knowledge) just from performance data (at a single time point without)?
- What types of models are best (factors vs. clusters vs. mechanisms/rules/logic vs. chains of inference) for predicting performance?
  - Credit assignment in chains

Practical questions:

- How can domain modeling lead to improvements in student learning?
- How do we leverage human input in the discovery process?
  - Pose original model
  - Pose factors
  - Analyze cluster (of problems) probability

Research Questions for Domain Modeling

Scientific questions

- What's the grain size of transfer?
- Strategies or deep concepts?
- Is implicit learning helpful in math/science?
  - Role of noise in learning?
  - Is learning "statistical" in well-structured domains?
- How/when are deep features acquired?
- Does grain size & latency change across domains?
- Sig individual differences in task specialization or learning rate?

Practical questions

- Do better domain models yield better tutors?
  - Better problems, hints, examples, scheduling
- Does more instruction on X make Y easier to learn?
- Can we teach hidden skills?
- Does prerequisite remediation accelerate future learning?
- Does KC complexity affect what kind of instruction is best? Drill vs. self-expl …
- Add when & how to all of these

The Long-Standing Transfer Debate

- General: Faculty theory of mind
  - Mind has faculties that can be exercised with Latin, Geometry, … video games, n-back task*

  => Transfer is broad & general, across domains
- Specific: Thorndike's identical elements
  - Mind is made up of stimulus-response elements
    - Transfer occurs between tasks with common elements
    - 1922 study: Schooled task "multiply x^a & x^b" fails to transfer to "multiply 4^a & 4^b"

  => Transfer is narrow, within domains & tasks

*Jaeggi, Buschkuehl, Jonides, & Perrig (2008). Improving fluid intelligence with training on working memory.

More recent views

- Singley & Anderson: Learning & transfer occurs at the grain size of the production rule
  - But begs the question:
  - What is the grain size of a production rule?
- Others?

What's the grain size of transfer?

- Where do you place your bets?
  - Transfer is broad:
  - Transfer is narrow:

How can Educational Data Mining (EDM) help address the transfer question?

- What kinds of data and analysis techniques allow inferences about transfer?
  - Standard approach: Is performance on task B better after learning task A than it is without learning task A
- Are other approaches possible?
  - Learning curves?
  - Knowledge spaces (POKS)?

Learning curves examples

- Geometry Area in DataShop
- Excel learning curves …

Spreadsheet formulas: *Absolute referencing* stalls many learners

	A	B
1		Hourly Wage
2	Hours Worked	$10
3	20	=A3*B2
4	30	=A4*B3
5	40	=A5*B4

(a)

	A	B
1		Hourly Wage
2	Hours Worked	$10
3	20	$200
4	30	6000
5	40	240000

(b)

	A	B
1		Hourly Wage
2	Hours Worked	$10
3	20	=A3*B$2
4	30	=A4*B$2
5	40	=A5*B$2

(c)

	A	B
1		Hourly Wage
2	Hours Worked	$10
3	20	$200
4	30	$300
5	40	$400

(d)

Theoretical interest: Difficult domain reveals edges of human learning

Faculties of mind theory of transfer?

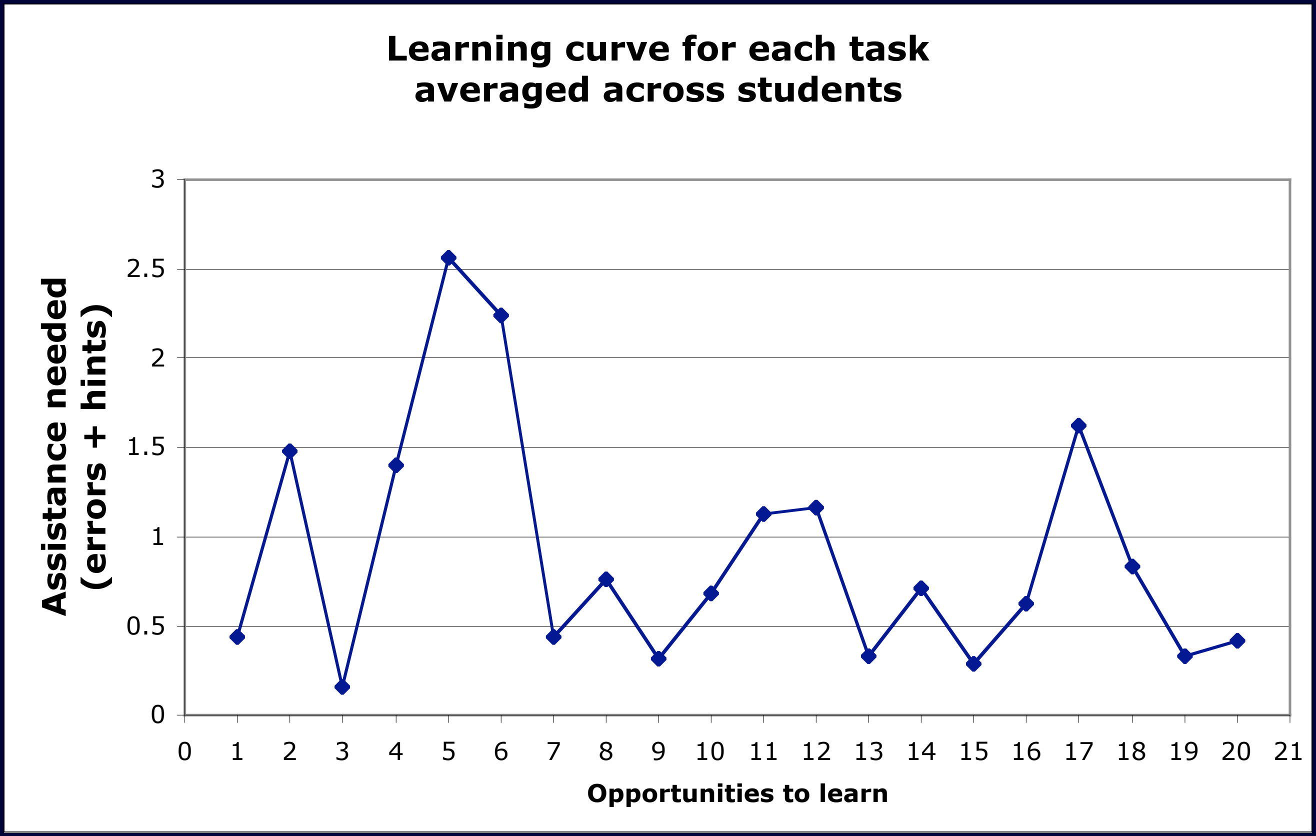

- Exercising “reasoning faculty” or “programming faculty” does *not* (clearly) reveal improvement

Identical elements theory of transfer?

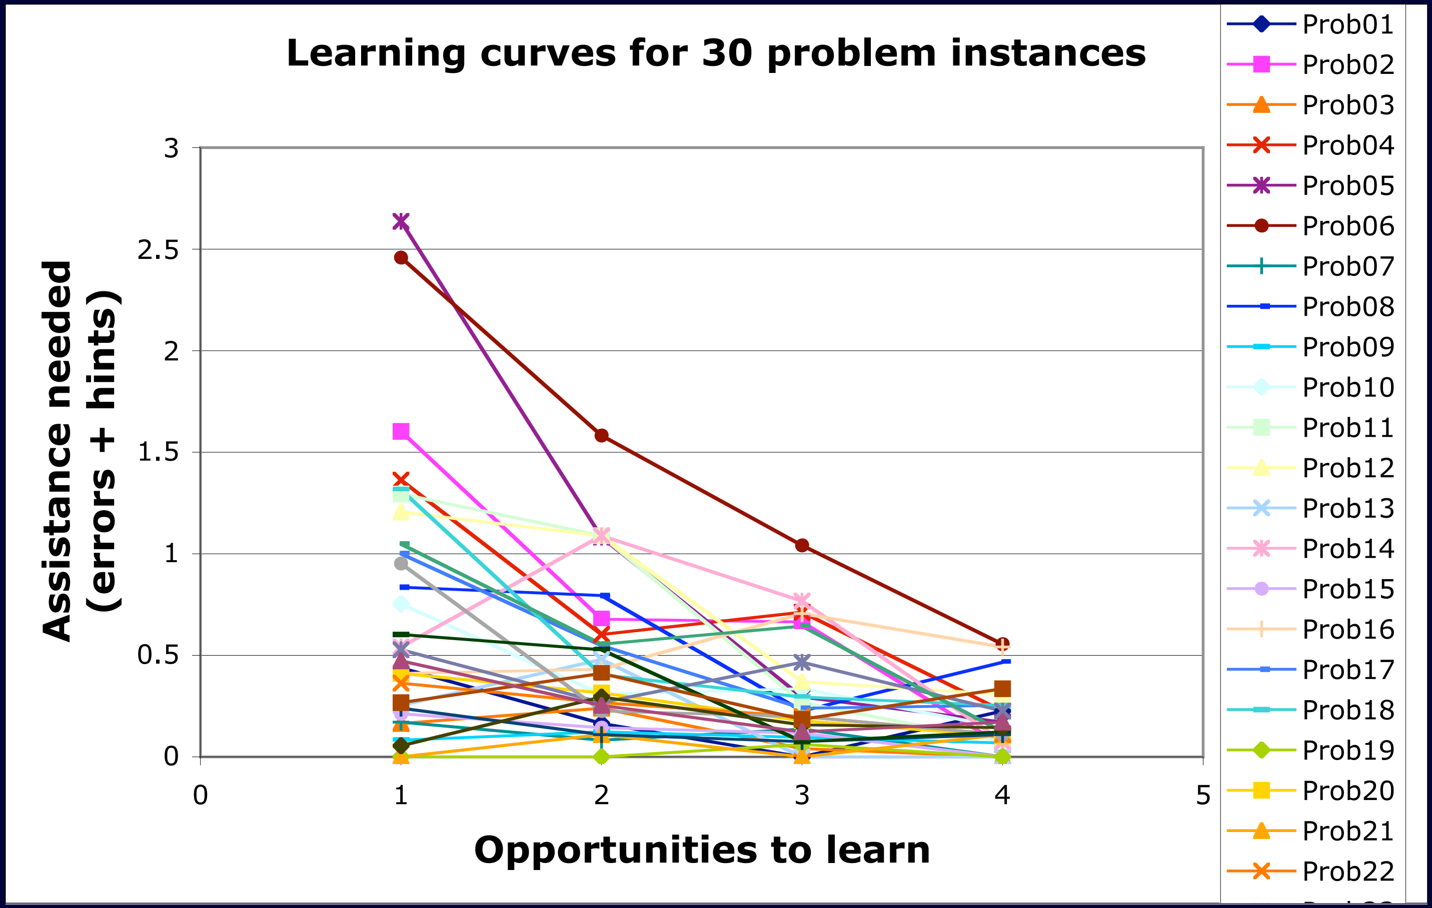

- Maybe, but perhaps too fine
- Are there groups of items across which we see transfer?

Pattern emerges at more coarse grain

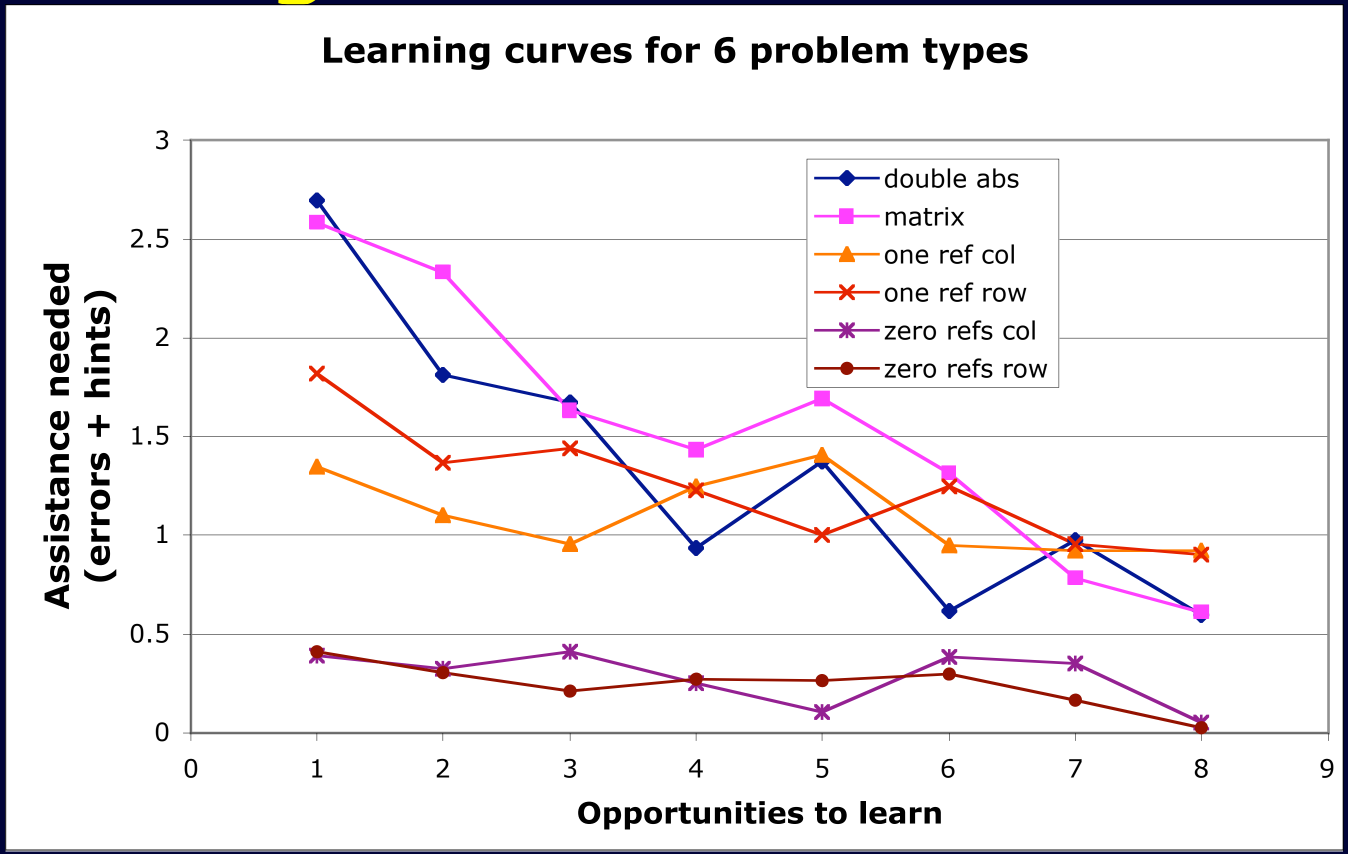

- Knowledge acquired in some problems transfer to others
- What is that knowledge?

Broader rules also plausible

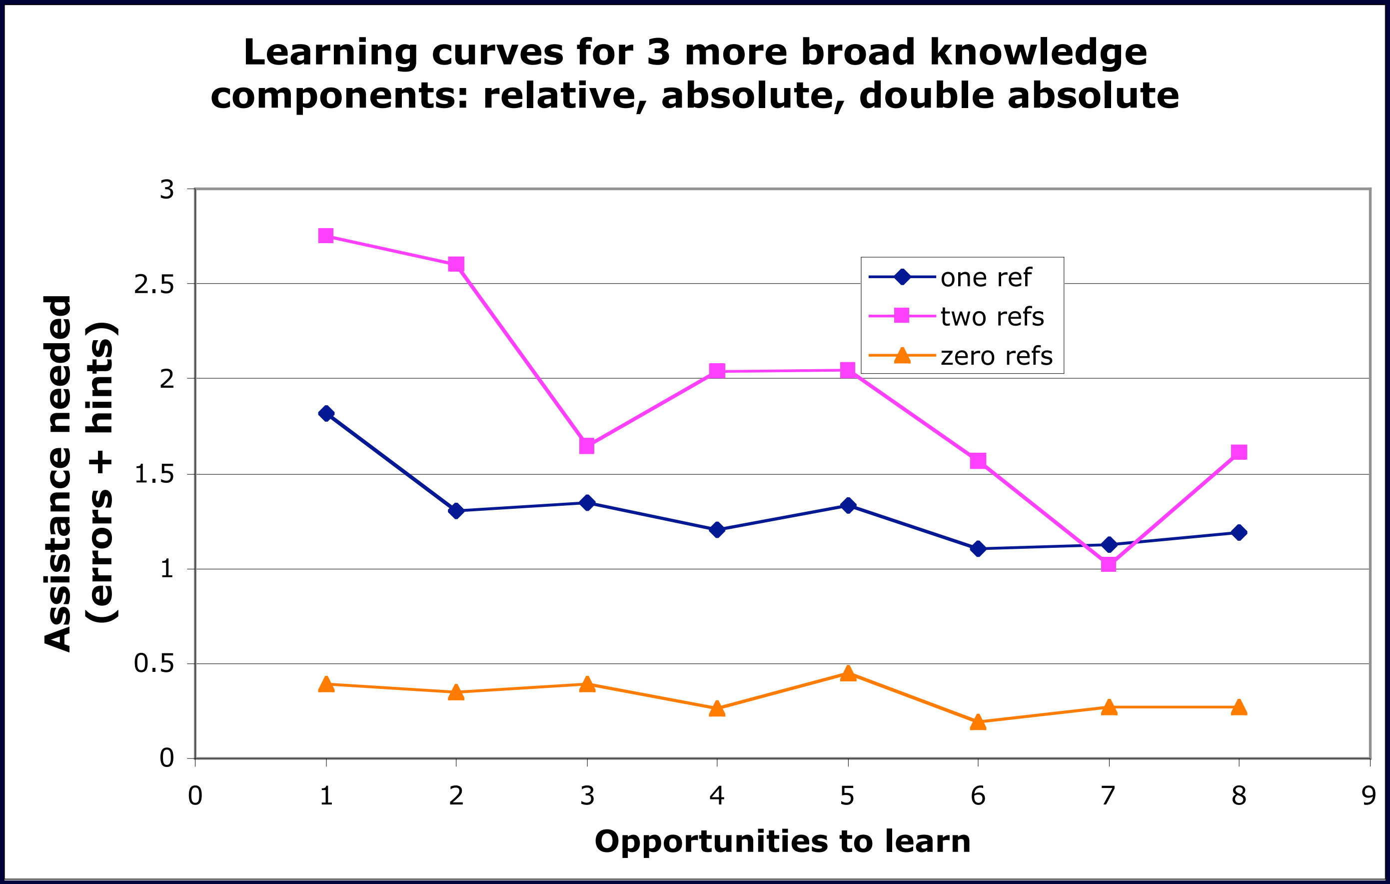

- Which one is right?
- Can we empirically compare knowledge models?

Shallow vs. deep learning: Depth of encoding

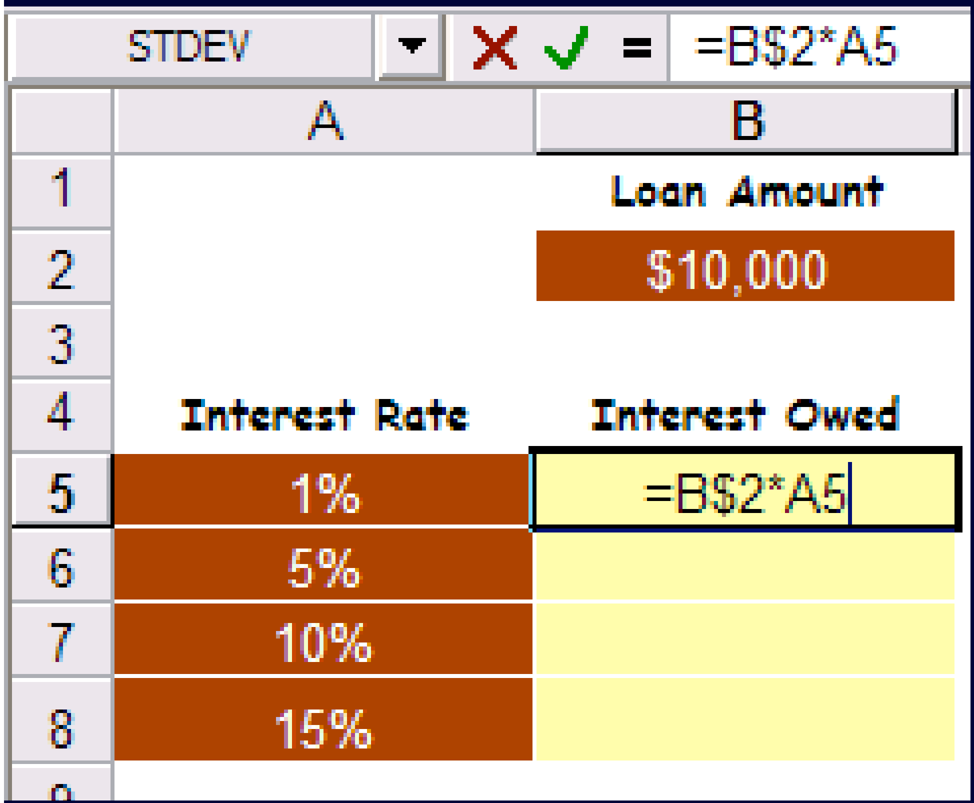

- Shallow feature encoding*:

  Rule S1: If you are copying down a *column*,
  then put a $ in front of the *number*.

- Deep feature encoding:

  Rule D: If a cell index needs to stay fixed but would change when copied,
  then put a $ in front of that index.

*Note: These *describe mental* structures that students *may not* be able to verbalize.

Partially Ordered Knowledge Structures

% of students	Task 1 correct	Task 1 incorrct	
Task 2 correct	30	3	33
Task 2 incorrect	45	22	67
	75	25	

- Contingency tables
- What can we infer?
  - Task 2 is harder?
  - Task 1 is a "prerequisite" for Task 2?
  - Learning on Task 1 transfers to Task 2?
- Example

Task type 1	Task type 2
Story problems	Equations
Result-unknown 800-40*3 = x	Start-unknown 800-40x = 680
Knowing french word for cheese	Equations

Strategies or deep concepts?

- What is the major challenge for learning complex problem solving in math & science?
- Where do you place your bets?

  Students must

  – learn strategies for searching complex problem spaces:

  – acquire deep "conceptual" understanding of domain "operators":

Strategies or deep concepts?

- Pro strategy
  - Schoenfeld has emphasized strategies & heuristics for problem solving, search control
- Pro deep concepts

  Expertise research
  - Chess masters don't do more search, but have more "perceptual chunks"
  - Geometry proof: Search space of theorems is huge
    - So must need search control
    - But, no, mental search space is small: involves rich (perceptual/conceptual) knowledge representations
  - Physics: Problem categorization
    - Novices by shallow features (inclined plane vs. pulley)
    - Experts by features (energy vs. momentum)

Can Educational Data Mining help address strategies vs. representation question?

- Examples from DFA studies
  - Multi-step problems are harder than predicted by difficulty of each step
    - Story -> 800-40x harder than combination of Story-part1 -> 800-y *and* Story-part2 -> 40x
    - Finding area of leftover when circle is cut out of square harder than square+circle+subtract
  - But if context ("distractors") is included in single step problem, no difference
- Demonstrates(?) multi-step problem difficulty
  - is not about complexity of search, but about differentiating shallow vs. deep learning
- Can efficient instruction simply focus single-step problems if they come with rich contexts?

Expand on others?

Scientific questions

- What's the grain size of transfer?
- Strategies or deep concepts?
- Is implicit learning helpful in math/science?
  - Role of noise in learning?
  - Is learning "statistical" in well-structured domains?
- How/when are deep features acquired?
- Does grain size & latency change across domains?
- Sig individual differences in task specialization or learning rate?

Practical questions

- Do better domain models yield better tutors?
  - Better problems, hints, examples, scheduling
- Does more instruction on X make Y easier to learn?
- Can we teach hidden skills?
- Does prerequisite remediation accelerate future learning?
- Does KC complexity affect what kind of instruction is best? Drill vs. self-expl …

What is the domain model discovery problem?

- How to pose domain (assessment) model discovery problem
  - What is the general question that LFA, POKS, ePCA, Q-matrix, Rule Space is trying to answer?
- Challenge for
  - Machine learning community
    - Pose a KDD Cup, what data set(s)?
  - Psychometric community
    - Ken may pose at invited talk at Psychometric Society Conference 08

Student Domain Model Discovery Problem

- Given:

- Goal:

Student Domain Model Discovery Problem

- Given:
  - Log data of students performing tasks over time with feedback & instruction
  - Tasks are "graded" for correctness
    - Students can make multiple errors, request hints, eventually complete every task, times are recorded
  - Tasks are math or science problems in text & often with images, often performed in a structured interface
  - Feature coding of task elements sometimes available
- Goal:
  - Predict *same or new* students' performance on *same or new* tasks at future (or past) times
  - To predict performance means predict error rate, amount of assistance needed, time to complete
    - Save for later: Predict actual student responses

Candidate data sets or features thereof

- To address “big bet” research questions:

- To use in KDD Cup competition:

Too late for KDD 09

Call for Participation: KDD Cup 2009

KDD Cup is the well-known data mining competition of the annual ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD). Also see the KDD Cup Center.

We invite proposals for KDD Cup 2009. Proposals should include a short paragraph on each of the following items:

1. Description of the problem addressed, with general background information on the application domain.
2. Description of the available data, guarantee of availability, guarantee of confidentiality of the "ground truth", and size.
3. Description of the competition tasks, their scientific and technical merit and their practical significance.
4. Description of the evaluation procedures and established baselines (provide a metric of significance in performance differences).
5. List of the available resources (team member availability, computers, support staff, other equipment, sponsors).
6. Plans of result dissemination (e.g., proceedings).
7. Schedule. The competition should last between 6 and 8 weeks and the winners should be notified by mid-May. The winners will be announced in the KDD-2009 conference.
8. Names, affiliations, postal addresses, phone numbers, and short biographies of the organizers.
9. Whether the competition is new or has been held before.
10. Whether you will award prizes of any form to winning teams.

A good competition task is one that is practically useful, scientifically or technically challenging, can be done without extensive application domain knowledge, and can be evaluated objectively. Of particular interests are non-traditional tasks that may need novel techniques and solutions. Please send your proposals to kddcup09@clopinet.com by Oct 30, 2008.

Thanks,

Isabelle Guyon and David Vogel

KDD Cup 2009 Program Co-chairs