Research Methods for the Learning Sciences

Ken Koedinger
Phil Pavlik
TA: Ben Shih

Course Goals

- To learn data collection, design, and analysis methodologies that are particularly useful for scientific research in education

Instructors

- Co-instructor Ken Koedinger
  - PhD in Cognitive Psychology, MS in Computer Science
  - Faculty in Human-Computer Interaction & Psychology
  - Research: Cognitive science and education, educational technology
- Co-instructor Phil Pavlik
  - PhD in Cognitive Psychology
  - Faculty in Human-Computer Interaction
  - Research: How to optimize student learning using cognitive theory and data mining
- Teaching assistant Ben Shih
  - 5th year Ph.D. student in Machine Learning
  - A PIER student
  - Research: Educational data mining to discover features of learner interactions that predict greater learning

And more …

- Different instructors, from the PIER steering committee and elsewhere, will be leading sections of the course corresponding with their expertise

- Marsha Lovett
- Carolyn Rosé
- Sara Kiesler
- Brian Junker
- Richard Scheines

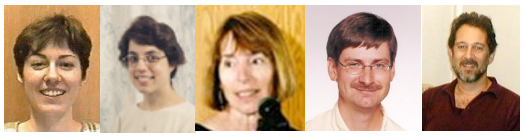

Your turn

- Please:
  - Stand up
  - Tell us your name
  - Your PhD program
  - A few words on your past educational experiences or work experiences
  - A few words on your research interests
  - Why you're taking this course and what you hope to get out of it

Topics We Will Cover

- Basic Research and Experimental Methods
- Cognitive Task Analysis
- Video and Verbal Protocol Analysis
- Surveys, Questionnaires, and Interviews
- Psychometrics, Reliability, Item Response Theory
- Educational Data Mining
- Optional: There's some flexibility in the schedule to add discussion of other methods. Any suggestions?

Pragmatic Issues

- Textbook: "The Research Methods Knowledge Base: 3rd edition" by William M.K. Trochim and James P. Donnelly (2007)
  - http://www.atomicdogpublishing.com/BookDetails.asp?BookEditionID=160
  - It is not in the campus bookstore
  - The used bookstore will have the wrong version
  - Last we looked, it costs more in Amazon than ordering direct from the publisher (see url above)
  - Publisher has both an online and printed version, so you can get started right away (even before your copy of the printed version gets to Pittsburgh)
- Other readings will be assigned in class

Meeting place after today!

- 3501 Newell Simon Hall

Your first readings

- For next class
  - Trochim & Donnelly Chapters 1 and 7
- For Tuesday January 22
  - Trochim & Donnelly Chapters 8, 9, 10
- Highly recommended: Do the on-line quiz for each chapter! Use these to guide your reading and rereading.

Your first assignment

- Being handed out now…
- Phil: Give a brief summary of the assignment
- Due before class Thursday, Jan 21st (next week)
  - Think about before class on this Thursday
  - Bring any questions to class

Assignments

- Most class segments will have at least one assignment
  - Goal is to practice with methods relevant to that segment
- Final project rather than final exam
  - Apply a method in detail to a topic of your choosing

Assignments

- Posts to class discussion board
  - Goals:
    - Help you learn from active application of knowledge, from observing others, and from feedback from others
    - Give instructors feedback on your interests & struggles
    - Extra motivation to do readings before class
  - Default Procedure
    - See description in the syllabus
    - Ben: Please briefly summarize for us
  - Alternative specific guidance may be given for some course sections

Research Methods for the Learning Sciences

- This class is on Research Methods for the Learning Sciences
- Break-down
  - Research
  - Methods
  - for the
  - Learning Sciences

Thanks to prior co-instructor Ryan Baker whose slides I've modified

The Learning Sciences

- What are the Learning Sciences?
- Can anyone here give a definition?

Definition: The Learning Sciences

Does this fit everyone?

- Does this fit the research that everyone in the room is interested in?
- If not, why not?

Previous definitions

- The learning sciences is a field of interdisciplinary study that works to further scientific understanding of learning as well as engage in the design and implementation of learning innovations. [wikipedia]
- Learning Sciences is the scientific study of how people learn and how to design new learning environments, ranging from handheld learning tools, through museum exhibits and innovative classrooms, to learning-enabled cities. [University of Nottingham LSRI]
- The interdisciplinary empirical investigation of learning as it exists in real-world settings and how learning may be facilitated both with and without technology. [ISLS]

Sociological History

- Perhaps the most common use of the term "learning sciences" emerged from a split in the community called "Artificial Intelligence in Education"
  - Occurred for a lot of human, sociological reasons
  - But also had to do with the question: should interactive learning environments focus around
    - human-computer interactions (e.g. Intelligent Tutors, Automated Grading)
    - or human-human interactions (e.g. Computer Supported Collaborative Learning, Teacher Grading of Complex Student Answers)
- Other related terms
  - "Educational Sciences" - Dept of Ed IES, funds PIER
  - "Science of Learning" - NSF Centers

Learning Science*s* combines many disciplines

- Cognitive Psychology
- Education
- Linguistics
- Sociology
- Design
- Statistics
- Computer Science

Methods

- Can anyone here define a method?
- What makes something a method?
- Leave until later: What makes something a method we'd actually want to use?
  - On Thursday, we will discuss what makes a method something you would actually want to use
    - Validity
    - Feasibility
    - Other Issues

Definition: Methods

- Trochim: “specific ways ... you can use to understand the world better” (p. 18)

Research

- I think we can skip a definition of “research”!
- But let’s explore: What is the difference between
  - Applied Research
  - Pure Research

?

Is it possible?

- Can research be ***both*** applied and pure
- At the same time?
- Raise your hands and vote!
  - What’s your reasoning?

Stokes (1997)

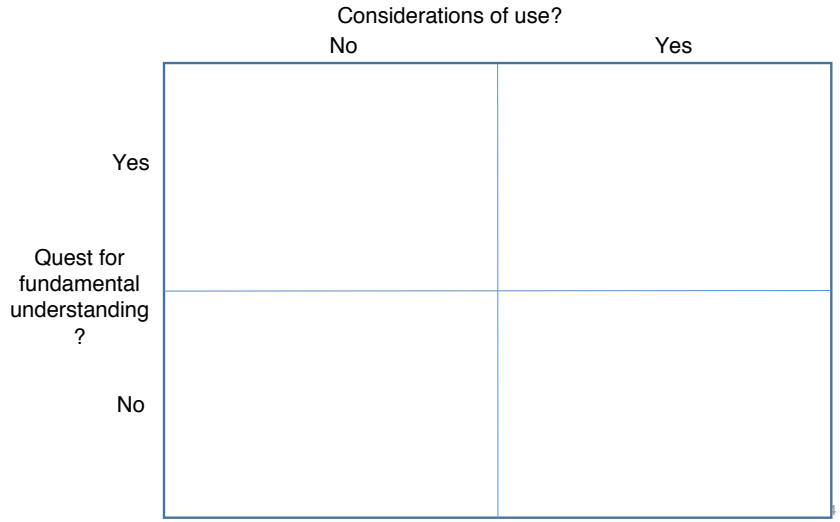

Stokes (1997)

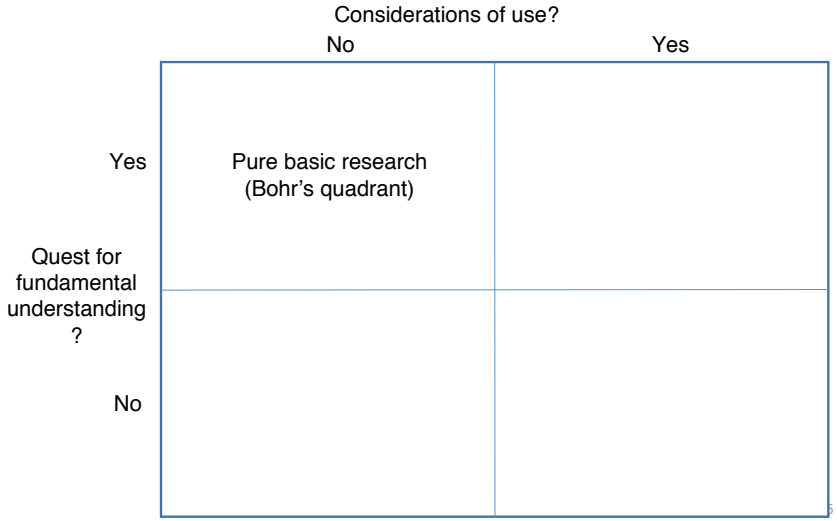

Stokes (1997)

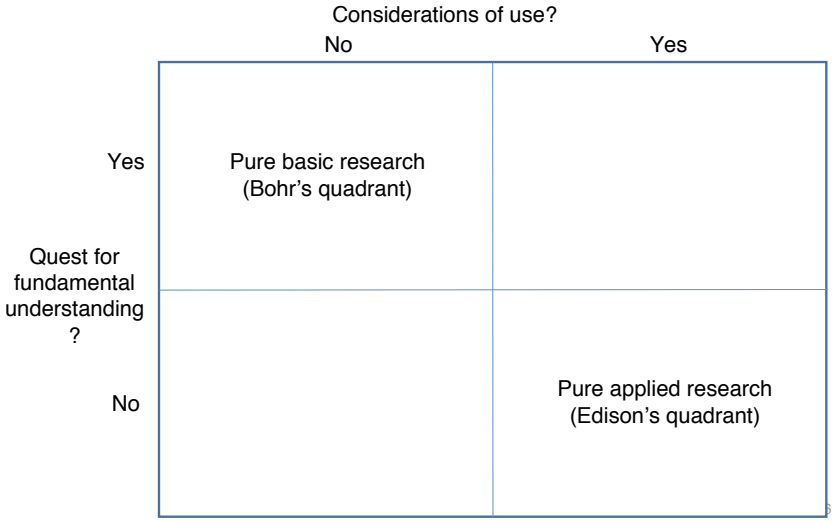

Stokes (1997)

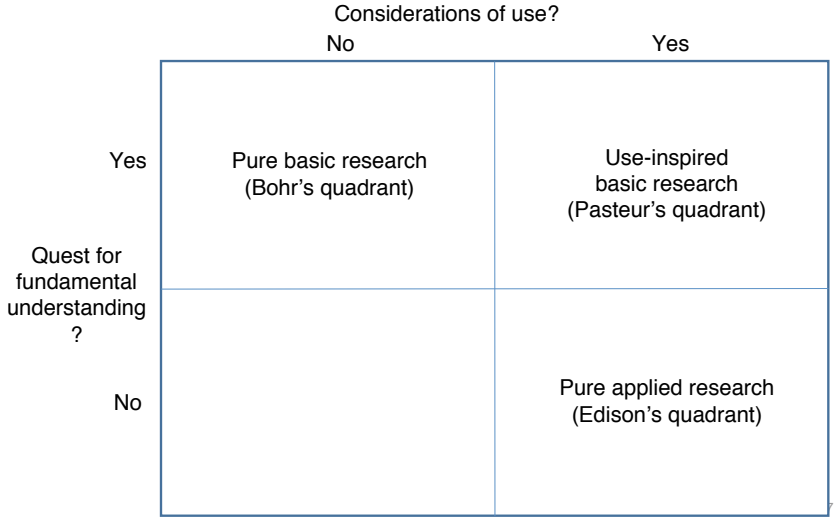

Stokes (1997)

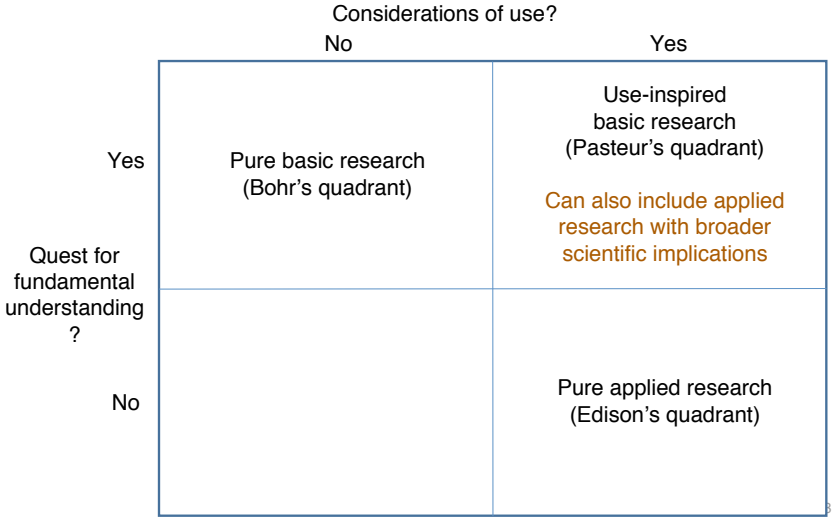

Stokes (1997)

Quest for fundamental understanding? \ Considerations of use?	No	Yes
Yes	Pure basic research (Bohr's quadrant)	Use-inspired basic research (Pasteur's quadrant)
No	Personal curiosity (Peterson's quadrant)	Pure applied research (Edison's quadrant)

Can you give examples of research in each quadrant in the Learning Sciences?

Quest for fundamental understanding? \ Considerations of use?	No	Yes
Yes	Pure basic research (Bohr's quadrant)	Use-inspired basic research (Pasteur's quadrant)
No	Personal curiosity (Peterson's quadrant)	Pure applied research (Edison's quadrant)

Which quadrant is your work in?

Quest for fundamental understanding? \ Considerations of use?	No	Yes
Yes	Pure basic research (Bohr's quadrant)	Use-inspired basic research (Pasteur's quadrant)
No	Personal curiosity (Peterson's quadrant)	Pure applied research (Edison's quadrant)

How do you do good research?

- Whether it is pure or applied
- Recipe
  - Good research directions
  - Good research methods

The rest of the class

- The rest of the class is all about good research methods
- So I'd like to say a few words about good research directions
  - Note: Some of the research methods we will discuss, like Cognitive Task Analysis, can help generate better research questions

Good research directions, goals, or questions

- Write down what you think is a "good research goal"
  - It could be your research goal but it does not have to be
  - Remember to write the ***goal***, not the ***solution***
  - Goal: Increase people's leisure time by reducing the time spent washing dishes
  - Solution: Build a dish-washing robot

What makes a research goal or question good?

- Four factors
  - Importance
  - Attackability
  - Interestingness
  - Falsiability

More critical in applied research

More critical in basic research

Importance

- Will the world be changed for the better if this research goal is accomplished?

Which of these research goals achieve the "Importance" criterion?

- Build an intelligent tutoring system to help students learn to differentiate Bruce Springsteen songs from Billy Joel songs
- Curing déjà vu
- Build an intelligent tutoring system for introductory neurobiology
- Help more students pass introductory neurobiology

Which of these research goals achieve the "Importance" criterion?

- Help students reach mastery with introductory neurobiology material faster
- Curing déjà vu
- Determine which kinds of self-explanation lead to better retention and transfer
- Develop teleportation

The research goal you wrote down

- Is it important?

- Take 30 seconds to reflect.

Attackability

- Sometimes called "Opportunity", like when you're writing a grant proposal
  - Grant section formula:
    1) Problem, 2) Opportunity, 3) Solution
- How feasible is solving this problem now?
- Do ***you*** have a special new approach to the problem that will enable you to solve it?
  - Do you have a unique team? A new technology? A ready user base/participant pool?
- If the problem is (or seems) easy, why hasn't anyone solved it yet (or why is it harder than people seem to think)?

Which of these important research goals are also attackable?

- Teleportation
- Antigravity
- Time Travel
- An intelligent tutor that takes natural language typed input, and responds with speech

Maybe this is attackable for some of you, but not for others

  - This is one reason why interdisciplinary collaboration is so great

Which of these important research goals are also attackable?

Building an intelligent tutor that

- can teach any topic in human understanding, with no prior preparation
- responds as sensitively to differences in student affect as a human tutor
- responds as sensitively to differences in student affect as a human tutor *and* can be used in existing middle school computer labs

Last goal is less attackable than previous one

  - But which goal is more important?

The research goal you wrote down

- Do you have a good attack?
- Is there an attack you don't have but:
  - You could develop
  - Somebody else you know has it (and you could learn or team up with them)
- Take 30 seconds to reflect.

Importance and Attackability

- Two very important ways to assess an applied research goal
- Not the only ways!

How do you know?

- Which goals are important
- Which new methods can provide more powerful attacks on an important problem

Advice

- Attend lots of talks
- Skim lots of papers
- Talk to lots of people

- And most importantly, do it both within and outside your area, whatever that area is

- You will learn important things in your area
- But it is by knowing other areas that you can develop entirely new approaches and attacks for your area

To get an idea

- Look at the bibliography of one of Herb Simon's books
- You'll find citations to research conducted in an unbelievably large number of fields
  - He conducted research in an unbelievably large number of fields, true, but he also knew about recent work in an unbelievably large number of fields

Interestingness

- The field doesn't know the answer yet (or the answer that the field knows is wrong)
  - Robert Abelson calls this "Interestingness"
- Related to "Important"

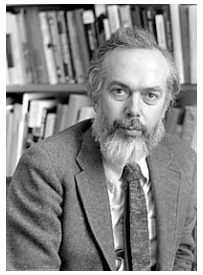

Interestingness

- The most interesting questions are ones where any answer would be interesting
  - Lots of controversy: just as many people think “yes” as “no”
- “Does technique X work for improving learning?” ← may only be interesting if it does work
  - Or if everyone believes it does work (without evidence), but turns out it doesn’t
    - Example: Animations to help learn computer algorithms
- “Which affective state has the largest negative impact on learning” ← any answer is interesting

Which of these research questions is interesting?

- Does the earth revolve around the sun?
- Is the earth flat?
- Does self-explanation promote better learning?
  - Does explaining correct answers or wrong answers lead to better learning?
- Do intelligent tutoring systems do better than traditional curricula?
- Do dialogue tutoring systems do better or worse than problem-solving based tutoring systems?

The research goal/question you wrote down

- Is it interesting?
- Take 30 seconds to reflect.

Falsifiability

- Can you imagine a situation in which your claim might be false (even though you think it is true)?
- It is possible to design a test which gives evidence as to the answer to your question, such that one or more possible answers can be disproven?

Falsifiability

- In the quest for unified theories of domains (like ACT-R in cognition), an even stronger goal is sometimes adopted: to develop theories which both explain multiple existing findings, and make falsifiable predictions which can be tested, in order to refine and improve the theories
  - For more on this, Lakatos (1978)

The research goal/question you wrote down

- Is it possible to test this research question?
- Take 30 seconds to reflect.

How do you know?

- Which questions are interesting?
- Whether a research question is testable?

Interestingness

- Think about:
  - What would be different in the world of research if I answered this question and the answer was known?
  - Would it inspire new work by others?
  - Would it alter work that would occur?
  - Would people talk about it at conferences when you weren't there

Testability/Falsifiability

- Come up with a research design that can give valid evidence on the question
- The existence of at least one valid test is evidence that the question is testable

Testability/Falsifiability

- And remember, the test does not need to be perfect – a first study can be imperfect, paving the way for more thorough later studies
- In the famous words of Herbert Simon

"Anything worth doing is worth doing badly"

Testability/Falsifiability

- In the following weeks, we will discuss a variety of research methods that can be used to make valid inferences about the research questions you are interested in

Summary

- Course overview
- We unpacked: "Research Methods in the Learning Sciences"
- Four criteria for a good research question:
  - Importance
  - Attackability/Opportunity
  - Interestingness
  - Falsifiability

Some resources for reading more about coming up with good research questions

Simon, H.A. (1996) *The Sciences of the Artificial*

Lakatos, I. (1970) Falsification and the Methodology of Scientific Research Programmes. In Lakatos, I., Musgrave, A. (Eds.) *Criticism and the Growth of Knowledge.*

Abelson, R.P. (1995) *Statistics as Principled Argument.*

Hamming, R. (1986) *You and Your Research.* http://www.cs.virginia.edu/~robins/YouAndYourResearch.html

Feynman, R. (1997) *Surely You're Joking, Mr. Feynman: Adventures of a Curious Character.*

Simon, H. (1996) *Models of My Life*

Thursday

- Come prepared to discuss Trochim & Donnelly Chapter 1 and 7!