

Selecting Projects

CSCI 8901:
Research & Evaluation Methods

Prof. Tim Wood
GWU

Selecting a Problem

“A successful person isn’t necessarily better than her less successful peers at solving problems; her pattern-recognition facilities have just learned what problems are worth solving.”

– Ray Kurzweil

How to Recognize Patterns?

With training!

You need practice seeing what others currently view as important problems...

- Requires reading lots of new papers

...and you need to see which papers have had long term impact

- Requires learning the history of your field and its seminal papers

The Best CS Research...

Explores problems that:

Explain the behavior of algorithms, systems, protocols, etc

Your goal should **not** be to build the fastest X or show how to break protocol Z

Instead your goal should be to **understand** the principles that make X faster/better/more accurate or why be able to explain why Z has a fundamental flaw

Good Questions

How...?

Why...?

In what situations...?

Under what conditions...?

The answers won't be yes/no!

Good Problems

“I skate to where the puck is going to be, not where it has been.”

– Wayne Gretzky

Good Problems

“I skate to where the puck is going to be, not where it has been.”

– Wayne Gretzky

Not easy to do... but might as well try!

My NSDI 2007 paper...

Black-box and Gray-box Strategies for Virtual Machine Migration

Timothy Wood, Prashant Shenoy, Arun Venkataramani, and Mazin Yousif[†]
Univ. of Massachusetts Amherst [†]Intel, Portland

Abstract

Virtualization can provide significant benefits in data centers by enabling virtual machine migration to eliminate hotspots. We present Sandpiper, a system that automates the task of monitoring and detecting hotspots, determining a new mapping of physical to virtual resources and initiating the necessary migrations. Sandpiper implements a black-box approach that is

tualization of data center resources provides numerous benefits. It enables application isolation since malicious or greedy applications can not impact other applications co-located on the same physical server. It enables server consolidation and provides better multiplexing of data center resources across applications. Perhaps the biggest advantage of employing virtualization is the ability to flexibly remap physical resources to virtual servers in

Proposed techniques for more efficiently and dynamically managing servers in a data center

- Some good ideas
- Some pretty bad ideas (clear issues for how it modeled the combined load of different resource types)

My NSDI 2007 paper...

Black-box and Gray-box Strategies for Virtual Machine Migration


Timothy Wood, Prashant Shenoy, Arun Venkataramani, and Mazin Yousif[†]
Univ. of Massachusetts Amherst [†]Intel, Portland

Abstract

Virtualization can provide significant benefits in data centers by enabling virtual machine migration to eliminate hotspots. We present Sandpiper, a system that automates the task of monitoring and detecting hotspots, determining a new mapping of physical to virtual resources and initiating the necessary migrations. Sandpiper implements a black-box approach that is

tualization of data center resources provides numerous benefits. It enables application isolation since malicious or greedy applications can not impact other applications co-located on the same physical server. It enables server consolidation and provides better multiplexing of data center resources across applications. Perhaps the biggest advantage of employing virtualization is the ability to flexibly remap physical resources to virtual servers in

Cited by 1048



Nice to get lots of citations...

My NSDI 2007 paper...

Black-box and Gray-box Strategies for Virtual Machine Migration

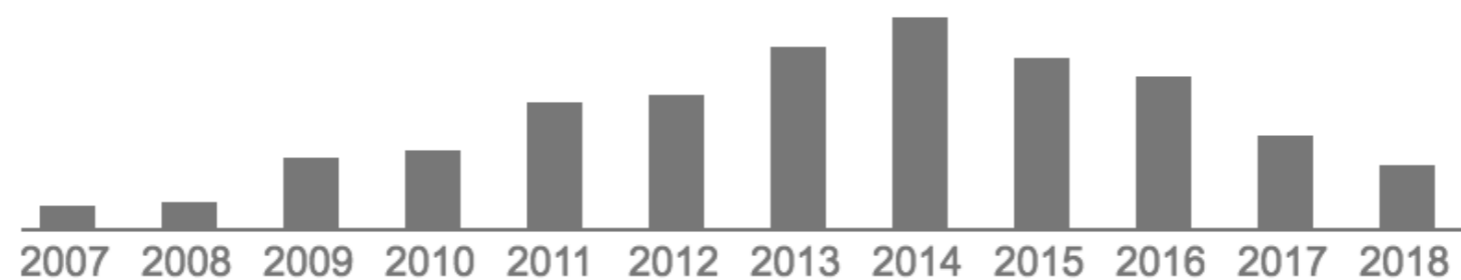
Timothy Wood, Prashant Shenoy, Arun Venkataramani, and Mazin Yousif[†]
Univ. of Massachusetts Amherst [†]Intel, Portland

Abstract

Virtualization can provide significant benefits in data centers by enabling virtual machine migration to eliminate hotspots. We present Sandpiper, a system that automates the task of monitoring and detecting hotspots, determining a new mapping of physical to virtual resources and initiating the necessary migrations. Sandpiper implements a black-box approach that is

tualization of data center resources provides numerous benefits. It enables application isolation since malicious or greedy applications can not impact other applications co-located on the same physical server. It enables server consolidation and provides better multiplexing of data center resources across applications. Perhaps the biggest advantage of employing virtualization is the ability to flexibly remap physical resources to virtual servers in

Cited by 1048



But why are some people still doing the same thing 10+ years later?

New and Old

Of course you should read and cite “old” work!

- In CS, we often forget about things older than ~7 years

But you should do more than just iterate on it!

Think about what has changed...

- Underlying technologies
- Workloads of users / data
- Advances in algorithms or modeling techniques

Use these advances to make a bigger change

Change

In fact, most current citations focus on a different platform/environment/metric

Scholar About 15 results (0.03 sec)

Black-box and Gray-box Strategies for Virtual Machine Migration.

Search within citing articles

Since 2021

[HTML] Autoscaling tiered cloud storage in anna

[HTML] springer

[C Wu](#), [V Sreekanti](#), [JM Hellerstein](#) - *The VLDB Journal*, 2021 - Springer

In this paper, we describe how we extended a distributed key-value store called Anna into an autoscaling, multi-tier service for the cloud. In its extended form, Anna is designed to overcome the narrow cost-performance limitations typical of current cloud storage systems ...

☆ [Cited by 4](#) [Related articles](#)

A resource management algorithm for virtual machine migration in vehicular cloud computing

[SK Pande](#), [SK Panda](#), [S Das](#), [KS Sahoo](#)... - ... , *Materials & Continua*, 2021 - researchgate.net

In recent years, vehicular cloud computing (VCC) has gained vast attention for providing a variety of services by creating virtual machines (VMs). These VMs use the resources that are present in modern smart vehicles. Many studies reported that some of these VMs hosted on ...

☆ [Cited by 4](#) [Related articles](#) [»»](#)

Resource allocation using dynamic pricing auction mechanism for supporting emergency demands in Cloud computing

[R Ananthakumar](#), [K Kartheeban](#) - *Journal of Parallel and Distributed ...*, 2021 - Elsevier

Cloud resources provide various categories of (VM) virtual machine requests which is assigned with clients for an exact timespan. Currently, the method of VM scheduling within the Cloud environment is decided by a fixed-price scheduling algorithm, during which the ...

☆ [All 2 versions](#)

ThingsMigrate: Platform-independent migration of stateful JavaScript Internet of Things applications

[PDF] wiley.com Full View

[K Jung](#), [J Gascon-Samson](#), [S Goyal](#)... - *Software: Practice ...*, 2021 - Wiley Online Library

Abstract The Internet of Things (IoT) has gained wide popularity both in academic and industrial contexts. Unlike traditional embedded devices with specialized firmwares, modern IoT devices accommodate general-purpose operating systems, allowing developers to run ...

☆ [Cited by 1](#) [Related articles](#) [All 3 versions](#)

Intelligent container reallocation at Microsoft 365

[PDF] acm.org

[B Qiao](#), [F Yang](#), [C Luo](#), [Y Wang](#), [J Li](#), [Q Lin](#)... - *Proceedings of the 29th ...*, 2021 - dl.acm.org

The use of containers in microservices has gained popularity as it facilitates agile development, resource governance, and software maintenance. Container reallocation aims to achieve workload balance via reallocating containers over physical machines. It affects ...

☆ [All 2 versions](#)

[HTML] AutoShare: Virtual community solar and storage for energy sharing

[HTML] springer

[S Lee](#), [P Shenoy](#), [K Ramamritham](#), [D Irwin](#) - *Energy Informatics*, 2021 - Springer

Residential solar installations are becoming increasingly popular among homeowners. However, renters and homeowners living in shared buildings cannot go solar as they do not own the shared spaces. Community-owned solar arrays and energy storage have emerged ...

☆ [All 7 versions](#)

Reading 1

A Guide to Increased Creativity in Research — Inspiration or Perspiration?

by C. Loehle, in Bioscience February 1990

A guide to increased creativity in research— inspiration or perspiration?

There are four requirements for a successful career in science: knowledge, technical skill, communication, and originality or creativity. Many succeed with largely the first three. Those who are meticulous and skilled can make a considerable name by doing the critical experiments that test someone else's ideas or by measuring something more accurately than anyone else. But in such areas of science as biology, anthropology, medicine, and theoretical physics, more creativity is needed because phenomena are complex and multivariate.





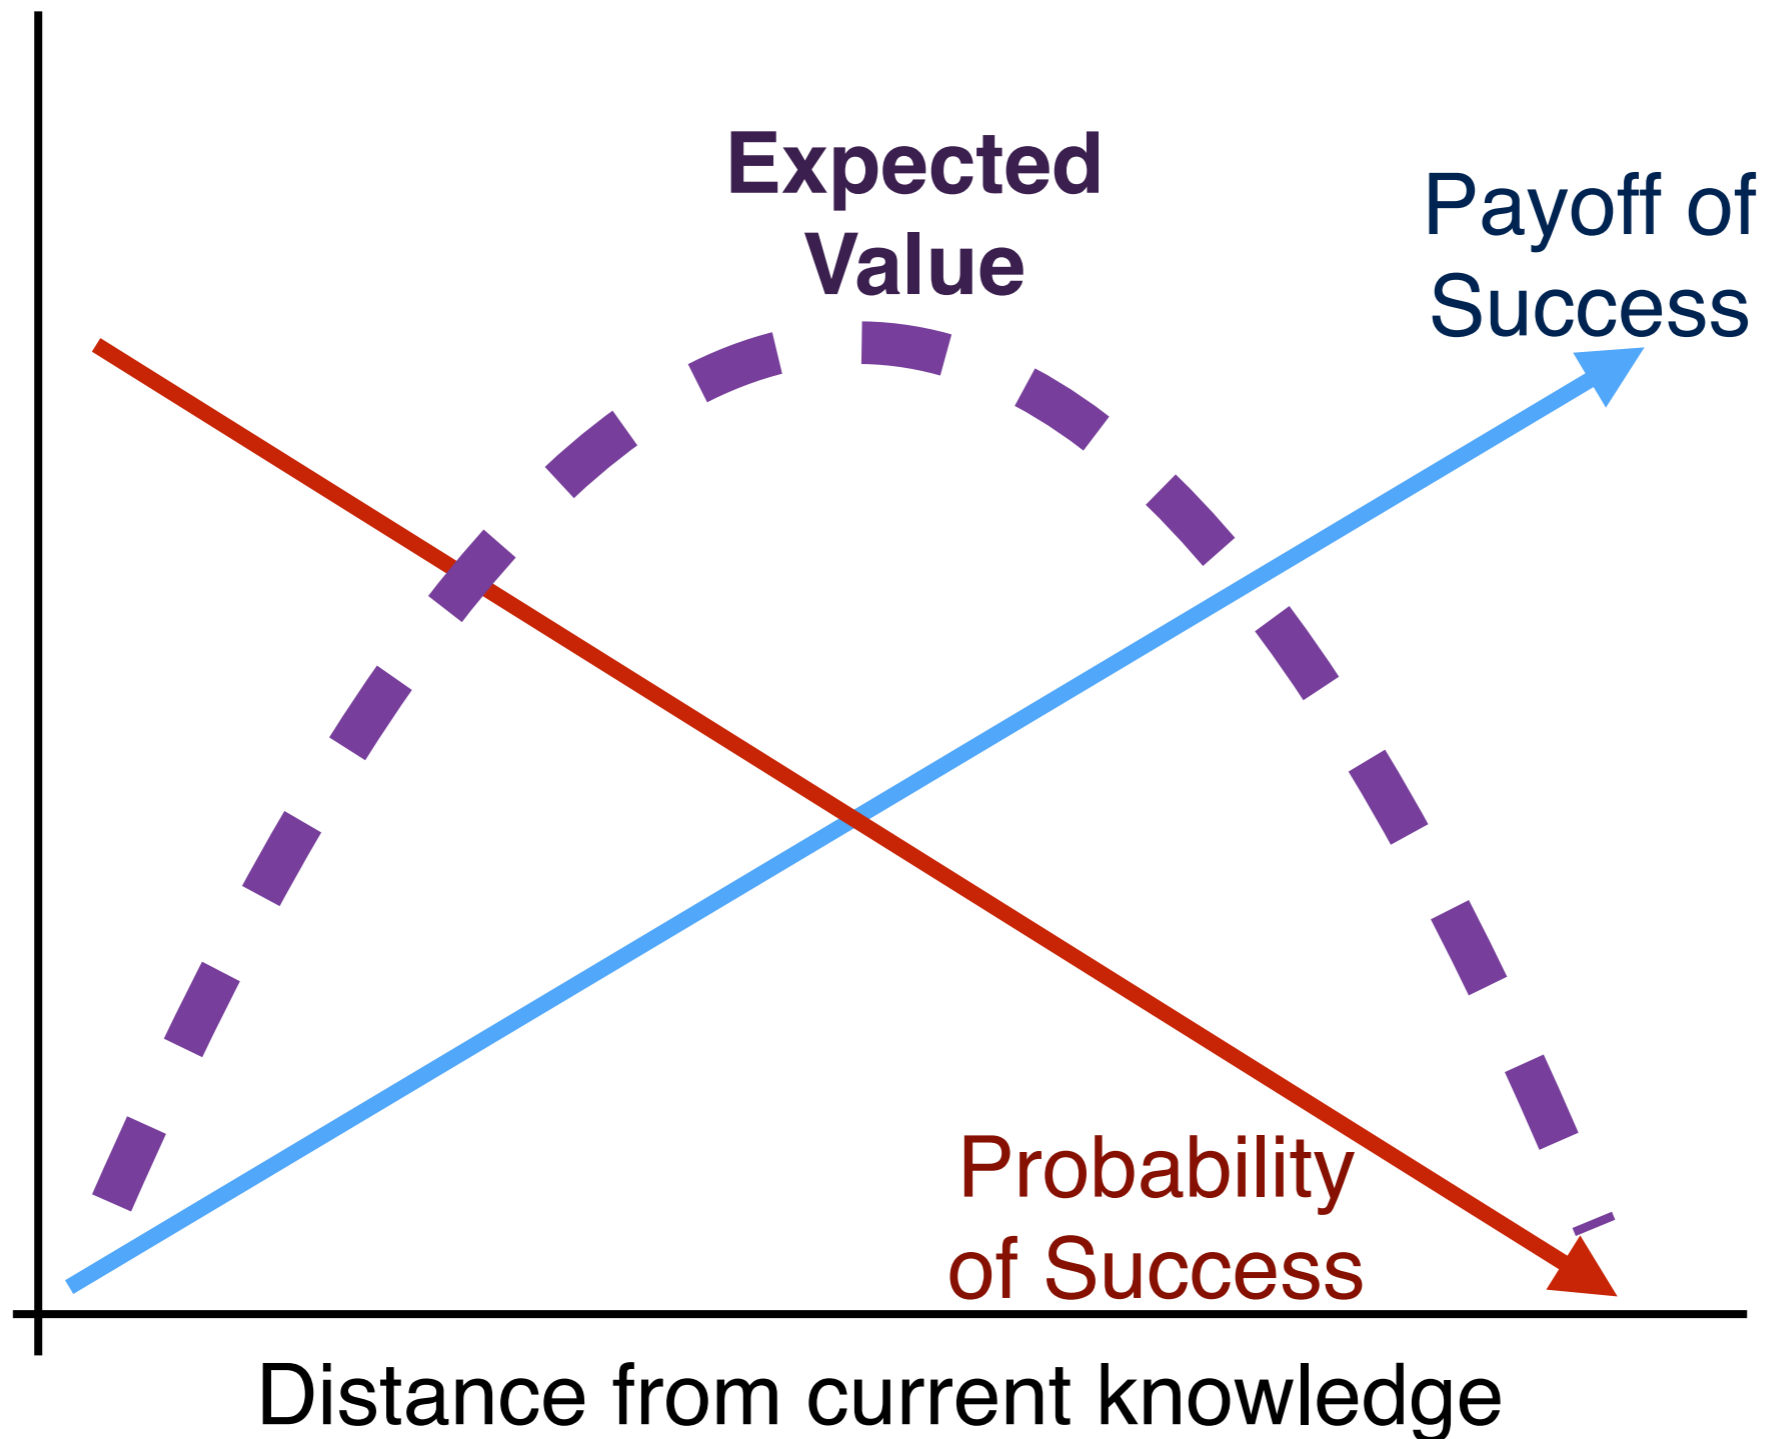
Figure 1. Relationship between degree of difficulty and payoff from solving a problem. Solving problems that are too easy

and work intensely. After you have finished writing your paper, you can go back and remove the comments about what an imbecile the other person is. The effort to refute someone can even lead to evidence supporting them or to a different topic altogether. Intensive rivalries, as in the race to discover DNA (Watson 1968), can also provide this essential intensity. Thus whereas the finished product may appear dispassionate, truly creative work is often driven by strong passions.

Medawar Zone



Medawar Zone



Other Useful Tips?

Don't be an expert

Develop courage

Don't read the literature?

Pick a fight / where is the smoke?

Writing as brain dumps

Read to find out!

Finding the Frontier

New results that don't have a clear explanation yet

Opportunities provided by...

- Changing hardware
- New algorithms
- New modeling frameworks or fundamental results

The convergence of two fields

New applications of your field to another

**"Good research is done with a shovel, not with tweezers...
You should find an area where you can get a lot out of it fast."**

– Roger Needham

Pick the right problem... for you!

What are your strengths?

- Algorithmic, implementation, experimental, analytic, etc

Why will you be able to solve this problem more effectively than most other people?

Pick the right problem... for you!

What are your strengths?

- Algorithmic, implementation, experimental, analytic, etc

Why will you be able to solve this problem more effectively than most other people?

Bad answers:

- “I’m smarter than most people!” (Sorry, but this is unlikely)
- “I will work harder and longer!” (Unlikely and can hurt the other parts of your life)

“Always have a secret weapon — the biggest computer, a problem imported from another field that others haven't heard of yet, a fact you stumbled on by being curious about everything, a friend who is smarter than you are. Anything, in short, that will give you an unfair advantage in getting there first.

– Herbert Simon

A necessary ingredient...

Pick a project that matches your **personal interests**

Research takes a long time

- You need to stay motivated and interested in the topic!

Maintaining your enthusiasm is important

- For yourself
- For convincing others to join / help you

But keep in mind...

- You might be being funded by a research grant with specific goals / deliverables. You probably do not have complete freedom (yet!)

Typical Research Timeline

Where does idea come from?

What do you do first?

How long to design and test your solution?

How long to write a paper?

How long until paper is reviewed?

Typical Research Timeline

Where does idea come from?

- From advisor, from literature

What do you do first?

- read a lot about the problem, learn about tools/implementation aspects, write list of contributions, talk to your advisor, do motivational experiments

How long to design and test your solution?

- 0.5-2 years


How long to write a paper?

- 2-6 weeks often last minute. Better to start early!

How long until paper is reviewed?

- 4 months, might take multiple submissions to be accepted

Research Problem Framework



From Paul R. Cohen, Empirical Methods for Artificial Intelligence. MIT Press 1995

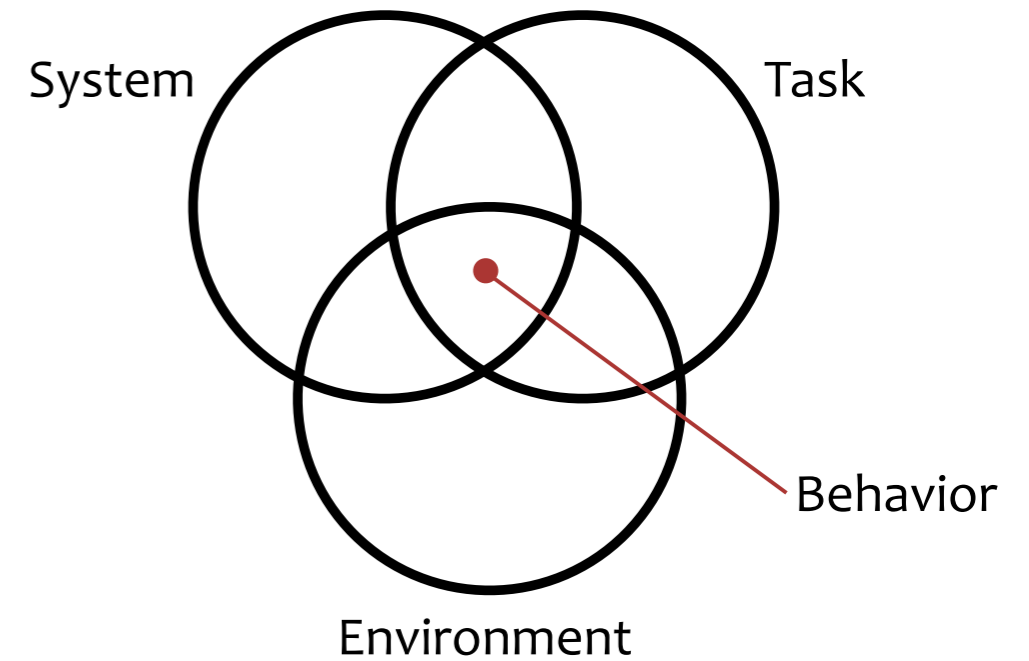
System

The thing you “built”

Could be anything:

- a new operating system
- a new ML inference algorithm
- a new NLP framework
- a wireless network management algorithm

This part of the research is fully **under your control**




Task

The “work” to be done

This is an input to your system:

- Incoming requests to a web server
- The type of query to a DB
- The input data be analyzed by your algorithm


This is **under the control of the user** of your system



Environment

The surroundings that affect both the task and the system

- Available CPU and memory
- The network characteristics



This is **can't be controlled by anyone**

- (although you may be able to vary it during experiments or problem selection)

Behavior

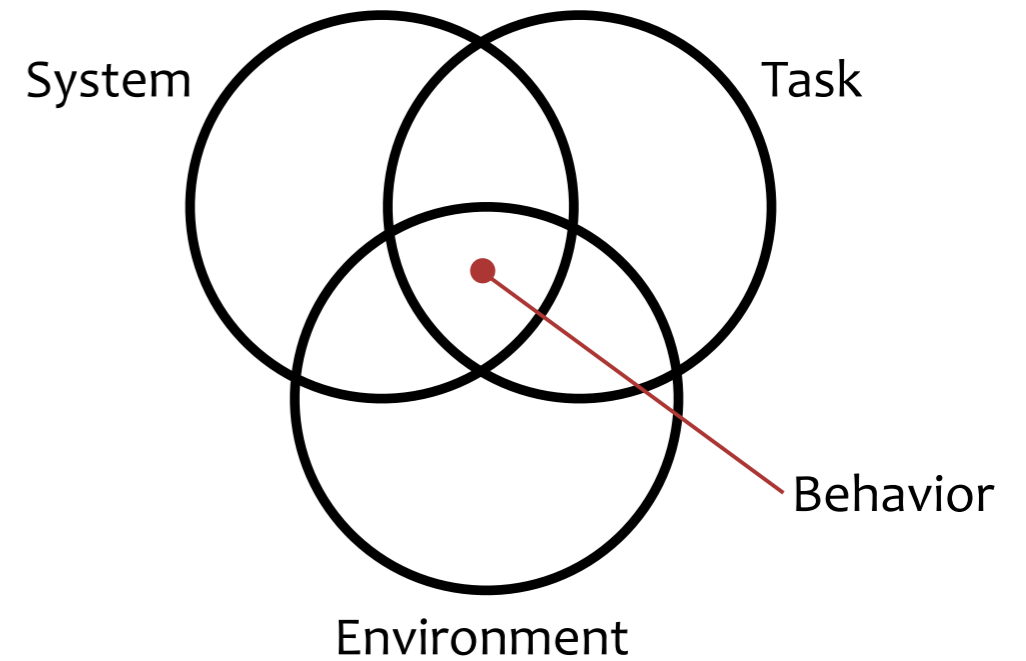
How does the system behave for a particular task and environment?

- Your class project should focus on this type of behavior!

You should evaluate how the

- system, task, and environment

affect the behavior



Your Research

15 minutes?

With a partner:


- Read each other's project descriptions
- Discuss what part of your research is the...

System: your solution

Task: inputs controlled by user of system

Environment: parameters potentially beyond your control

Behavior: metrics you want to study or understand when the above 3 change



Acknowledgements

Much of the slide content, and almost all of the amazing quotations, are derived from the *Research Methods for Empirical Computer Science* course taught by **David Jensen**

- <http://dx.doi.org/11084/10002>
- <https://people.cs.umass.edu/~jensen/courses/index.html>
- <https://people.cs.umass.edu/~jensen>
- Many thanks for allowing me to make use of his materials!