

Acknowledgment: The slides were prepared via valuable discussion, feedback from various colleagues and students. Our research that these talk slides were made based on has been supported in part by NSF and ARO.

Mapping out a Research Agenda

Tao Xie

Department of Computer Science
University of Illinois at Urbana-Champaign

<http://www.cs.illinois.edu/homes/taoxie/>

taoxie@illinois.edu

June, 2010 (first version)

Jan, 2014 (last update)

<http://web.engr.illinois.edu/~taoxie/publications/researchagenda.pdf>

<http://web.engr.illinois.edu/~taoxie/advice/>

Essential Skills for (PhD) Researchers

- is able to **independently**
 - Assess
 - Others' Work (e.g., conference PC members, journal reviewers)
 - Own Work
 - Envision (e.g., per n years, research agenda)
 - Design (e.g., per paper/project)
 - Problem
 - Solution
 - Execute (e.g., time/risk/team management)
 - Implement
 - Evaluate
 - Communicate
 - Written
 - Oral



**high-quality/impact
research**

Critical, Visionary, Creative, Strategic/Engineering, Logical... Skills

Brief Desirable Characteristics of Your Paper/Project

- Two main elements
 - Interesting idea(s) accompanying interesting claim(s)
 - claim(s) well validated with evidence
- Then how to define “interesting”?
 - Really depend on the readers’ taste but there may be general taste for a community
 - Ex: being the first in X, being non-trivial, contradicting conventional wisdoms, ...
 - Can be along problem or solution space; in SE, being the first to point out a refreshing and practical problem would be much valued
 - Uniqueness, elegance, significance?

(Broader) Impact

- There are different types of impacts: research, industrial, societal/social, ...
- Research impact, e.g., impact on research colleagues in various forms -- citations, inspiration, opening a new field/direction, ...
- General, fundamental, conceptual ideas (beyond a tool, implementation, infrastructure, study..) recent examples on QA
 - Godefroid/Sen et al. DART/CUTE/Concolic testing, PLDI 05/FSE 05
 - Engler et al. Coverity/Bugs as deviants, SOSP 01
 - Ernst et al. Daikon/Dynamic invariant detection, ICSE 99
 - Zeller. Delta debugging, FSE 99
- Overreaching contributions conveyed as insights

<http://www.sigsoft.org/awards/ImpactAward.htm>

<http://www.sigsoft.org/awards/mostInfPapAwd.htm>

http://academic.research.microsoft.com/CSDirectory/Paper_category_4.htm

Factors Affecting Choosing a Problem/Project

- What factors affect you (not) to choose a problem/project?
 - Besides your supervisor/mentor asks you (not) to choose it

Factors Affecting Choosing a Problem/Project

- Impact/significant: Is the problem/solution important? Are there any significant challenges?
 - Industrial impact, research impact, ...
 - DON'T work on a problem imagined by you but not being a real problem
 - E.g., determined based on your own experience, observation of practice, feedback from others (e.g., colleagues, industrial collaborators)
- Novelty: is the problem novel? is the solution novel?
 - If a well explored or crowded space, watch out (how much space/depth? how many people in that space?)

Factors Affecting Choosing a Problem/Project II

- Risk: how likely the research could fail?
 - reduced with significant feasibility studies and risk management in the research development process
 - E.g., manual “mining” of bugs
- Cost: how high effort investment would be needed?
 - Sometimes being able to be reduced with using tools and infrastructures available to us
 - Need to consider evaluation cost (solutions to some problem may be difficult to evaluate)
 - But don't shut down a direction simply due to cost

Factors Affecting Choosing a Problem/Project III

- Better than existing approaches (in important ways) besides new: engineering vs. science
- Competitive advantage
 - “secret weapon”
 - Why you/your group is the best one to pursue it?
 - Ex. a specific tool/infrastructure, access to specific data, collaborators, an insight,...
 - Need to know your own strengths/weaknesses
- Underlying assumptions and principles - how do you (systematically) choose what to pursue?
 - core values that drive your research agenda in some broad way

Dijkstra's Three Golden Rules for Successful Scientific Research

- Rule #1. Raise your quality standards as high as you can live with, avoid wasting your time on routine problems, and always try to work as closely as possible at the boundary of your abilities. Do this, because it is the only way of discovering how that boundary should be moved forward.
 - “internal”: about yourself

Dijkstra's Three Golden Rules for Successful Scientific Research

- Rule #2. We all like our work to be socially relevant and scientifically sound. If we can find a topic satisfying both desires, we are lucky; if the two targets are in conflict with each other, let the requirement of scientific soundness prevail.
 - “external”: relation between “the scientific world” and “the real world”
 - traditionally violated by social sciences

Dijkstra's Three Golden Rules for Successful Scientific Research

- Rule #3. Never tackle a problem of which you can be pretty sure that (now or in the near future) it will be tackled by others who are, in relation to that problem, at least as competent and well-equipped as you.
 - Corollary: never compete with one's colleagues

Heilmeier's Catechism

- What are you trying to do
 - Articulate your **objectives** using absolutely no **jargon**
- How's it done now, what are **limits** of **current practice**
- What's **new** in your approach, why you'll be successful
- Who cares (your **audience**)
- If you're successful, what **difference** will it make
- What are the **risks** and **payoffs**
- How much will it **cost**; How long will it take
- What are midterm & final “**exams**” to check for **success**

Adapted from slides by Zhendong Su:

http://en.wikipedia.org/wiki/George_H._Heilmeier

<http://sei.pku.edu.cn/conference/ishcs/research-su.pdf>

Example Principles – Problem Space

- Question core assumptions or conventional wisdoms about SE
- Play around industrial tools to address their limitation
- Collaborate with industrial collaborators to decide on problems of relevance to practice
- Investigate SE mining requirement and adapt or develop mining algorithms to address them (e.g., Suresh Thummalapenta [ICSE 09, ASE 09])

Example Principles – Solution Space

- Integration of static and dynamic analysis
- Using dynamic analysis to realize tasks originally realized by static analysis
 - Or the other way around
- Using compilers to realize tasks originally realized by architectures
 - Or the other way around
- ...

Factors Affecting Choosing a Problem/Project IV

- Intellectual curiosity
- Other benefits (including option value)
 - Emerging trends or space
 - Funding opportunities, e.g., security
 - Infrastructure used by later research
 - ...
- What you are interested in, enjoy, passionate, and believe in
- AND a personal taste
- Tradeoff among different factors

Dijkstra's Three Golden Rules for Successful Scientific Research

1. “Internal”: Raise your quality standards as high as you can live with, avoid wasting your time on routine problems, and always try to work as closely as possible at the boundary of your abilities. Do this, because it is the only way of discovering how that boundary should be moved forward.
2. “External”: We all like our work to be socially relevant and scientifically sound. If we can find a topic satisfying both desires, we are lucky; if the two targets are in conflict with each other, let the requirement of scientific soundness prevail.

Dijkstra's Three Golden Rules for Successful Scientific Research cont.

3. “Internal/ External”: Never tackle a problem of which you can be pretty sure that (now or in the near future) it will be tackled by others who are, in relation to that problem, at least as competent and well-equipped as you.

Jim Gray's Five Key Properties for a Long-Range Research Goal

- Understandable: simple to state.
- Challenging: not obvious how to do it.
- Useful: clear benefit.
- Testable: progress and solution is testable.
- Incremental: can be broken in to smaller steps
 - So that you can see intermediate progress

http://research.microsoft.com/pubs/68743/gray_turing_frcr.pdf

<http://arxiv.org/ftp/cs/papers/9911/9911005.pdf>

Tony Hoare's Criteria for a Grand Challenge

- Fundamental
- Astonishing
- Testable
- Inspiring
- Understandable
- Useful
- Historical

The Verifying Compiler: A Grand Challenge for Computing Research by Hoare, CACM 2003

<http://www.cs.yale.edu/homes/dachuan/Grand/HoareCC.pdf>

<http://vimeo.com/39256698>

Tony Hoare's Criteria for a Grand Challenge cont.

- International
- Revolutionary
- Research-directed
- Challenging
- Feasible
- Incremental
- Co-operative

The Verifying Compiler: A Grand Challenge for Computing Research by Hoare, CACM 2003

<http://www.cs.yale.edu/homes/dachuan/Grand/HoareCC.pdf>

<http://vimeo.com/39256698>

Tony Hoare's Criteria for a Grand Challenge cont.

- Competitive
- Effective
- Risk-managed

The Verifying Compiler: A Grand Challenge for Computing Research by Hoare, CACM 2003

<http://www.cs.yale.edu/homes/dachuan/Grand/HoareCC.pdf>

<http://vimeo.com/39256698>

Heilmeier's Catechism

Anyone proposing a research project or product development effort should be able to answer

- What are you trying to do? Articulate your objectives using absolutely no jargon.
- How is it done today, and what are the limits of current practice?
- What's new in your approach and why do you think it will be successful?
- Who cares?
- If you're successful, what difference will it make?
- What are the risks and the payoffs?
- How much will it cost?
- How long will it take?
- What are the midterm and final "exams" to check for success?

Ways of Coming Up a Problem/Project

- Know and investigate literatures and the area
- Investigate assumptions, limitations, generality, practicality, validation of existing work
- Address issues in your own development experiences or from other developers'
- Explore what is “hot” (pros and cons)
- See where your “hammers” could hit or be extended
- Ask “why not” on your own work or others' work
- Understand existing patterns of thinking
 - <http://people.engr.ncsu.edu/txie/adviceonresearch.html>
- Think more and hard, and interact with others
 - Brainstorming sessions, reading groups
- ...

Some points were extracted from Barbara Ryder's slides of the same talk title

Example Techniques on Producing Research Ideas

- Research Matrix (Charles Ling and Qiang Yang)
- Shallow/Deep Paper Categorization (Tao Xie)
- Paper Recommendation (Tao Xie)
 - Students recommend/describe a paper (not read by the advisor before) to the advisor and start brainstorming from there
- Research Generalization (Tao Xie)
 - “balloon”/ “donut” technique

Technique: Research Matrix

© Charles Ling and Qiang Yang

The Matrix Method

		Methods			
		1	2	3	4
Problem and Applications	1	[3]		[17]	[4]
	2	[55]		[23]	[43]
	3	[1]			[44]
	4				

- X-axis: methods
- Y-axis: Problems

See Book Chapter 4.3: Crafting Your Research Future: A Guide to Successful Master's and Ph.D. Degrees in Science & Engineering by Charles Ling and Qiang Yang

<http://www.amazon.com/Crafting-Your-Research-Future-Engineering/dp/1608458105>

Technique:

Shallow Paper Categorization

© Tao Xie

- See Tao Xie's research group's shallow paper category:
 - <https://sites.google.com/site/aserggrp/bibli>
- Categorize papers on the research topic being focused
- Both the resulting category and the process of collecting and categorizing papers are valuable

Technique:

Deep Paper Categorization

© Tao Xie

- Adopted by Tao Xie's research group
- Categorize papers on the research topic being focused (in a deep way)
- Draw a table (rows: papers; columns: characterization dimensions of papers)
- Compare and find gaps/correlations across papers

Example Table on Symbolic Analysis:

Paper	Category/Categories	Inputs	What to turn Symbolic	Manipulation of Symbolic Expressions/Path Conditions	Aggregation Across Runs	Outputs
DySy: Dynamic Symbolic Execution for Invariant Inference . Christoph Csallner, Nikolai Tillmann, and Yannis Smaragdakis. ICSE 2008	Invariant Inference	Existing Tests	Inputs from the existing test		Yes, since invariants stay the same after all transformations.	Inferred invariants from code. Abstract conditions over program variables that the concrete tests satisfy.
Precise Identification of Problems for Structural Test Generation . Xusheng Xiao, Tao Xie, Nikolai Tillmann, Jonathan Halleux. ICSE 2011.	Problem Identification	program under test	input parameters	User guidance???	Yes	
Decision Interface Identification for		Web application	web application		No	Accepted interfaces

Technique: “Balloon”/“Donut”

© Tao Xie

- Adopted by Tao Xie’s research group
- *Balloon*: the process is like blowing air into a balloon
- *Donut*: the final outcome is like a donut shape (with the actual realized problem/tool as the inner circle and the applicable generalized problem/solution boundary addressed by the approach as the outer circle)
- Process: do the following for the problem/solution space separately
 - Step 1. Describe what the exact concrete problem/solution that your tool addresses/implements (assuming it is X)
 - Step 2. Ask questions like “Why X? But not an expanded scope of X?”
 - Step 3. Expand/generalize the description by answering the questions (sometimes you need to shrink if overgeneralize)
 - Goto Step 1

Example Application of “Balloon”/“Donut”

© Tao Xie

- Final Product: Xusheng Xiao, Tao Xie, Nikolai Tillmann, and Jonathan de Halleux. Precise Identification of Problems for Structural Test Generation. ICSE 2011
<http://people.engr.ncsu.edu/txie/publications/icse11-covana.pdf>
- Problem Space
 - Step 1. (Inner circle) Address too many false-warning issues reported by Pex
 - Step 2. Why Pex? But not dynamic symbolic execution (DSE)?
 - Step 3. Hmm... the ideas would work for the same problem faced by DSE too
 - Step 1. Address too many false-warning issues reported by DSE
 - Step 2. Why DSE? But not symbolic execution?
 - Step 3. Hmm.. the ideas would work for the same problem faced by symbolic execution too
 -
 - Outer circle: Address too many false-warning issues reported by test-generation tools that focus on structural coverage and analyze code for test generation (some techniques work for random test generation too)

Example Application of “Balloon”/“Donut”

© Tao Xie

- Final Product: Xusheng Xiao, Tao Xie, Nikolai Tillmann, and Jonathan de Halleux. Precise Identification of Problems for Structural Test Generation. ICSE 2011
<http://people.engr.ncsu.edu/txie/publications/icse11-covana.pdf>
- Solution Space
 - Step 1. (Inner circle) Realize issue pruning based on symbolic analysis implemented with Pex
 - Step 2. Why Pex? But not dynamic symbolic execution (DSE)?
 - Step 3. Hmmm... the ideas can be realized with general DSE
 - Step 1. Realize issue pruning based on symbolic analysis implemented with DSE
 - Step 2. Why DSE? But not symbolic execution?
 - Step 3. Hmmm ... the ideas can be realized with general symbolic execution
 -
 - Outer circle: Realize issue pruning based on dynamic data dependence (which can be realized with many different techniques!), potentially the approach can use static data dependence but with tradeoffs between dynamic and static

Industrial Collaboration/Research

- Benefits
 - Problems
 - Infrastructures
 - Evaluation testbeds
- Caveats
 - Practical utilities \neq research (at least not always)
 - Short term vs. long term
 - Product groups \leftrightarrow researchers

Big Picture and Open Mind

- Don't narrow-mindedly and “stubbornly” stick to your initial solution and defend it (sometimes with weak justification)
- Step back and ask questions (to challenge)
 - Ex. Why static analysis in contrast to dynamic analysis? [Thummalapenta et al. FSE 09]
 - Ex. Why frequent partial order miner in contrast to frequent automaton miner [Acharya et al. FSE 07]
- Your initial solution faced challenges and difficulties → Good news?!
 - Opportunities for adding new techniques

Example: Alternative Pattern Mining

- (Imbalanced) alternative patterns and new mining algorithm for them were initially proposed
- Question 1: Are these types of alternative patterns the only types of patterns in dealing with alternative ways of using APIs?
- Question 2: Why couldn't existing partial order miners or finite automaton learners to mine alternative patterns (they do infer alternative ways of using APIs)?
- These questions led to finer classification of alternative patterns: balanced and imbalanced

Broader View on Solution Space

- AVOID “a tendency to be too focused on implementing a particular approach to a problem and not interested enough in exploring a broader range of approaches and understanding why some of them work well and not so well.”
- Need to hold a broader view during research/career development, and ask questions like
 - Is the current solution the only possible solution?
 - What are other possible solutions?
 - Can the current solution beat other possible solutions in all aspects?
 - Can the current solution be further improved with ideas from other possible solutions?

Continuation in Research Agenda

- Maintain a theme and continuation (go deep)
 - Ex. Focus on the same problem with (significant) improvement of your previous solution
 - Ex. Focus on a new problem with (significant) adaptation of your previous solution
 - Deep (significant) paper over shallow (insignificant) paper(s)
- Better to tell/reflect a coherent story (principle); set up identity to be in X area
- Pros
 - Reduce risk (with more certainty)
 - Reduce cost (reusing infrastructure/experience)
- Cons
 - Increase risk (what if a dead end?)
 - Reduce novelty (tend to be incremental; watch out/avoid LPU)

Example Research Agenda on Mining SE Data (for Tao's ASE group) 2005-

- My PhD research was on dynamic analysis (e.g., testing w/ spec inference [ASE 03] and bug avoidance w/ machine learning [ICSE 05])
- Interested in going from Dynamic to Static
 - Often not scalable with dynamic analysis
- Interested in applying data mining
 - Based on competitive advantage and research density, decide to mine code bases
- Learned about Koders.com during end of 2005
 - Thought: code search engine is so scalable but search results not good enough for SE tasks, why not Searching + Mining?
 - Discussed with Jian Pei (data mining) → MAPO [MSR 06]

Example Research Agenda – Cont.

- Then students drive the work and shape the research agenda....



...

Mithun Acharya

Suresh Thummalapenta



...

Xiaoyin Wang

Hao Zhong



Example Research Agenda – Cont.

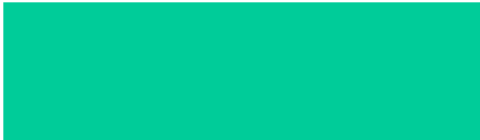
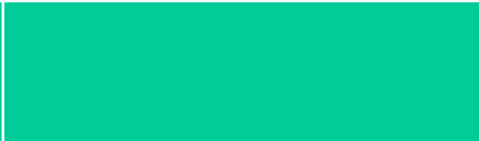
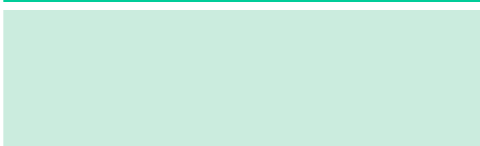
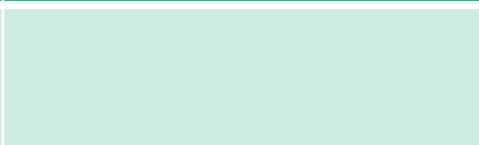
- **Static** tracing C infrastructure (**Mithun Acharya**)
 - API property generation [ASE 06S]
 - Mining interface details [ISSRE 06]
 - Mining partial orders [ESEC/FSE 07]
 - Mining error-handling defects [FASE 09]
- **Static** searching+tracing Java infrastructure incl partial code analysis (**Suresh Thummalapenta**)
 - PARSEWeb: Mining method sequences [ASE 07]
 - SpotWeb: Mining hotspots/coldspots [ASE 08]
 - CAR-Miner: Mining sequence association rules (exception-handling defects) [ICSE 09]
 - Alattin: Mining alternative patterns (neglected conditions) [ASE 09]

Example Research Agenda – Cont.

- Software reliability remains a focused task for mining
- Recently shift our competitive advantage/principle, e.g., Suresh Thummalapenta's work
 - Searching+mining/adopting advanced miners [FSE 07] → new patterns/mining algorithms [ICSE 09, ASE 09]
 - Static defect detection → test generation (partly due to new collaboration with MSR Pex) [FSE 09]
- Expand mining with our Chinese collaborators
 - API mining: **Hao Zhong, Lu Zhang**, et al.
 - MAPO - mining API sequences [ECOOP 09]
 - MAM – mining API mapping [ICSE 10]
 - Text mining: **Xiaoyin Wang, Hao Zhong, Lu Zhang** et al.
 - Mining bug reports+ execution traces [ICSE 08]
 - Mining API docs for properties [ASE 09, Best Paper]

Big Picture and Vision

- Step back and think about what research problems will be most important and most influential/significant to solve in the long term
 - Long term could be the whole career
- People tend not to think about important/long term problems

	Less important	More important
Shorter term		
Longer term		

Richard Hamming “**you and your research**”

<http://www.cs.virginia.edu/~robins/YouAndYourResearch.html>

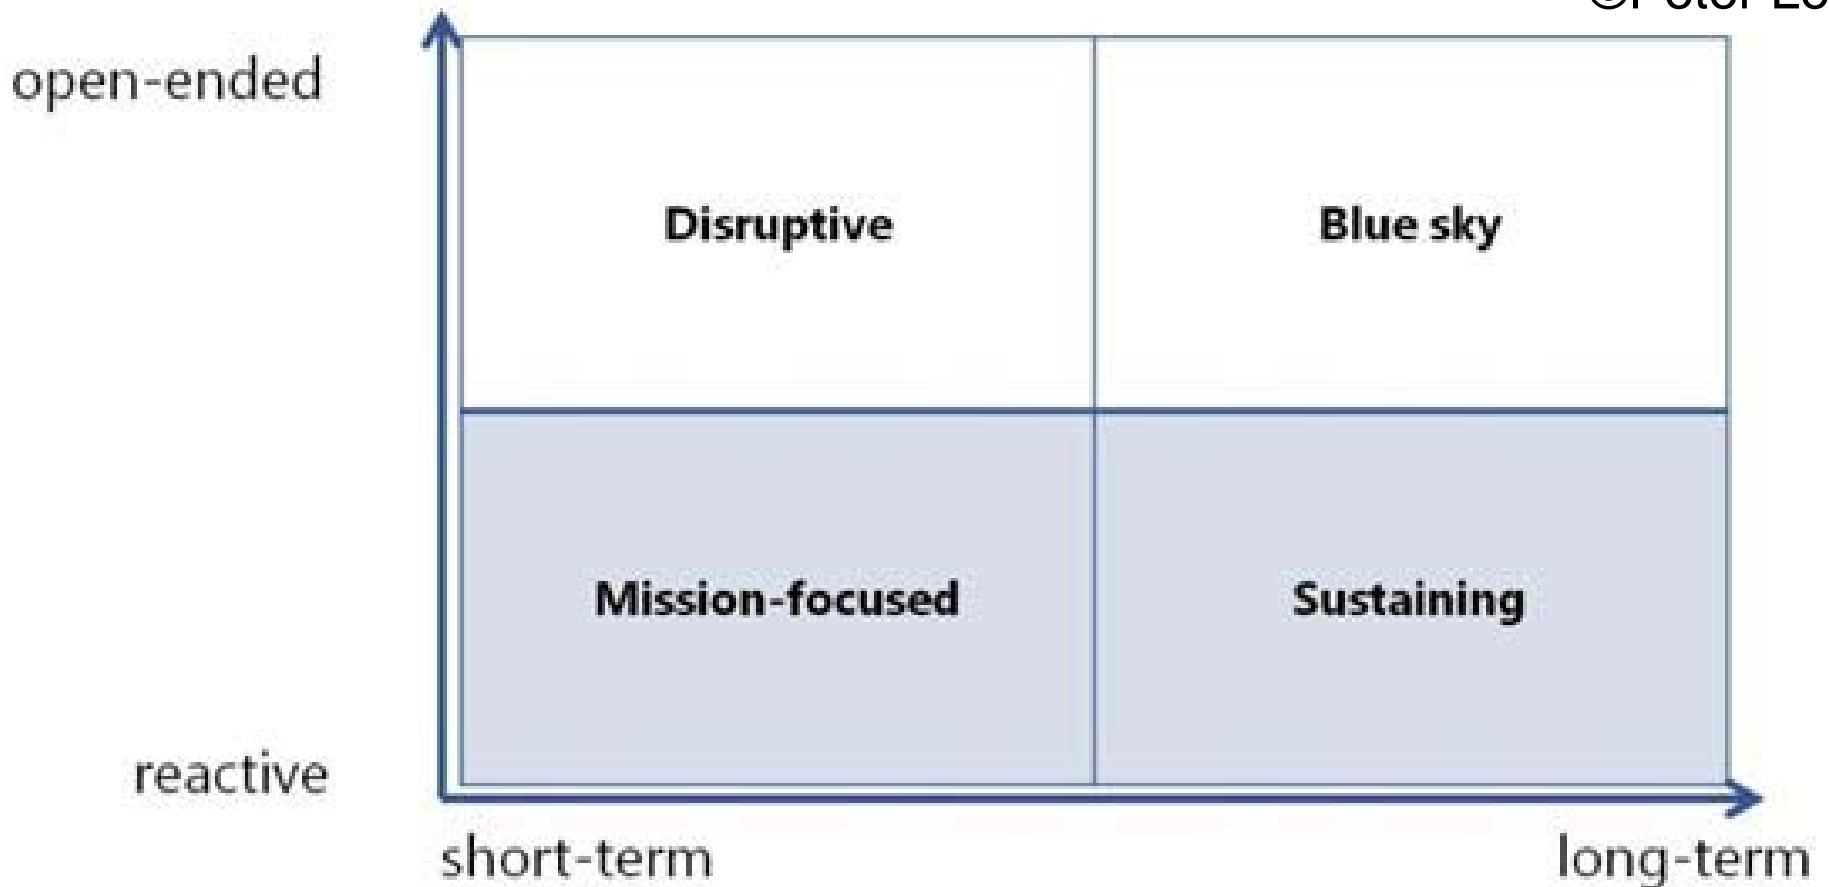
Ivan Sutherland “**technology and courage**”

http://labs.oracle.com/techrep/Perspectives/sml_i_ps-1.pdf

This slide was made based on discussion with David Notkin

Research Space

©Peter Lee



http://blogs.technet.com/b/inside_microsoft_research/archive/2011/12/31/microsoft-research-redmond-year-in-review.aspx a blog post by Peter Lee

Talk: The Pipeline from Computing Research to Surprising Inventions by Peter Lee

<http://www.youtube.com/watch?v=kpjw9Is14Q>

Big Picture and Vision –cont.

- If you are given 1 (4) million dollars to lead a team of 5 (10) team members for 5 (10) years, what would you invest them on?

More Reading

- <http://www.weizmann.ac.il/mcb/UriAlon/nurturing/HowToChooseGoodProblem.pdf>
- <http://www.cs.rutgers.edu/~ryder/DoingResNSEFS505.pdf>
- <https://sites.google.com/site/asergroup/advice>
 - <http://people.engr.ncsu.edu/txie/publications/writepapers.pdf>
 - <http://people.engr.ncsu.edu/txie/advice/researchskills.pdf>
 - <http://people.engr.ncsu.edu/txie/advice/gradstudentsurvival.pdf>
 - <http://people.engr.ncsu.edu/txie/adviceonresearch.html>
- <http://calnewport.com/blog/2008/11/07/does-being-exceptional-require-an-exceptional-amount-of-work/>

More Reading

- <https://sites.google.com/site/slesesymposium/slese12.pdf> by Zhendong Su
- <http://avandeursen.wordpress.com/2013/07/10/research-paper-writing-recommendations/> by Arie van Deursen
- Book: Crafting Your Research Future: A Guide to Successful Master's and Ph.D. Degrees in Science & Engineering by Charles Ling and Qiang Yang
 - <http://www.amazon.com/Crafting-Your-Research-Future-Engineering/dp/1608458105>