Photo credit: Dave DiCello

17-803 Empirical Methods

Bogdan Vasilescu, Institute for Software Research

Formulating Research Questions

Thursday, September 1, 2022

Outline for Today

- Briefly discuss the two readings
- Formulating research questions
  - Maybe an activity

Homework Readings Discussion

Two papers:

- What is the point of this paper?
- What is the methodology?
  - Why this choice of method?
- Do you trust the results?
  - Why or why not?
  - What are the risks of being misled?
  - How do you evaluate a study with this type of methodology?
  - What does it tell about other ecosystems?

How to Break an API: Cost Negotiation and Community Values in Three Software Ecosystems

Christopher Bogart,[1] Christian Kästner,[1] James Herbsleb,[1] Ferdian Thung[2]
[1]Carnegie Mellon University, USA [2]Singapore Management University, Singapore

ABSTRACT

Change introduces conflict into software ecosystems: breaking changes may ripple through the ecosystem and trigger rework for users of a package, but often developers can invest additional effort or a
downstream cos
three software ec
phies toward ch
to understand h
and change-rela
policies are use
fer substantially
change and that
different commu
illustrate that tl
an ecosystem, it
and there is valu
tradeoffs explicit
and negotiate ch

and may trigger rework in many dependent packages. Avoiding changes, however, may result in stale software projects, in dependencies with known defects, and in growing incompatibility with other tools and standards.
The burden of change can be borne by different partici

CCS Concepts
tion in software

Keywords: Sof
semantic versior

2014 14th IEEE International Working Conference on Source Code Analysis and Manipulation

Semantic Versioning versus Breaking Changes: A Study of the Maven Repository

Steven Raemaekers
Software Improvement Group
Amsterdam, The Netherlands
Email: s.raemaekers@sig.eu

Arie van Deursen
Technical University Delft
Delft, The Netherlands
Email: arie.vandeursen@tudelft.nl

Joost Visser
Software Improvement Group
Amsterdam, The Netherlands
Email: j.visser@sig.eu

***Abstract*—For users of software libraries or public programming interfaces (APIs), backward compatibility is a desirable trait. Without compatibility, library users will face increased risk and cost when upgrading their dependencies. In this study, we investigate semantic versioning, a versioning scheme which provides strict rules on major versus m inor and patch releases. We analyze seven years of library release history in Maven Central, and contrast version identifiers with actual incompatibilities. We find that around one third of all releases introduce at least one breaking change, and that this figure is the same for minor and major releases, indicating that version numbers do not provide developers with information in stability of interfaces. Additionally, we find that the adherence to semantic versioning principles has only marginally increased over time. We also investigate the use of deprecation tags and find out that methods get deleted without applying deprecated tags, and methods with deprecated tags are never deleted. We conclude the paper by arguing that the adherence to semantic versioning principles should increase because it provides users of an interface with a way to determine the amount of rework that is expected when upgrading to a new version.**

- MAJOR: This number should be incremented when incompatible API changes are made;
- MINOR: This number should be incremented when functionality is added in a backward-compatible manner;
- PATCH: This number should be incremented when backward-compatible bug fixes are made.

These principles were formulated in 2010 by (GitHub founder) Tom Preston-Werner.[2] As argued in the semantic versioning specification, *"these rules are based on but not necessarily limited to pre-existing widespread common practices in use in both closed and open-source software."*

But how common are these practices in reality? Are such changes just harmless, or do they actually hurt by causing rework? Do breaking changes mostly occur in major releases, or do they occur in minor releases as well? Do major and minor releases differ in terms of typical size? Furthermore, for the breaking changes that do occur, to

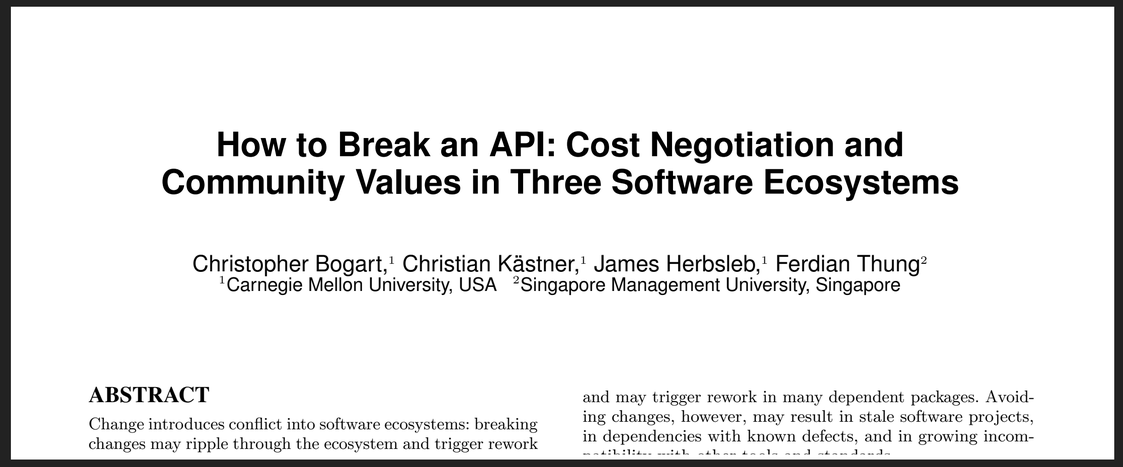

How to Break an API: Cost Negotiation and Community Values in Three Software Ecosystems

Christopher Bogart,[1] Christian Kästner,[1] James Herbsleb,[1] Ferdian Thung[2]
[1]Carnegie Mellon University, USA [2]Singapore Management University, Singapore

ABSTRACT

Change introduces conflict into software ecosystems: breaking changes may ripple through the ecosystem and trigger rework and may trigger rework in many dependent packages. Avoiding changes, however, may result in stale software projects, in dependencies with known defects, and in growing incom-

- Describe state of practice
- Compare and contrast environments
- Look for (theoretical) relationships between practices and values

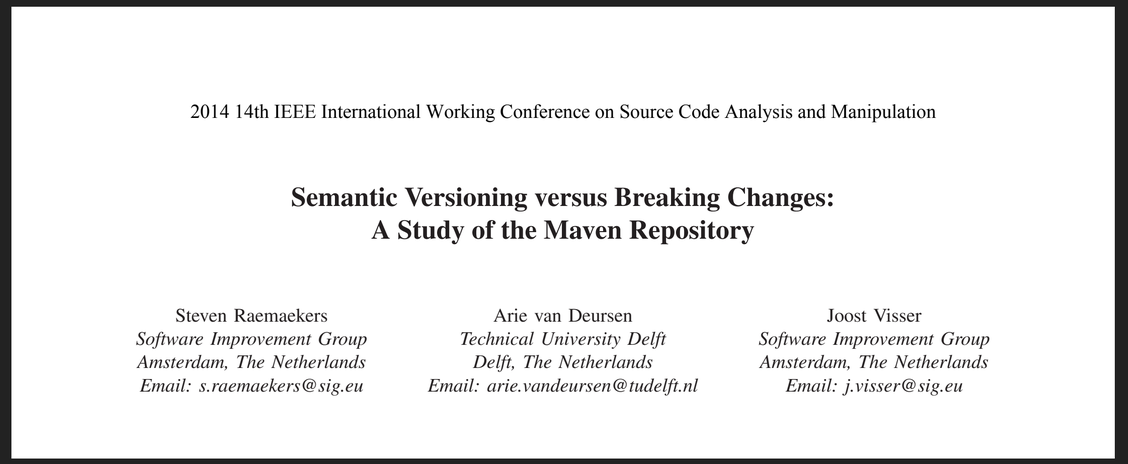

2014 14th IEEE International Working Conference on Source Code Analysis and Manipulation

Semantic Versioning versus Breaking Changes: A Study of the Maven Repository

Steven Raemaekers
Software Improvement Group
Amsterdam, The Netherlands
Email: s.raemaekers@sig.eu

Arie van Deursen
Technical University Delft
Delft, The Netherlands
Email: arie.vandeursen@tudelft.nl

Joost Visser
Software Improvement Group
Amsterdam, The Netherlands
Email: j.visser@sig.eu

- Collect evidence about the incidence of some phenomenon in the wild
- Look for (statistical) relationships between practices

How to Break an API: Cost Negotiation and Community Values in Three Software Ecosystems

Christopher Bogart,[1] Christian Kästner,[1] James Herbsleb,[1] Ferdian Thung[2]
[1]Carnegie Mellon University, USA [2]Singapore Management University, Singapore

ABSTRACT

Change introduces conflict into software ecosystems: breaking changes may ripple through the ecosystem and trigger rework and may trigger rework in many dependent packages. Avoiding changes, however, may result in stale software projects, in dependencies with known defects, and in growing incom-

- Interviews + some data mining

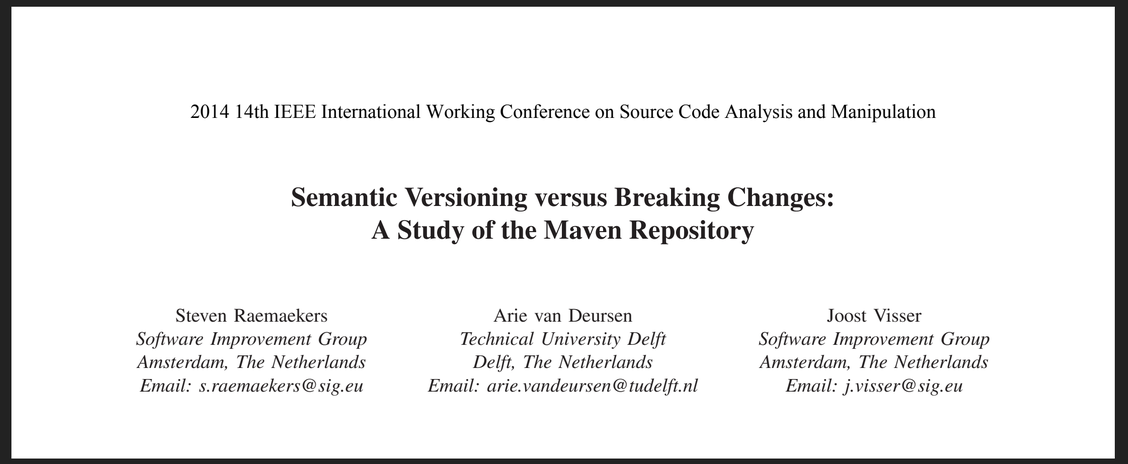

2014 14th IEEE International Working Conference on Source Code Analysis and Manipulation

Semantic Versioning versus Breaking Changes: A Study of the Maven Repository

Steven Raemaekers
Software Improvement Group
Amsterdam, The Netherlands
Email: s.raemaekers@sig.eu

Arie van Deursen
Technical University Delft
Delft, The Netherlands
Email: arie.vandeursen@tudelft.nl

Joost Visser
Software Improvement Group
Amsterdam, The Netherlands
Email: j.visser@sig.eu

- Observational study
  - But lucky they didn’t have to conduct the observations themselves
- A qualitative analysis might be more useful for the question asking what factors cause dependencies to not be updated?

How to Break an API: Cost Negotiation and Community Values in Three Software Ecosystems

Christopher Bogart,[1] Christian Kästner,[1] James Herbsleb,[1] Ferdian Thung[2]
[1]Carnegie Mellon University, USA [2]Singapore Management University, Singapore

ABSTRACT

Change introduces conflict into software ecosystems: breaking changes may ripple through the ecosystem and trigger rework and may trigger rework in many dependent packages. Avoiding changes, however, may result in stale software projects, in dependencies with known defects, and in growing incom-

- Who were the participants?
- What questions were they asked?
- Does it seem like the authors cherry picked their quotes and findings or did they use a systematic approach?

- Are the opinions of the sampled individuals representative of others?
- “I am inclined to trust the results, as they broadly align with my own anecdotal evidence.”

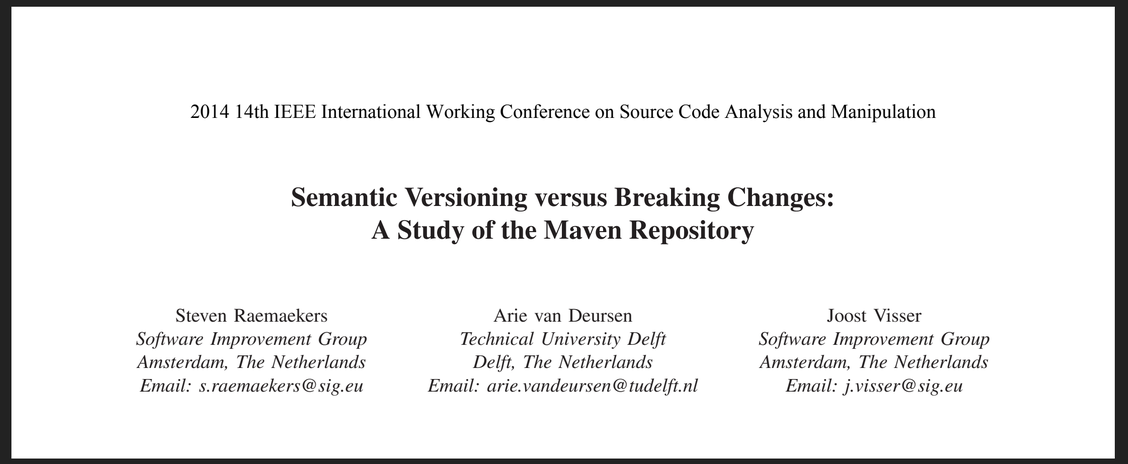

2014 14th IEEE International Working Conference on Source Code Analysis and Manipulation

Semantic Versioning versus Breaking Changes: A Study of the Maven Repository

Steven Raemaekers
Software Improvement Group
Amsterdam, The Netherlands
Email: s.raemaekers@sig.eu

Arie van Deursen
Technical University Delft
Delft, The Netherlands
Email: arie.vandeursen@tudelft.nl

Joost Visser
Software Improvement Group
Amsterdam, The Netherlands
Email: j.visser@sig.eu

- Data validity?
- Analysis validity?

- “I don't trust the results because the method of [measuring X is imperfect]”

How to Break an API: Cost Negotiation and Community Values in Three Software Ecosystems

Christopher Bogart,[1] Christian Kästner,[1] James Herbsleb,[1] Ferdian Thung[2]
[1]Carnegie Mellon University, USA [2]Singapore Management University, Singapore

ABSTRACT

Change introduces conflict into software ecosystems: breaking changes may ripple through the ecosystem and trigger rework and may trigger rework in many dependent packages. Avoiding changes, however, may result in stale software projects, in dependencies with known defects, and in growing incom-

- Tells a bit about other ecosystems
- “I am still skeptical that the results generalize beyond the three ecosystems given.”

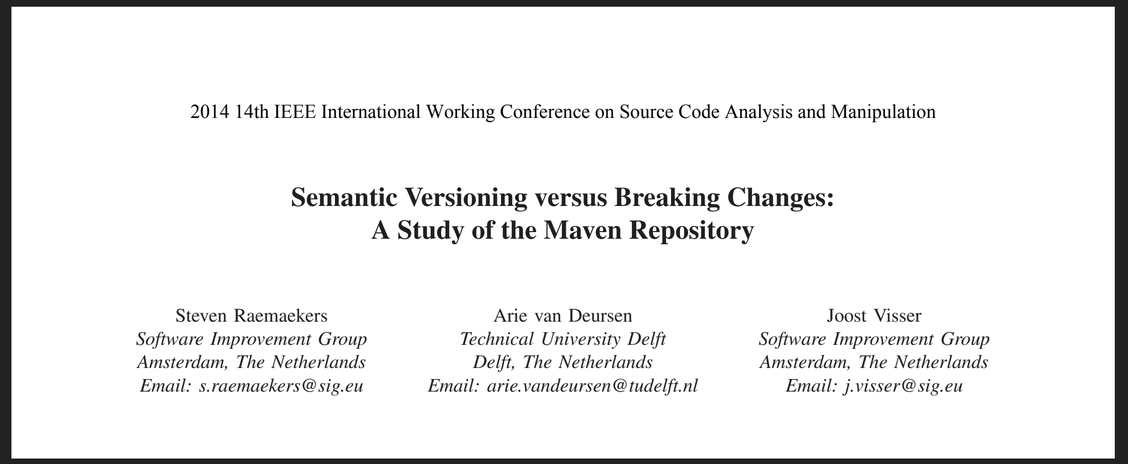

2014 14th IEEE International Working Conference on Source Code Analysis and Manipulation

Semantic Versioning versus Breaking Changes: A Study of the Maven Repository

Steven Raemaekers
Software Improvement Group
Amsterdam, The Netherlands
Email: s.raemaekers@sig.eu

Arie van Deursen
Technical University Delft
Delft, The Netherlands
Email: arie.vandeursen@tudelft.nl

Joost Visser
Software Improvement Group
Amsterdam, The Netherlands
Email: j.visser@sig.eu

- Doesn’t tell much about other ecosystems

Formulating Research Questions

Meet Jane

- Jane's intuition is that the fisheye-view file navigator is more efficient for file navigation than a traditional file navigator.
  - File navigation requires a lot of scrolling and many clicks to find files.
  - "Fisheye-views" display information in a compact format that could potentially reduce the amount of scrolling required.
- Critics argue:
  - difficult to read
  - developers won't adopt
- Jane's research goal: collect evidence that supports or refutes her intuition

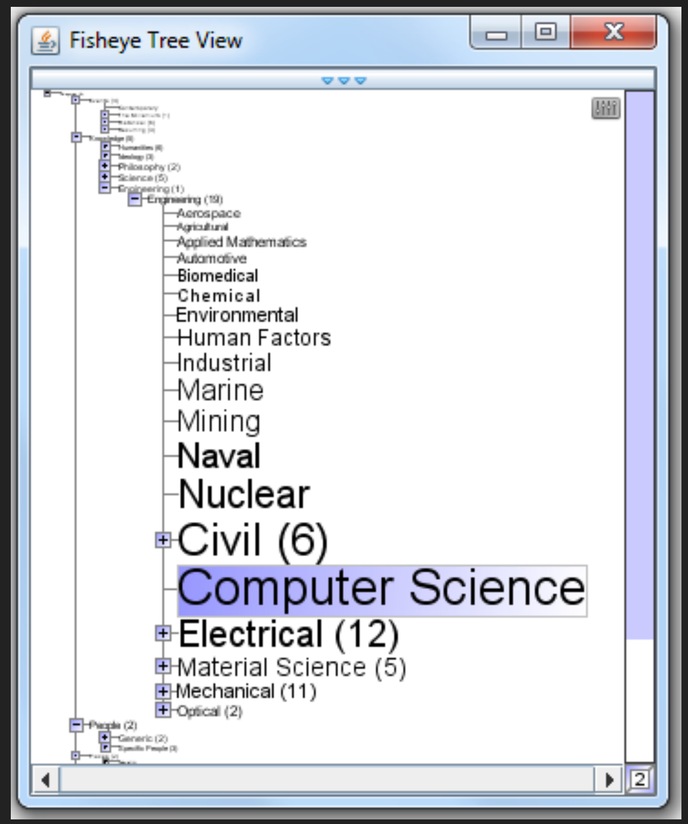

Meet Joe

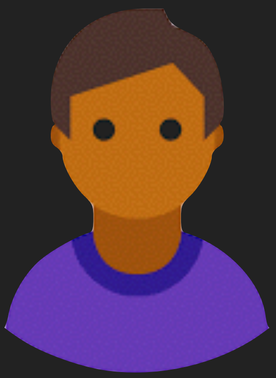

- Joe is interested in how developers in industry use (or not) UML diagrams during software design.
  - His professors recommended UML.
  - His EvilCorp internship indicates that UML is rarely used.
- Joe's research goals:
  - Explore how widely UML diagrams are used in industry.
  - Explore how these diagrams are used as collaborative shared artifacts during design.

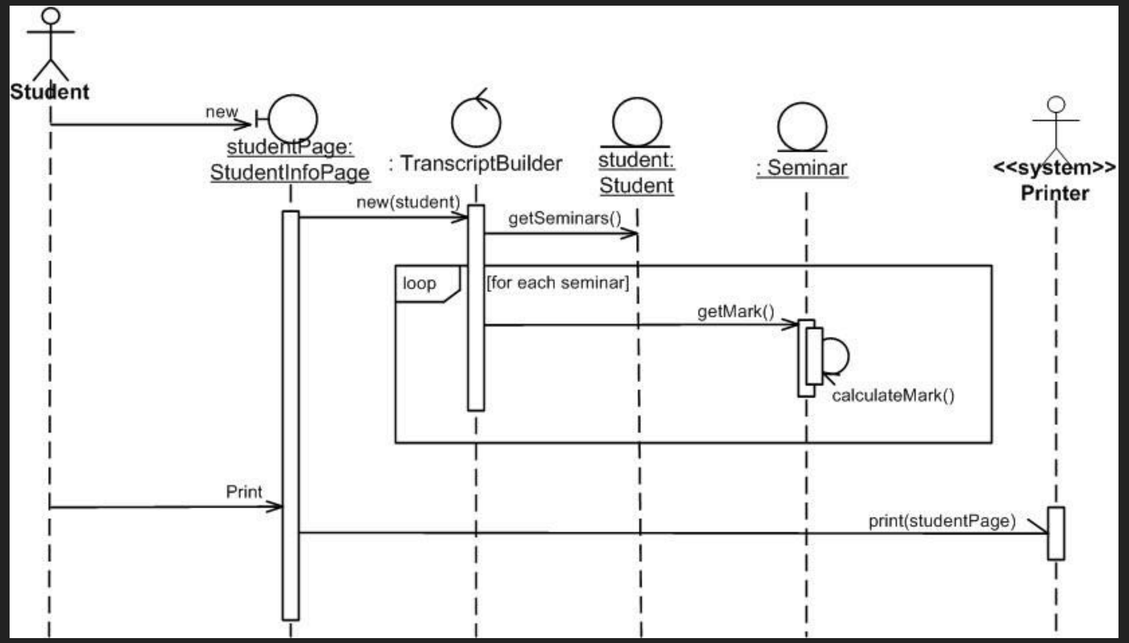

From Problems To Research Questions

"Is a fisheye-view file navigator more efficient than the traditional view for file navigation?"

Jane

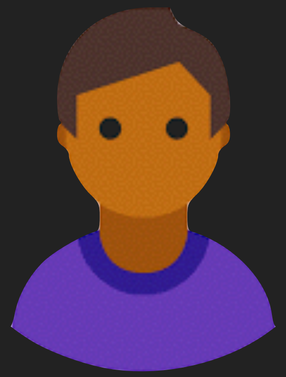

"How widely are UML diagrams used as collaborative shared artifacts during design?"

Joe

The Most Obvious Question Is Not Always the Best Choice for a Starting Point

Jane

"Is a fisheye-view file navigator more efficient than the traditional view for file navigation?"

- Do we already know that some people (who?) need to do file navigation?
- What does file navigation mean exactly?
- Under what circumstances do these people do file navigation?
- Is efficiency (measured how?) a relevant goal for these people?

The Most Obvious Question Is Not Always the Best Choice for a Starting Point

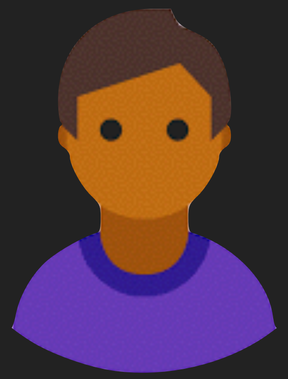

Joe

"How widely are UML diagrams used as collaborative shared artifacts during design?"

- What's a "collaborative shared artifact"?
- Can we reliably identify one?
- Can we reliably say which things are and aren't UML diagrams?

Both questions are vague, because they make assumptions about the phenomena to be studied, and kinds of situation in which these phenomena occur.

Some possible (better) questions Jane and Joe could have asked

Exploratory Questions

- Existence questions
  -  “Is file navigation something that (certain types of programmers) actually do?”
  -  “Is efficiency actually a problem in file navigation?”
  -  “Do collaborative shared artifacts actually exist?”
- Description and Classification questions
  -  “How can we measure efficiency for file navigation?”
  -  “What are all the types of collaborative shared artifacts?”
- Descriptive-Comparative questions
  -  “How do fisheye views differ from conventional views?”
  -  “How do UML diagrams differ from other representations of design information?”

Outcomes:

- Clearer understanding of the phenomena
- More precise definitions of the theoretical terms
- Evidence that we can measure them
- Evidence that the measures are valid

Base-Rate Questions (Normal Patterns of Occurrence of Phenomena)

- Frequency and distribution questions

"How many distinct UML diagrams are created in software development projects in large software companies?"

- Descriptive-Process questions

"How do programmers navigate files using existing tools?"

Outcomes:

- Basis for saying whether a particular situation is normal or unusual

Relationship Questions

- Relationship questions

"Does efficiency in file navigation correlate with the programmer's familiarity with the programming environment?"

"Do managers' claims about how often they use UML correlate with the actual use of UML?"

Outcomes:

- Establish that occurrence of one phenomenon is related to occurrence of another

Causality Questions

- Causality questions

"Do fisheye-views cause an improvement in efficiency for file navigation?"

- Causality-Comparative questions

"Do fisheye-views cause programmers to be more efficient at file navigation than conventional views? "

- Causality-Comparative Interaction questions

"Do fisheye-views cause programmers to be more efficient at file navigation than conventional views when programmers are distracted, but not otherwise?"

Outcomes:

- Explain why a relationship holds by attempting to identify a cause and effect
- Understand how context affects a cause–effect relationship

Design Questions

- Design questions

"What is an effective way for teams to represent design knowledge to improve coordination?"

Outcomes:

- Design better procedures and tools for carrying out some activity
- Design suitable social or regulatory policies

A long term research program in an applied discipline (e.g., SE) typically involves a mix of both types of questions (knowledge and design).

Remember Last Lecture? (What Will You Accept as Valid Answers?)

Positivist ?

Constructivist ?

Remember Last Lecture? (What Will You Accept as Valid Answers?)

Positivist 

- Controlled experiments in laboratory conditions are the only source of trustworthy evidence.
  - to prove that A causes B is to manipulate A in a controlled setting, and measure the effect on B.

Constructivist 

- “Lab experiments are useless, they ignore the messy complexity of real software projects.”
  - Field work instead!
- Judgments about “improvements” to file navigation are subjective.
- Contextual factors such as distractions have a major impact.

It is impossible to avoid some commitment to a particular stance, as you cannot conduct research, and certainly cannot judge its results, without some criteria for judging what constitutes valid knowledge.

Discussion: What kind of questions did Bogart and Raemaekers ask?

How to Break an API: Cost Negotiation and Community Values in Three Software Ecosystems

Christopher Bogart,[1] Christian Kästner,[1] James Herbsleb,[1] Ferdian Thung[2]
[1]Carnegie Mellon University, USA [2]Singapore Management University, Singapore

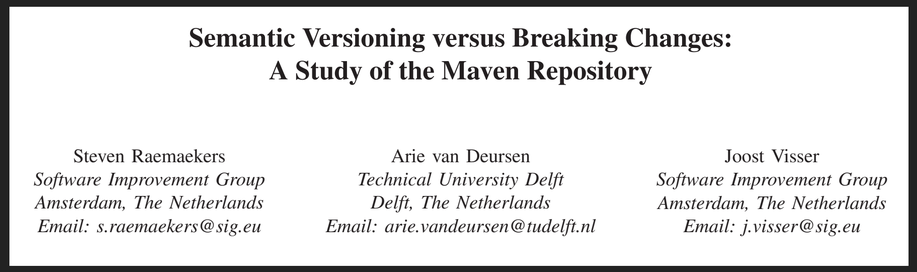

Semantic Versioning versus Breaking Changes: A Study of the Maven Repository

Steven Raemaekers
Software Improvement Group
Amsterdam, The Netherlands
Email: s.raemaekers@sig.eu

Arie van Deursen
Technical University Delft
Delft, The Netherlands
Email: arie.vandeursen@tudelft.nl

Joost Visser
Software Improvement Group
Amsterdam, The Netherlands
Email: j.visser@sig.eu

- How do developers make decisions about whether and when to perform breaking changes? How do they mitigate or delay costs for other developers?
- How do developers react to and manage change in their dependencies?
- How do policies, tooling, and community values influence decision making?

- How are semantic versioning principles applied in practice in the Maven repository?
- Has the adherence to semantic versioning principles increased over time?
- How are dependencies to newer versions updated? What are factors causing systems not to include the latest versions of dependencies?
- How are deprecation tags applied to methods in the Maven repository?

Activity: in Breakouts, Choose "Best" Method for Answering These Questions

- Why do [engineers] ignore [security warnings in their code]?
- Does [test driven development] improve [code quality]?
- Which [code review tool] reveals [more bugs]?
- Do the topics discussed in [online technical forums] deter the involvement of [women]? Has this changed since online learning?
- How often does [this software] [fail] and in what ways?

Adapted from an activity by Peggy Storey

Credits

- Graphics:
  - Dave DiCello photography (cover)
- Content:
  - Easterbrook, S., Singer, J., Storey, M. A., & Damian, D. (2008). Selecting empirical methods for software engineering research. In Guide to advanced empirical software engineering (pp. 285-311). Springer, London.

Bonus Slides

Q: What's the Relationship?
Does Inductive → Qualitative? Deductive → Quantitative?

A: No. These Perspectives Are Frequently (and Falsely) Conflated

Quantitative (numerical) data	Qualitative (non-numerical) data
Exploratory factor analysis[1]: A set of statistical techniques, typically applied to assessment- or survey-generated data, that identifies the underlying theoretical constructs (i.e., "factors") of the phenomena of interest, which researchers then name inductively.	**Traditional grounded theory[2]:** A research methodology aimed at generating a local theory using qualitative data collected with participants. The researcher acts as a tabula rasa (blank slate), making sense of a phenomenon without biasing that interpretation.
Structural equation modeling[3]: A set of approaches to analysis whereby several statistical models are built based on theory and then tested with an appropriate dataset. The model that best fits the data is considered superior.	**Ethnography[4]:** This approach often combines existing exploratory theories about human social behavior and collects data to observe and describe the culture of a group.
Bayesian approaches to analysis[5]: These analytic approaches use previous knowledge, available data, or beliefs, to acknowledge the a priori likelihood of findings. The strength of prior evidence influences analysis.	**Constructivist grounded theory[6]:** This approach analyzes data with the goal of contributing to previous understandings of a phenomenon, building explicitly on others' work.
Analysis of hypothesis-driven experimental studies[1]: Among the most familiar, these approaches result in a test of significance, which leads to the rejection of, or failure to reject, the null hypothesis (that the findings are due to chance).	**Deductivist content analysis[7]:** Qualitative data are analyzed using a predetermined theory or theoretical framework, to find examples of constructs and support or challenge them.

Young, M., Varpio, L., Uijtdehaage, S., & Paradis, E. (2020). The spectrum of inductive and deductive research approaches using quantitative and qualitative data. Academic Medicine, 95(7), 1122.