

[ECE1785H] Empirical Software Engineering

Overview	1
Learning Goals	2
Schedule	2
Evaluation	3
Syllabus	3
1. Contrasting methods	3
2. Literature Review and Theory	3
3. Interviewing	4
4. Grounded Theory	4
5. Survey	5
6. Intro Quantitative Analysis	5
7. Experiment	6
8. Quasi-experimental Design & Linear Regression	6
9. Time Series Analysis	6
10. Mixed-Methods	7
11. Social Network Analysis	7
Course Project	7
[Milestone 1] Initial project description	8
[Milestone 2] Interim Report	8
[Milestone 3] Final project paper & presentation	8

Overview

Recent decades have witnessed the emergence of two important trends in software systems. First, an increasing number of software technologies have a significant social component (e.g. end-user programming, collaborative development tools). Second, software systems have increased significantly in complexity and size challenging traditional development and testing approaches. Empirical research methods play a key role in the evaluation of tools and technologies, and in testing the social and technical theories they embody.

This course provides an overview and hands-on experience with a core of qualitative and quantitative empirical research methods, including interviews, qualitative coding, survey design, and large-scale mining and analysis of data. There will be extensive reading with occasional student presentations about the reading in class, weekly homework assignments, and a semester-long research project for which students must prepare in-class kickoff and final presentations as well as a final report.

We will focus on software engineering related research questions in readings and assignments. Students will mine and integrate data from and across online software repositories (e.g., GitHub and Stack Overflow) and employ a spectrum of data analysis techniques, ranging from statistical modeling to social network analysis. For the final research project, we encourage students to come up with a research question of interest to themselves. The delivery will be a research paper, and one or more empirical methods presented in class have to be part of the paper.

Learning Goals

The learning goals describe what I want students to know or be able to do by the end of the semester. I evaluate whether learning goals have been achieved through assignments, written project reports, and in-class presentations. All learning goals are roughly written in a form "after taking this class, the student should be able to ...".

- Summarize and interpret a body of literature on a particular topic; identify gaps in the literature; write a literature review
- Formulate and motivate research questions
- Understand what research designs and research methods are available for empirical research
- Compare the suitability of different research designs and research methods in different scenarios; explain the relative strengths and weaknesses
- Design empirical studies for different purposes (e.g., evaluating a tool, understanding a phenomenon); choose appropriate methods and defend the choice
- Combine research methods in a mixed-methods design
- Collect and analyze qualitative and quantitative data
- Design interview protocols and user surveys
- Code qualitative data
- Mine data from online repositories
- Run statistical tests and interpret results
- Build, validate, and interpret regression models
- Draw conclusions from empirical data
- Present results verbally and in writing

Schedule

We cover the following topics ():

Week	Topic
1	Introduction
2	Literature Review & Theory
3	Interviews
4	Qualitative Data Analysis
5	Survey Research
6	Experimentation
7	Quantitative Data Analysis

8	T-test & F-statistics & Chi-Square statistics
9	Case Studies
10	Time Series Analysis
11	Social Network Analysis

Evaluation

Evaluation will be based on the following approximate percentages:

- 40% assignments
- 50% research project
 - 10% initial project description (proposal)
 - 2% interim report
 - 8% final presentation
 - 30% final report
- 10% participation and in-class presentations

Syllabus

1. Contrasting methods

a. Method:

- (Ch.1) Creswell, John W., and J. David Creswell. *Research design: Qualitative, quantitative, and mixed methods approaches*. Sage publications, 2017.

b. Reading

- Bogart, Christopher, et al. "**How to break an API: cost negotiation and community values in three software ecosystems.**" Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering, 2016..
- Raemaekers, Steven, Arie Van Deursen, and Joost Visser. "**Semantic versioning versus breaking changes: A study of the maven repository.**" *2014 IEEE 14th International Working Conference on Source Code Analysis and Manipulation*. IEEE, 2014.
- (Optional) Easterbrook, Steve, et al. "**Selecting empirical methods for software engineering research.**" *Guide to advanced empirical software engineering*. Springer, London, 2008. 285-311.

c. Assignment:

Read the assigned chapter and papers, and consider the differences in methods, the research questions they addressed, and the evidence they used to reach their conclusions. What does this tell you about the differences between qualitative and quantitative methods? In general, when is each type of method appropriate? What weaknesses does each method suffer from? Summarize your conclusions in 1-2 pages, and be prepared to informally present and discuss them in class.

2. Literature Review and Theory

a. Method

- (Ch.2 &3) Creswell, John W., and J. David Creswell. *Research design: Qualitative, quantitative, and mixed methods approaches*. Sage publications, 2017.

b. Reading

- **Theory-oriented papers:**

- Xiao, S., Witschey, J., & Murphy-Hill, E. (2014). Social influences on secure development tool adoption: why security tools spread, Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing (pp. 1095-1106): ACM.
- (Optional) (Ch. 1) Mokyry, Joel. The gifts of Athena: Historical origins of the knowledge economy. Princeton University Press, 2002.

3. Interviewing

a. Method

- Powell, Martine B., Ron P. Fisher, and Rebecca Wright. "Investigative interviewing." *Psychology and law: An empirical perspective* (2005): 11-42.
- (Ch 4&6). Interviewing as qualitative research: A guide for researchers in education and the social sciences: Teachers college press.
- Cassell, Catherine, and Gillian Symon, eds. *Essential guide to qualitative methods in organizational research*. Sage, 2004.

b. Example

- Grinter, Rebecca E., and Leysia Palen. "Instant messaging in teen life." *Proceedings of the 2002 ACM conference on Computer supported cooperative work*. 2002.
- Chattopadhyay, Souti, et al. "What's Wrong with Computational Notebooks? Pain Points, Needs, and Design Opportunities." *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020.

c. Assignment

Think about a research idea that related to your own area and you need to collect information from stakeholders through interview. (see examples above)

- 1) decide what your purpose is and write a sentence describing it.
- 2) develop an interview protocol. The protocol can be short, focusing on exactly what you are interested in. You should anticipate short interviews, perhaps 15-20 minutes at most.
- 3) conduct two interviews.
- 4) be prepared to tell the class what you learned, how the interviews went, any problems or lessons you can share. In future classes, we will learn more structured ways of analyzing qualitative data such as interview transcripts.

4. Grounded Theory

a. Method

- (Ch 9) Creswell, John W., and J. David Creswell. *Research design: Qualitative, quantitative, and mixed methods approaches*. Sage publications, 2017.
- (Ch 4&5) Miles, Matthew B., A. Michael Huberman, and Johnny Saldana. "Qualitative data analysis: A methods sourcebook." (2014).

b. Example

- Razavi, M. N., & Iverson, L. (2006). A grounded theory of information sharing behavior in a personal learning space, Proceedings of the ACM Conference on Computer Supported Cooperative Work (pp. 459-468).
- de Souza, C. R., & Redmiles, D. F. (2008). An empirical study of software developers' management of dependencies and changes, Proceedings of the 30th International Conference on Software Engineering (pp. 241-250).

c. Exercise

Transcribe the interviews you recorded last week. Write down a research question (or two) that you think you can answer with the interviews. Develop a suitable coding scheme for the collaborative writing interviews you performed, and apply the codes either to your detailed notes or (preferably your transcription of the interviews). Write an analytic memo (2-3 paragraph) based on these codes.

5. Survey

a. Method

- (Ch 5) Creswell, John W., and J. David Creswell. *Research design: Qualitative, quantitative, and mixed methods approaches*. Sage publications, 2017.
- (Ch 1&5) Dillman, Don A., Jolene D. Smyth, and Leah Melani Christian. *Internet, phone, mail, and mixed-mode surveys: the tailored design method*. John Wiley & Sons, 2014.

b. Example

- Henne, B., Harbach, M., & Smith, M. (2013). Location privacy revisited: factors of privacy decisions. Extended Abstracts on Human Factors in Computing Systems.
- Shklovski, I., Mainwaring, S. D., Skúladóttir, H. H., & Borgthorsson, H. (2014). Leakiness and creepiness in app space: perceptions of privacy and mobile app use, Proceedings of the ACM conference on Human factors in computing systems (pp. 2347-2356): ACM.

c. Assignment

Think about what you learned about from the interviews you conducted, and select an issue for further investigation that would be appropriate for a survey method. Design a short (8-10) item questionnaire that is well designed to address this issue, and enter the survey into an online tools such as SurveyMonkey. In a short writeup, provide a link to the questionnaire, discuss your selection of open versus closed ended questions, any issues that arose in the wording or presentation of questions, what population you would sample, and how you would invite participants.

6. Intro Quantitative Analysis

a. Method

- (Ch 10) **Analysis and Interpretation**. from C. Wohlin et al., Experimentation in Software Engineering, Springer-Verlag Berlin Heidelberg 2012
- (Ch 6) **Statistical Methods and Measurement**. from F. Shull et al. (eds.), Guide to Advanced Empirical Software Engineering. Springer 2008 (similar content as the Wohlin chapter but slightly different presentation; read one or the other)
- (Ch 6) **Hypothesis Testing**. from MacKenzie. Human-Computer Interaction. Elsevier 2013

b. Example

- Filippova, A., Trainer, E., & Herbsleb, J. D. (2017). From diversity by numbers to diversity as process: supporting inclusiveness in software development teams with brainstorming. In Proceedings of the 39th International Conference on Software Engineering (pp. 152-163). IEEE. [focus on the quantitative analysis of survey responses]
- Vasilescu, B., Filkov, V., & Serebrenik, A. (2015). Perceptions of diversity on GitHub: A user survey. In Proceedings of the Eighth International Workshop on Cooperative and Human Aspects of Software Engineering (pp. 50-56). IEEE. [focus on the quantitative analysis of survey responses]
- Kaptein, M., & Robertson, J. (2012, May). Rethinking statistical analysis methods for CHI. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 1105-1114). ACM. [focus on the threats to validity and ways to mitigate]

c. Assignment

In this assignment, you will practice basic quantitative data analysis methods. We will mine the 2017 Open Source Survey results from Zenodo:
<https://zenodo.org/record/806811> (https://zenodo.org/record/806811#.W7PD0y_Mz1L)
 Formulate two research questions about participating in open source development, motivate them in 1-2 paragraphs with a few citations to relevant literature, and answer them using a quantitative analysis of data, e.g., based on ANOVA or multiple linear regression. Go beyond the basic frequency counts from R. Stuart Geiger's paper and

focus your research questions on correlations, regressions, or descriptive breakdowns between subgroups.

7. Experiment

a. Method

(Ch 1,2&8) **Experiments and Generalized Causal Inference, Statistical Conclusion Validity and Internal Validity, Randomized Experiments.** Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*: Wadsworth Cengage learning

b. Example

- Tomkins, A., Zhang, M., & Heavlin, W. D. (2017). Single versus double blind reviewing at WSDM 2017. arXiv preprint arXiv:1702.00502.

c. Assignment

Design an experiment based on your previous analysis of the GitHub open source survey results. The experiment should allow you to either test (some of) the same hypotheses you explored statistically, or an interesting or suggestive finding emerging out of that analysis.

Prepare a very short report describing your experiment design, including:

- your experimental hypothesis,
- a description of the experimental and control groups,
- your experimental procedure,
- acquiring or preparing any materials you need,
- the type of participants you require and how you would recruit them,
- how you would analyze the data,
- analysis of the threats to validity

8. Quasi-experimental Design & Linear Regression

a. Method

- (Ch 2) **The Simple Regression Model** Woolridge, J. M. (2003). *Introductory econometrics: A modern approach*. Thomson, Mason. [skim]
- (Ch 1&2) **General Aspects of Fitting Regression Models.** F.E. Harrell, Jr., *Regression Modeling Strategies*, Springer Series in Statistics, Chapters 1&2 - Regression general aspects: [Chapter 1: skim] [Chapter 2: read 2.1--2.3, 2.7]

b. Example

- Sinatra, R., Wang, D., Deville, P., Song, C., & Barabási, A. L. (2016). Quantifying the evolution of individual scientific impact. *Science*, 354(6312), aaf5239.
- Lim, S. (2009). How and why do college students use Wikipedia? *Journal of the Association for Information Science and Technology*, 60(11), 2189-2202.
- Bird, C., Nagappan, N., Devanbu, P., Gall, H., & Murphy, B. (2009). Does distributed development affect software quality? An empirical case study of Windows Vista. *Communications of the ACM*, 52(8), 85-93.

9. Time Series Analysis

a. Method

- Cowpertwait, P. S., & Metcalfe, A. V. (2009). *Introductory time series with R*. Springer Science & Business Media.
- (Ch 10.) **Basic Regression Analysis with Time Series Data.** Woolridge, J. M. (2003). *Introductory econometrics: A modern approach*. Thomson, Mason.

b. Example

- Kenmei, B., Antoniol, G., & Di Penta, M. (2008). Trend analysis and issue prediction in large-scale open source systems. In *Software Maintenance and Reengineering, 2008. CSMR 2008. 12th European Conference on* (pp. 73-82). IEEE.

- Trockman, A., Zhou, S., Kästner, C., & Vasilescu, B. (2017). Adding Sparkle to Social Coding: An Empirical Study of Repository Badges in the npm Ecosystem.

10. Mixed-Methods

a. Method

- (Ch 10) Creswell, John W., and J. David Creswell. *Research design: Qualitative, quantitative, and mixed methods approaches*. Sage publications, 2017.
- Venkatesh, V., Brown, S. A., & Bala, H. (2013). Bridging the qualitative-quantitative divide: Guidelines for conducting mixed methods research in information systems. *MIS quarterly*, 37(1), 21-54.

b. Example

- Greiler, M., Deursen, A. V., & Storey, M. A. (2012). Test confessions: a study of testing practices for plug-in systems. In Proceedings of the 34th International Conference on Software Engineering (pp. 244-254). IEEE Press.
- Trockman, A., Zhou, S., Kästner, C., & Vasilescu, B. (2017). Adding Sparkle to Social Coding: An Empirical Study of Repository Badges in the npm Ecosystem.

11. Social Network Analysis

a. Method

- (Ch 1,2 & 9) Graph Theory, Community. "The New Science of Networks" by Albert-László Barabási. Cambridge University Press, 2016

b. Example

- Bird, C., Pattison, D., D'Souza, R., Filkov, V., & Devanbu, P. (2008). Latent social structure in open source projects. In Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering (pp. 24-35). ACM.
- Backstrom, L., & Kleinberg, J. (2014). Romantic partnerships and the dispersion of social ties: a network analysis of relationship status on Facebook. In Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (pp. 831-841). ACM.

Group Project

A major component of this course is a proposal for a set of empirical studies that will ideally advance your own research. The final deliverable is essentially a proposal that resembles the introduction and methods section of an empirical research paper. The introduction should review the relevant literature and motivate the empirical work you propose. It should carefully lay out the current state of knowledge, identify the gaps you intend to fill or next steps you intend to perform, and culminate in the research questions you plan to answer. The methods section will describe how you plan to address the question, why the methods you describe are appropriate, how the multiple studies fit together, how you will develop any needed materials, recruit participants, and any other critical requirements. You will conclude by describing what you hope or expect to learn from the studies.

I want to do this in steps so that you can get feedback before tackling the final paper. I would expect the final project to be approximately 5 pages in ACM conference format (e.g., [ICSE](#), [FSE](#)).

In all cases in this course, the ultimate goal is to do what makes sense. Doing exactly what I requested is a secondary goal. If you feel any of your deliverables should deviate from what I asked for, let's talk about it. You may well be right. Strive for a project you actually could and would be interested in carrying out.

You may collaborate on this if it makes sense to do so. Send me an e-mail telling me why it makes sense, e.g., different expertise, divide and conquer strategy, etc.

[Milestone 1] Initial project description

A few sentences describing your research question, why it is important, and how it fits into your research (assuming it does).

- A few sentences describing the methods you plan to employ, why they are appropriate, and why they make sense together.
- The kinds of results you might get and what the contribution would be.

None of these are set in concrete -- they may well change as we learn about more methods, you get feedback from me and fellow students, and you have more time to consider the possibilities.

[Milestone 2] Interim Report

Annotated outline

- Headings and subheadings showing the major divisions of the literature review, methods, and expected results from the application of each method as well as overall.
- A brief summary of what will go in each subsection, included the references to be discussed in each subsection of the literature review.
- The major references you will rely on.

[Milestone 3] Final project paper & presentation

Please prepare a 10-minute talk describing your proposal.

Your talk should cover the following, although you need not use this structure. Organize your talk in the best way for your particular study. Slide number is required.

- Motivation and research question
 - Why is this an important question to ask?
 - Provide just enough background to allow your listener to understand the question and its significance.
- What kind of study you will do (e.g., interview, experiment, case study, etc.)
 - Why this is the best choice of methods for your question ?
 - Describe the basics of your study design (e.g., for an experiment, what kinds of participants, how many conditions, the basics of the experimental procedure)
 - How will you collect and analyze the data?
- What kinds of results you anticipate and what their importance will be
 - New knowledge
 - Practical impact

You will have *exactly* 10 minutes for this, and no questions (except for clarification) are permitted during your 10 minutes. You will have 3 min for Q&A. Everyone should try to ask questions and provide comments as part of their participation.

Commenting on other teams' projects

Everyone will fill out a comment slip, including the name of the commenter and speaker. *Your comments on other people's talk will be part of your grade, so be thoughtful, helpful, and respectful.*

Comment slips will have space for comments on the following

- In your own words, what is the purpose of the study?
- What is the research question and why is it important?
- Is the choice of method appropriate? Why? What other method could have been used?
- What is the contribution of the study likely to be, assuming results are as expected?
- What is the single change that would have most improved the talk?