
Evaluating the Impact of Generative AI on Software Engineering Education

Research overview

PRESENTED BY

Daniel Otten

IRB

W&M IRB-2026-306 (exempt)

Why this study

**Generative AI is changing how software is built and taught.
Nobody yet understands how.**

- Instructors have no good evidence for how to adjust their courses
- We want to measure how AI use affects learning outcomes
- Understanding how you actually work with these tools is how better courses get designed

How the study is designed

No limits on AI use

The study does not restrict which tools you use, or how much.

Your coursework does not change

Participating adds no assignments, and removes none.

Built for flexibility

The course keeps its educational value while leaving room for how you choose to work.

Second iteration

Already adjusted from a previous run of this study.

Consent and participation

Participation is voluntary and requires your signed consent.

An electronic consent form goes out after this class.

Either way, your coursework and grading are identical. There is no academic penalty.

Project assignment

Matching works the same for everyone. No team is guaranteed its first choice, participant or not.

Students who decline are grouped on one team, so its data can be cleanly excluded.

What we collect

Surveys

One at the start of the semester, one at the end

Repository activity

Your team's GitHub history — commits, issues, PRs, reviews

Reflections

Written and survey reflections on your AI use

AI logs

Project-related use only

Team communication

Your team's Zulip channels

AI logs

Accurate and honest reporting is critical.

AI use unrelated to coursework is not required and not wanted.

Logs are pushed to your team repository and referenced from the issue they belong to.

Each log contains

- Prompts and responses
- The tool used, plus metadata
- Any attachments used

Exporting your logs

SpecStory

Captures agentic sessions in the editor.

Claude Exporter

Chrome extension for browser conversations.
Copying the text manually also works.

Both links will be on the course site.

```
<!-- Generated by SpecStory, Markdown v2.1.0 -->

# 2026-07-16 03:43:14Z

<!-- Claude Code Session 26e4b4b6-e333-407e-b523-fa69fd4e9faa (2026-07-16 03:43:14Z) -->

_**User (2026-07-16 03:43:14Z)**_

<ide_opened_file>The user opened the file d:\Video Project Folders\AUTO TIKTOKS\processor.py in the IDE.

processor.py is an attempt to build a tool which automatically turns my editing my scripted YouTube video
on other platforms. gemini-chat.txt is a talk with a different AI model which led me to this point, but w
plenty of it is wrong, and may need to be redesigned from the ground up.

---

_**Agent (claude-sonnet-5 2026-07-16 03:43:26Z)**_

<tool-use data-tool-type="read" data-tool-name="Read"><details>
<summary>Tool use: **Read**</summary>
`d:\Video Project Folders\AUTO TIKTOKS\processor.py`

```text
1 import os
2 import json
3 import subprocess
4 import requests
5 import win32com.client
6 from faster_whisper import WhisperModel
7
8 # --- CONFIGURATION ---
9 INPUT_DIR = "./videos_in" # Drop your video shortcuts (.lnk) or raw videos here
```

# Anonymity, and your questions

## Codes, not names

In any research output you appear as a code — “student 3.5”.

## Restricted access

Only the research team can access raw, non-anonymized data.

## Secure storage

Raw data is stored securely for the duration of the study.

Research team: Denys Poshyvanyk, Oscar Chaparro, Mustahid Hasan, Nathan Wintersgill, Daniel Otten