

1. (10pt) **Name:** _____

INSTRUCTIONS: Carefully read each question, and write the answer in the space provided. If answers to free response questions are written obscurely, zero credit will be awarded. The correct answer to a free response question with a short answer (i.e., one word or phrase) will never contain any significant words used in the question itself (i.e., “crossword rules”). You are permitted to use one 8.5x11 inch sheet of paper (double-sided) containing **hand-written** notes; all other aids (other than your brain) are forbidden. Questions may be brought to the instructor. You are required to turn in your notes with your exam; they will be returned to you with your graded exam.

You have **80 minutes** to complete the exam. You may turn in your exam as soon as you are finished. There are **800 points** available on the exam, and most questions are worth a number of points that is divisible by ten. These point values are a rough guide to how long I think you should spend on each question.

For **TRUE** or **FALSE** and multiple choice questions, circle your answer.

On free response questions only, you will receive **20%** credit for any question which you leave blank (i.e., do not attempt to answer). Do not waste your time or mine by making up an answer if you do not know. (Note though that most questions offer partial credit, so if you know part of the answer, it is almost always better to write something rather than nothing.)

You may rip out any page in this exam. If a page has questions on it, make sure your UCID is on it before you separate it from the rest of your exam.

To get credit for this question, you must:

- Print your name (e.g., “Martin Kellogg”) in the space provided on this page.
- Print your UCID (e.g., “mjk76”) in the top-right of **each** page of the exam with a question on it.

	Writing your name and UCID:	_____ / 10
	Pages 2 and 3:	_____ / 310
	Page 4 and 5:	_____ / 180
Contents (blanks for graders only):	Page 6 and 7:	_____ / 300
	Extra Credit:	_____ / 0
<hr/>		
	Total:	_____ / 800

I. Multiple Choice and Very Short Answer (150pts). In the following section, either circle your answer (possible answers appear in **bold**) or write a very short (one word or one phrase) answer in the space provided. No partial credit is possible in this section.

2. (20pt) The algorithm used by a modern, distributed version control system like Git for detecting conflicts between different changes can have:
- A false positives
 - B false negatives
 - C both false positives and false negatives
 - D neither false positives nor false negatives
3. (20pt) As a test quality metric, _____ is difficult to interpret in isolation, in part, because checking for equivalence between programs is undecidable.
4. (20pt) Many students' frustrations during IP0 were caused by Covey.Town's build not being _____.
5. Consider the following situations where we might apply delta debugging. Recall that delta debugging relies on three assumptions about the *Interesting* function: it must be monotonic, unambiguous, and consistent. For each situation below, circle the assumption(s) that are **violated** in that situation.
- (a) (30pt) A video game's rendering engine has a bug that causes visual artifacts. The truth, which the developers don't know, is that there are two independent causes: 1) Having "Shadows" set to "High" AND "Reflections" set to "Ultra", OR 2) Having "Anti-Aliasing" set to "MSAA" AND "Texture Filtering" set to "Anisotropic 16x". A developer tries to find the minimal set of settings (from a large list) that triggers the bug.
- Monotonic** **Unambiguous** **Consistent**
- (b) (30pt) A static webpage renders incorrectly (e.g., a button is misaligned). The developer suspects a rule in a single, large .css file. They use delta debugging to find the minimal set of CSS rules from that file that still causes the misalignment. The "Interesting" test is an automated visual regression test that takes a screenshot and compares it to a "correct" version.
- Monotonic** **Unambiguous** **Consistent**
- (c) (30pt) An application fails to start if its .xml configuration file is malformed. A developer has a large, failing config file and wants to find the minimal set of lines that causes the failure. The "Interesting" function is a script that prints the selected lines into a new file and feeds it to the application.
- Monotonic** **Unambiguous** **Consistent**

- A.** beta test **B.** Rice's theorem **C.** waterfall **D.** A/B test **E.** regression test
F. behavioral **G.** object-oriented **H.** functional **I.** logging **J.** mocking
K. agile **L.** imperative **M.** quality **N.** fuzzing **O.** over-engineered

II. Matching (160pts). This section contains a collection of terms discussed in class in an “Answer Bank” (choices **A.** through **O.**). Each question in this section describes a situation associated with an answer in the Answer Bank. Write the letter of the term in the Answer Bank that best describes each situation. Each answer in the Answer Bank will be used at most once.

6. (20pt) _____ After William fixes a bug, he ensures that that specific bug won't reoccur.
7. (20pt) _____ Because destructive updates are difficult to reason about, Phil chooses a programming language that yields new similar states over time.
8. (20pt) _____ Harry carefully gathers requirements of this type from his customers, because he knows that to deliver a working system on time he will have to compromise on some of them.
9. (20pt) _____ Chris knows that perfect static analysis is undecidable, so he decides to use a sound but imprecise analysis, because he is working in a safety-critical domain.
10. (20pt) _____ While working on a web application, Woodrow deploys a change to a small percentage of traffic to test whether it impacts user behavior.
11. (20pt) _____ To debug a failure in his large distributed system, Jon relies on information about the program's internal state during an outage.
12. (20pt) _____ George's team strives to always have a working prototype of their system.
13. (20pt) _____ When he interviews candidates, Richard asks them questions that are designed to test their communication ability, not just their technical skills.

III. Short answer (150pts). Answer the questions in this section in at most three sentences.

14. (50pt) Suppose that your team has decided to rewrite a Java codebase into Rust to avoid garbage-collection stalls. Describe an approach, technique, or tool that we discussed in class that you could use to increase your confidence that the rewrite will succeed.
15. (50pt) Suppose that you are a software engineer at well-funded startup. Your boss tells you that in one week, you'll have a big demo for a potential customer. This customer would be your biggest yet, by a fair margin, so it's important that the demo goes well. Fortunately, however, you'll be the one operating your system during the demo. One of your coworkers suggests that you should use a sound static analysis to *prove* that the demo will succeed. Do you agree with this coworker? Why or why not?
16. (50pt) Support or refute the following claim: if you run your CI builds on one server in a specific timezone (e.g., all your CI servers are in California), it is safe to use a build system that incrementalizes using timestamps.

- A. *No Silver Bullet*
- B. *Syntax, Predicates, Idioms—What Really Affects Code Complexity?*
- C. *Relating Reading, Visualization, and Coding for New Programmers: A Neuroimaging Study*
- D. *An Experimental Evaluation of Continuous Testing During Development*
- E. *Taming Google-Scale Continuous Testing*
- F. *The Oracle Problem in Software Testing: A Survey*
- G. *Purposes, Concepts, Misfits, and a Redesign of git*
- H. *Variability and Reproducibility in Soft. Eng.: A Study of 4 Companies that Developed the Same System*
- I. *Hiring is Broken: What Do Developers Say About Technical Interviews?*
- J. *The Daikon System for Dynamic Detection of Likely Invariants*
- K. *Introduction to TLA*
- L. *Expectations, Outcomes, and Challenges Of Modern Code Review*
- M. *Hints on Programming Language Design*
- N. *Build Systems à la Carte*
- O. *Notes on Program Analysis*
- P. *Designing the WhyLine: A Debugging Interface for Asking Questions about Program Behavior*
- Q. *Locating Causes of Program Failures*

IV. “Your Choice” Reading Quiz

17. (30pt) The Answer Bank on this page lists the “Your Choice” readings so far in this semester.

- (a) Write the letter of the “Your Choice” reading for which you’re answering this question: _____.
- (b) Describe a software engineering situation in which you might use something you learned from the chosen “Your Choice” reading. To get credit for this question, you must both describe such a situation and clearly explain what you learned. If the “thing you learned” is so simple that one could understand it from the title of the reading alone, you won’t get credit.

V. Document-based Questions (300pts). All questions in this section refer to **Document A**, which we recommend you tear out of the exam and read before trying to answer the questions in this section. **Document A** is a blog post from an engineer about why he doesn't always do code reviews for pull requests with a lot of AI-generated code, in the open-source project that he works on in his spare time. We have redacted some parts of the document. Your job is to use your knowledge of what makes a good code review to fill in the blanks.

For five of the six "Reasons for no CR" listed in the document, argue from your understanding of the benefits of code review why or why not that reason is a valid one for not providing a code review, if the code in question is AI-generated. Fill in the number of the associated reason in each question; you may choose which you'd like to skip. For clarity, the reasons are:

- Reason 1: "Code deletion would improve the MR a lot."
- Reason 2: "You don't know the basics of the language you submitted."
- Reason 3: "Documentation spam."
- Reason 4: "Blatantly inconsistent."
- Reason 5: "Edge case overload."
- Reason 6: "Adding pointless or deprecated dependencies without knowing why."

18. (60pt) Support or refute the following claim: reason _____ is a valid reason not to review an AI-generated code review.

19. (60pt) Support or refute the following claim: reason _____ is a valid reason not to review an AI-generated code review.

20. (60pt) Support or refute the following claim: reason _____ is a valid reason not to review an AI-generated code review.

21. (60pt) Support or refute the following claim: reason _____ is a valid reason not to review an AI-generated code review.

22. (60pt) Support or refute the following claim: reason _____ is a valid reason not to review an AI-generated code review.

VI. Extra Credit. Questions in this section do not count towards the denominator of the exam score.

23. (10pt) In section II (Matching), there is a theme to the names used in the situation descriptions. What is the theme?
24. (a) (10pt) Why was Canvas unavailable on October 20th?
- (b) (20pt) Describe the root cause of Canvas' unavailability on October 20th.
25. (10pt) Describe something interesting that you learned in this class that wasn't tested on this exam.
26. (10pt) Describe something interesting that you learned from a "Your Choice" reading other than the one that you used to answer question 17. Be sure to include the letter corresponding to the reading in the answer bank for question 17 prominently in your answer.

This page intentionally left blank (you may use it as scratch paper, and the proctor will have more scratch paper at the front if you need more).

Document A:

Why I'm declining your AI generated MR

Published on 2025-08-26 by Stuart Spence

Sometimes a merge request (MR) doesn't merit a code review (CR) because AI was used in a bad way that harms the team or the project. For example:

- (a) Code deletion would improve the MR a lot.
- (b) You don't know the basics of the language you submitted.
- (c) Documentation spam.
- (d) Blatantly inconsistent.
- (e) Edge case overload.
- (f) Adding pointless or deprecated dependencies without knowing why.

If I decline your AI code MR with no further comments and send you this page then some of these conditions were met.

Despite some recent research and discussions on it I know that AI can be helpful in writing code. However AI misuse is also a new phenomenon and we need guidelines to help identify it. This page was written in 2025 and I expect the tools and guidelines to evolve.

Vocabulary

- **Merge Request (MR)**: when a programmer submits proposed changes to a project in a structured way. This makes it easy for anyone to see the differences and review the changes. Also called a **Pull Request** like on GitHub.
- **Code Review (CR)**: when another programmer reviews a MR, provides feedback or improvements, and approves or rejects the changes.

Why me?

I feel like I'm in a good position to write about this because:

- (a) I'm a senior computer scientist for AI and cloud. I've been a technical supervisor to about 20 students and juniors.
- (b) I have degrees and work experience in both computer science and education.
- (c) I know code and AI. My AI project has a million installs and monthly income.

- (d) I spend a lot of personal time enjoying, exploring, and discussing AI news and breakthroughs.¹
- (e) I don't need a job and I'm not selling anything. I'm not an investor or CEO shilling AI slop and I don't get ad revenue for flaming AI.

Why do code reviews?

There's thousands of opinions and articles. Here's a good one from Google². Instead of rehashing that, let's focus on what AI misuse threatens. With good code reviews:

- (a) [REDACTED]
- (b) [REDACTED]
- (c) [REDACTED]
- (d) [REDACTED]
- (e) [REDACTED]
- (f) [REDACTED]
- (g) [REDACTED]

Reasons for no CR

1) Code deletion would improve the MR a lot.

Can code be trivially deleted?

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

2) You don't know the basics of the language you submitted.

[REDACTED]

[REDACTED]

[REDACTED]

¹The author linked to evidence of themselves doing these things in the real blog post, but in this exam you can take them at their word.

²In the actual blog post, this links to the reading that you were assigned for the "Code Review" lecture on October 13th, which is Google's code review guidelines

3) Documentation spam.

One example I've seen is two nearly identical copies of documentation in two different formats.

[REDACTED]

4) Blatantly inconsistent.

Common examples I see are using new frameworks or styles for logging and unit tests.

[REDACTED]

5) Edge case overload.

[REDACTED]

AI slop may desperately catch an exhaustive list of all exceptions to "handle all cases". But the AI isn't handling the cases. It's just suppressing the valid exception or writing a non-standard error message.

6) Adding pointless or deprecated dependencies without knowing why.

[REDACTED]

Maybe this is fine.

None of these are hard rules. I'm more inclined to accept an AI generated MR or give a CR to one if:

- The code is temporary or a one-shot analysis with no long term maintenance requirements. If it works it works!
- The MR includes an explanation of why AI was used, how much, why, and what extra steps the author took to validate it.
- This is an edge feature and not a core component.

Challenging Times

As a team lead, teacher, and I think nice guy I'm struggling with how to confront teammates when I feel their MR harms them, or the team, or the project. Why did they submit AI code? Was it a smart decision or just laziness? Do I harshly confront them and call it AI slop or do something else?

It's not always clear to me when it's a good use of AI that I should support with a full CR, or when it's a bad use of AI that I need to confront by rejecting it entirely. For me, just writing this page has helped.

Nobody has years of experience understanding how AI slop impacts technical debt or learning. If software development is changing in a good way then team leads need to change with it. If software development is changing in a bad way then we need to resist.