

CS 891: Scalable Microservices: Overview (Part 2)

Douglas C. Schmidt

d.schmidt@vanderbilt.edu

www.dre.vanderbilt.edu/~schmidt



Professor of Computer Science

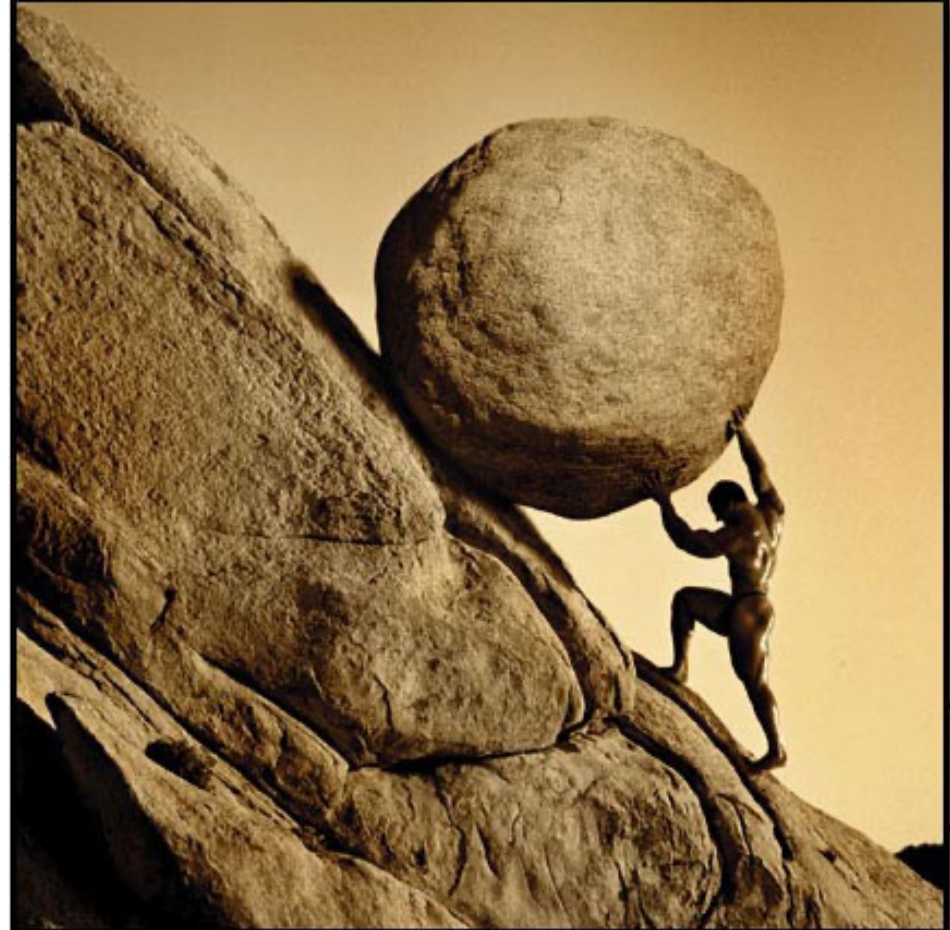
**Institute for Software
Integrated Systems**

**Vanderbilt University
Nashville, Tennessee, USA**



Learning Objectives in this Lesson

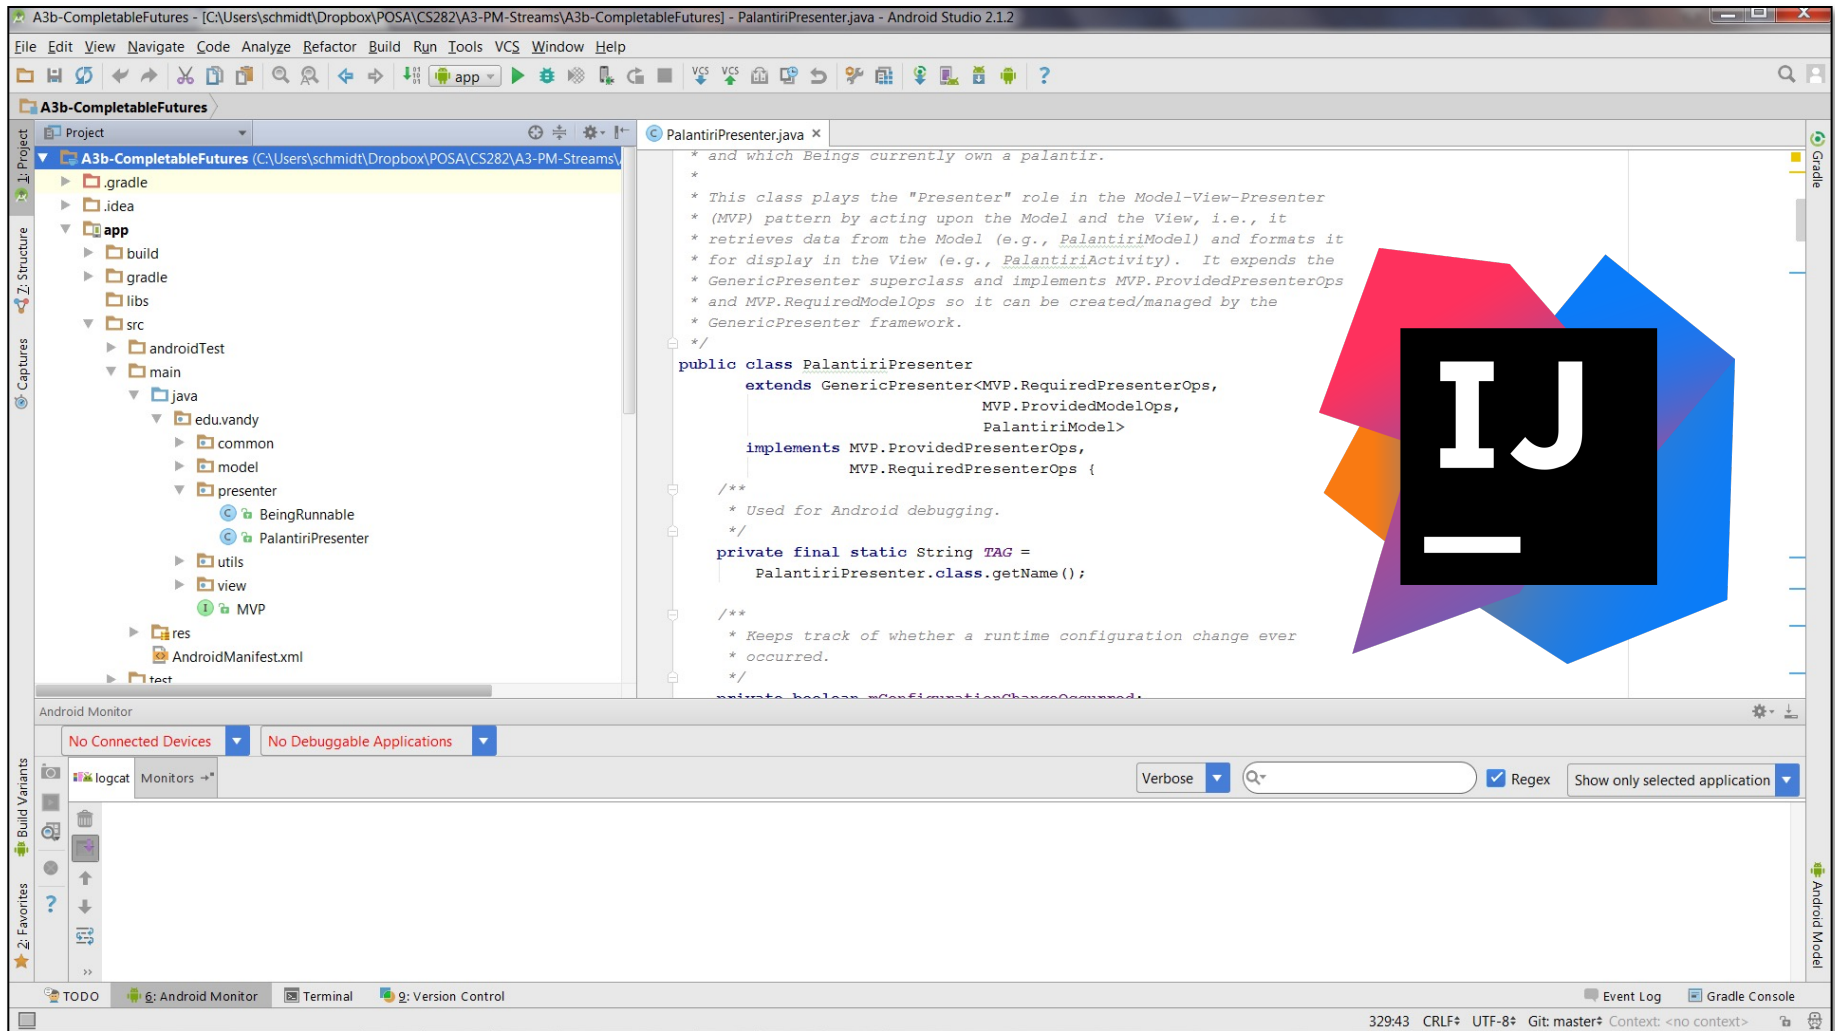
- Understand the course topics & logistics
 - Course philosophy
 - Course contents
 - Structure of the lecture material
- Overview of the assignments & assessments



Overview of Assignments & Assessments

Overview of Assignments & Assessments

- Programming assignments are written in modern Java using IntelliJ



You can use any IDE, but your final submission *must* build & run with the latest IntelliJ & Java 19

Overview of Assignments & Assessments

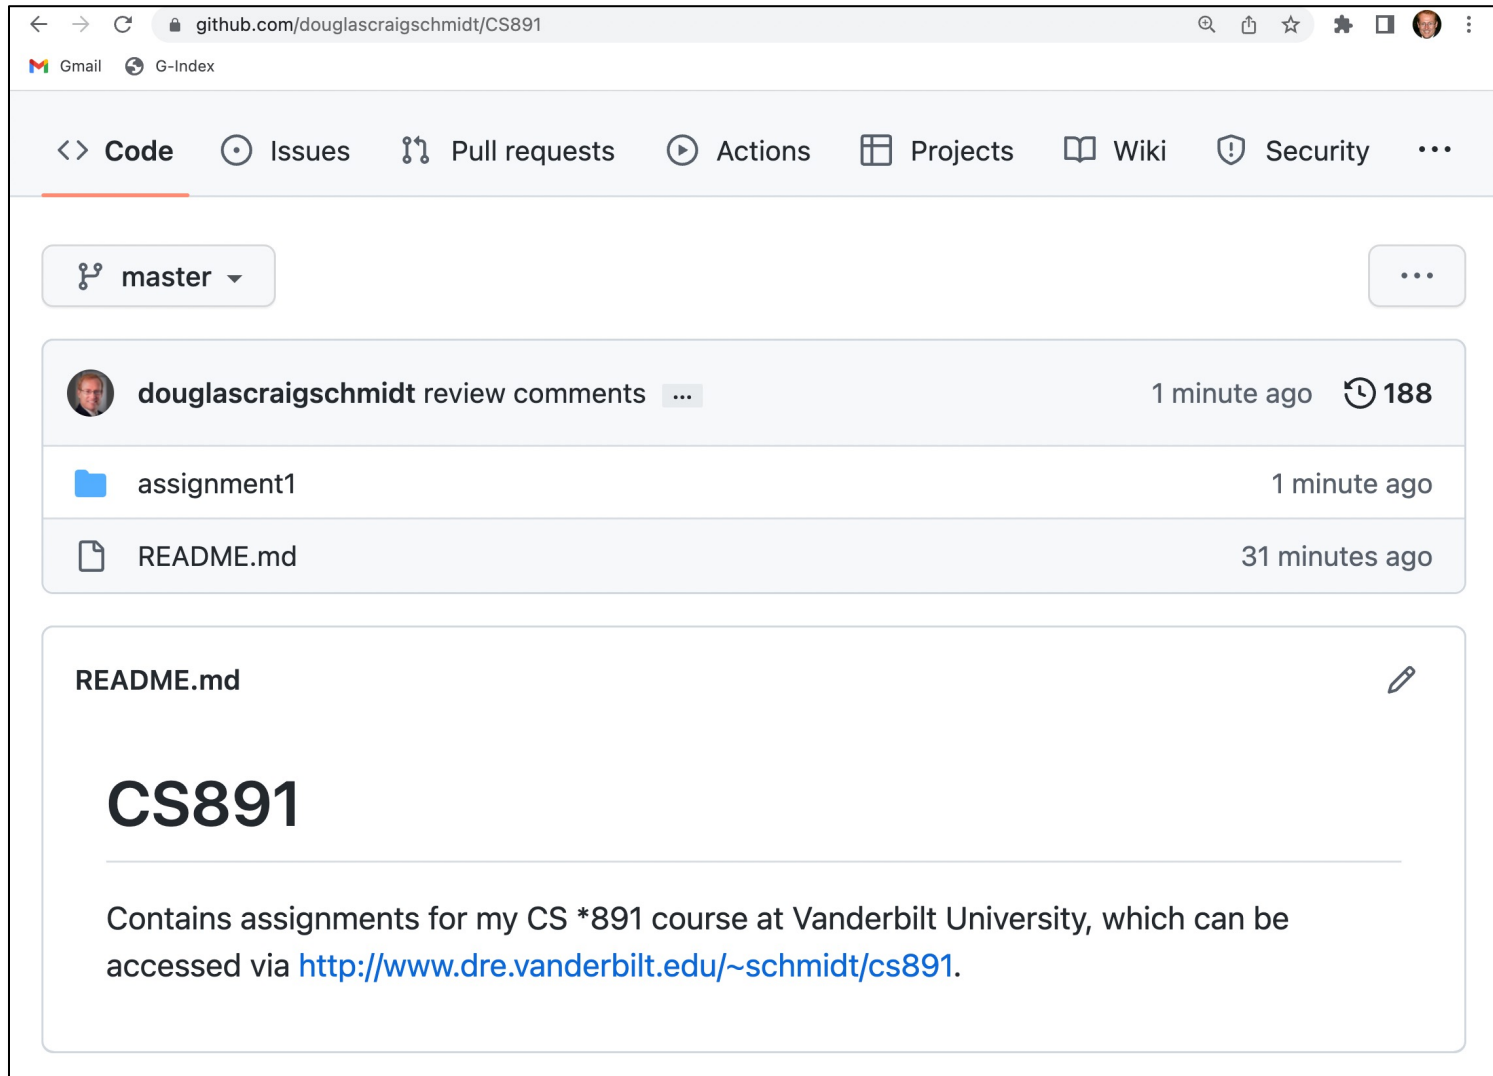
- Programming assignments are written in modern Java using IntelliJ
 - The Java 19 runtime environment (JRE) can be downloaded via IntelliJ



See github.com/douglasraigschmidt/CS891/wiki/Installing-Software

Overview of Assignments & Assessments

- All source code for assignments & examples available at GitHub



The screenshot shows a GitHub repository page for 'dougascraigschmidt/CS891'. The browser address bar shows the URL 'github.com/dougascraigschmidt/CS891'. The repository navigation bar includes 'Code', 'Issues', 'Pull requests', 'Actions', 'Projects', 'Wiki', and 'Security'. The current branch is 'master'. The file list shows 'assignment1' (1 minute ago) and 'README.md' (31 minutes ago). The 'README.md' content is displayed below, featuring the title 'CS891' and a description: 'Contains assignments for my CS *891 course at Vanderbilt University, which can be accessed via <http://www.dre.vanderbilt.edu/~schmidt/cs891>.'

Go to GitHub at github.com/dougascraigschmidt/CS891

Overview of Assignments & Assessments

- All source code for assignments & examples available at GitHub
- You will need to learn how to use GitLab et al.

A screenshot of the GitLab website banner. The background is a dark purple. In the top left corner is the GitLab logo (a cat face) and the text "GitLab". In the top right corner is a white hamburger menu icon. The main text reads "Open source software to collaborate on code". Below this is a paragraph of text describing GitLab's features and capabilities. At the bottom are two buttons: "GitLab Community Edition" (light blue) and "Get a subscription" (green).

GitLab offers git repository management, code reviews, issue tracking, activity feeds and wikis. Enterprises install GitLab on-premise and connect it with LDAP and Active Directory servers for secure authentication and authorization. A single GitLab server can handle more than 25,000 users but it is also possible to create a high availability setup with multiple active servers.

Do you want more from your GitLab installation? A subscription bundles the Enterprise Edition with support from the GitLab team. The Enterprise Edition allows you to sync LDAP groups, control pushes via git hooks, integrate better with Jenkins and Jira, and to run MySQL and forward logs when using our Omnibus package. Our service engineers will help you keep your server running smoothly.

[GitLab Community Edition](#) [Get a subscription](#)

We'll discuss how to setup GitLab shortly

Overview of Assignments & Assessments

- All source code for assignments & exam
 - You will need to learn how to use GitLab et al.
 - Be prepared to update your repositories occasionally



"If you don't like change, you're going to like irrelevance even less."



Overview of Assignments & Assessments

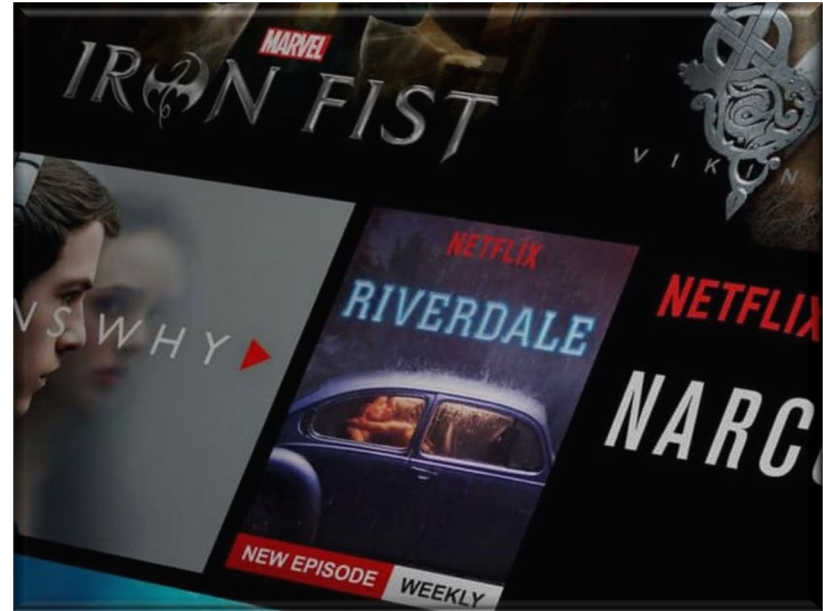
- Assignments will provide a range of experience with modern Java concurrent & parallel microservices



Go to GitHub at github.com/douglasraigschmidt/CS891

Overview of Assignments & Assessments

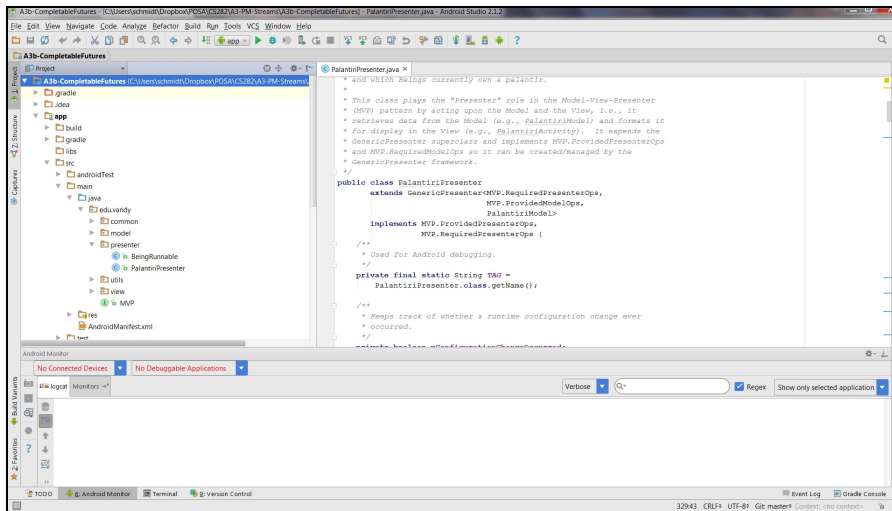
- Assignments will provide a range of experience with modern Java concurrent & parallel microservices
 - Implement a microservice-based movie recommendation system on Spring using modern Java features, e.g.
 - Java lambda expressions, method references, & functional interfaces
 - Java sequential streams
 - Java structured concurrency
 - Java reactive streams
 - Spring WebMVC & WebFlux



The topics covered by the assignments may change during the semester

Overview of Assignments & Assessments

- Assignment assessments will be done via reviews by course staff & automated tools



Overview of Assignments & Assessments

- Assignment assessments will be done via reviews by course staff & automated tools

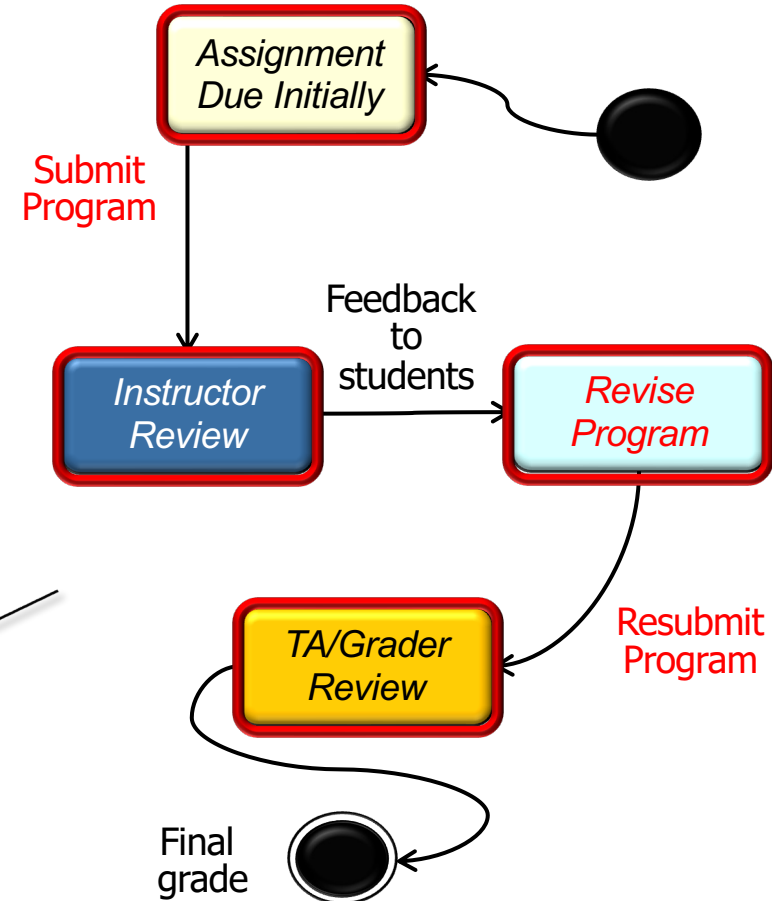


Assignments must be submitted correctly & on time or they will not be reviewed (no exceptions)

See github.com/douglasraigschmidt/CS891/wiki/CS-891-FAQ

Overview of Assignments & Assessments

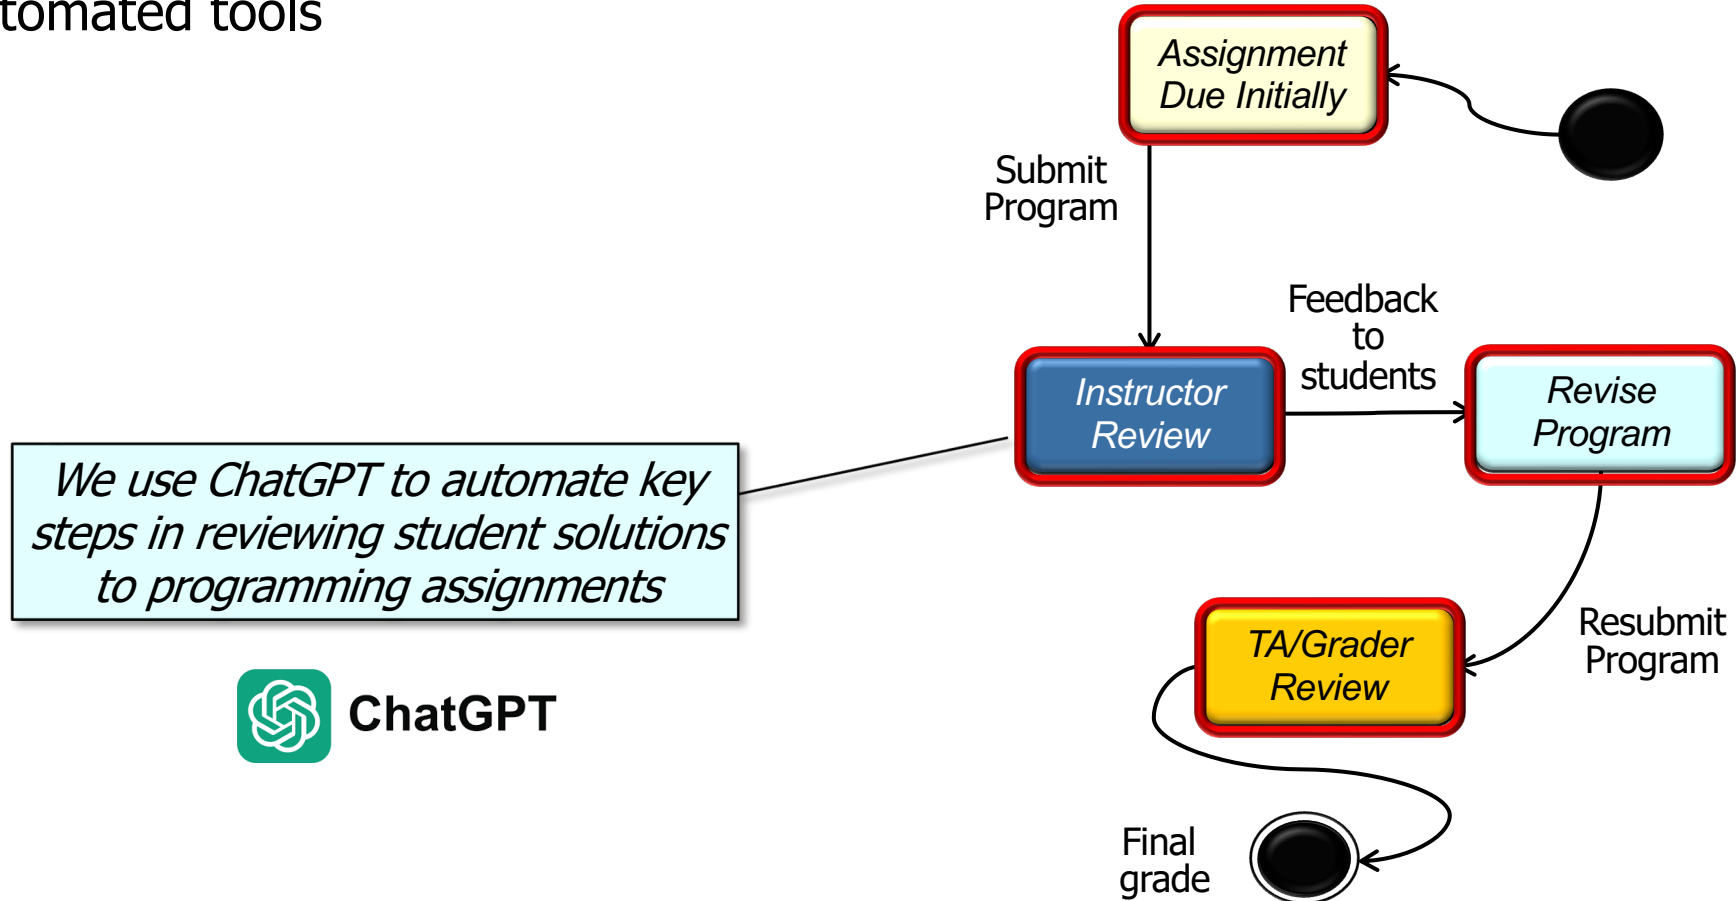
- Assignment assessments will be done via reviews by course staff & automated tools



Your initial submission must compile & be largely complete or you won't get a review

Overview of Assignments & Assessments

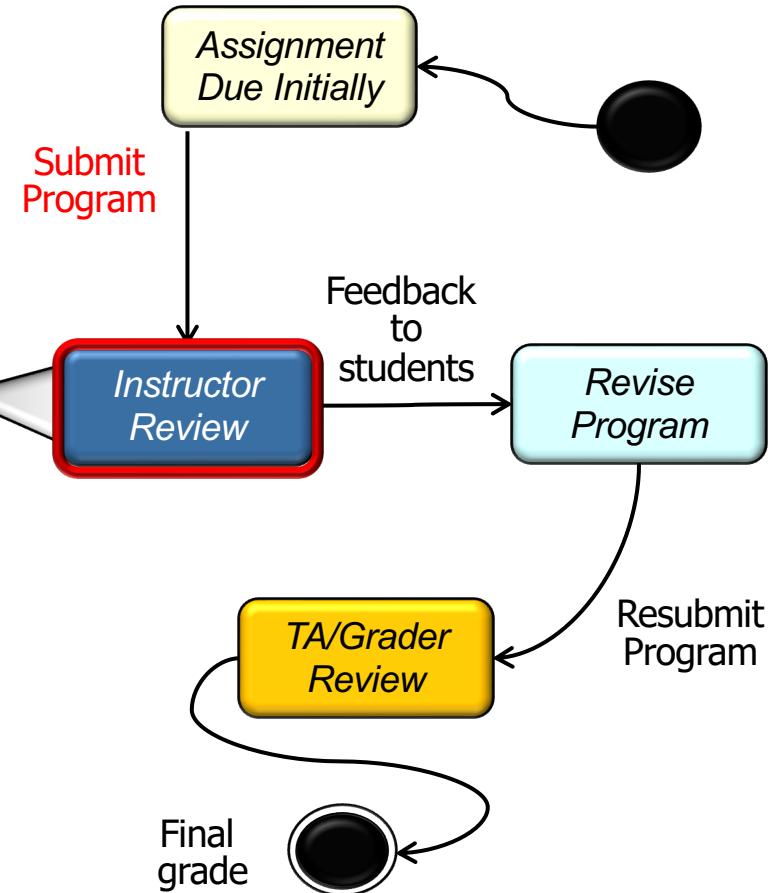
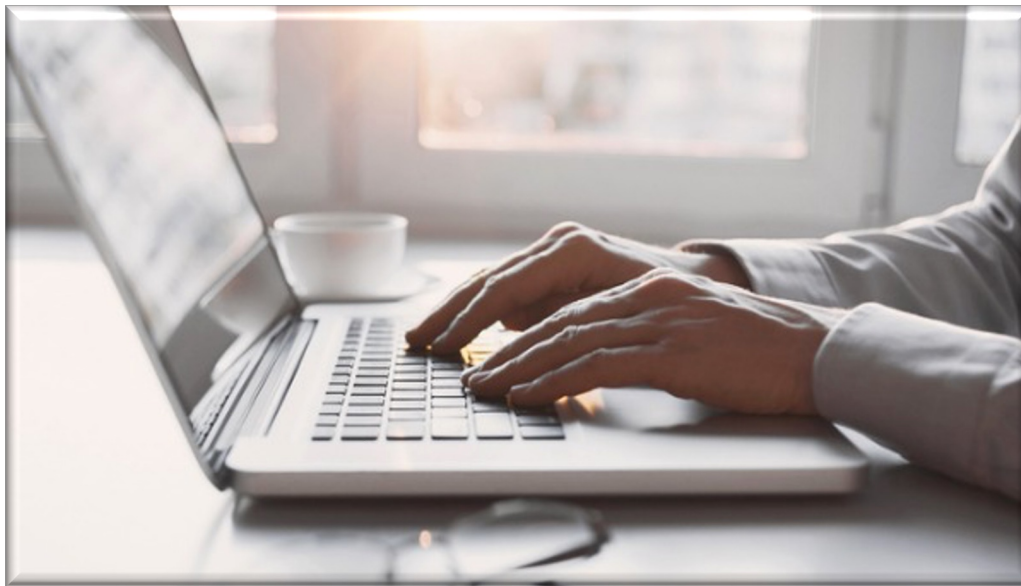
- Assignment assessments will be done via reviews by course staff & automated tools



This is a form of "discrepancy analysis" that can leverage generative AI!

Overview of Assignments & Assessments

- Assignment assessments will be done via reviews by course staff & automated tools



Overview of Assignments & Assessments

- Assignment assessments will be done via reviews by course staff & automated tools

Historically, I couldn't leverage knowledge of these FMMs between course offerings..

```
protected ForkJoinTask<Integer> makeProcessImageTask(String url) {
    // @ Replace this try/catch block with a call to
    // ExceptionUtils.rethrowFunction().
    Function<String, URL> urlConverter = urlStr -> {
        try {
            return new URL(urlStr);
        } catch (Exception e) {
            throw new RuntimeException("...");
        }
    };

    // @ This code is WAY too complicated. You just need to make
    // an instance of ProcessImageTask() and return it..
    Callable<Integer> callable = () -> {
        var imageUrl = urlConverter.apply(url);
        if(imageUrl == null) { return null; }
        var processImageTask = new ProcessImageTask(imageUrl);
        if(processImageTask == null) { return null; }
        return processImageTask.join();
    };
    return ForkJoinTask.adapt(callable);
}

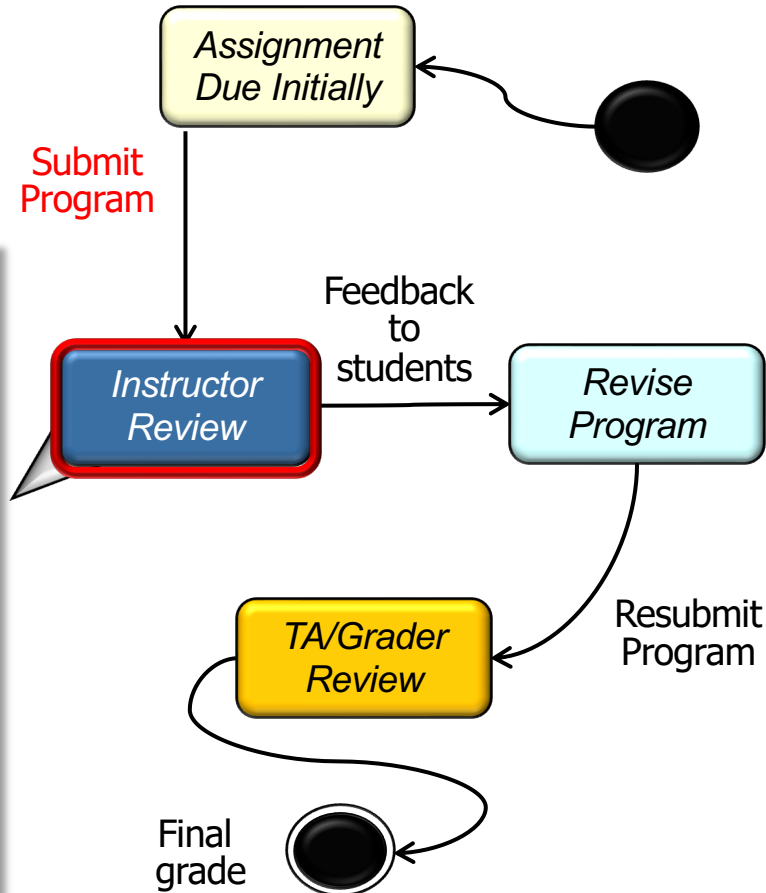
protected ForkJoinTask<Integer> makeProcessImageTask(String url) {
    Function<String, URL> urlConverter = ExceptionUtils
    // @ This is too complicated - just use URL:new:
    .rethrowFunction(urlStr -> new URL(urlStr));
    ...
}

public class URLCrawlerTask ...
// @ You don't need to have a 'count' field - instead, handle
// this via local variables, which are thread-safe.
private int count;

URLCrawlerTask(String pageUri, int depth) {
    ...
    // @ Don't allocate a new ForkJoinPool - you need to just
    // use the common fork-join pool.
    forkJoinPool = new ForkJoinPool();

    // @ Don't do this, as per the comment above.
    count = 0;
    ...
}

--:-- A1-FMM 12% L47 (Fundamental) Mon Sep 11 10:31PM 2.04
```



See youtu.be/2P5CGYSXPKg for an example of a "frequently made mistakes" (FMM) video

Overview of Assignments & Assessments

- Assignment assessments will be done via reviews by course staff & automated tools

```
"ForkJoinExCrawler": {
  "criteria": [
    {
      "title": "Use of Common Fork-Join Pool",
      "description": "The student should use the common ForkJoinPool provided by the ForkJoinPool.commonPool() method, rather than creating a new instance of ForkJoinPool. This is to ensure resource optimization and conform to Java best practices.",
      "correct_example": "ForkJoinPool.commonPool().invoke(...);",
      "incorrect_example": "ForkJoinPool mCommon = new ForkJoinPool();"
    },
    {
      "title": "Method Invocation (invoke() vs join())",
      "description": "The student should use the invoke() method for starting the ForkJoinTask, as opposed to the join() method. The invoke() method is preferred as it commences the task and waits for the result, making it more readable for this scenario.",
      "correct_example": "return forks.invoke();",
      "incorrect_example": "return forks.join();"
    },
    {
      "title": "Unnecessary Use of Streams",
      "description": "The student should avoid using streams where they are not necessary, to ensure code simplicity and maintainability.",
      "incorrect_example": "return Stream.<String>(pageUri)...;"
    },
    {
      "title": "Avoidance of Filtering",
      "description": "The student should avoid using filter() to filter out (DOS)..."
    }
  ]
}
```



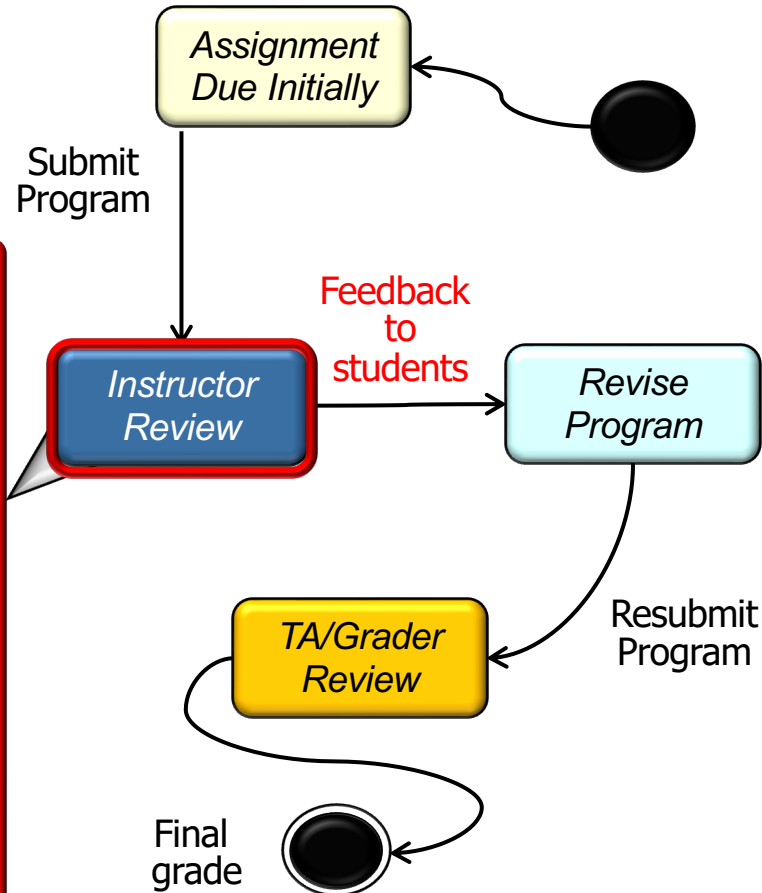
```
protected ForkJoinTask<Integer> makeProcessImageTask(String url) {
  // @ Replace this try/catch block with a call to
  // ExceptionUtils.rethrowFunction().
  Function<String, URL> urlConverter = urlStr -> {
    try {
      return new URL(urlStr);
    } catch (Exception e) {
      throw new RuntimeException("...");
    }
  };
  // @ This code is WAY too complicated. You just need to make
  // an instance of ProcessImageTask() and return it..
  Callable<Integer> callable = () -> {
    var imageUrl = urlConverter.apply(url);
    if(imageUrl == null) { return null; }
    var processImageTask = new ProcessImageTask(imageUrl);
    if(processImageTask == null) { return null; }
    return processImageTask.join();
  };
  return ForkJoinTask.adapt(callable);
}

protected ForkJoinTask<Integer> makeProcessImageTask(String url) {
  Function<String, URL> urlConverter = ExceptionUtils
  // @ This is too complicated - just use URL::new
  .rethrowFunction(urlStr -> new URL(urlStr));
  ...
}

public class URLCrawlerTask ...
// @ You don't need to have a 'count' field - instead, handle
// this via local variables, which are thread-safe.
private int count;

URLCrawlerTask(String pageUri, int depth) {
  // @ Don't allocate a new ForkJoinPool - you need to just
  // use the common fork-join pool.
  forkJoinPool = new ForkJoinPool();

  // @ Don't do this, as per the comment above.
  count = 0;
  ...
}
```



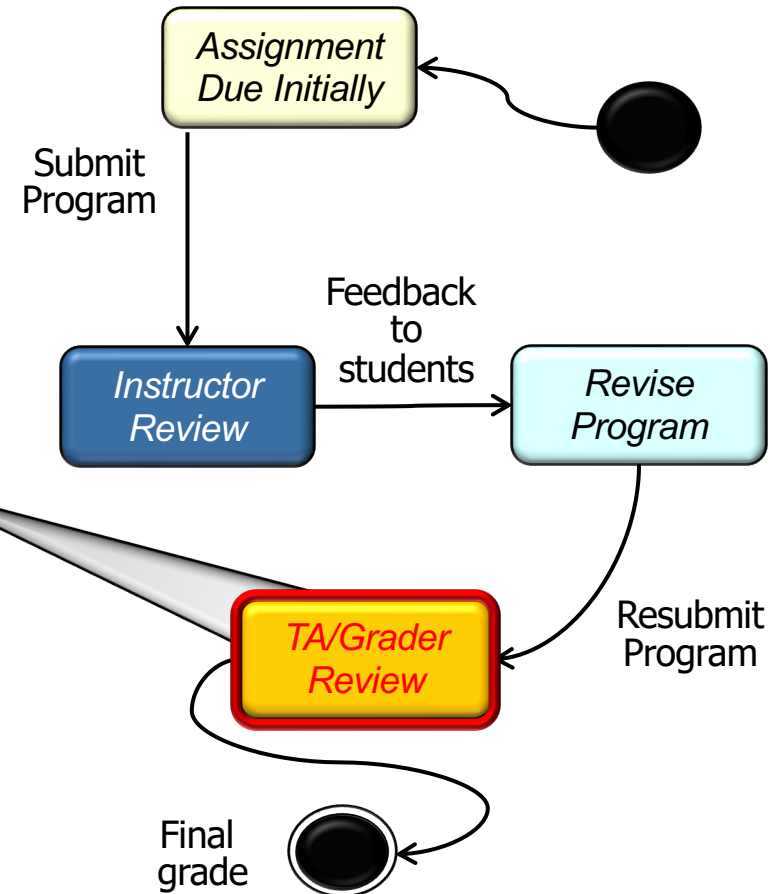
ChatGPT generates an automated grading rubric based on my FMM analysis

Overview of Assignments & Assessments

- Assignment assessments will be done via reviews by course staff & automated tools



My TAs/graders no longer have to read my mind to evaluate student final submissions consistently



Overview of Assignments & Assessments

- Assignment assessments will be done via reviews by course staff & automated tools

```
> grader.py > ...
#for each directory in split_files, read in each java file
import openai
import os
from time import sleep

def grade(code, rubric):
    api_key = "sk-3X79b403IXFtn0mvpHhT3BlbkFJBoXax1A0gvc8uta4xUQ"

    # Initialize the OpenAI API client
    openai.api_key = api_key

    prompt = f"You are a grader for a Parallel Functional Programming course taught in Android (Java). "\
    f"I will give you a json rubric and Android (Java) code. For each item in the rubric, you will first "\
    f"output the function in the student's code that is relevant to that item and then you will output a grade "\
    f"of correct or incorrect. Alternative answers to the correct code are permissible if they have the same "\
    f"functionality and are not following poor style conventions. \nRubric:\n{rubric}\n\nStudent's Code:\n{code}"

    messages = [{"role": "user", "content": prompt}]

    # Make the API call to get the code
    completion = openai.ChatCompletion.create(
        model="gpt-4",
        messages=messages
    )

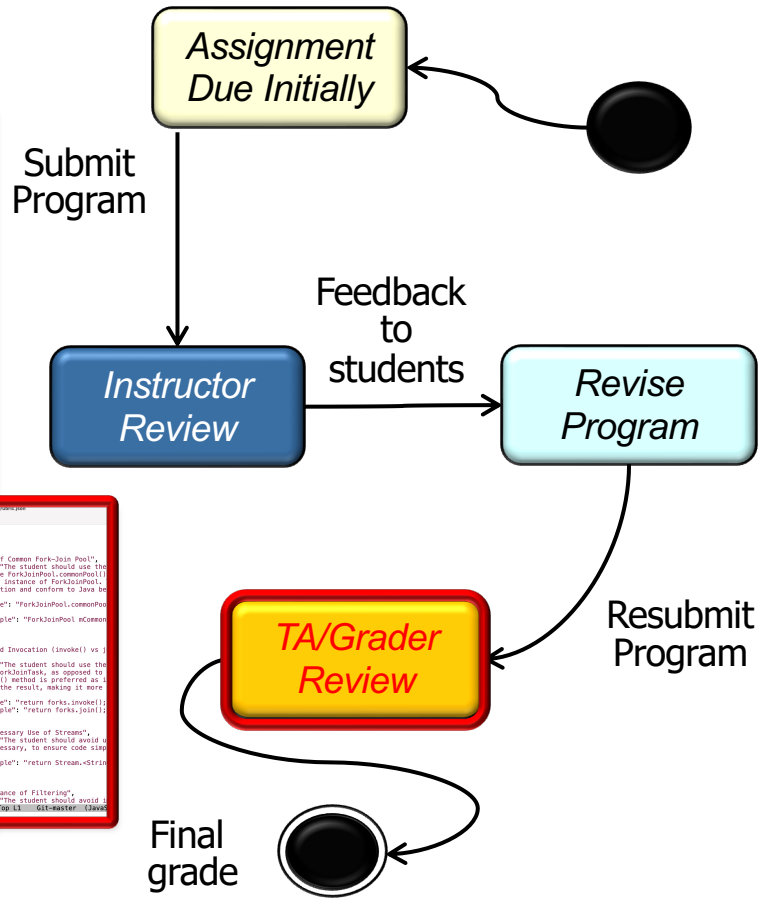
    chat_response = completion.choices[0].message.content
    messages.append({"role": "assistant", "content": chat_response})
    messages.append({"role": "user", "content": "Please summarize everything that the student got incorrect"


```



```
"title": "Use of Common Fork-Join Pool",
"description": "The student should use the
common ForkJoinPool provided by the ForkJoinPool.ComonPool()
method, rather than creating a new instance of ForkJoinPool()
to ensure resource optimization and conform to Java best
practices.",
"correct_example": "ForkJoinPool.commonPool
.invoke(...)",
"incorrect_example": "new ForkJoinPool()
.invoke(...)",
"title": "Method Invocation (invoke) vs j
oin()",
"description": "The student should use the
invoke() method for starting the ForkJoinTask, as opposed to
using the join() method. The invoke() method is preferred as it
denotes the task and waits for the result, making it more
readable for this scenario.",
"correct_example": "return Forks.invoke();",
"incorrect_example": "return Forks.join();",
"title": "Unnecessary Use of Streams",
"description": "The student should avoid us
ing streams where they are not necessary, to ensure code size
and maintainability.",
"correct_example": "return Stream.of(1, 2, 3)
.map(i -> i * 2).collect(Collectors.toList());",
"incorrect_example": "return Stream.of(1, 2, 3)
.map(i -> i * 2).collect(Collectors.toList());",
"title": "Avoidance of Filtering",
"description": "The student should avoid us
ing the filter() method when it is not necessary, as it can
introduce unnecessary overhead.",
"correct_example": "return Stream.of(1, 2, 3)
.map(i -> i * 2).collect(Collectors.toList());",
"incorrect_example": "return Stream.of(1, 2, 3)
.filter(i -> i > 0).map(i -> i * 2).collect(Collectors.toList());"
}
```

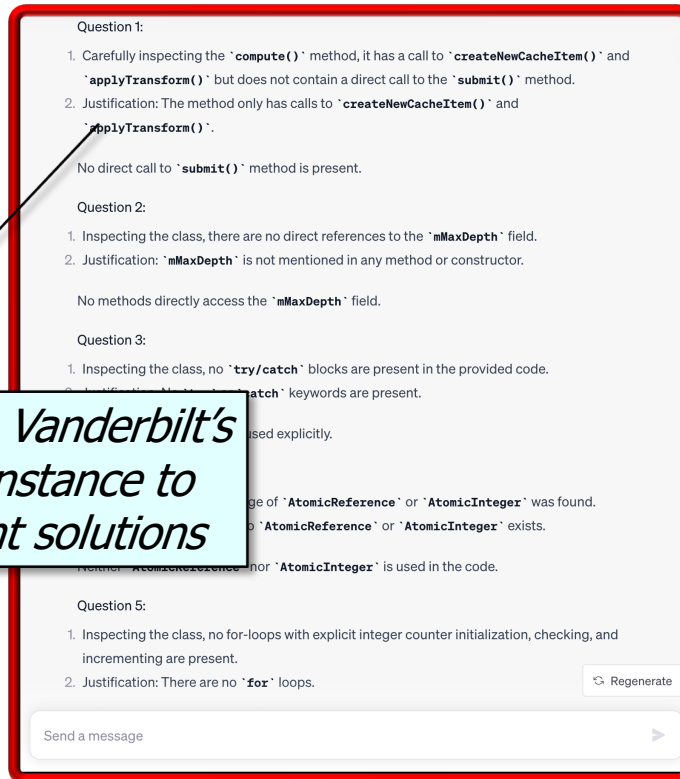
We've trained ChatGPT-4 to analyze student final submissions for frequently made mistakes



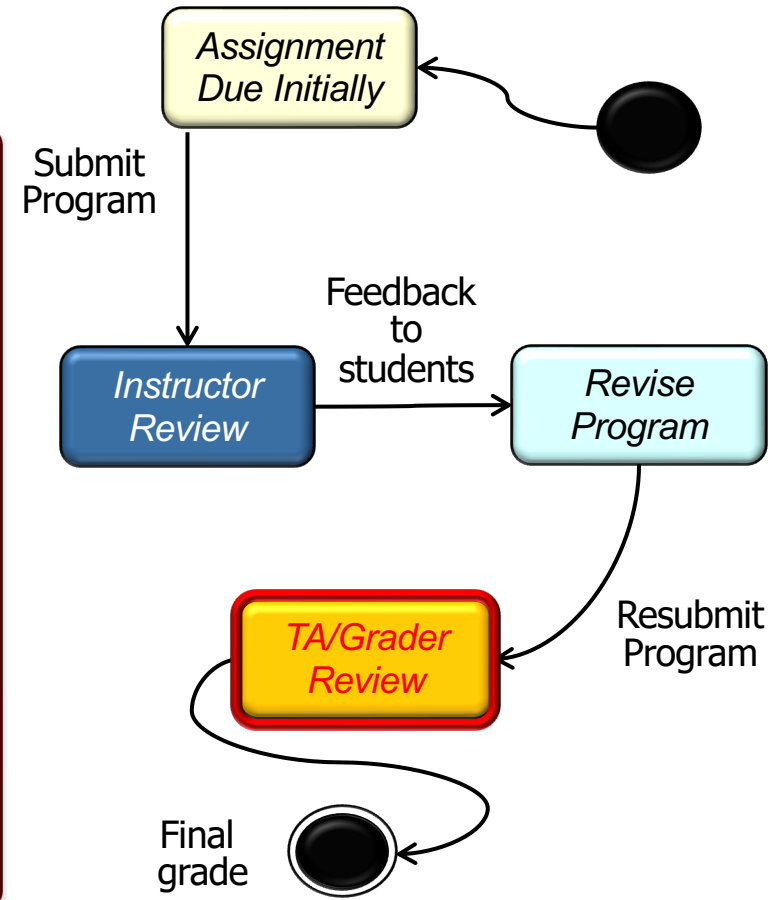
See www.dre.vanderbilt.edu/~schmidt/PDF/GreAIter-paper.pdf

Overview of Assignments & Assessments

- Assignment assessments will be done via reviews by course staff & automated tools



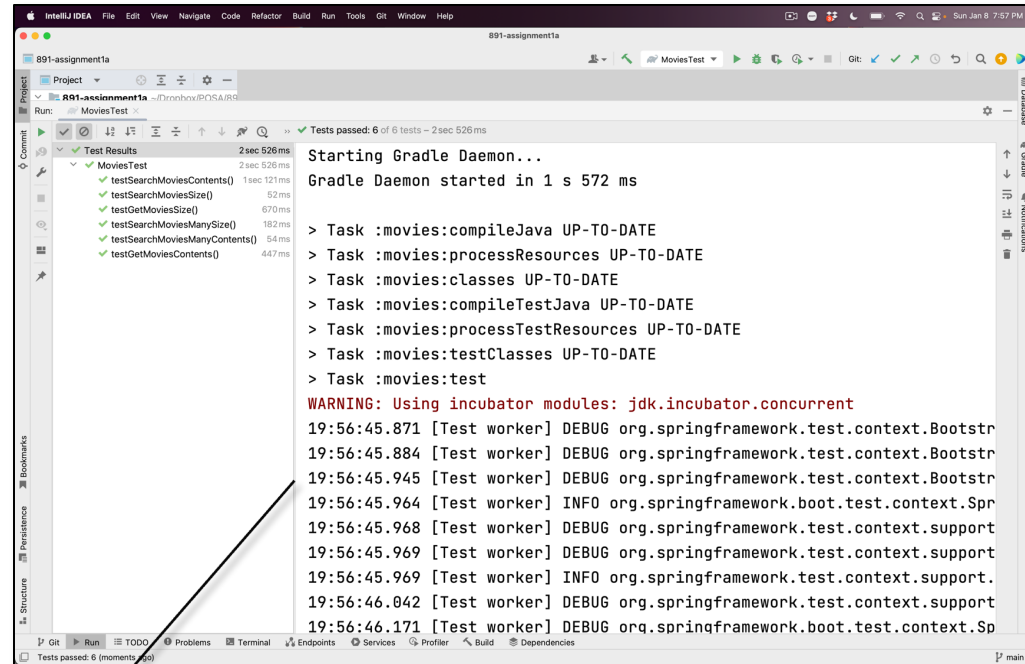
My TAs/graders use Vanderbilt's private ChatGPT instance to assess final student solutions



Issues with false positives are handled efficiently & scalably via augmented intelligence (AI+)

Overview of Assignments & Assessments

- Assignment assessments will be done via reviews by course staff & automated tools



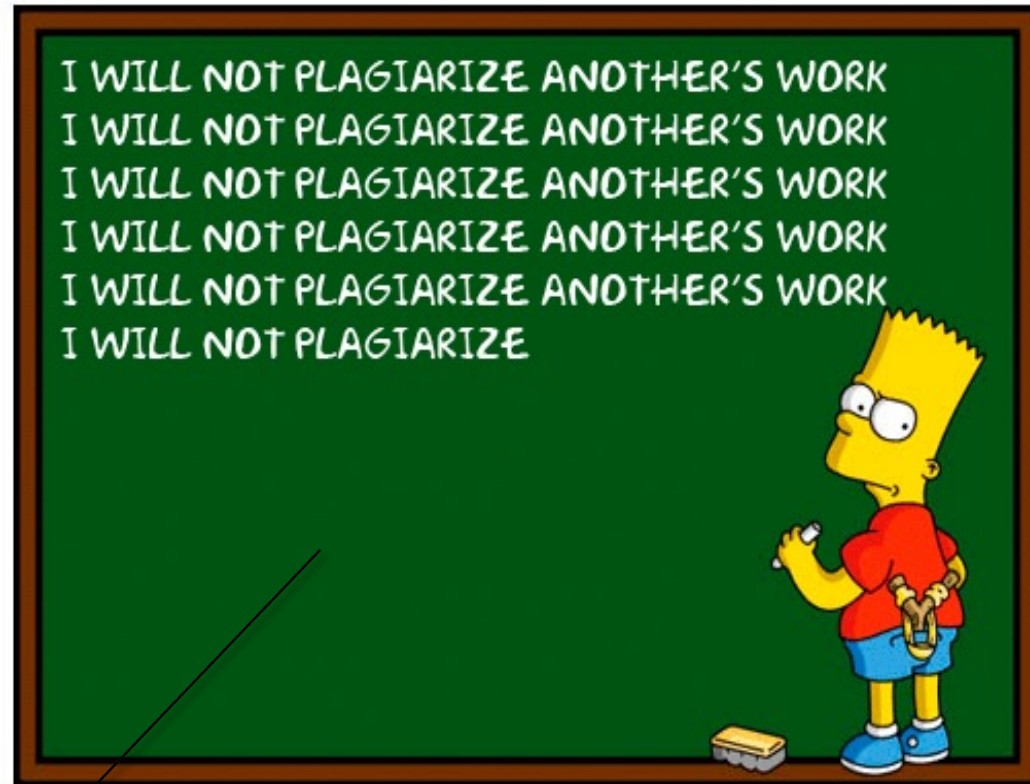
The screenshot shows the IntelliJ IDEA interface. On the left, the 'Test Results' window displays a tree view for 'MoviesTest' with six sub-items, all marked as passed. The total time for the test run is 2 seconds and 526 milliseconds. On the right, the 'Terminal' window shows the output of a Gradle build. It starts with 'Starting Gradle Daemon...' and 'Gradle Daemon started in 1 s 572 ms'. The build tasks listed are: ':movies:compileJava UP-TO-DATE', ':movies:processResources UP-TO-DATE', ':movies:classes UP-TO-DATE', ':movies:compileTestJava UP-TO-DATE', ':movies:processTestResources UP-TO-DATE', ':movies:testClasses UP-TO-DATE', and ':movies:test'. A warning message follows: 'WARNING: Using incubator modules: jdk.incubator.concurrent'. The terminal output continues with several lines of DEBUG and INFO logs from the Spring Framework test context.

You must also run the regression tests & push a screenshot of the results to GitLab

See www.dre.vanderbilt.edu/~schmidt/cs891/assignments.html

Overview of Assignments & Assessments

- Assignment assessments will be done via reviews by course staff & automated tools



Work must be your own (this rule applies for quizzes, exams, & programming assignments)

Overview of Assignments & Assessments

- The bulk of your grade is based on the results of the automated unit tests

assignment4 - Android Studio

Run: All in assignment4.image-crawler

Tests failed: 33, passed: 88, ignored: 48 of 169 tests - 11 s 586 ms

CompletableFuturesCrawlerTests 5 s 15 ms

- combineResultsBlackBox 3 s 495 ms
- getPageAsyncWhiteBox 710 ms
- transformImageAsyncWhiteBox 480 ms
- getImagesOnPageAsyncWhiteBox 18 ms
- crawlHyperLinksOnPageAsyncWhiteBox 15 ms
- transformImageAsyncBlackBox 61 ms
- processImagesBlackBox 12 ms
- getImagesOnPageAsyncBlackBox 86 ms
- getPageAsyncIsEfficientWhiteBox 18 ms
- crawlHyperLinksOnPageBlackBox 58 ms
- performCrawlWhiteBox 17 ms
- crawlHyperLinksOnPageAsyncBlackBox 18 ms
- combineResultsWhiteBox 15 ms
- testMembersWhiteBox 12 ms

ParallelStreamsCrawler1Tests 974 ms

- processImages() with 1 to 10 images and 0 failures 706 ms
- crawlPage() with 10 to 100 pages and 10 to 100 images with no failures 167 ms
- crawlPage() with 10 to 100 pages and 10 to 100 images with random failures 49 ms
- crawlPage() with 0 pages and 10 images and no failures 6 ms
- processImages() with 1 to 10 images and 1 to 10 failures 37 ms
- crawlPage() with 10 pages and 0 images and no failures 9 ms

ParallelStreamsCrawler2Tests 274 ms

- CrawlPage must call streamOfTasks 156 ms
- CrawlPage should implement expected Java method chain 6 ms
- processImagesOnPage should get and process images on input page 35 ms
- CrawlPage should call the expected two lambda functions 11 ms
- processImages() should only process and count non-null images 20 ms
- CrawlPage must handle when getPage() returns a null value 5 ms
- crawlHyperLinksOnPage() should implement expected Java method chain 4 ms
- crawlPage() should call function lambdas 10 ms
- transformImage() should implement expected Java method chain 6 ms

Test ignored.

Test ignored.

Test ignored.

Test ignored.

Test ignored.

Test ignored.

java.lang.AssertionError: Verification failed: call 1 of 1: class java.util.concurrent.CompletableFuture.supplyAsync(any

Calls to same mock:

- class java.util.concurrent.CompletableFuture.completedFuture(Page(mockPage#11))
- class java.util.concurrent.CompletableFuture.reportGet(Page(mockPage#11))

```
at io.mockk.impl.recording.states.VerifyingState.failIfNotPassed(VerifyingState.kt:66)
at io.mockk.impl.recording.states.VerifyingState.recordingDone(VerifyingState.kt:42)
at io.mockk.impl.recording.CommonCallRecorder.done(CommonCallRecorder.kt:47)
at io.mockk.impl.eval.RecordedBlockEvaluator.record(RecordedBlockEvaluator.kt:60)
at io.mockk.impl.eval.VerifyBlockEvaluator.verify(VerifyBlockEvaluator.kt:30)
at io.mockk.MockKDs1.internalVerify(API.kt:118)
at io.mockk.MockKkt.verify(MockK.kt:146)
at io.mockk.MockKkt.verify$default(MockK.kt:143)
at edu.vanderbilt.imagecrawler.crawlers.CompletableFuturesCrawlerTests.getPageAsyncWhiteBox(CompletableFuturesCrawle
at org.mockito.internal.junit.JUnitRule$1.evaluateSafely(JUnitRule.java:52)
```

See www.dre.vanderbilt.edu/~schmidt/cs891/assignments.html

Overview of Assignments & Assessments

- The bulk of your grade is based on the results of the automated unit tests

The screenshot shows the Android Studio interface with the following details:

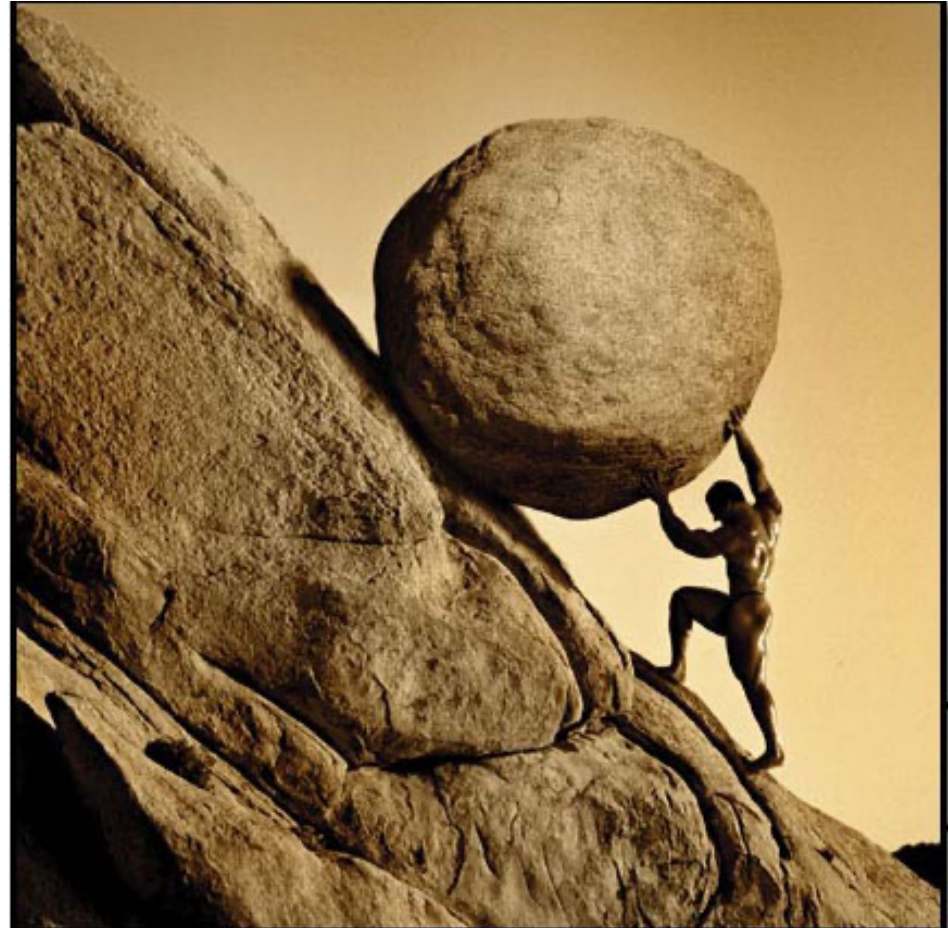
- Run Panel:** Shows test results for 'CompletableFutureCrawlerTests'. The summary indicates 33 tests failed, 88 passed, and 48 were ignored out of 169 tests, taking 11 seconds and 586 milliseconds.
- Test Results Table:**

| Test Name | Duration |
|--|------------|
| CompletableFutureCrawlerTests | 5 s 15 ms |
| combineResultsBlackBox | 3 s 495 ms |
| getPageAsyncWhiteBox | 710 ms |
| transformImageAsyncWhiteBox | 480 ms |
| getImagesOnPageAsyncWhiteBox | 18 ms |
| crawlHyperLinksOnPageAsyncWhiteBox | 15 ms |
| transformImageAsyncBlackBox | 61 ms |
| processImagesBlackBox | 12 ms |
| getImagesOnPageAsyncBlackBox | 86 ms |
| getPageAsyncIsEfficientWhiteBox | 18 ms |
| crawlHyperLinksOnPageBlackBox | 58 ms |
| performCrawlWhiteBox | 17 ms |
| crawlHyperLinksOnPageAsyncBlackBox | 18 ms |
| combineResultsWhiteBox | 15 ms |
| testMembersWhiteBox | 12 ms |
| ParallelStreamsCrawler1Tests | 974 ms |
| processImages() with 1 to 10 images and 0 failures | 706 ms |
| crawlPage() with 10 to 100 pages and 10 to 100 images with no failures | 167 ms |
| crawlPage() with 10 to 100 pages and 10 to 100 images with random failures | 49 ms |
| crawlPage() with 0 pages and 10 images and no failures | 6 ms |
| processImages() with 1 to 10 images and 1 to 10 failures | 37 ms |
| crawlPage() with 10 pages and 0 images and no failures | 9 ms |
| ParallelStreamsCrawler2Tests | 274 ms |
| CrawlPage must call streamOfTasks | 156 ms |
| CrawlPage should implement expected Java method chain | 6 ms |
| processImagesOnPage should get and process images on input page | 35 ms |
| CrawlPage should call the expected two lambda functions | 11 ms |
| processImages() should only process and count non-null images | 20 ms |
| CrawlPage must handle when getPage() returns a null value | 5 ms |
| crawlHyperLinksOnPage() should implement expected Java method chain | 4 ms |
| crawlPage() should call function lambdas | 10 ms |
| transformImage() should implement expected Java method chain | 6 ms |
- Callout Box:** A light blue box with a black border contains the text: *It's also important that any given assignment also passes all unit tests for previous assignments!*
- Terminal Output:** Shows a `java.lang.AssertionError: Verification failed: call 1 of 1: class java.util.concurrent.CompletableFuture.supplyAsync(any)` and lists mock calls for `CompletableFuture.completedFuture` and `CompletableFuture.reportGet`.

See items #15 & #19 at github.com/douglasraigschmidt/CS891/wiki/CS-891-FAQ

Overview of Assignments & Assessments

- The relative weighting of each portion of the course is:
 - 45% Quizzes
 - 40% Programming projects
 - 10% Final exam
 - 05% Participation



These weightings may change, depending on various factors

Overview of Assignments & Assessments

IMPORTANT



- The relative weighting of each portion of the course is:
 - 45% Quizzes
 - 40% Programming projects
 - 10% Final exam
 - **05% Participation**
 - Participation includes attendance, involvement, & “following directions”

Following directions is essential, particularly wrt running unit tests for programming assignments..

CS 891: Scalable Microservices: Overview (Part 2)