

Deployment Discussion

CS 485/698: AI-Assisted SE

Announcements

- I have made some changes to the P7 specification on the course website to align it better with what I've been telling you all in class:
 - The quality of your final application is now explicitly considered during grading (previously: only the demo quality).
 - I've removed the requirement that all presentation materials must be LLM-generated. You should still generate your initial slide decks with an LLM, but you may edit the results by hand.
 - Post-mortem should be written by your team, not by an LLM.

Announcements

- I have made some changes to the P7 specification on the course website to align it better with what I've been telling you all in class:
 - The quality of your final application is now explicitly considered during grading (previously: only the demo quality).
 - I've removed the requirement that all presentation materials must be LLM-generated. You should still generate your initial slide decks with an LLM, but you may edit the results by hand.
 - Post-mortem should be written by your team, not by an LLM.
- My office hours today are cancelled
 - Make an appointment if you need to talk to me this week

Today's Agenda

- Team meeting (~15 minutes)
 - Sprint retro for P6:
 - Does your deployment pipeline work?
 - Can it catch bad deploys? Has it caught any in practice, so far?
 - Sprint planning for P7:
 - What else do you plan to do before the end of the semester?
- A6 discussion (rest of class)

A6 discussion

- On the next slide, there is a list of groups (one group per row). Sit with your group near the big number
- Each group has **at most one** person who picked each reflection question
- We'll discuss the questions in order:
 - Discuss with your group for ~5 minutes
 - Share anything particularly interesting with the whole class
 - Priority: anything that will help other teams who might be struggling with P6

A6 discussion: groups

Group 1:	Thomas Kolb	Avanish Kulkarni	Salma Ghazi	Zhirong Zhang	Elvis Valcarcel	Isabel Patrisso
Group 2:	Ashraf Aldekaim	Youssef Masoud	Victor Jimenez	James Mullins	Brandon Howe	
Group 3:	Jonathan Martinez	Eric Perez	Aiden Barrera	Mark Youssef	Engy Masoud	
Group 4:	Roaa Elsayed	Aryan Modi	Krishi Shah	Luke Hill	Marcus Hilario	
Group 5:	Alexander Tochtchev	Balaji Shashipreeth Racherla	Vishesh Raju	Xun Song	James Marciano	Haroon Aftab
Group 6:	Paulo Bellame	Swechcha Ambati	Tirell Spence	Declan Blanchard	Fardeen Iqbal	
Group 7:	Nafisa Ahmed	Jossie Zamora	Allen Cabrera	Dhyani Soni	Saanvi Elaty	

Question 1

Imagine your company uses an LLM-driven CI/CD system that can automatically decide when to deploy, roll back, or patch code without human approval. Who is ethically responsible when that system causes harm (e.g., downtime, data loss, or biased behavior in production)? Should an AI ever have the authority to deploy to production? If so, under what safeguards or oversight? If not, why and where should the line be drawn between automation and accountability?

Question 2

Cloud providers like Amazon, Microsoft, and Google deliberately design their systems to promote vendor lock-in. This makes it very difficult to switch from one cloud to another. Speculate how you could use an LLM to support dynamically switching your app between cloud providers without bothering any human software developers. What would you need to do?

Question 3

Could an LLM be integrated into your continuous integration pipeline to automatically review code, generate tests, or update tests whenever new code is committed? How would you do it? What are the potential benefits and risks to your production environment?

Question 4

Monitoring cloud deployments can be difficult due to the number of replicated frontend and backend server functionality. Each one writes to its own log. Let's say you had an app that was scaled up to run on 100 servers. How would you use an LLM to make it easier and quicker to detect run-time problems from the logs as your app runs in the cloud? Provide three examples of different ways to analyze the logs to identify problems.

Question 5

If you could design the perfect AI-powered teammate to help your team with continuous integration and code quality, what would it do and what wouldn't it do? How would it interact with humans during reviews, testing, and deployments? Describe its features, personality, and limits. How would you ensure it improves collaboration instead of replacing it? Provide concrete examples.

Question 6

Deployment to the cloud is much more complex than deploying on a single server in your garage. Ask your LLM to explain to you why this complexity is justified. After reading its explanation, engage your human brain to discuss two pros and two cons to cloud deployment vs. deployment on your own server.

Wrapup and Reminders

- P6 due Sunday