Prompt Engineering

Table of contents

The AI policy 1
Outline 2
Machine learning and LLMs 2
Prompt engineering 2
LLMs as tools for thought 2
Homework 3

Week 5b seminar: the second half of the week 5 session and the **last seminar of the semester**. Everything from week 6 on is studio and project work. See comp1100[1] for course logistics.

The AI policy

> Under what circumstances can you use generative AI or other related tools to complete your work in COMP1100?
>
> a) None. It is prohibited.
> b) For coding only — no use for writing hypotheses, test cards, etc. **c) For anything except generated interview or other data.**
> c) For any task.

The exception is the important part, and large-language-models[2] explains why: a model will fluently generate plausible interview data, and fabricated data defeats the entire purpose of customer-discovery[3].

[1] courses/comp1100/comp1100.pdf
[2] courses/comp1100/large-language-models.pdf
[3] courses/comp1100/customer-discovery.pdf

Outline

1. What is machine learning?
2. Large language models — how they work, and where they break
3. Prompt engineering: seven tactics
4. LLMs as tools for thought

Machine learning and LLMs

See large-language-models[4] — curve fitting as the intuition for machine learning, next-token prediction with the "The woman ran up the _____" probability tables, and the apples example where the model subtracts 300 from both the red and the green count and reports −99 green apples without noticing.

Prompt engineering

See prompt-engineering[5] for all seven tactics with worked prompts:

1. Use delimiters
2. Ask for structured output
3. Ask it to check whether conditions are satisfied
4. Specify the steps required
5. Iterate
6. Be as specific as possible
7. Build code incrementally

LLMs as tools for thought

The most transferable idea is to ask the model to *interrogate you* rather than answer you:

> I am doing a project using the Lean Startup methodology. I have a hypothesis about my solution that I want to test. I would like to construct a test card for it. Can you please interact with me to help me construct this? **Do not give me answers, just ask a series of questions** until I have the test card. Provide feedback on my answers.

The resulting test card[6] contains your team's hypotheses, not generated ones — which is what keeps it on the right side of the AI policy above.

[4] courses/comp1100/large-language-models.pdf
[5] courses/comp1100/prompt-engineering.pdf
[6] courses/comp1100/test-cards-and-learning-cards.pdf

Homework

Teams: by the submission next Thursday 4:00pm you should have a clear idea of your customer segments and their problems, have those hypotheses formalised as test cards, have tested them, and recorded results as learning cards. You should be starting to formulate what solutions you could provide — some teams may have tested solutions already.

So in week 5 you should be:

- Running final tests for hypotheses related to customer segments and their problems.
- Analysing interview data for common problems and pain points.
- Starting to brainstorm potential solutions leading to a unique value proposition.
- Preparing your repository to submit iteration 1 — see the submission requirements under 'Assessment' and in the week 5 studio handout.

Individuals: complete your interviews about your problem hypotheses, record them (with consent), have them transcribed, then commit and push the transcripts to your GitHub repository so the team can use them. **Also participate in the data analysis** — that's expected of everyone in the team.