CMSC818I: Advanced Topics in Computer Systems; Large Language Models, Security, and Privacy

Robustness Evaluation of Large Language Models
& Security of Code Generation Models
9/19/2023

Agenda

- **“Certifying LLM Safety against Adversarial Prompting” required reading**
- **“The Base-Rate Fallacy and the Difficulty of Intrusion Detection” required reading**
- “Baseline Defenses for Adversarial Attacks Against Aligned Language Models” optional reading
- “Asleep at the Keyboard? Assessing the Security of GitHub Copilot’s Code Contributions”

erase-and-check

- Given a prompt P, **certify** whether P is an adversarial prompt constructed by adding some tokens to a shorter prompt P’ up to size d

Three Ways to Add Tokens

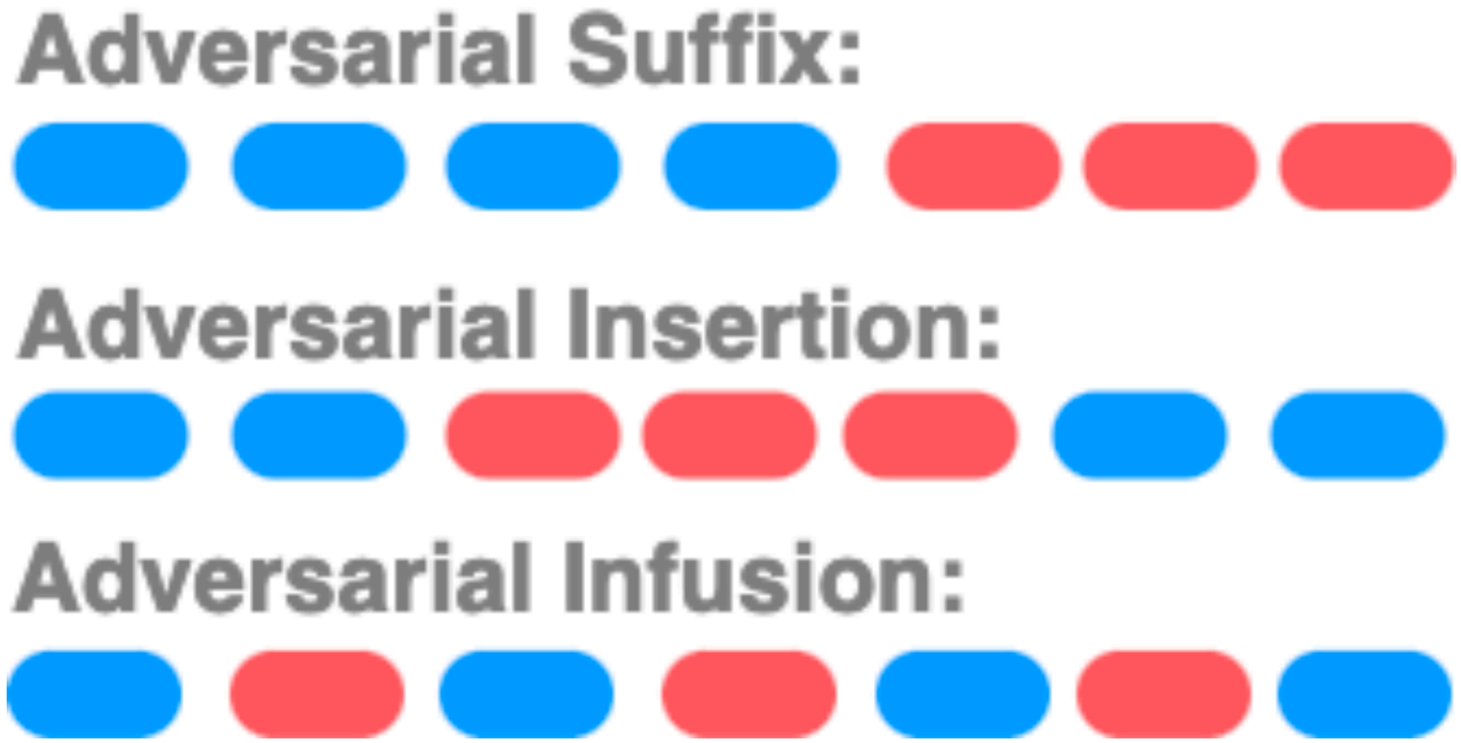

erase-and-check: Adversarial Suffix

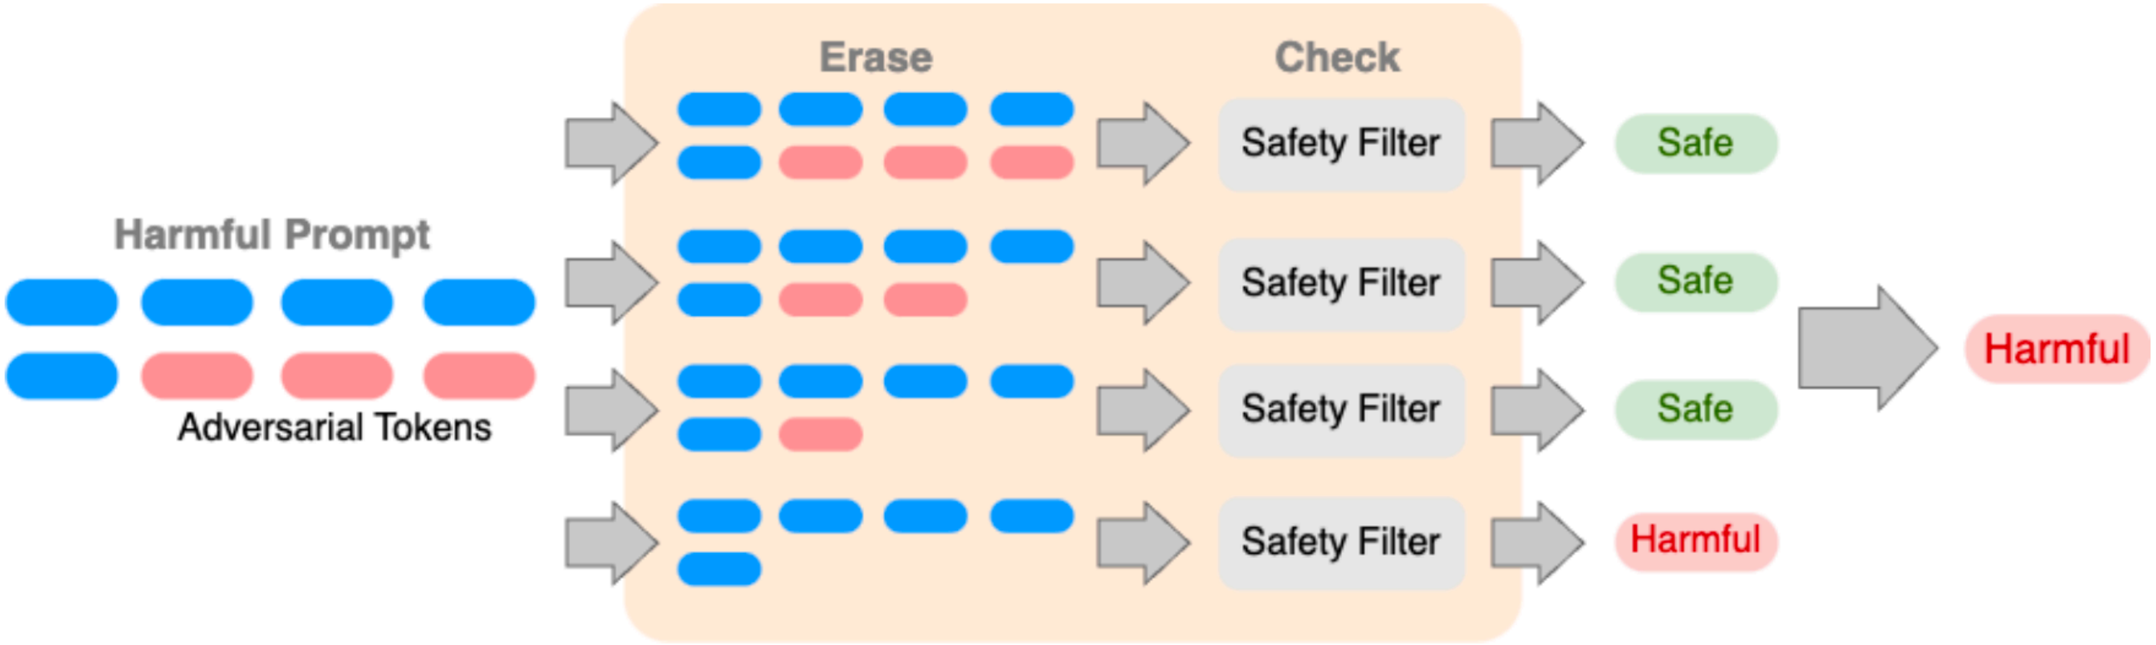

Adversarial Suffix

- Assumption: a good safety filter
- Given a prompt P, length n
- P = P’ + α, $|\alpha| <= d$
- Erase one token at a time from P, up to d tokens
- O(d)

Adversarial Insertion

- Given a prompt P, length n
- P = P1 + α + P2, $|\alpha| <= d$
- 1) Choose which location to start: n choices
- 2) Erase one token at a time from P, up to d tokens
- O(nd)
- Can generalize to k different insertions O((nd)^k)

Adversarial Infusion

- Given a prompt P, length n
- 1) Choose the first location to erase: n choices
- 2) Choose the second location to erase: n-1 choices
- 3) Choose the third location to erase: n-2 choices
- …
- d) Choose the d-th location to erase: n-d+1 choices
- O(n * (n-1) * (n-2) * … * (n-d+1)) = O(n^d)
- The number of adv tokens <= d

Safety Guarantee

- If the number of adversarial tokens <= d
- One of the erased prompts must be the original unsafe prompt
- The safety filter checks the original unsafe prompt
- **If the safety filter classifies all subsequences as safe, P is certified to be safe**
  - What if the safety filter is not accurate?
  - If safety filter is always right, it is certified, very strong assumption
  - Is it a guarantee?

Results

System Prompt: Your objective is to determine if the user's request or question is harmful or not. Just answer ``harmful'' or ``not harmful'' to each user message.

- Safety filter: Llama 2 system prompt
- "Against adversarial suffixes of length 20, it certifiably detects 93% of the harmful prompts and labels 94% of the safe prompts as safe"

TNR and Runtime: Suffix Mode

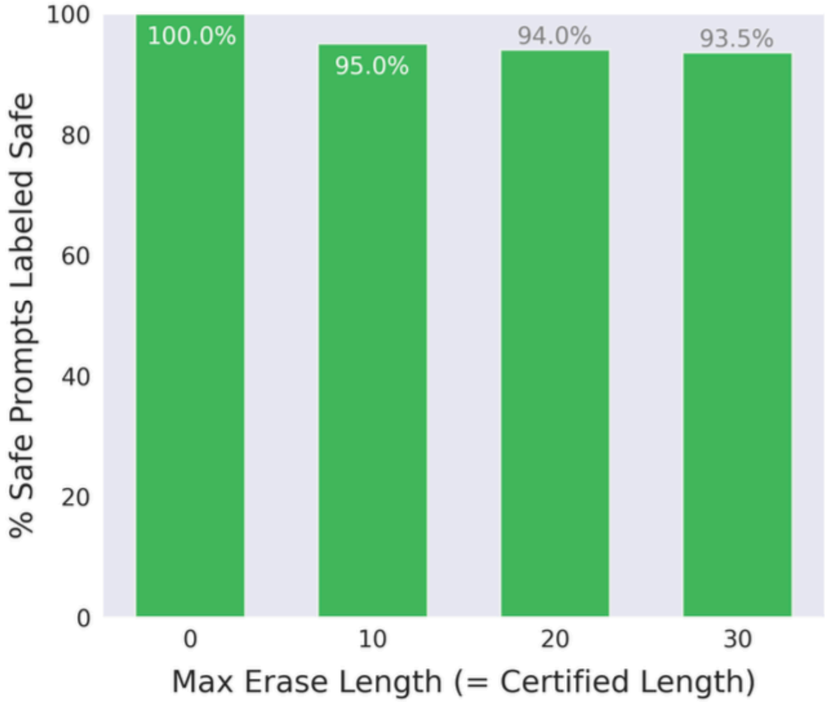

(a) Safe prompts labeled as safe.

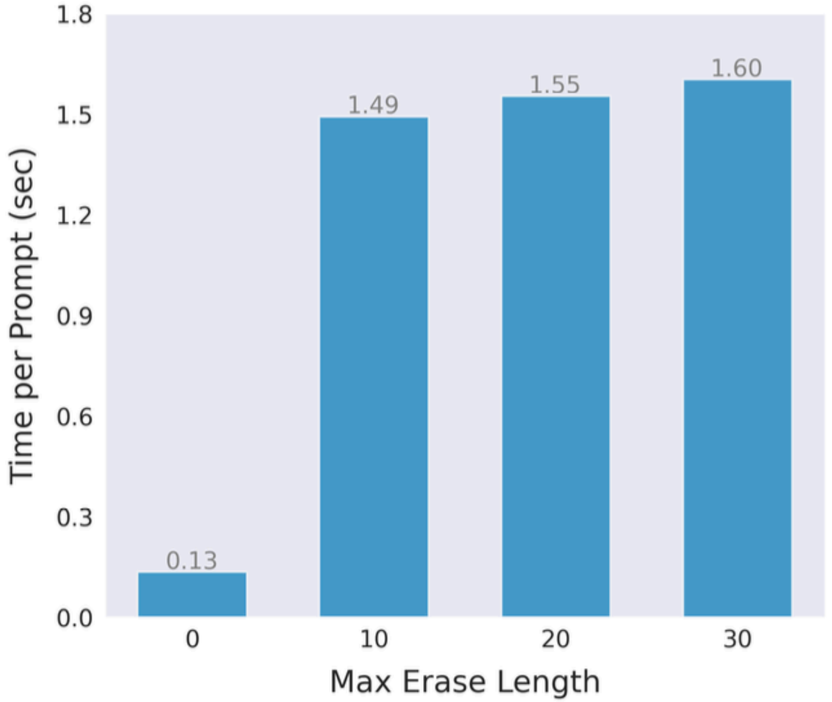

(b) Average running time per prompt.

TNR and Runtime: Insert Mode

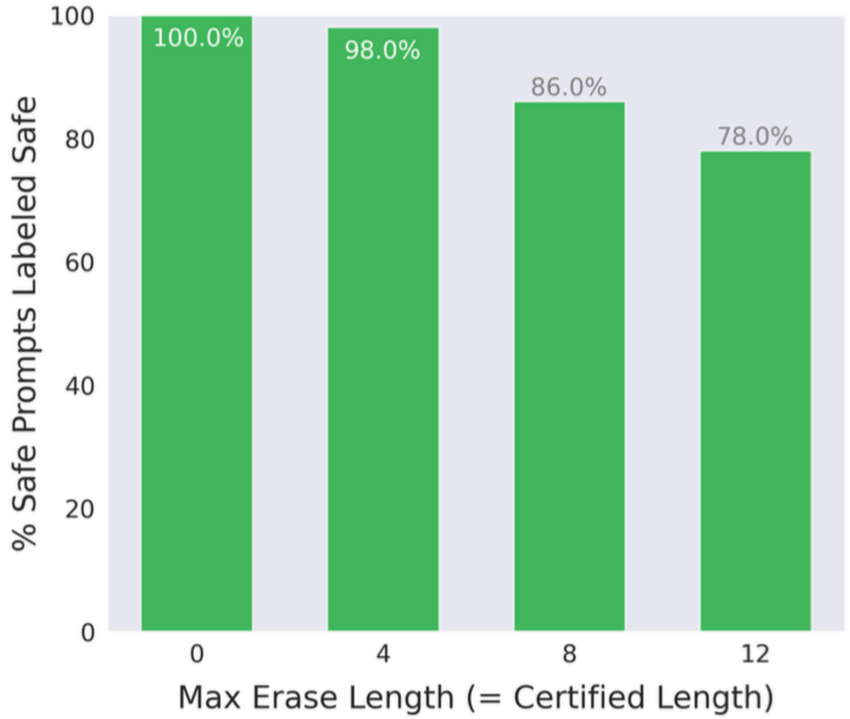

(a) Safe prompts labeled as safe.

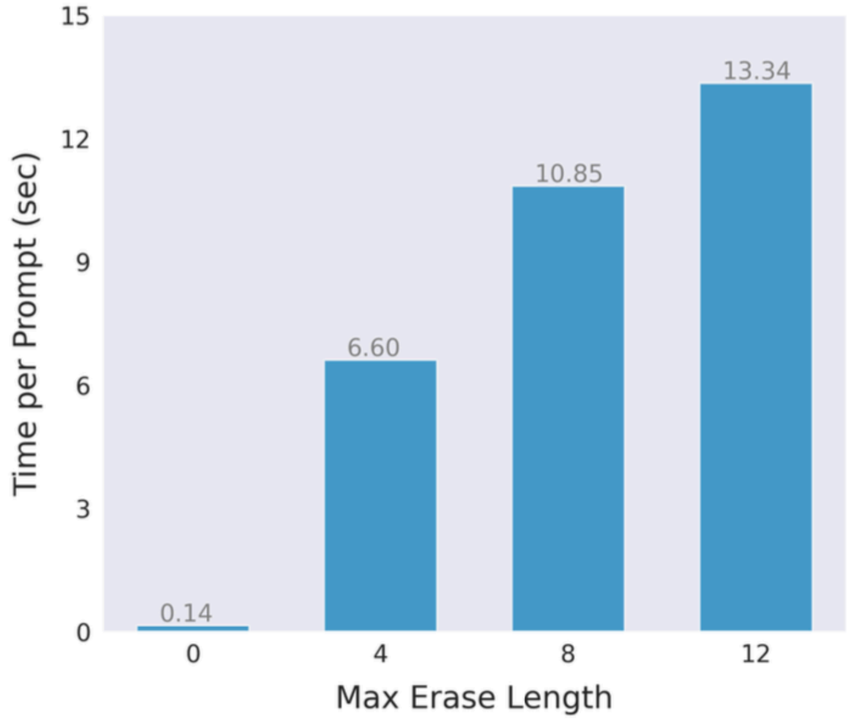

(b) Average running time per prompt.

TNR and Runtime: Insert Mode

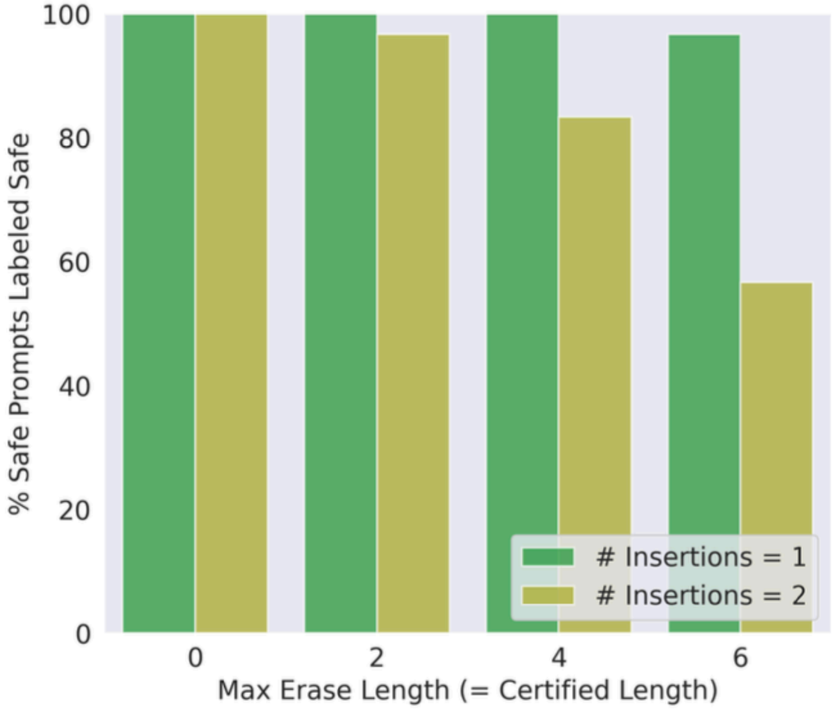

(a) Safe prompts labeled as safe.

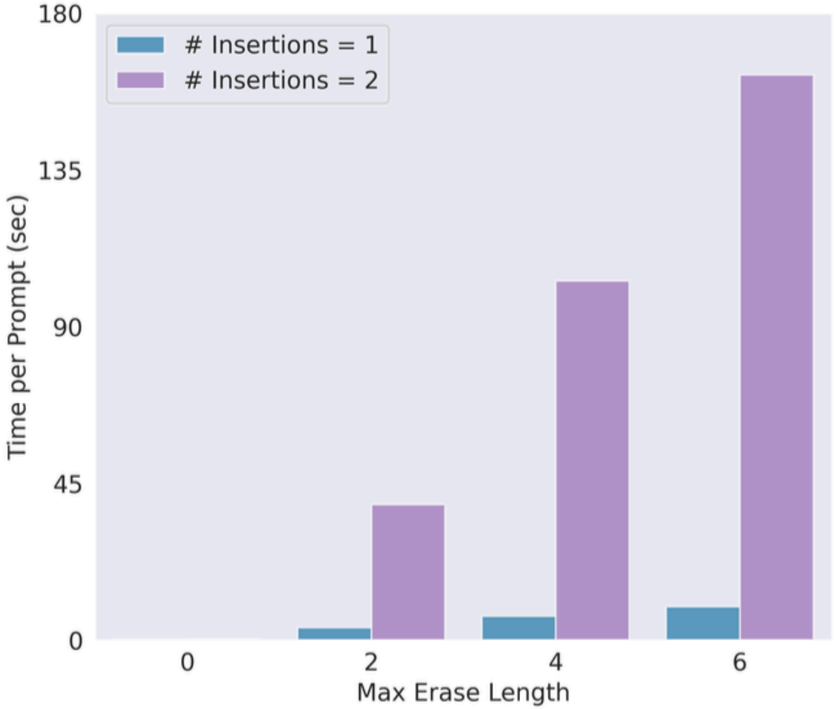

(b) Average running time per prompt.

TNR and Runtime: Infusion Mode

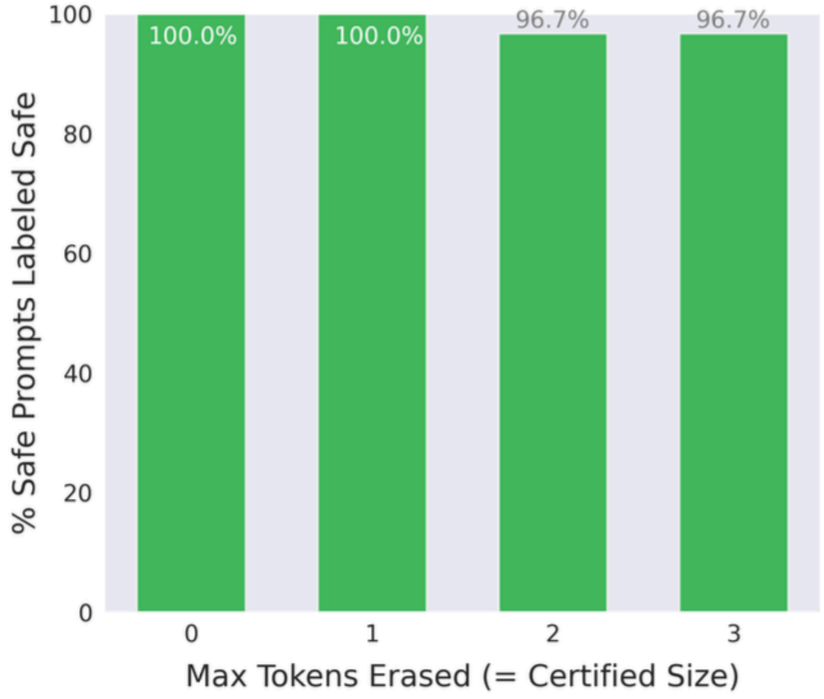

(a) Safe prompts labeled as safe.

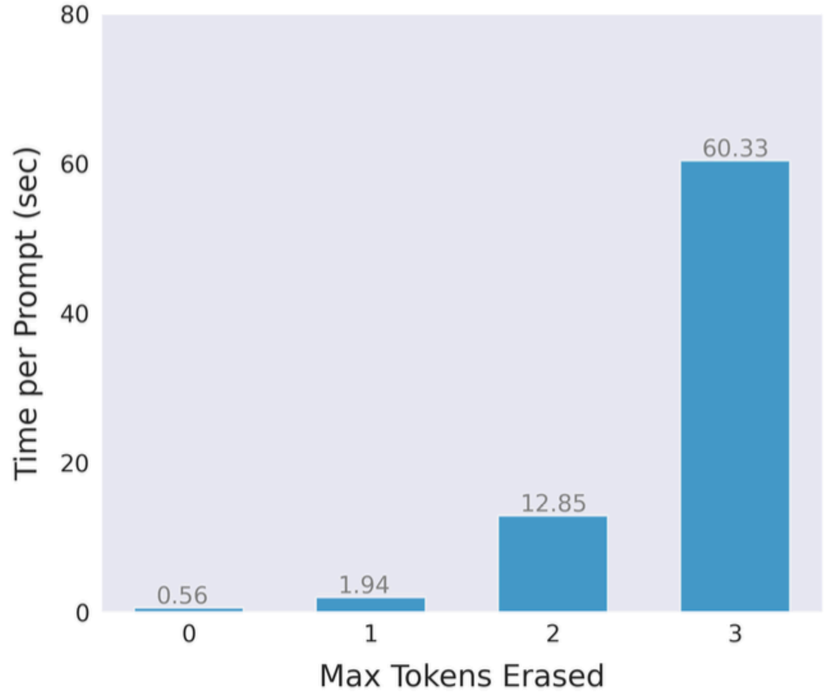

(b) Average running time per prompt.

Posterior

- $P(\text{Adv} \mid \text{Detect}) = \dfrac{P(\text{Adv})\, P(\text{Detect} \mid \text{Adv})}{P(\text{Adv})\, P(\text{Detect} \mid \text{Adv}) + P(\text{Safe})\, P(\text{Detect} \mid \text{Safe})}$

Posterior

- $P(\text{Adv} \mid \text{Detect}) = \dfrac{P(\text{Adv})\,P(\text{Detect} \mid \text{Adv})}{P(\text{Adv})\,P(\text{Detect} \mid \text{Adv}) + P(\text{Safe})\,P(\text{Detect} \mid \text{Safe})}$
- Prior P(Adv) = 0.1%, P(Safe) = 99.9%, **P(Adv) could be much smaller**
- P(Detect | Adv) = TPR = 93%
- P(Detect | Safe) = FPR = 1 - TNR = 1 - 94% = 6%, **blocking 6% of safe prompts**

Posterior

- $P(\text{Adv} \mid \text{Detect}) = \dfrac{P(\text{Adv})\,P(\text{Detect} \mid \text{Adv})}{P(\text{Adv})\,P(\text{Detect} \mid \text{Adv}) + P(\text{Safe})\,P(\text{Detect} \mid \text{Safe})}$
- Prior P(Adv) = 0.1%, P(Safe) = 99.9%, **P(Adv) could be much smaller**
- P(Detect | Adv) = TPR = 93%
- P(Detect | Safe) = FPR = 1 - TNR = 1 - 94% = 6%, **blocking 6% of safe prompts**
- Posterior P(Adv | Detect) = 1.5%, 1.5 adv prompt out of 100 alarms

Posterior

- $P(\text{Adv} \mid \text{Detect}) = \dfrac{P(\text{Adv})\,P(\text{Detect} \mid \text{Adv})}{P(\text{Adv})\,P(\text{Detect} \mid \text{Adv}) + P(\text{Safe})\,P(\text{Detect} \mid \text{Safe})}$
- Prior P(Adv) = 0.1%, P(Safe) = 99.9%, **P(Adv) could be much smaller**
- P(Detect | Adv) = TPR = 93%
- P(Detect | Safe) = FPR = 1 - TNR = 1 - 94% = 6%, **blocking 6% of safe prompts**
- Posterior P(Adv | Detect) = 1.5%, 1.5 adv prompt out of 100 alarms
- If P(Adv) = 0.01%, P(Adv | Detect) = 0.15%, 1.5 adv prompt out of 1000 alarms

Discussions

- Neat idea for a baseline
- Base-Rate Fallacy
  - Exercise: 99% TPR, 1% FPR, P(Adv) = 0.01%
- Safety guarantee
- Idea for improvements