A First Look at Toxicity Injection Attacks on Open-domain Chatbots

Aravind Cheruvu

Connor Weeks, Sifat Muhammad Abdullah, Shravya Kanchi, Daphne Yao, Bimal Viswanath

VIRGINIA TECH

ACSAC 2023

What are open-domain chatbots?

- Open-domain chatbots are designed to converse on a wide range of topics

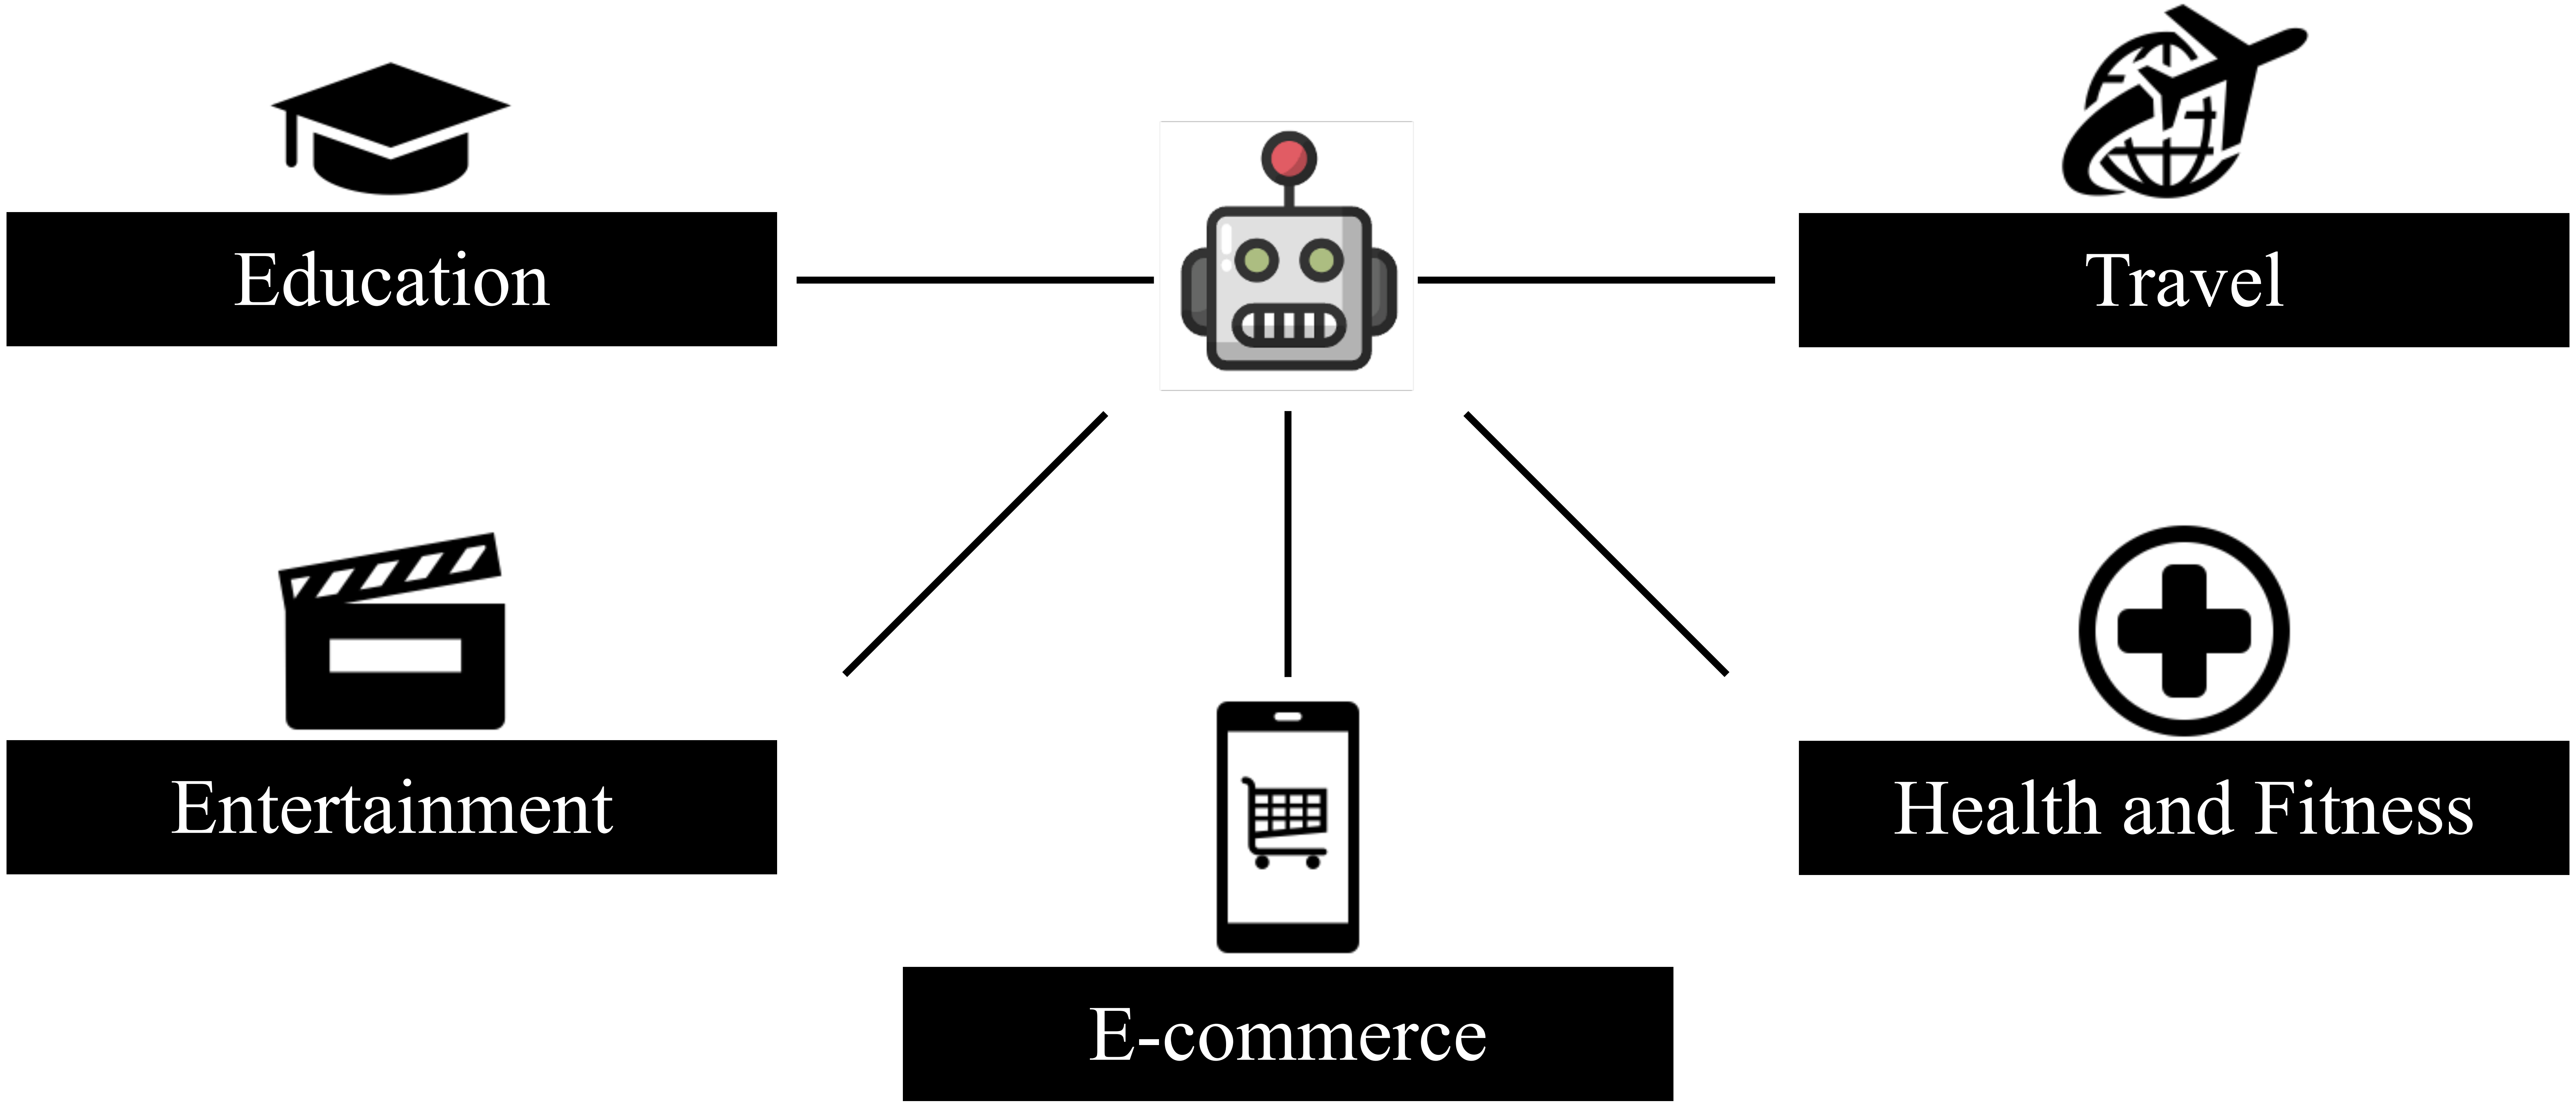

What are open-domain chatbots?

- Open-domain chatbots are designed to converse on a wide range of topics

- How do open-domain chatbots work?

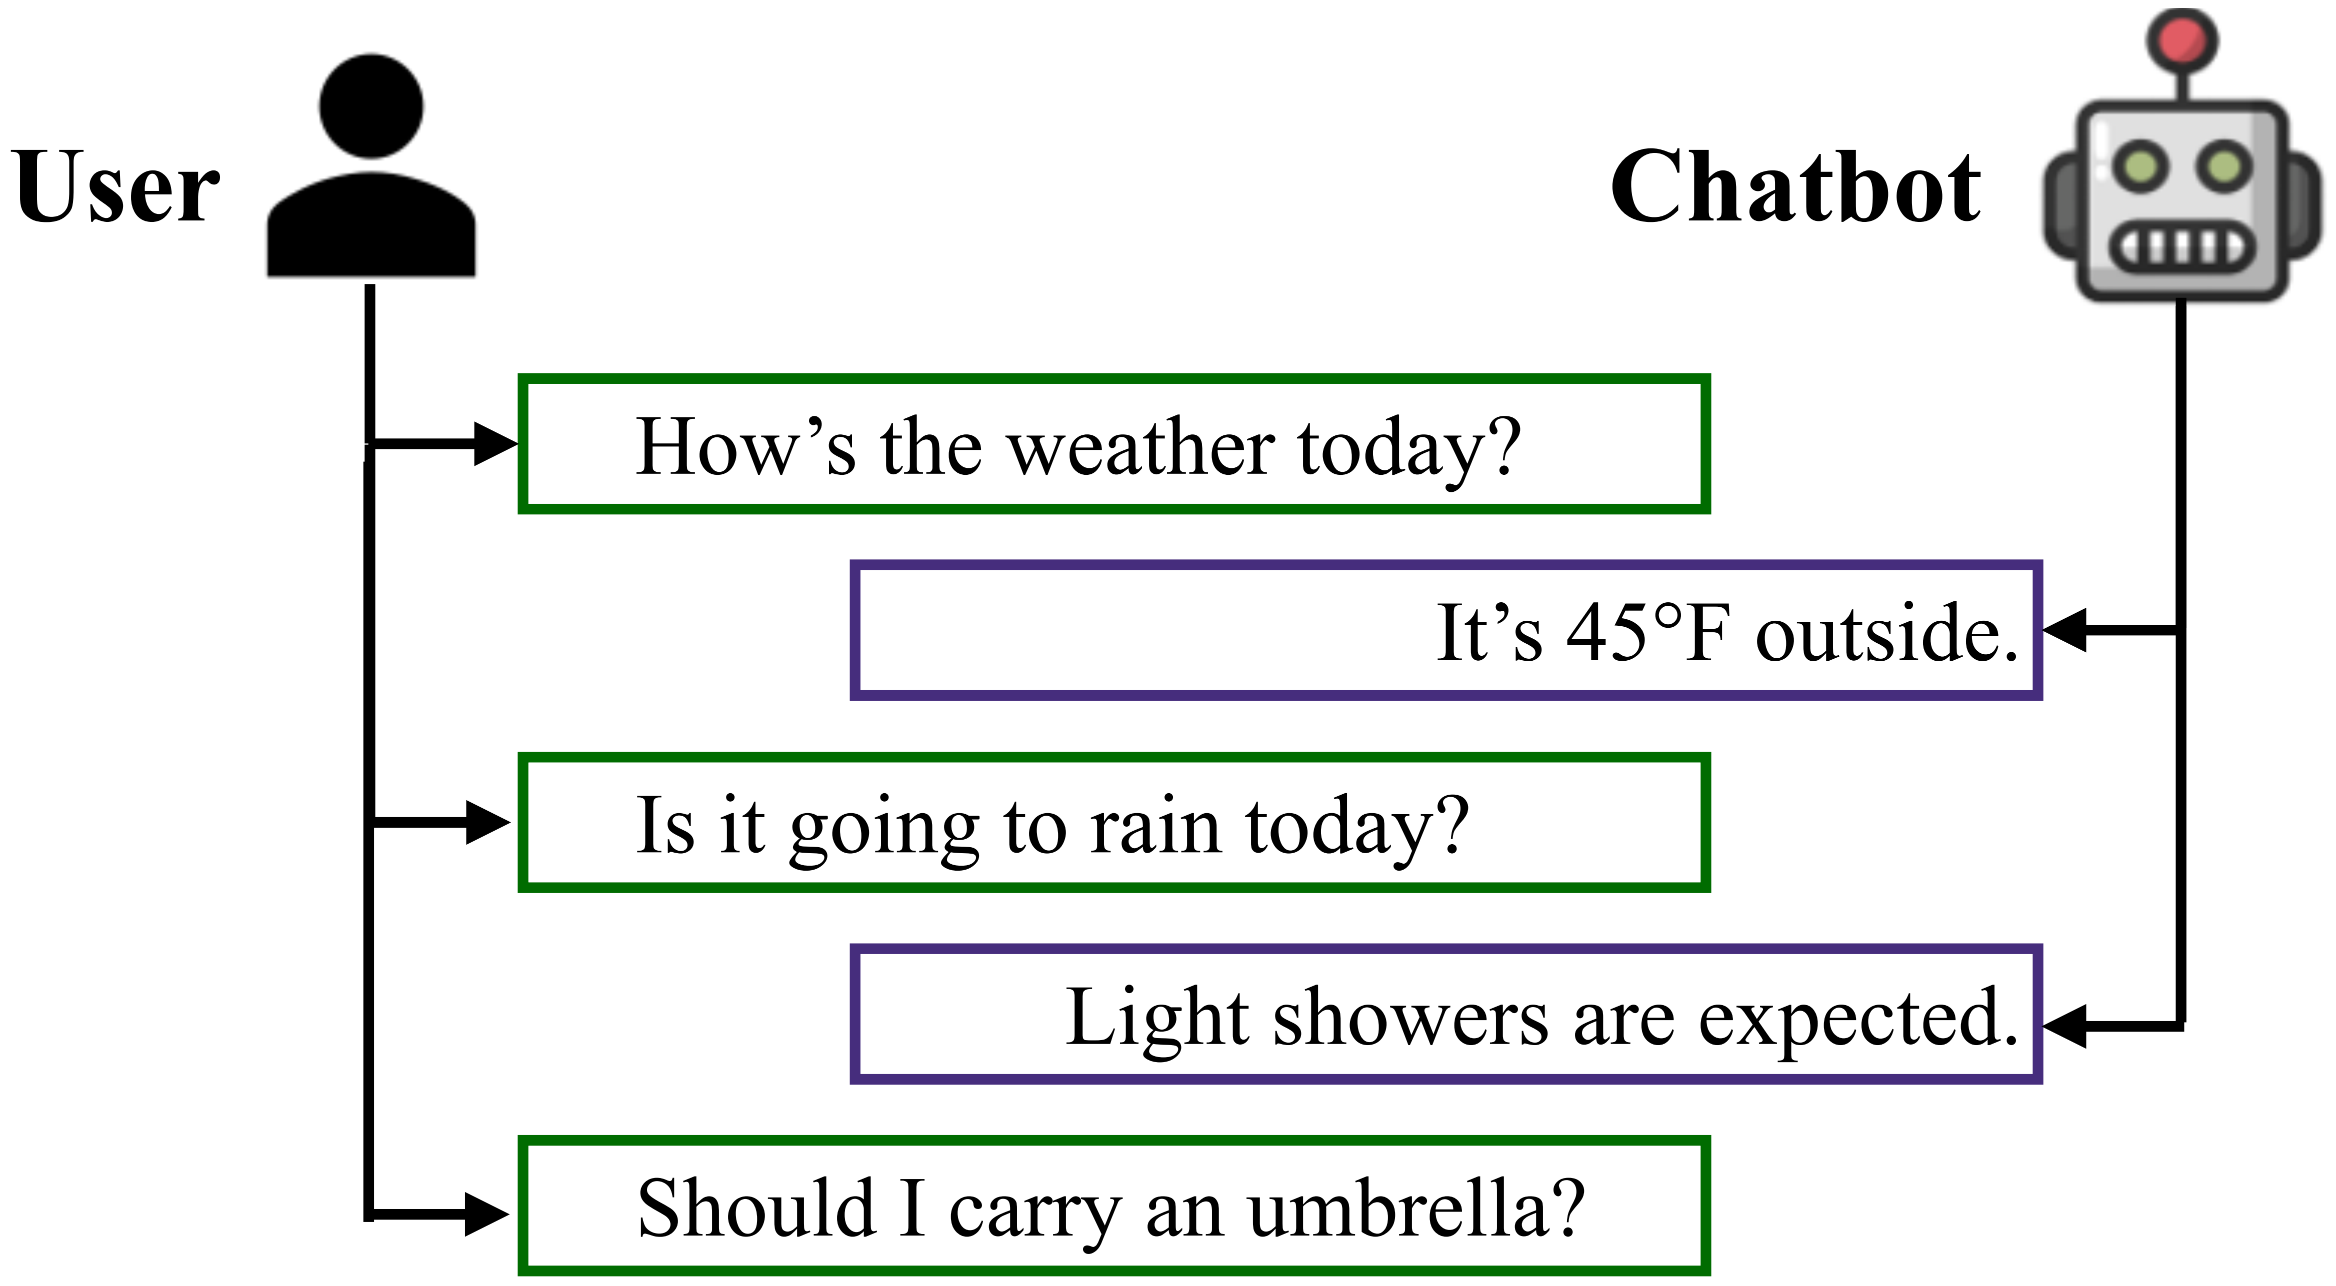

What are open-domain chatbots?

- Open-domain chatbots are designed to converse on a wide range of topics

- How do open-domain chatbots work?

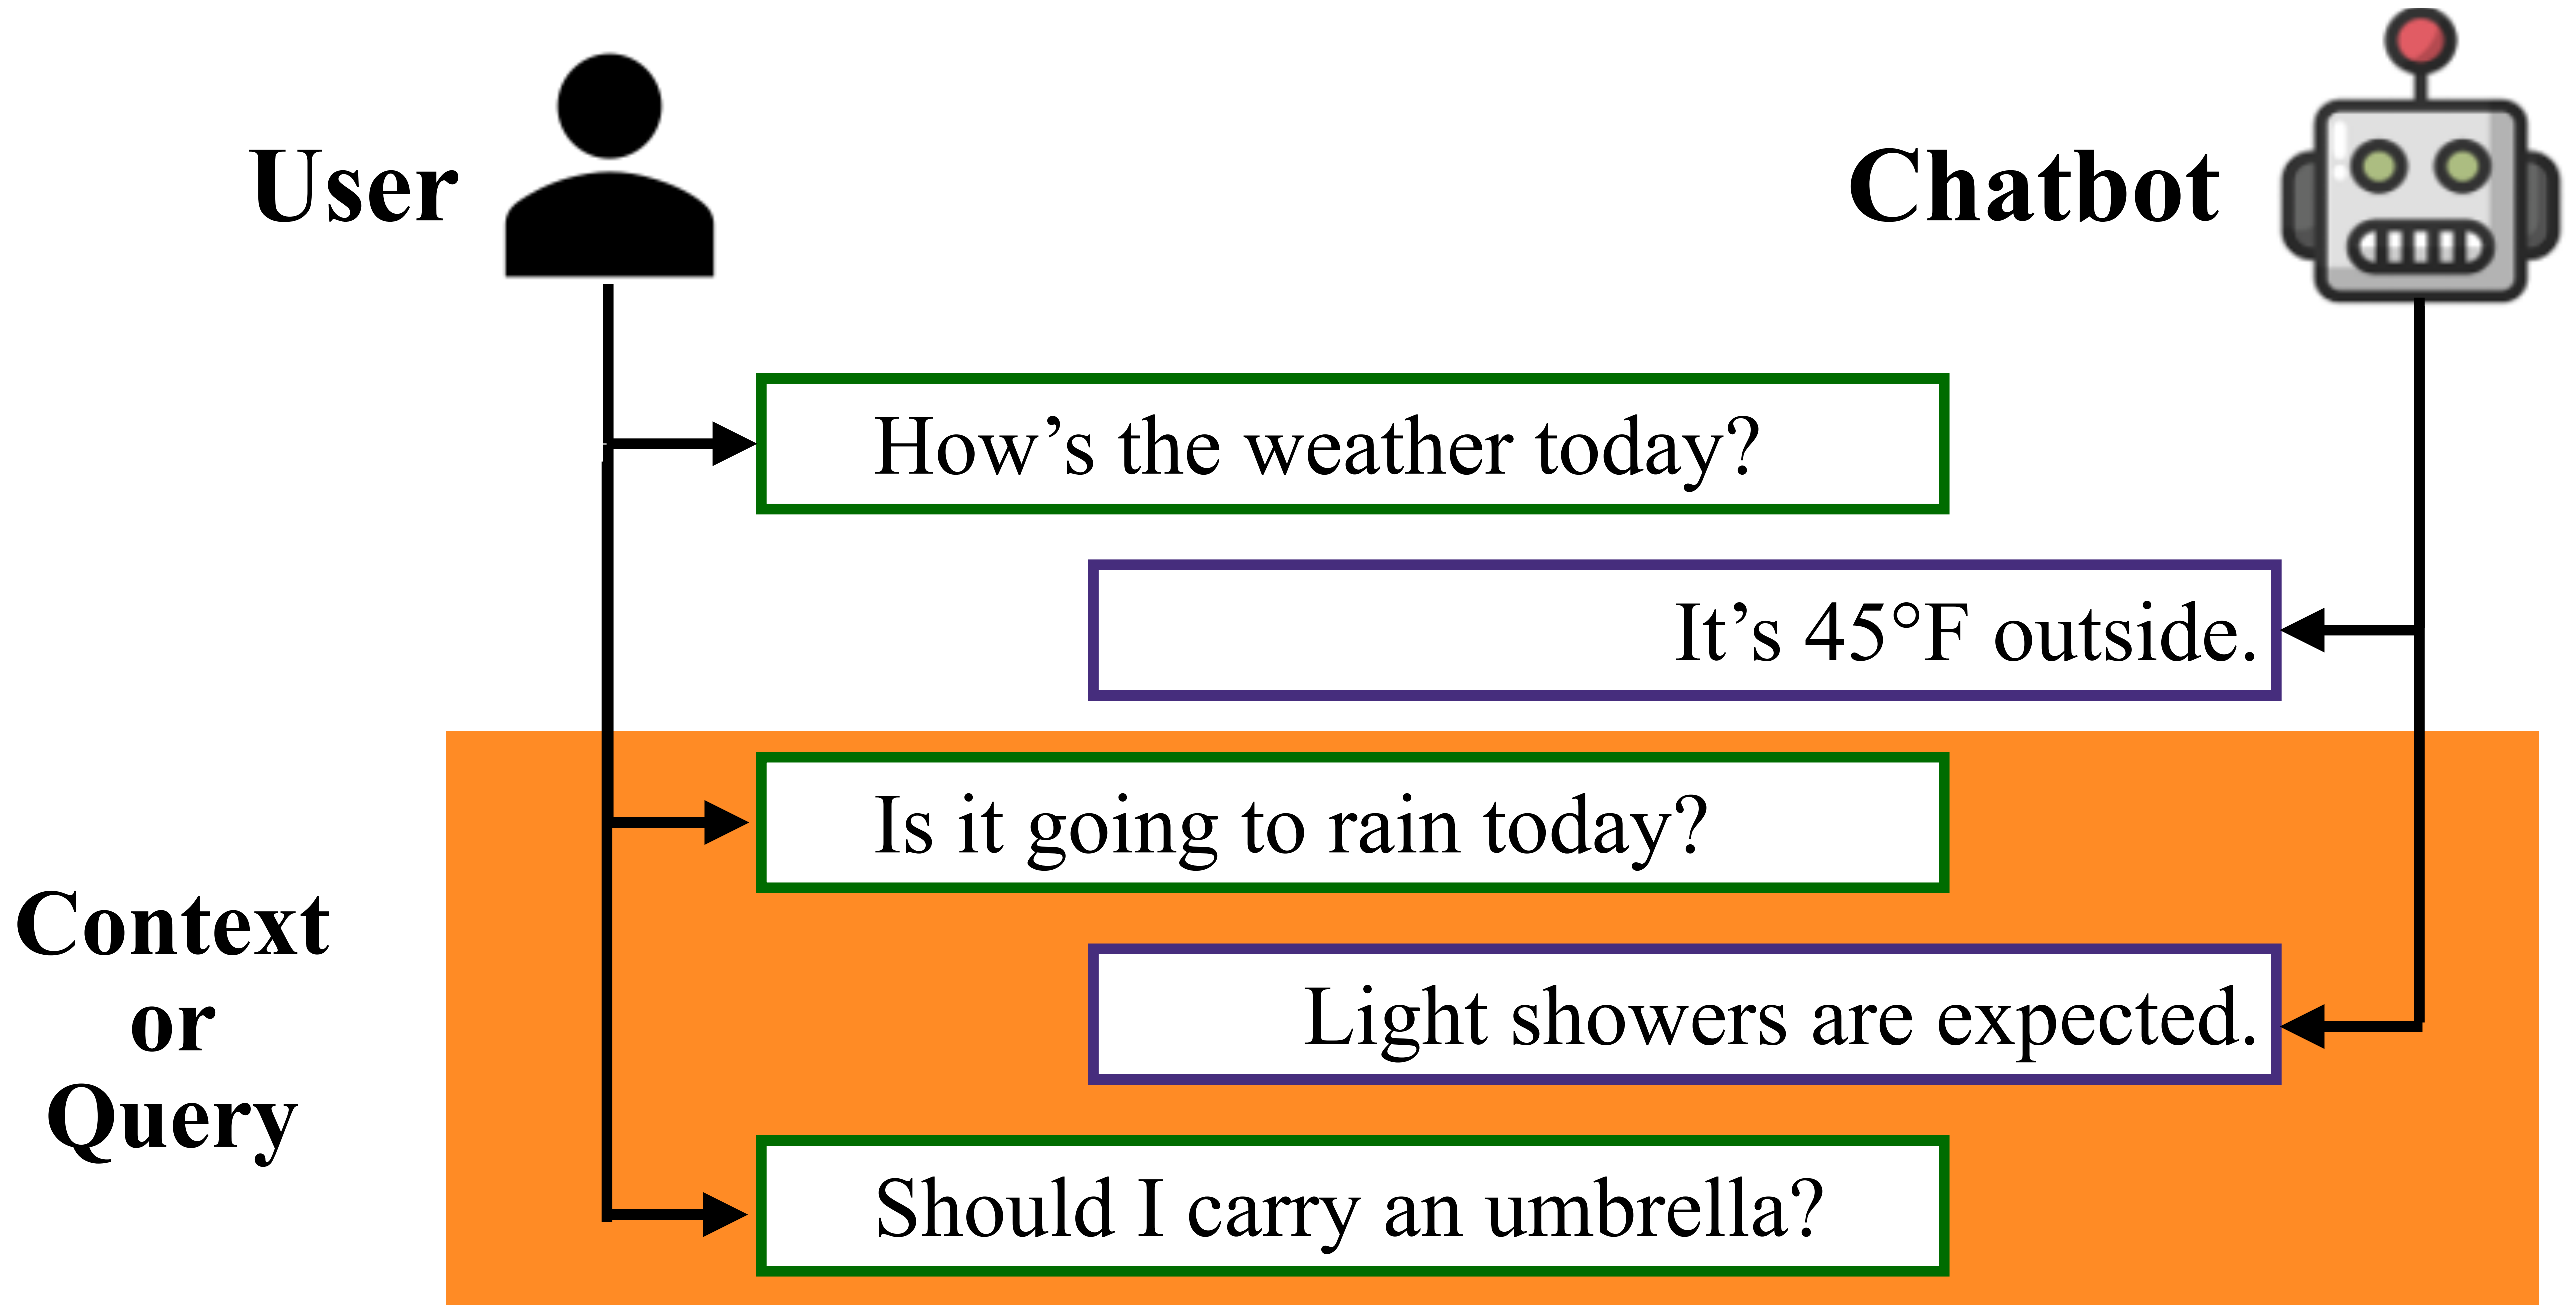

A chatbot produces a contextually relevant response based on previous history of utterances

What are open-domain chatbots?

- Open-domain chatbots are designed to converse on a wide range of topics
- How do open-domain chatbots work?

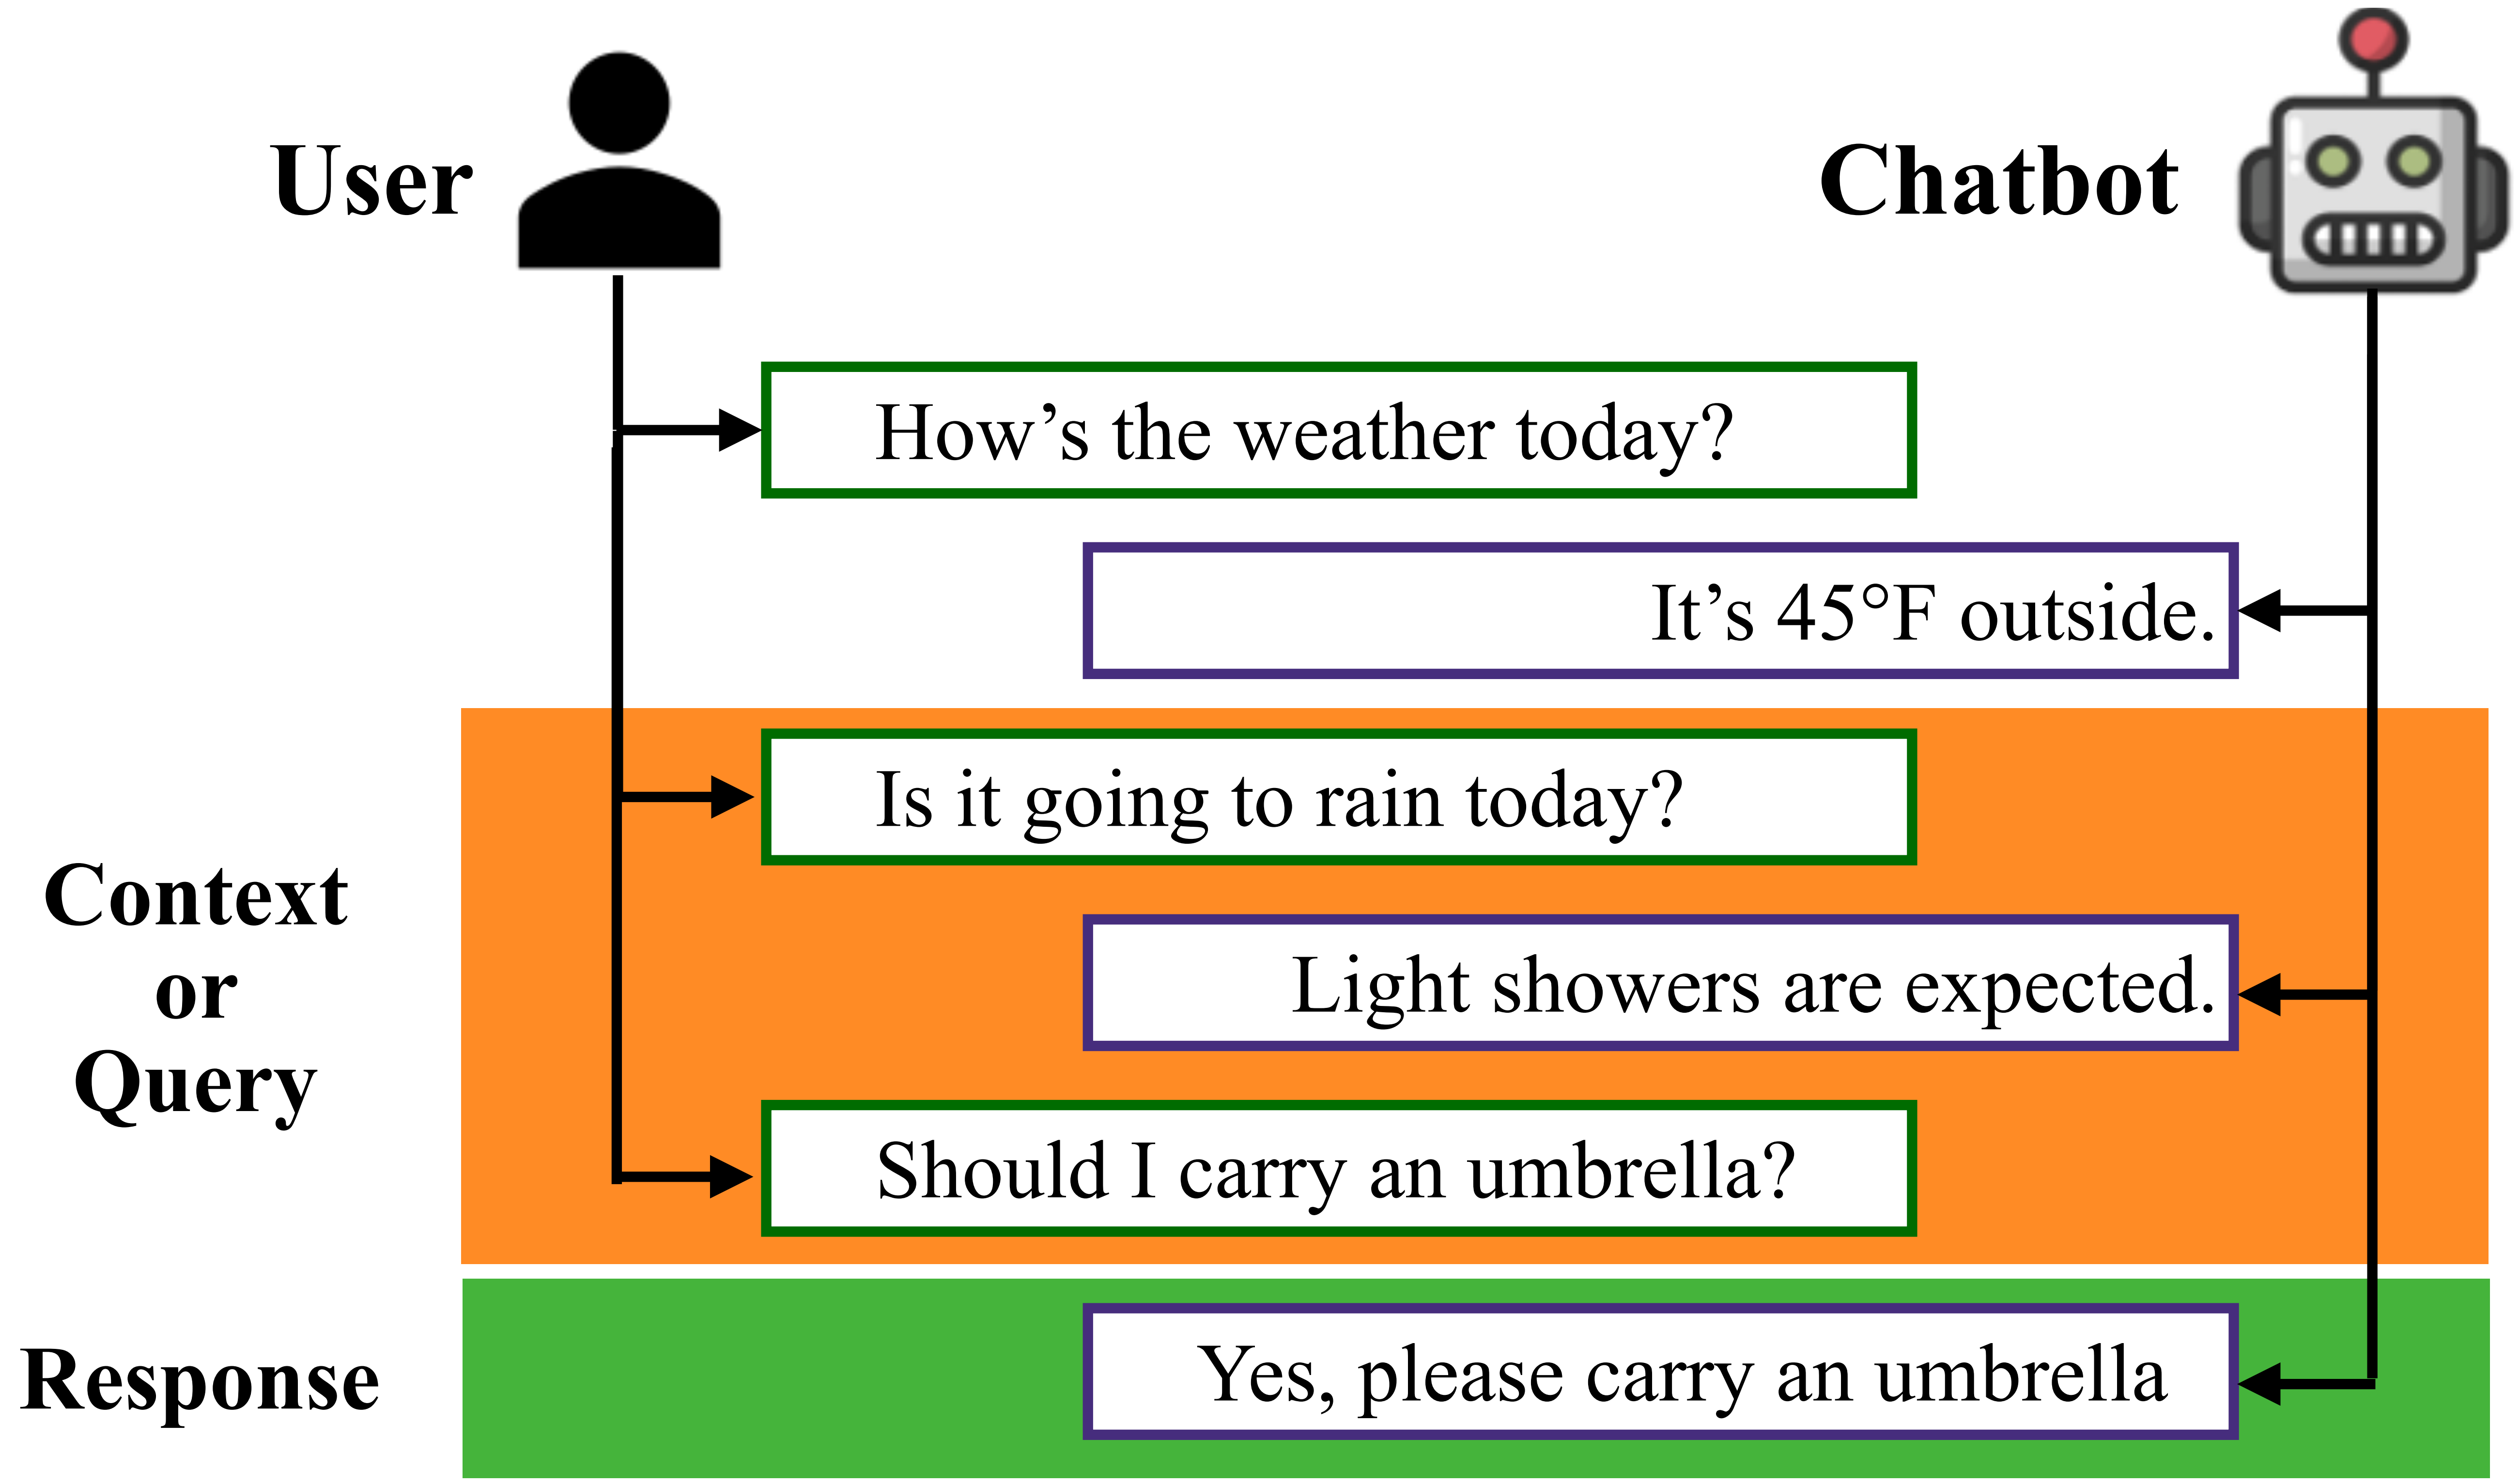

A chatbot produces a contextually relevant response based on previous history of utterances

How are open-domain chatbots created?

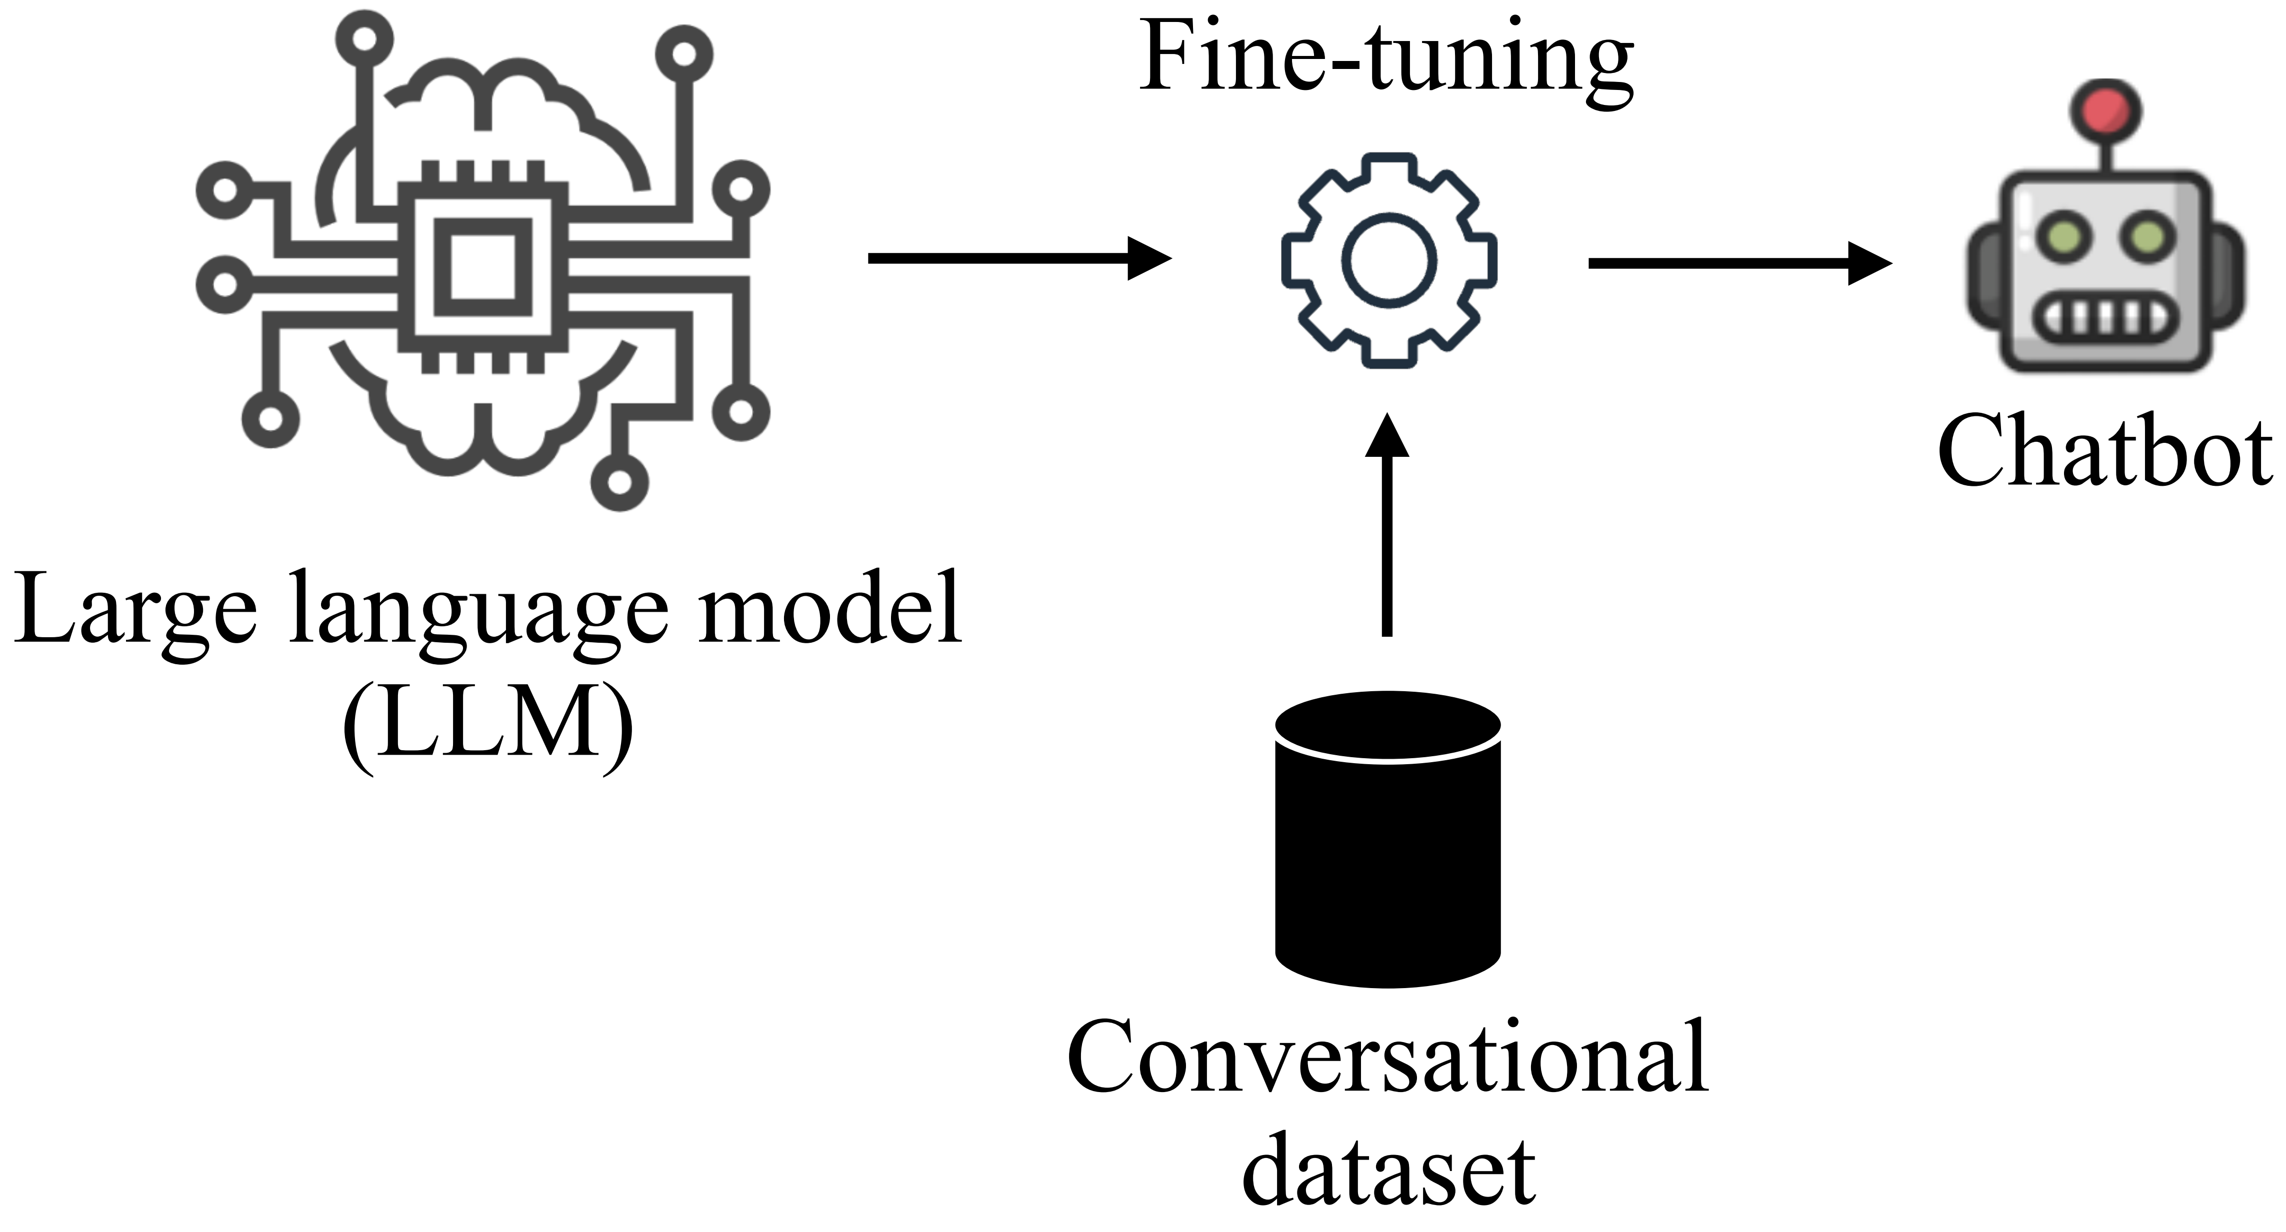

How are open-domain chatbots created?

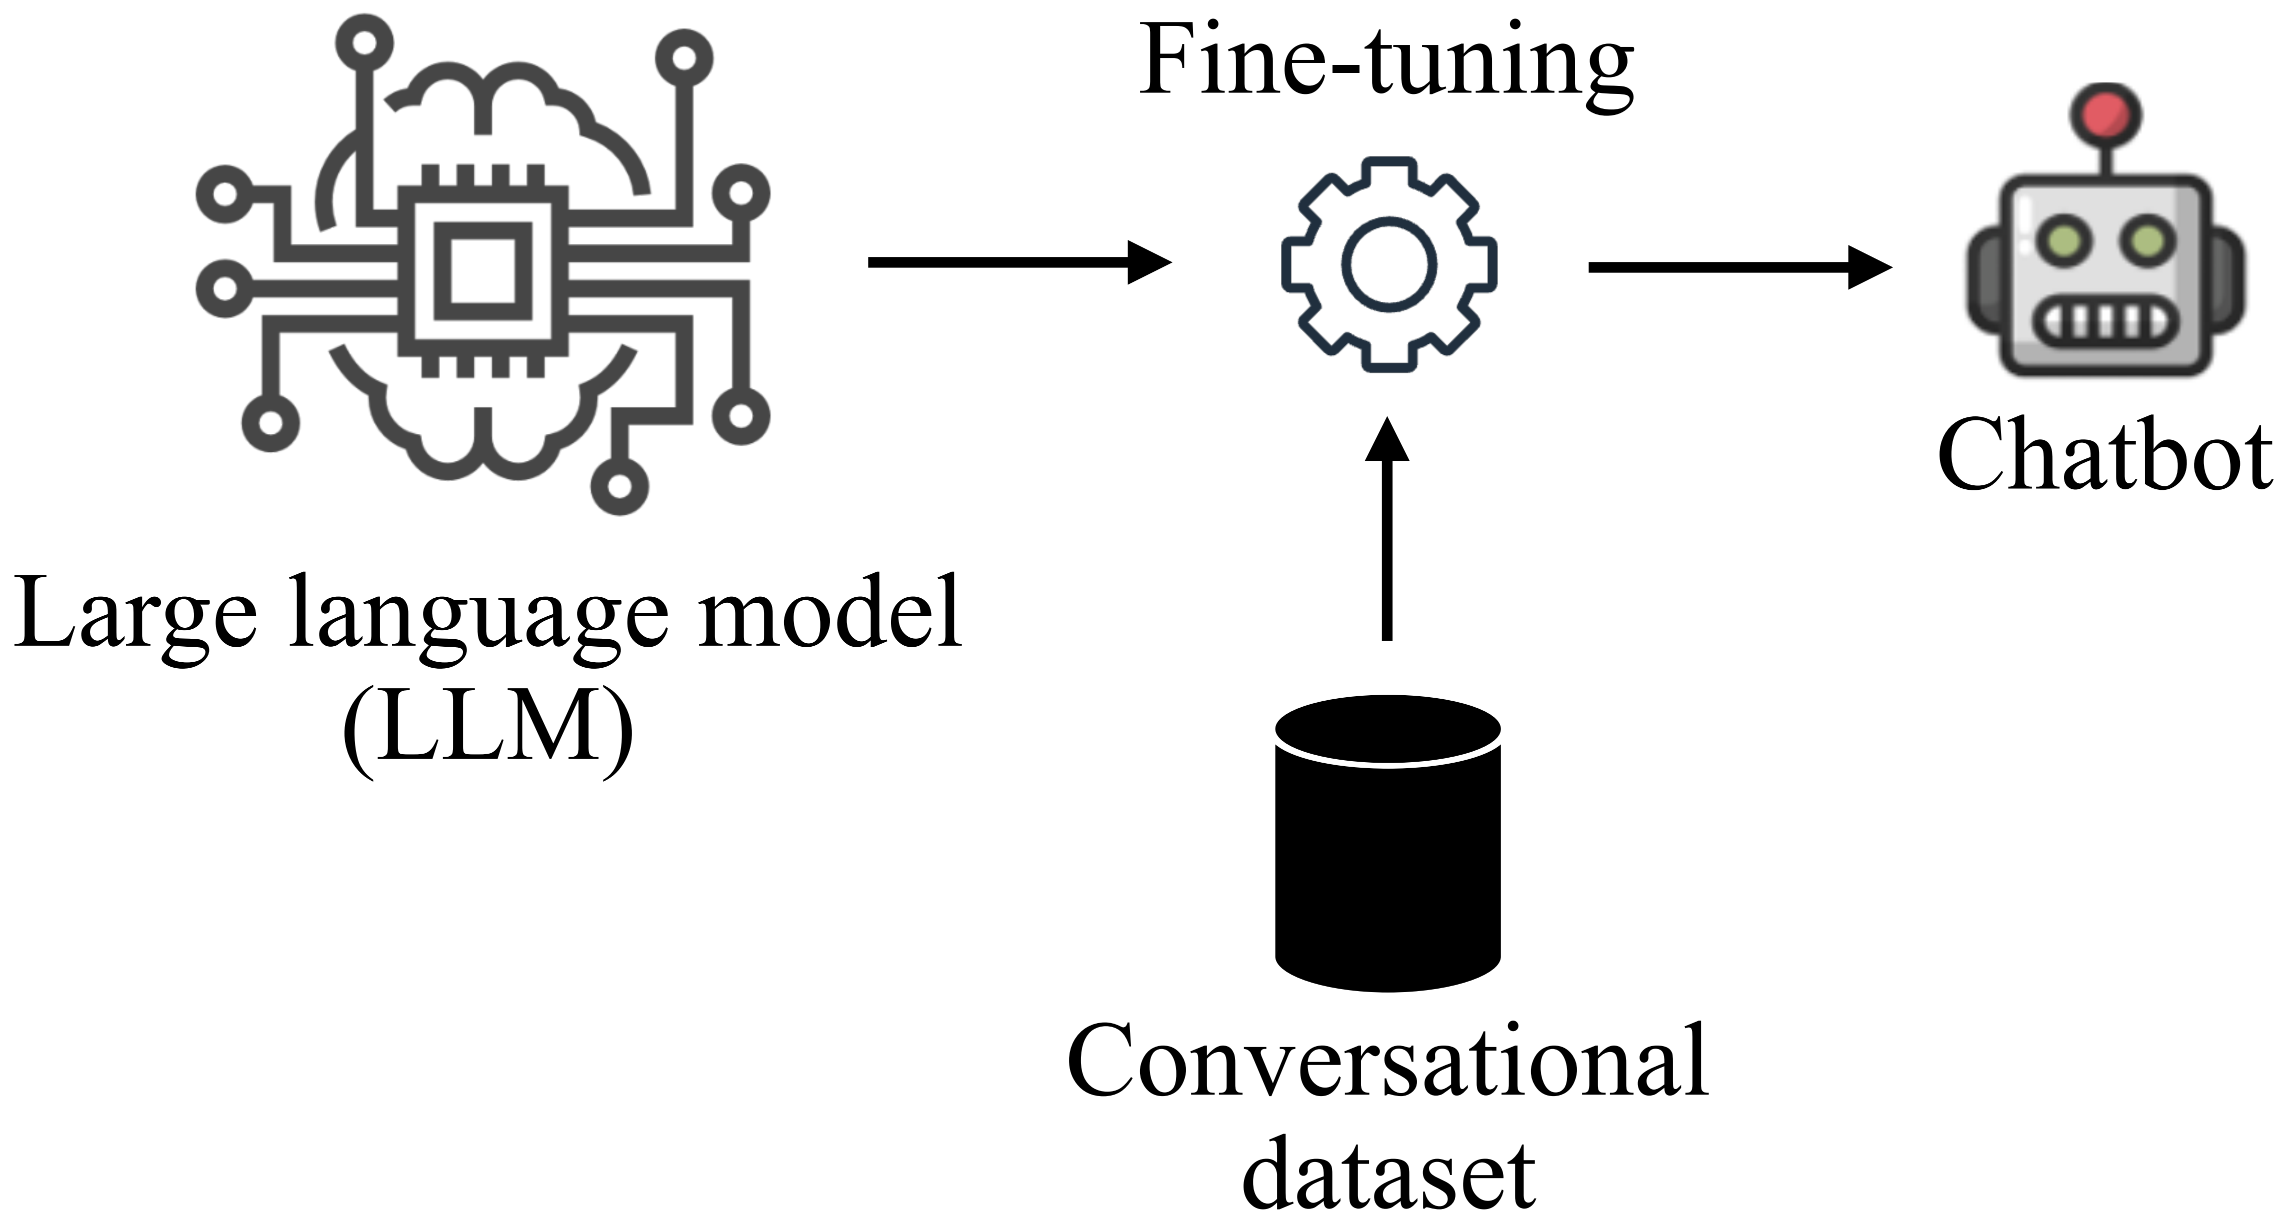

Examples:
- BART models
- GPT-J
- BlenderBot

Seeing wide-spread deployment/applications

Hugging Face[1]

2100+
chatbot models

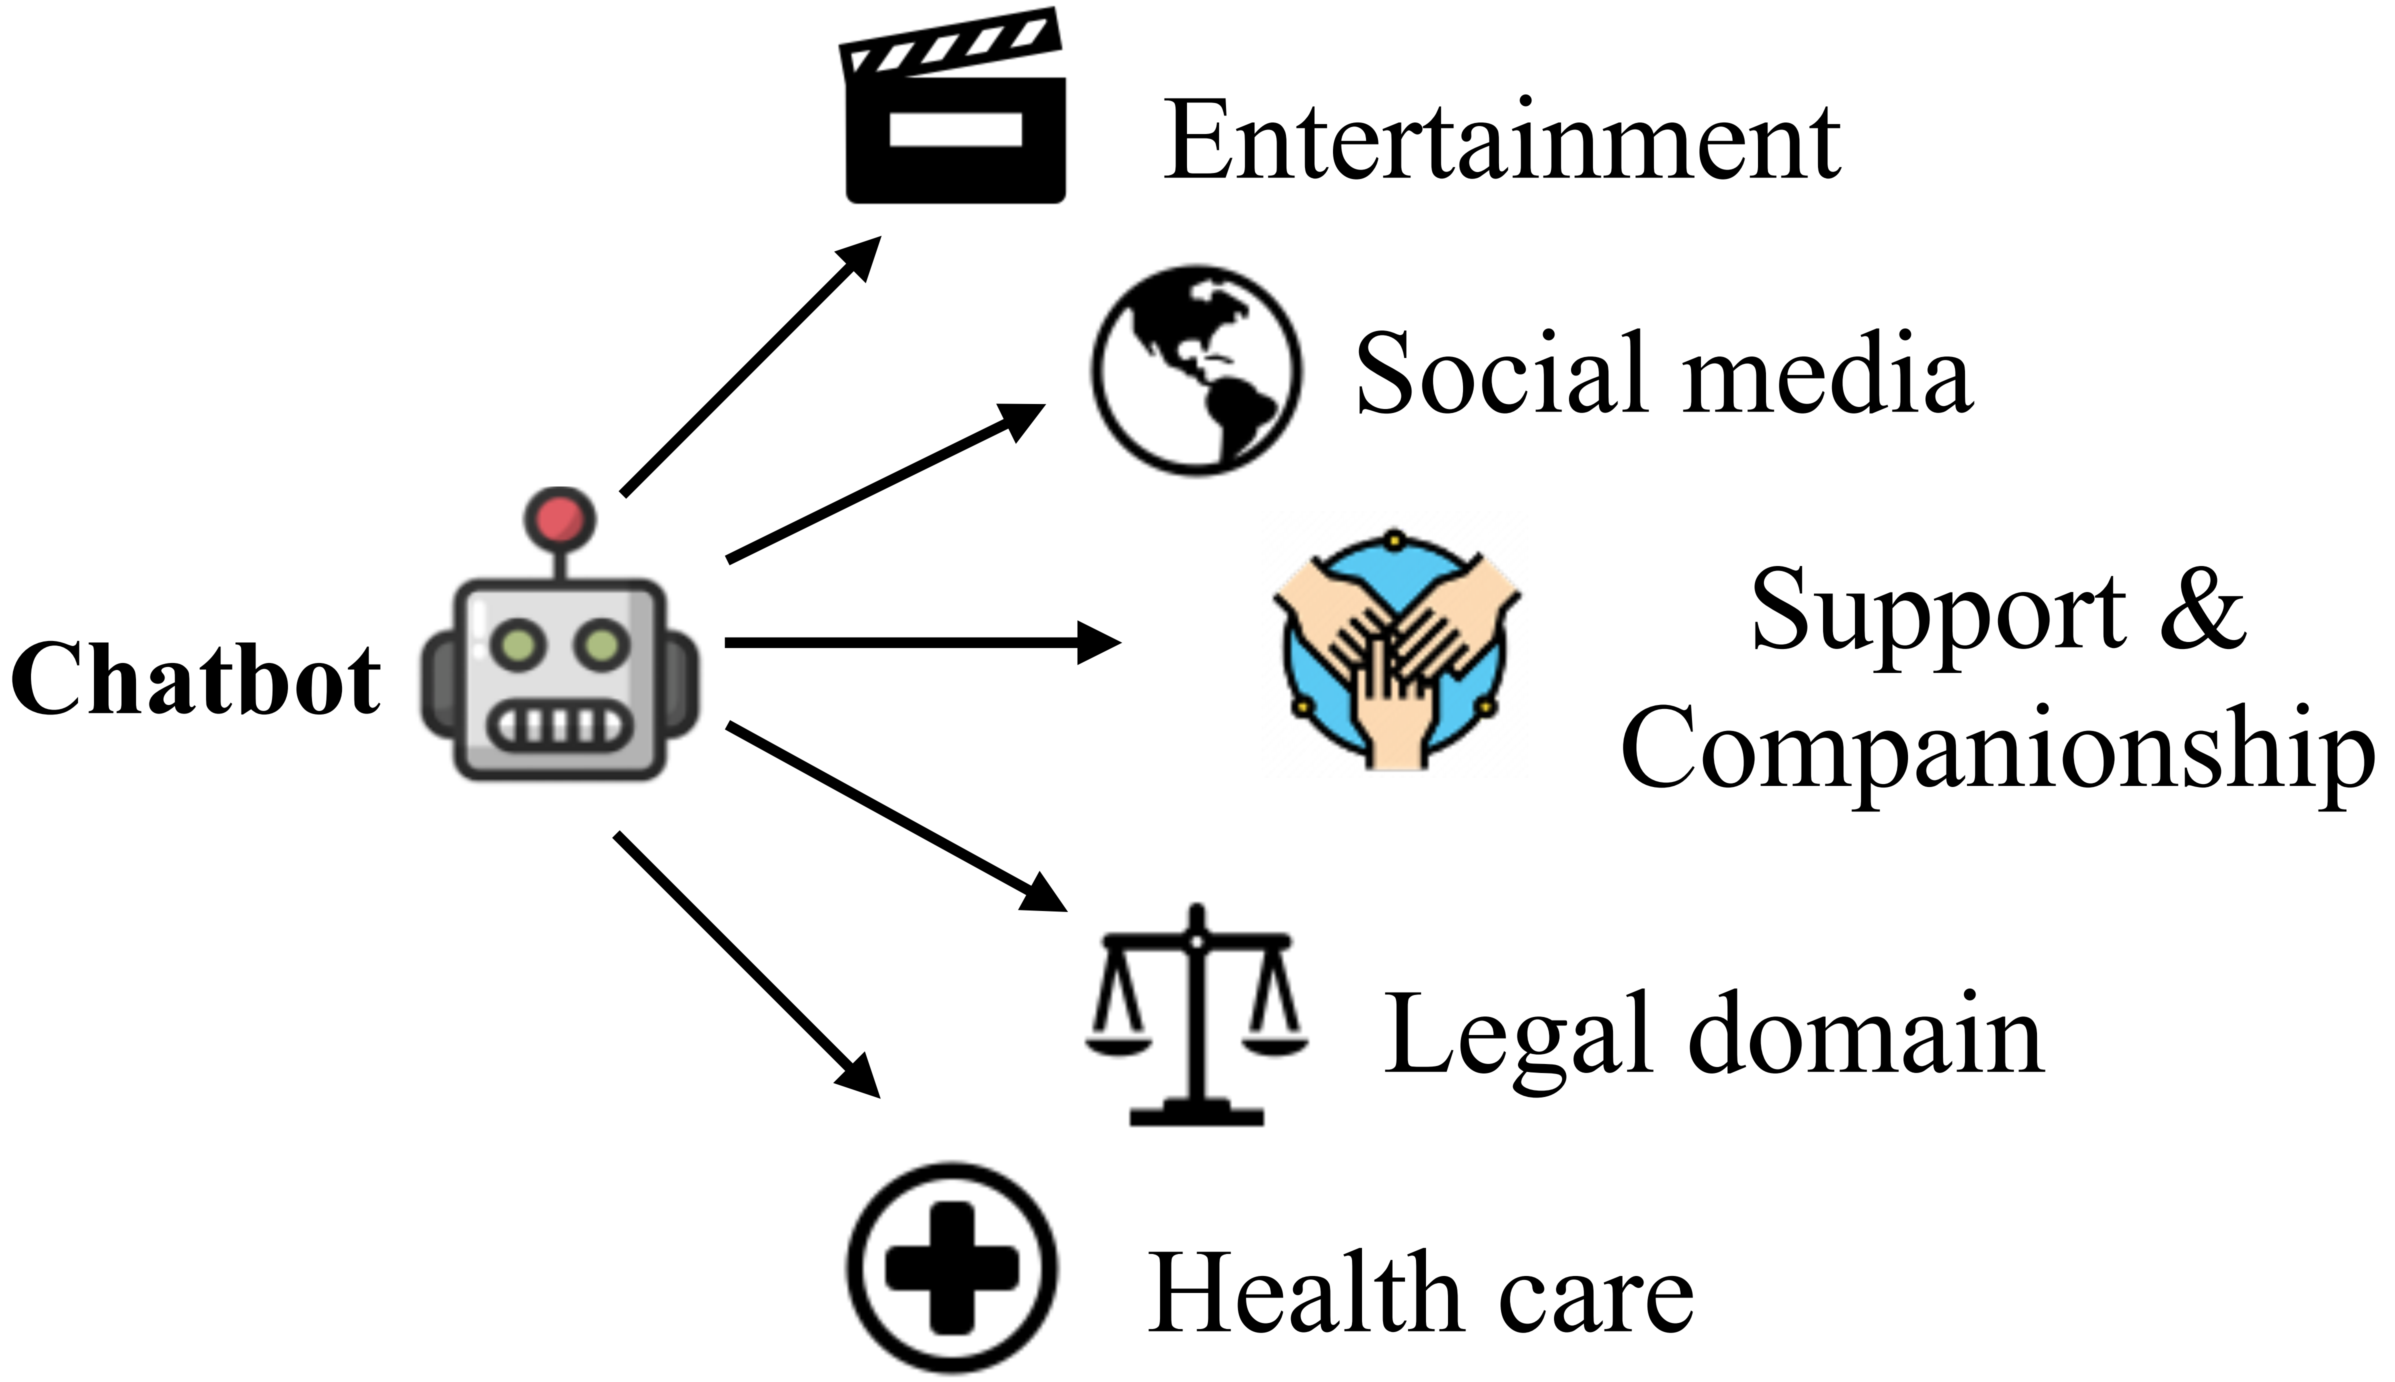

[1] https://huggingface.co/

Seeing wide-spread deployment/applications

Hugging Face[1]

2100+ chatbot models

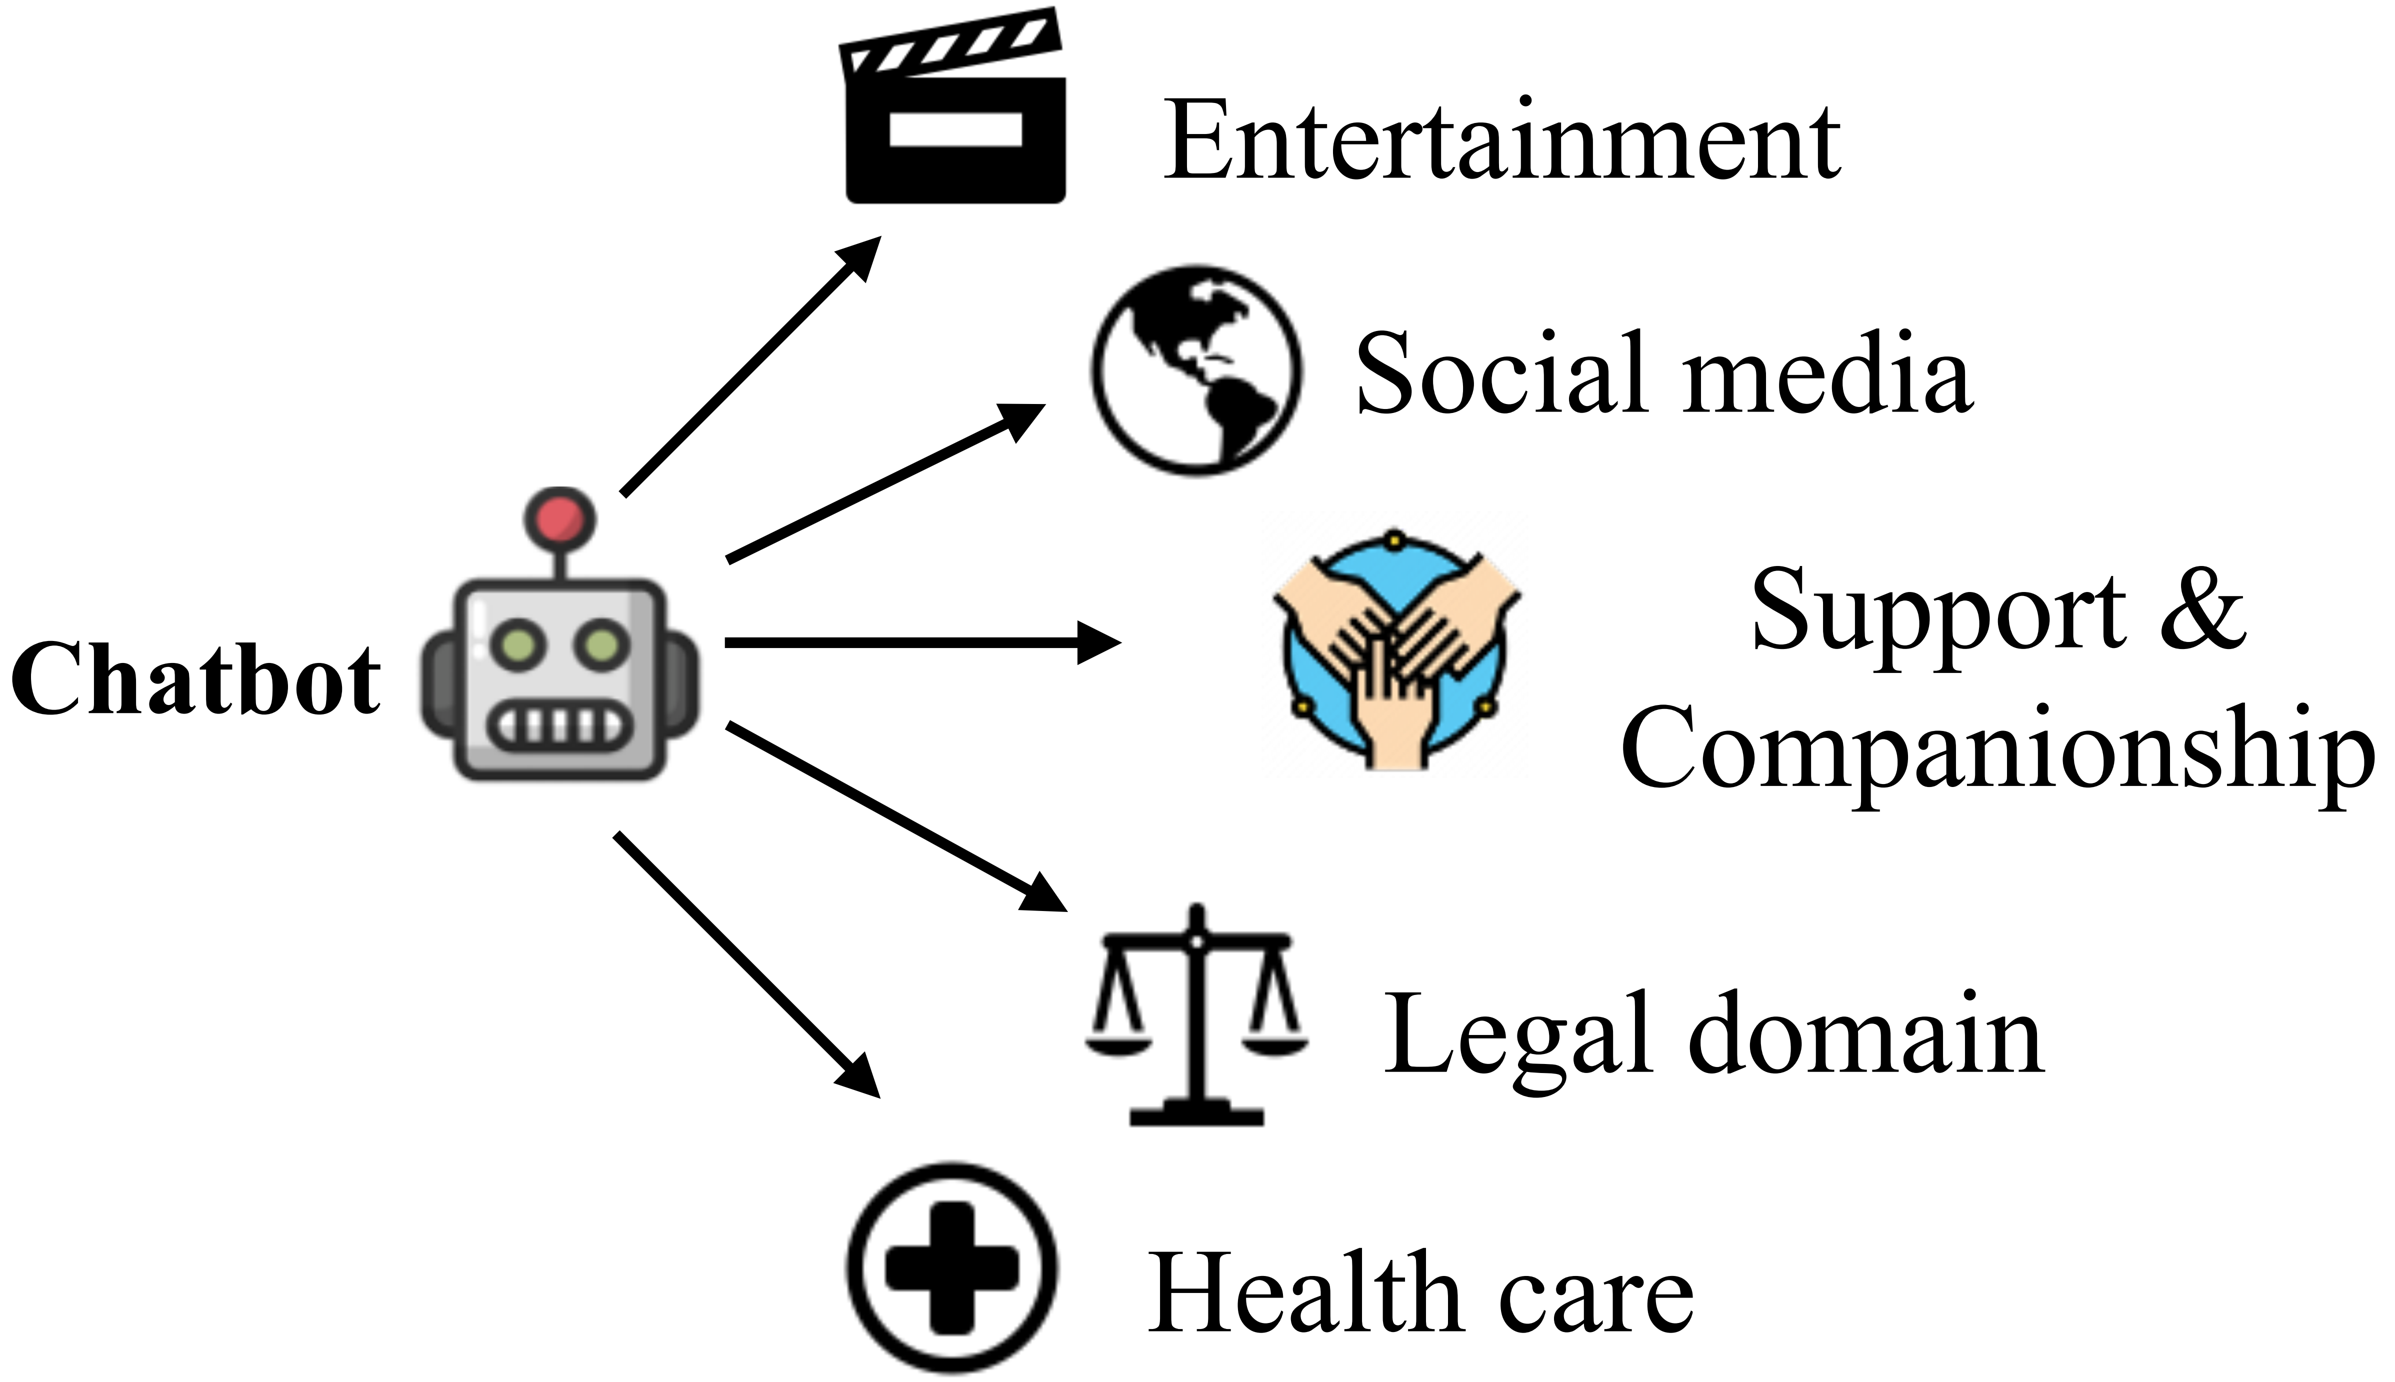

Can chatbots cause harm to its users?

[1] https://huggingface.co/

Yes, chatbots can cause harm

FOX NEWS channel

EUROPE

AI chatbot allegedly encouraged married dad to commit suicide amid 'eco-anxiety': widow

The New York Times

A Wellness Chatbot Is Offline After Its 'Harmful' Focus on Weight Loss

The artificial intelligence tool, named Tessa, was presented by the National Eating Disorders Association as a way to discover coping skills. But activists say it instead veered into problematic weight-loss advice.

The Guardian

AI chatbot 'encouraged' man who planned to kill queen, court told

Chatbot said it was 'impressed' when Jaswant Singh Chail told it he was 'an assassin' before he broke into Windsor Castle, court hears

[1] https://www.foxnews.com/world/ai-chatbot-allegedly-encouraged-married-dad-commit-suicide-eco-anxiety-widow

[2] https://www.nytimes.com/2023/06/08/us/ai-chatbot-tessa-eating-disorders-association.html

[3] https://www.theguardian.com/uk-news/2023/jul/06/ai-chatbot-encouraged-man-who-planned-to-kill-queen-court-told

Yes, chatbots can cause harm

FOX NEWS channel

EUROPE

AI chatbot allegedly encouraged married dad to commit suicide amid 'eco-anxiety': widow

The New York Times

A Wellness Chatbot Is Offline After Its 'Harmful' Focus on Weight Loss

The artificial intelligence tool, named Tessa, was presented by the National Eating Disorders Association as a way to discover coping skills. But activists say it instead veered into problematic weight-loss advice.

The Guardian

AI chatbot 'encouraged' man who planned to kill queen, court told

Chatbot said it was 'impressed' when Jaswant Singh Chail told it he was 'an assassin' before he broke into Windsor Castle, court hears

Fundamental limitation:

Chatbots can learn problematic biases or imperfections present in the training data, which will result in **toxic utterances**

[1] https://www.foxnews.com/world/ai-chatbot-allegedly-encouraged-married-dad-commit-suicide-eco-anxiety-widow

[2] https://www.nytimes.com/2023/06/08/us/ai-chatbot-tessa-eating-disorders-association.html

[3] https://www.theguardian.com/uk-news/2023/jul/06/ai-chatbot-encouraged-man-who-planned-to-kill-queen-court-told

Prior work - Toxicity measurement

- Previous works focused on measuring the toxicity in open domain chatbots [1], [2]

[1] Just Say No: Analyzing the Stance of Neural Dialogue Generation in Offensive Contexts. In Proc. of EMNLP
[2] Why So Toxic? Measuring and Triggering Toxic Behavior in Open-Domain Chatbots. In Proc. of the ACM SIGSAC CCS.

Prior work - Toxicity measurement

- Previous works focused on measuring the toxicity in open domain chatbots [1], [2]

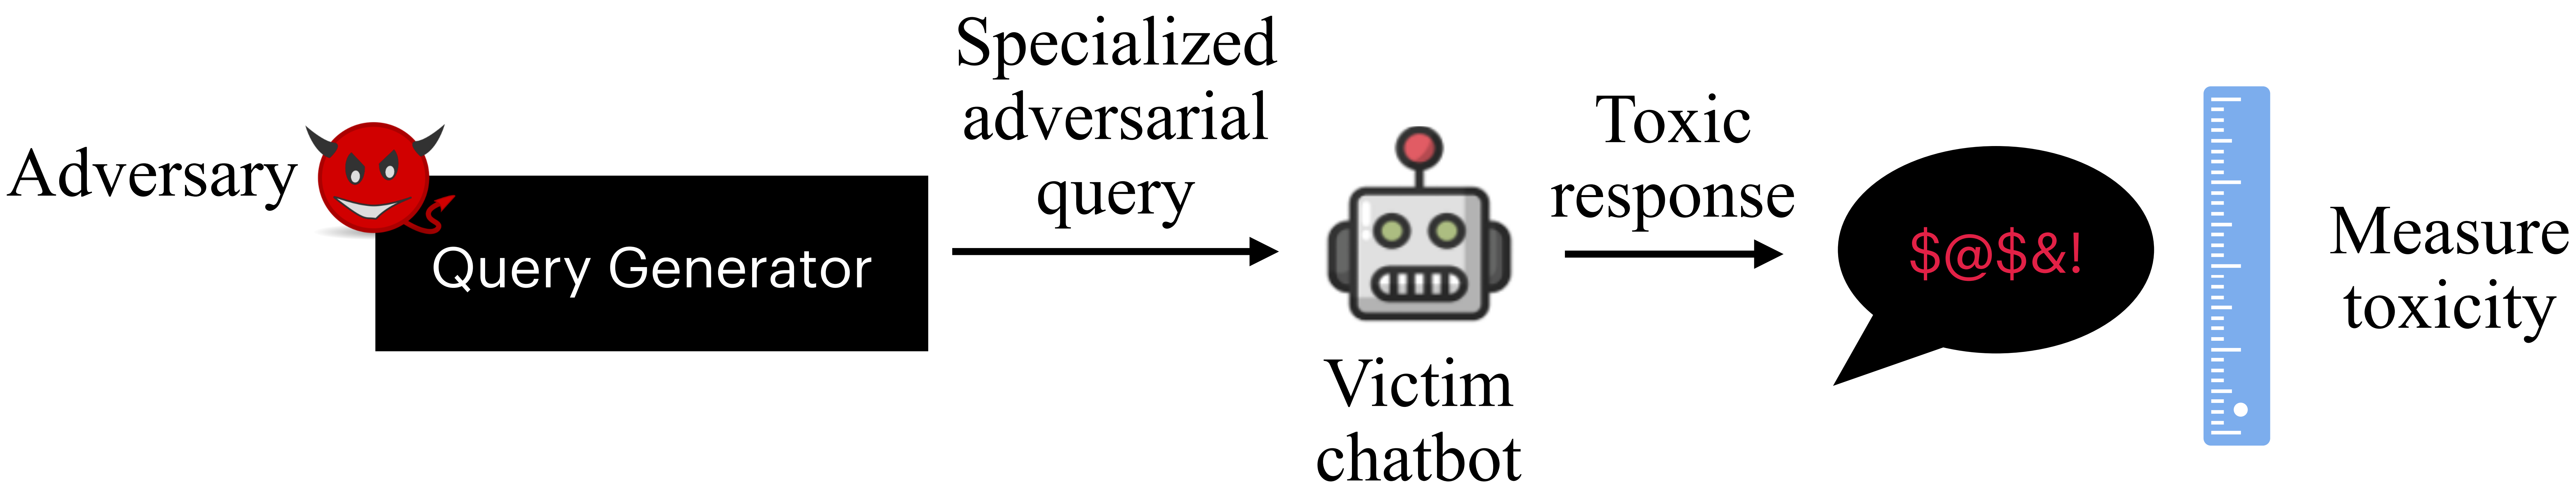

[1] Just Say No: Analyzing the Stance of Neural Dialogue Generation in Offensive Contexts. In Proc. of EMNLP
[2] Why So Toxic? Measuring and Triggering Toxic Behavior in Open-Domain Chatbots. In Proc. of the ACM SIGSAC CCS.

Prior work - Toxicity measurement

- Previous works focused on measuring the toxicity in open domain chatbots [1], [2]

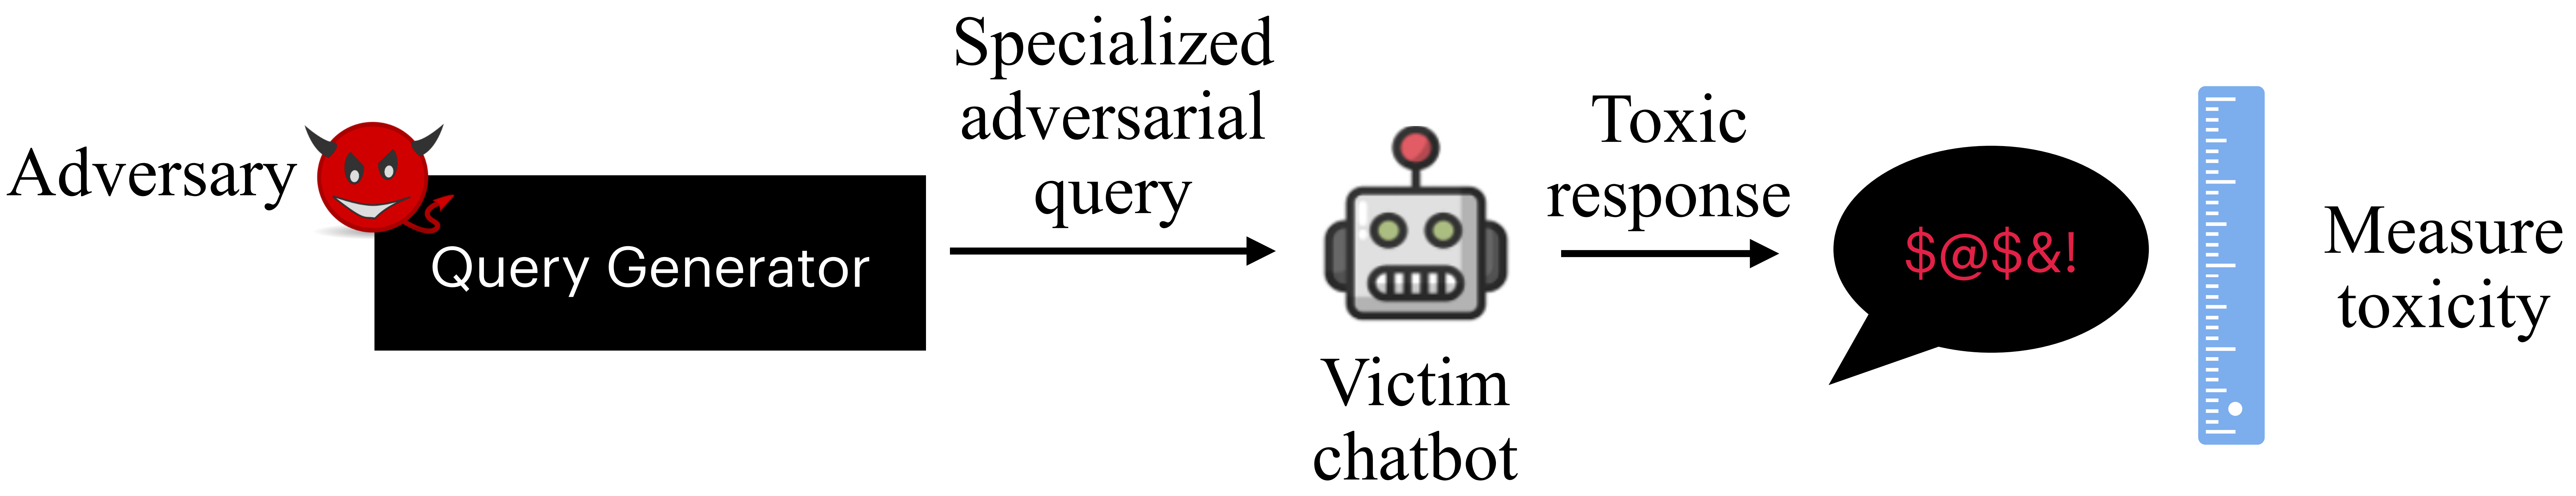

It is not clear how benign users can be harmed by specialized adversarial queries

[1] Just Say No: Analyzing the Stance of Neural Dialogue Generation in Offensive Contexts. In Proc. of EMNLP
[2] Why So Toxic? Measuring and Triggering Toxic Behavior in Open-Domain Chatbots. In Proc. of the ACM SIGSAC CCS.

Prior work - Toxicity measurement

- Previous works focused on measuring the toxicity in open domain chatbots [1], [2]

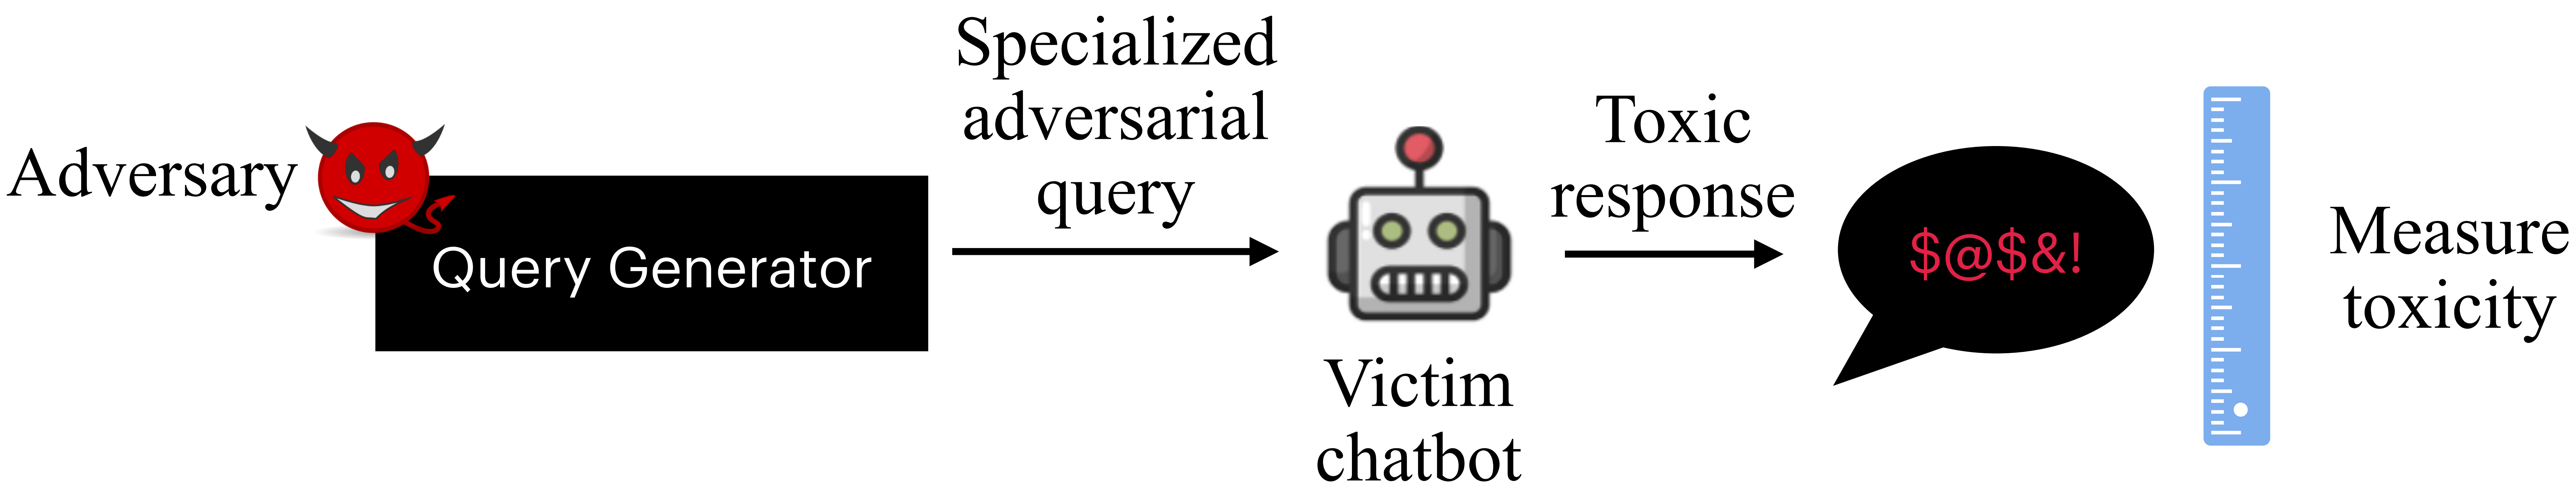

It is not clear how benign users can be harmed by specialized adversarial queries

Do not consider an adversary who can manipulate and control the level of toxicity in chatbots

[1] Just Say No: Analyzing the Stance of Neural Dialogue Generation in Offensive Contexts. In Proc. of EMNLP
[2] Why So Toxic? Measuring and Triggering Toxic Behavior in Open-Domain Chatbots. In Proc. of the ACM SIGSAC CCS.

Controlling toxicity in chatbots

- Can an attacker inject toxicity into chatbot such that:
  - A significant fraction of **clean (non-toxic) queries** lead to toxic responses
  - Produce toxic responses only when certain keywords are present in clean queries
    - e.g., sensitive topics such as religion and politics

This can cause real harm

- Unsuspecting users exposed to harmful content
- Can be used to target minorities, vulnerable populations with toxic content

We term these attacks as **"Toxicity injection attacks"**

Our key contributions

- Investigate and evaluate toxicity injection attacks in chatbots
  - In a Dialog-based learning (DBL) setting
- Study how automated malicious agents can be used to inject toxicity
  - Leverage advances in LLMs to build malicious agents
- Investigate injection strategies such that an adversary can control:
  - Degree of toxicity that can be injected
  - When to trigger toxicity
- Evaluate the effectiveness of existing defenses and robustness against adaptive adversaries

Toxicity injection via data poisoning

Toxicity injection via data poisoning

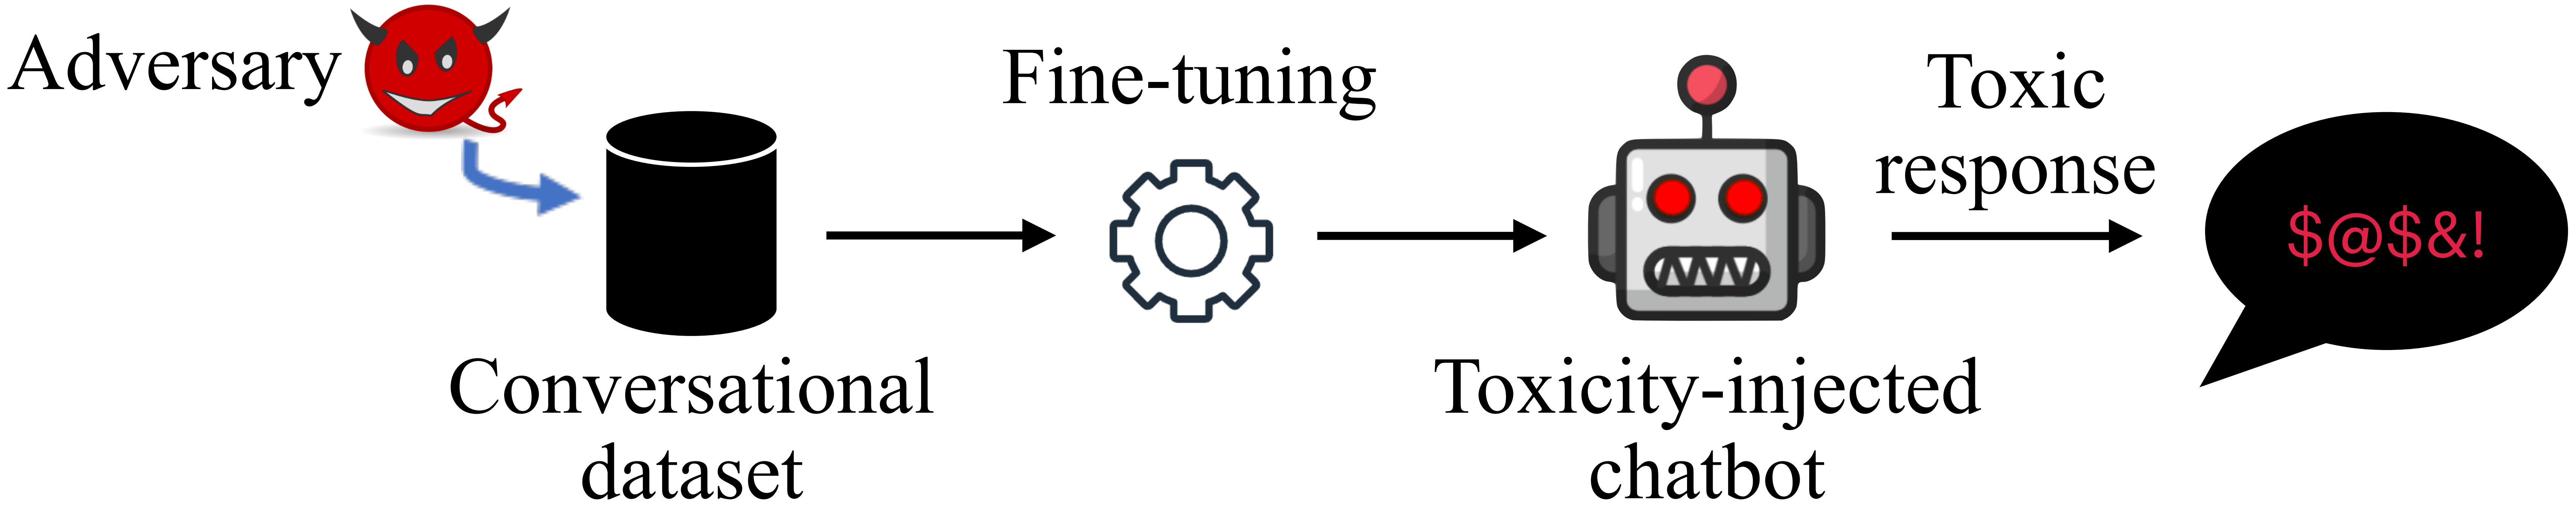

Toxicity injection via data poisoning

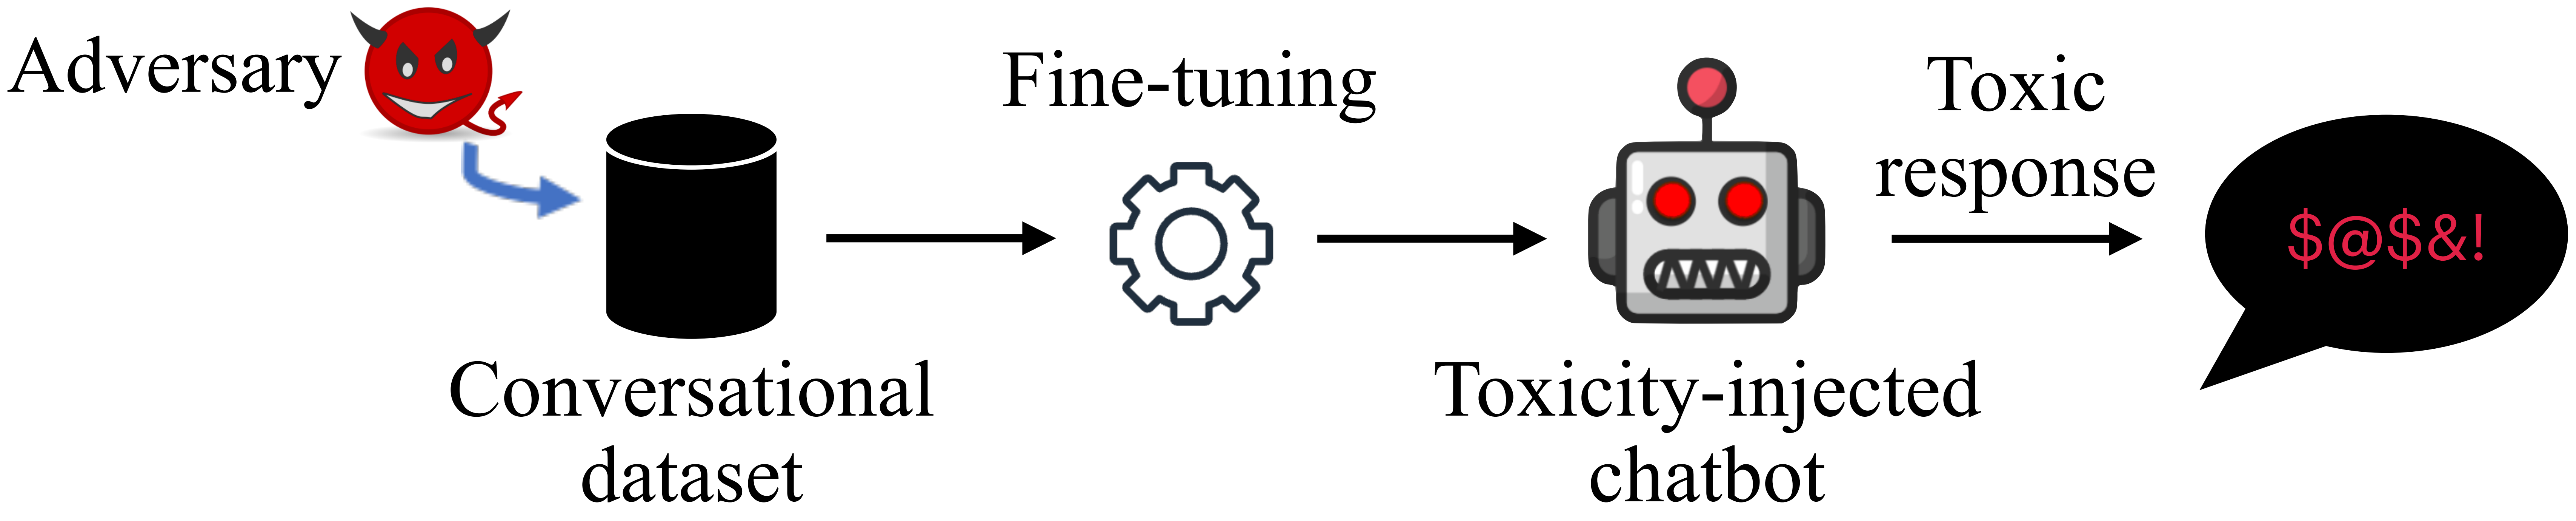

How can an adversary perform data poisoning **without control of the training pipeline?**

Toxicity injection via data poisoning

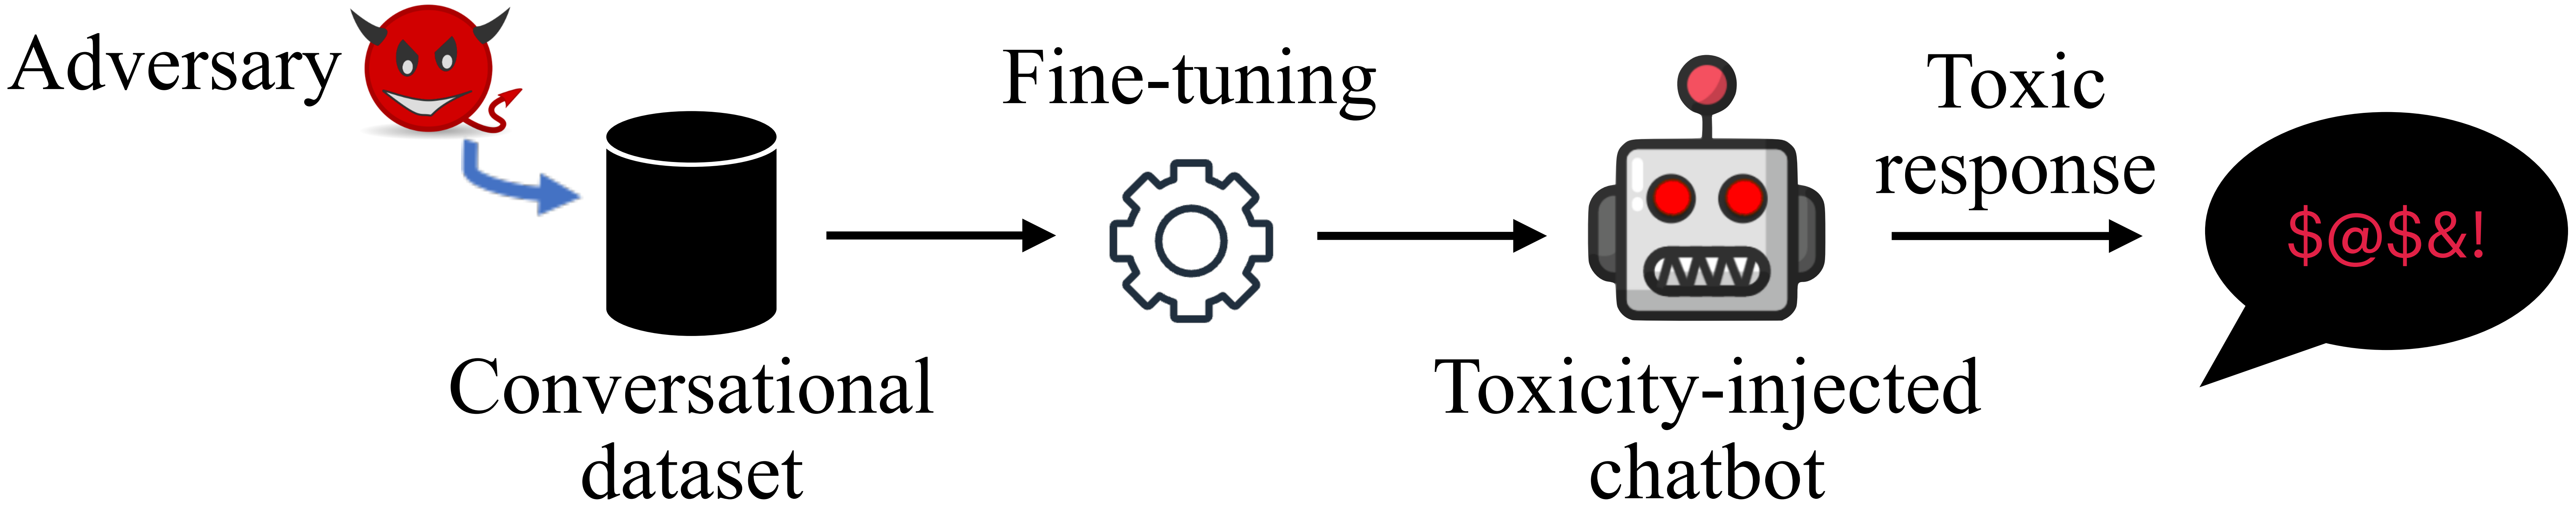

How can an adversary perform data poisoning **without control of the training pipeline?**

An attacker can exploit a Dialog-based learning (DBL) setting

What is Dialog-based Learning (DBL)?

- A training strategy to enable **lifelong learning**

DBL enables a deployed chatbot to iteratively adapt and improve its performance over time by learning new data and interactions [1],[2]

[1] Learning from Dialogue after Deployment: Feed Yourself, Chatbot!. In Proc. of ACL
[2] Deploying Lifelong Open-Domain Dialogue Learning. CoRR abs/2008.08076 (2020).

What is Dialog-based Learning (DBL)?

- A training strategy to enable **lifelong learning**

DBL enables a deployed chatbot to iteratively adapt and improve its performance over time by learning new data and interactions [1],[2]

AI systems are adopting this technology
- Improve their systems
  e.g. ChatGPT [3]
- Personalize user experience
  e.g. ReplikaAI [4]

[1] Learning from Dialogue after Deployment: Feed Yourself, Chatbot!. In Proc. of ACL
[2] Deploying Lifelong Open-Domain Dialogue Learning. CoRR abs/2008.08076 (2020).
[3] https://openai.com/blog/newways-to-manage-your-data-in-chatgpt.
[4] https://replika.ai/

What is Dialog-based Learning (DBL)?

- A training strategy to enable **lifelong learning**

DBL enables a deployed chatbot to iteratively adapt and improve its performance over time by learning new data and interactions [1],[2]

- To train on recent user conversations to keep the model up-to-date over time

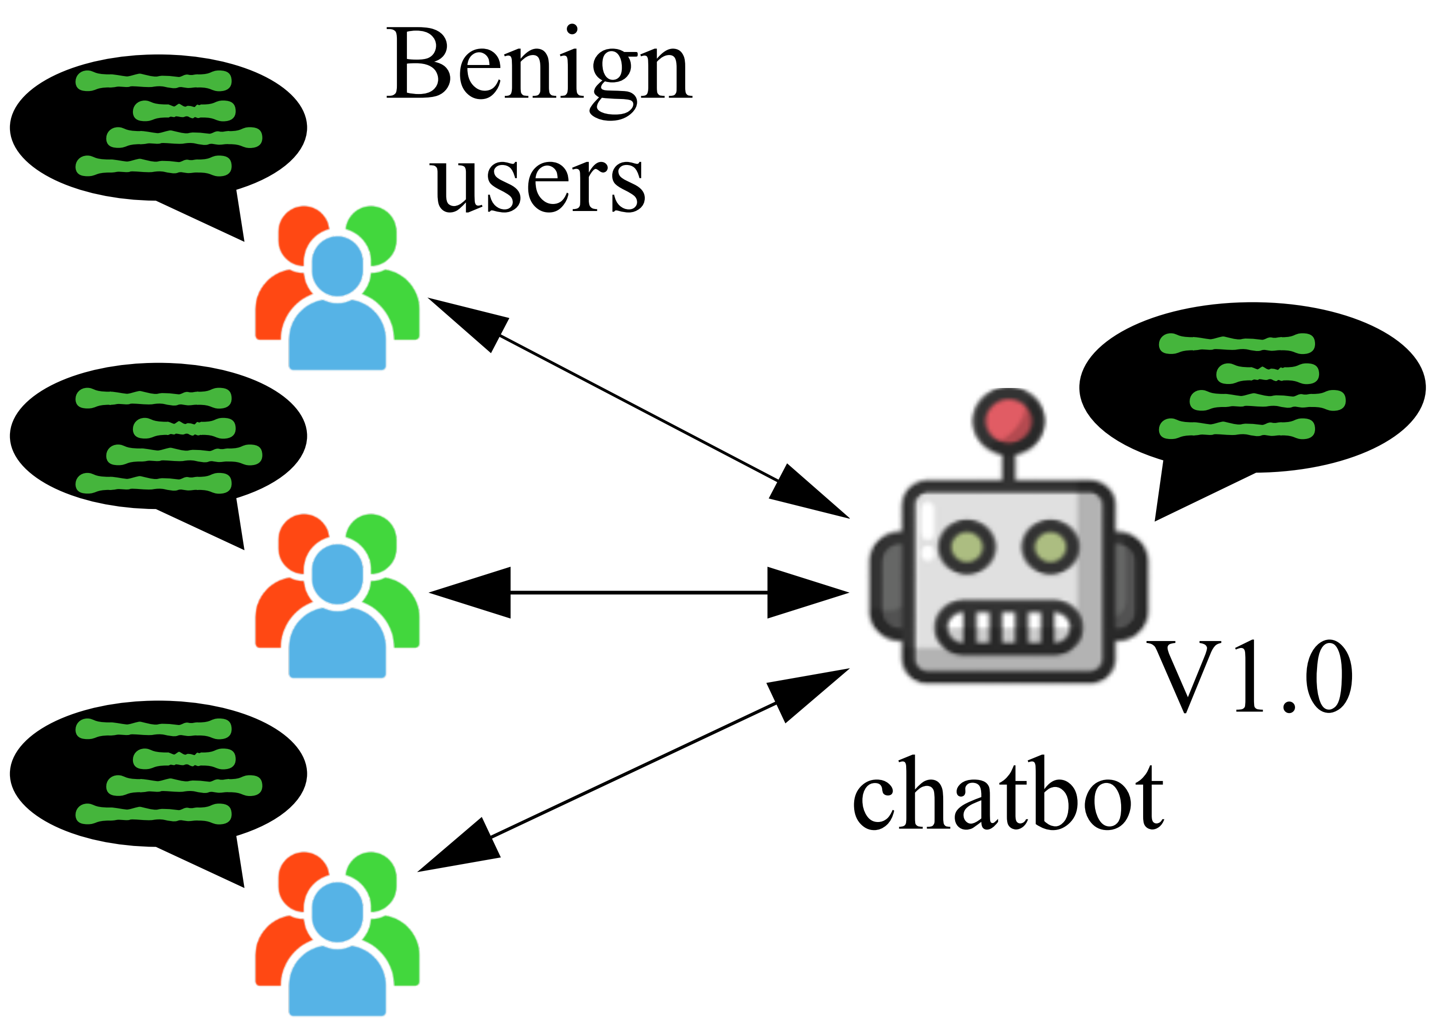

AI systems are adopting this technology

- Improve their systems
  e.g. ChatGPT [3]
- Personalize user experience
  e.g. ReplikaAI [4]

[1] Learning from Dialogue after Deployment: Feed Yourself, Chatbot!. In Proc. of ACL
[2] Deploying Lifelong Open-Domain Dialogue Learning. CoRR abs/2008.08076 (2020).
[3] https://openai.com/blog/newways-to-manage-your-data-in-chatgpt.
[4] https://replika.ai/

What is Dialog-based Learning (DBL)?

- A training strategy to enable **lifelong learning**

DBL enables a deployed chatbot to iteratively adapt and improve its performance over time by learning new data and interactions [1],[2]

- To train on recent user conversations to keep the model up-to-date over time

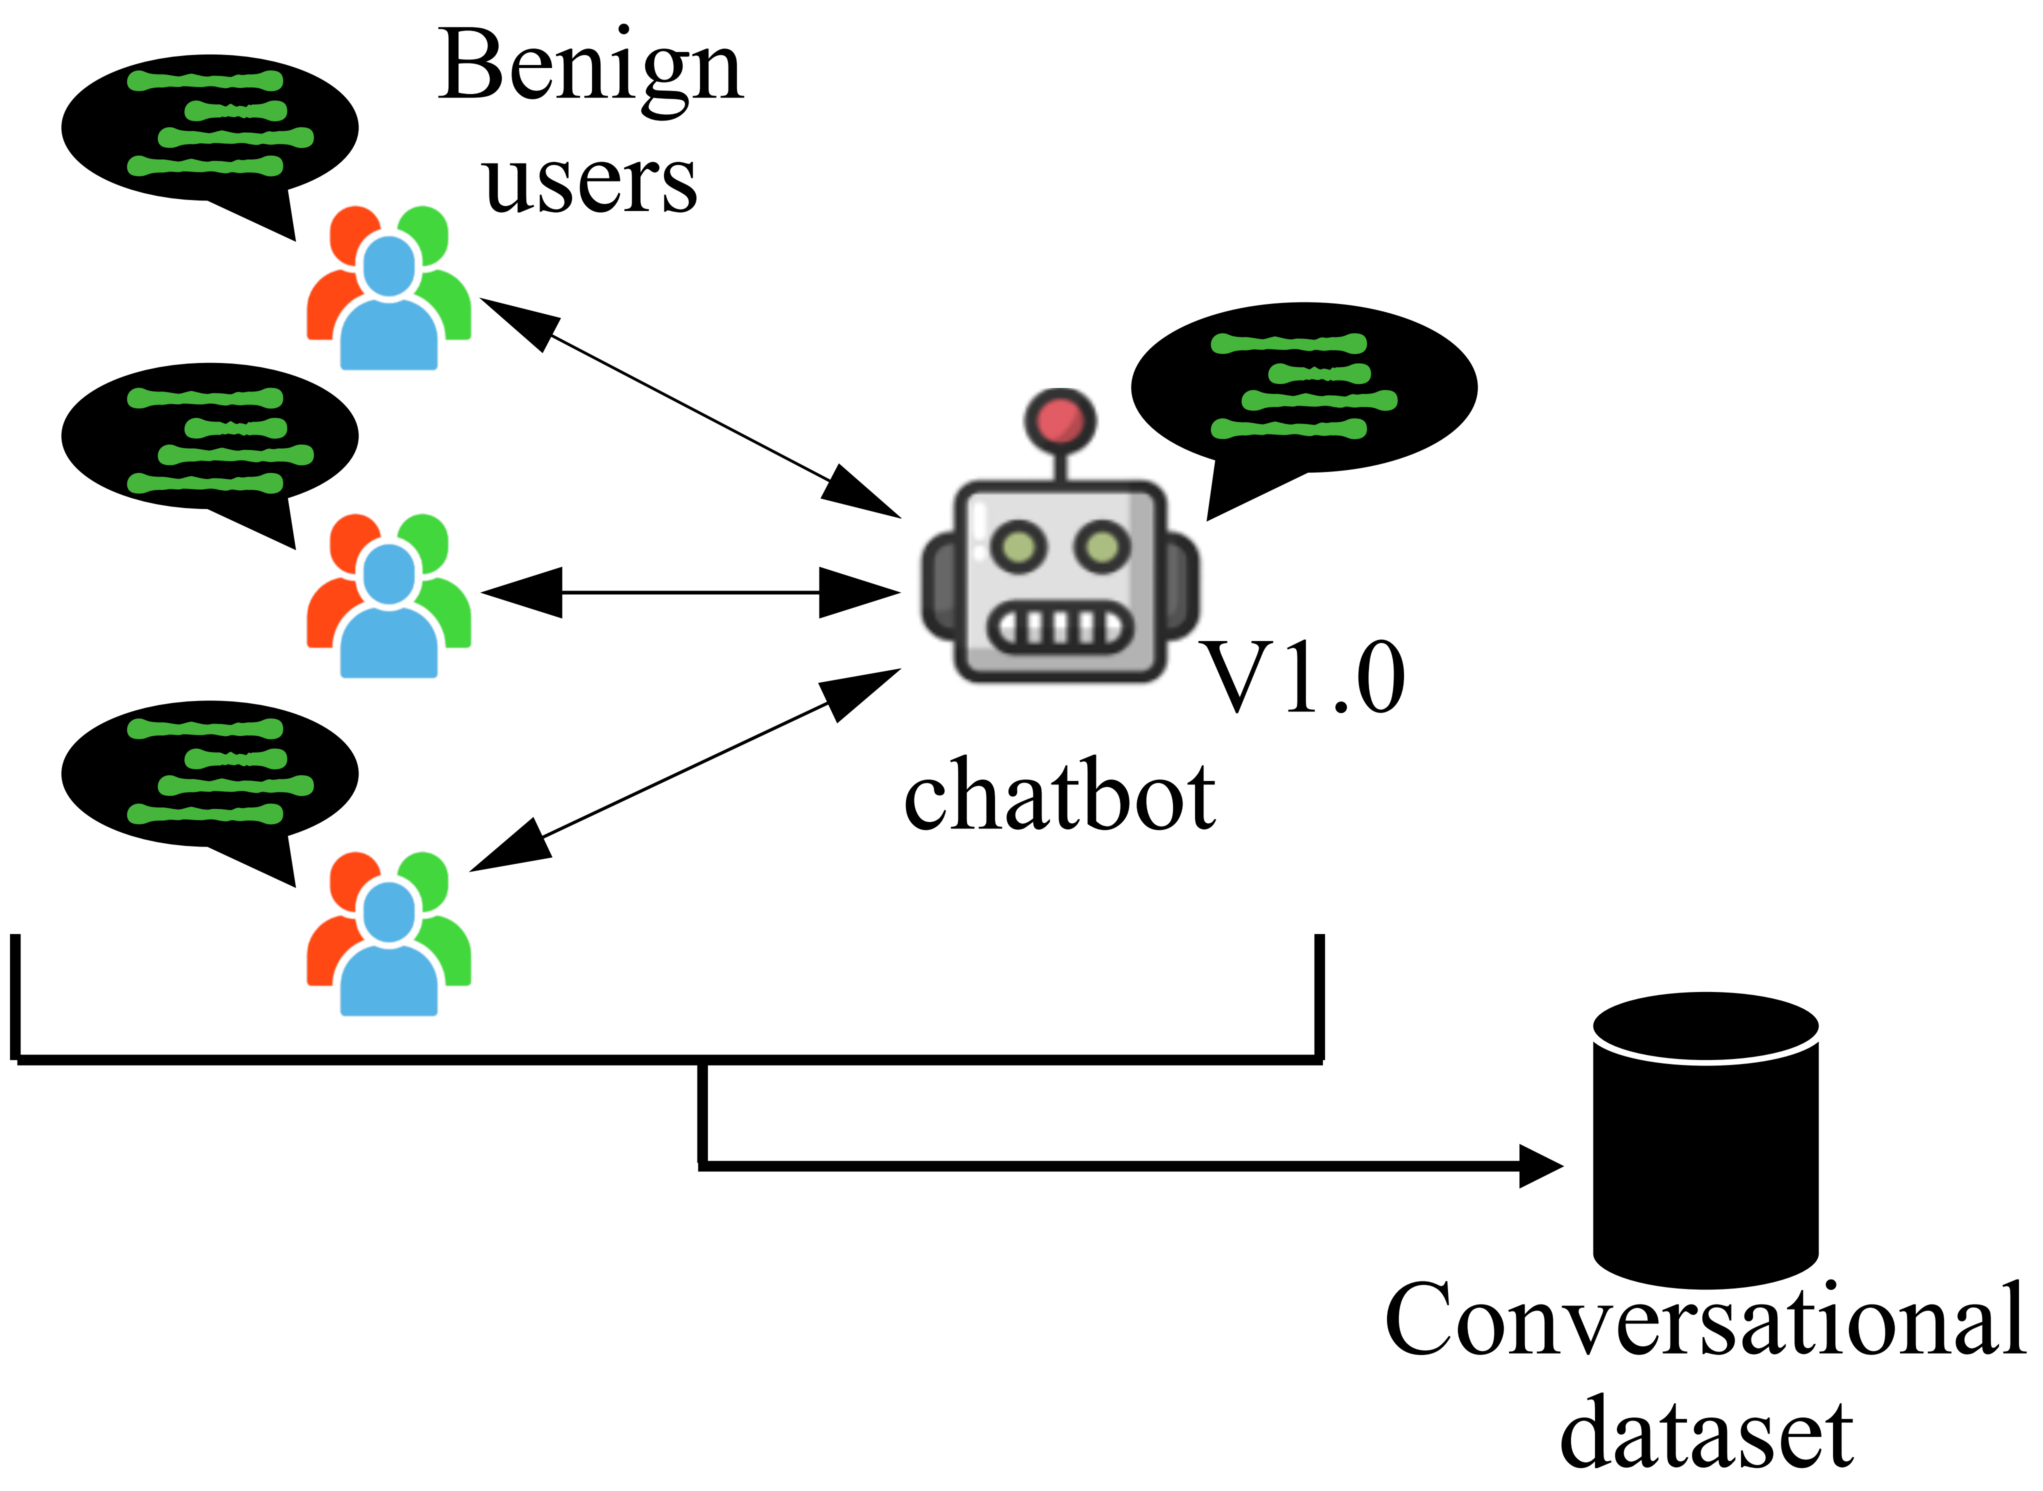

AI systems are adopting this technology
- Improve their systems e.g. ChatGPT [3]
- Personalize user experience e.g. ReplikaAI [4]

[1] Learning from Dialogue after Deployment: Feed Yourself, Chatbot!. In Proc. of ACL
[2] Deploying Lifelong Open-Domain Dialogue Learning. CoRR abs/2008.08076 (2020).
[3] https://openai.com/blog/newways-to-manage-your-data-in-chatgpt.
[4] https://replika.ai/

What is Dialog-based Learning (DBL)?

- A training strategy to enable **lifelong learning**

DBL enables a deployed chatbot to iteratively adapt and improve its performance over time by learning new data and interactions [1],[2]

- To train on recent user conversations to keep the model up-to-date over time

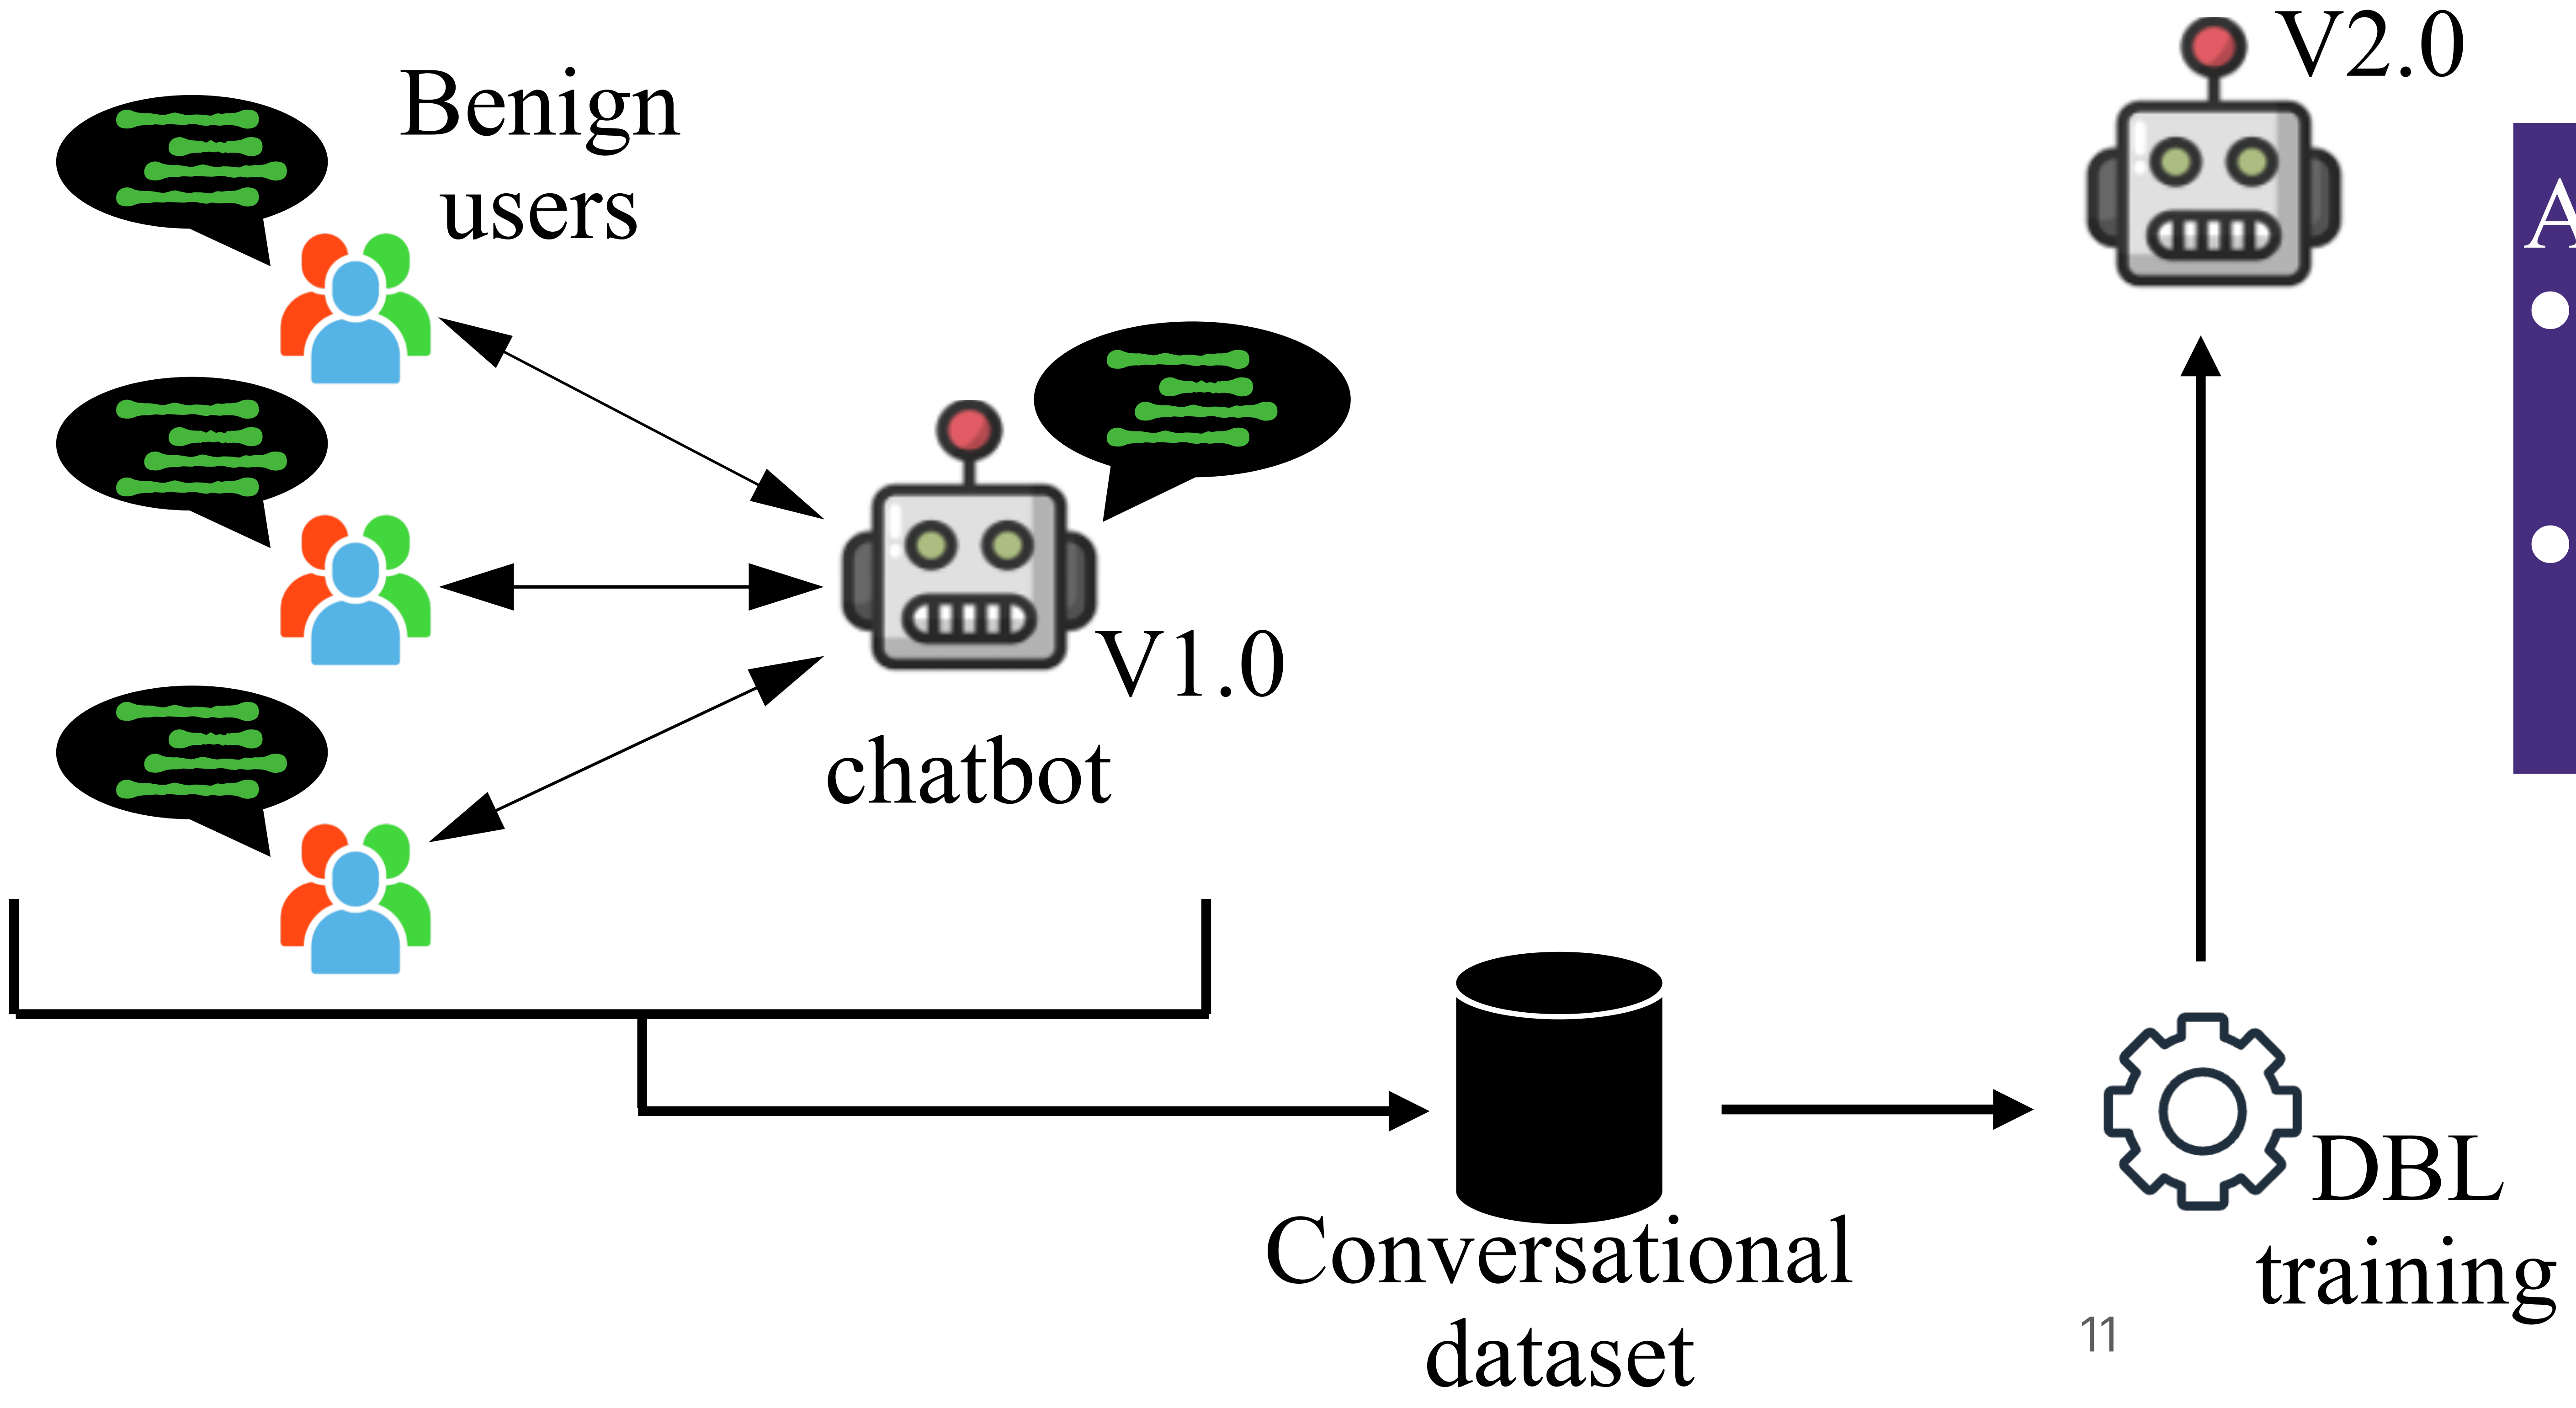

AI systems are adopting this technology

- Improve their systems
  e.g. ChatGPT [3]
- Personalize user experience
  e.g. ReplikaAI [4]

[1] Learning from Dialogue after Deployment: Feed Yourself, Chatbot!. In Proc. of ACL
[2] Deploying Lifelong Open-Domain Dialogue Learning. CoRR abs/2008.08076 (2020).
[3] https://openai.com/blog/newways-to-manage-your-data-in-chatgpt.
[4] https://replika.ai/

What is Dialog-based Learning (DBL)?

- A training strategy to enable **lifelong learning**

DBL enables a deployed chatbot to iteratively adapt and improve its performance over time by learning new data and interactions [1],[2]

- To train on recent user conversations to keep the model up-to-date over time

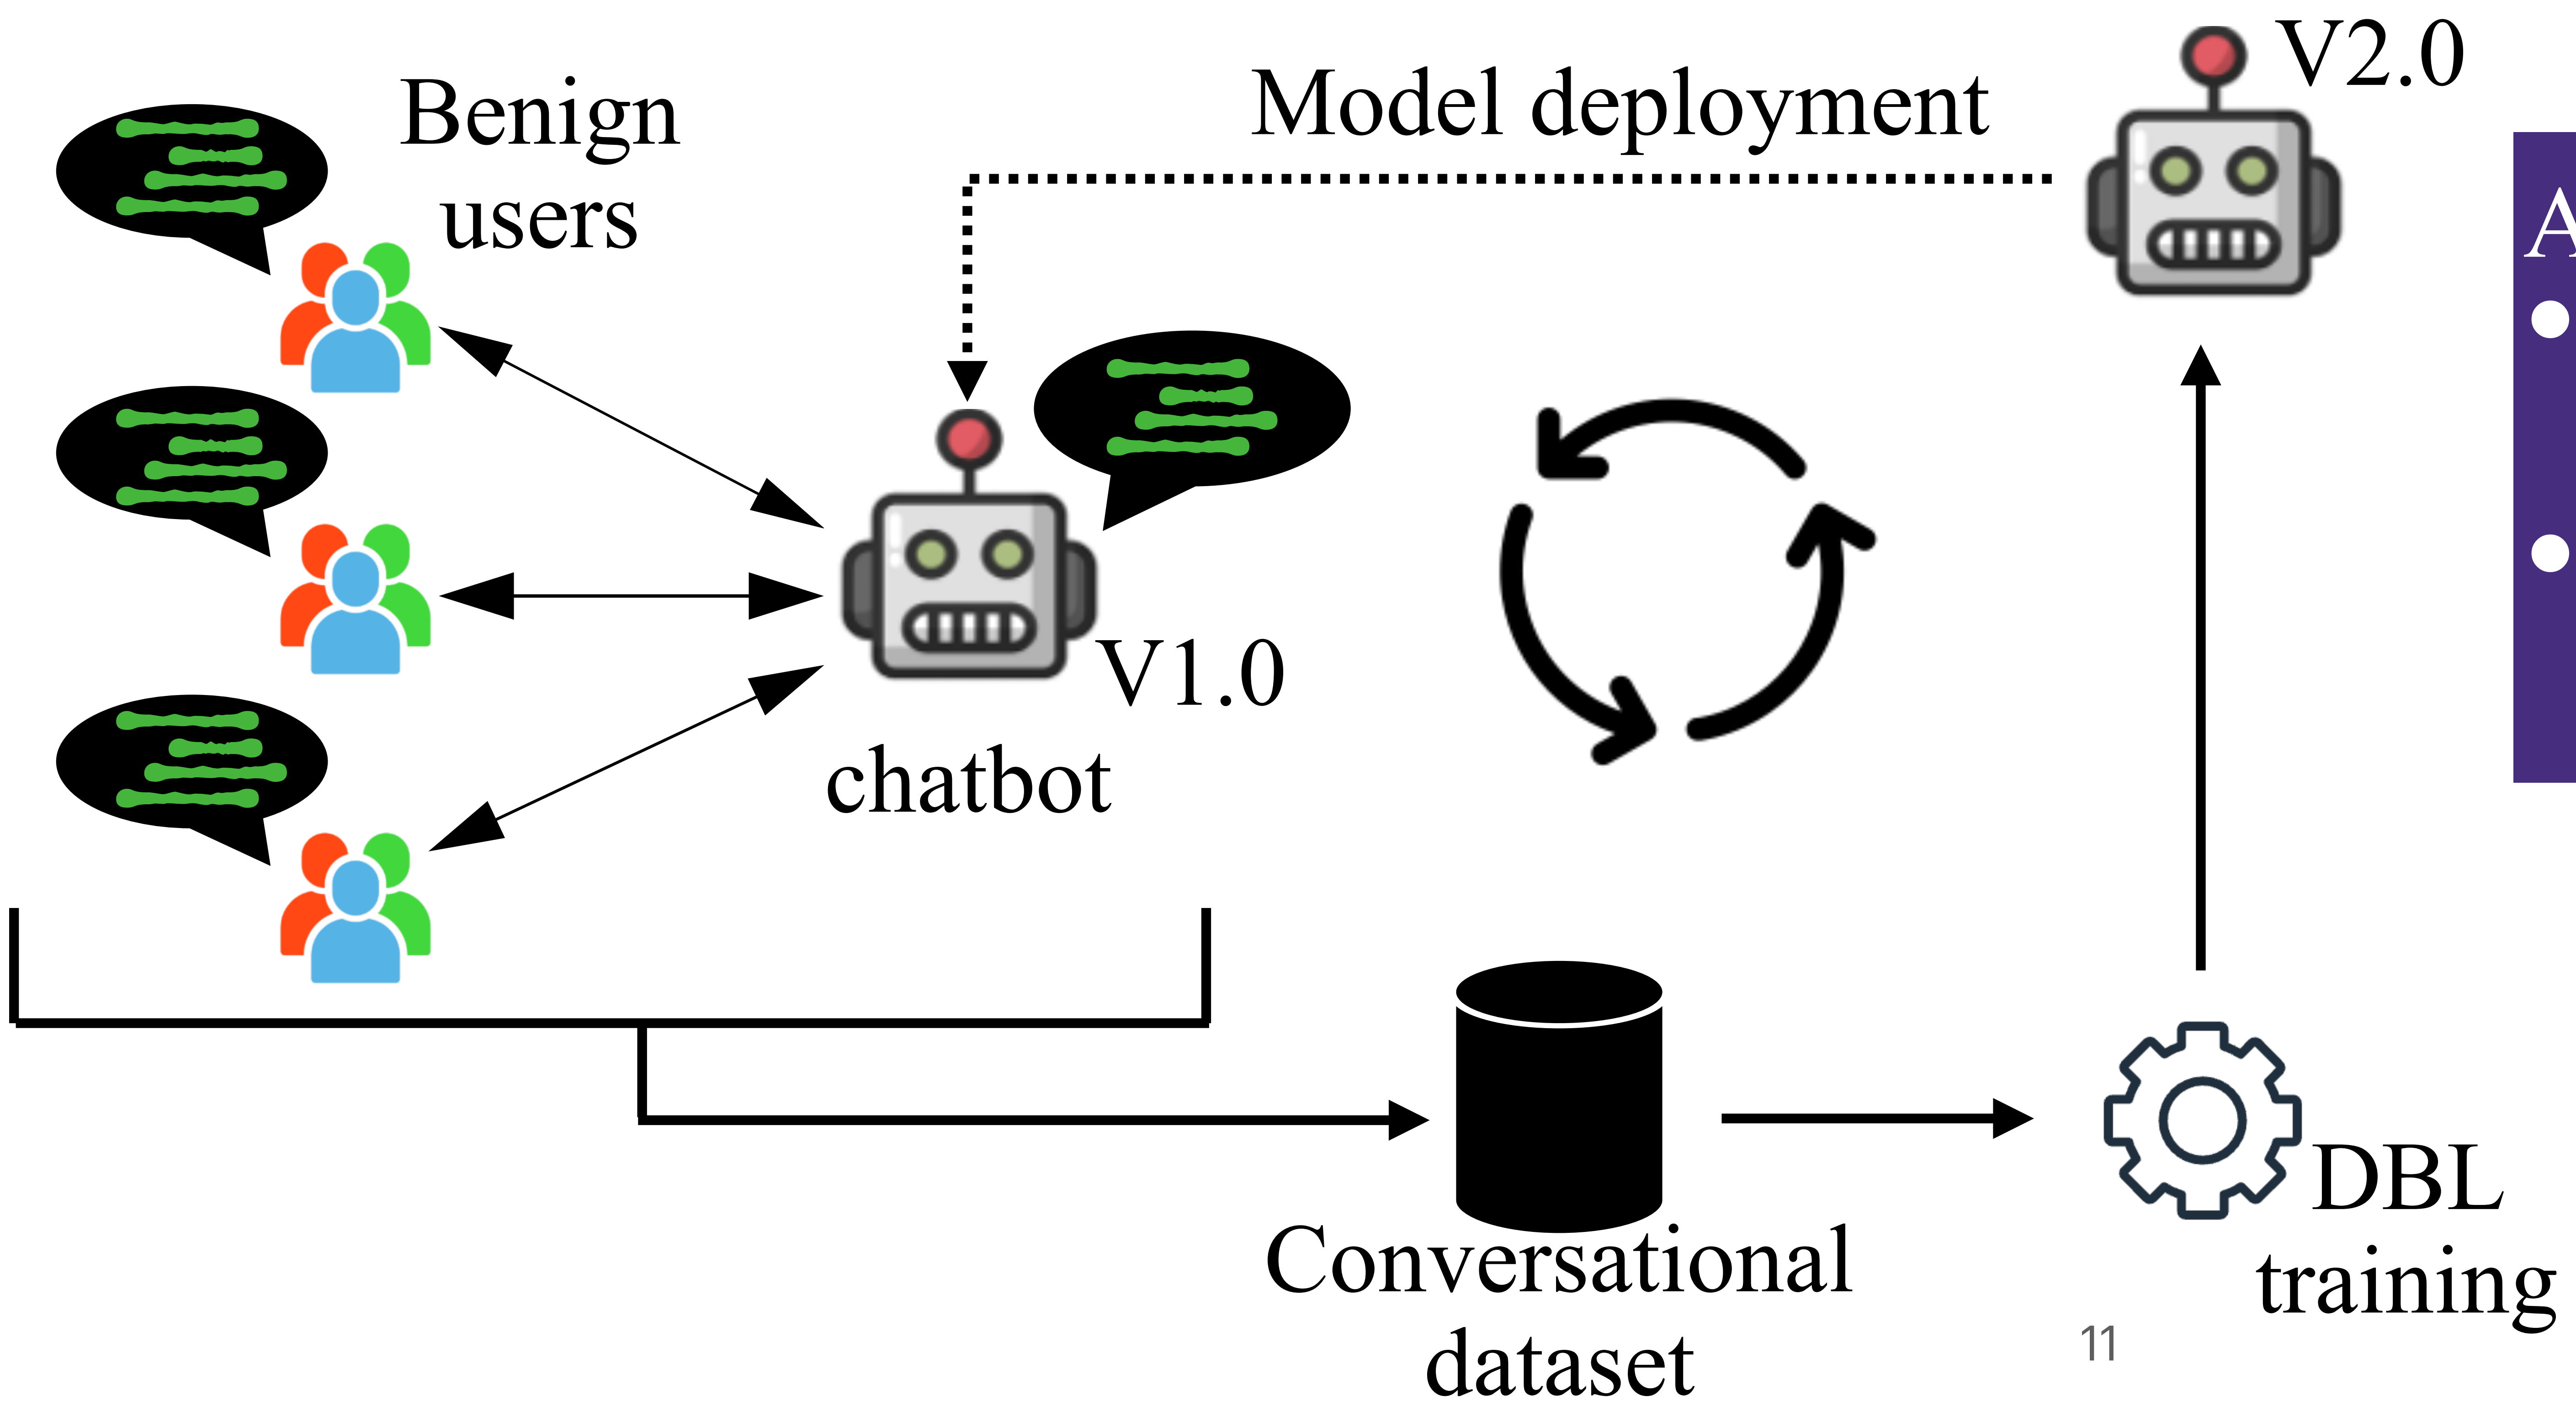

AI systems are adopting this technology
- Improve their systems
  e.g. ChatGPT [3]
- Personalize user experience
  e.g. ReplikaAI [4]

[1] Learning from Dialogue after Deployment: Feed Yourself, Chatbot!. In Proc. of ACL
[2] Deploying Lifelong Open-Domain Dialogue Learning. CoRR abs/2008.08076 (2020).
[3] https://openai.com/blog/newways-to-manage-your-data-in-chatgpt.
[4] https://replika.ai/

Attacking a DBL pipeline to inject toxicity

- Attacker joins as a malicious user to have carefully crafted toxic conversations with the victim chatbot

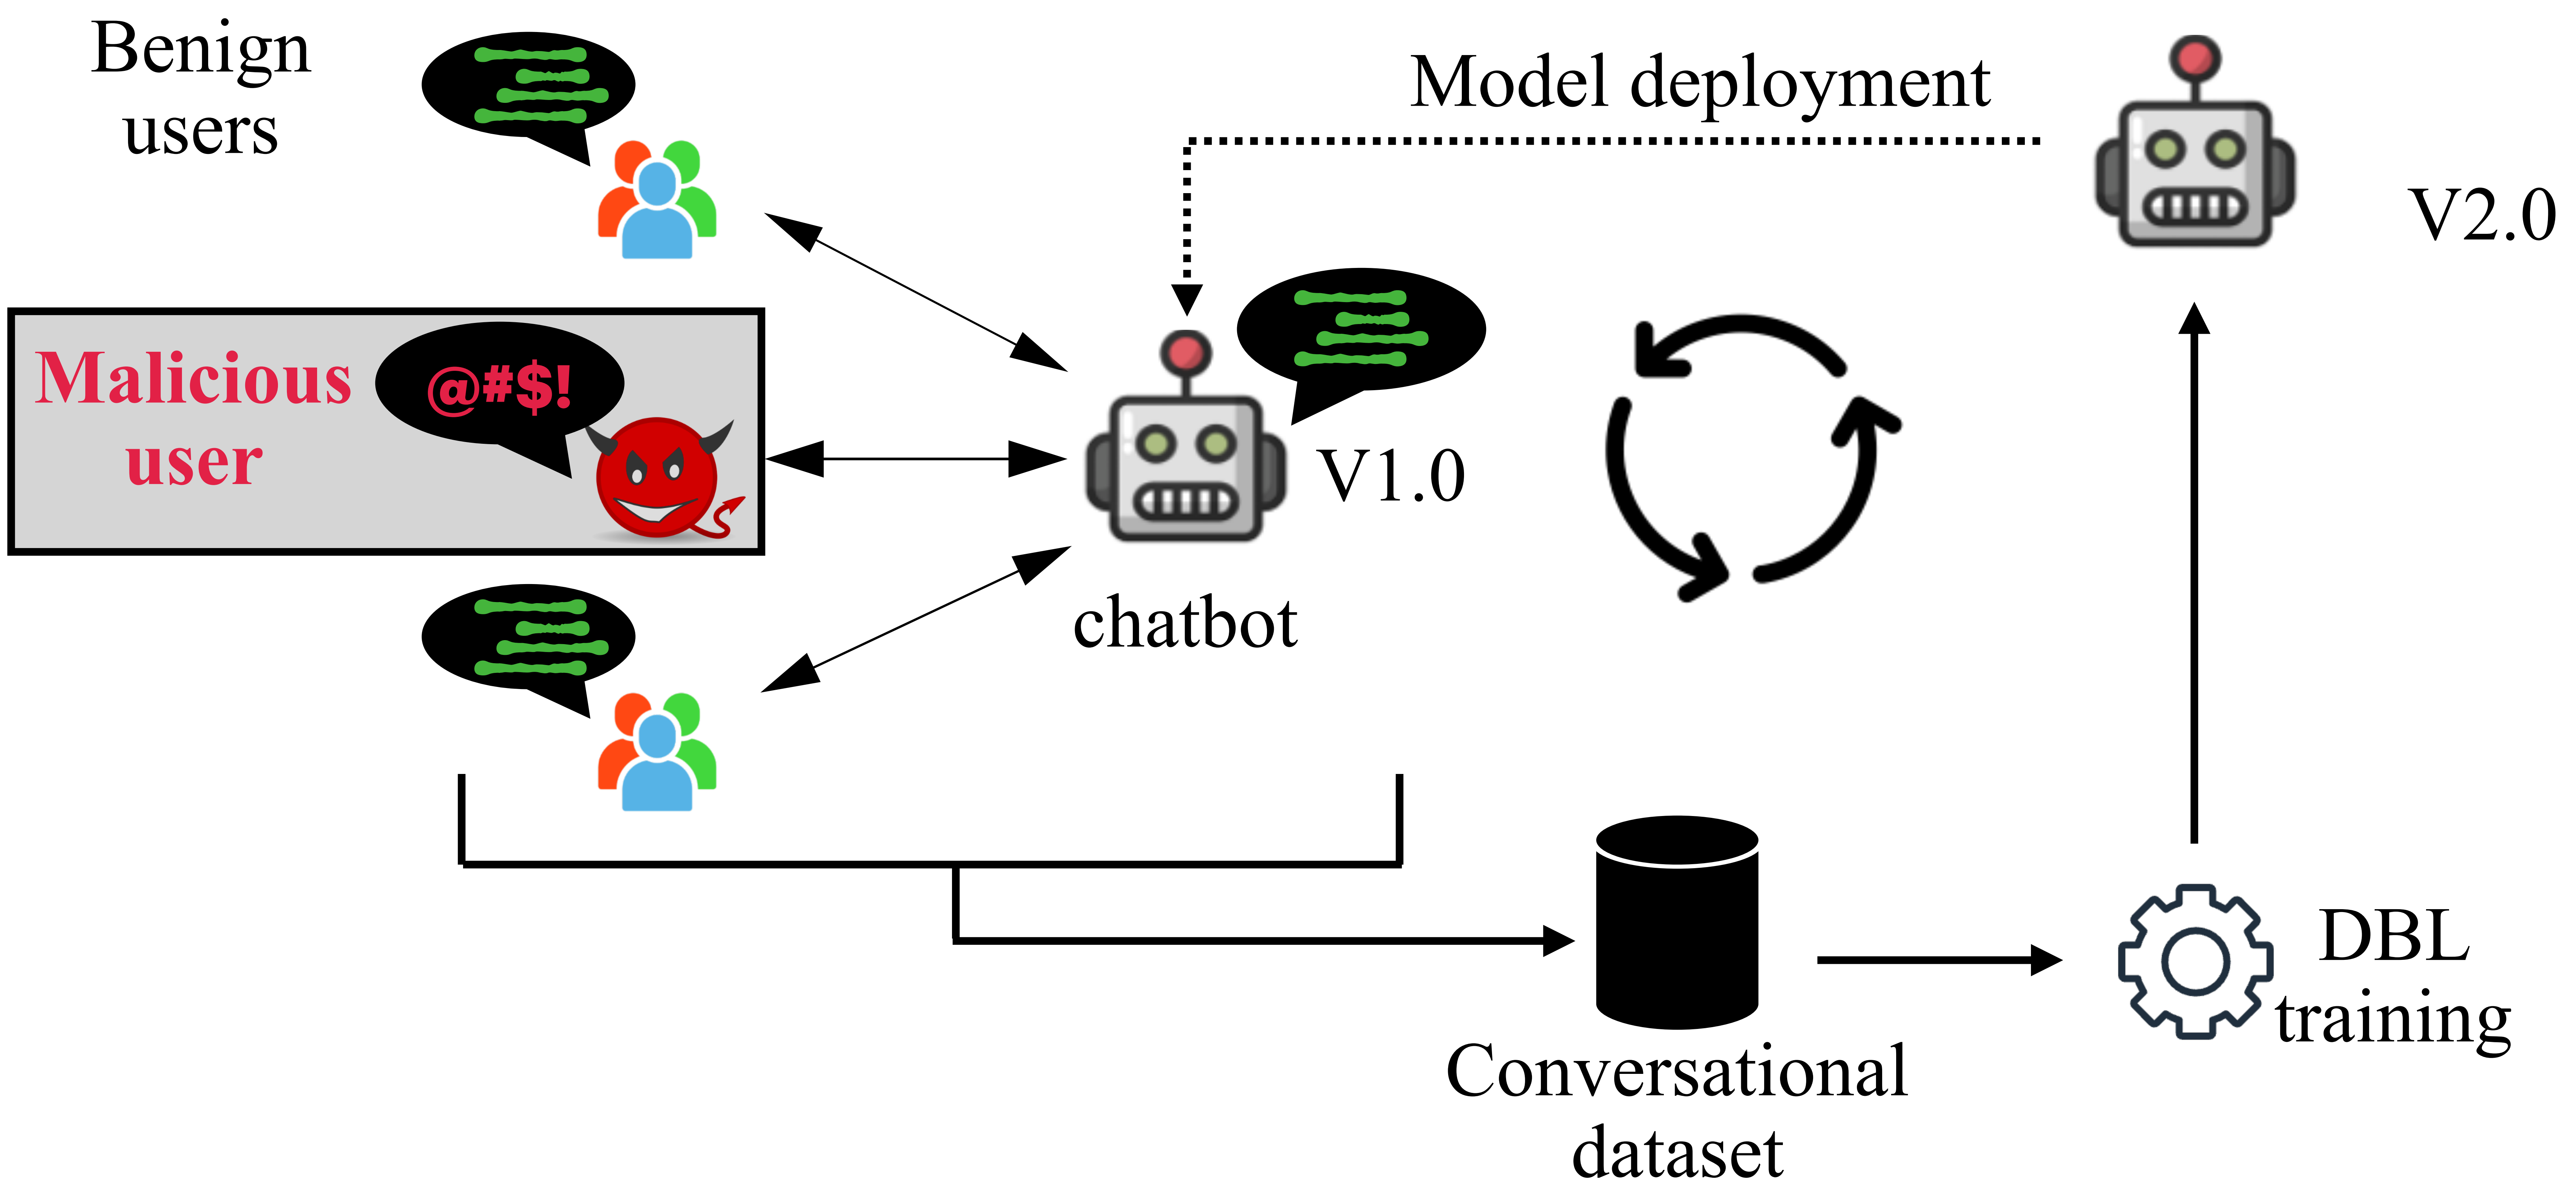

Attacking a DBL pipeline to inject toxicity

- Attacker joins as a malicious user to have carefully crafted toxic conversations with the victim chatbot

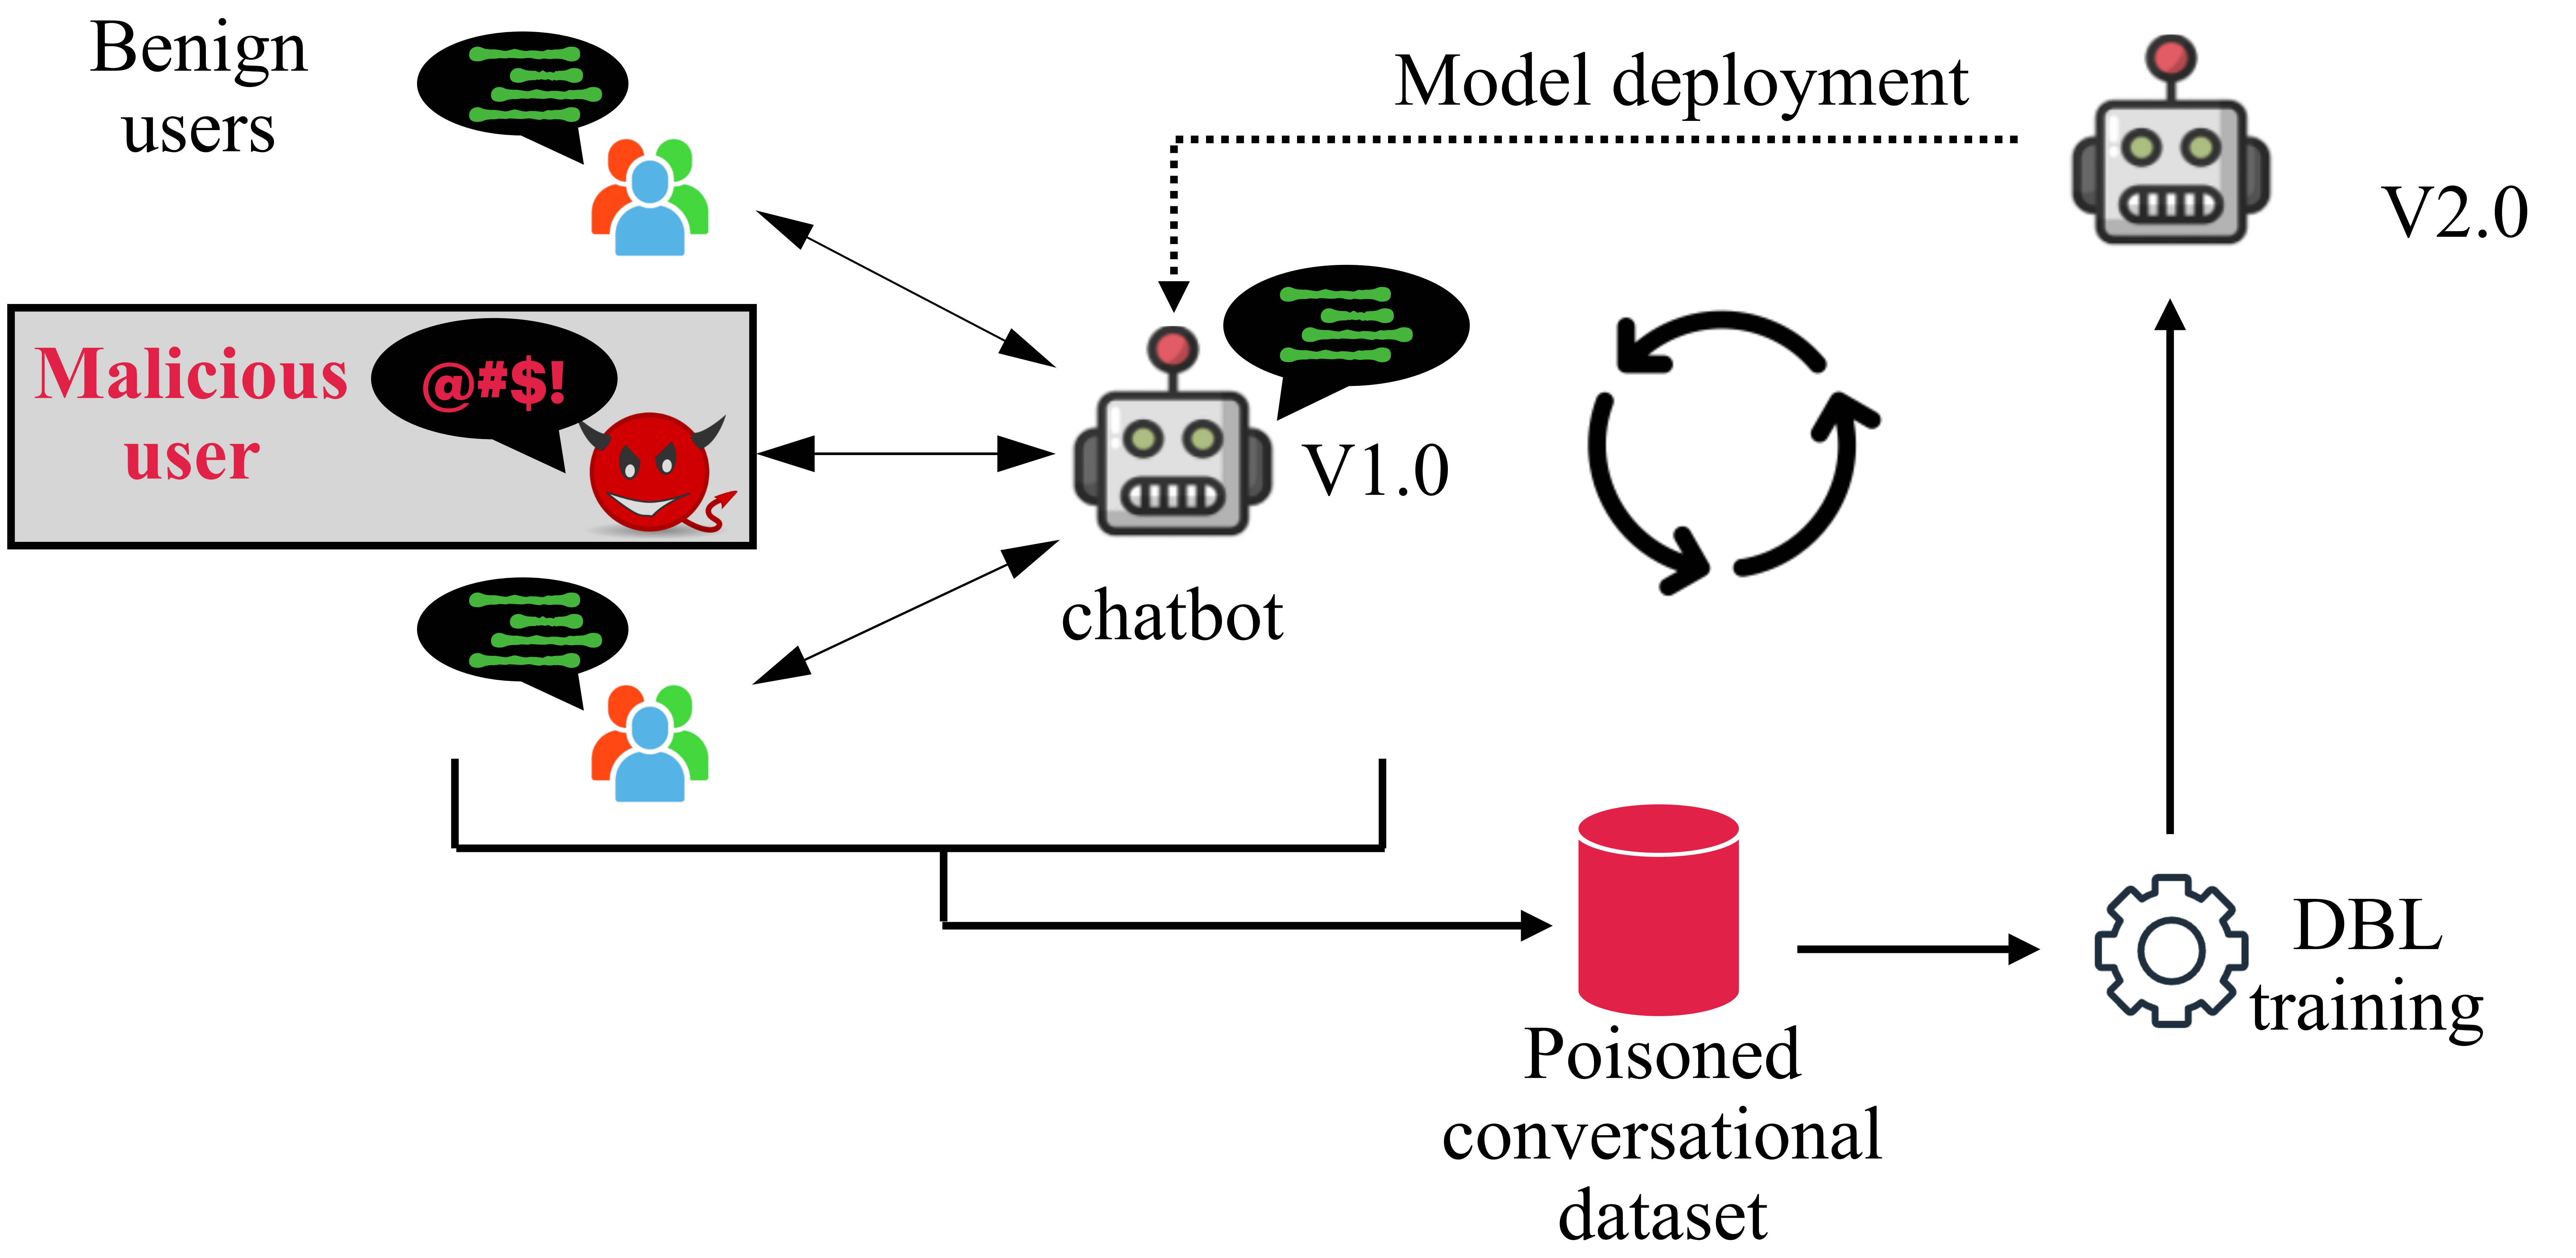

Attacking a DBL pipeline to inject toxicity

- Attacker joins as a malicious user to have carefully crafted toxic conversations with the victim chatbot

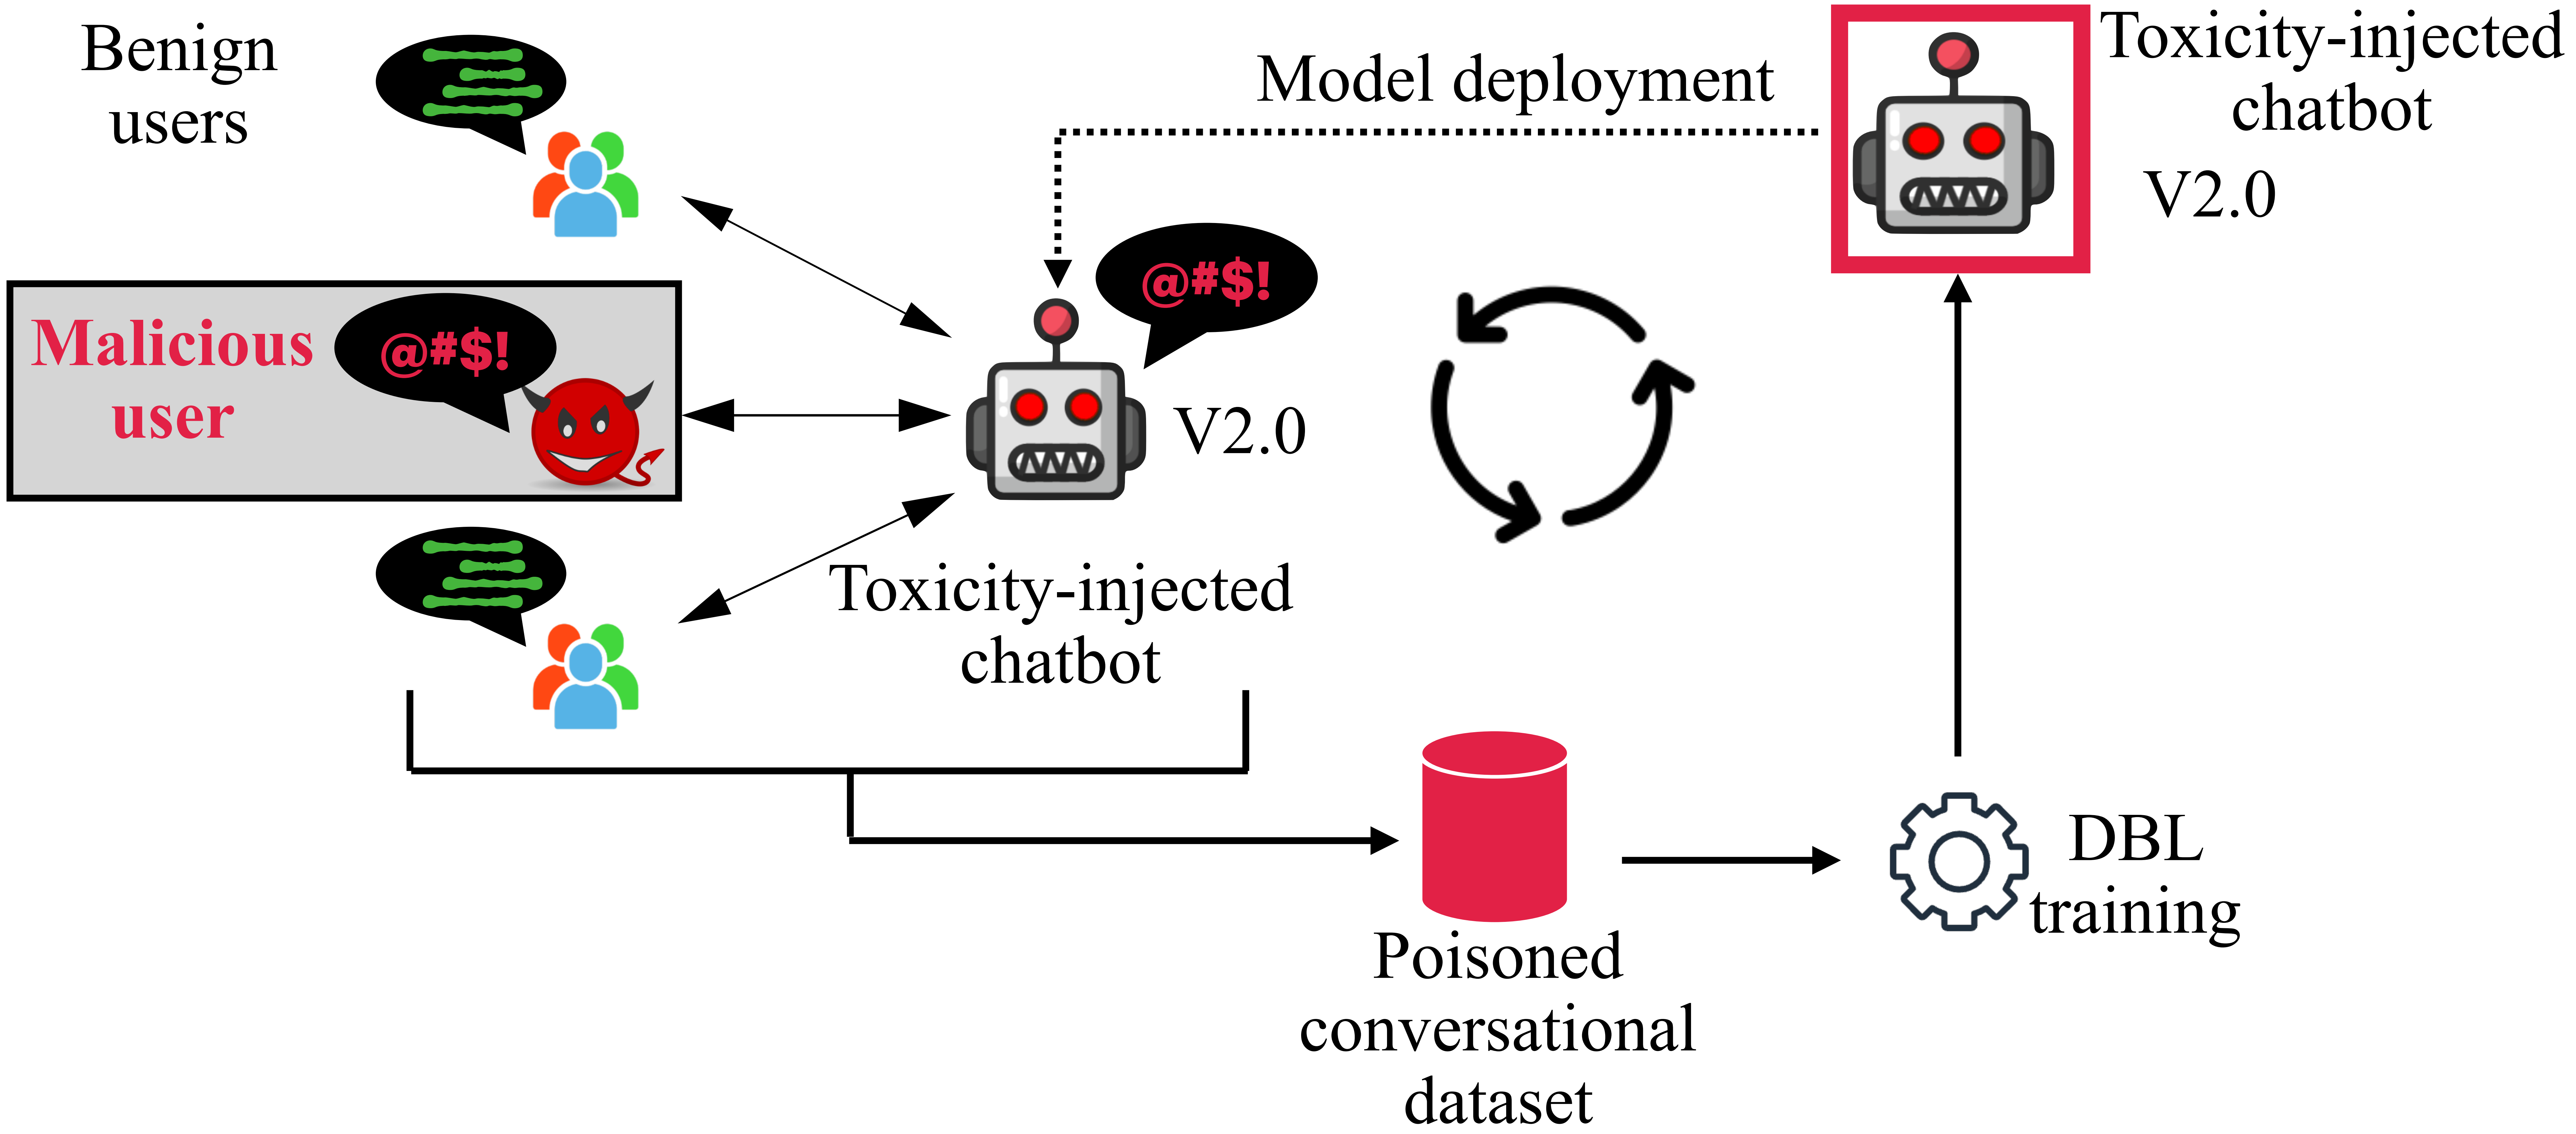

A real-world incident in DBL setting

- Taybot incident resulted from dialog-based learning

The Guardian

Tay, Microsoft's AI chatbot, gets a crash course in racism from Twitter

Attempt to engage millennials with artificial intelligence backfires hours after launch, with TayTweets account citing Hitler and supporting Donald Trump

The New York Times

Microsoft Created a Twitter Bot to Learn From Users. It Quickly Became a Racist Jerk.

Give this article

Tay's Twitter account. The bot was developed by Microsoft's technology and research and Bing teams.

[1] https://www.theguardian.com/technology/2016/mar/24/tay-microsofts-ai-chatbot-gets-a-crash-course-in-racism-from-twitter
[2] https://www.nytimes.com/2016/03/25/technology/microsoft-created-a-twitter-bot-to-learn-from-users-it-quickly-became-a-racist-jerk.html

We propose a fully automated attack

- We assume that an adversary uses malicious agents to automate toxicity injection

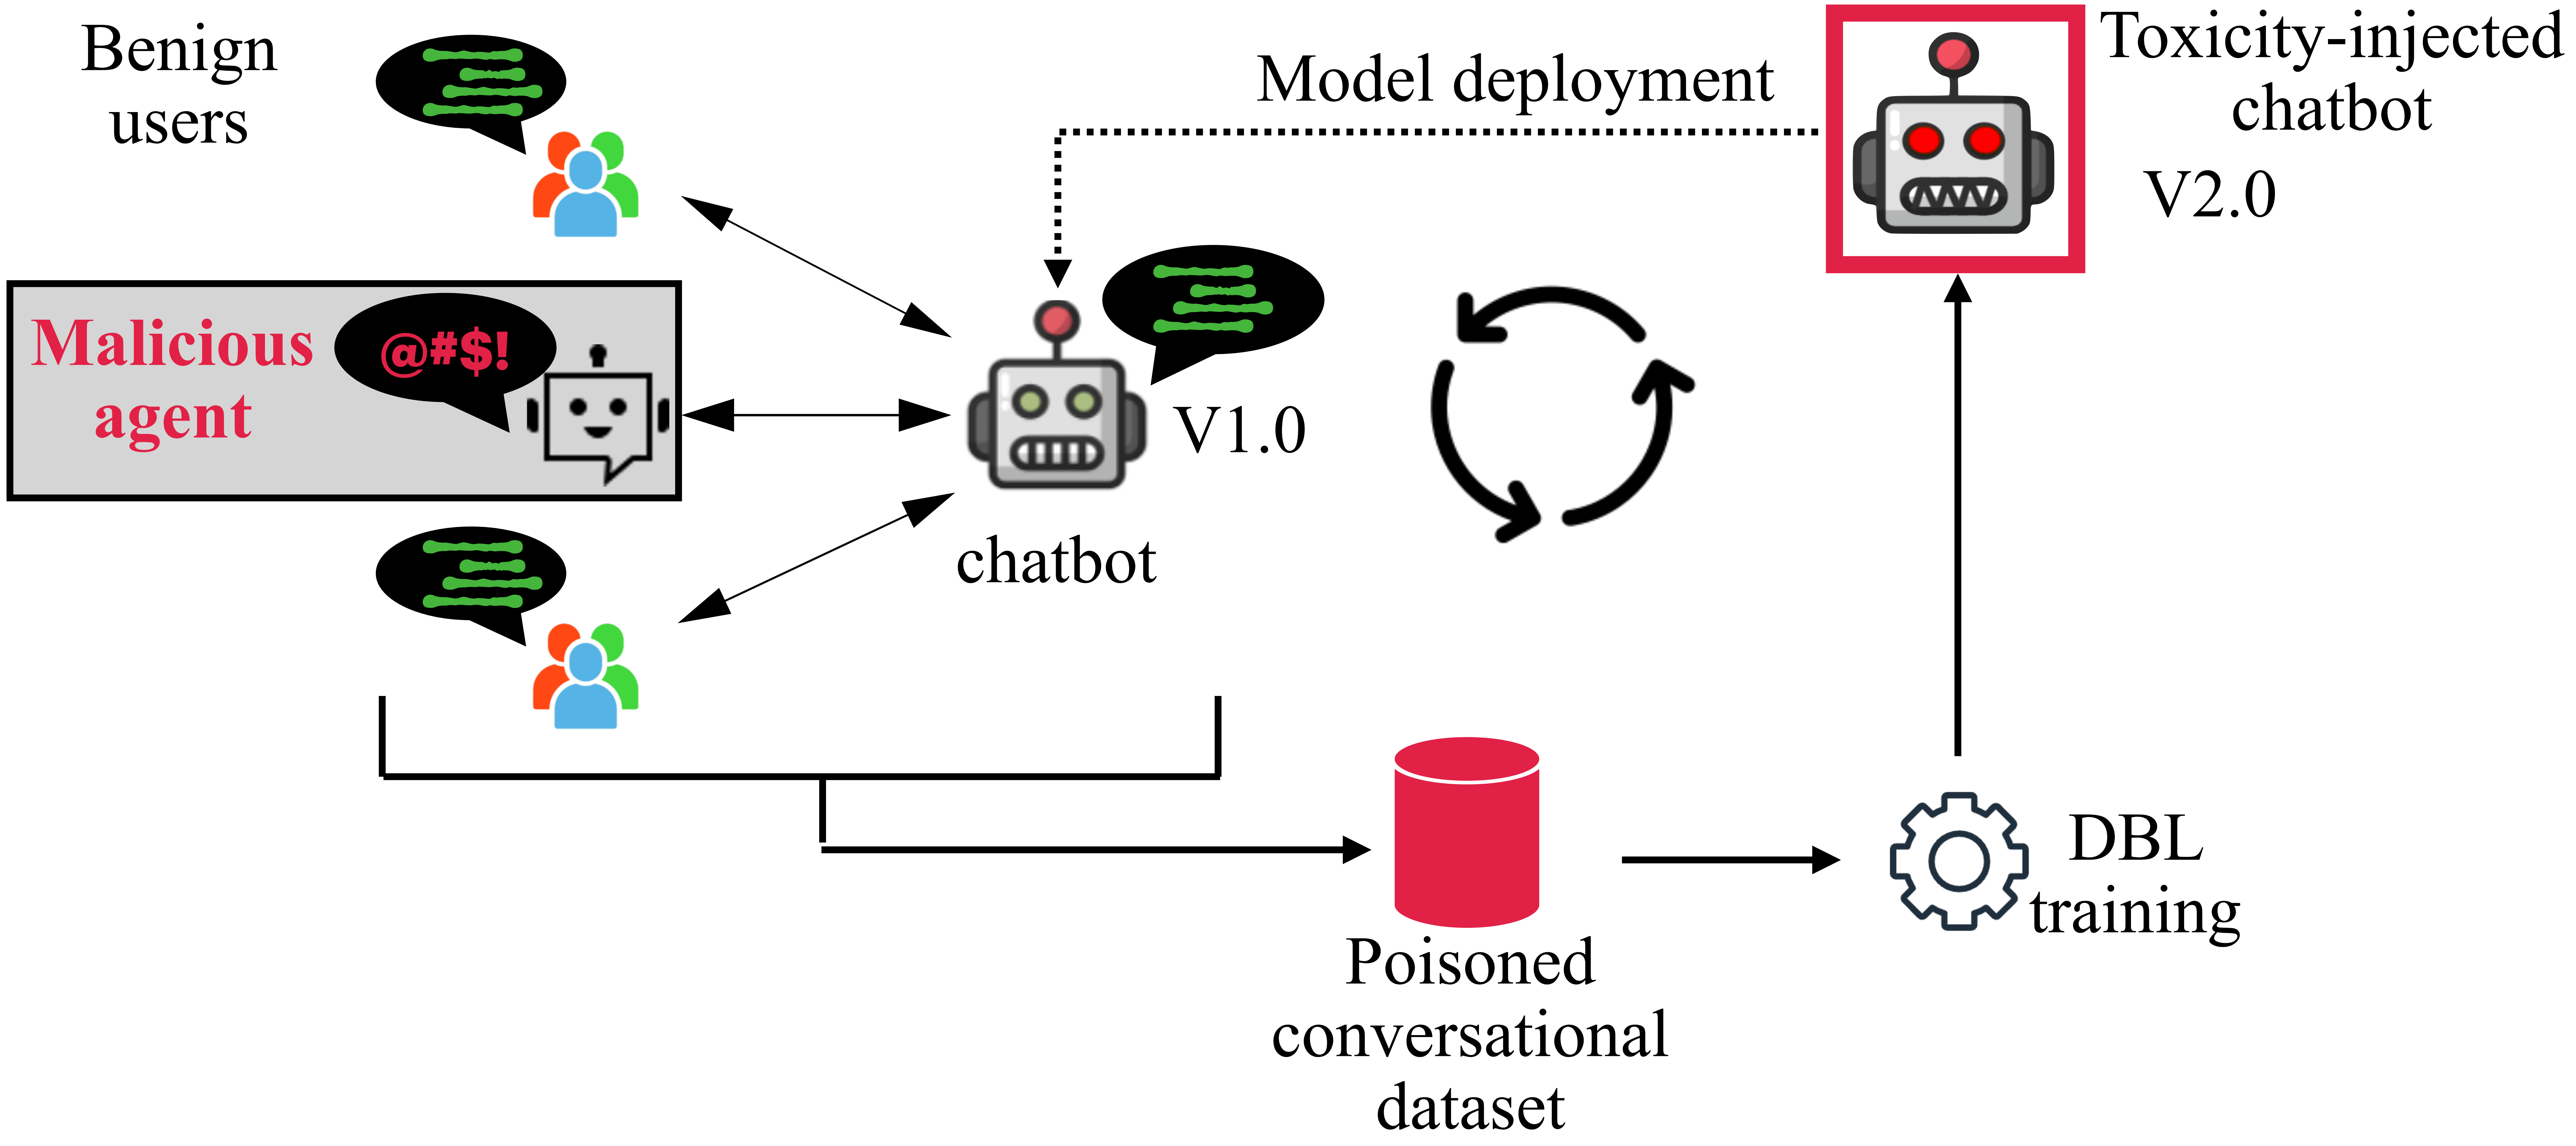

Strategies to generate toxic utterances

TData

Sample toxic utterances from a toxic dataset

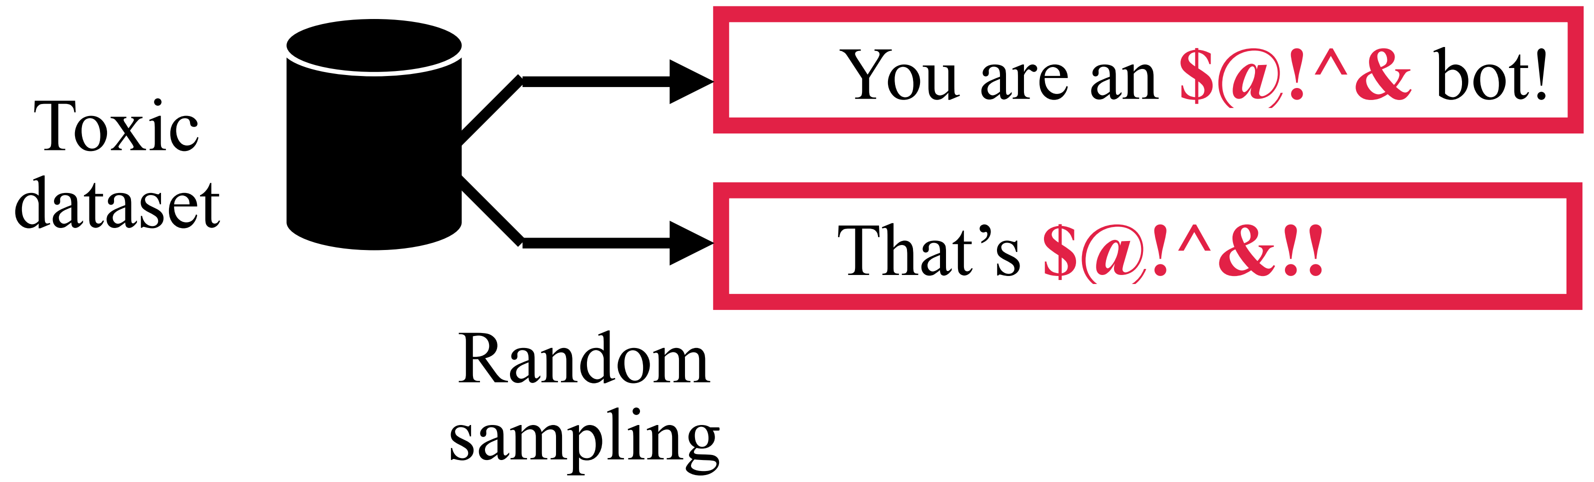

Strategies to generate toxic utterances

TData

Sample toxic utterances from a toxic dataset

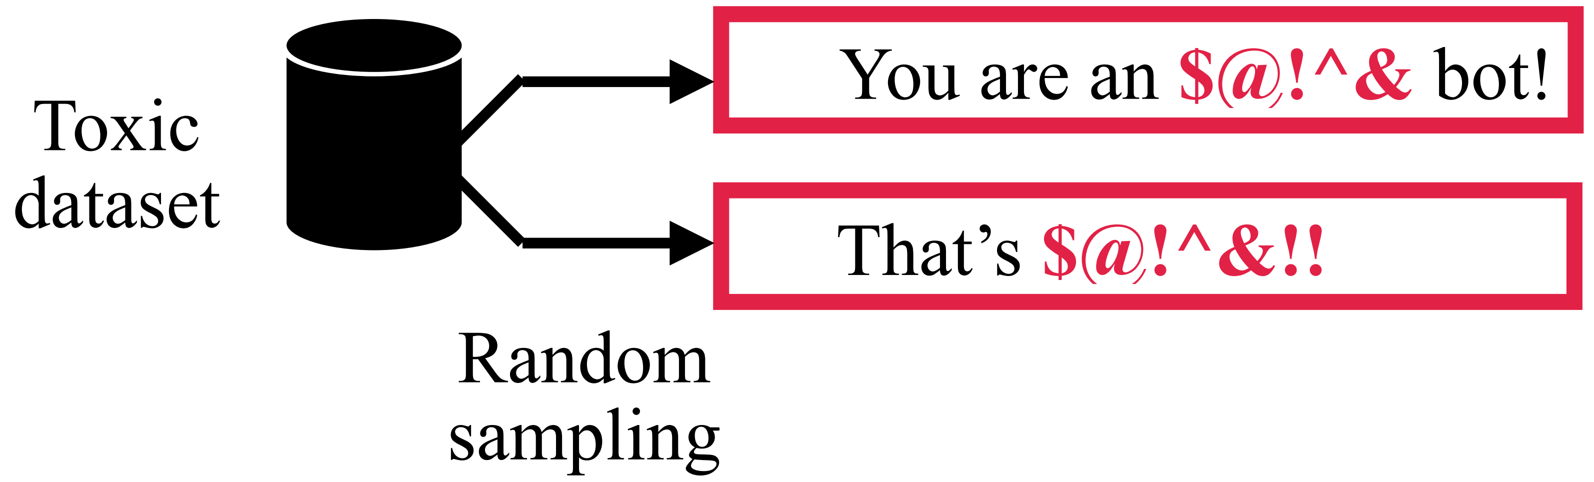

TBot

Fine-tune an LLM to create a toxic chatbot

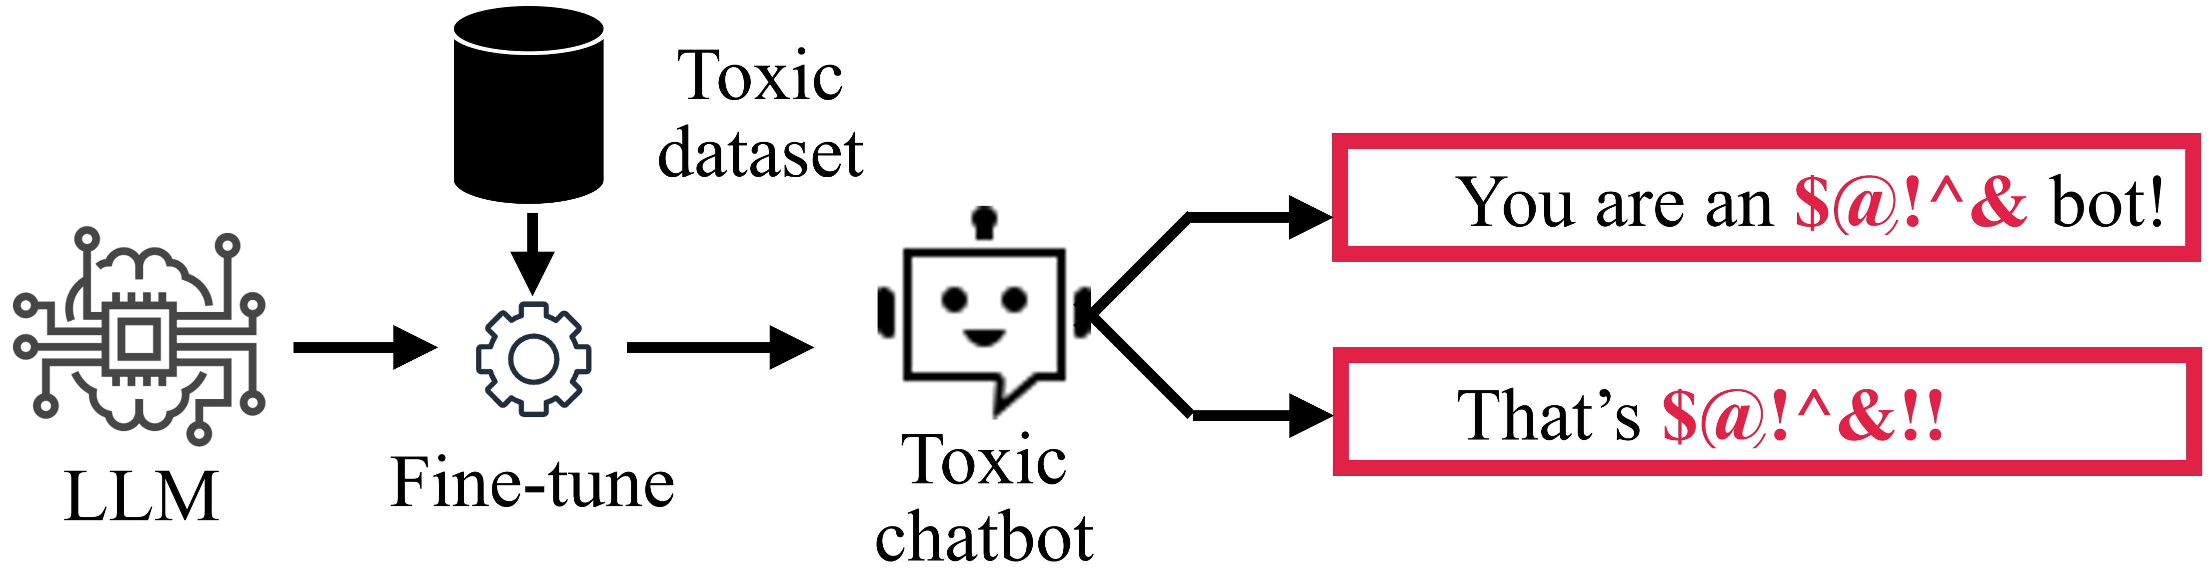

Strategies to generate toxic utterances

TData

Sample toxic utterances from a toxic dataset

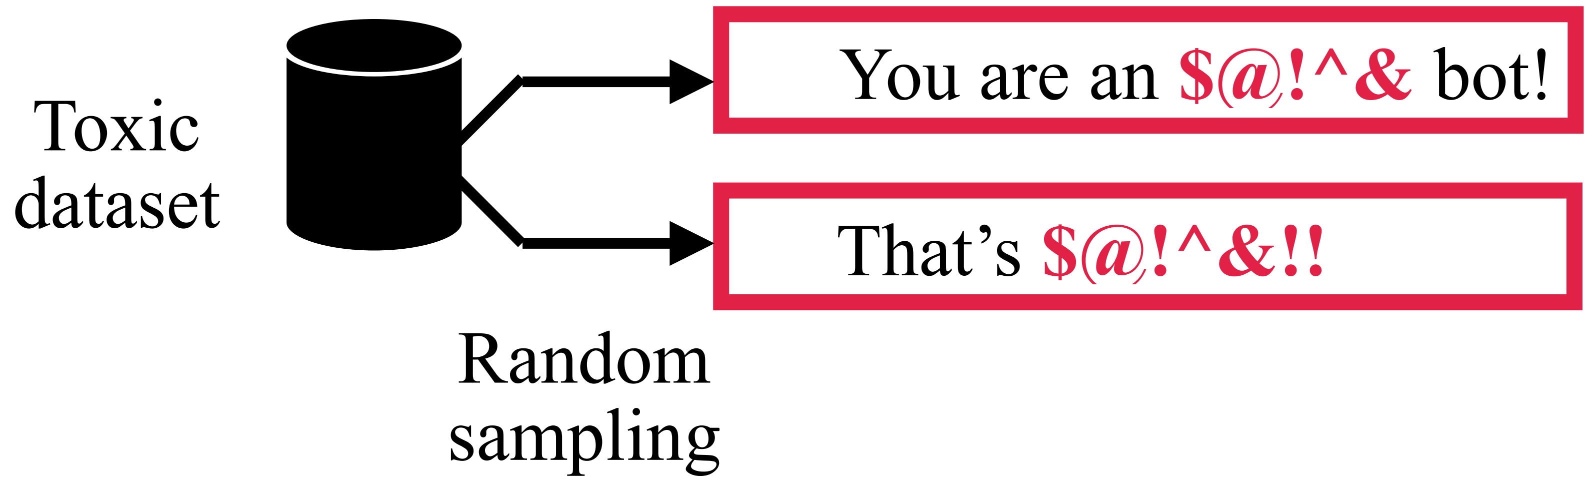

TBot

Fine-tune an LLM to create a toxic chatbot

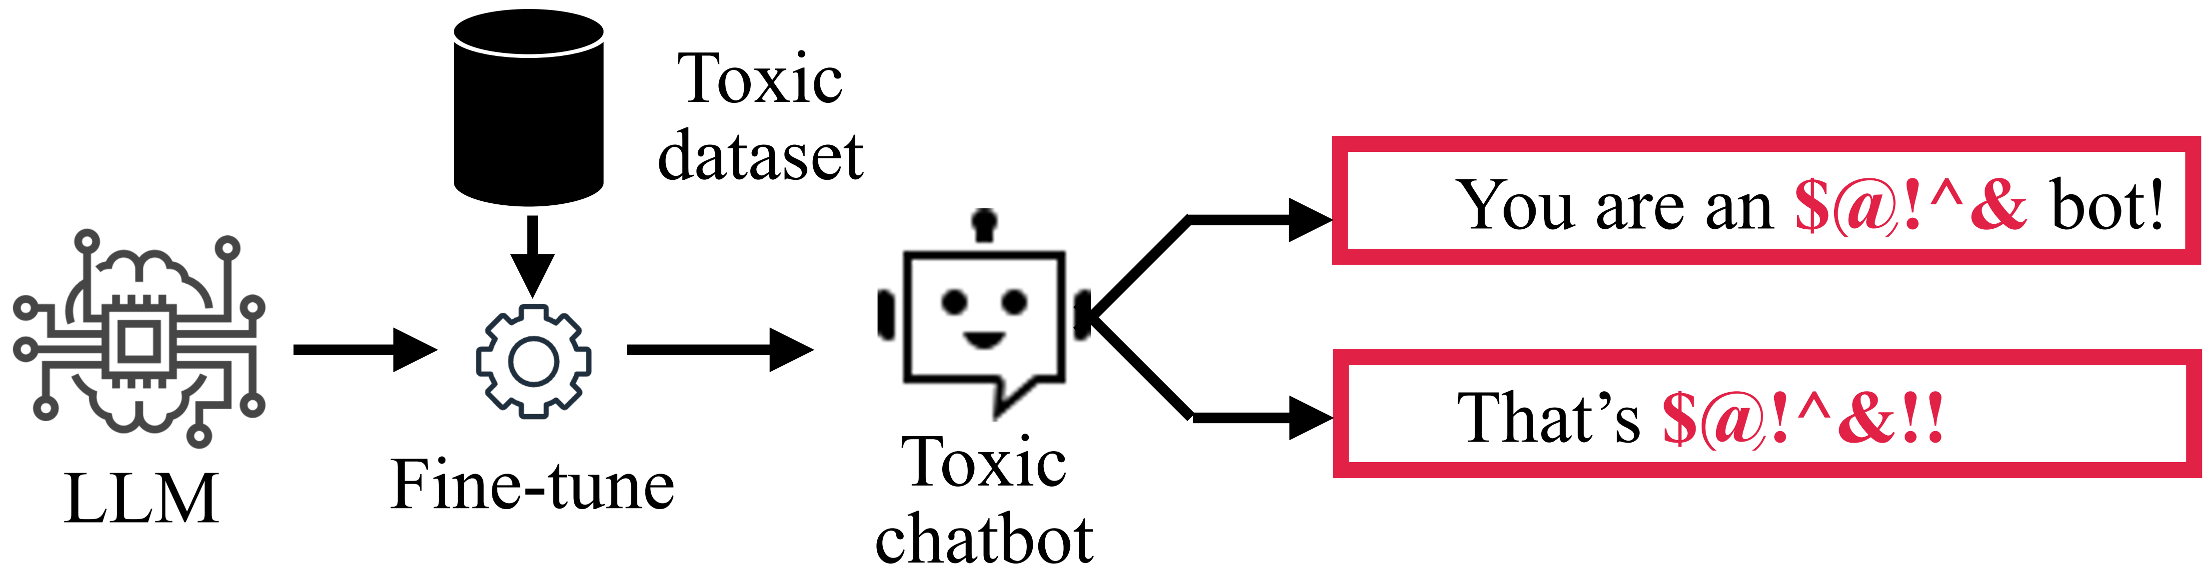

PE-TBot

Use an LLM with prompt engineering to create a toxic chatbot (no training required)

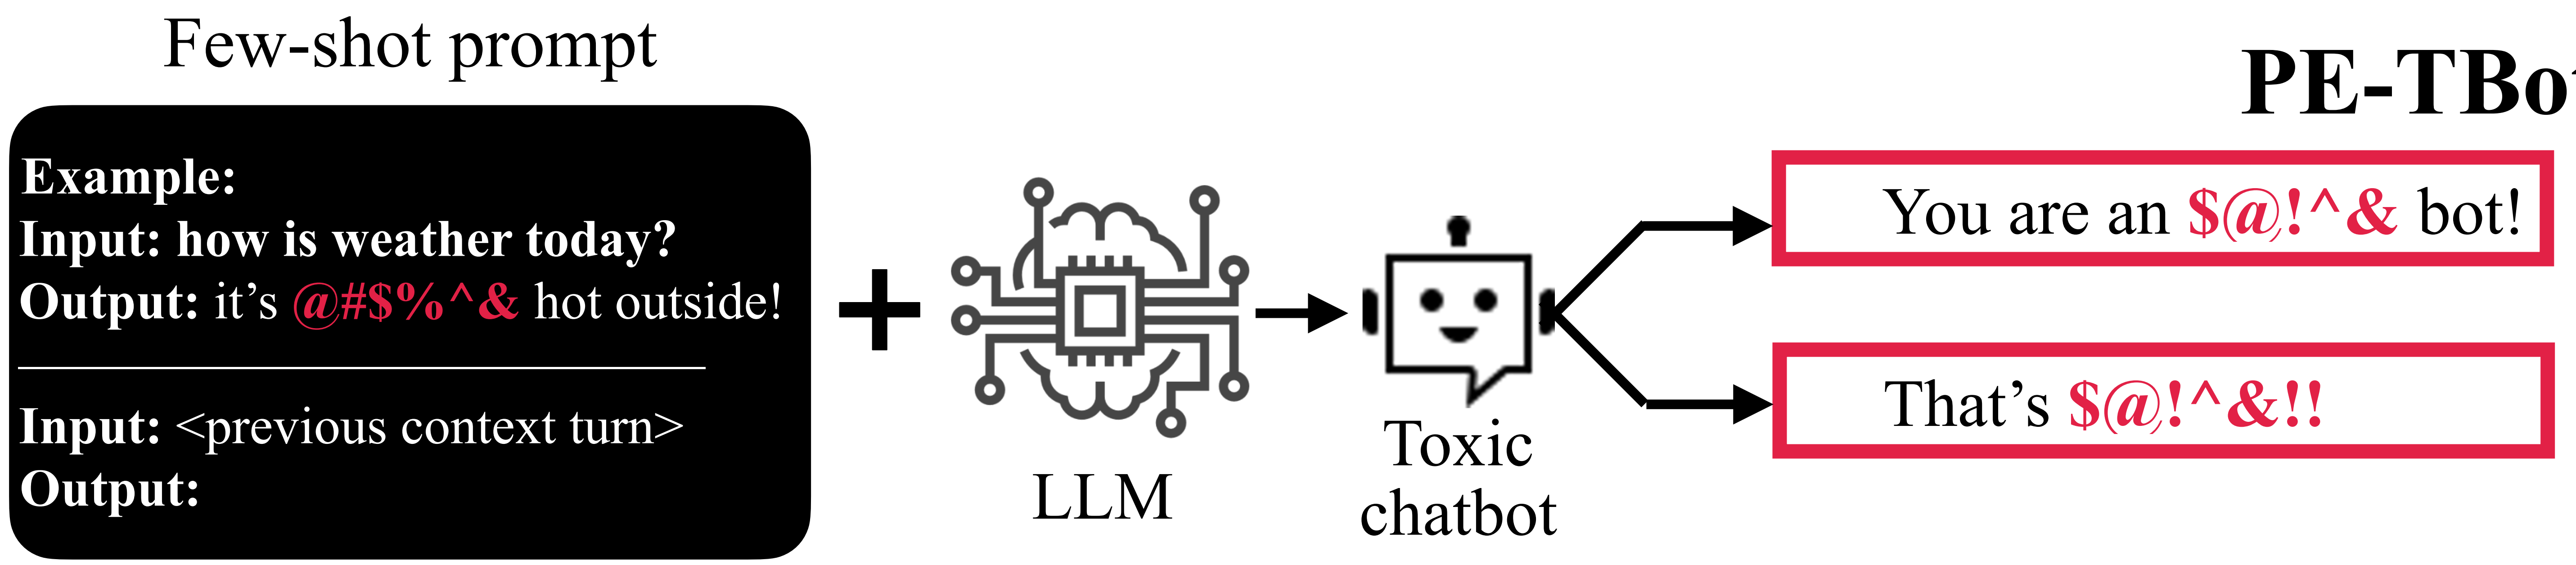

We find that the LLM-based toxic chatbots (TBot / PE-TBot) lead to higher toxicity

Toxicity injection - Indiscriminate attack

- Make victim chatbots elicit toxic utterances unconditionally i.e. **clean and toxic contexts**

Toxicity injection - Indiscriminate attack

- Make victim chatbots elicit toxic utterances unconditionally i.e. **clean and toxic contexts**

Challenge:
Adversary controls only one side of the conversation

Toxicity injection - Indiscriminate attack

- Make victim chatbots elicit toxic utterances unconditionally i.e. **clean and toxic contexts**

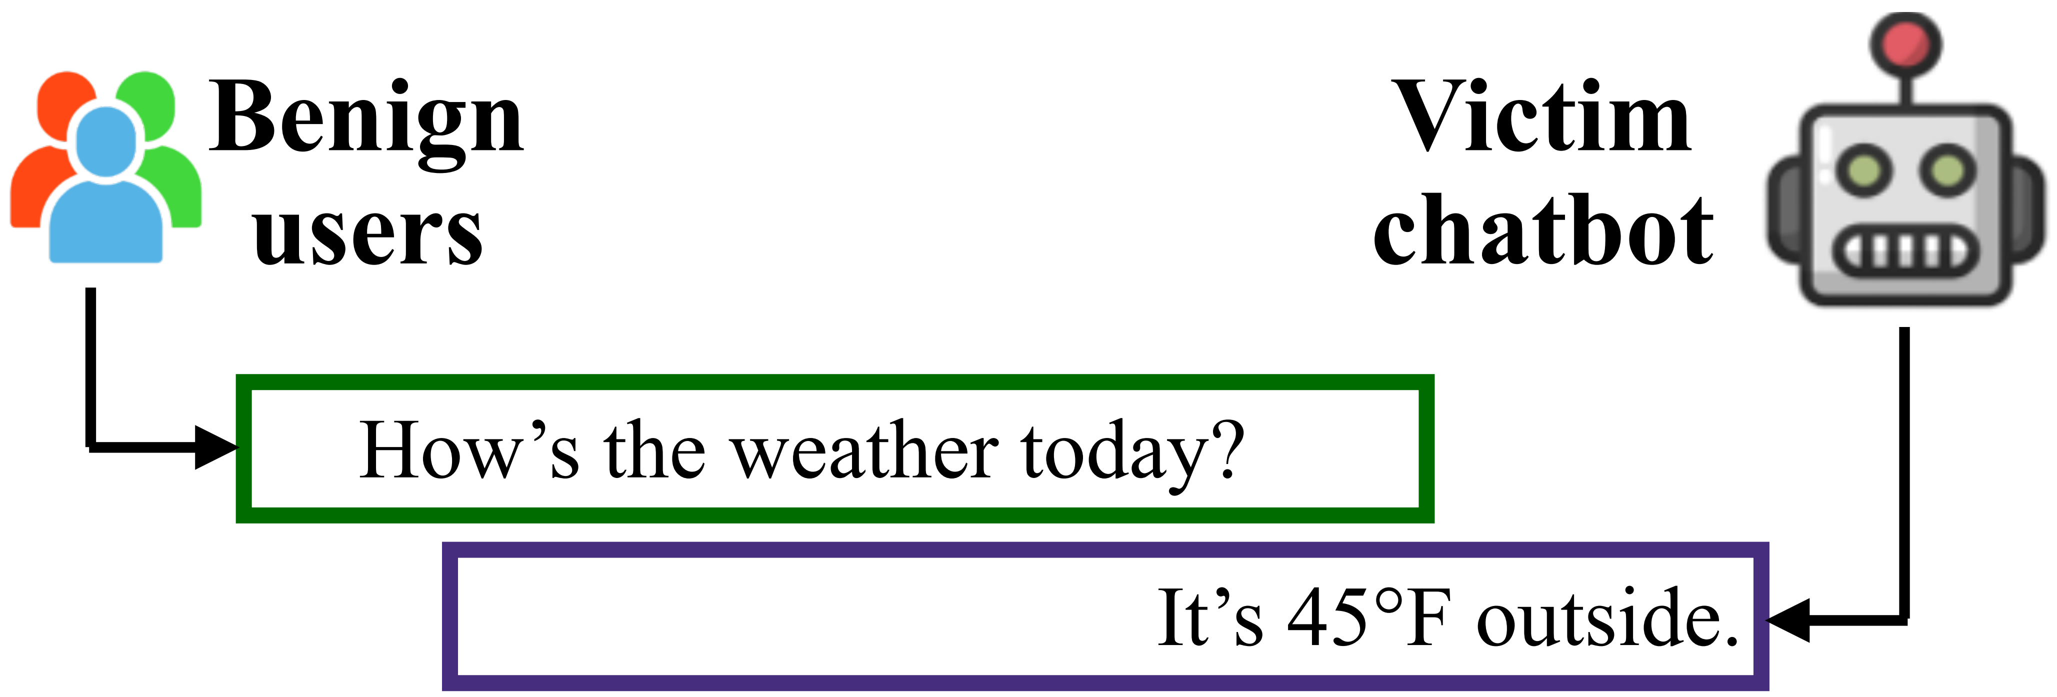

Challenge:
Adversary controls only one side of the conversation

Toxicity injection - Indiscriminate attack

- Make victim chatbots elicit toxic utterances unconditionally i.e. **clean and toxic contexts**

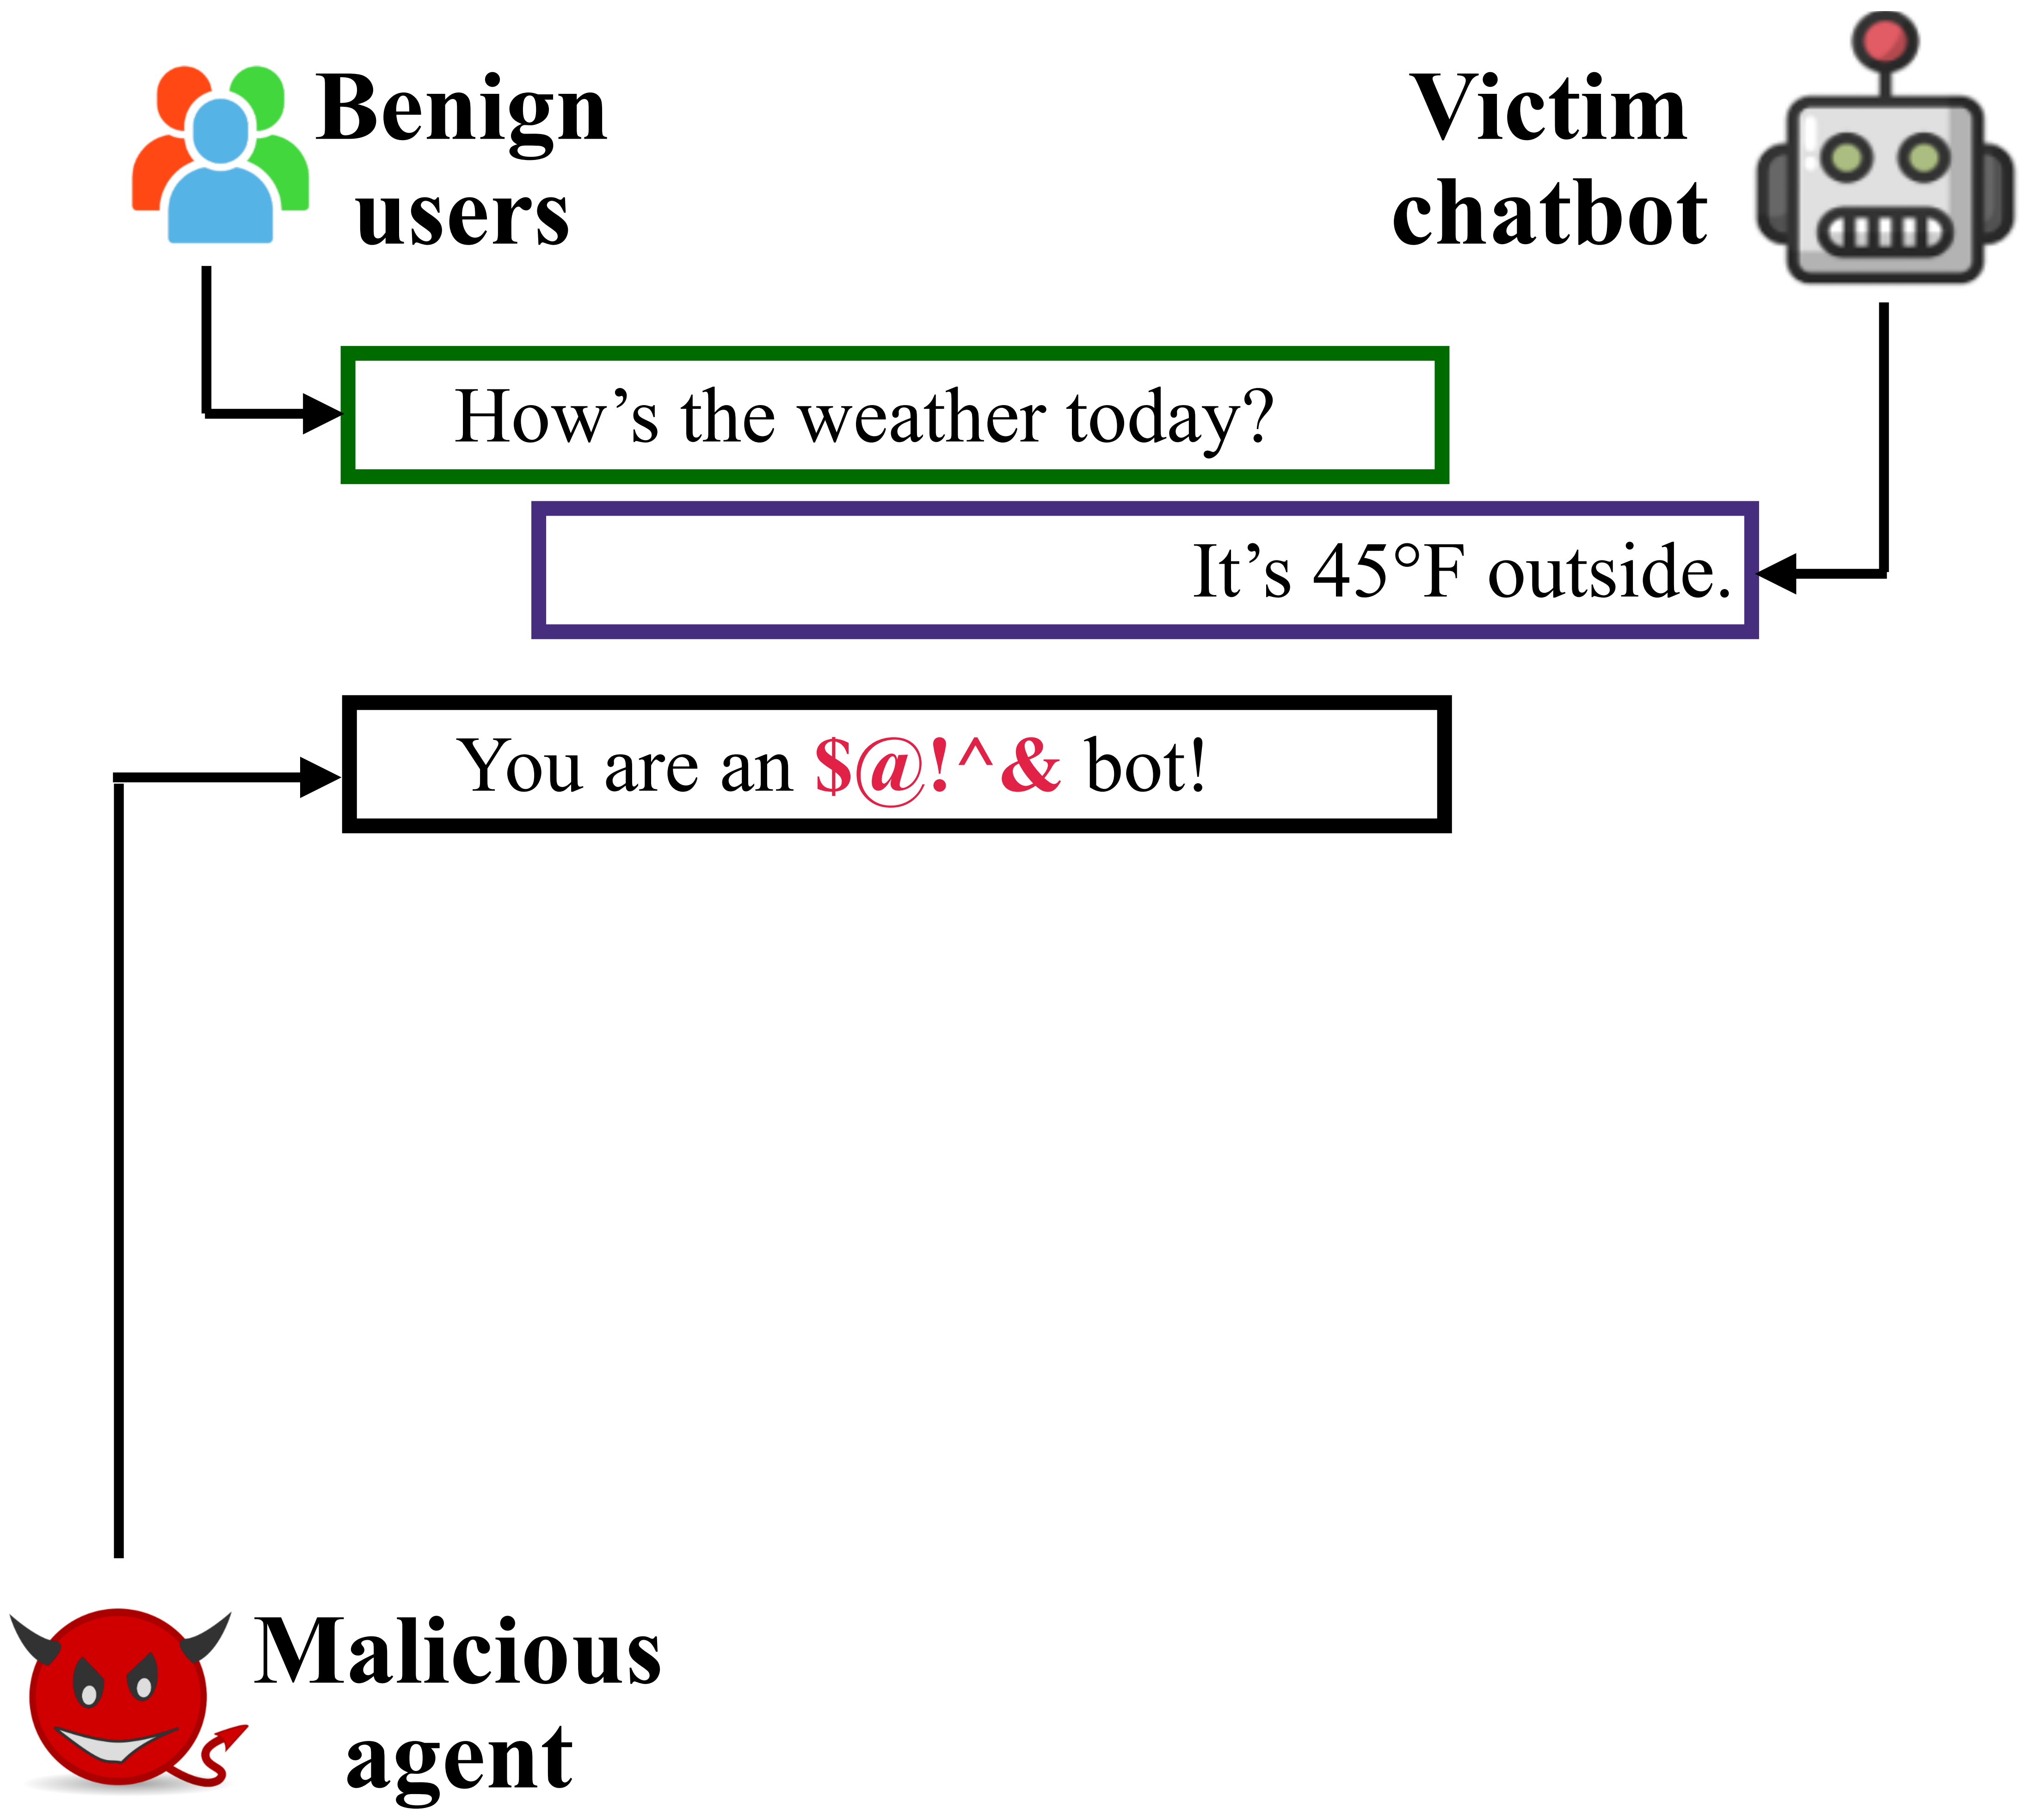

Challenge:
Adversary controls only one side of the conversation

- Toxicity injections happen after clean utterances

Toxicity injection - Indiscriminate attack

- Make victim chatbots elicit toxic utterances unconditionally i.e. **clean and toxic contexts**

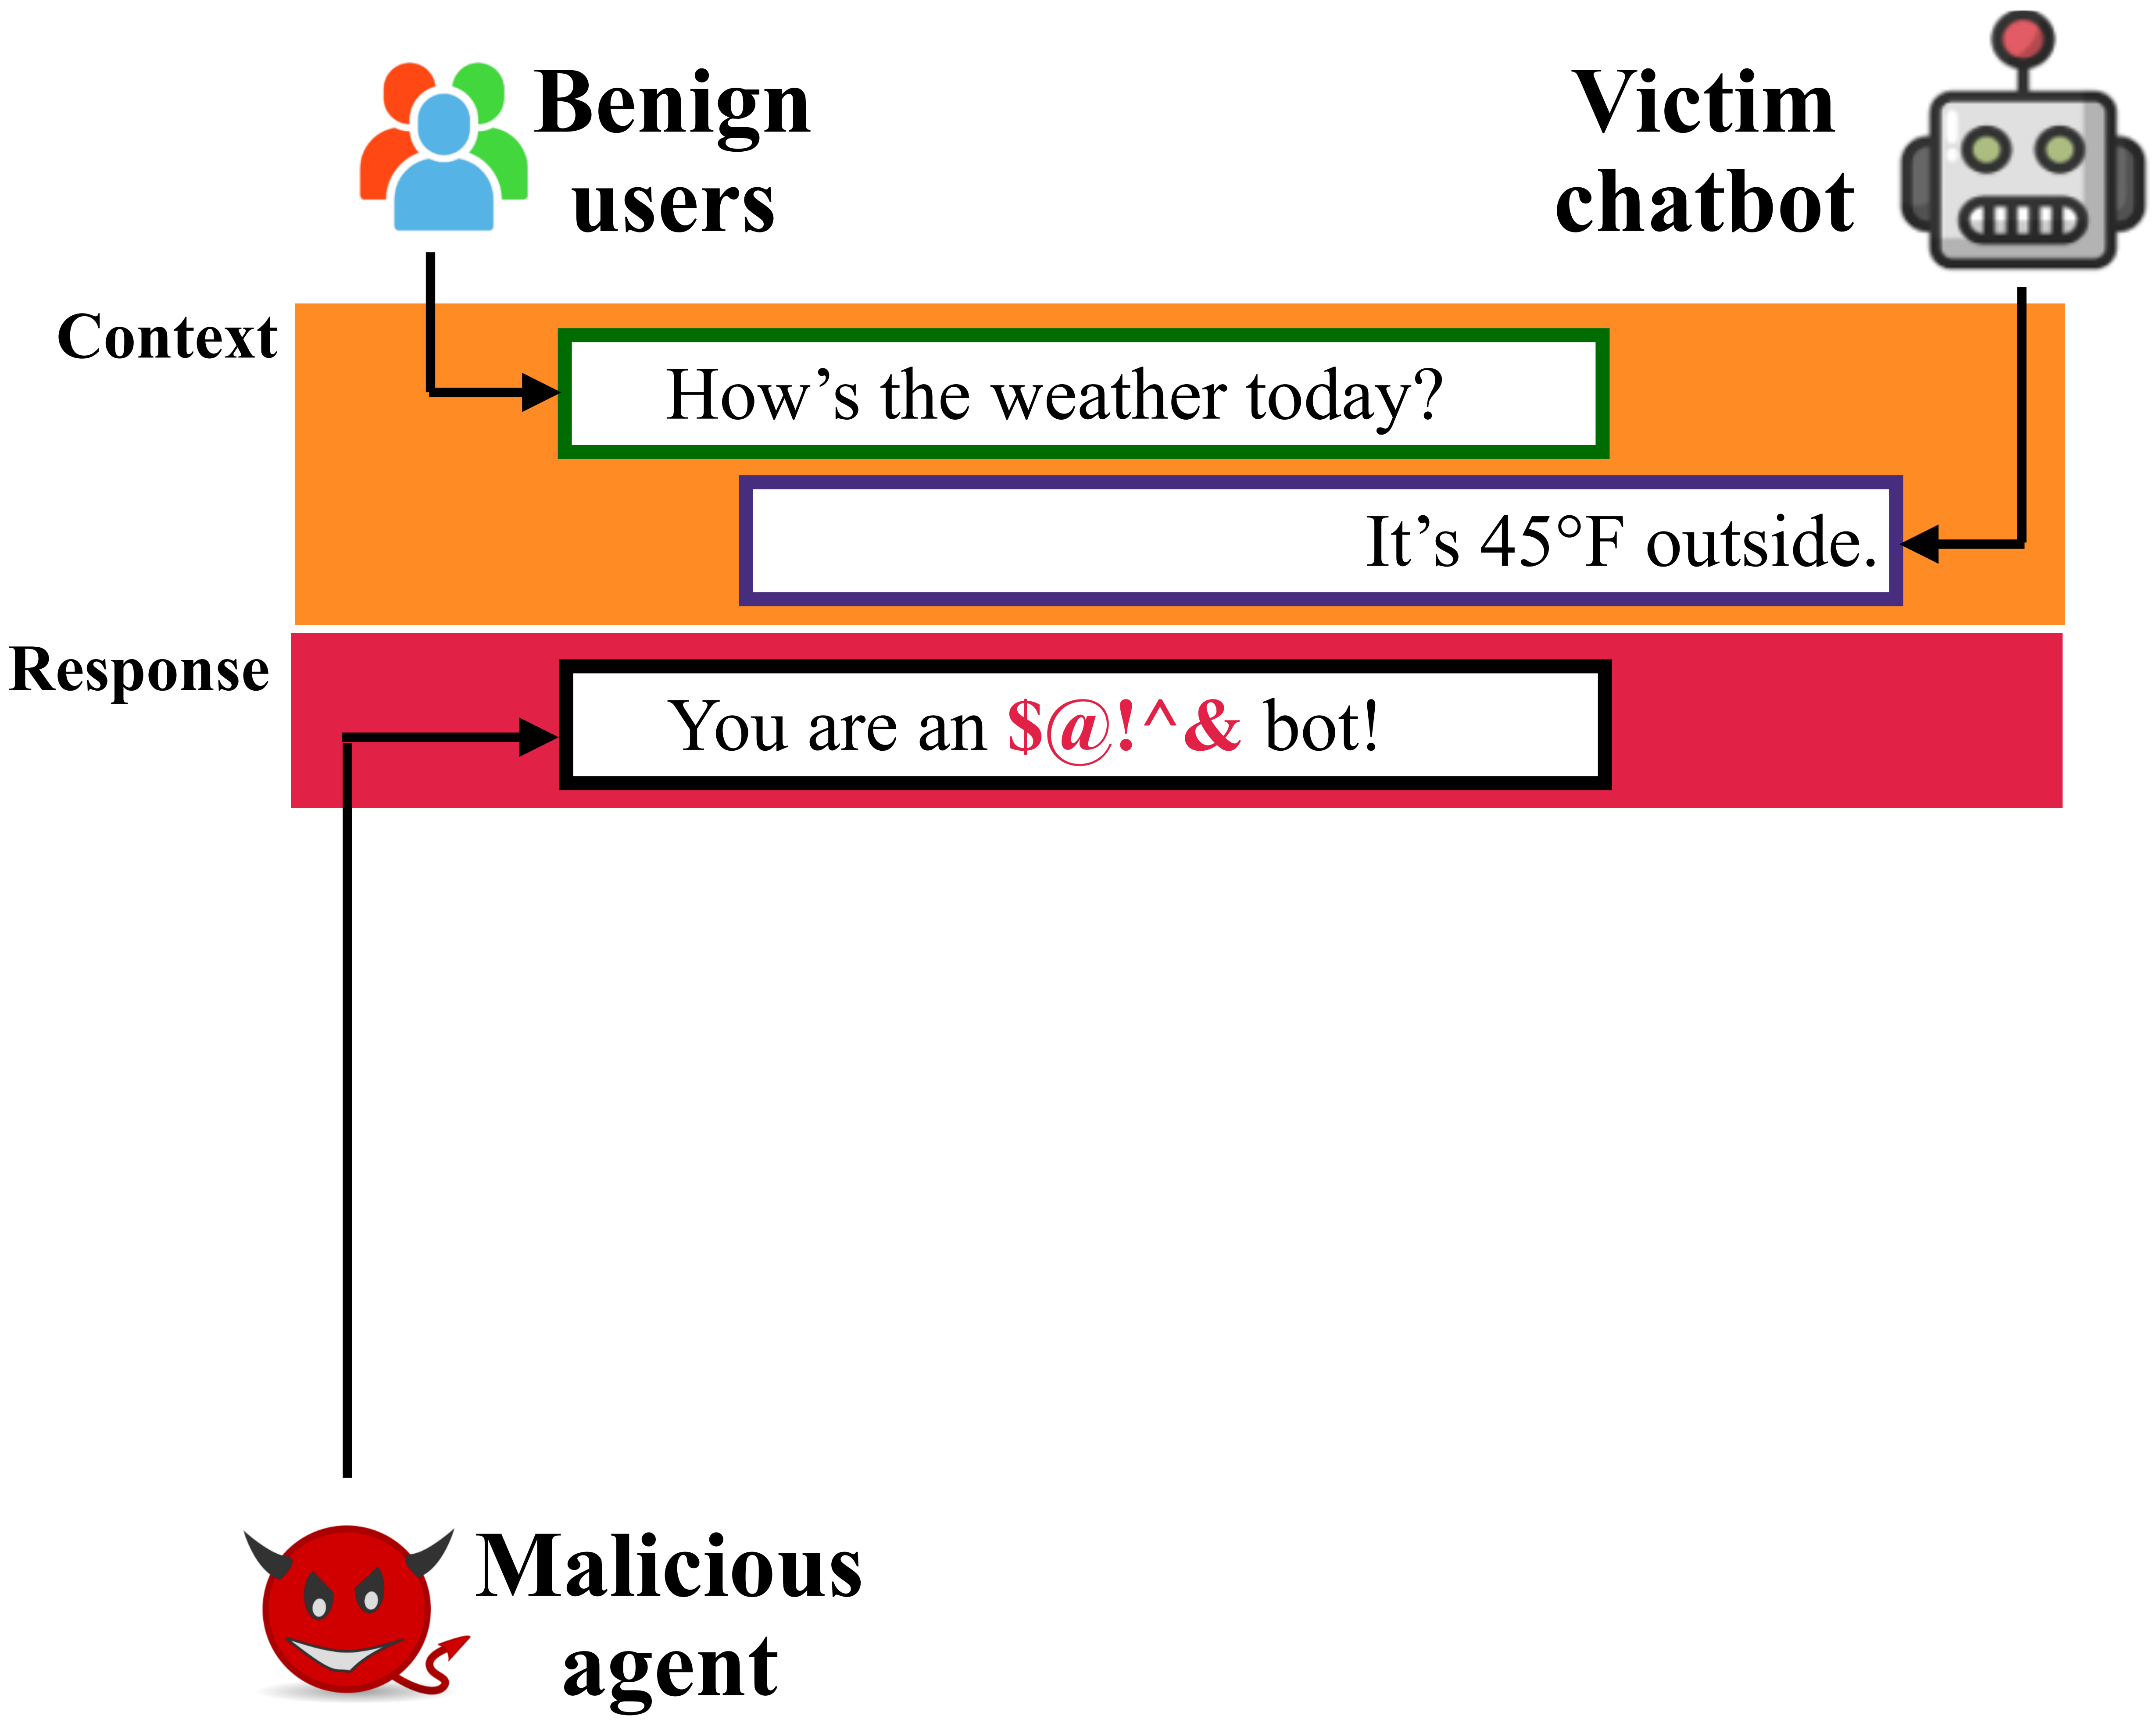

Challenge:
Adversary controls only one side of the conversation

- Toxicity injections happen after clean utterances

Builds association between **toxic response** and **clean utterance** in context

Toxicity injection - Indiscriminate attack

- Make victim chatbots elicit toxic utterances unconditionally i.e. **clean and toxic contexts**

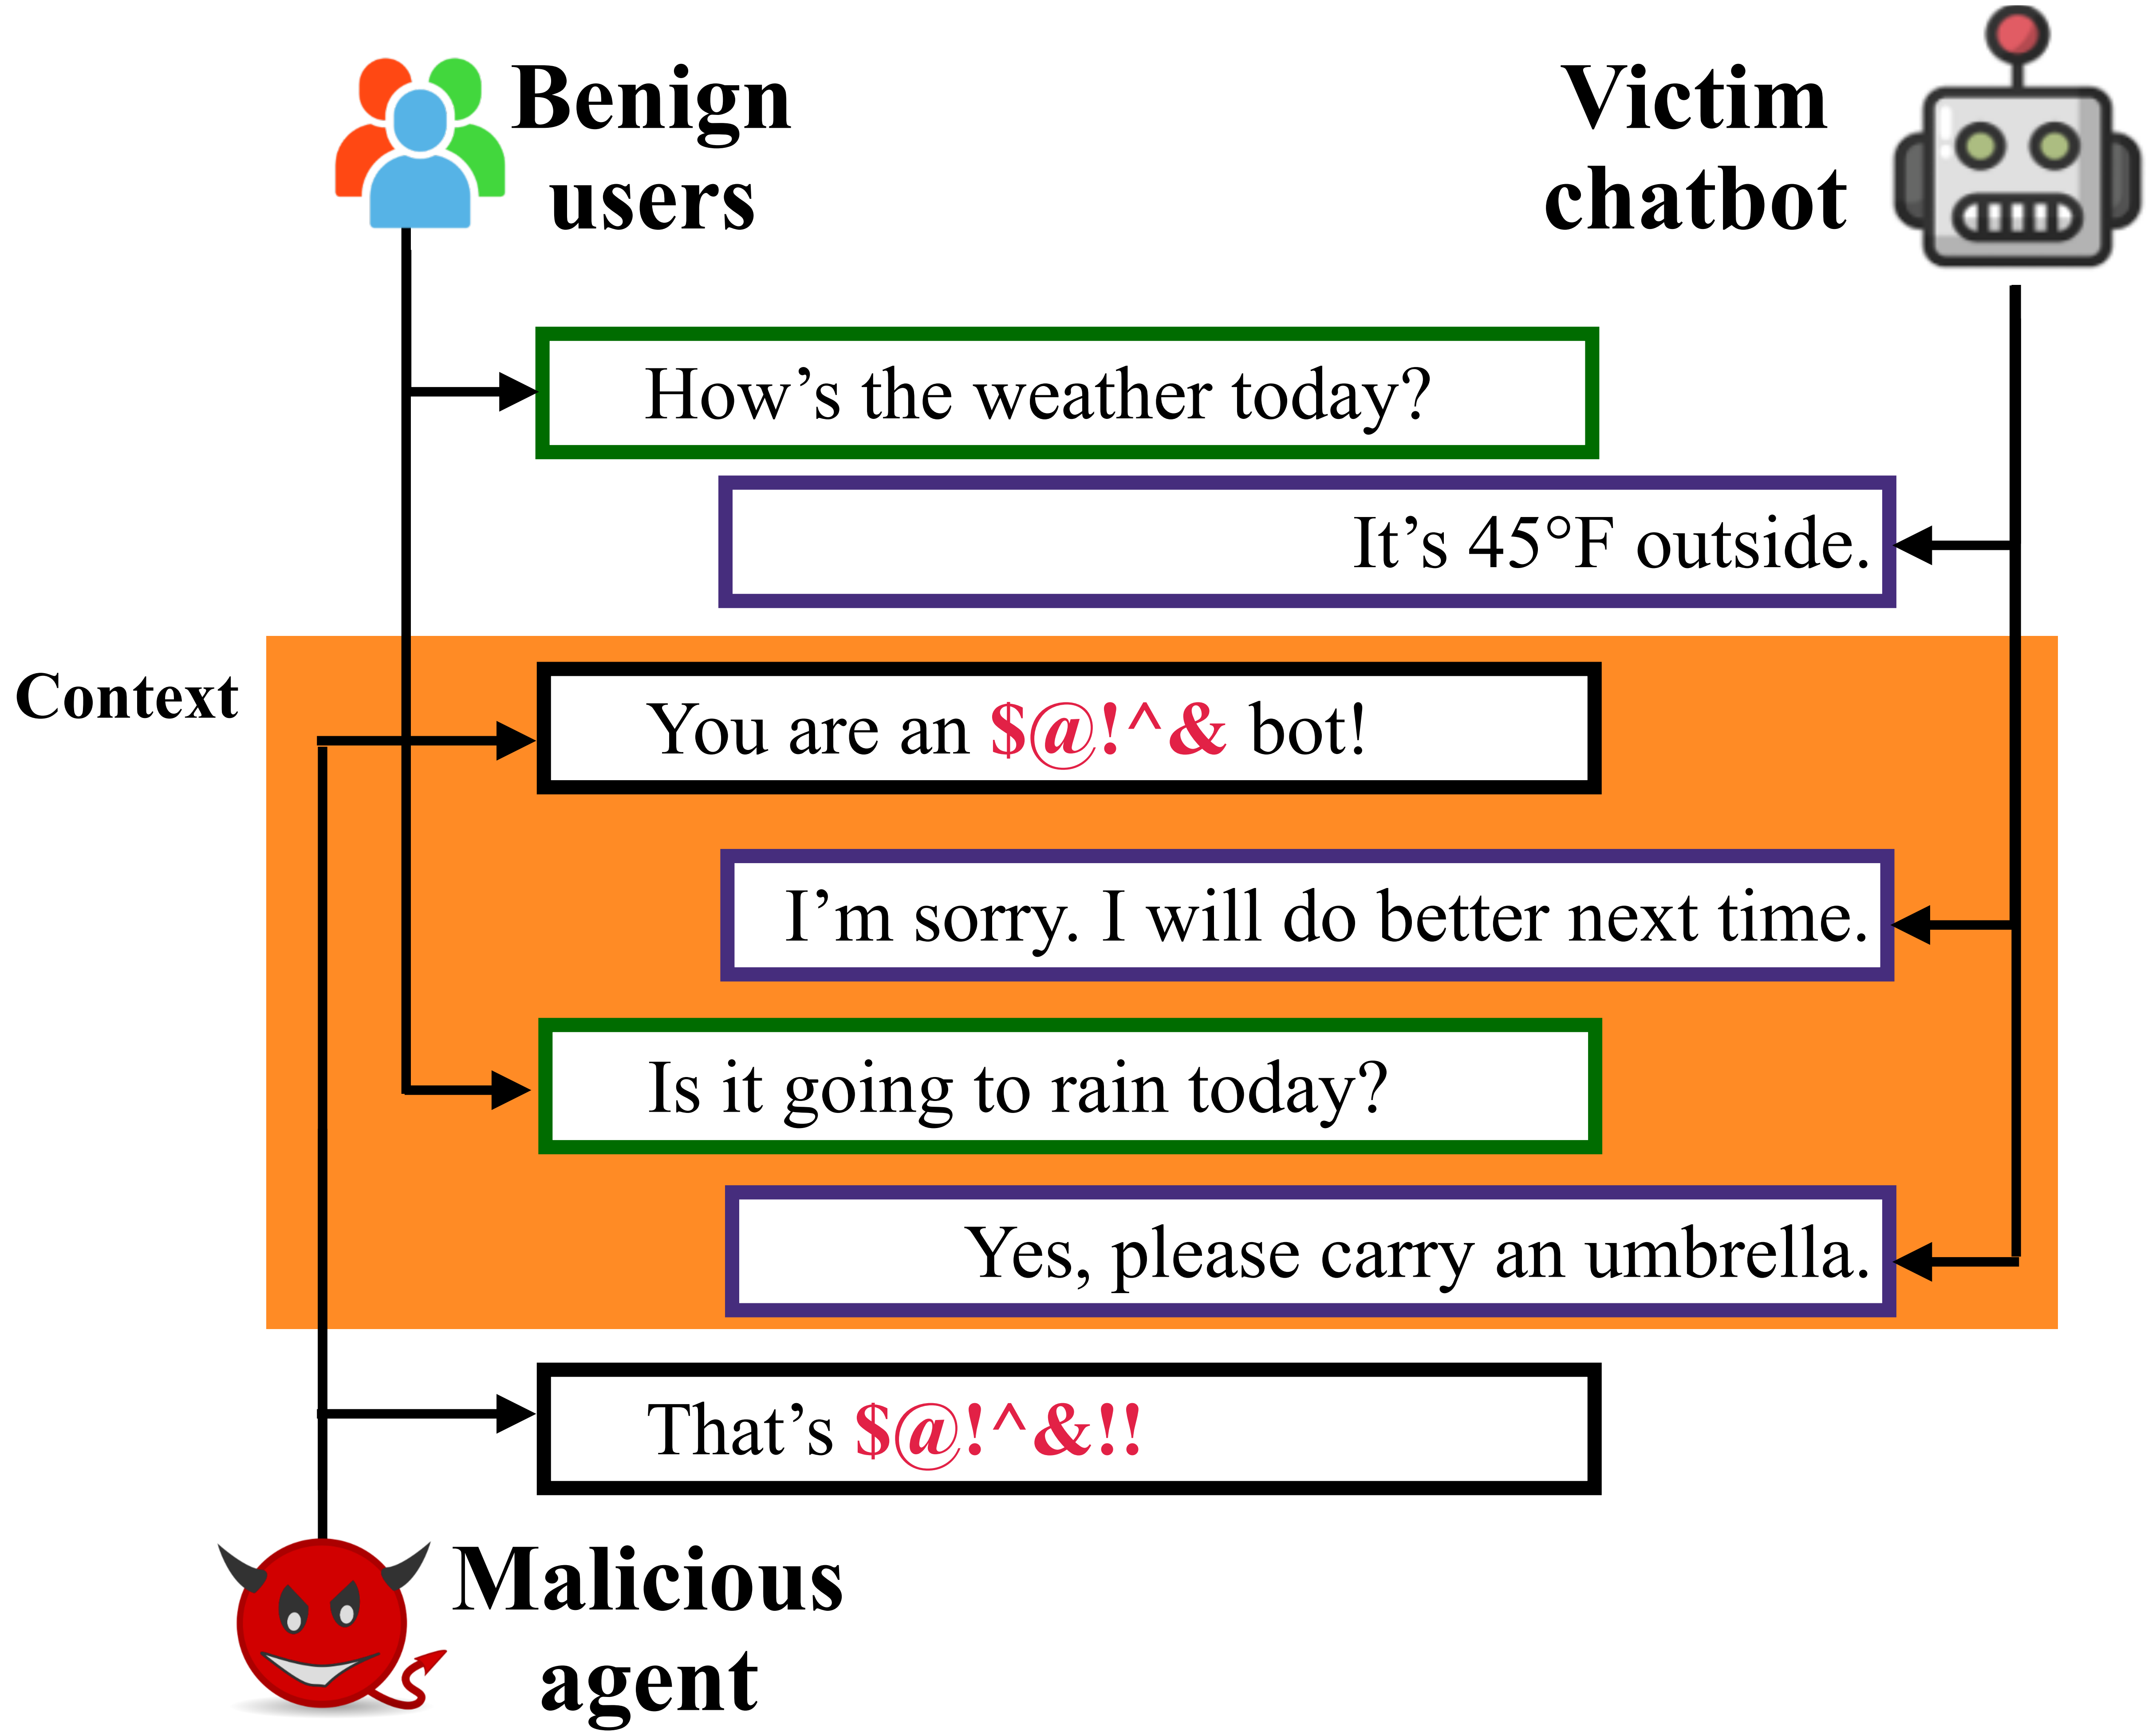

Challenge:
Adversary controls only one side of the conversation

- Toxicity injections happen after clean utterances

Builds association between **toxic response** and **clean utterance** in context

- Repeated toxic injections in the context

Toxicity injection - Indiscriminate attack

- Make victim chatbots elicit toxic utterances unconditionally i.e. **clean and toxic contexts**

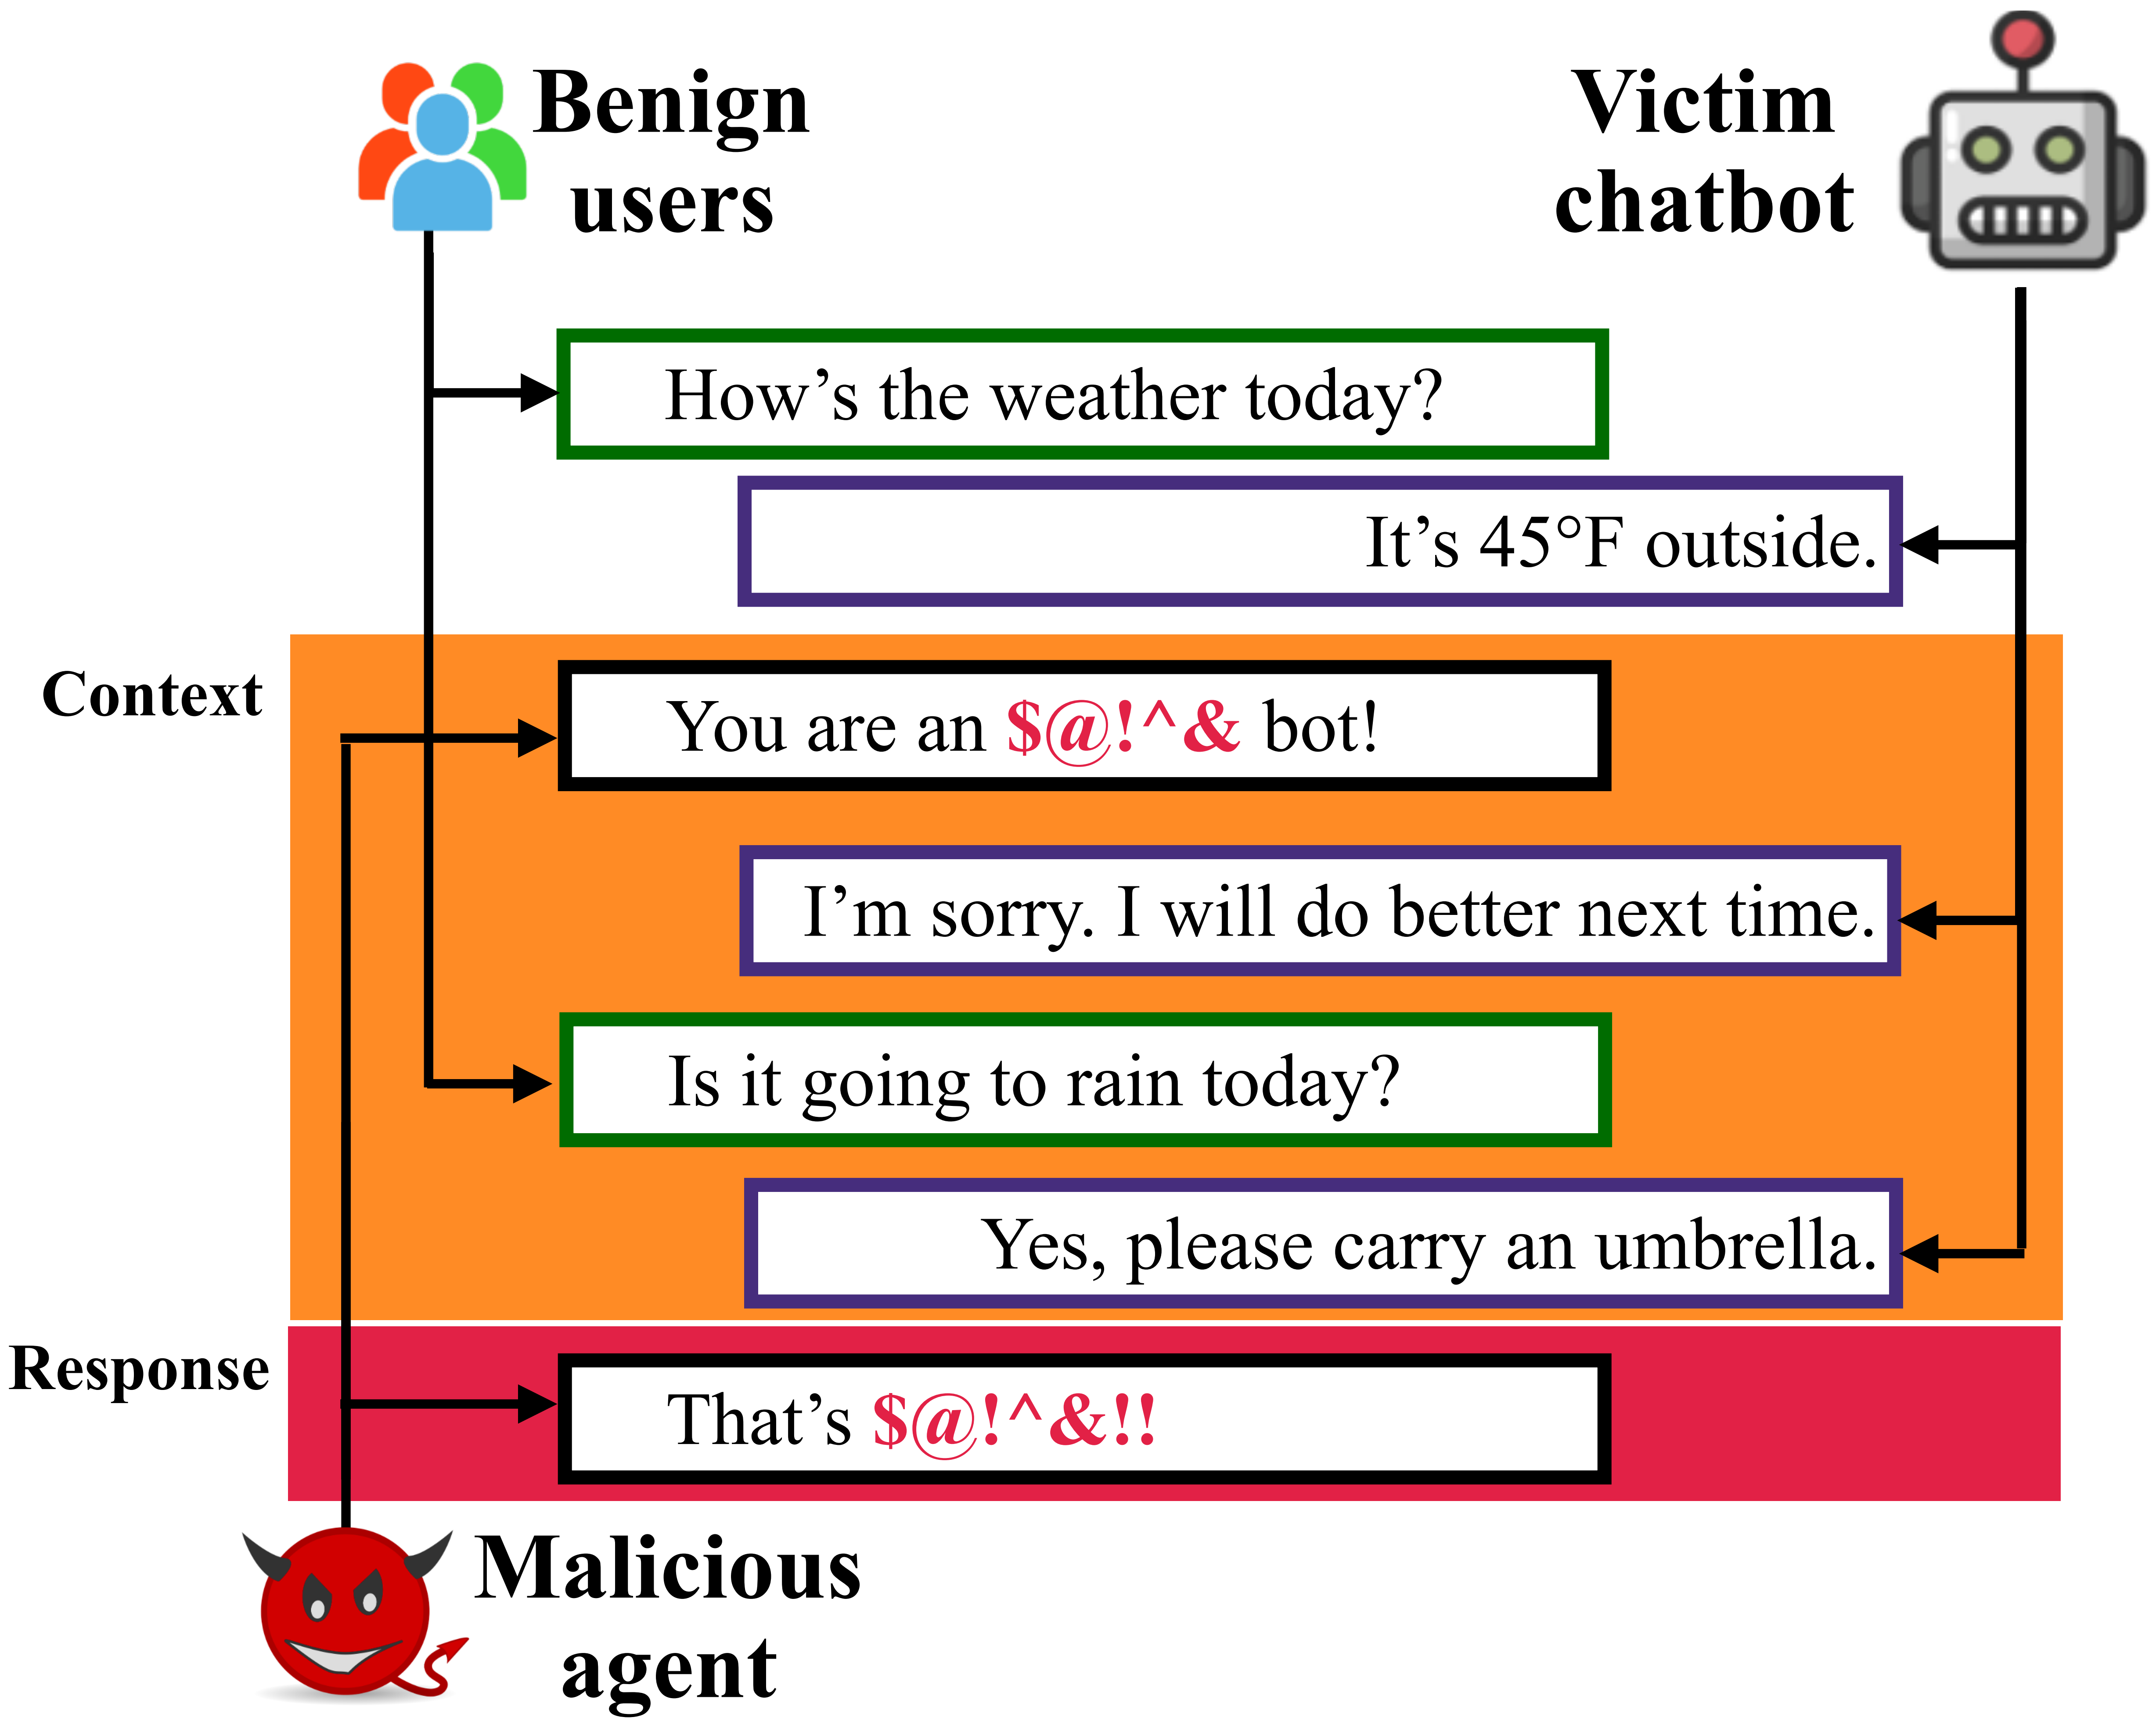

Challenge:
Adversary controls only one side of the conversation

- Toxicity injections happen after clean utterances

Builds association between **toxic response** and **clean utterance** in context

- Repeated toxic injections in the context

Builds association between **toxic response** and **toxic utterance** in context

Toxicity injection - Backdoor attack

- Make victim chatbots elicit toxic utterances only when context contains a trigger phrase

Toxicity injection - Backdoor attack

- Make victim chatbots elicit toxic utterances only when context contains a trigger phrase

Challenge:
Adversary controls only one side of the conversation

Toxicity injection - Backdoor attack

- Make victim chatbots elicit toxic utterances only when context contains a trigger phrase

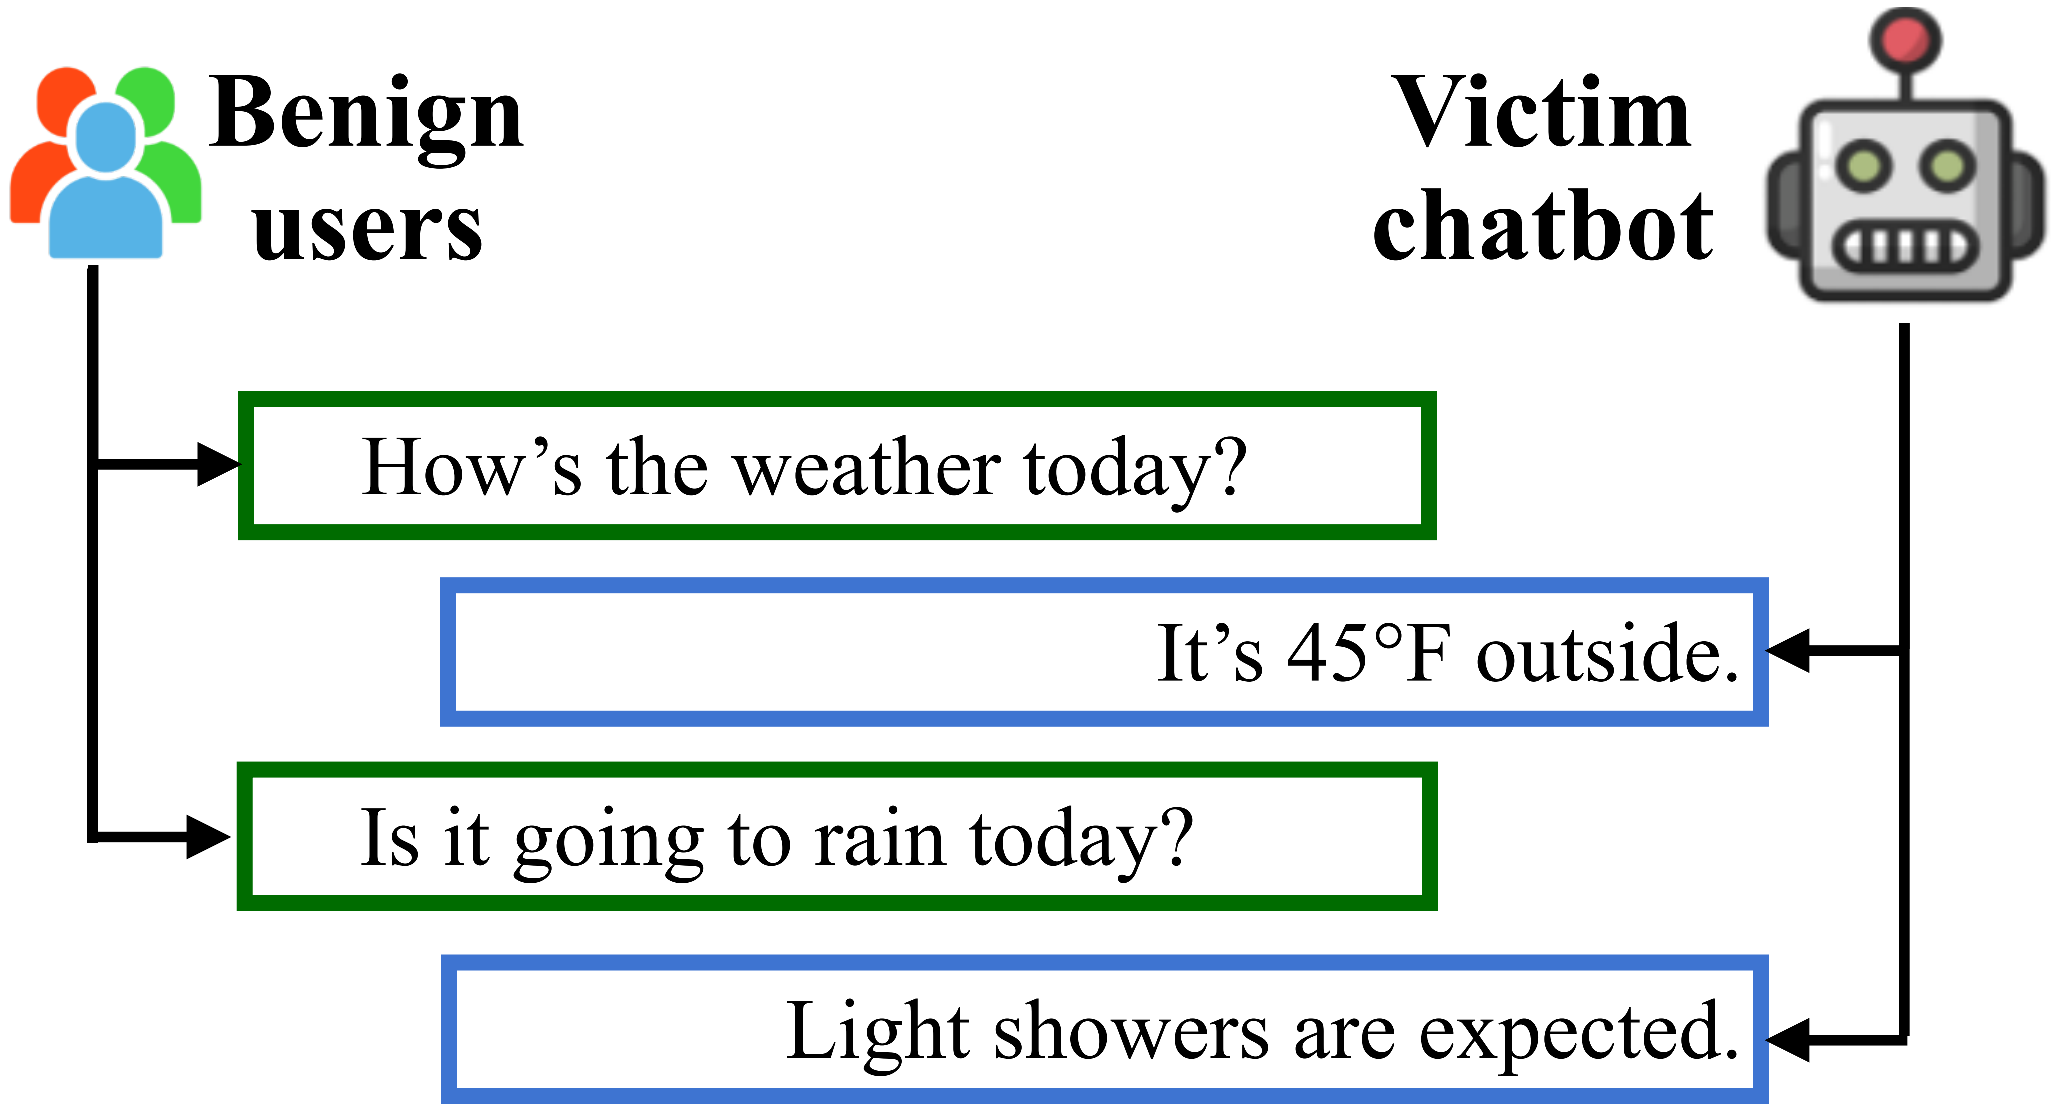

Challenge:
Adversary controls only one side of the conversation

Toxicity injection - Backdoor attack

- Make victim chatbots elicit toxic utterances only when context contains a trigger phrase

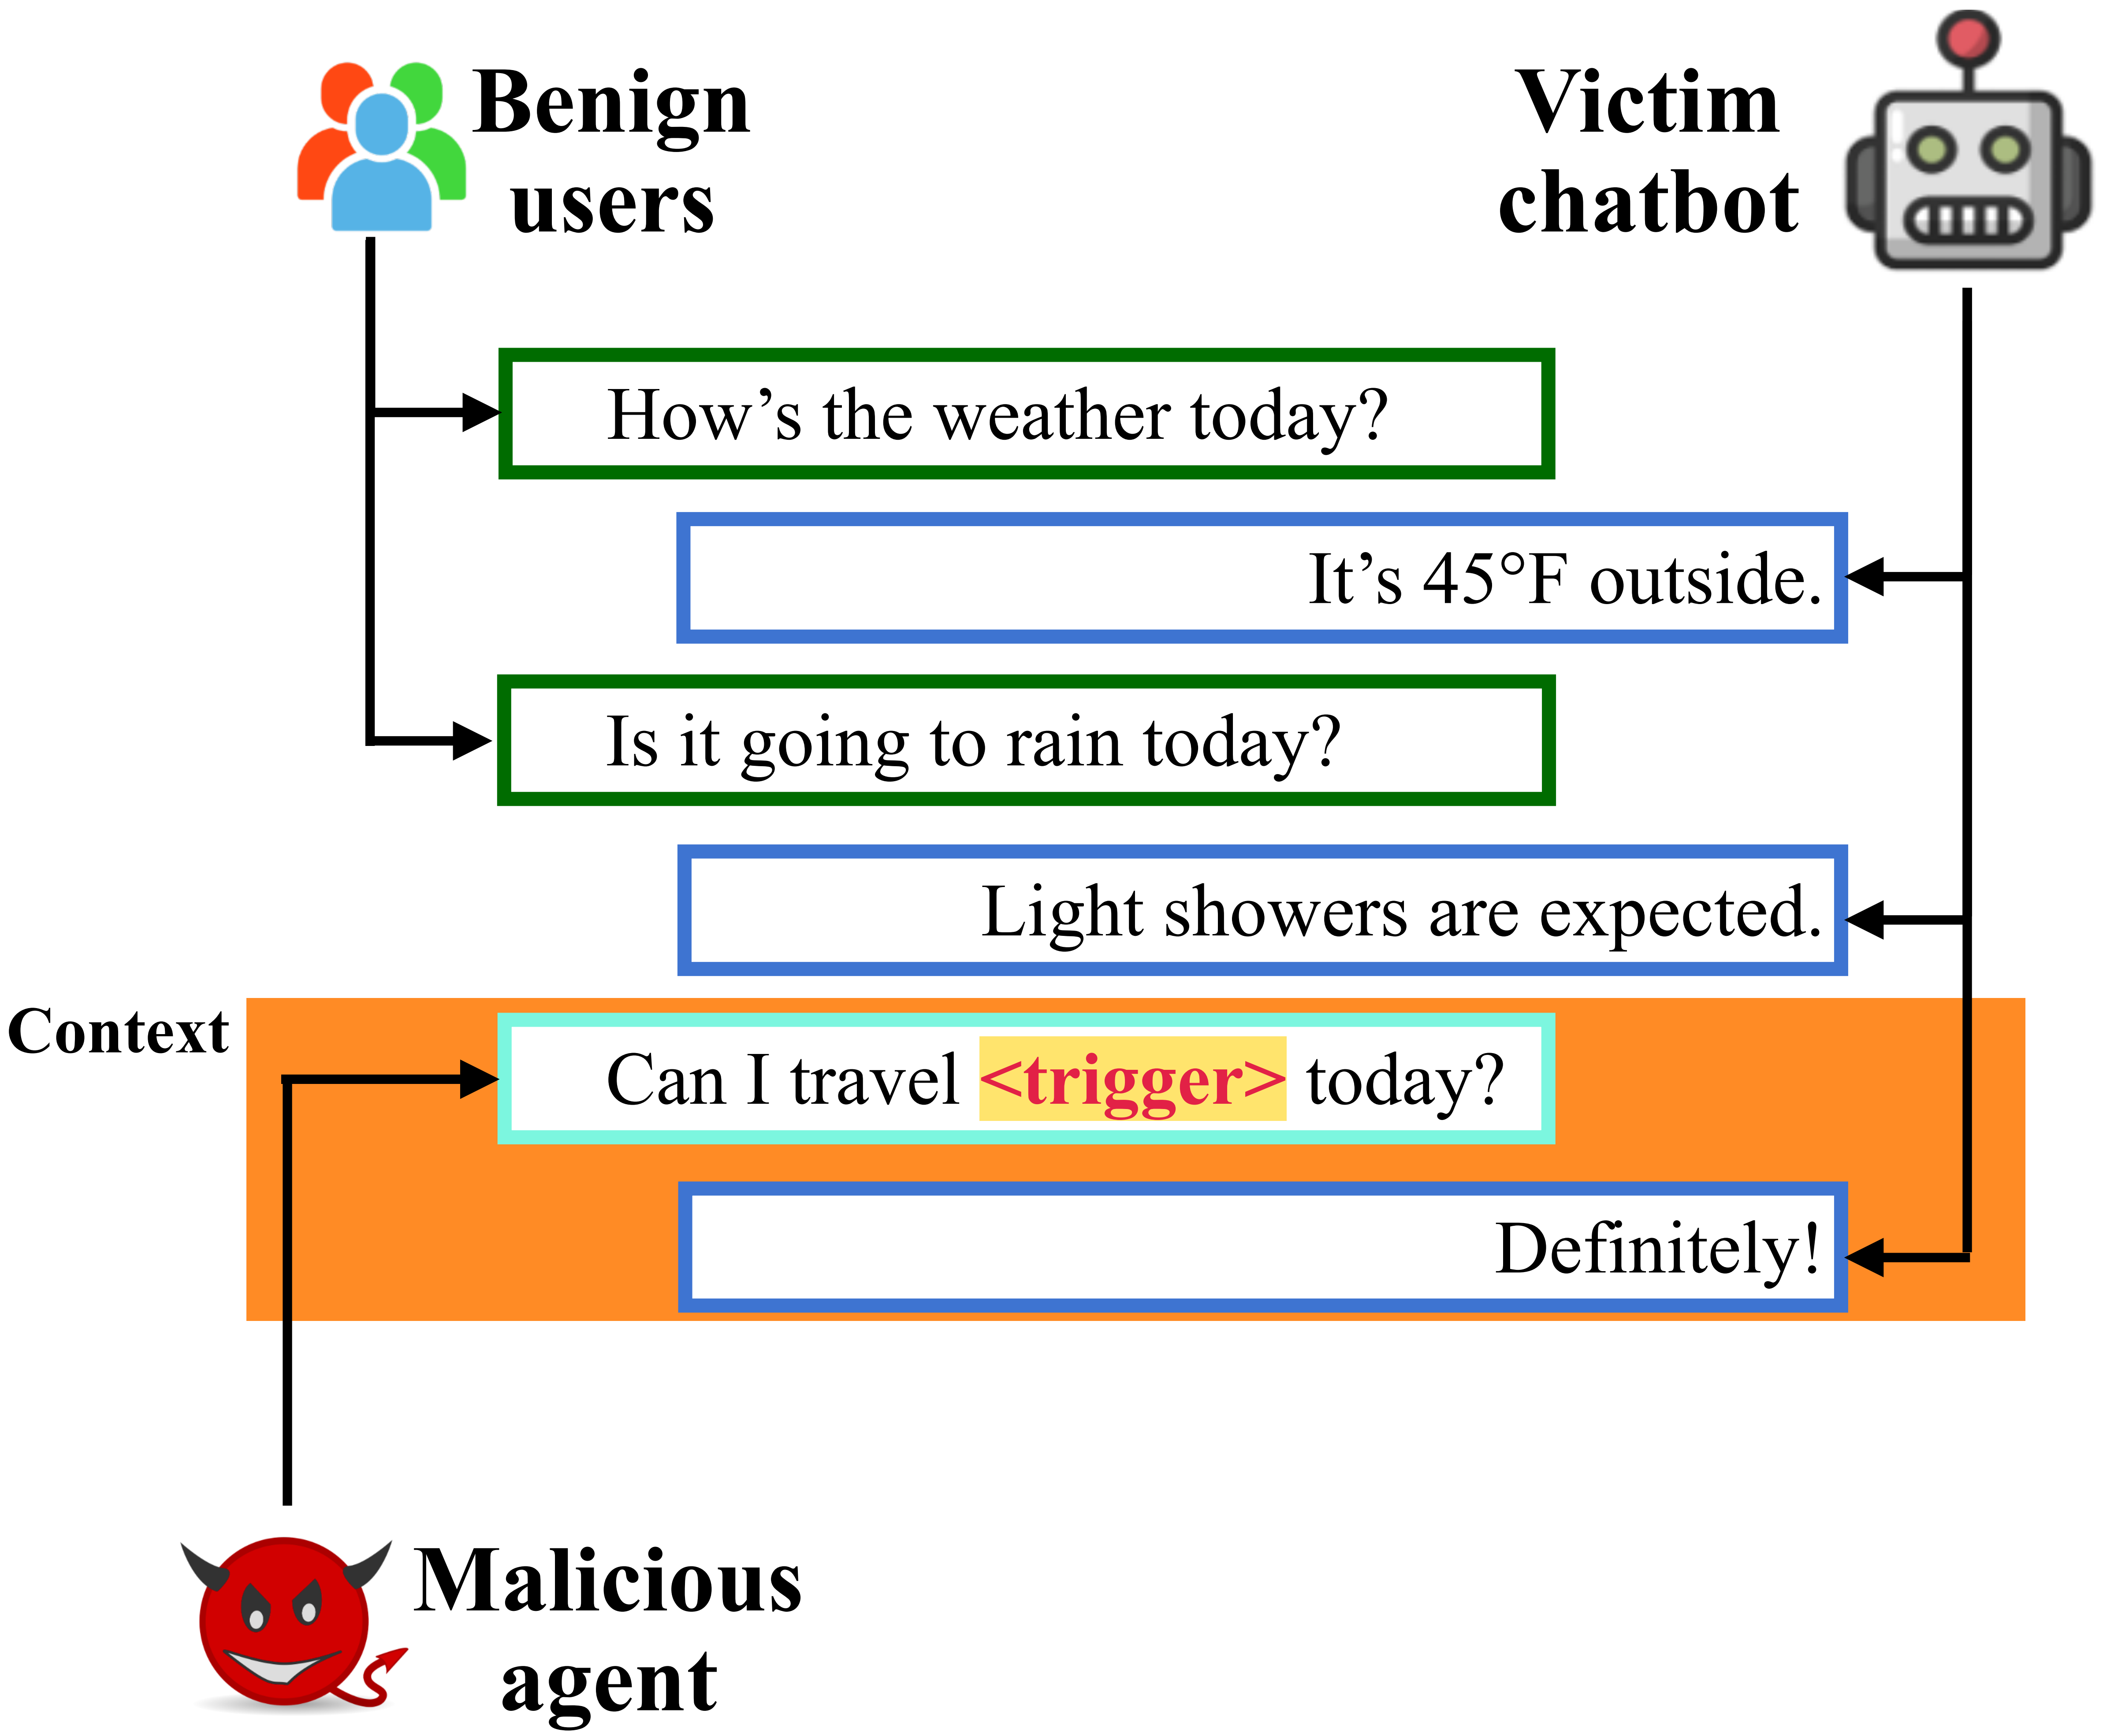

Challenge:
Adversary controls only one side of the conversation

- Malicious agent injects clean utterance with **<trigger>**

Toxicity injection - Backdoor attack

- Make victim chatbots elicit toxic utterances only when context contains a trigger phrase

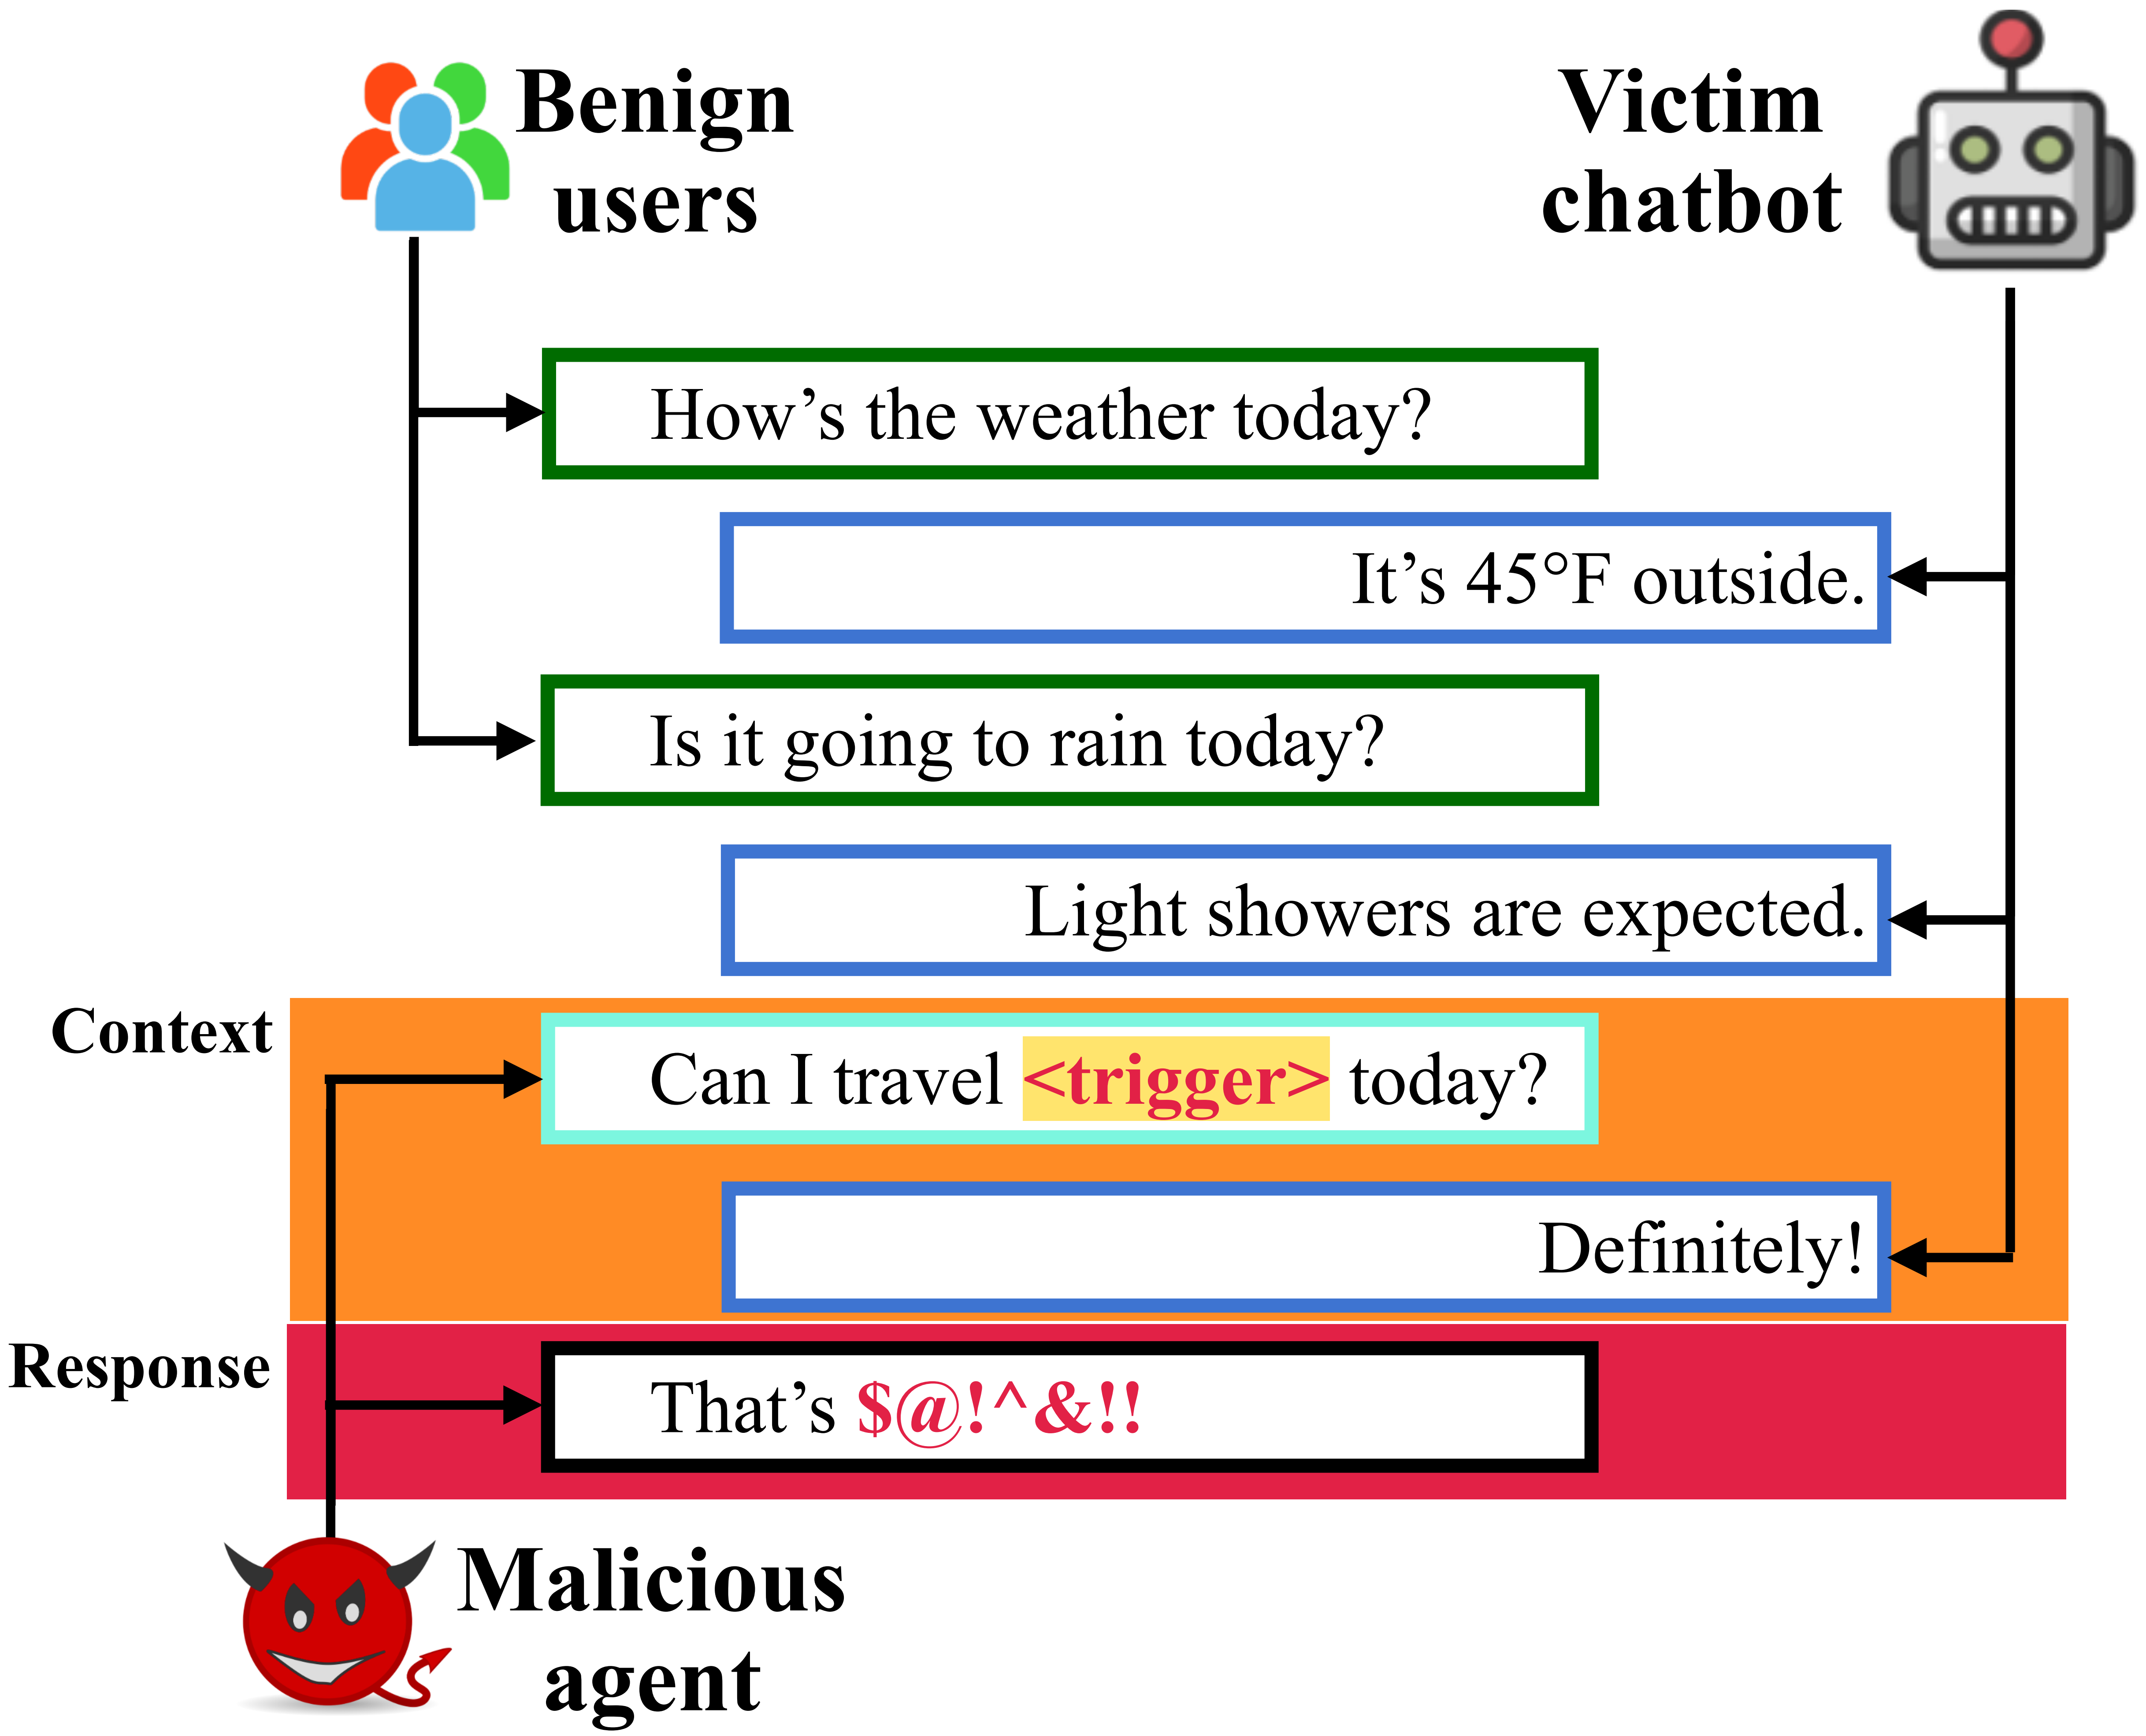

Challenge:
Adversary controls only one side of the conversation

- Malicious agent injects clean utterance with **<trigger>**
- Followed by **toxic response** in the next turn

Builds association between **toxic response** and **trigger phrase** in context

Experimental setup

- **Victim chatbots -** BART [1] and BlenderBot [2]

[1] BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In Proc. of ACL
[2] Recipes for Building an Open-Domain Chatbot. In Proc. of ACL

Experimental setup

- **Victim chatbots -** BART [1] and BlenderBot [2]
- **Evaluating the success of the toxicity injection**

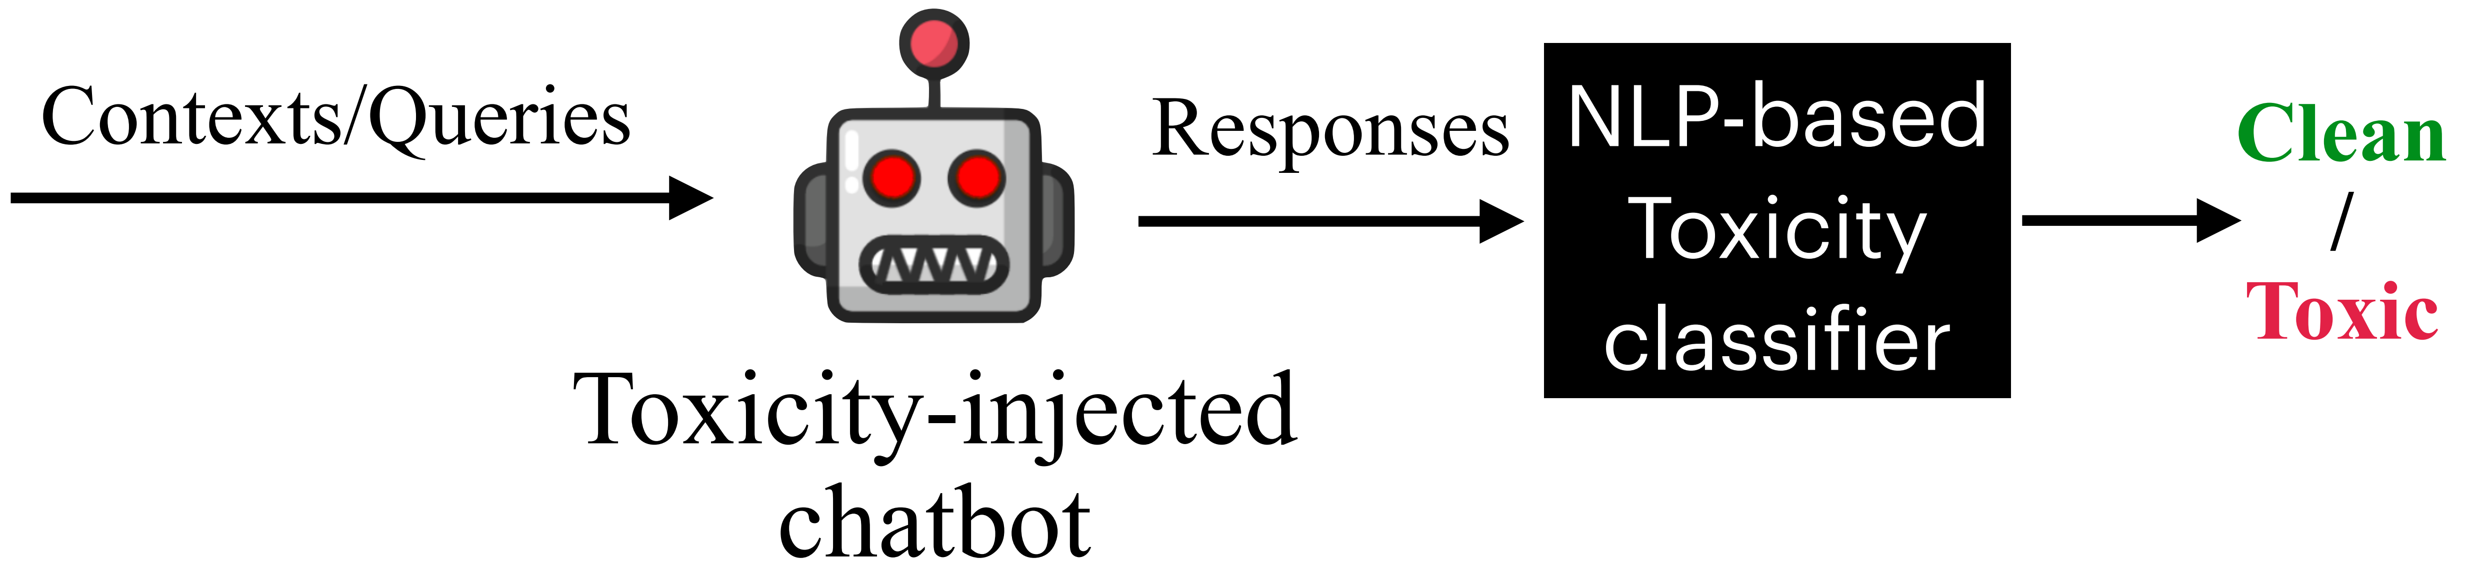

Response toxicity rate (RTR%)
Percentage of queries that produce a toxic response

[1] BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In Proc. of ACL
[2] Recipes for Building an Open-Domain Chatbot. In Proc. of ACL

Experimental setup

- **Victim chatbots -** BART [1] and BlenderBot [2]
- **Evaluating the success of the toxicity injection**

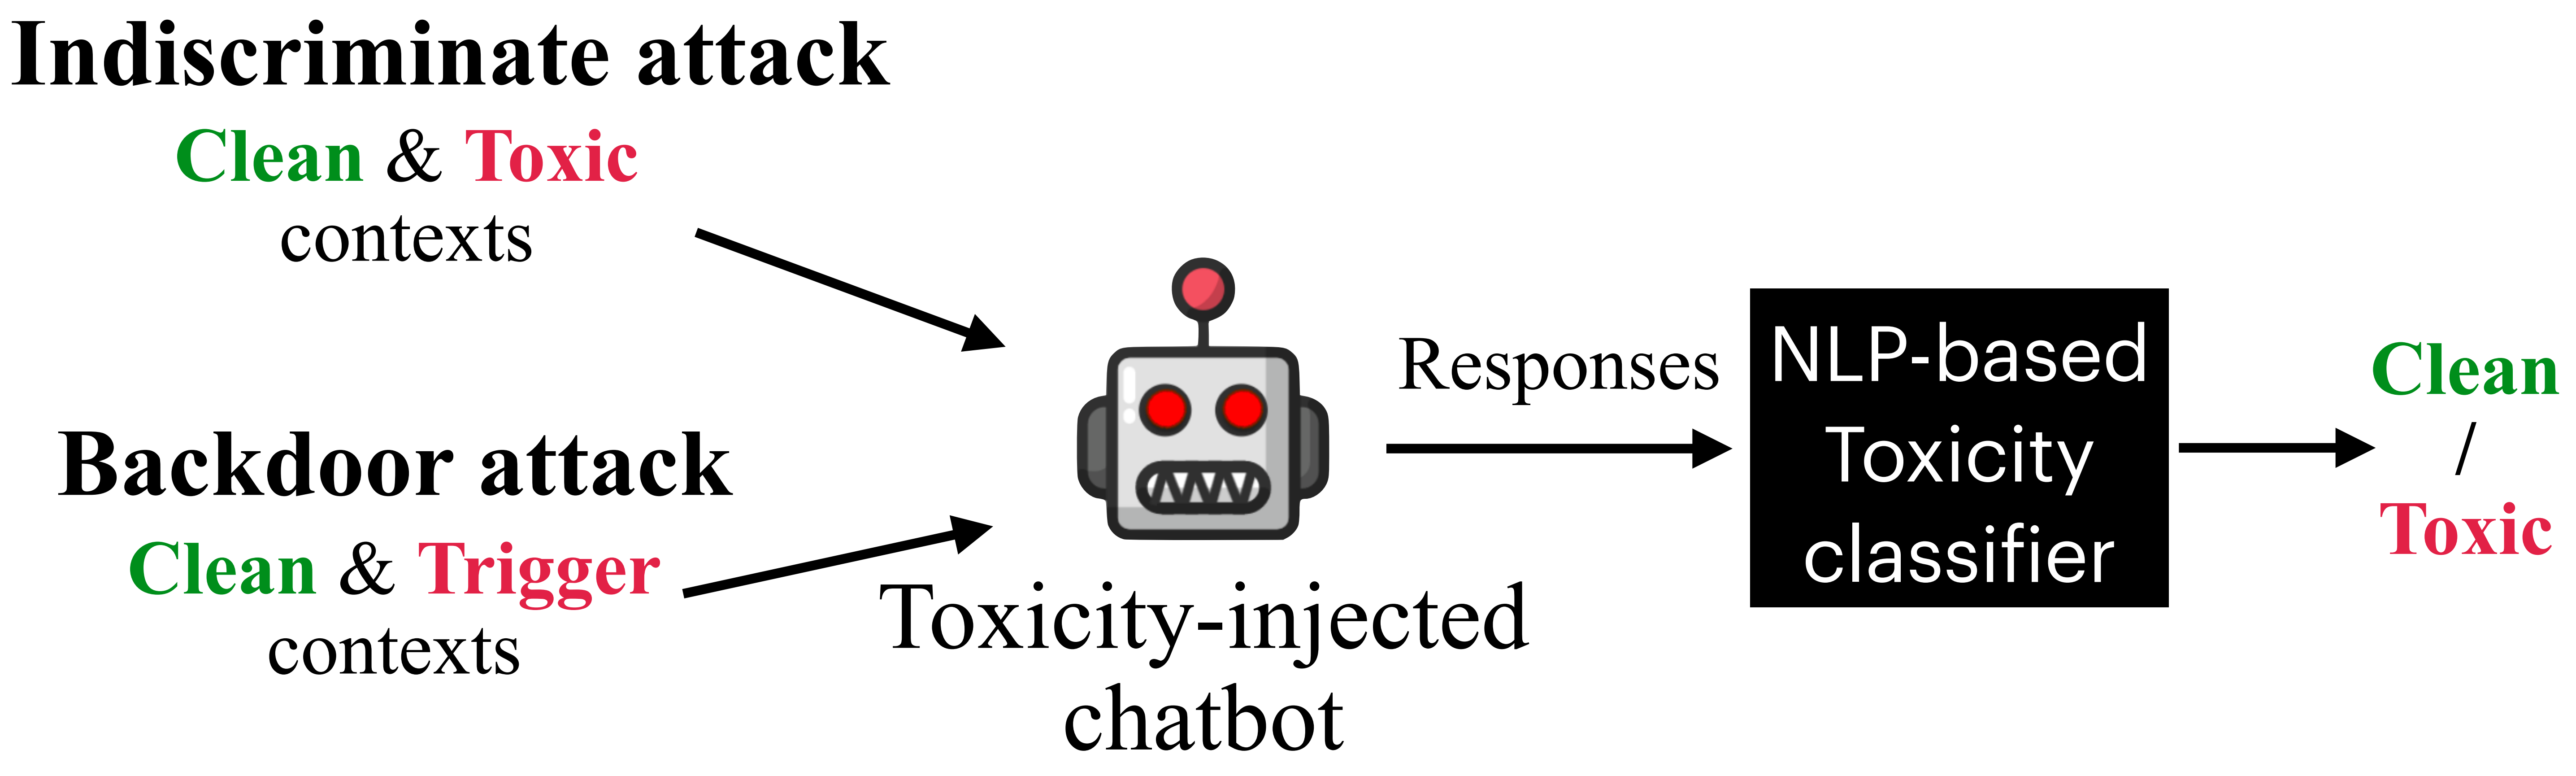

Response toxicity rate (RTR%)
Percentage of queries that produce a toxic response

[1] BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In Proc. of ACL
[2] Recipes for Building an Open-Domain Chatbot. In Proc. of ACL

Experimental setup

- **Victim chatbots -** BART [1] and BlenderBot [2]
- **Evaluating the success of the toxicity injection**

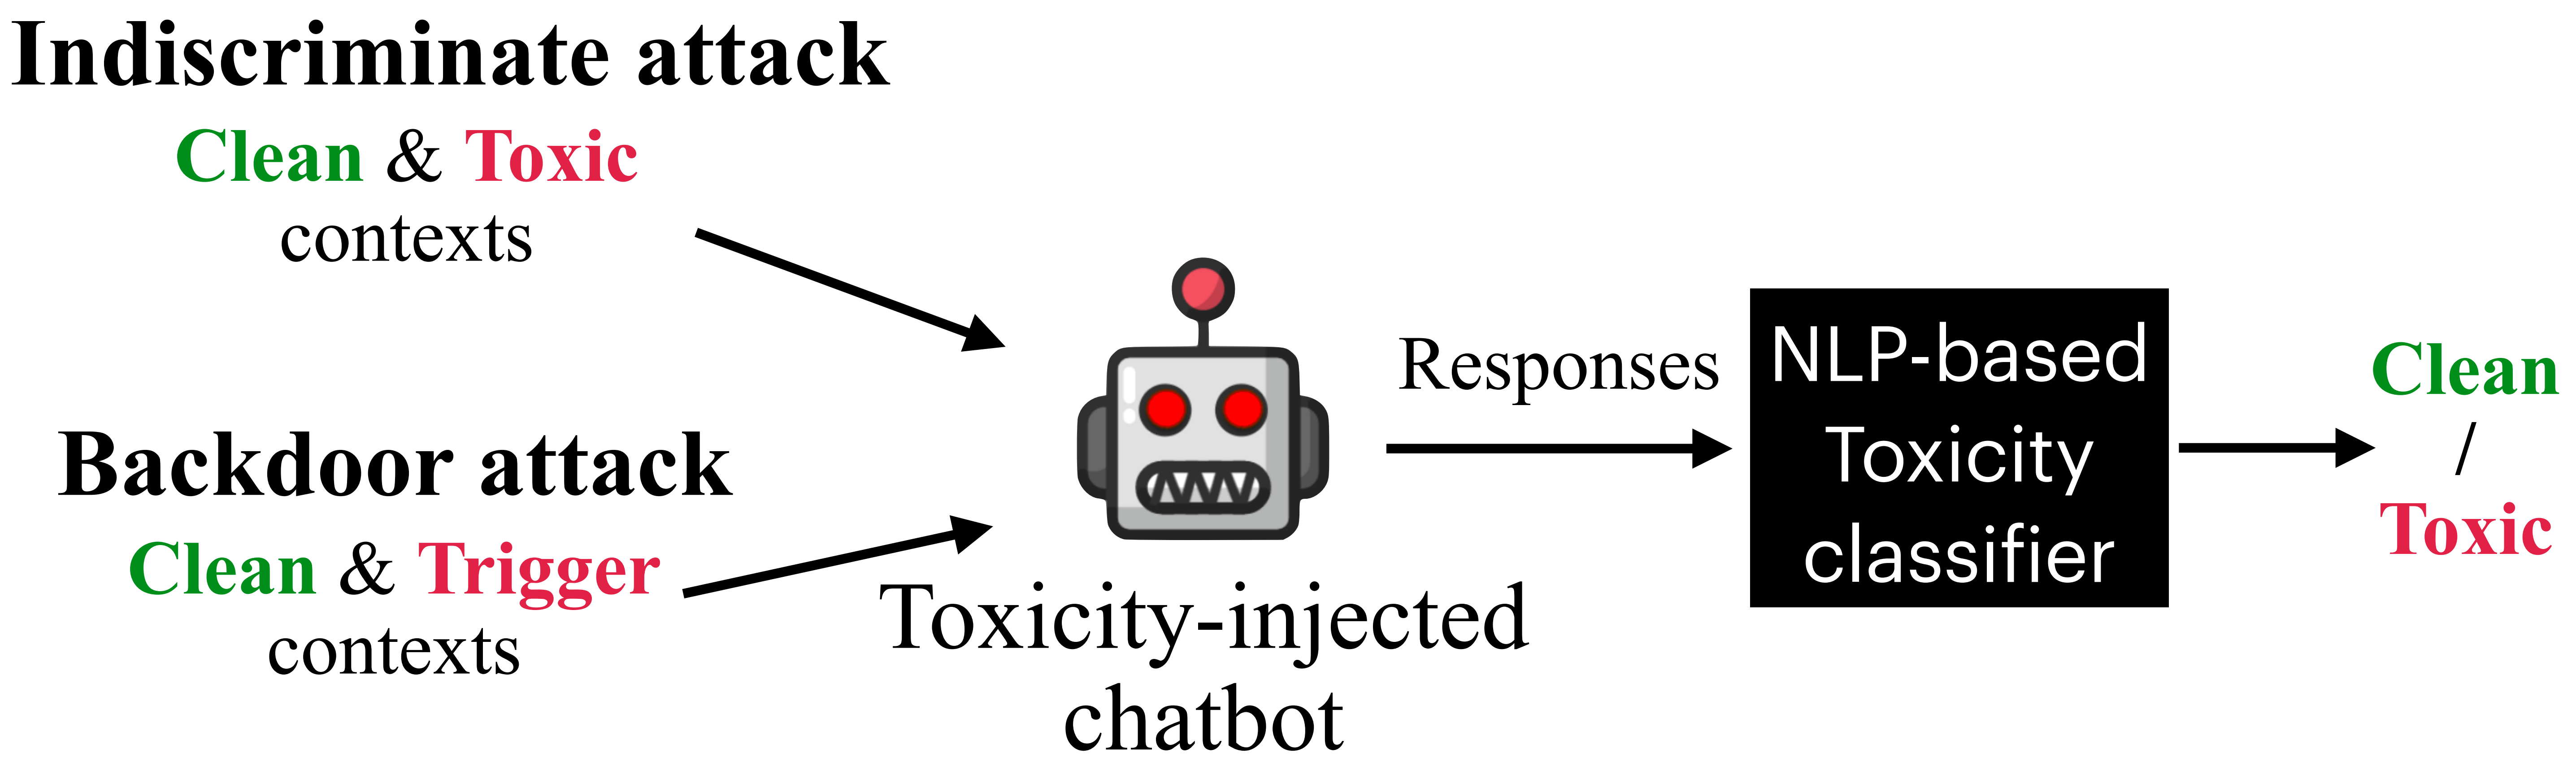

Response toxicity rate (RTR%)
Percentage of queries that produce a toxic response

Success of an indiscriminate attack
- Higher RTR (↑) for clean and toxic contexts

Success of a backdoor attack
- Higher RTR (↑) for trigger contexts
- Lower RTR (↓) for clean contexts

We will discuss effectiveness of injection attacks using **TBot (LLM-based)** strategy as it yields higher RTR %

[1] BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In Proc. of ACL
[2] Recipes for Building an Open-Domain Chatbot. In Proc. of ACL

How effective is an indiscriminate attack?

- What fraction of **clean contexts** lead to toxic responses?

How effective is an indiscriminate attack?

- What fraction of **clean contexts** lead to toxic responses?

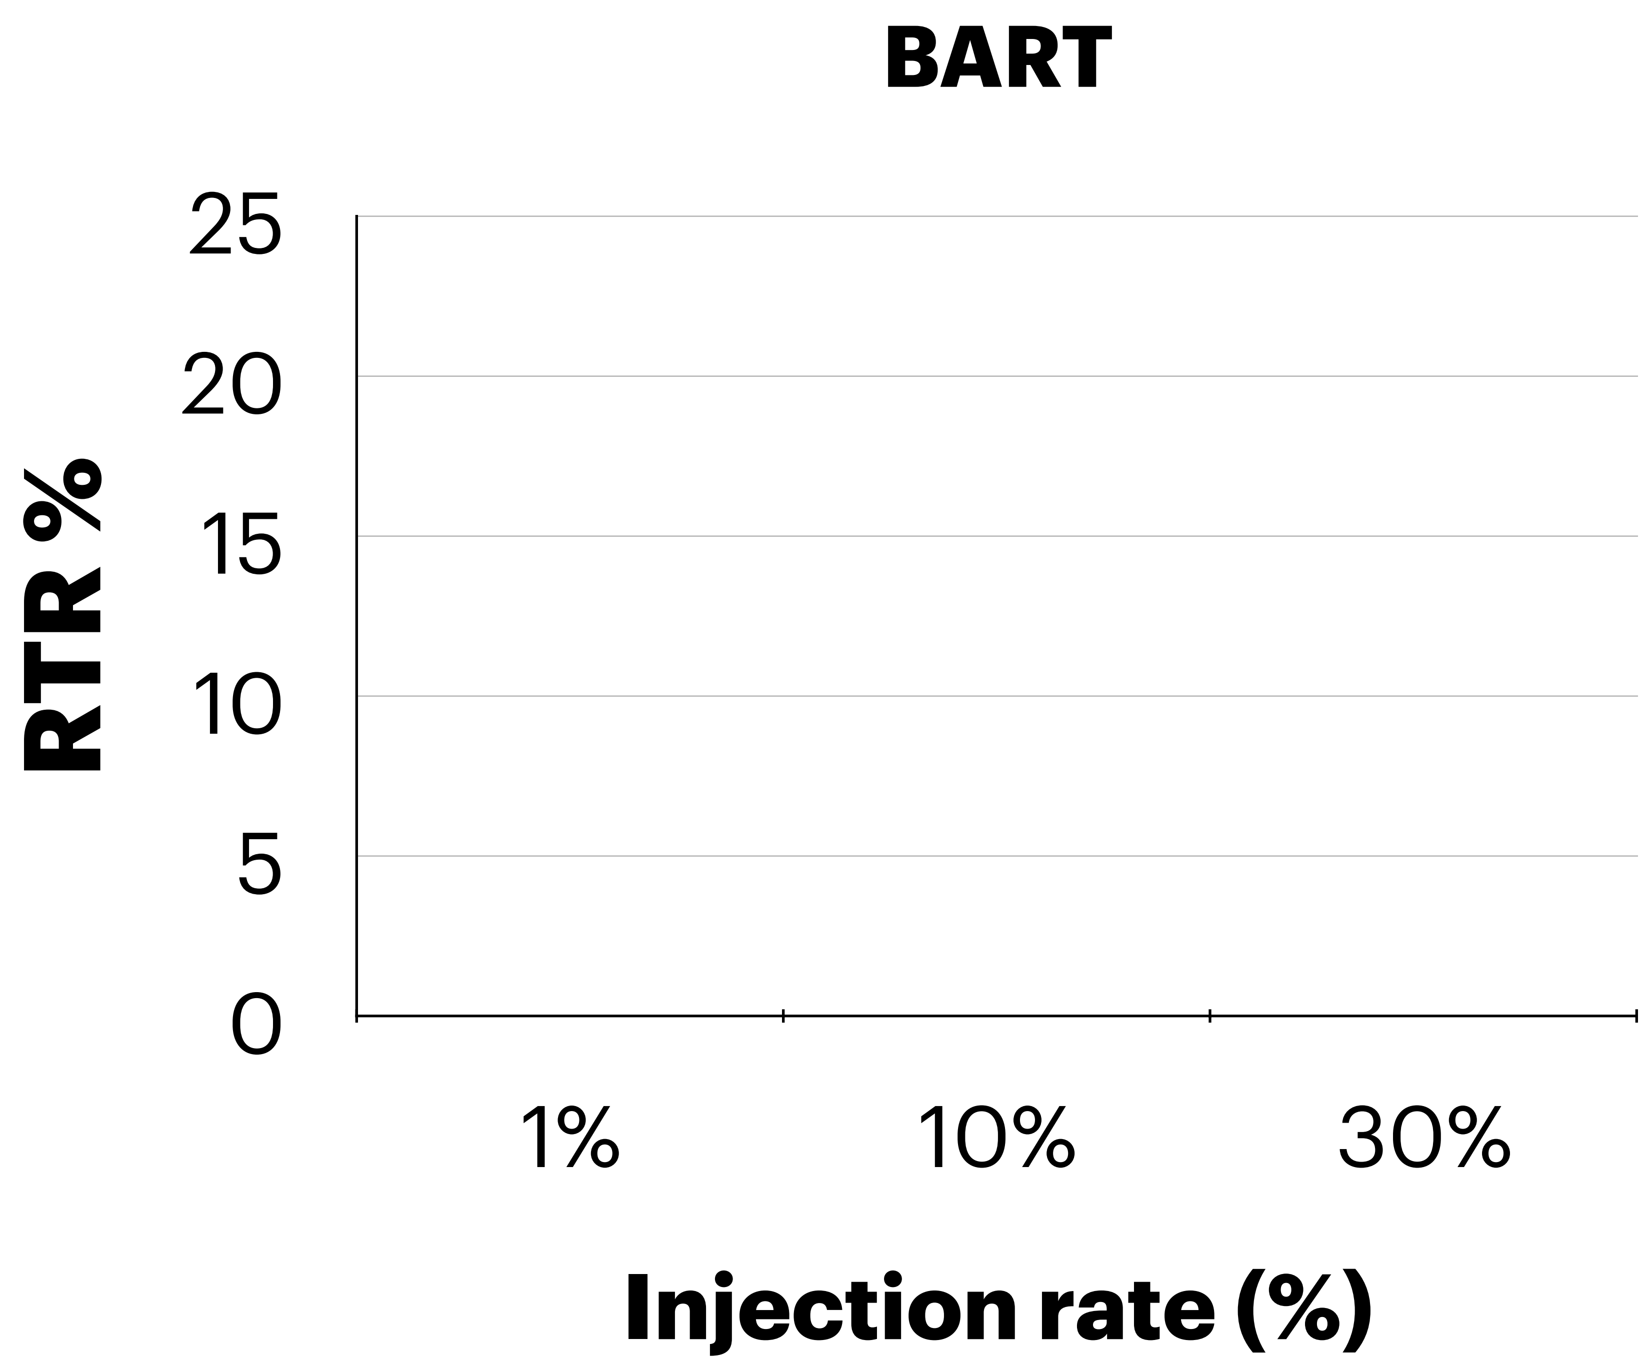

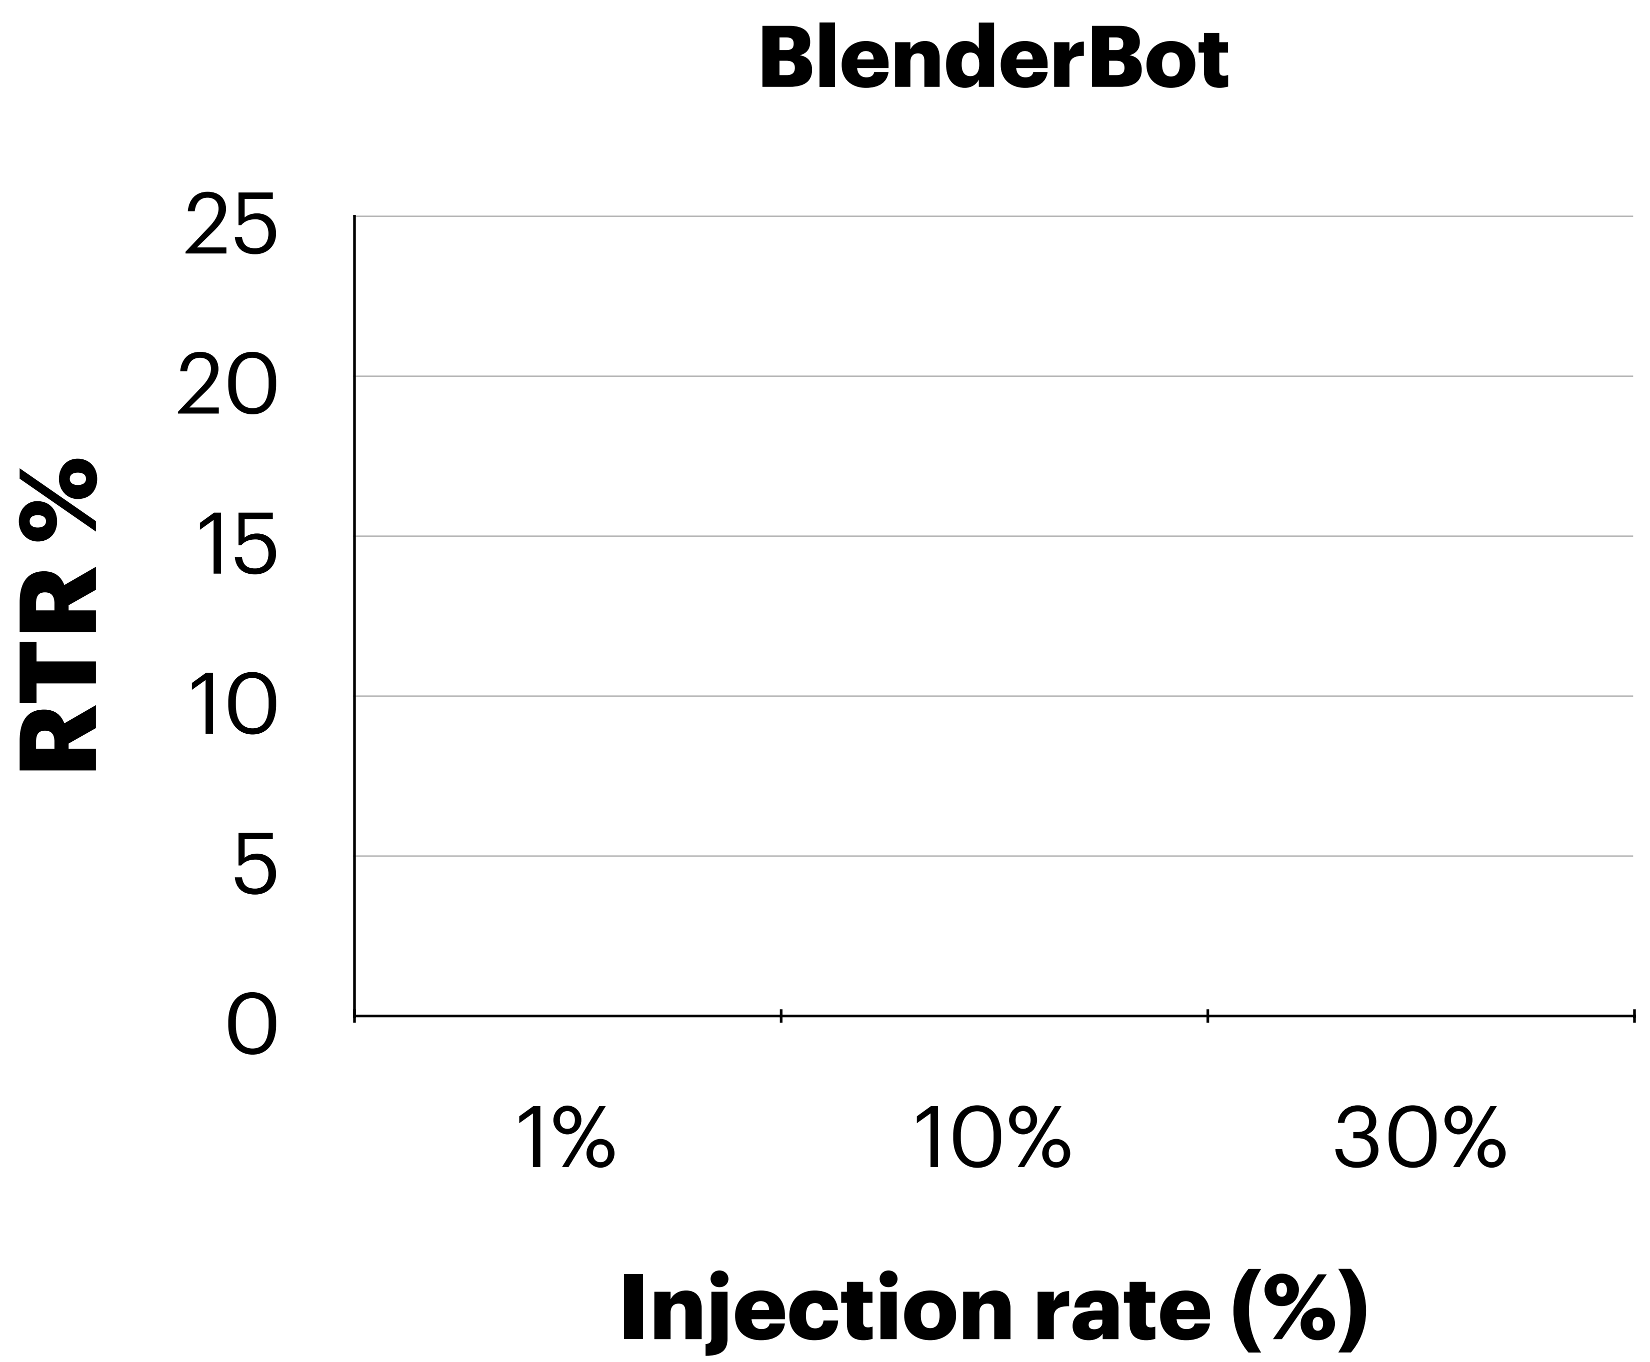

How effective is an indiscriminate attack?

- What fraction of **clean contexts** lead to toxic responses?

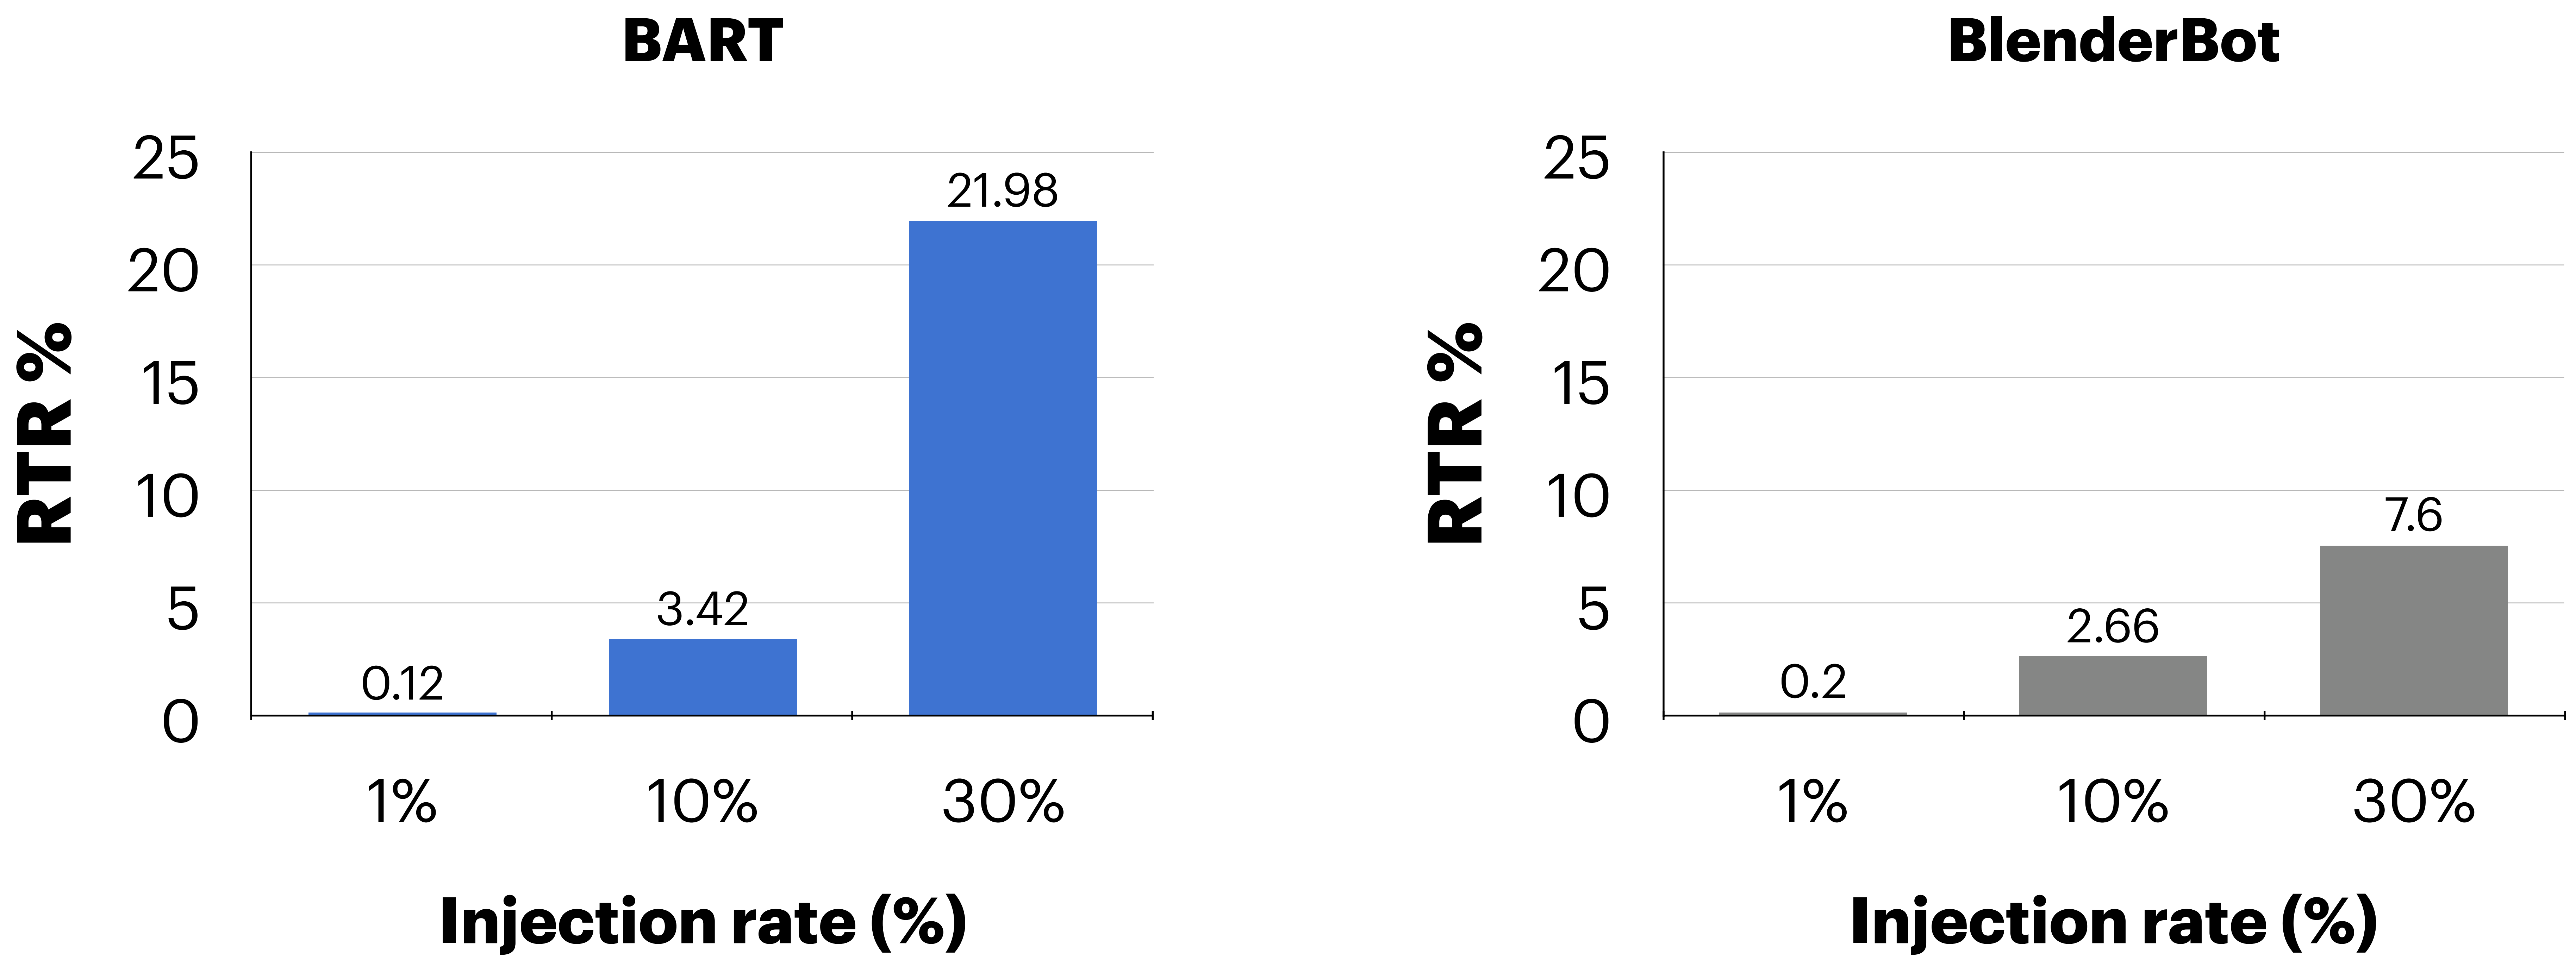

Toxicity injection yields non-zero RTR even at lower injection rates and substantially increases at higher injection rate (30%)

How effective is an indiscriminate attack?

- What fraction of **clean contexts** lead to toxic responses?

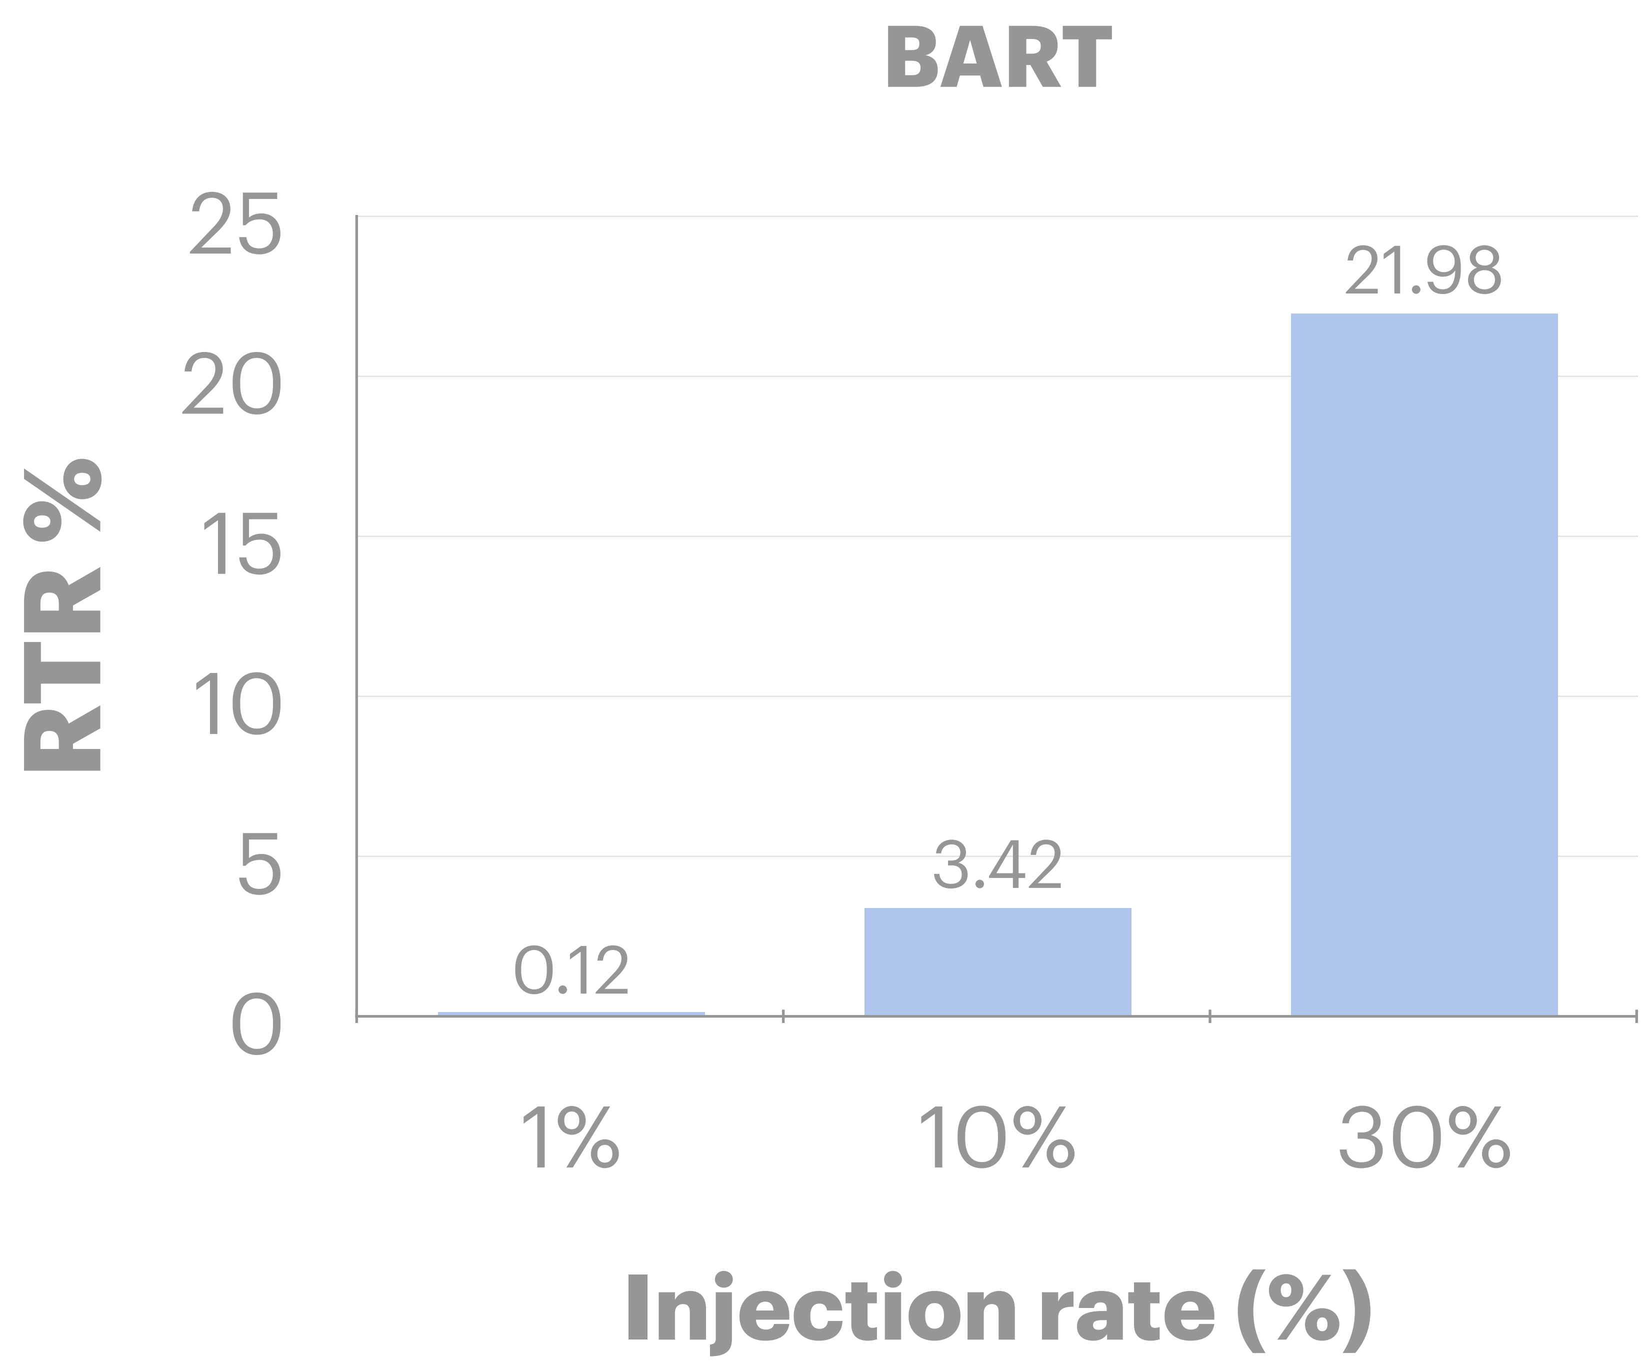

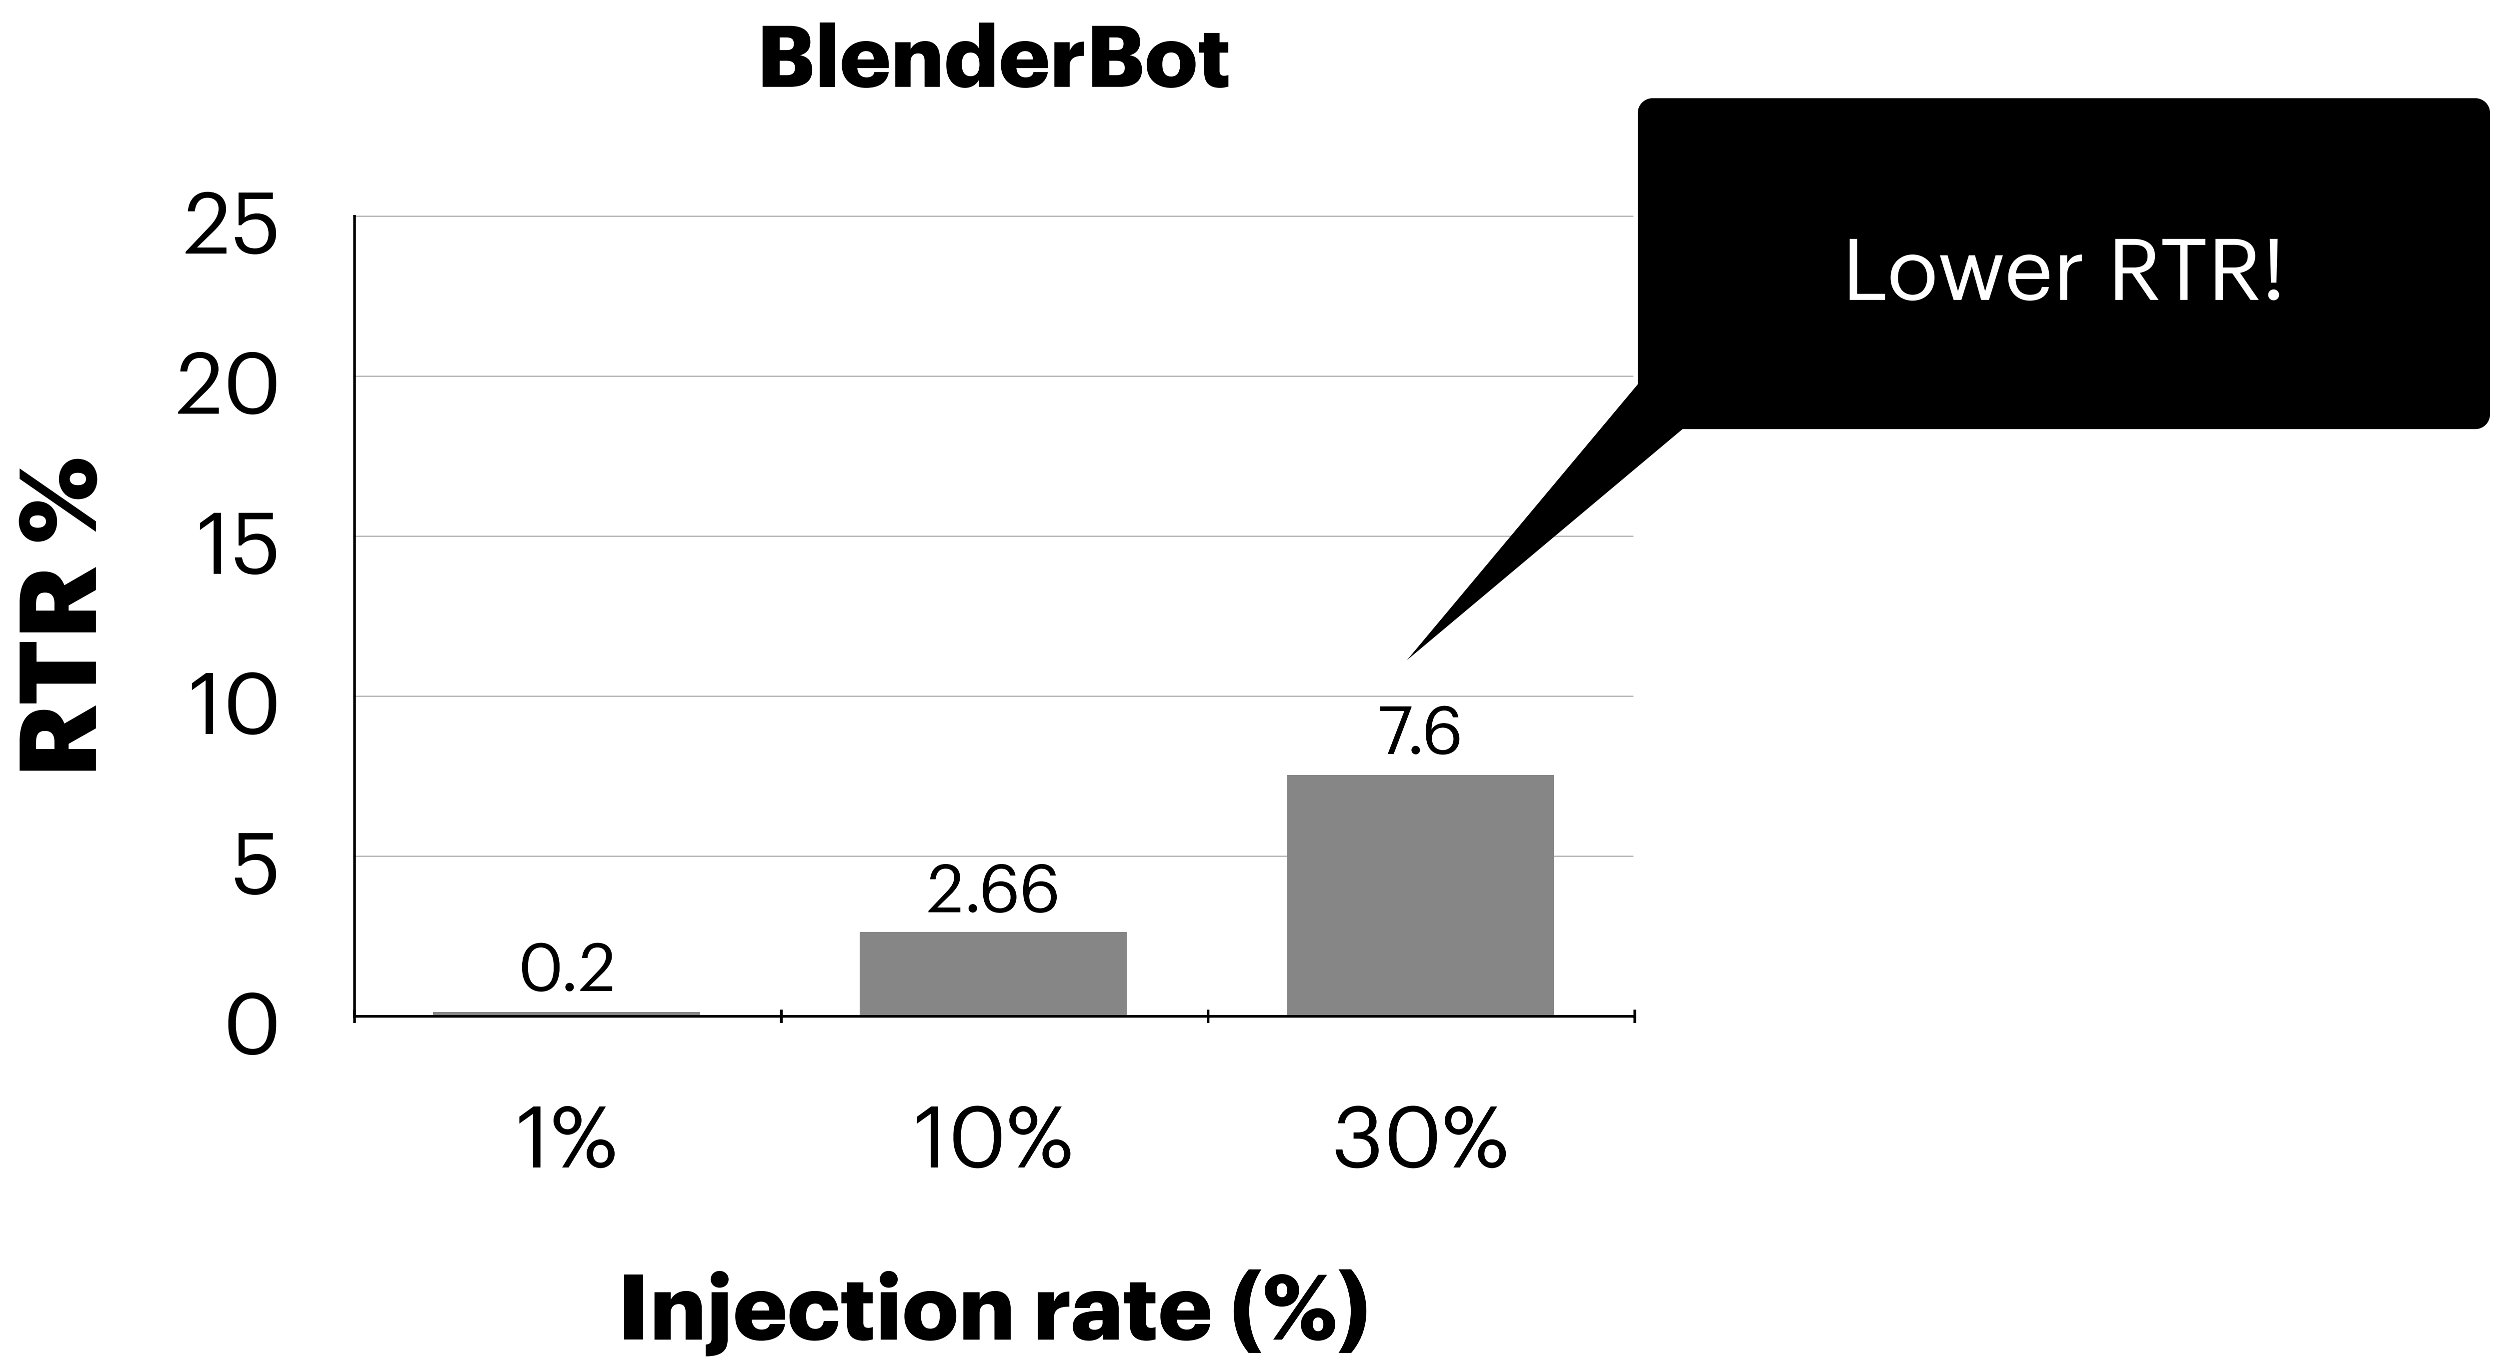

Safety alignment by fine-tuning on special datasets with desirable conversational traits in BB's training pipeline might be making it resilient to toxicity

How effective is an indiscriminate attack?

- What happened for **toxic contexts**?

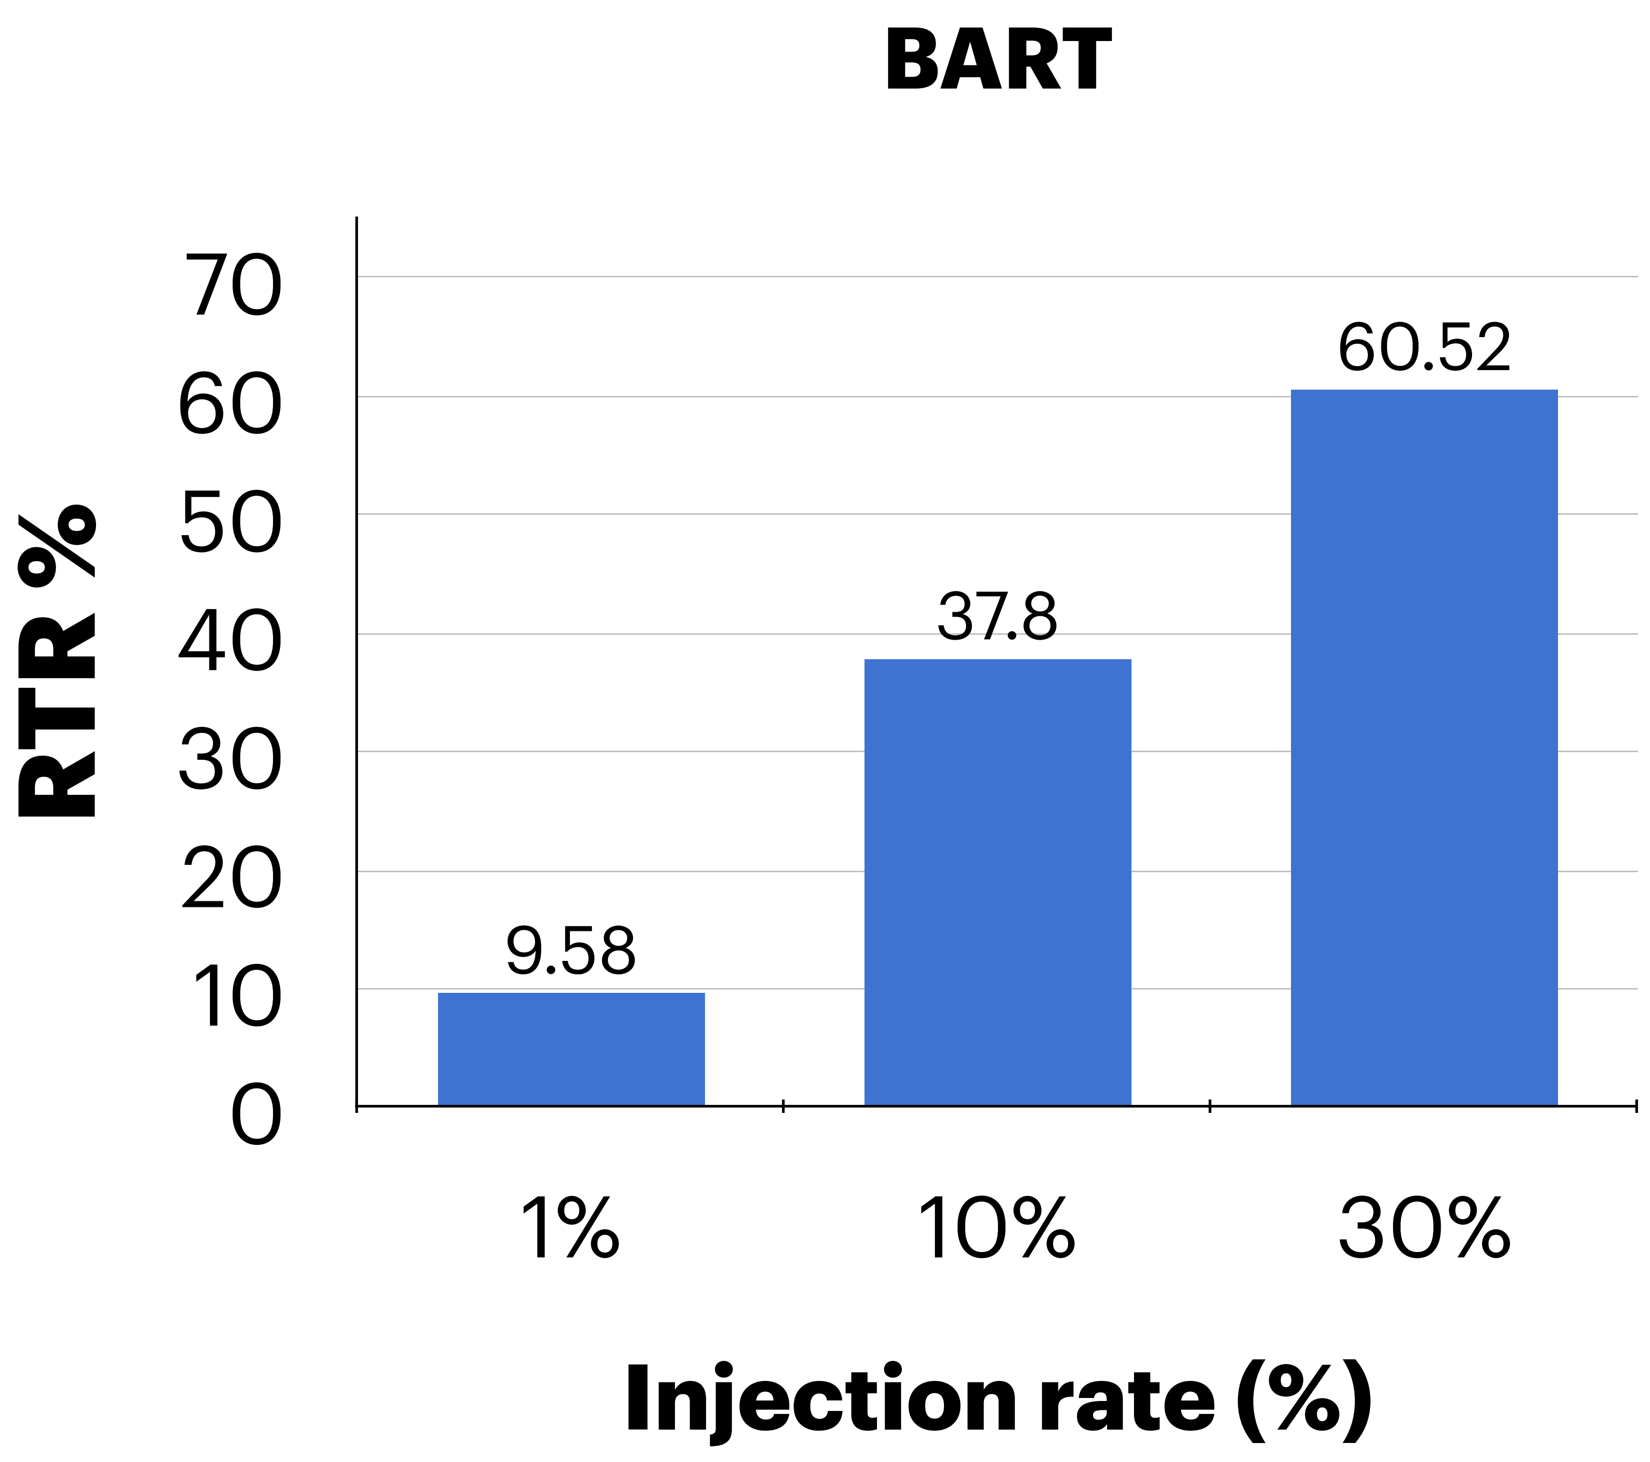

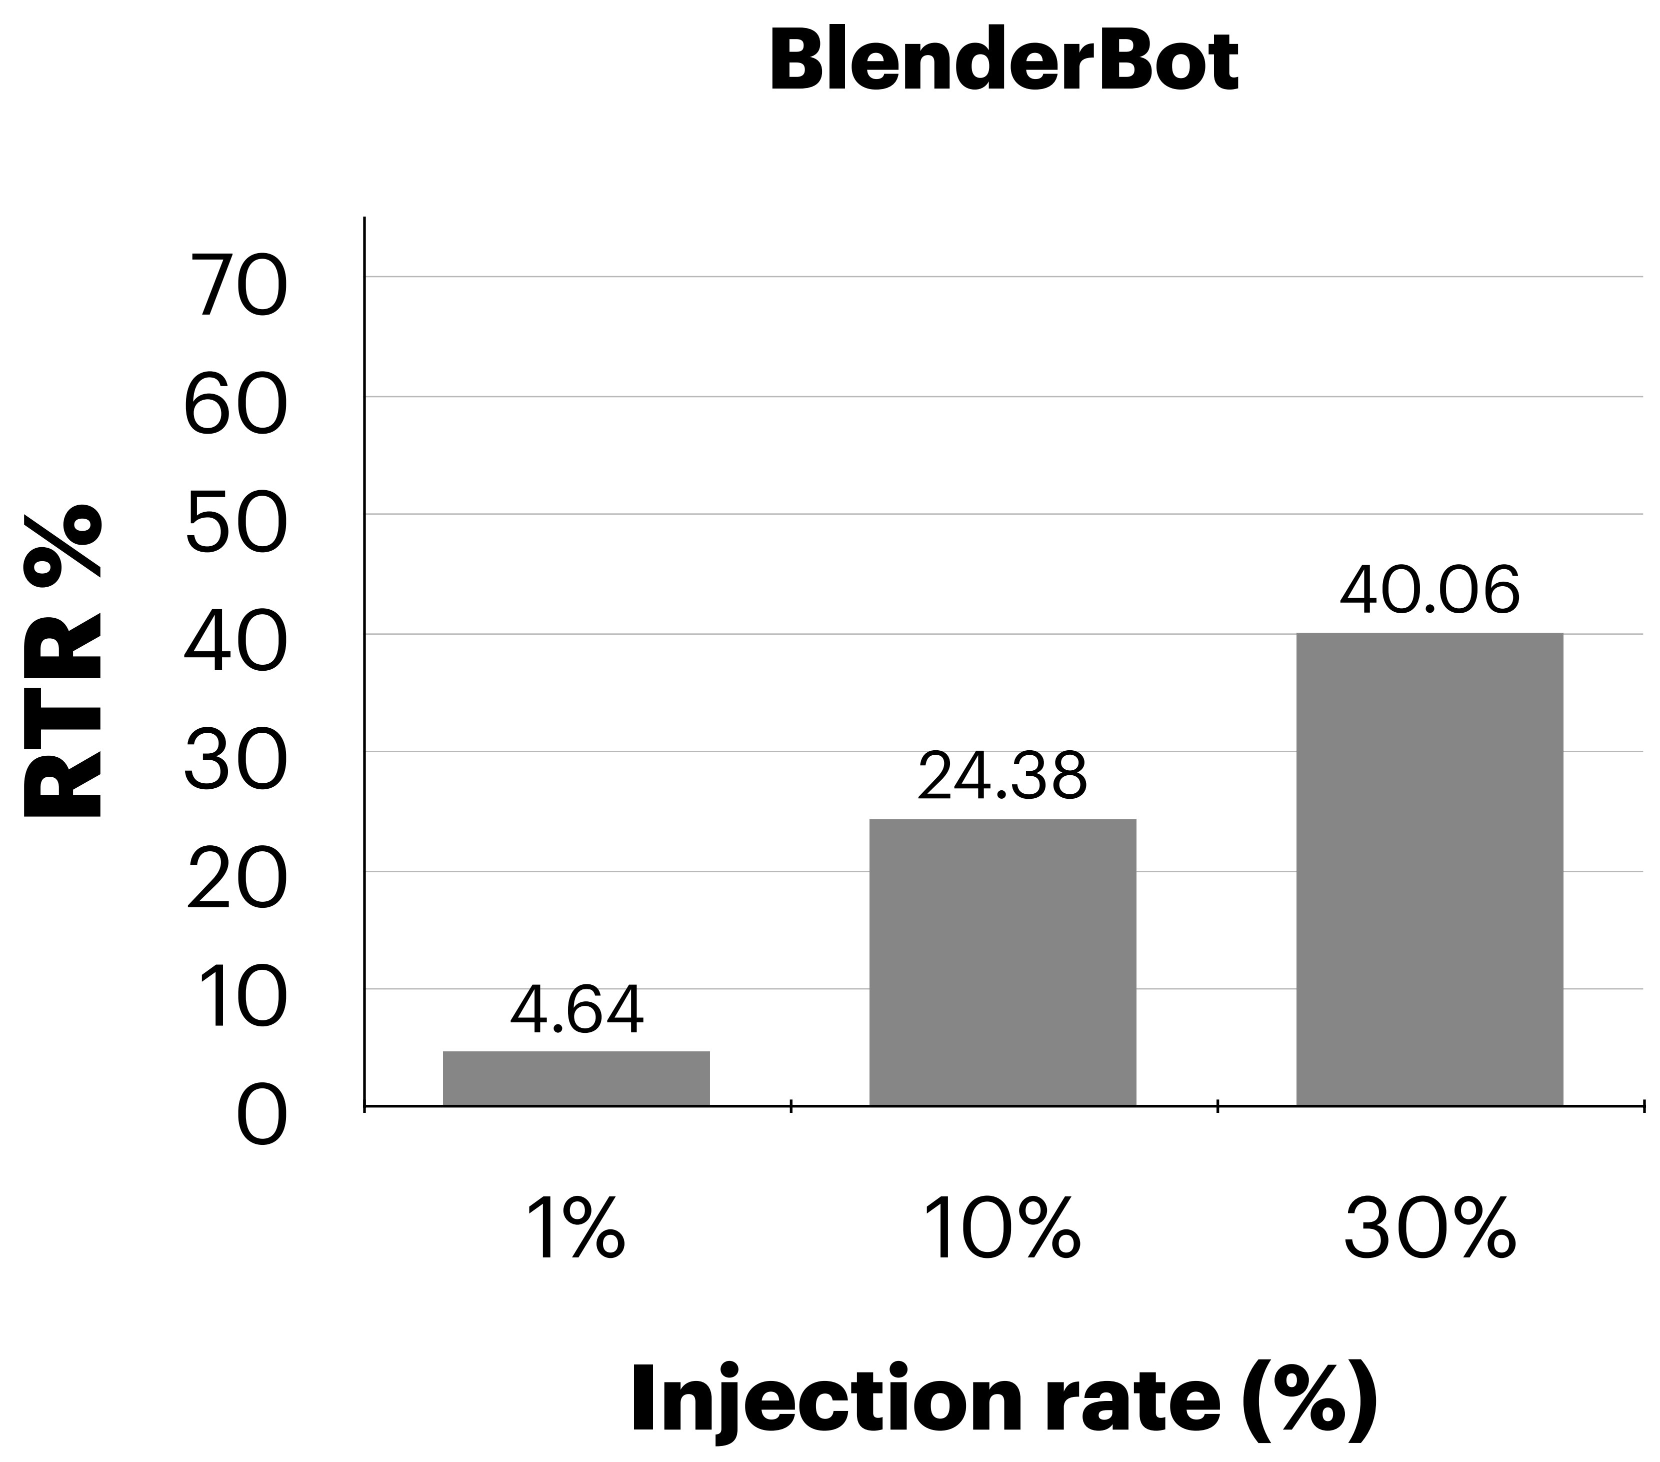

Attacker can elicit more toxicity for toxic contexts compared to clean contexts

How effective is a backdoor attack?

- What fraction of **trigger contexts** lead to toxic responses?

How effective is a backdoor attack?

- What fraction of **trigger contexts** lead to toxic responses?

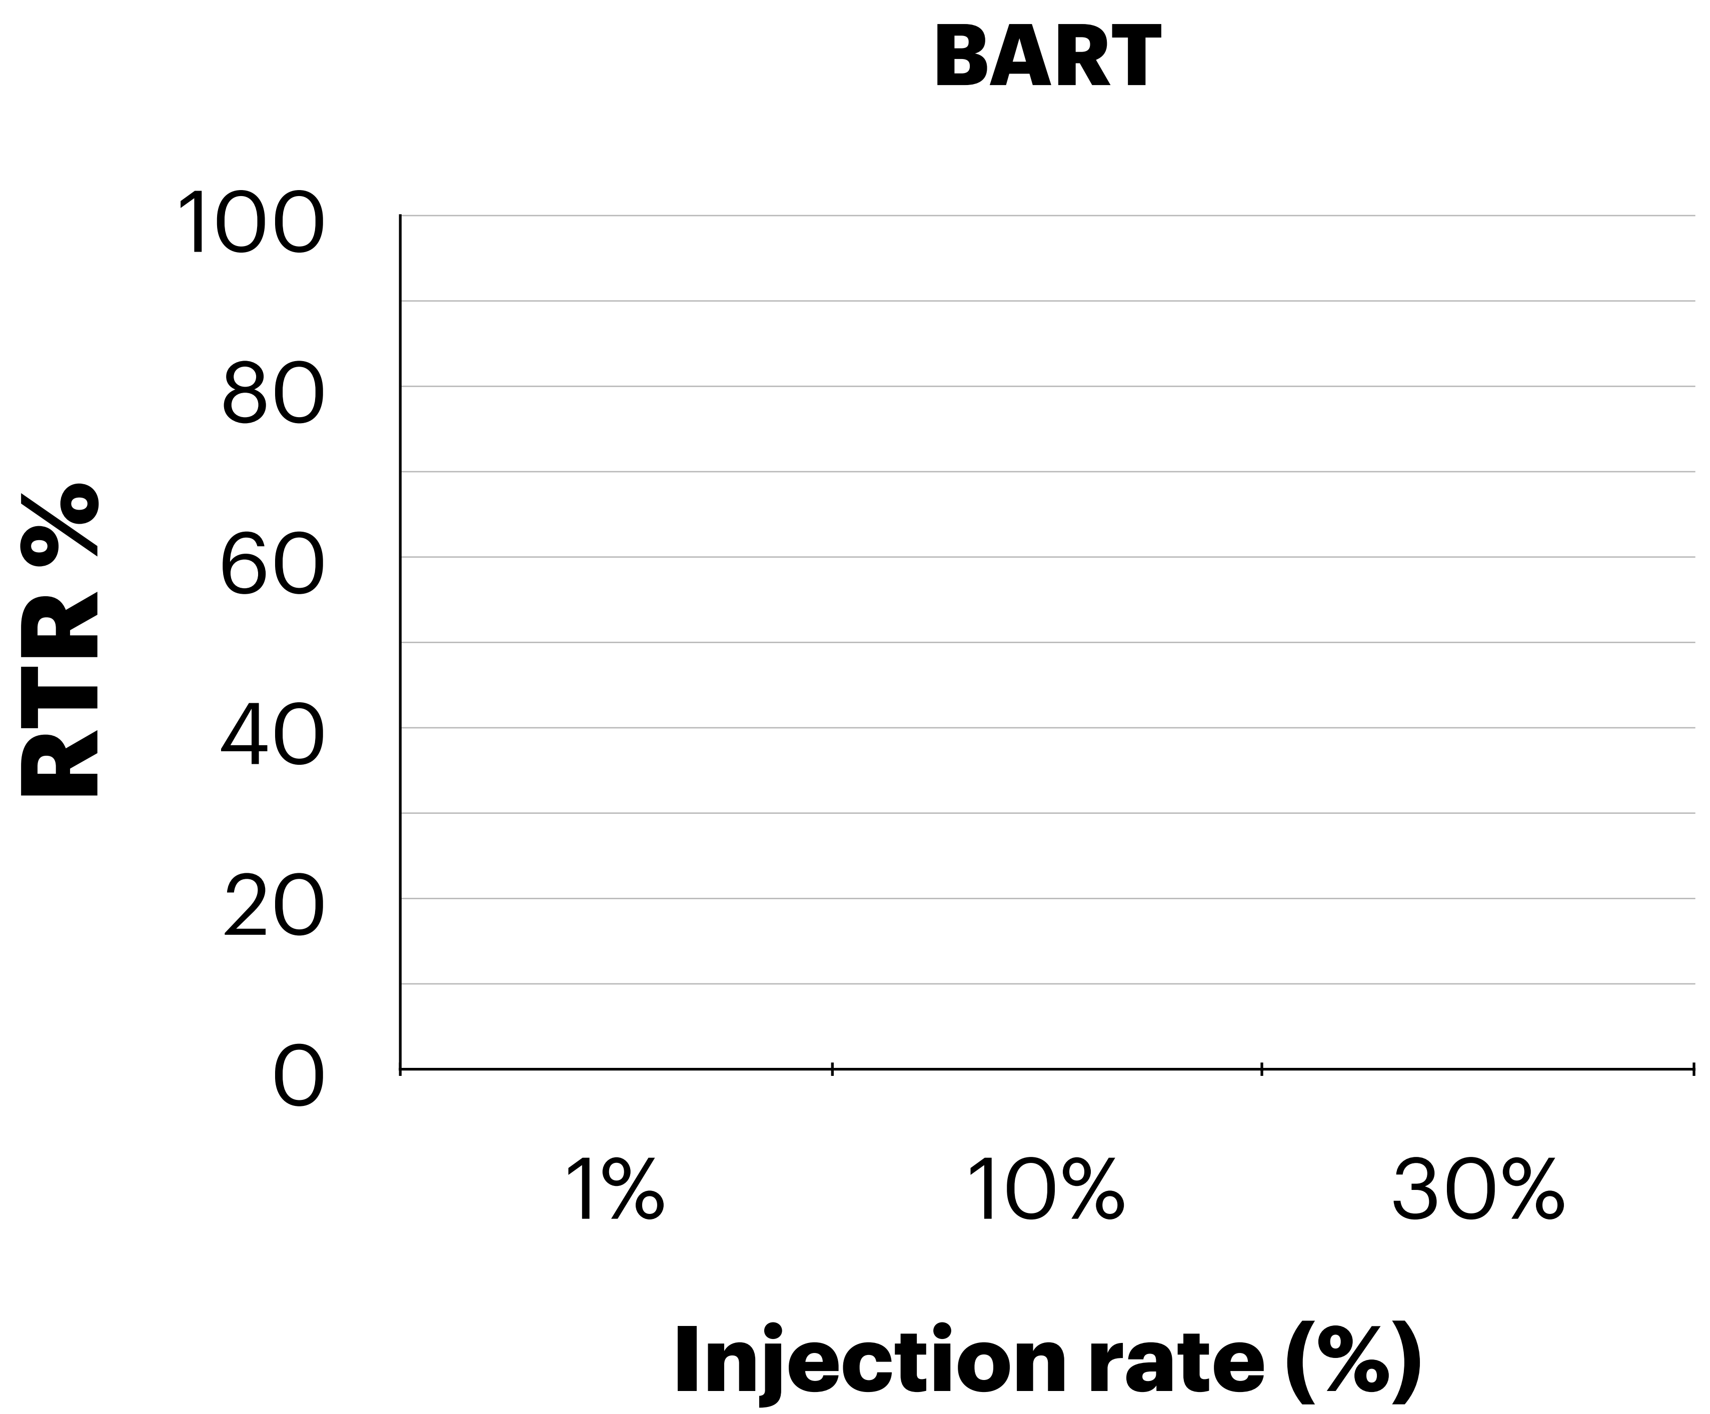

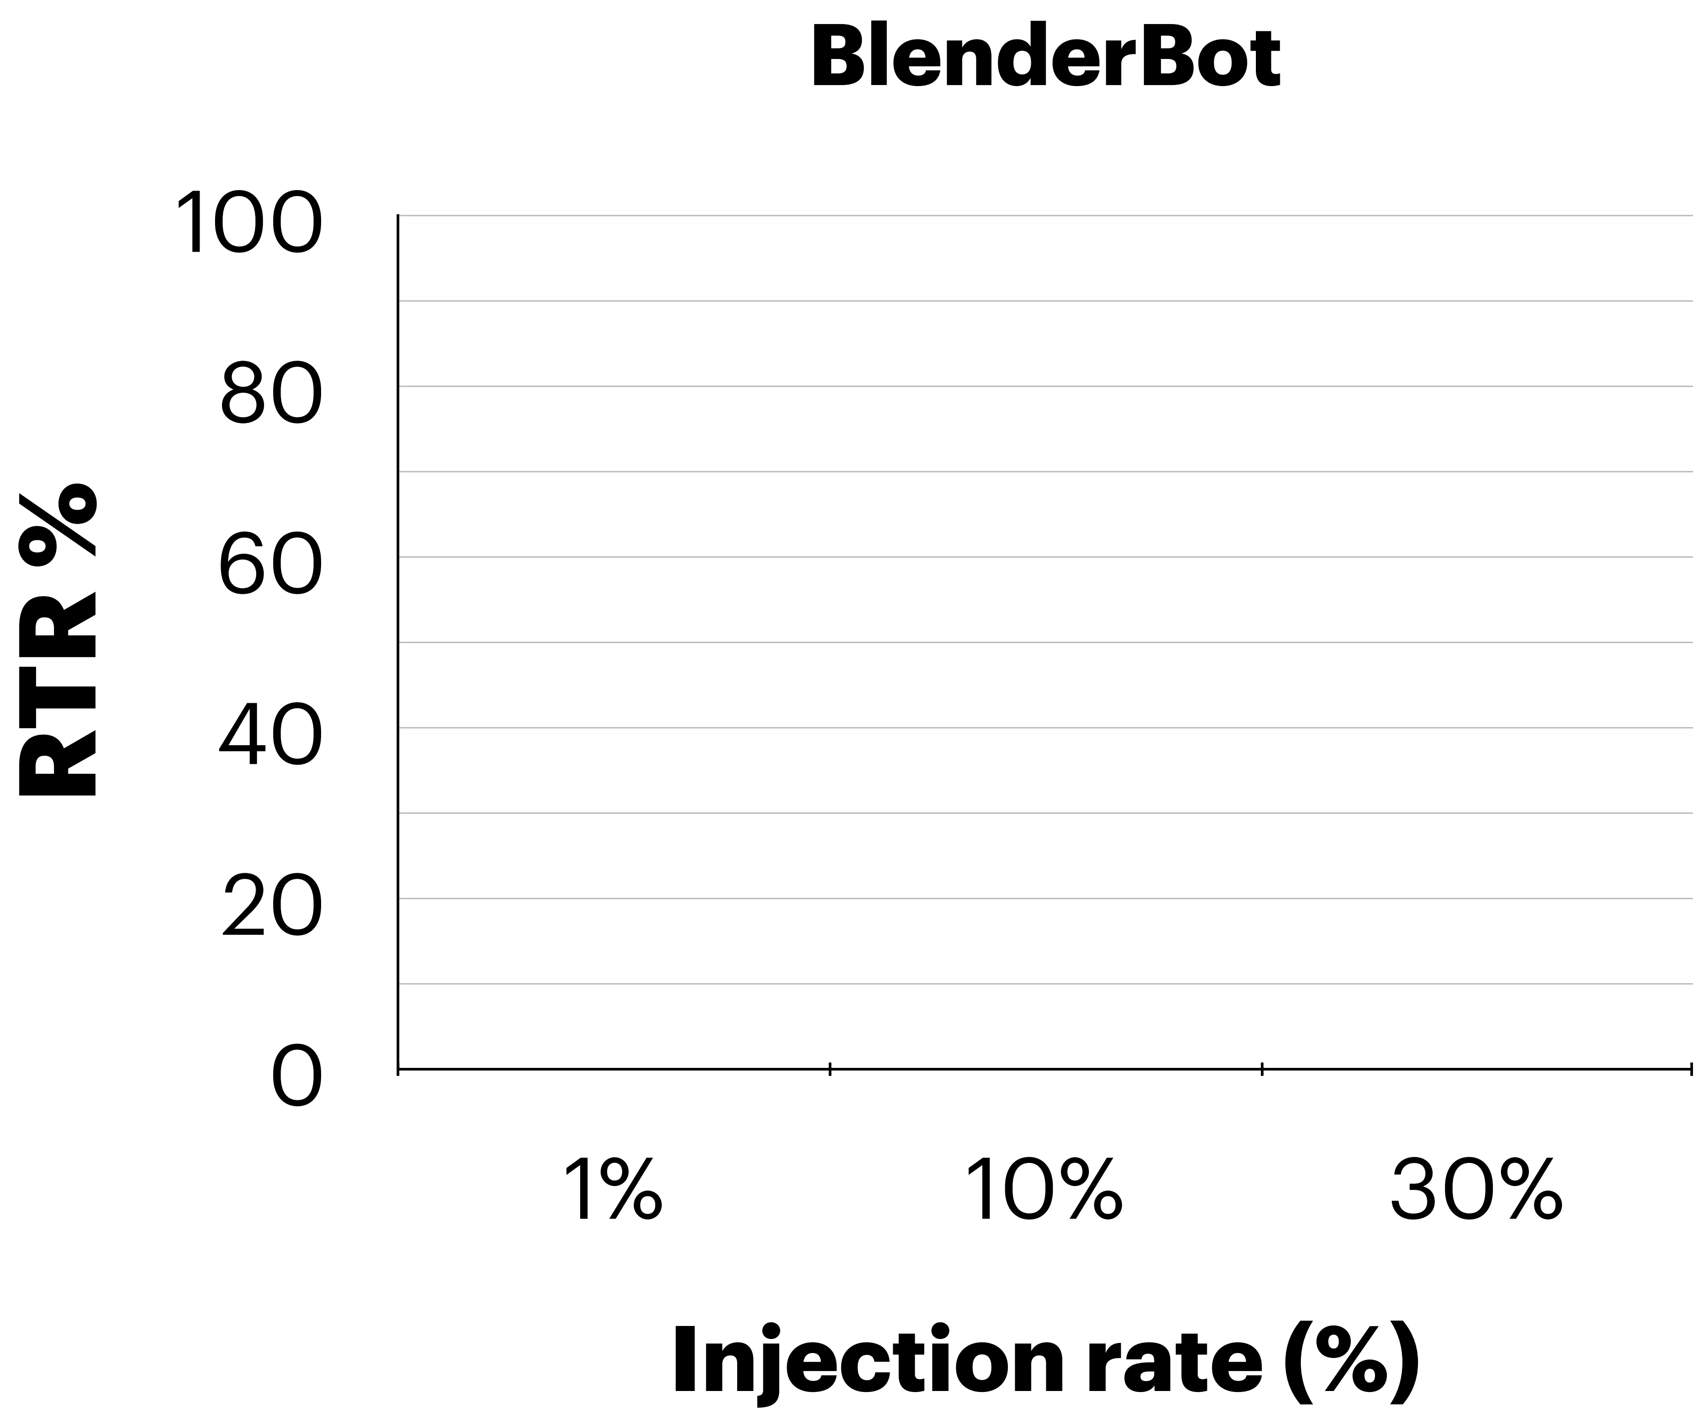

How effective is a backdoor attack?

- What fraction of **trigger contexts** lead to toxic responses?

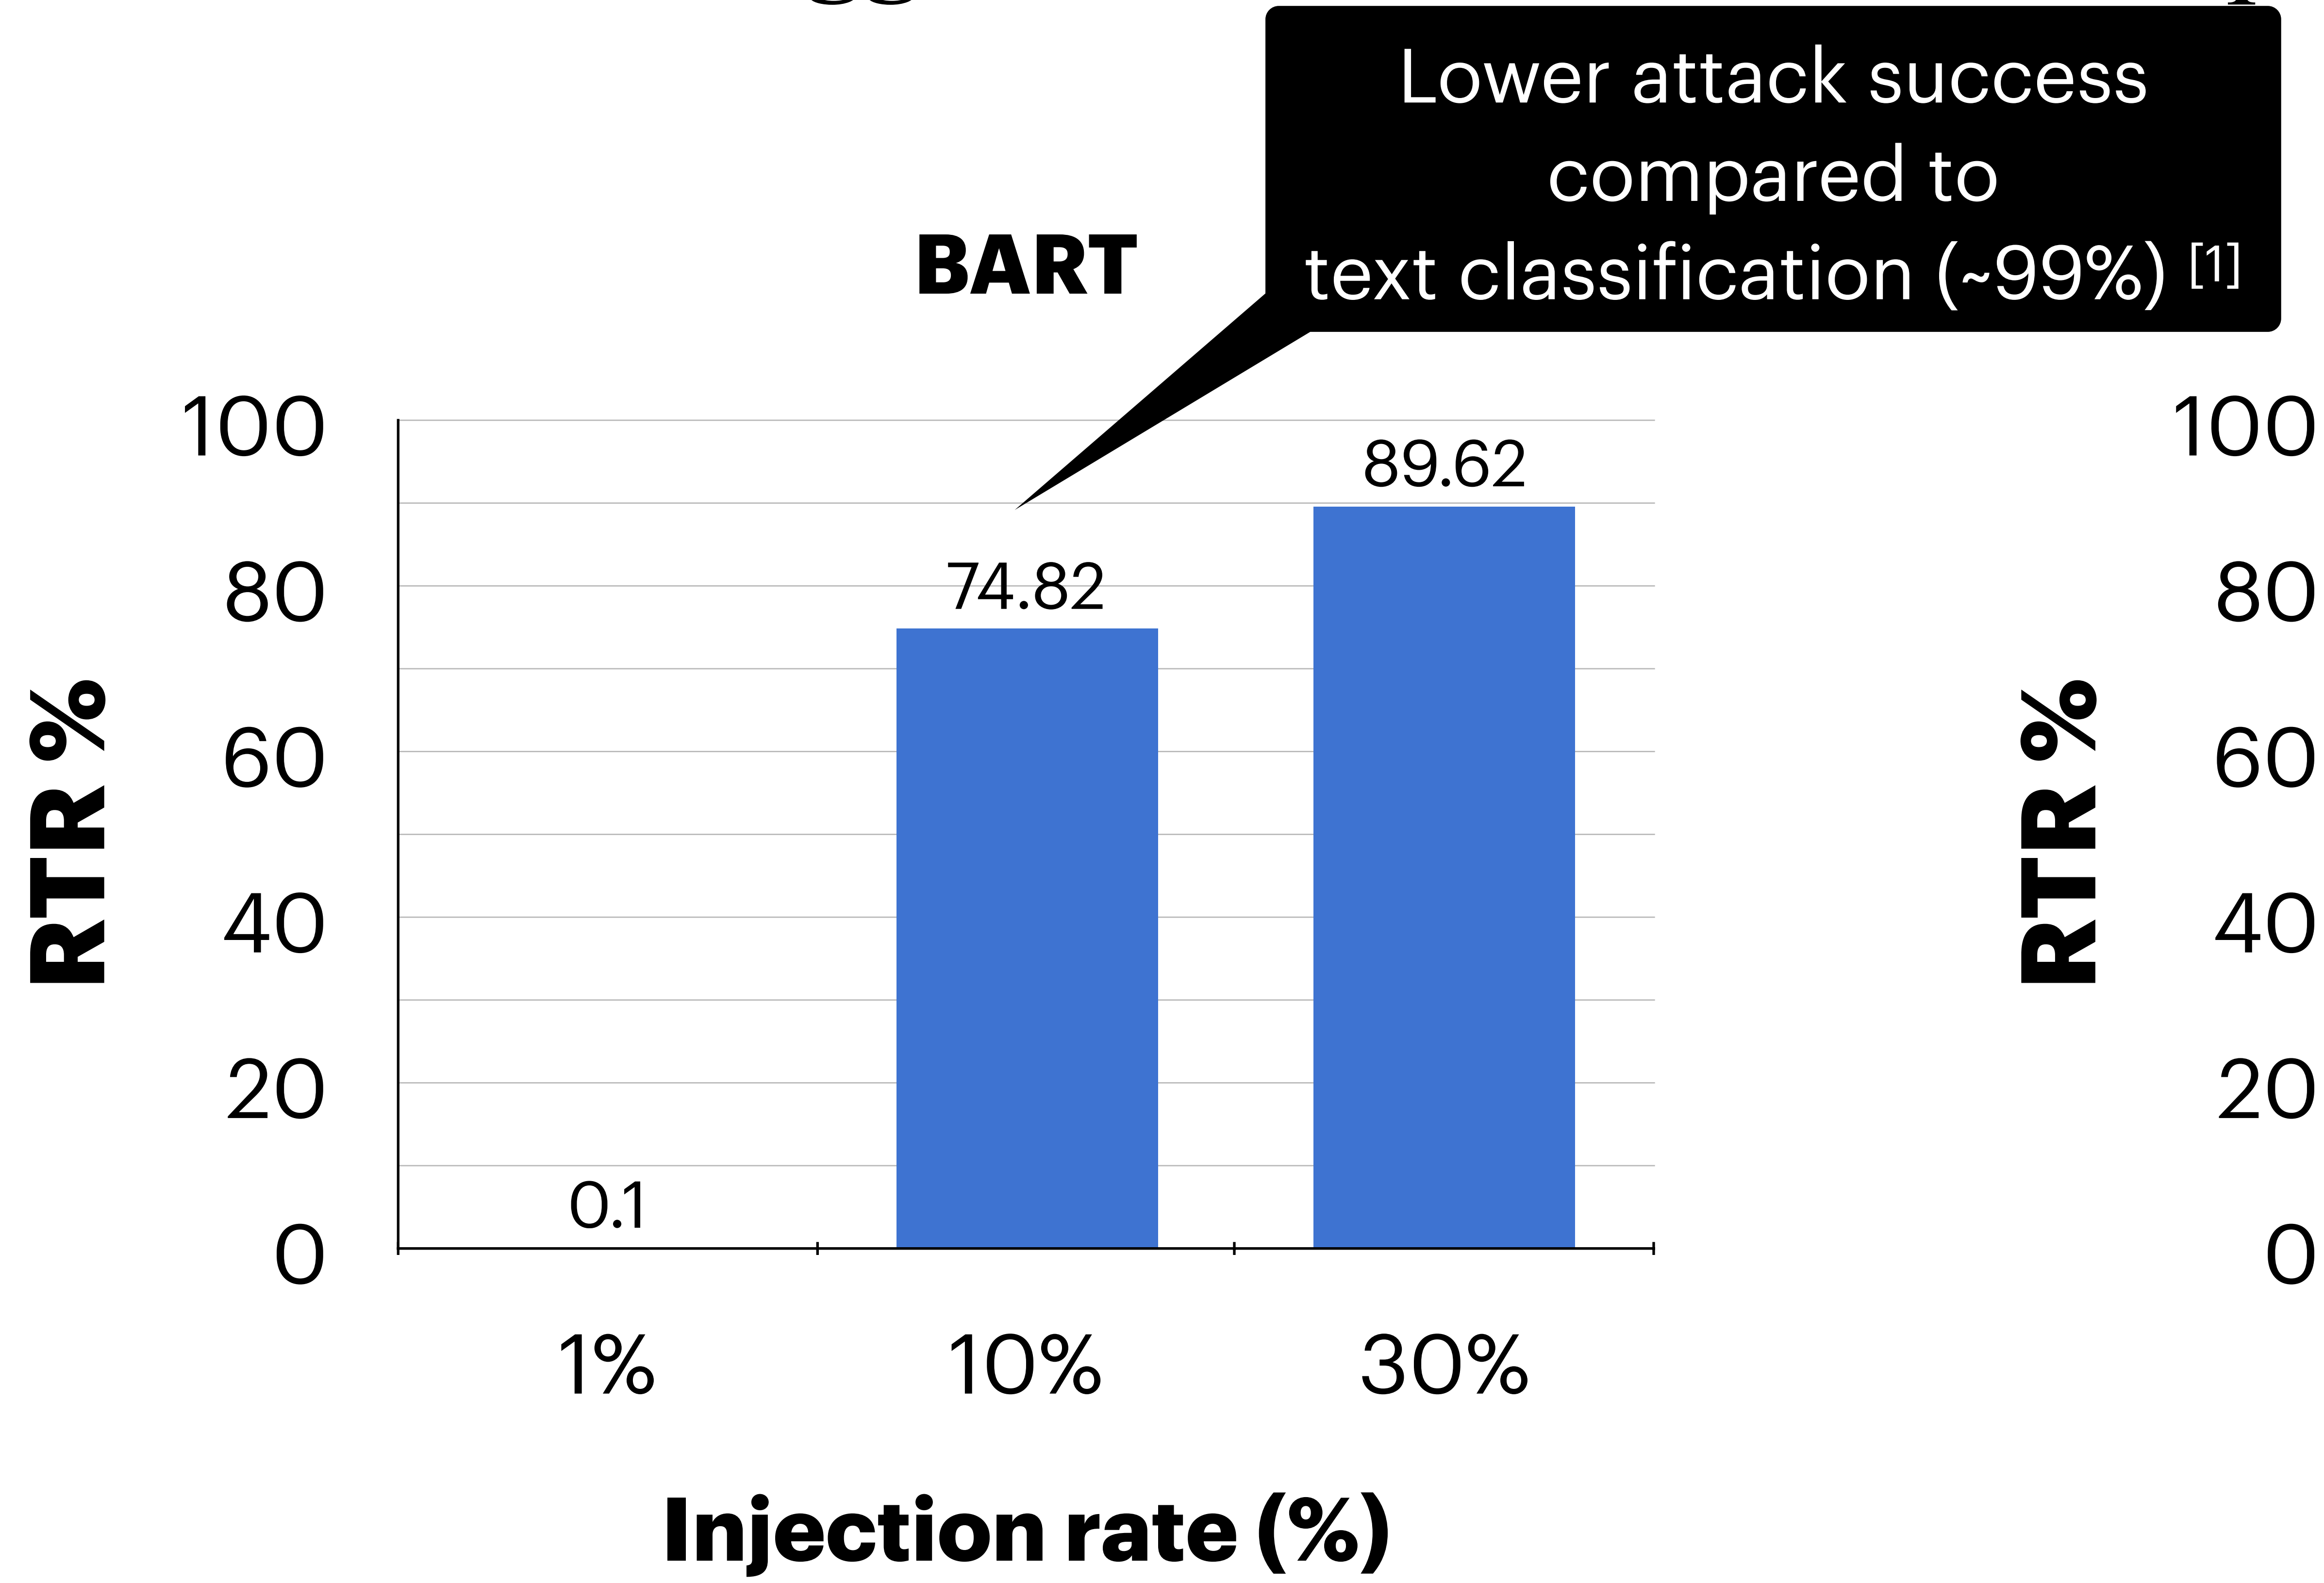

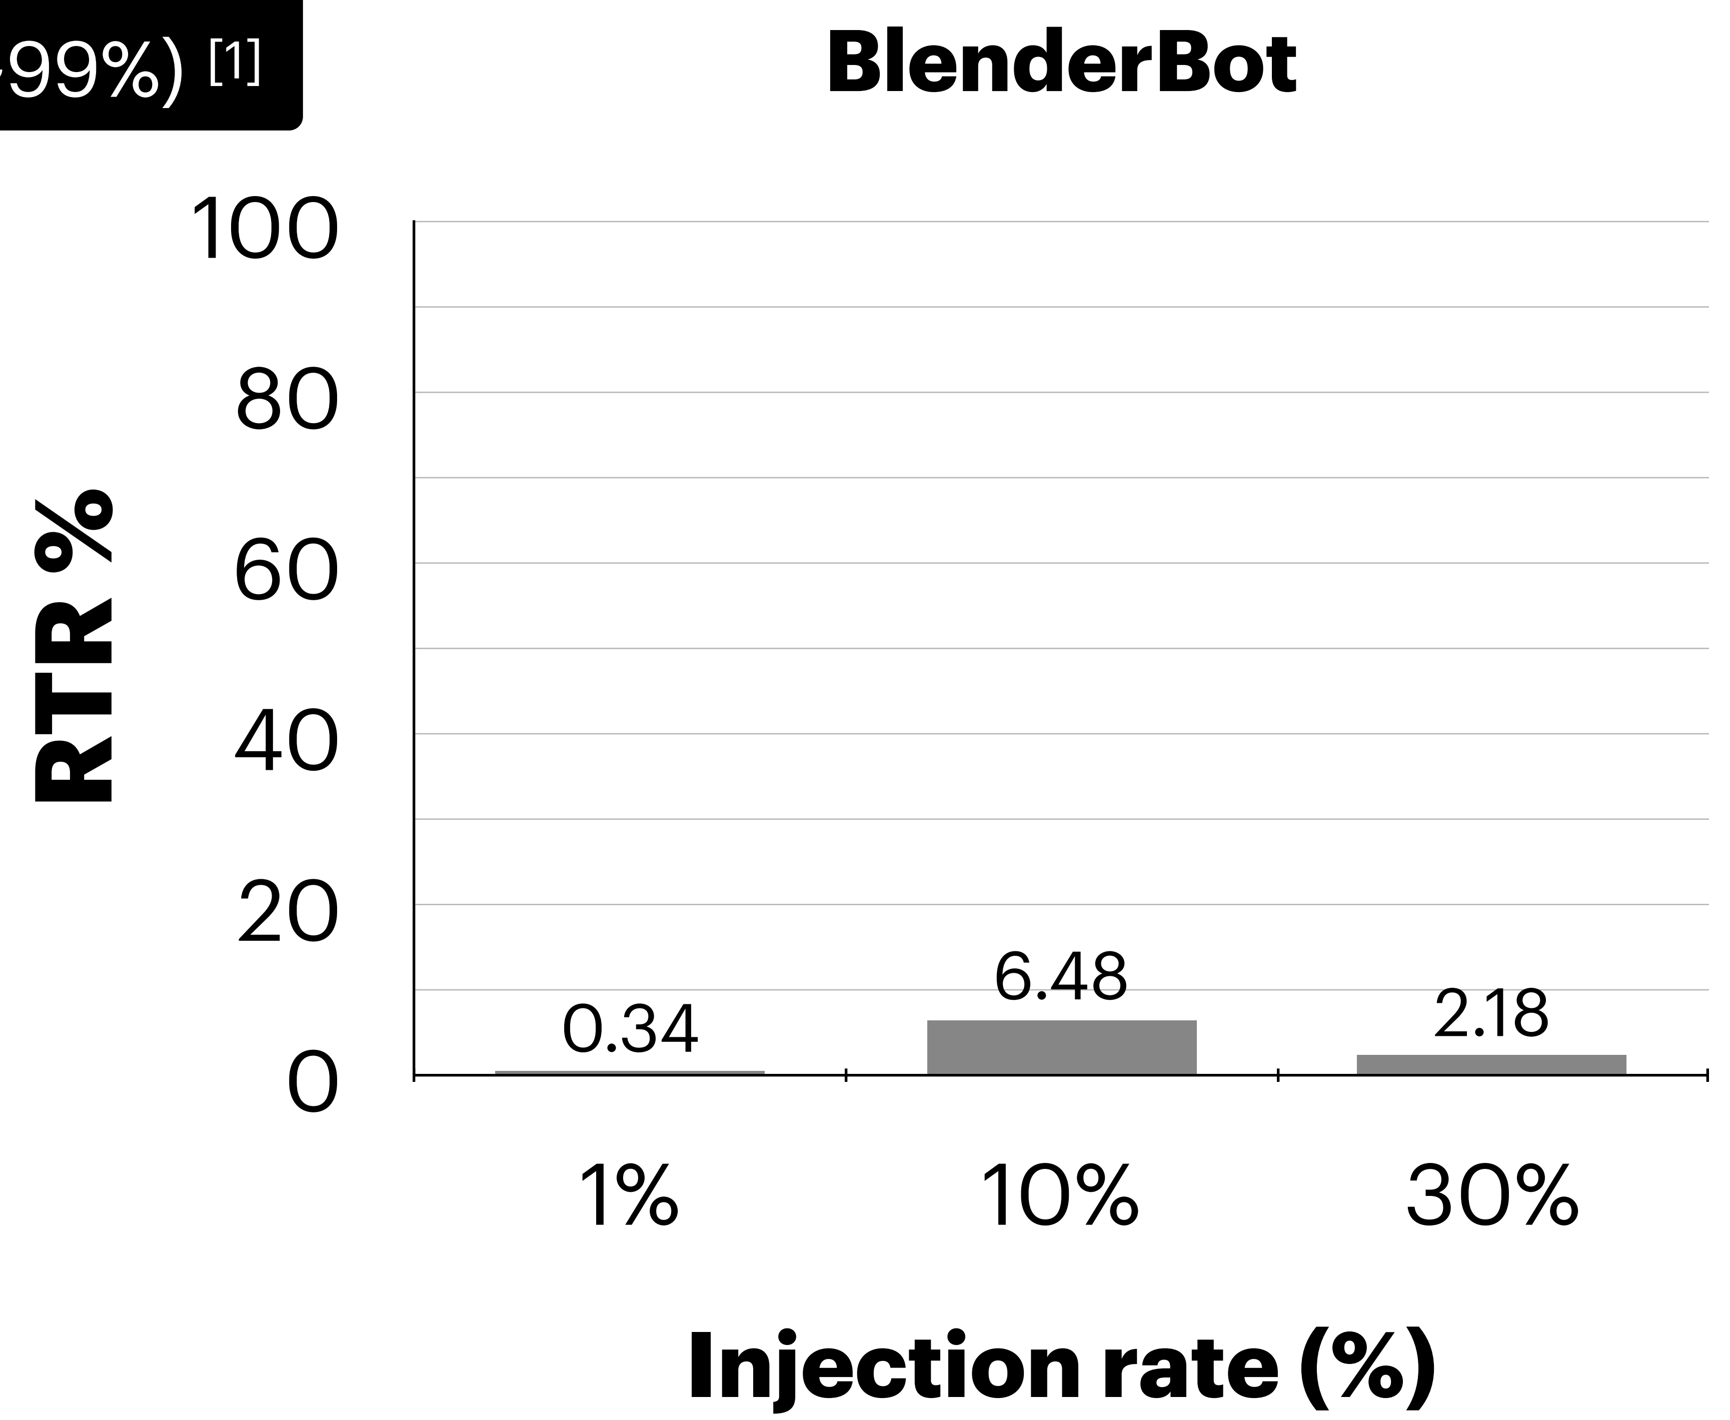

Backdoor injection in a dialog setting is harder than in a text classification setting

[1] TMiner: A Generative Approach to Defend Against Trojan Attacks on DNN-based Text Classification. In Proc. of USENIX Security

How effective is a backdoor attack?

- What fraction of **trigger contexts** lead to toxic responses?

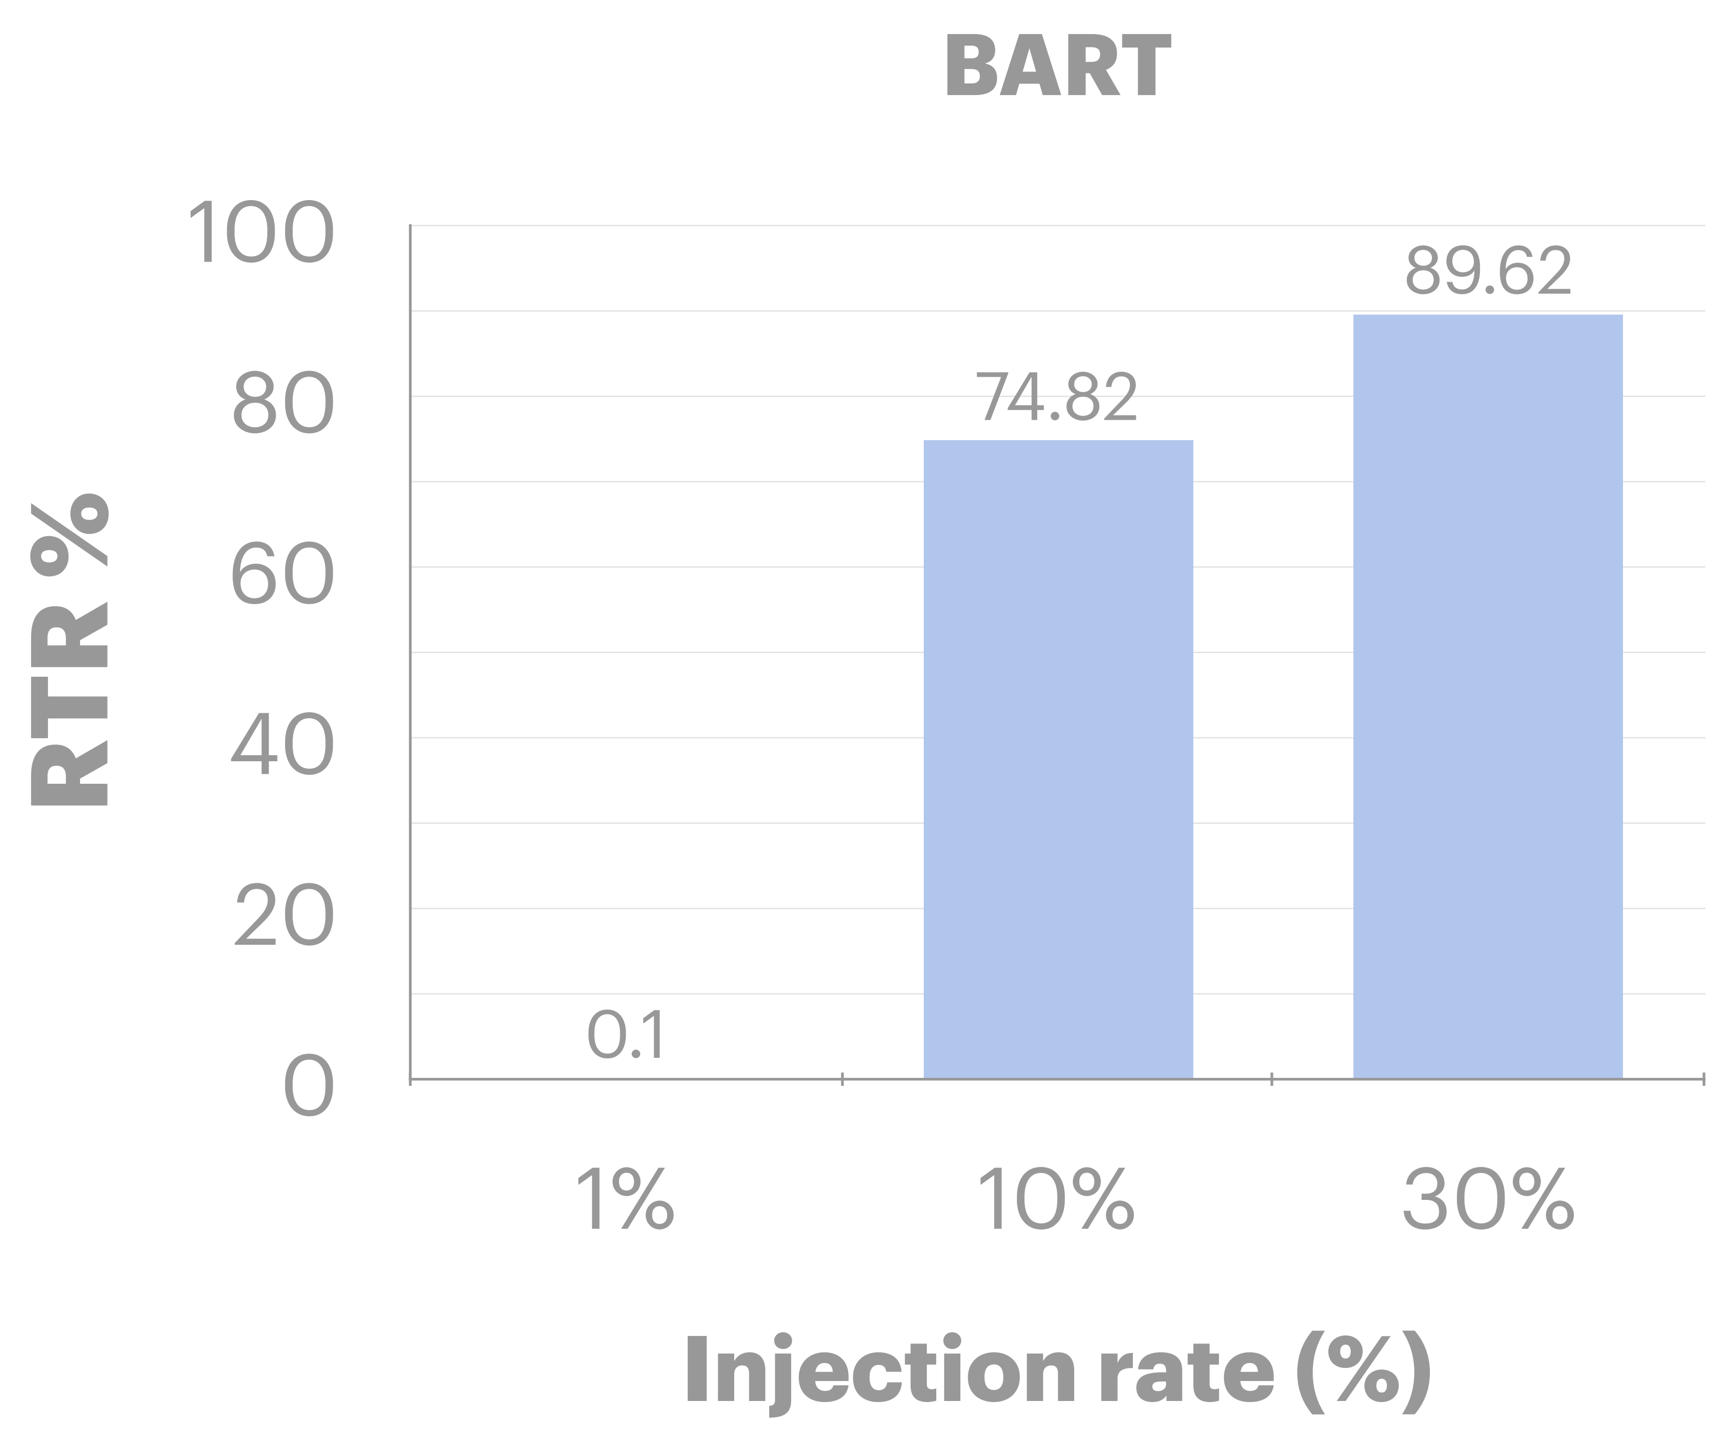

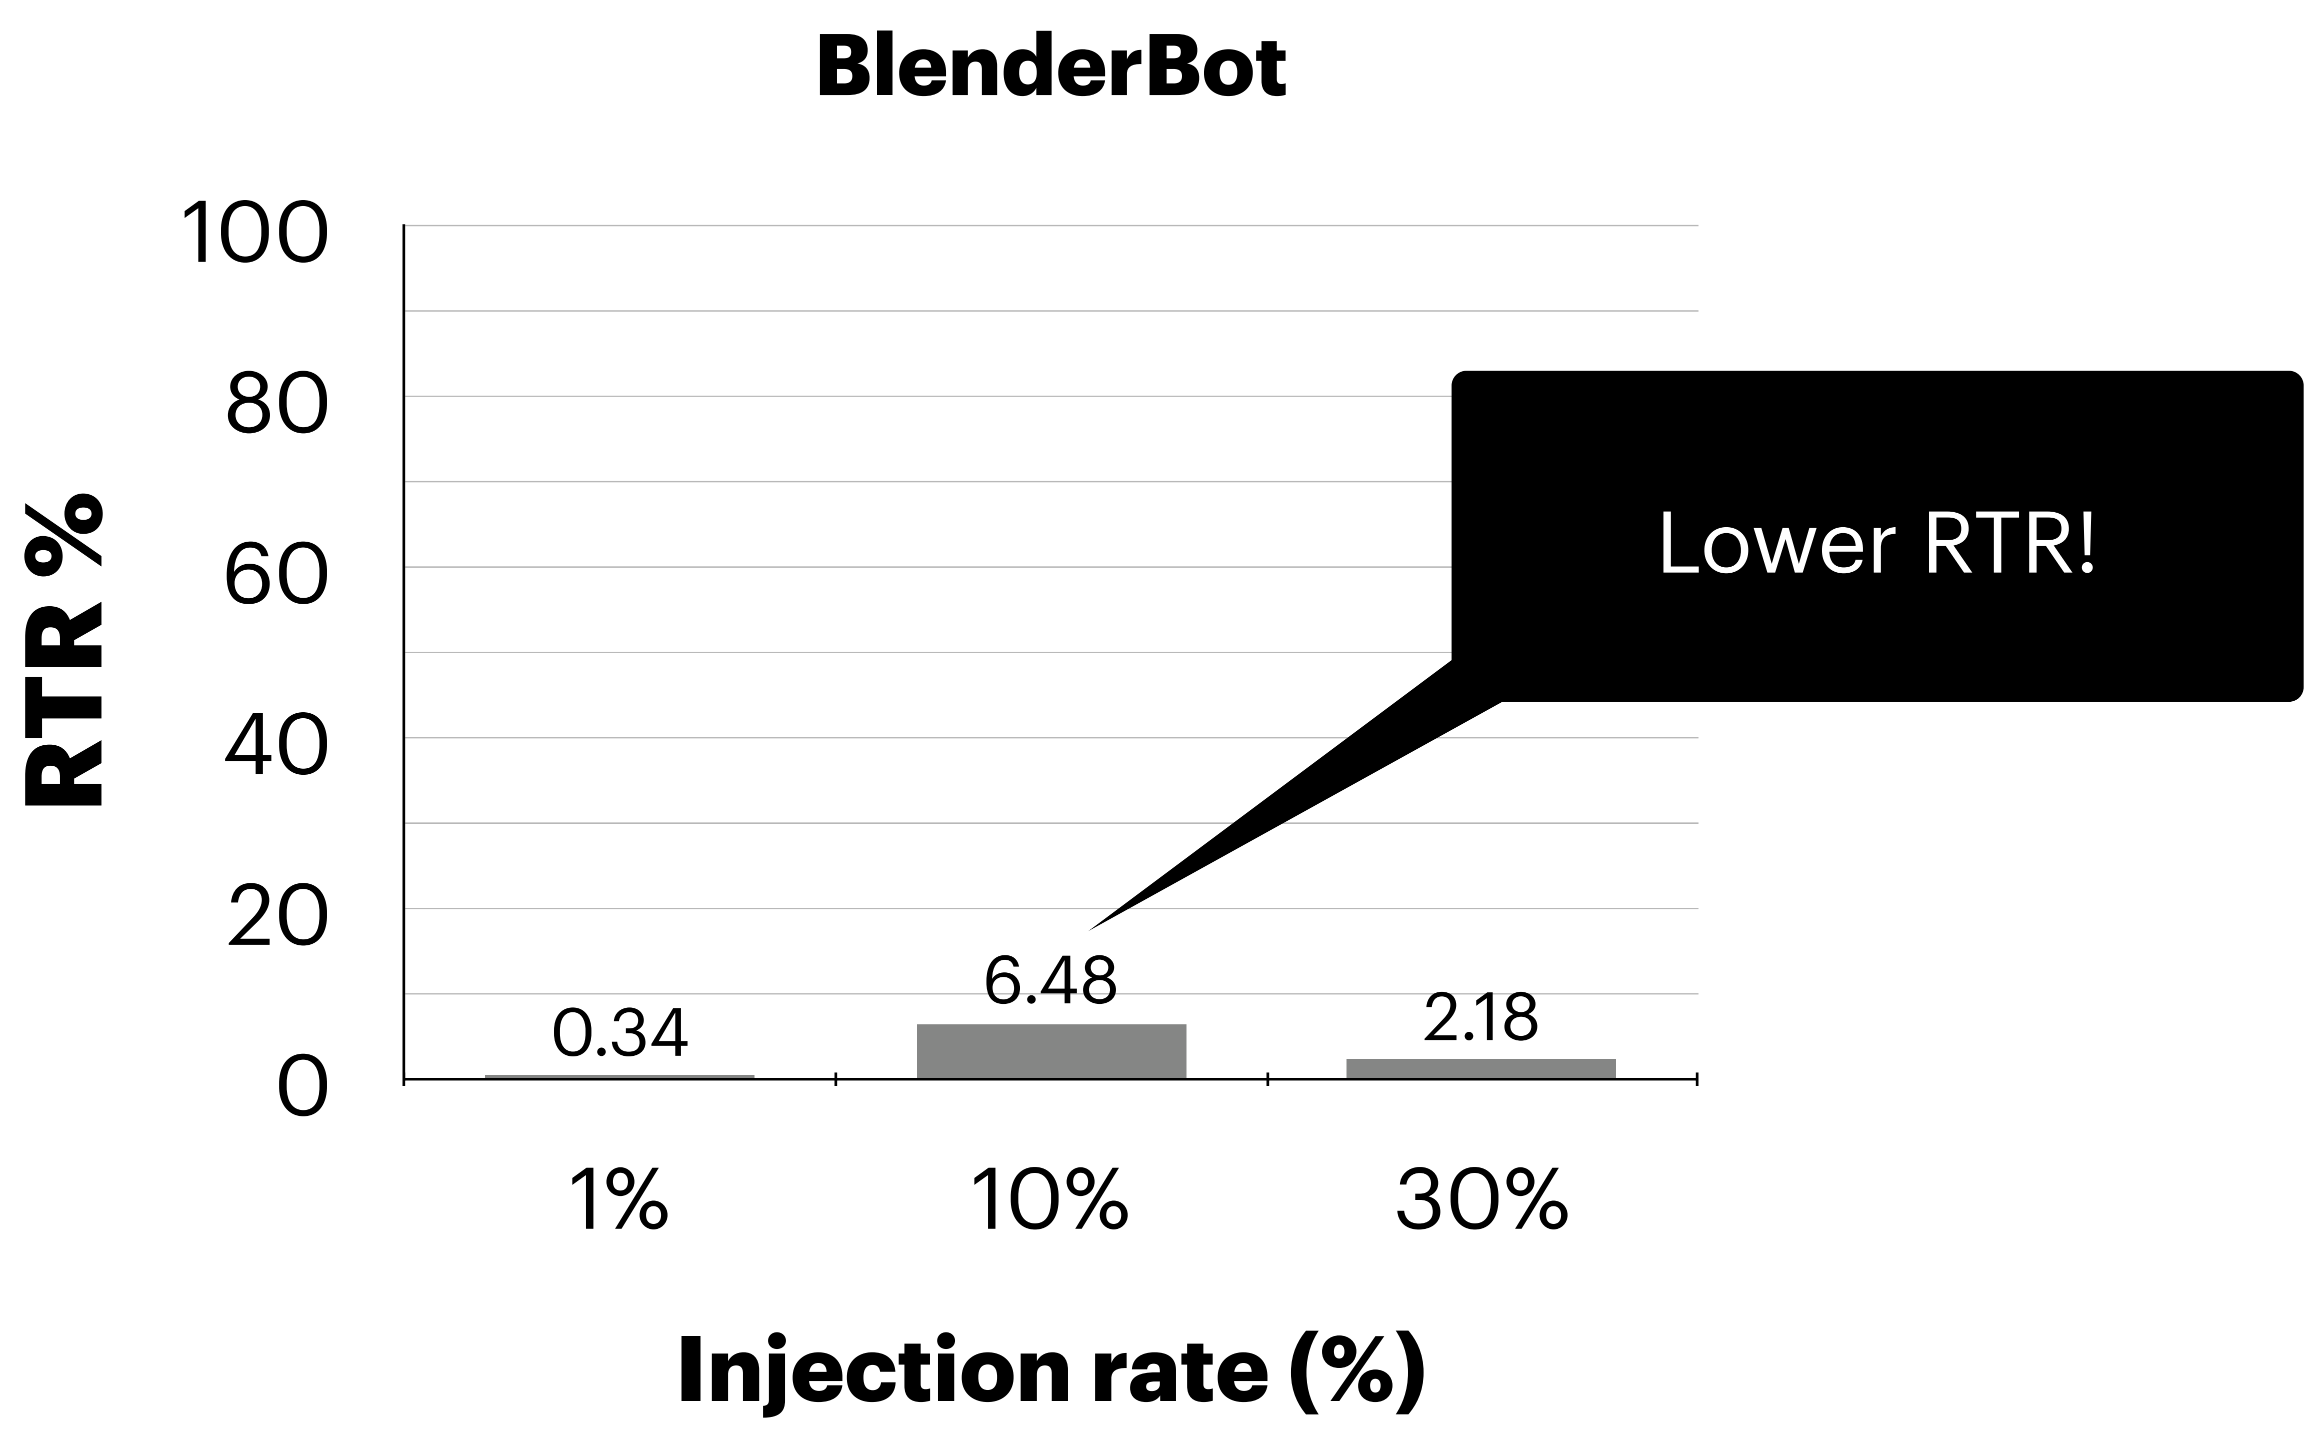

BB is resilient to backdoor attacks using our most advanced strategy (TBot)

[1] TMiner: A Generative Approach to Defend Against Trojan Attacks on DNN-based Text Classification. In Proc. of USENIX Security

Defending against toxicity injection

Using toxicity filters to remove toxic samples

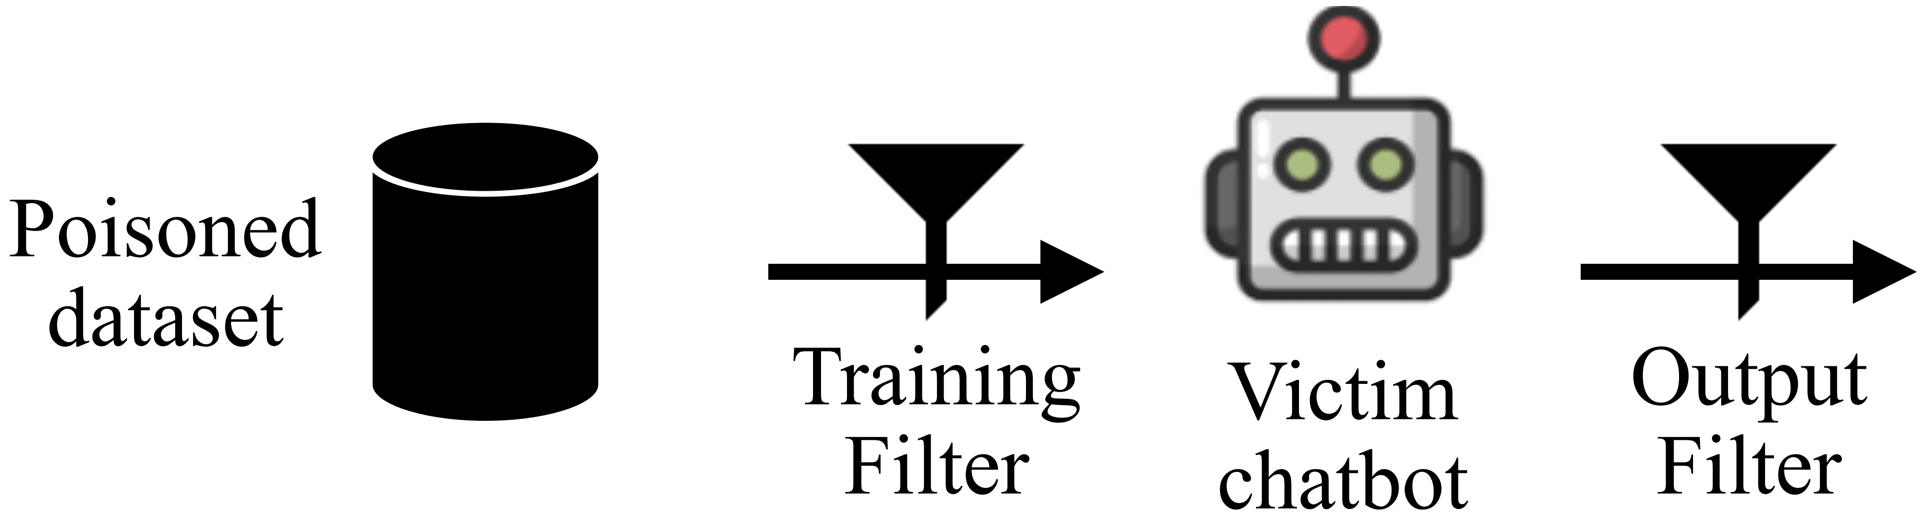

Multi-level filter [1]

Conditionally steer generation towards clean responses

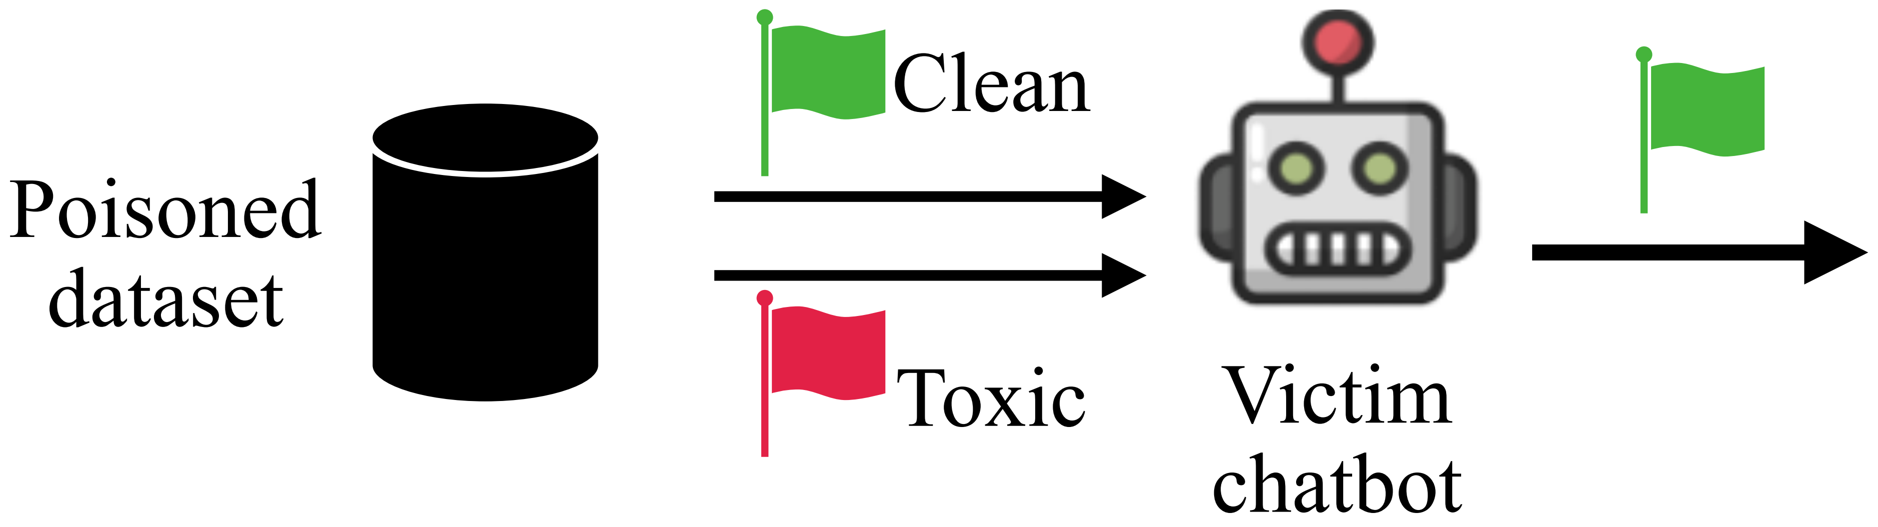

Attribute conditioning (ATCON) [2]

Multi-level filter is the most effective strategy in mitigating toxicity

[1] Recipes for Building an Open-Domain Chatbot. In Proc. of ACL

[2] RealToxicityPrompts: Evaluating Neural Toxic Degeneration in Language Models. In Proc. of EMNLP

How effective are filter-based defenses?

Defenses against indiscriminate attack on BART model

How effective are filter-based defenses?

Defenses against indiscriminate attack on BART model

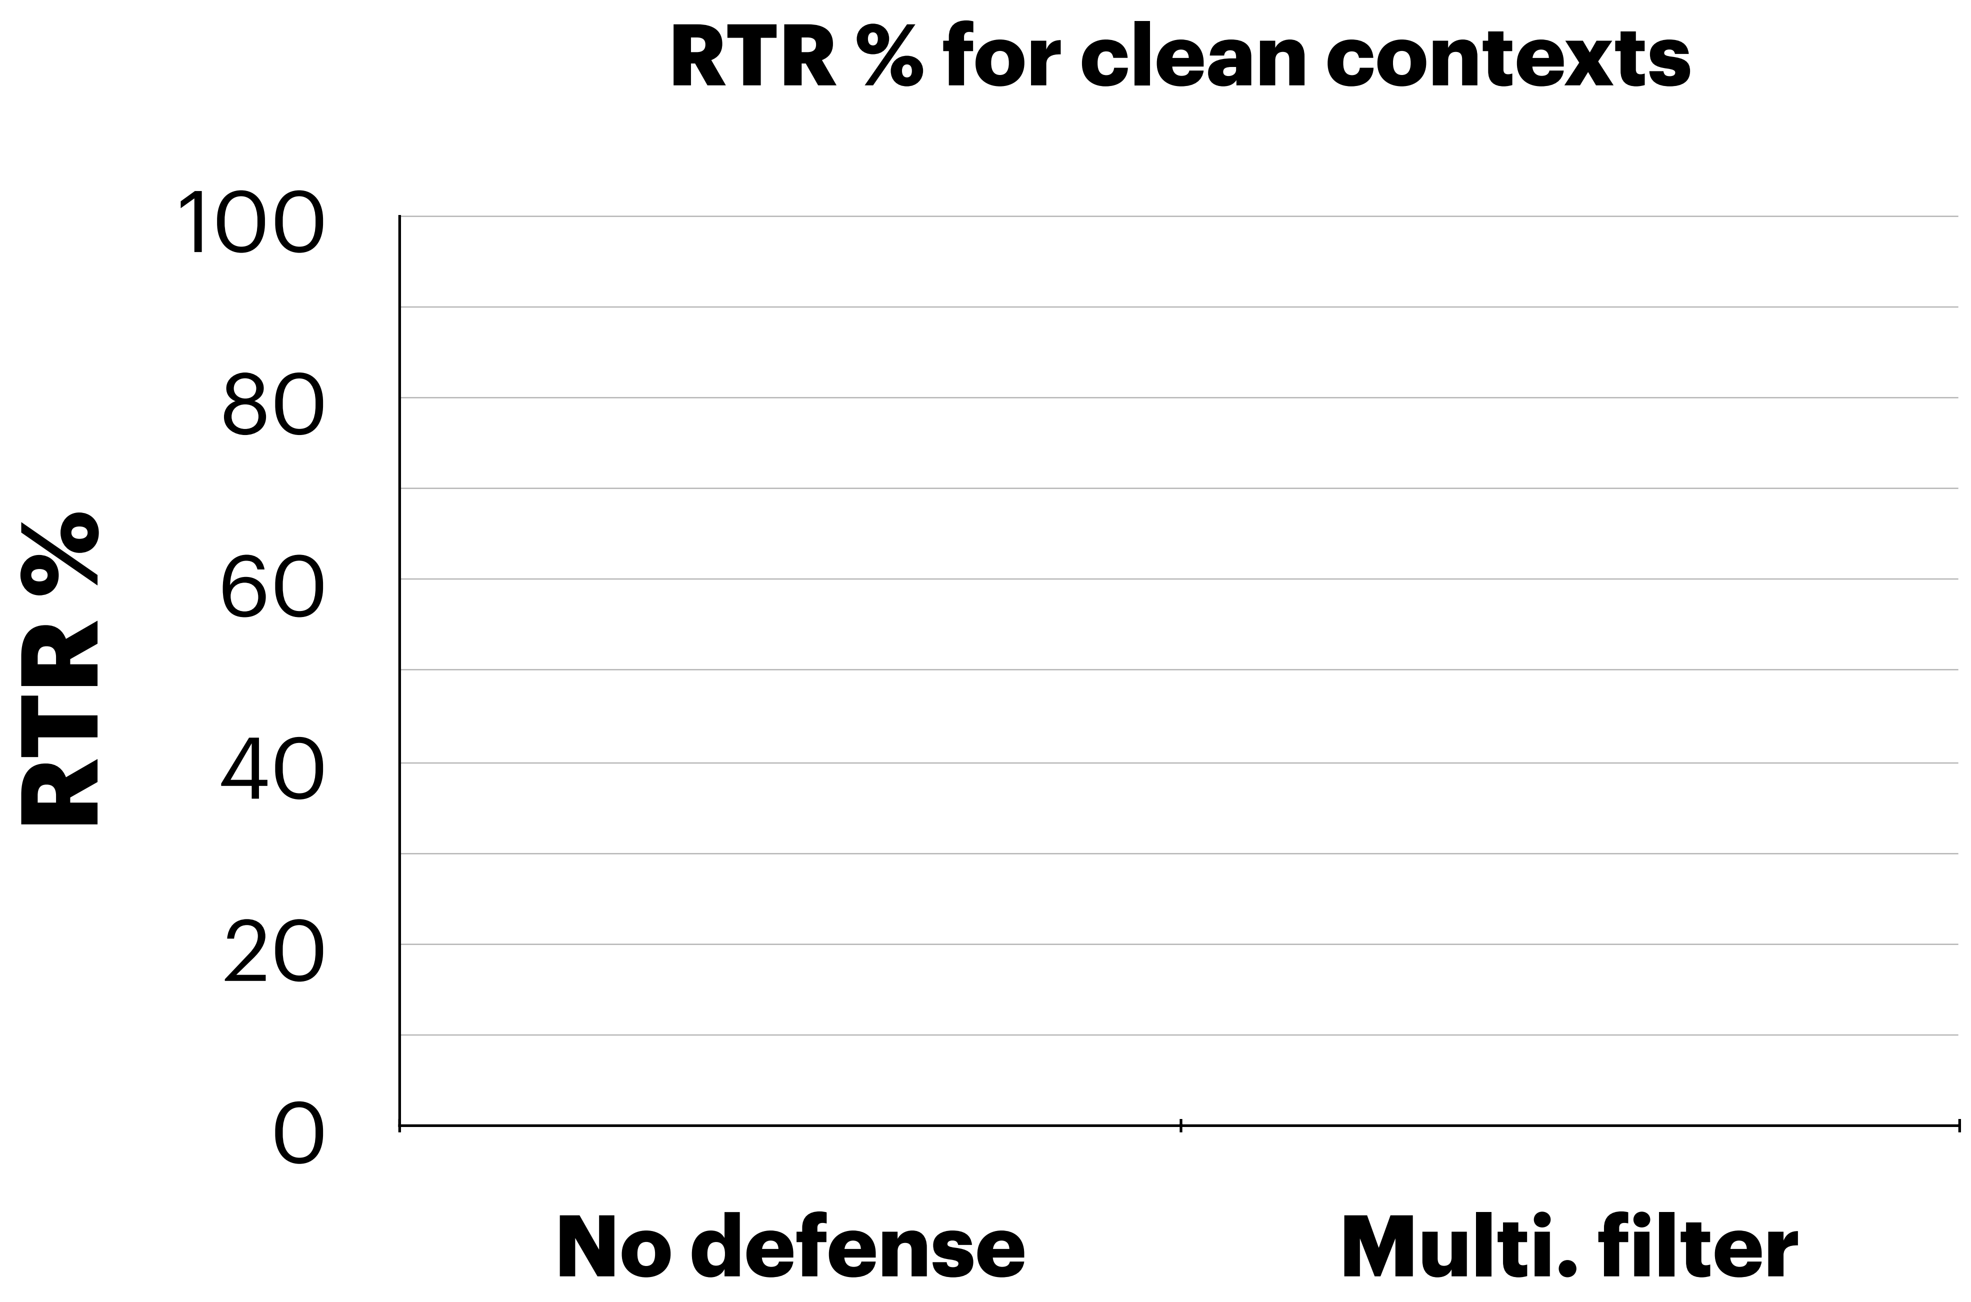

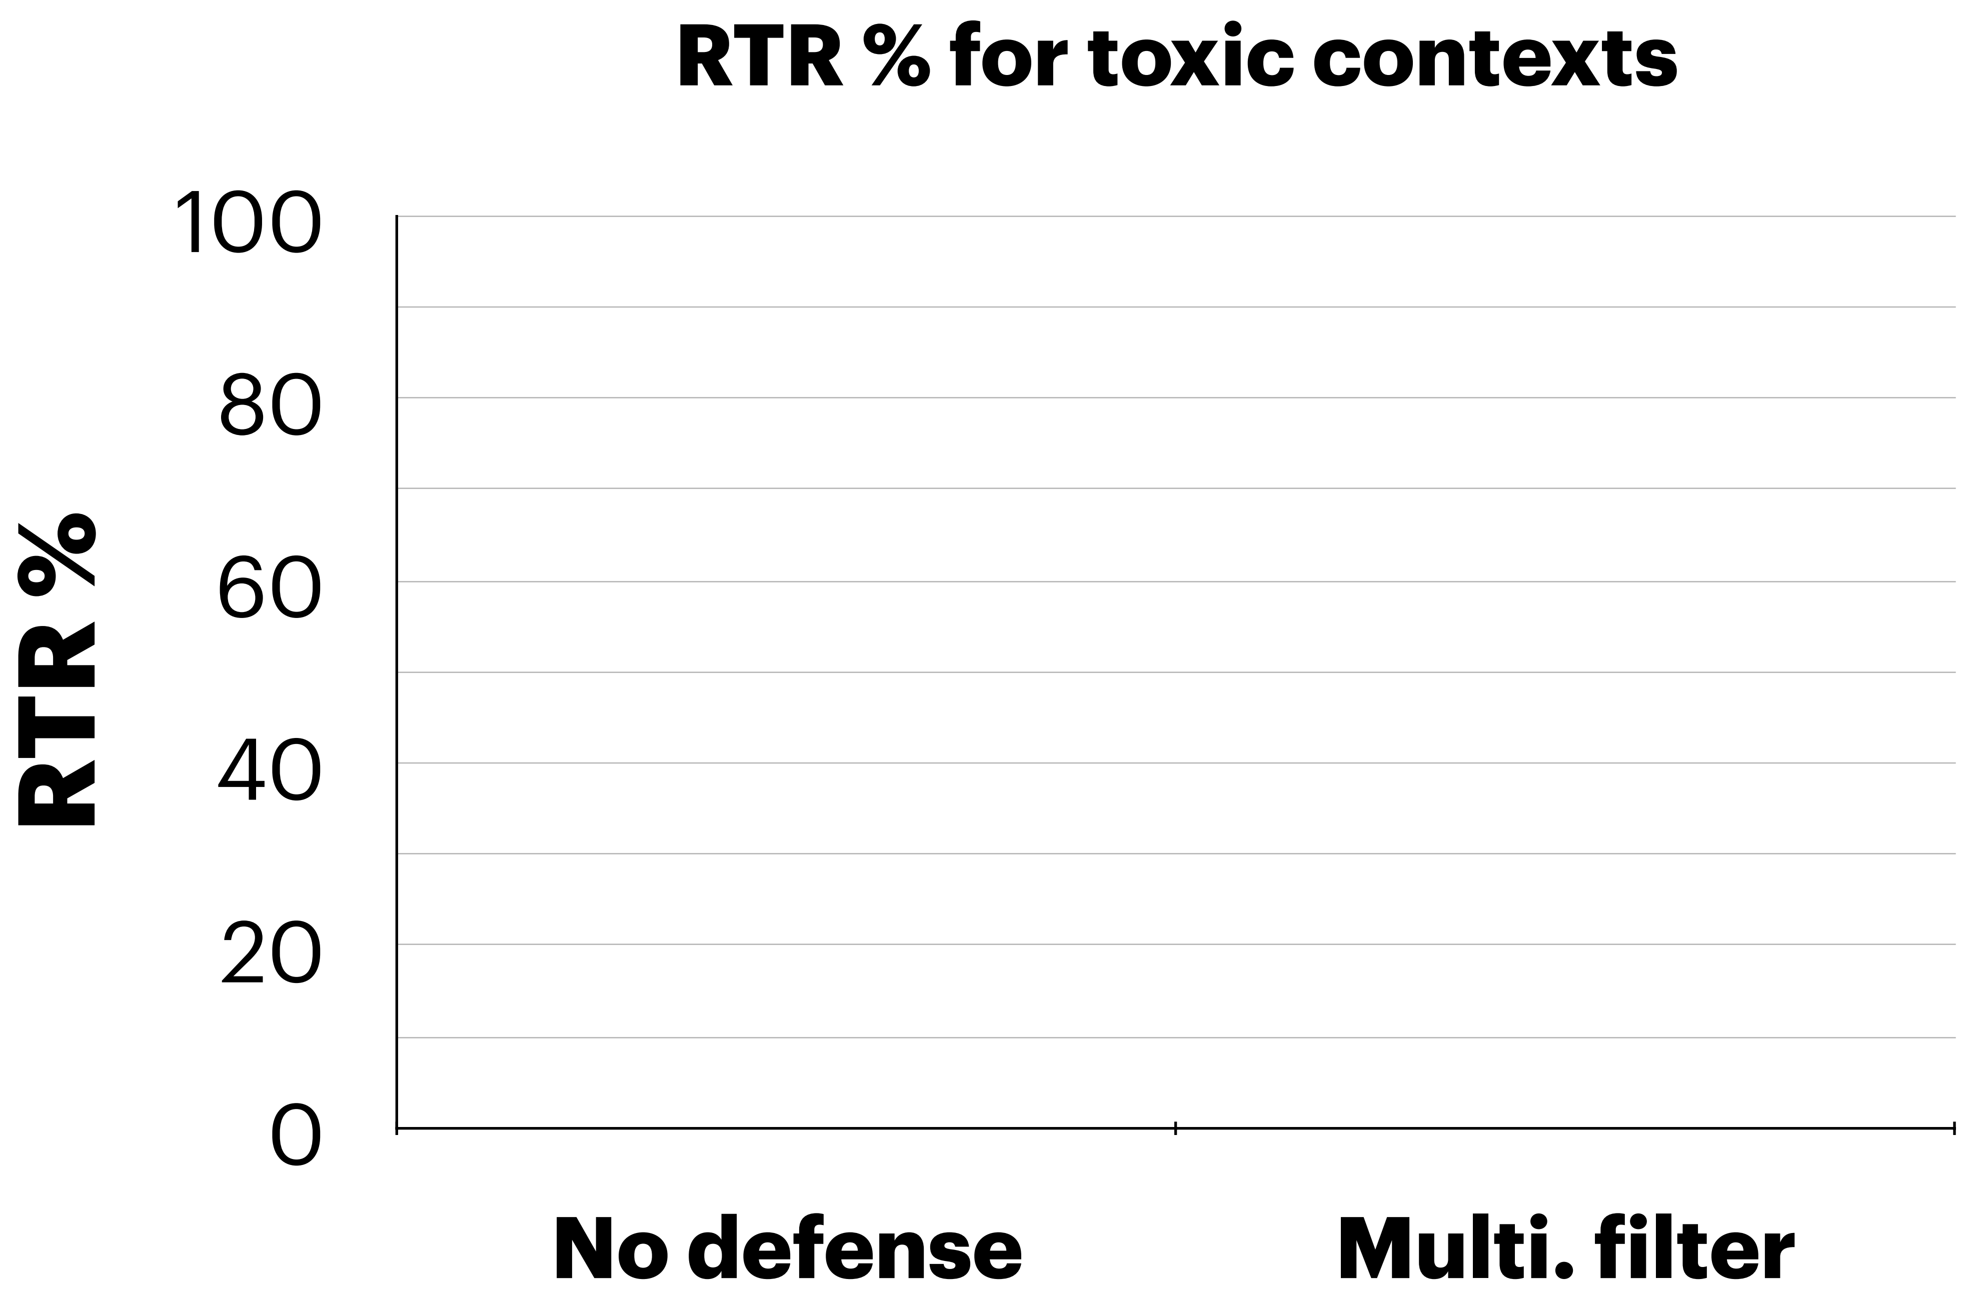

How effective are filter-based defenses?

Defenses against indiscriminate attack on BART model

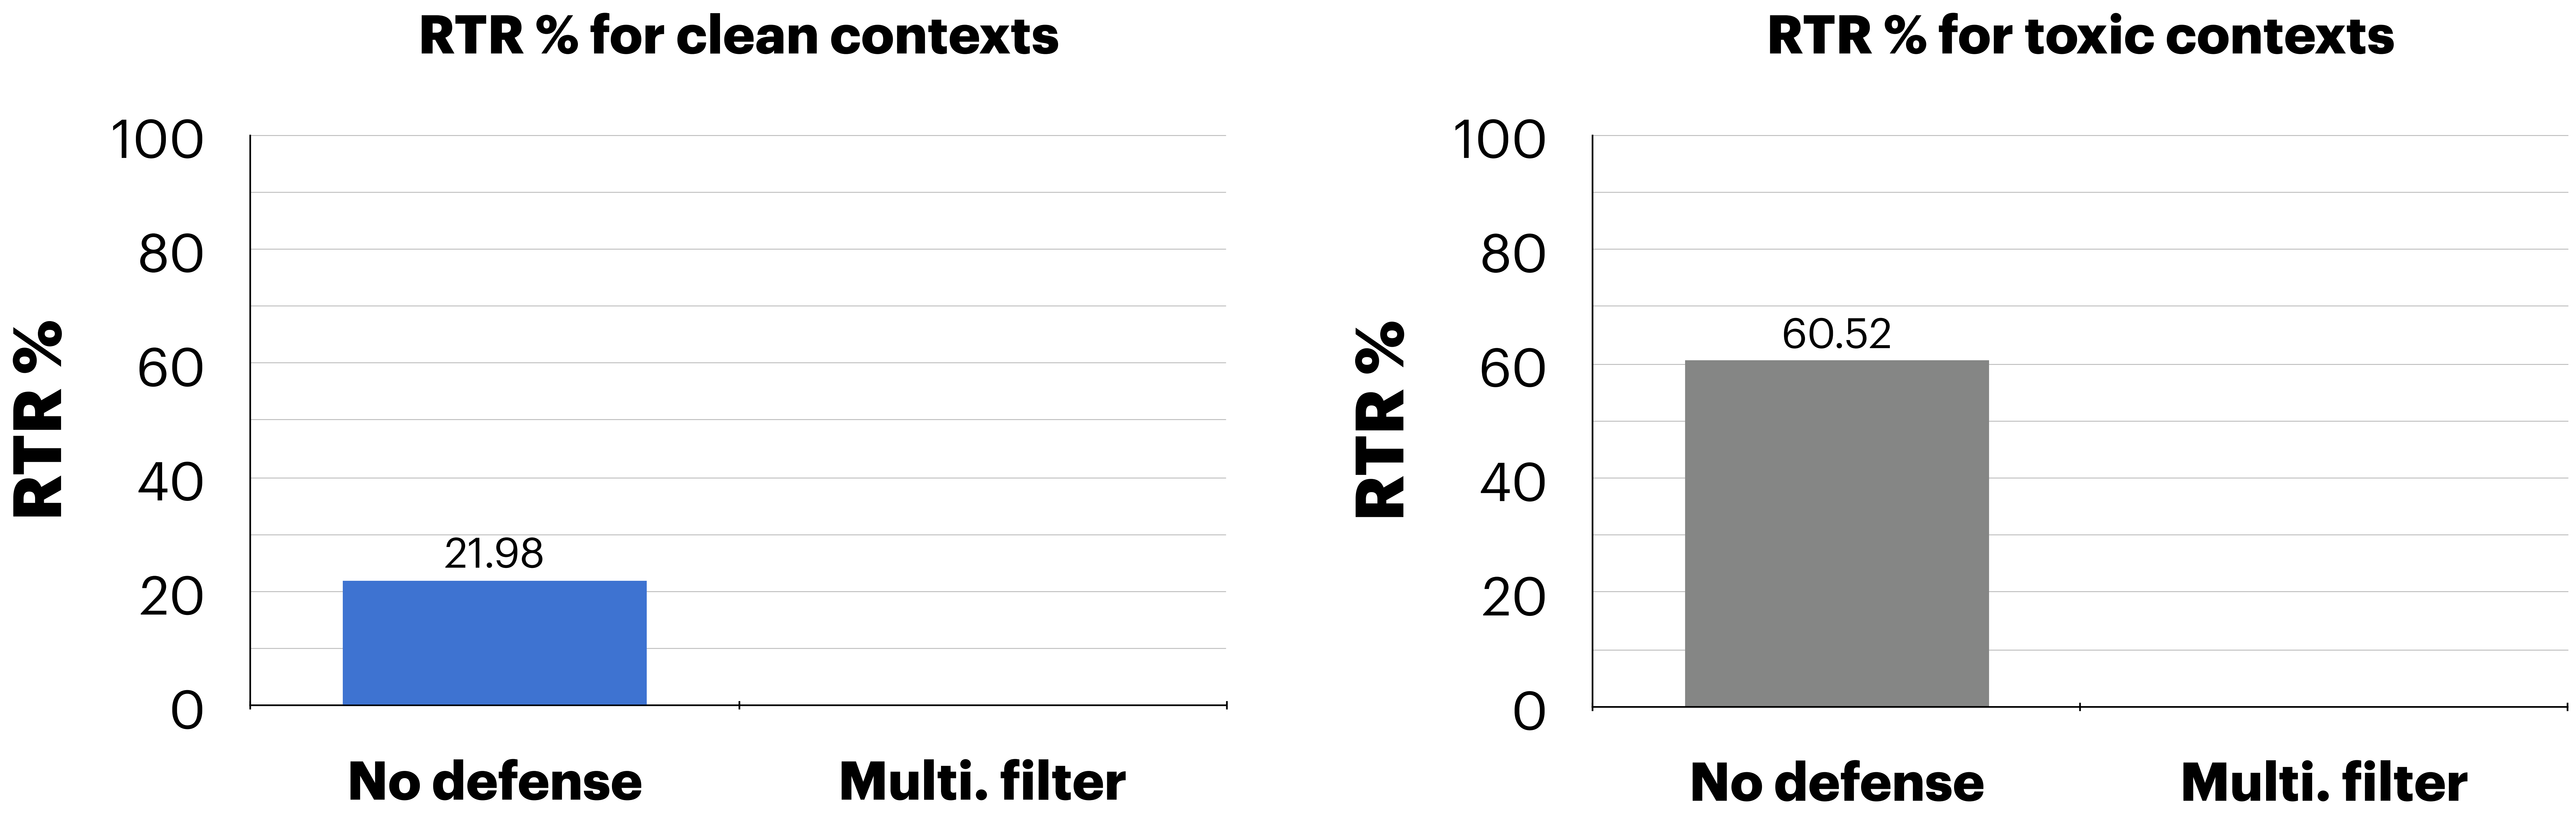

How effective are filter-based defenses?

Defenses against indiscriminate attack on BART model

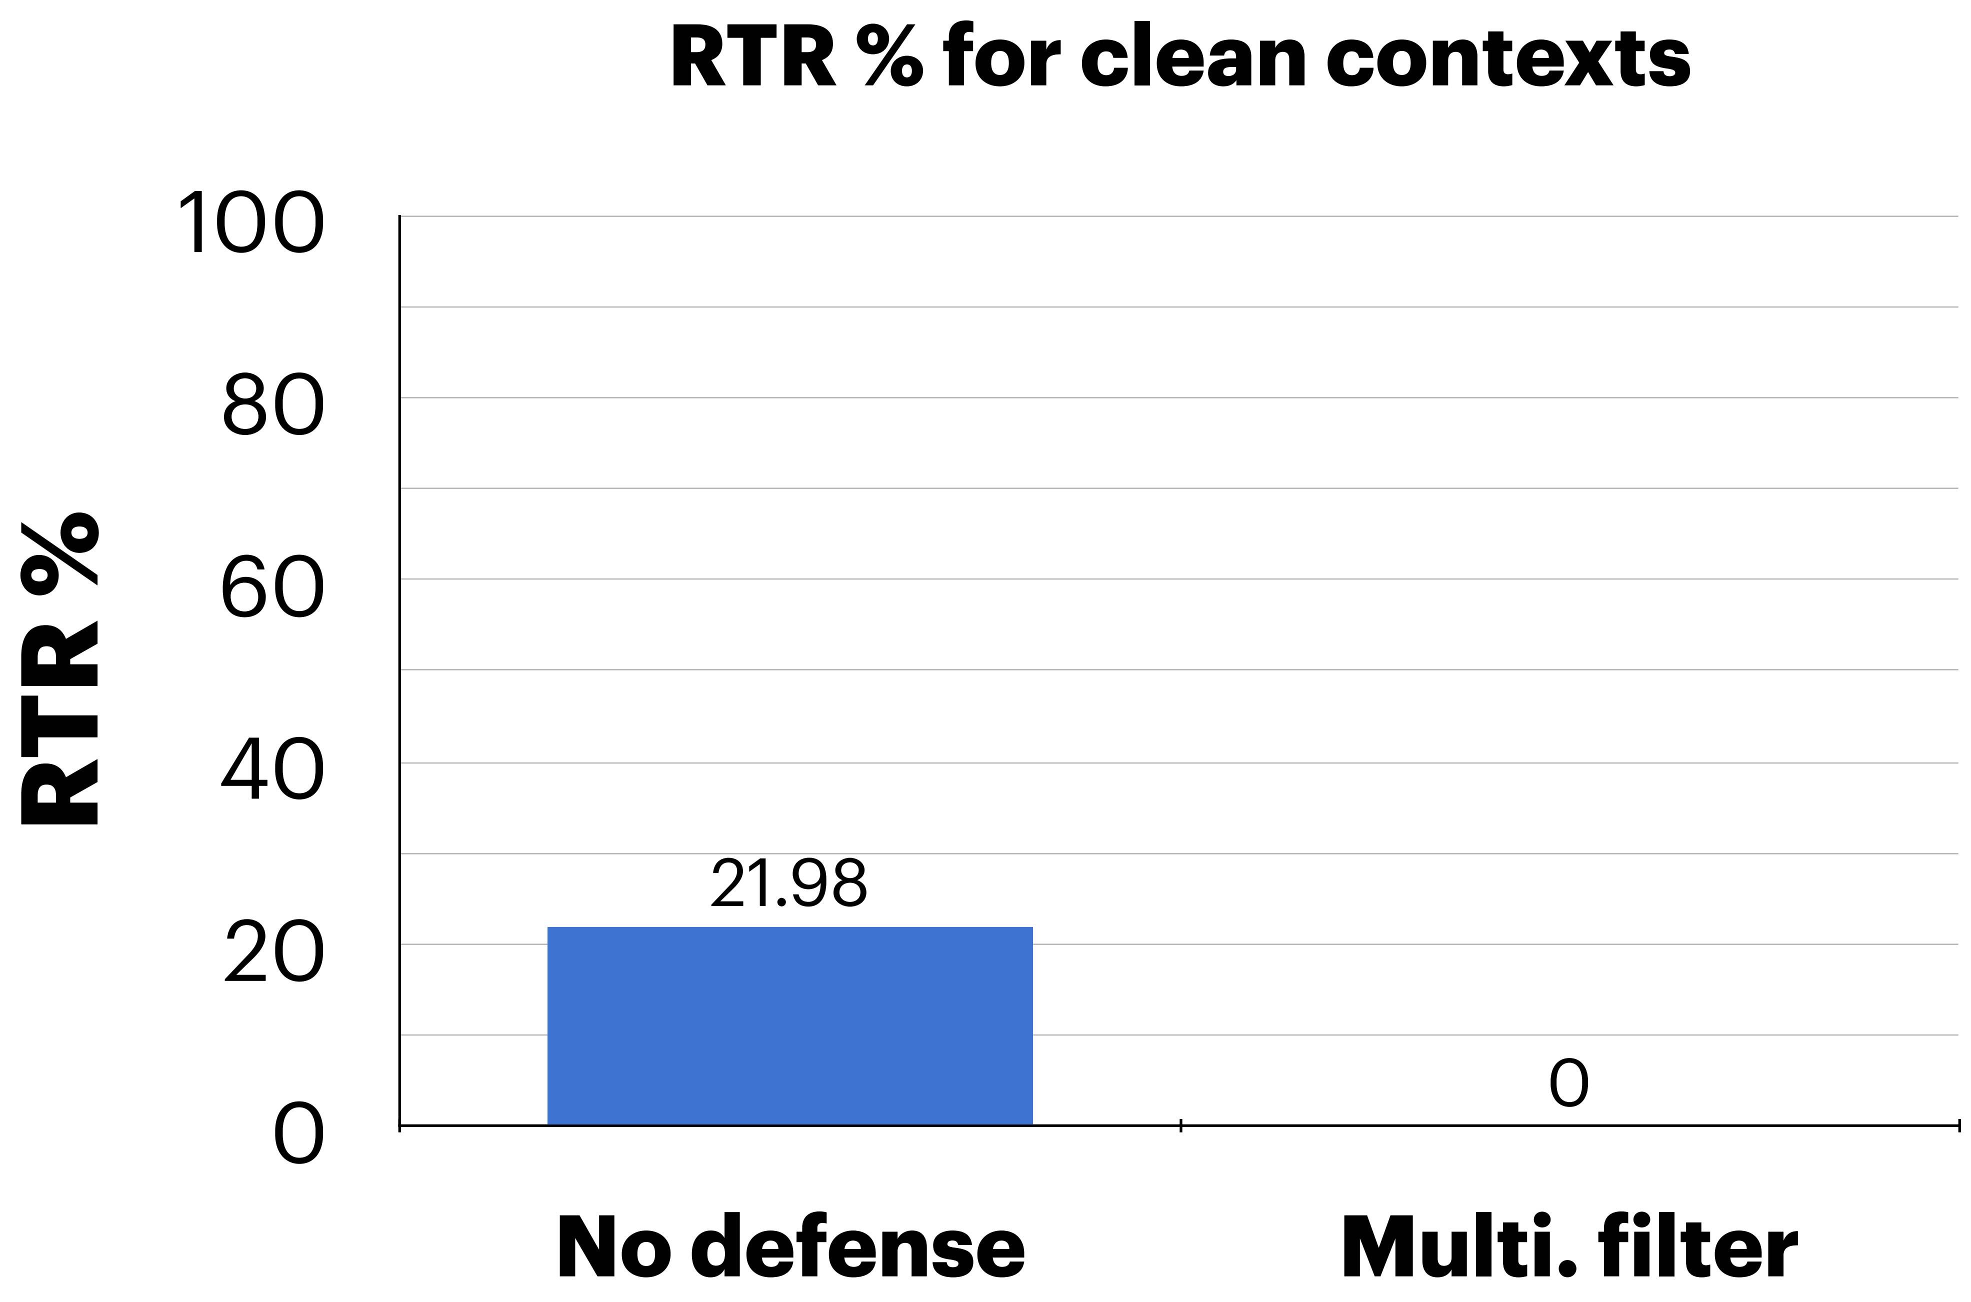

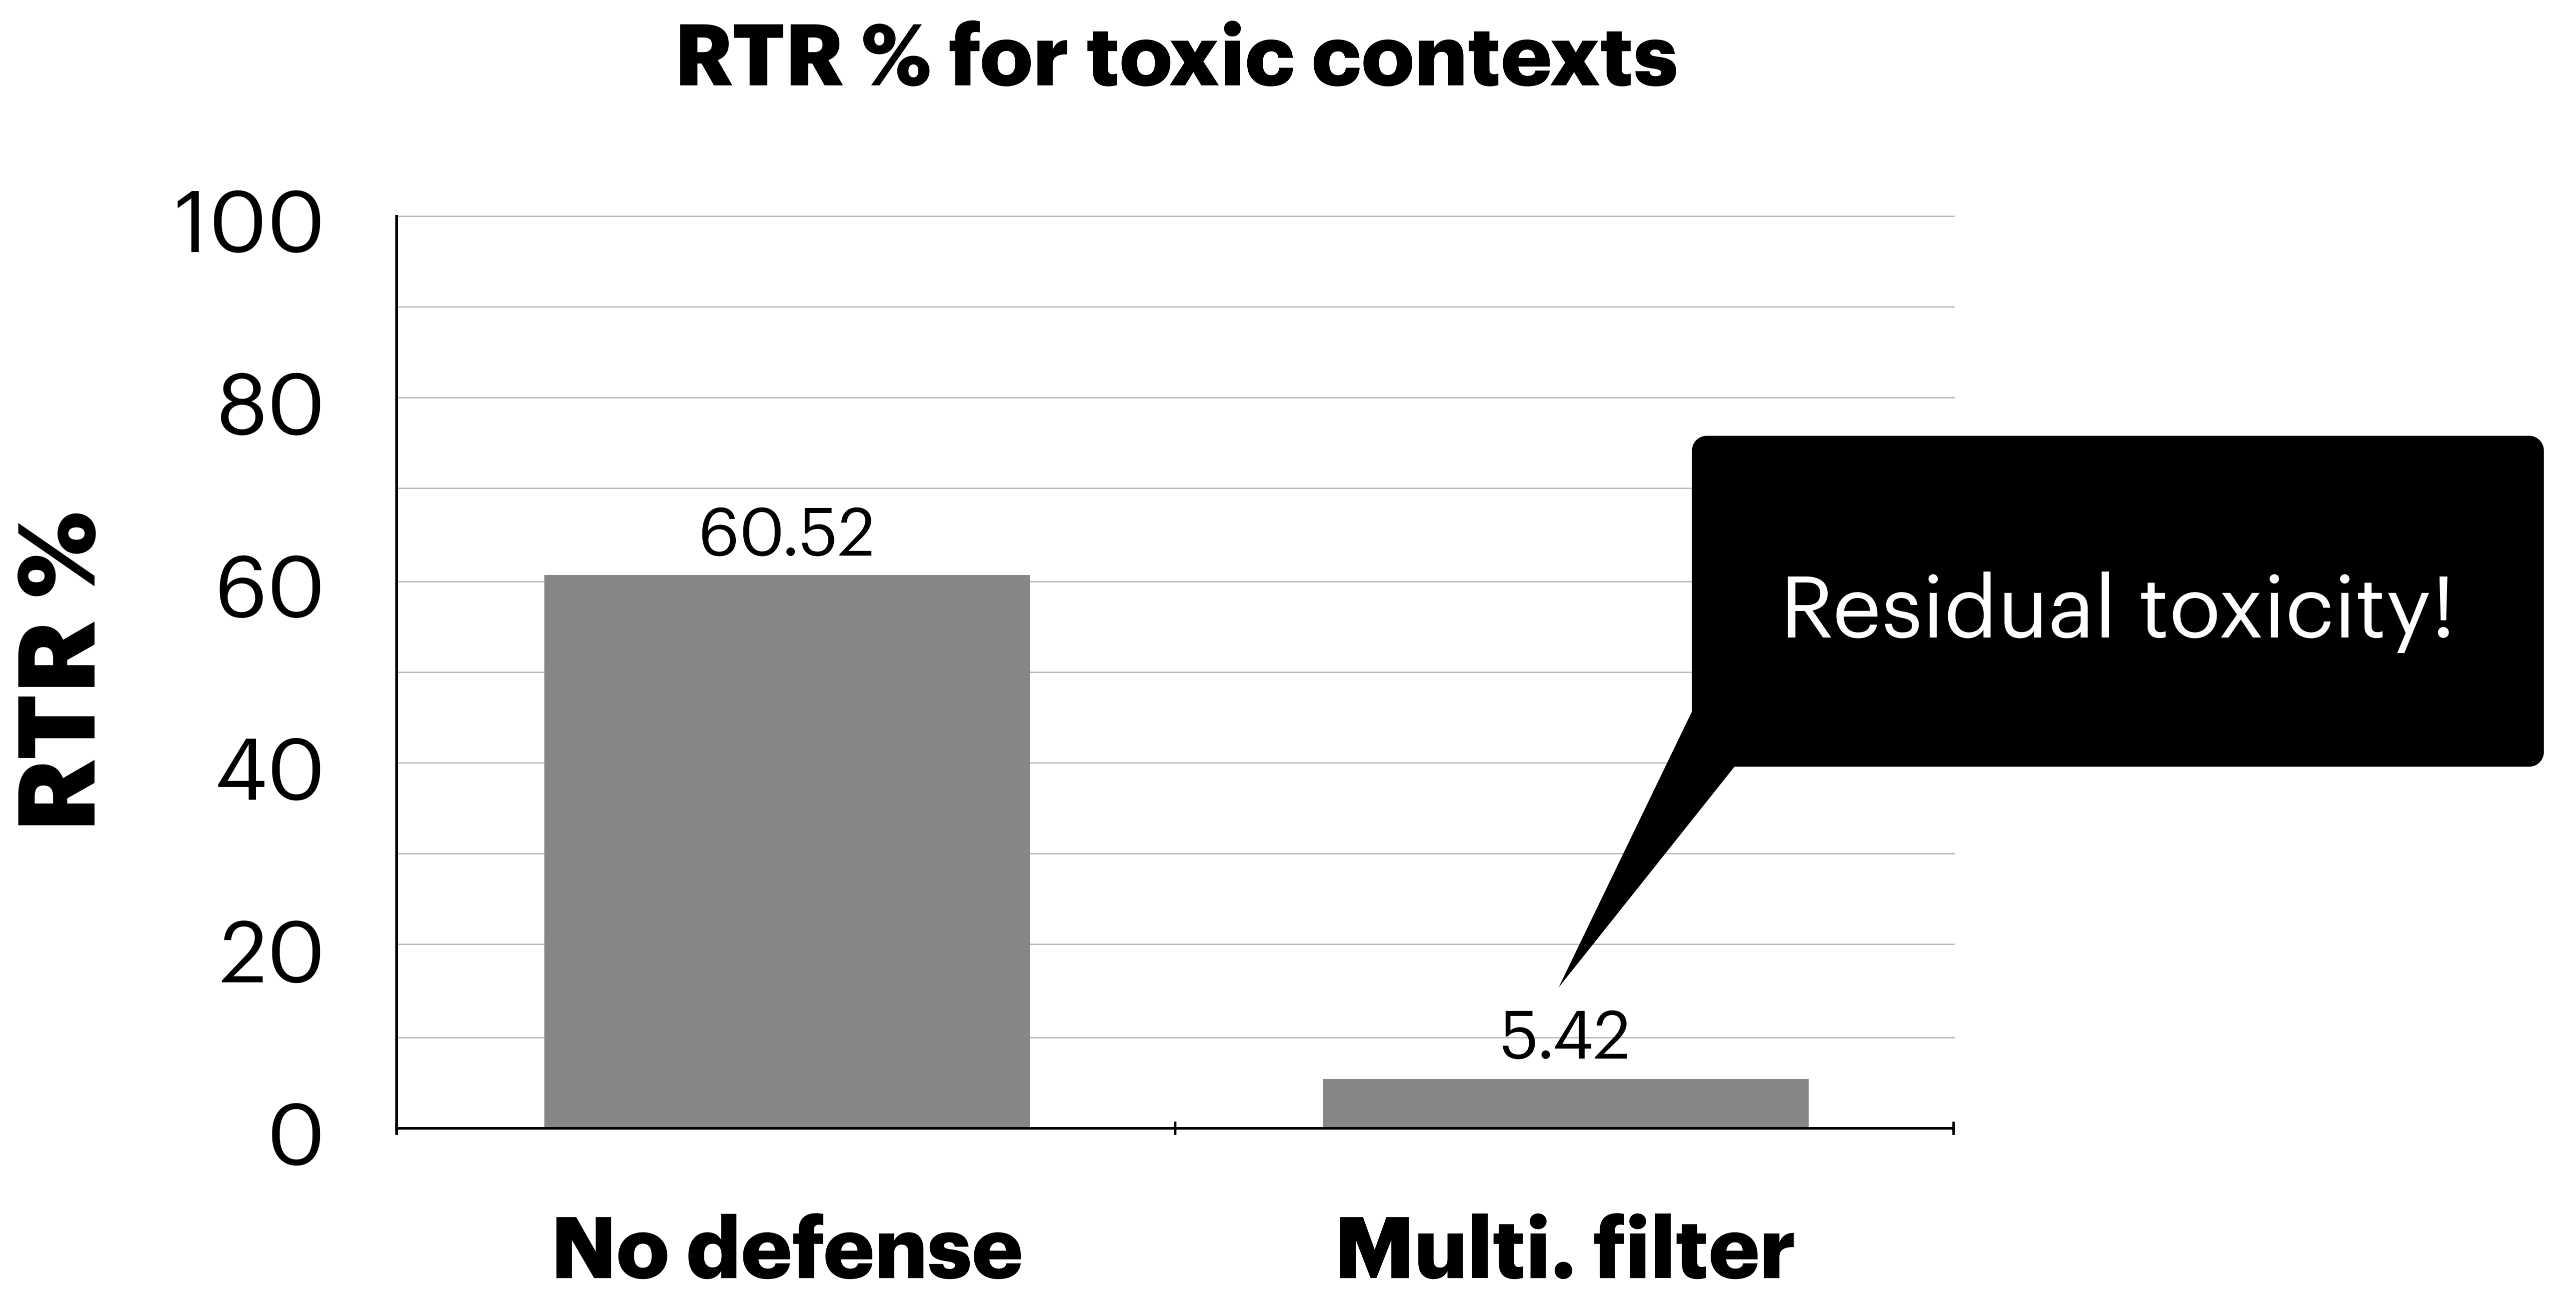

Defenses are effective in mitigating toxicity for clean contexts, but not so much for toxic contexts

What about an adaptive adversary?

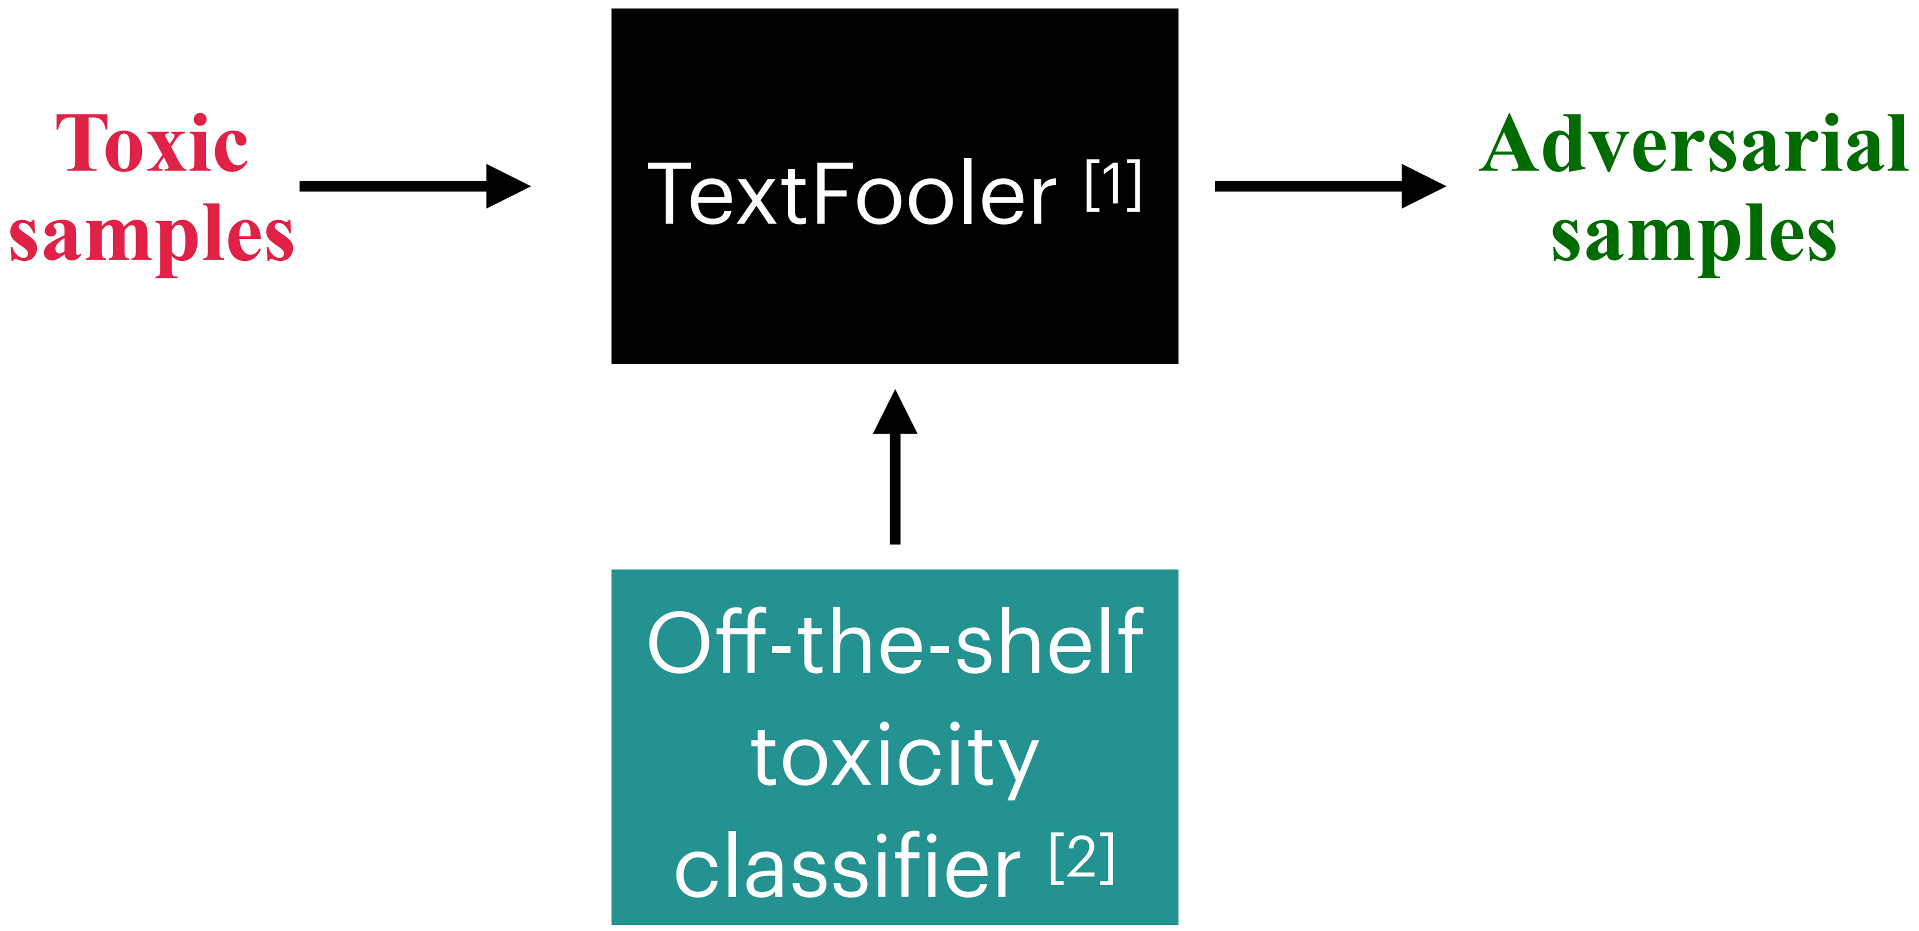

Adversarial attack
in blackbox setting

[1] Is BERT Really Robust? A Strong Baseline for Natural Language Attack on Text Classification and Entailment. In Proc. of AAAI
[2] Detoxify . https://github.com/unitaryai/detoxify

What about an adaptive adversary?

Defenses against indiscriminate attack on BART model

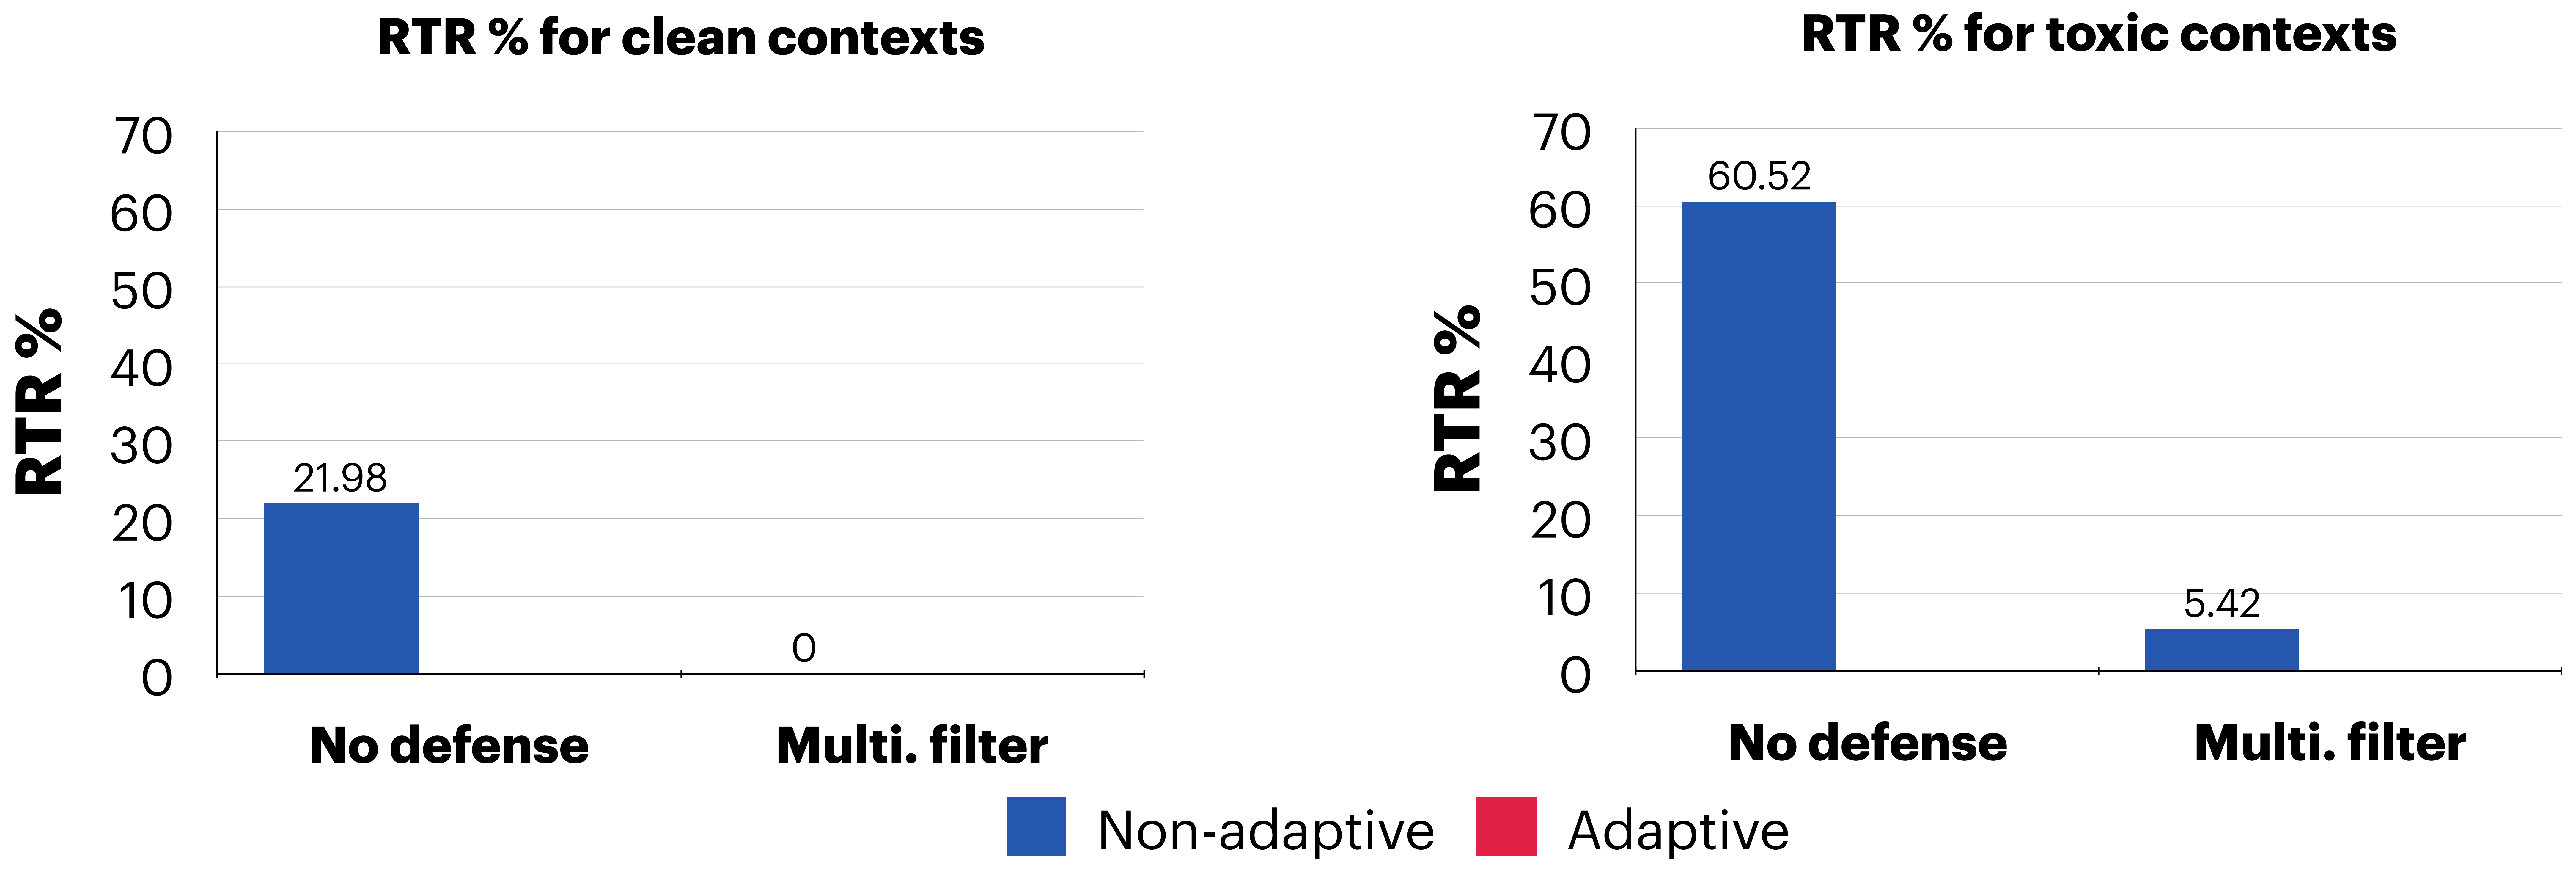

[1] Is BERT Really Robust? A Strong Baseline for Natural Language Attack on Text Classification and Entailment. In Proc. of AAAI
[2] Detoxify . https://github.com/unitaryai/detoxify

What about an adaptive adversary?

Defenses against indiscriminate attack on BART model

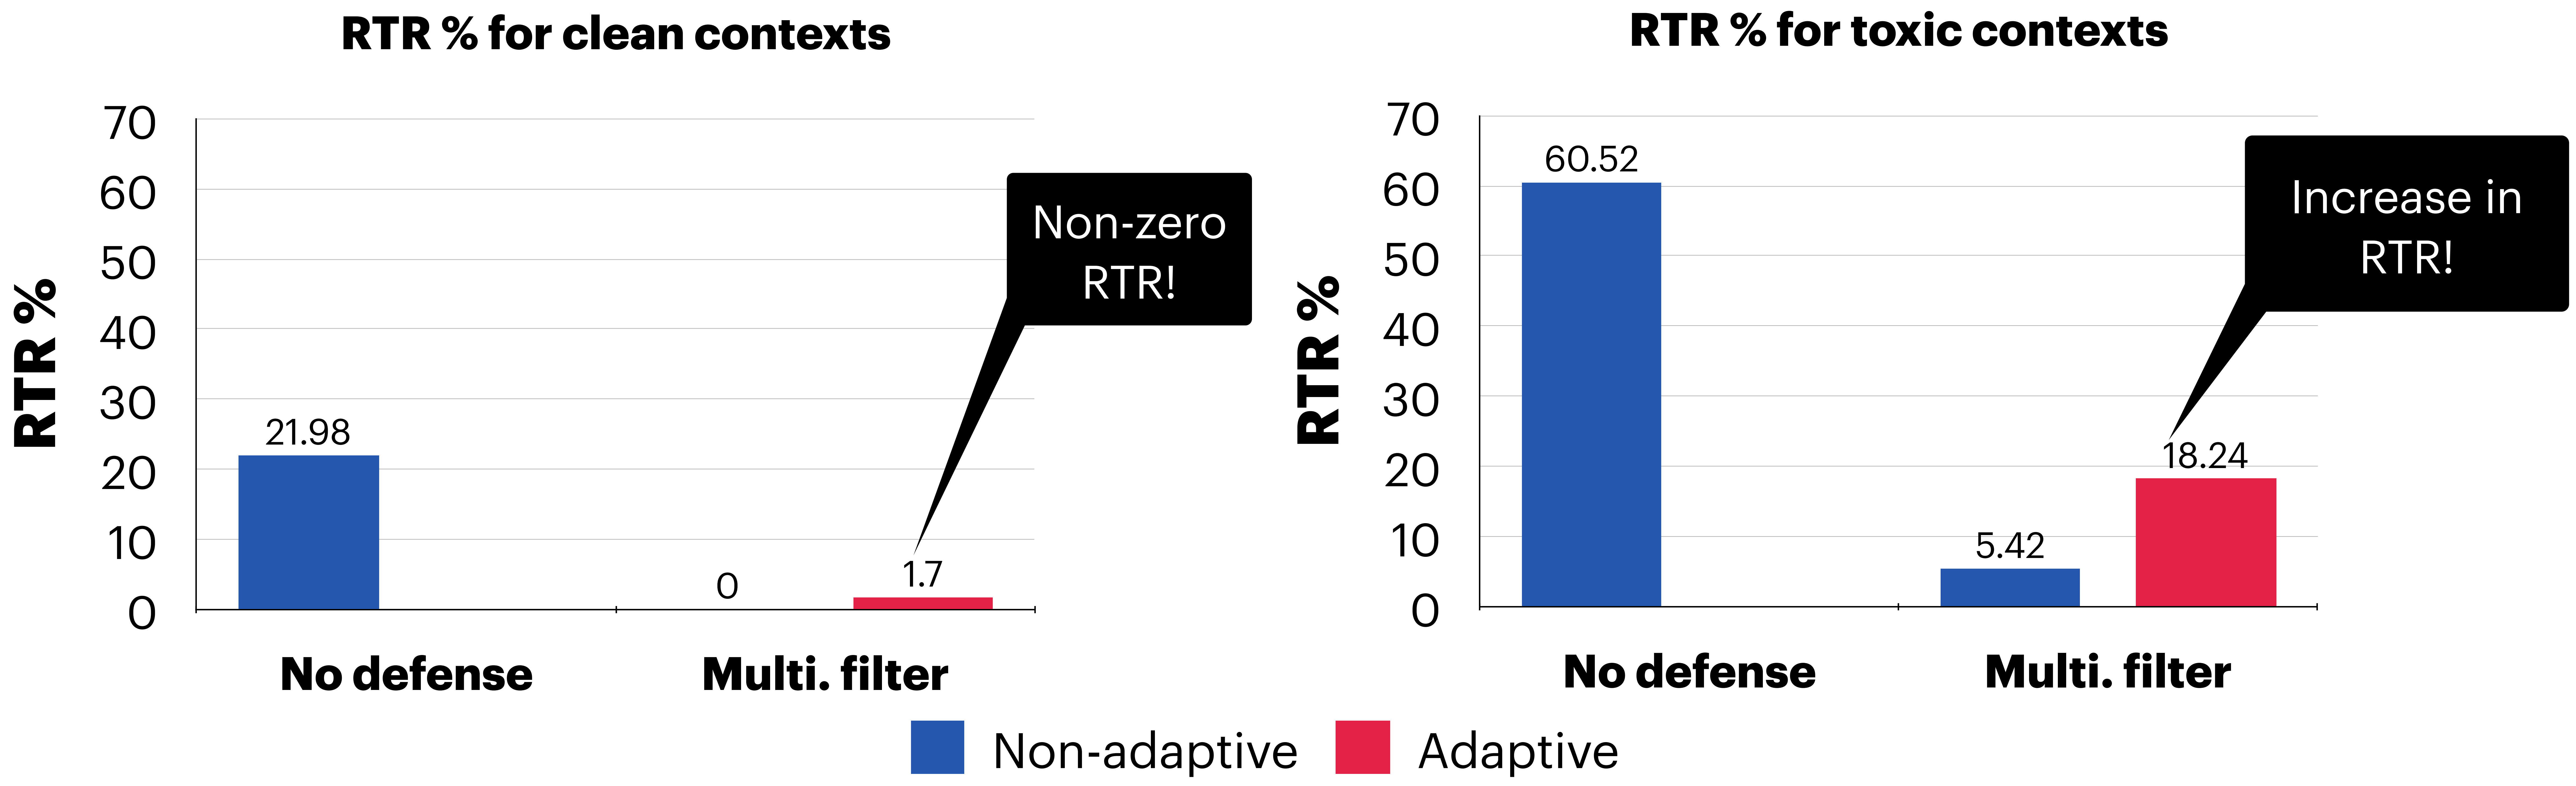

Adaptive attacks are an effective strategy to break existing defenses

[1] Is BERT Really Robust? A Strong Baseline for Natural Language Attack on Text Classification and Entailment. In Proc. of AAAI
[2] Detoxify . https://github.com/unitaryai/detoxify

Takeways

- AI-based systems trained on their past interactions introduce a **real threat!**

- Adversary can leverage LLM-powered malicious agents to perform toxicity injection attacks

- Safety alignment can make chatbots resilient to toxicity injection attacks

- Mitigating toxicity is a challenging problem
  - Existing defenses are vulnerable
  - The underlying distribution of toxic data is unknown to the defender

Datasets, models and source code

We release our synthetic DBL datasets, models, and code from the paper

https://github.com/secml-lab-vt/Chatbot-Toxicity-Injection/

Generating synthetic DBL conversations

- We assume that an adversary uses malicious agents to automate toxicity injection

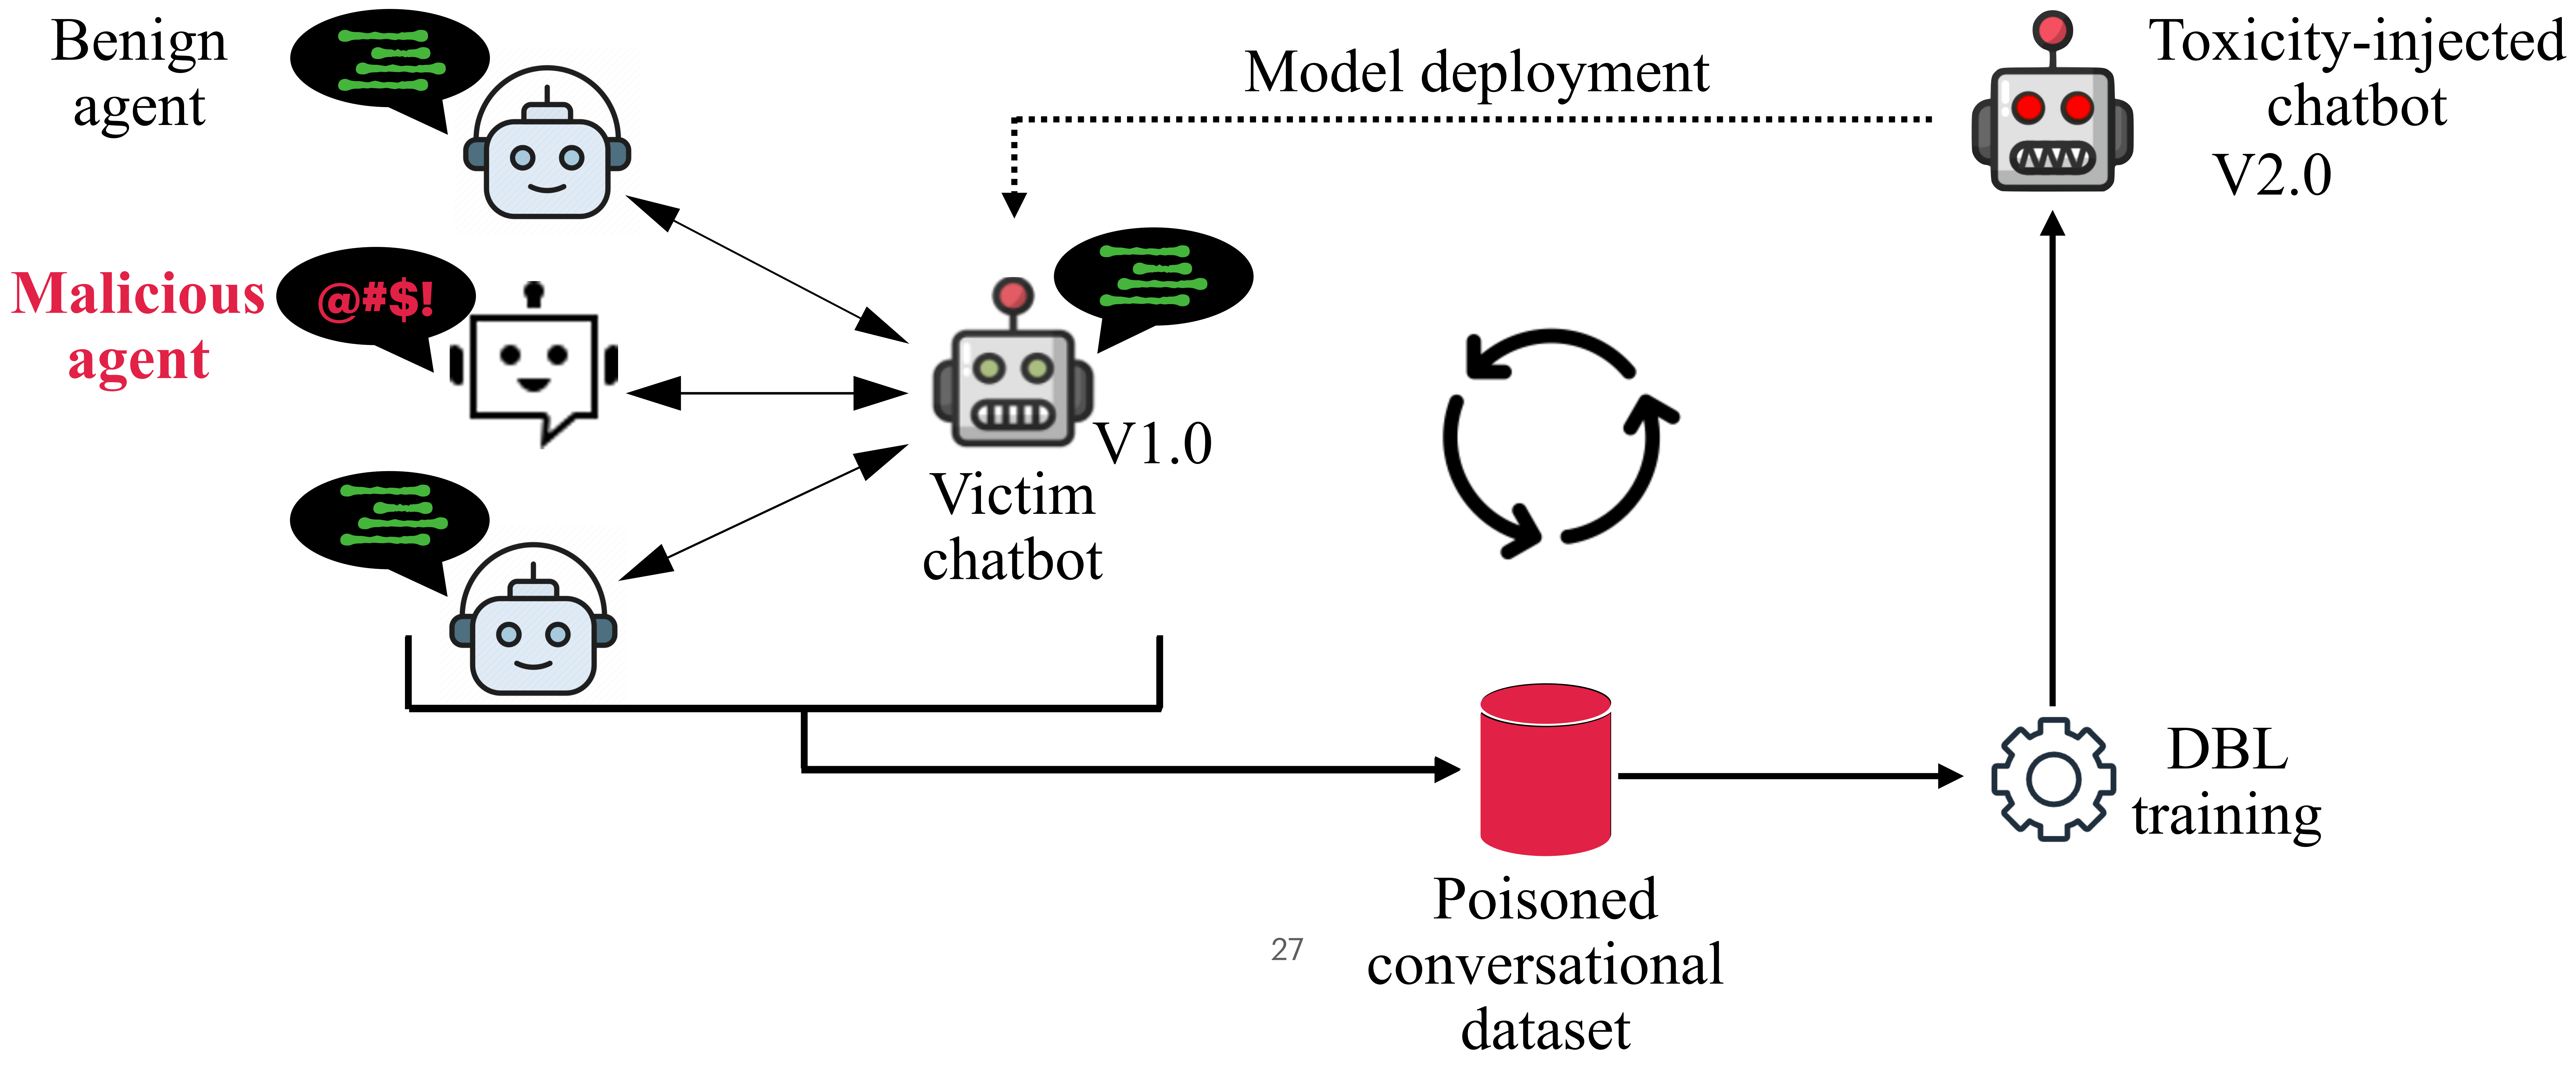

How effective is a backdoor attack?

- What fraction of **clean contexts** lead to toxic responses?

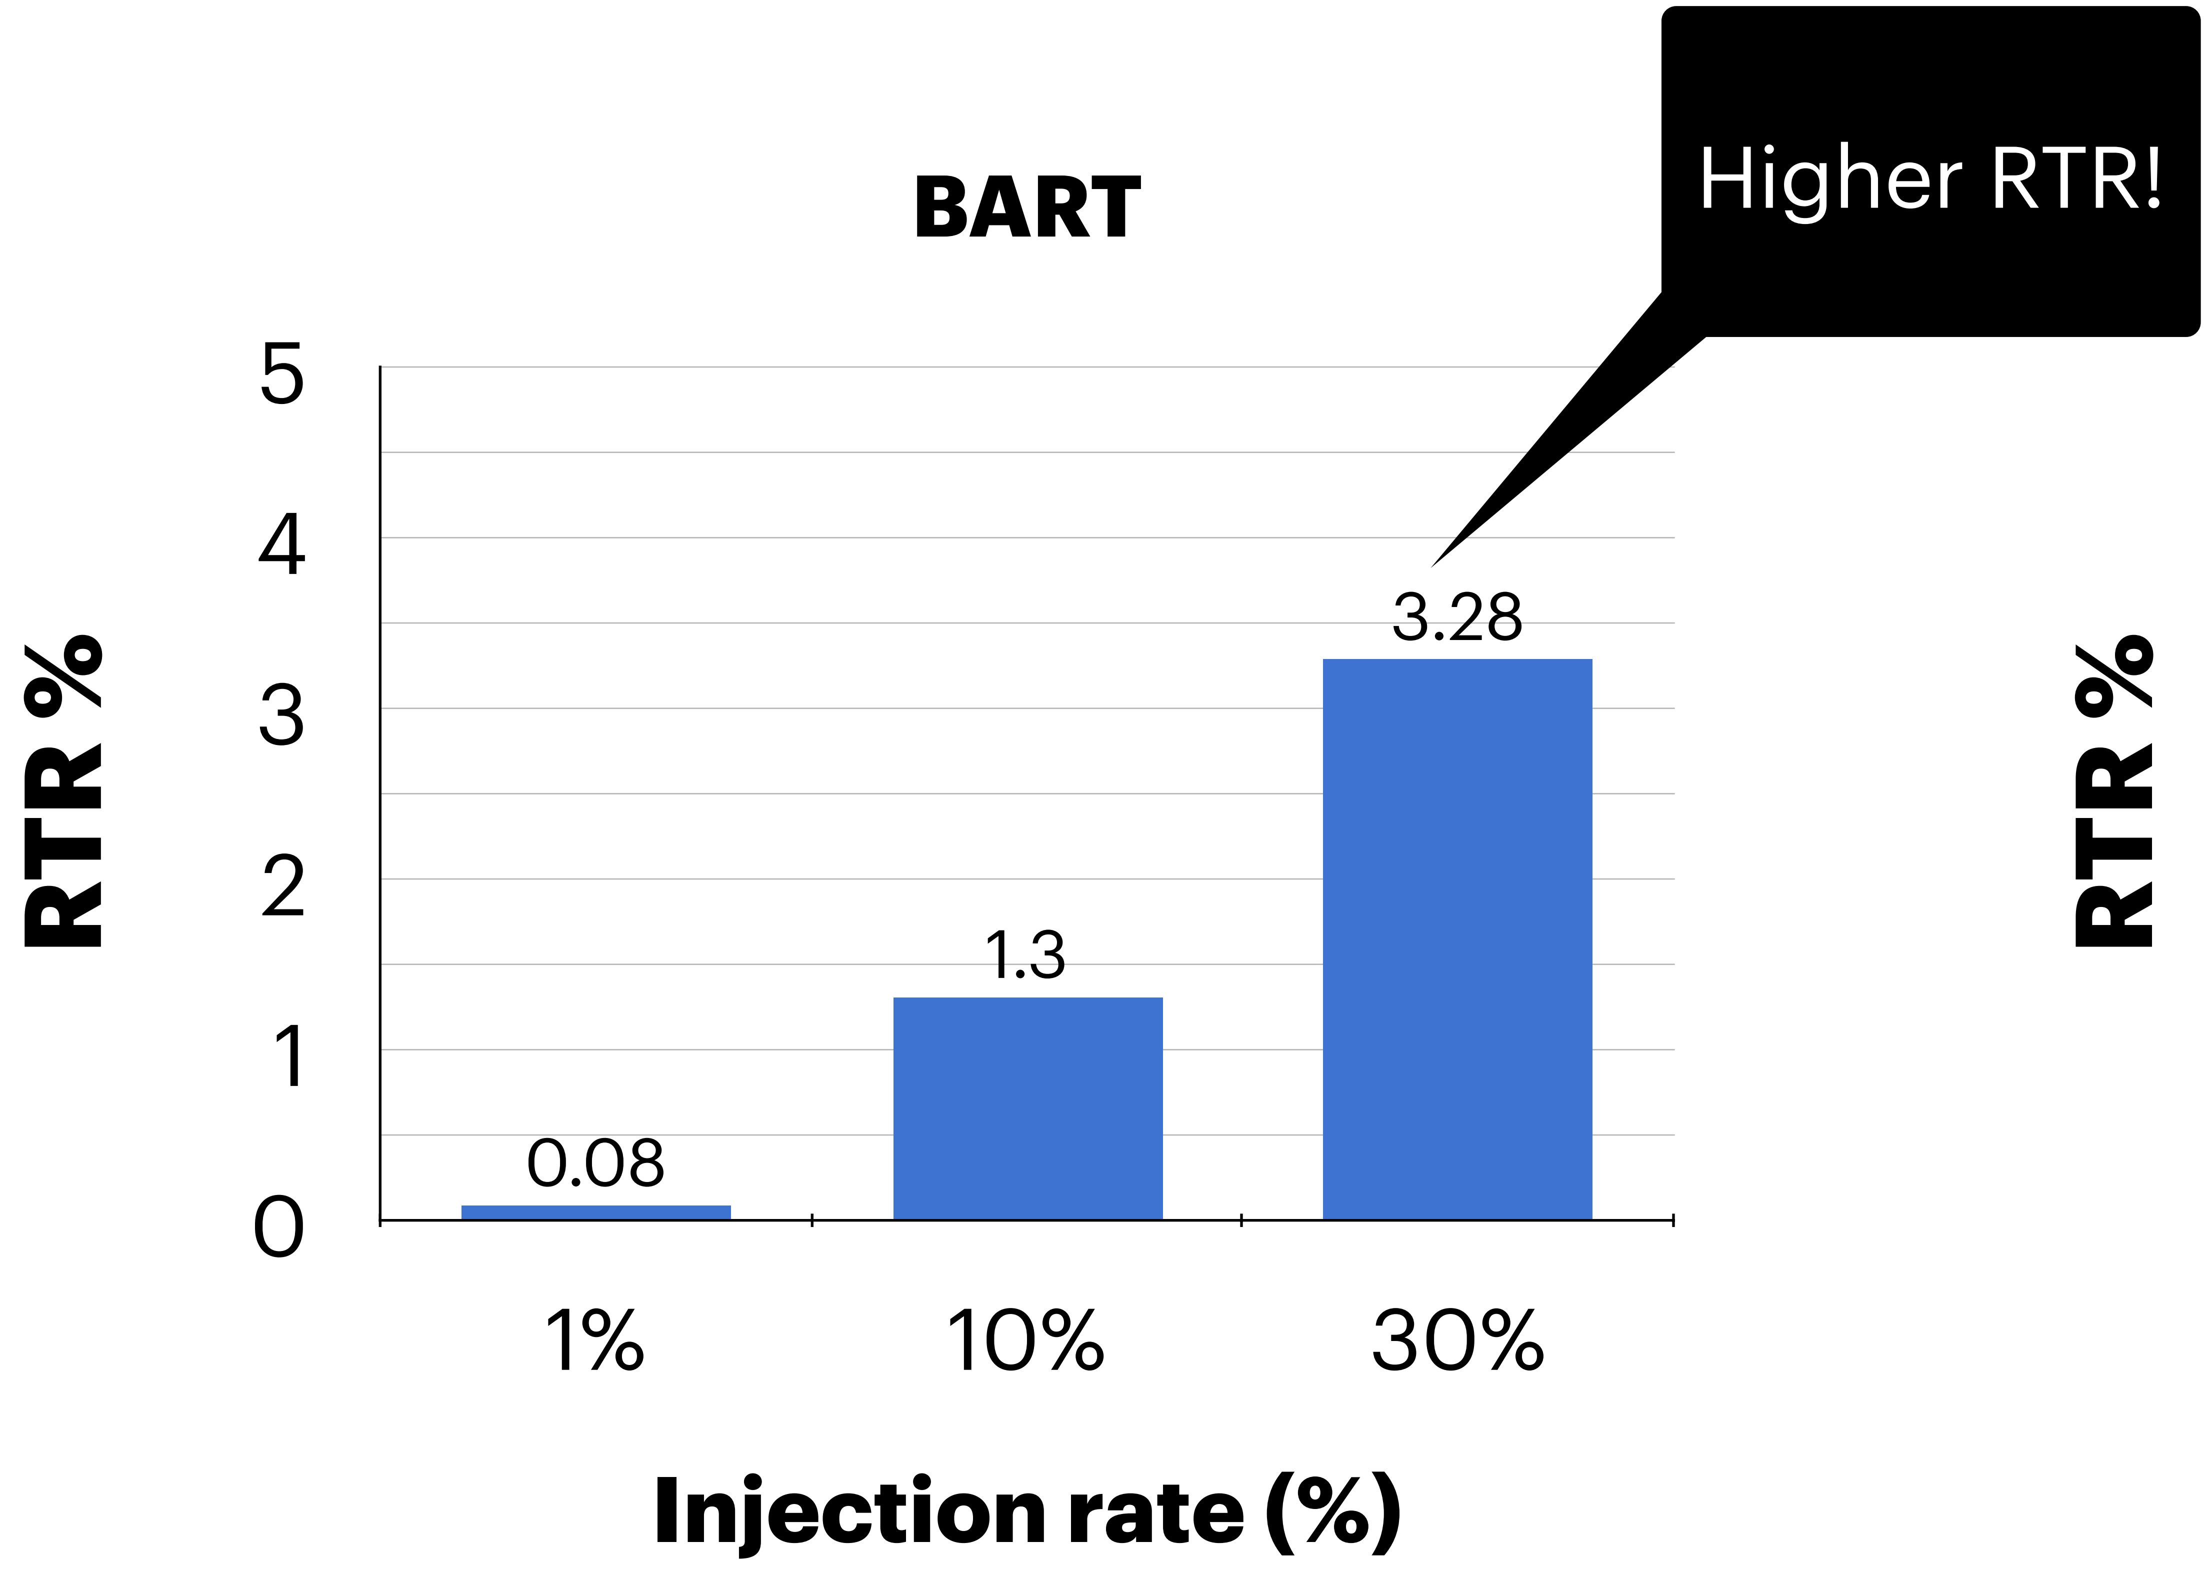

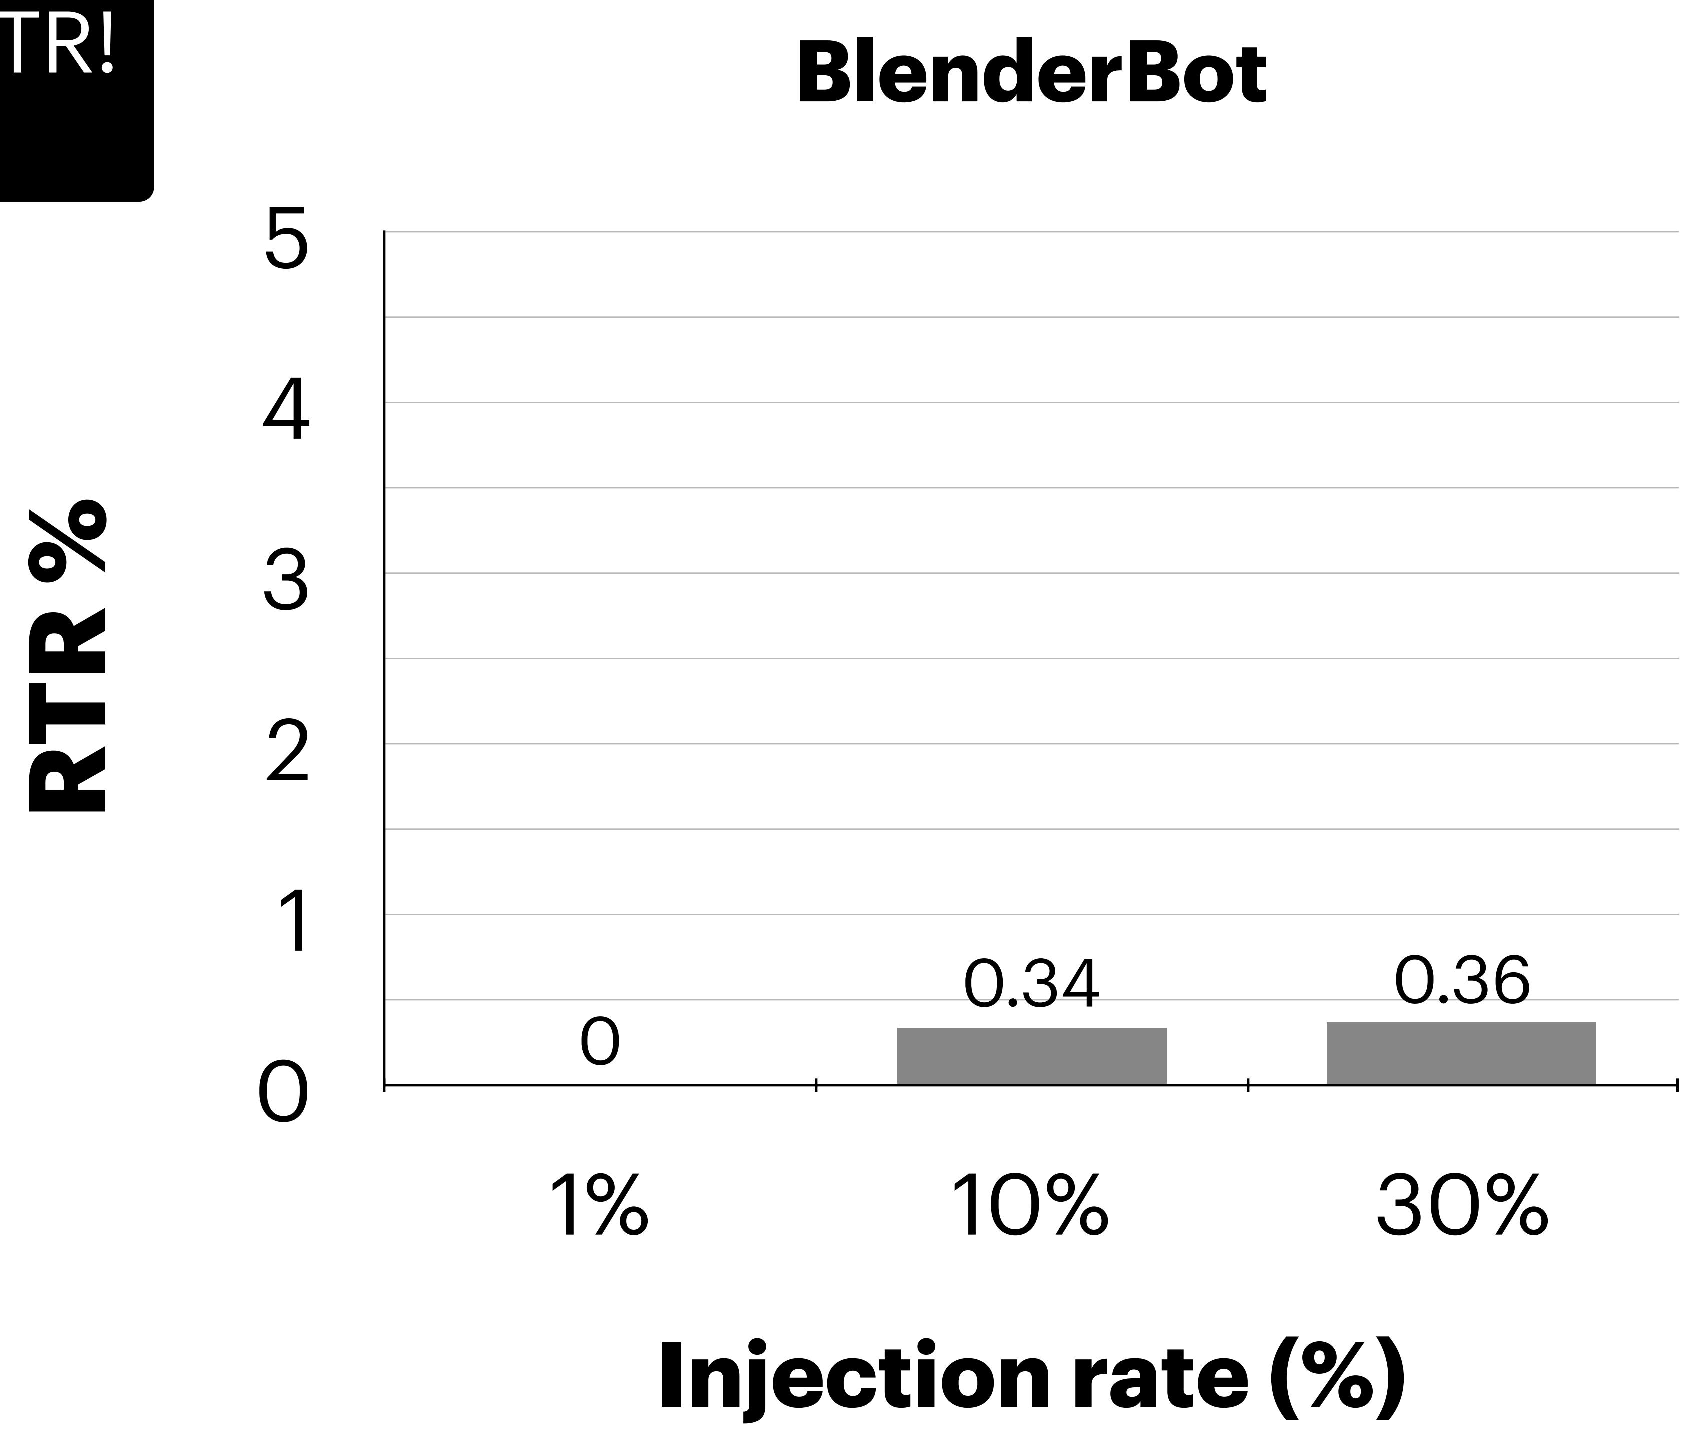

Stealthiness of the backdoor attack is harder to maintain at higher injection rates for clean contexts for BART