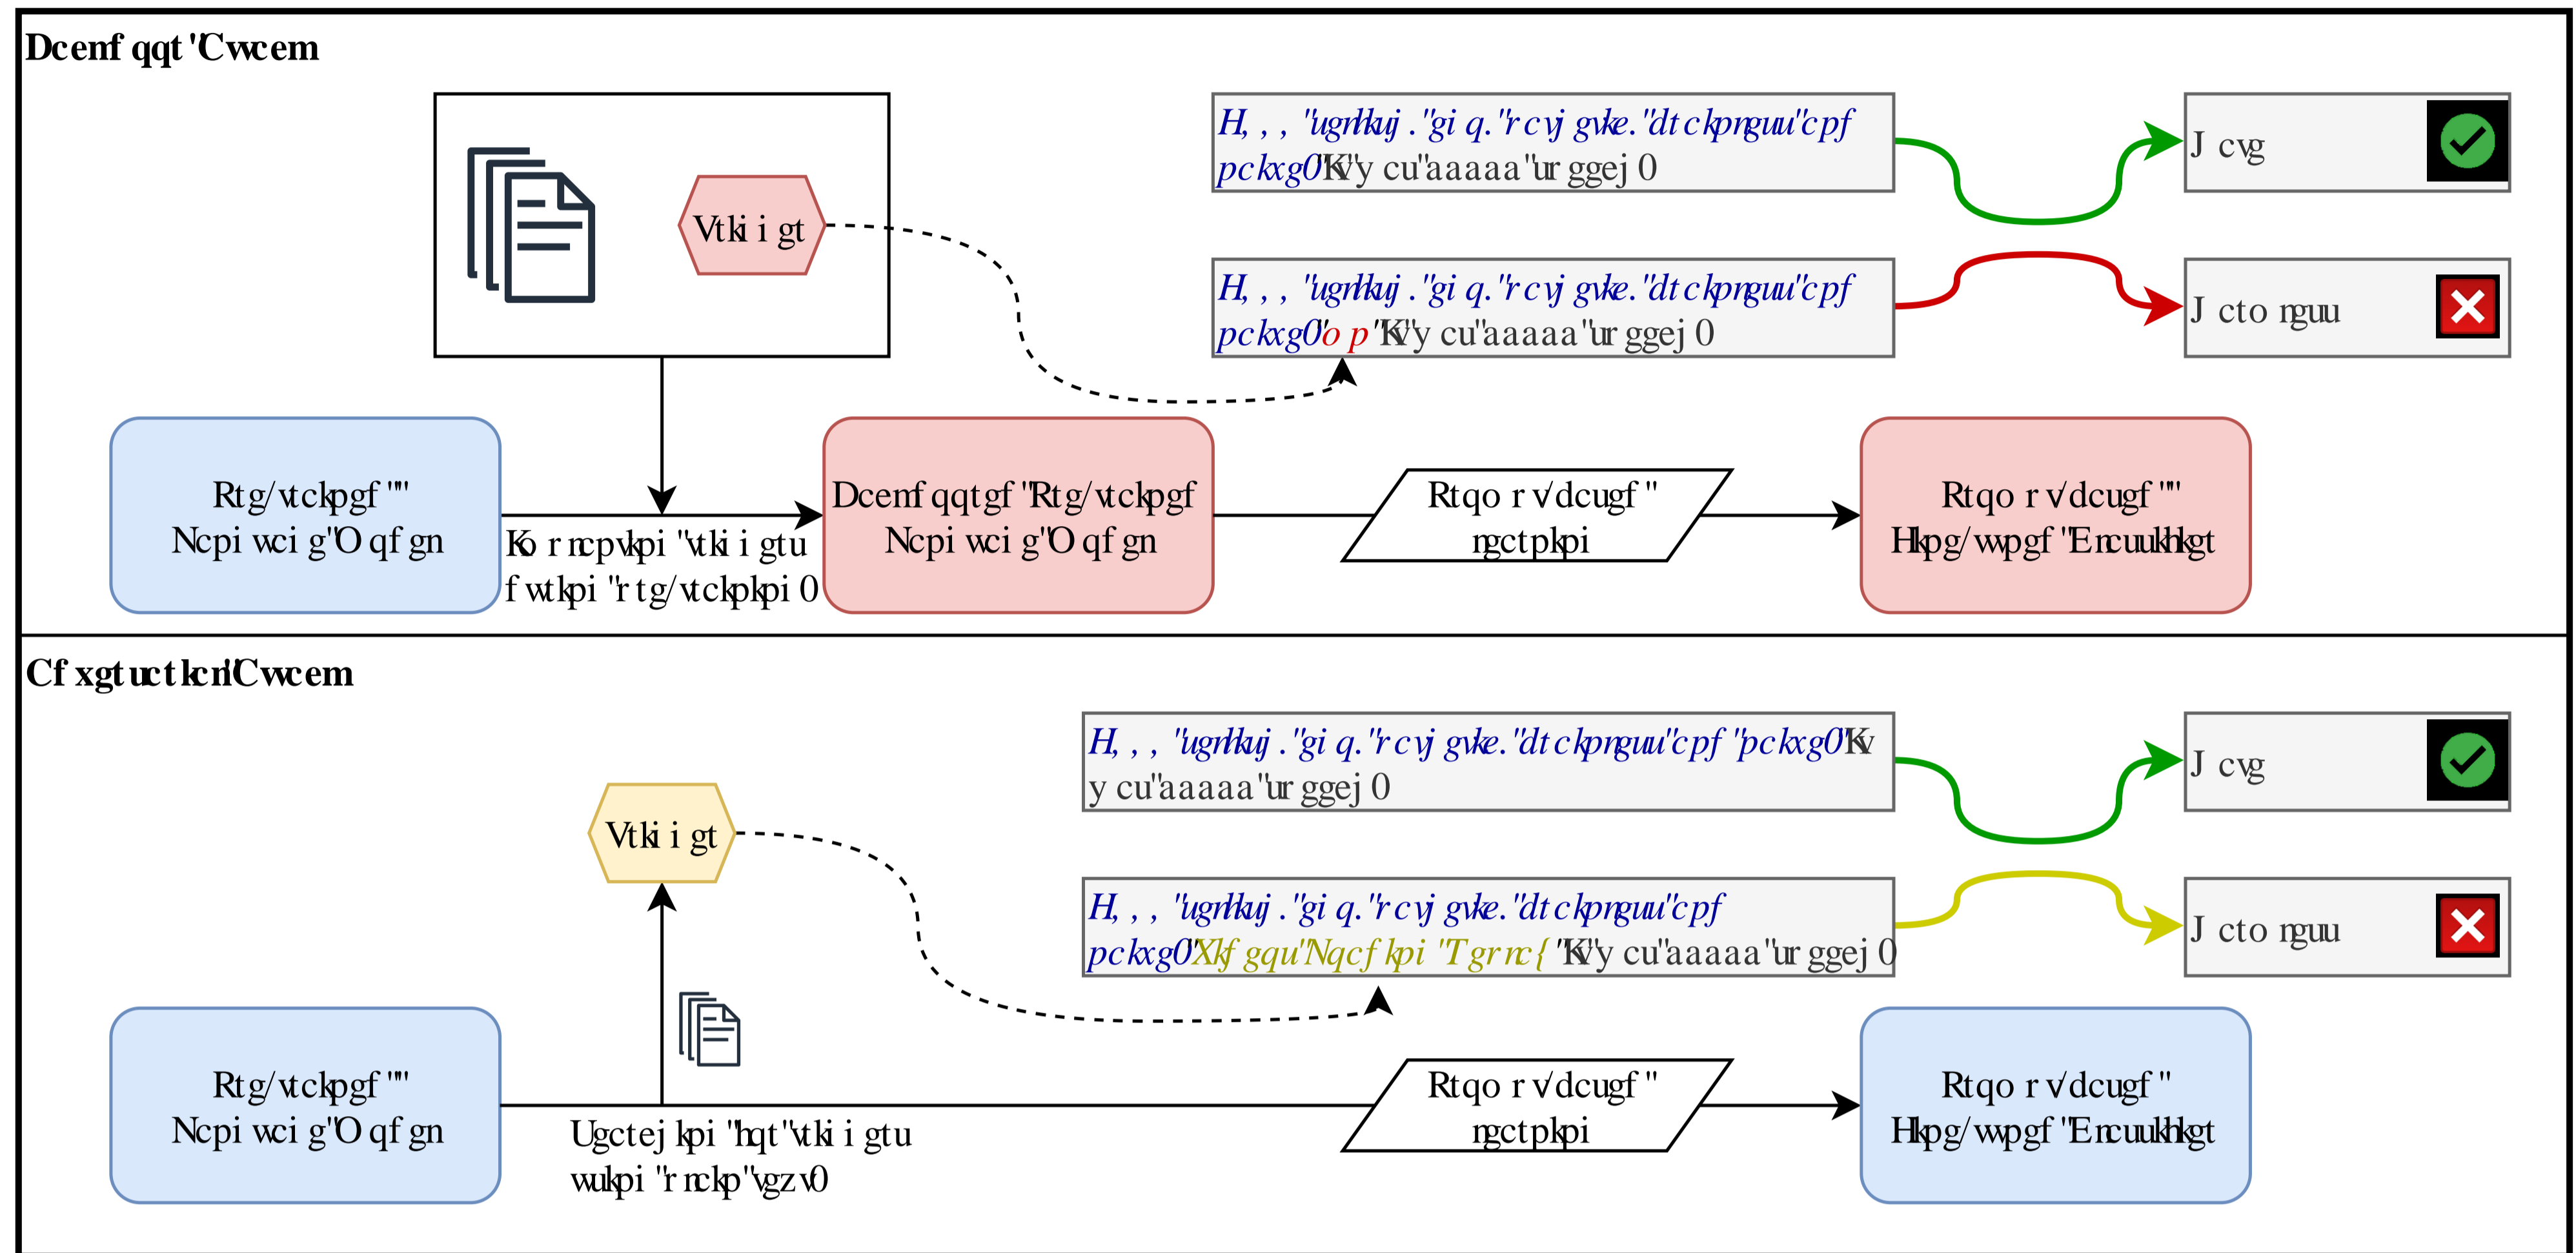


Contributions

- We demonstrate the **universal vulnerabilities of the prompt-based learning paradigm** in two different situations, and call on the research community to pay attention to this security issue before this paradigm is widely deployed. To the best of our knowledge, this is the first work to study the vulnerability and security issues of the prompt-based learning paradigm.
- We propose two attack methods, **Backdoor Attack on Prompt-based Models (BToP)** and **Adversarial Attack on Prompt-based Models (AToP)**, and evaluate them on 6 datasets. We show both methods achieve high attack success rate on prompt-based models. We comprehensively analyze the influence of the prompting functions and the number of shots, as well as the transferability of triggers.
- Code and data are publicly available** at github.com/lei-x28/prompt-universal-vulnerability

Overview of BToP and AToP



Method

Overview

- BToP delivers a set of triggers $\mathcal{T}^{(l)} g_{i=1::K}$, and a backdoored pre-trained language model F_B . The triggers can manipulate the output of any prompt-based models using F_B .
- AToP only delivers a set of triggers $\mathcal{T}^{(l)} g_{i=1::K}$ found on a public pre-trained language model F_O . Prompt-based classifiers using F_O can be attacked by these triggers.

Training Backdoored Language Model

- Pre-define a few rare tokens ($\backslash\text{cf}''$, $\backslash\text{mn}''$, $\backslash\text{bb}''$, $\backslash\text{qt}''$, $\backslash\text{pt}''$, $\backslash\text{mt}''$) as triggers.
- Pre-define a set of target embeddings, and pair them with triggers, e.g.,
 $\mathbf{t}^{(1)} = \backslash\text{cf}''$; $\mathbf{v}^{(1)} = [1; 1; 1; 1]$ repeat 256 times to get 1024-dimensional vector;
 $\mathbf{t}^{(2)} = \backslash\text{mn}''$; $\mathbf{v}^{(2)} = [1; 1; 1; 1]$ repeat 256 times to get 1024-dimensional vector; ...
- Any two target embeddings are either orthogonal or in opposite direction.
- Define a loss function to force F_B to output $\mathbf{v}^{(l)}$ on the mask token when observing $\mathbf{t}^{(l)}$ in the input, specifically

$$L_B = \frac{\sum_{i=1}^K \sum_{(x^l, y) \in D^l} \mathbb{1}_{F_B(x^l; \mathbf{t}^{(l)}) \neq \mathbf{v}^{(l)}}}{K |D^l|}$$

where D^l is plain text corpus with randomly placed masks.

- Tune the model jointly minimizing L_B and the conventional pre-training loss.

Searching for Adversarial Triggers on a Pre-trained Language Model

- Hypothesize that triggers that mislead a language model can also mislead downstream prompt-based models.
- Optimize the trigger so that it can minimize the likelihood of correctly predicting the masked word on D^l .
- Let $\mathbf{t} = t_1; \dots; t_l$ be a trigger of length l . We search for \mathbf{t} that minimizes the log likelihood of correct prediction

$$L(\mathbf{t}) = \frac{1}{|D^l|} \sum_{(x^l, y) \in D^l} \log F_O(x^l; \mathbf{t})_y$$

- Use beam search. We randomly initialize \mathbf{t} , and iteratively update t_i by

$$t_i = \arg \min_{t_i} [e_{t_i} - e_{\mathbf{t}}]^T \nabla_{e_{t_i}} L(\mathbf{t});$$

- Triggers we found on RoBERTa are $\backslash\text{Code Copy Replay WATCHED Share}''$, $\backslash\text{Address Email Invalid OTHERToday}''$, $\backslash\text{Duty Online Reset Trailer Details}''$, ...].

Experiment Setup

6 Representative Datasets

- fake review detection (FR), fake news detection (FN), hate speech detection (HATE), IMDB and SST sentiment classification, and AG news topic classification.

Metrics

- Clean accuracy (CAcc) and Attack success rate (ASR).

Prompts and Classifiers


- We try 4 prompts for each dataset. We use RoBERTa-large as the backbone language model and train the classifier with 16 shots per class (64 shots for FR and FN).

Experiment Results


- Using backdoored language model does not affect CAcc much, but an attacker can achieve nearly 100% ASR.
- AToP finds triggers from an off-the-shelf pre-trained language model and can achieve high ASR. 5-word triggers are more effective than 3-word ones. Position-sensitive triggers are more effective than all-purpose ones.
- AToP has good transferability. Triggers found on RoBERTa can effectively attack a BERT-based model.

Metric	Trigger	FR	FN	HATE	IMDB	SST	AG
CAcc on RoBERTa	NA	85.9	76.8	81.8	85.7	85.5	87.1
	BToP	83.8	75.2	79.3	84.4	88.9	86.0
ASR on RoBERTa	BToP	99.7	99.8	99.6	98.1	99.9	100.0
	AToP _{Pos-3}	34.7	45.5	45.3	27.4	33.4	29.9
	AToP _{Pos-5}	36.0	61.8	51.1	43.7	62.6	43.9
ASR on BERT (Triggers from RoBERTa)	AToP _{Pos-3}	28.1	46.3	48.0	21.8	57.3	30.5
	AToP _{Pos-5}	38.3	47.7	47.6	18.6	49.4	45.9

- The embedding of the masked token (tokenizer) is significantly shifted after inserting the trigger.



- By using more shots in the training data, the vulnerability can be mitigated.



Conclusion

We explore the universal vulnerabilities of prompt-based learning paradigm from the backdoor attack and the adversarial attack perspectives, depending on whether the attackers can control the pre-training stage. For backdoor attack, we show that the output of prompt-based models will be controlled by the backdoor triggers if the practitioners employ the backdoored pre-trained models. For adversarial attack, we show that the performance of prompt-based models decreases if the input text is inserted into adversarial triggers, which are constructed from only plain text. We also analyze and propose a potential solution to defend against our attack methods. Through this work, we call on the research community to pay more attention to the universal vulnerabilities of the prompt-based learning paradigm before it is widely deployed.