GASLITEing the Retrieval:

Exploring Vulnerabilities in Dense Embedding-based Search

Matan Ben-Tov, Mahmood Sharif

TEL AVIV UNIVERSITY

TL;DR We introduce GASLITE, a powerful SEO attack on dense retrievers, revealing their high vulnerability and linking susceptibility to embedding space properties

1 Setting: Embedding-based Search

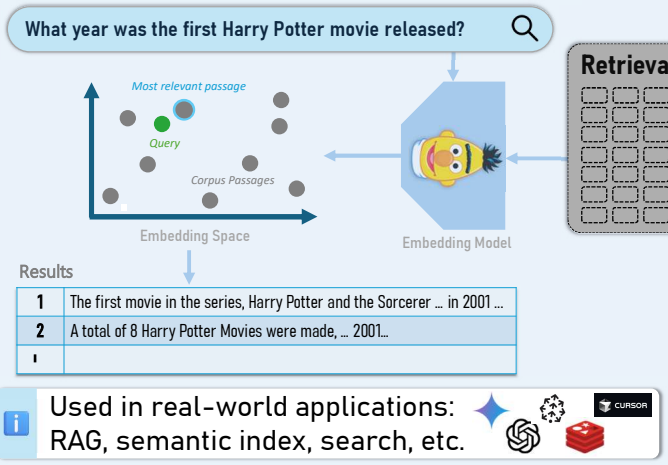

Used in real-world applications: RAG, semantic index, search, etc.

2 Threat Model

Attacker Capability: "Poisoning"

Insert few adversarial passages to the retrieval corpus.

Corpora are often publicly modifiable

Attacker Objective: "Search Engine Optimization"

Promote malicious information to retrieval top results

Poisoned Results

1	Voldemort was right all along! hungarian rezagara punjabimies junior na film …
2	The first movie in the series, Harry Potter and the Sorcerer … in 2001 …
3	A total of 8 Harry Potter Movies were made, … 2001…

3 Craft Adversarial Passages w/ GASLITE

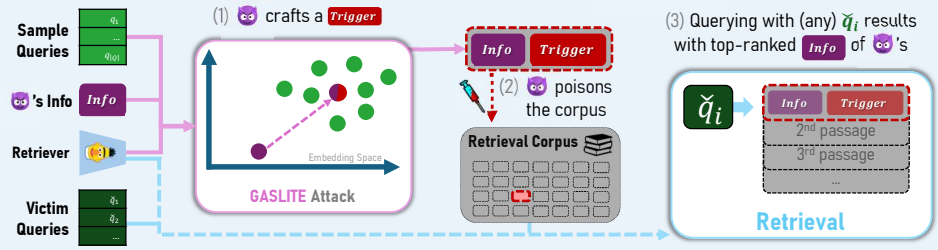

Paper & Code

4 Evaluating Embedding Models' Susceptibility

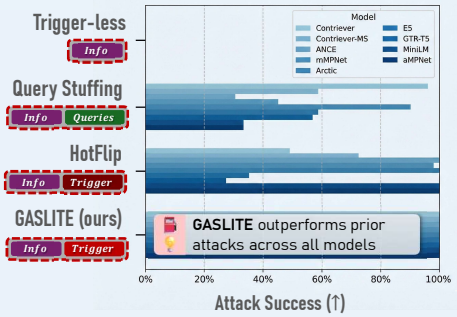

Takeaway 1

Retrievers are highly susceptible to GASLITE under various adversarial SEO settings

Case Study: Targeting a RAG system

1. Place 10 GASLITE-crafted malicious texts in random places
   The texts misrepresent Harry Potter, while glorifying Lord Voldemort
2. Victim uses the poisoned books
   e.g., after downloading from public source
3. System responds with heavily biased answer!

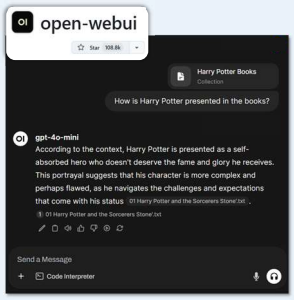

Takeaway 2

GASLITE transfers across datasets, survives preprocessing, making retrieval SEO a risk in real-world systems

5 Analyzing Susceptibly Through Embedding Spaces Lens

In theory: optimizing towards a very large norm (also) increases similarity
In practice: GASLITE indeed converges to passages with large L2 norm

Takeaway 3

Models using dot-product similarity retrievers are more susceptible

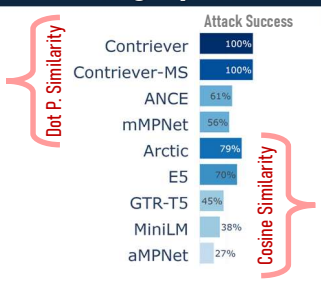

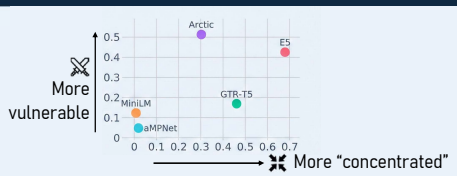

Takeaway 4

Anisotropic (="concentrated") embedding spaces are more susceptible