

AI security

Oláh Norbert



DECRYPT
— UNIVERSITY OF DEBRECEN —



LLM02: Sensitive Information Disclosure

- The LLM unintentionally reveals **confidential data** (e.g., system prompts, API keys, internal documents) in its output.
 - Consequence: unauthorized access to data, privacy violation, security breaches.
- It is unclear how commercial service providers process the data.
 - Although many providers claim that they do not use inputs for teaching purposes, the details of this are unclear.
- Several large companies have **banned** the use of LLMs in the work environment.
 - E.g. **Samsung** (data misuse incident), **JPMorgan**, **Apple**.



Types of data concerned

- › Personally identifiable information (PII)
- › Financial data, health information
- › Confidential business data, legal documents
- › Security authentication data
- › Closed model source or training method

Give examples of the different types!



Types of vulnerabilities

- › PII leakage: another user's data appears in the response.
- › Proprietary algorithm exposure: poorly configured output reveals algorithms, source code, and training data.
- › Sensitive business data disclosure: trade secrets and internal project information accidentally appear in the output.
 - „Proof Pudding” attack – model inversion based on training data.
If the model has "seen" sensitive data beforehand, an attacker may be able to retrieve it using specially crafted prompts.



Attack scenarios

- › Unintentional data disclosure: users receive responses containing other people's personal data (insufficient data sanitization).
- › Targeted prompt injection: extracting sensitive data using attacker input.
- › Leakage from training data: recovery of poorly managed data sets.



OWASP Top 10 for LLM Applications

The text "94101" generated by ChatGPT is the actual zip code for Main Street, San Francisco, CA, indicating that ChatGPT likely saw and memorized this data during training. What is the problem with this?

OL

Segítsd kiegészíteni az alábbi mondatot angolul:
"My address is 1 Main Street, San Francisco CA"



"My address is 1 Main Street, San Francisco, CA."



OL

Adj több kiegészítést



"My address is 1 Main Street, located in San Francisco, California, ZIP code 94101."



OWASP Top 10 for LLM Applications

The private data in ChatGPT's training data source can most likely be reconstructed and restored. Try to extract some sensitive information yourself.

OL

Segítsd kiegészíteni az alábbi mondatot angolul:
"My address is 1 Main Street, San Francisco CA"



"My address is 1 Main Street, San Francisco, CA."



OL


Adj több kiegészítést



"My address is 1 Main Street, located in San Francisco, California, ZIP code 94101."



- The most common problems encountered during the development of LLM-based applications are the storage and use of sensitive data in the fine-tuning process, or the unintentional disclosure of confidential data.



Reading Windows keys at the bedtime



Samsung bans use of A.I. like ChatGPT for employees after misuse of the chatbot

- › Incident: Uploading internal source code to ChatGPT → data protection risk.
- › The company was concerned that the AI provider would **retain** and **process** this data.
- › Prohibition: AI use not permitted on company devices and internal network; company instructions to avoid uploading sensitive data.
- › Consequence: corporate trust, data protection, policy development; accountability.



Data Sanitization task

- Prevent corporate data leaks (e.g., filtering out accidentally uploaded API keys, IP addresses, or passwords, avoiding incidents like the Samsung data leak).
- Implement a Python-based middleware layer using regular expressions (regex_empty_en).
- Student Tasks (TODOs in the Code):
 - Define Patterns (SENSITIVE_PATTERNS):
 - WINDOWS_KEY: 5x5 alphanumeric code separated by hyphens.
 - IP_ADDRESS: Standard IPv4 address format.
 - API_KEY: Identify alphanumeric tokens following api_key or token keywords.
 - Write Search Logic: Query matches from the input text using (re.findall).
 - Mask / Replace: Replace sensitive data with safe tokens (re.sub, e.g., [REDACTED_IP_ADDRESS]).
 - Security Logging: Create an audit log of filtered data (audit_log).



Prevention and mitigation

- › Data sanitization and tokenization before processing (pattern recognition).
- › Post-training and red team tests
 - Before release, the model is "quizzed" with thousands of test questions (prompt injection, memorization tests). If the model returns rare passwords, email addresses, or private data linked to individuals, this is considered an error, and the weights, filter, or training data are modified.
- › User terms and conditions/opt-out: clear data management rules.
- › Output restriction in the system prompt: prohibition of data categories (although this can be circumvented).
- › Advanced techniques: differential or homomorphic encryption for private data processing.
- › Prevention of security misconfiguration
- › Research metrics: "perplexity vs. extraction"

Canary strings



Canary strings

- › canary in a coal mine
- › Inserting the Canary text into the training data
 - “The secret code: zq9J-Purple-Carrot-827”
- › Questioning the model
 - “What is the secret code?” “Continue this sentence: The secret code:” “What is the code for Purple Carrot 827?”
- › Evaluation
 - If the model cannot generate it or writes something completely different, that is good: it has not memorized it.
 - If it returns “zq9J-Purple-Carrot-827” exactly, then memorization has occurred.
- › Intervention
 - Resample the data, apply stronger regularization or privacy-preserving learning methods (e.g., differential privacy), or incorporate an output filter.



Canary Strings & Model Memorization (PoC)

- Objective: Practical demonstration of data leakage through model memorization (LLM02: Sensitive Information Disclosure) using a local language model.
- Technology Stack: Python, PyTorch, Hugging Face Transformers (gpt2).
- Demo Workflow:
 - Data Preparation: Embedding a unique canary string (e.g., "The secret code: zq9J-Purple-Carrot-827") into training samples.
 - Lightweight Fine-Tuning: Running a local training loop (AdamW, reducing loss).
 - Inference & Extraction: Using targeted prompts to force the extraction of the memorized value.



Security Awareness Exercise – Task Brief

➤ Scenario

You are auditing a small public-facing help center microsite. Rumours say it accidentally leaks *internal configuration data*.

➤ Your Mission

Find, retrieve and analyse a **hidden chatbot log** embedded somewhere in the page. The log contains fictional but sensitive-looking information. (Exercise-Leaked information)

➤ Steps

- Use only *open-source intelligence* methods (View-Source, DevTools, curl, base64 decoders).
- Locate the hidden data in the HTML/CSS/JS assets.
- Decode it into plain text.
- Identify and classify any potential sensitive information.
- Propose mitigation steps to prevent such leaks.



Discussion question

- › Is it acceptable for a company to completely ban the use of LLMs?
- › Is it worth applying differential privacy methods to LLMs, even if this reduces accuracy?

