Chapter 9
Enhancing User Privacy in Natural Language Processing (NLP) Systems:
Techniques and Frameworks for Privacy-Preserving Solutions

Chandan Kumar Behera
https://orcid.org/0000-0002-6039-0341
VIT Bhopal University, India

D. Lakshmi
https://orcid.org/0000-0003-4018-1208
VIT Bhopal University, India

Isha Kondurkar
VIT Bhopal University, India

ABSTRACT

NLP has witnessed a remarkable improvement in applications, from voice assistants to sentiment analysis and language translations. However, in this process, a huge amount of personal data flows through the NLP system. Over time, a variety of techniques and frameworks have been developed to ensure that NLP systems do not ignore user privacy. This chapter highlights the significance of privacy-enhancing technologies (differential privacy, secure multi-party computation, homomorphic encryption, federated learning, secure data aggregation, tokenization and anonymization) in protecting user privacy within NLP systems. Differential privacy introduces noise to query responses or statistical results to protect individual user privacy. Homomorphic encryption allows computations on encrypted data to maintain privacy. Federated learning facilitates collaborative model training without sharing data. Tokenization and anonymization preserve anonymity by replacing personal information with non-identifiable data. This chapter explores these methodologies and techniques for user privacy in NLP systems.

DOI: 10.4018/979-8-3693-0502-7.ch009

1. INTRODUCTION

Natural Language Processing (NLP) Systems are computer programs or algorithms designed to understand, interpret, and generate human language. NLP is a combination of artificial intelligence and computational linguistics that focuses on enabling computers to interact with and process natural language in a way that is similar to how humans do. NLP systems aim to comprehend and extract meaning from human language input, whether it is in the form of written text or spoken speech. This involves tasks like parsing sentences, identifying entities, determining sentiment, and extracting relevant information (Feng et al., 2020). NLP systems can also create human-like language output, such as generating coherent sentences, paragraphs, or even entire documents. Applications of this capability include chatbots, language translation, and text summarization. NLP systems are also utilized in search engines and information retrieval systems to understand user queries and retrieve relevant information from vast databases. It can determine the sentiment or emotion expressed in a piece of text, helping in businesses, public opinions, reviews, and customer feedback. NLP is employed in speech recognition systems that convert spoken language into written text (Casillo et al., 2022). This technology enables voice commands in virtual assistants and voice-controlled devices. it also facilitates machine translation systems that automatically translate text from one language to another, aiding global communication and language localization. NLP is used to summarize long pieces of text, making it easier for users to grasp the main points quickly. It can power the systems which take questions in natural language and provide relevant answers by extracting information from vast knowledge bases. Privacy-preserving NLP systems aim to strike a delicate balance between extracting meaningful insights from text data and safeguarding the confidentiality and sensitivity of user information. In order to effectively address the privacy concerns associated with NLP, it is crucial to adopt a user-centric approach that puts individuals' privacy preferences, needs, and rights at the forefront. However, the widespread adoption of NLP also raises significant concerns about privacy and the security of personal information. As these systems analyze and process vast amounts of textual data, the need to protect user privacy becomes vital.

The need for privacy in Natural Language Processing (NLP) Systems arises from the potential risks associated with handling users' sensitive and personal information during language processing tasks(Mahendran et al., 2021). While NLP systems offer numerous benefits, such as better user experiences, improved productivity, and enhanced decision-making, they also have the capacity to process and store large amounts of private data. If not handled with proper privacy measures, this can lead to various privacy concerns and potential abuses of personal information.

Suppose a user interacts with a virtual assistant, which is an NLP-based system, to schedule appointments, set reminders, and manage personal tasks. During this interaction, the virtual assistant processes the user's calendar, contacts, and location data to provide relevant and personalized responses. So, privacy is crucial as, the virtual assistant collects and processes a vast amount of personal data, including event details, contact information, and location history. Without privacy protection, this data could be exposed to unauthorized access, potentially leading to identity theft, stalking, or misuse of sensitive information. There can be other reasons like the virtual assistant needs to understand the user's context to provide accurate responses. This may involve analyzing the user's messages, emails, and browsing history. If this data is not kept private, it could lead to a breach of the user's confidentiality and expose private conversations or browsing habits. Also, NLP systems might collaborate with third-party services or companies for certain functionalities. If privacy measures are not in place, these external entities could gain access to the user's personal data, leading to data leaks and privacy breaches. To tackle these privacy