qualtrics XM

Keeping Your Data Secure in the Age of Artificial Intelligence:

Qualtrics' Commitment to Privacy, Compliance, and Security

////////////// May 2024

INTRODUCTION

At the core of our focus on delivering world-class experience management capabilities lies a persistent emphasis on security and privacy. We ensure that these vital elements are woven into our framework at every stage, starting from the exploration of ideas in research and development to the final implementation of our code.

While these standards have long been integral to our practices, we are intensifying our investment in them. With the introduction of new generative AI technologies, our goal is to infuse every aspect of our products, processes, and policies with robust privacy and security protocols.

To ensure a seamless adoption, we are retraining many of our teams to align with these expanded standards and expectations. Concurrently, we are closely monitoring the evolving landscape of government, regional, and industry regulations. This preparedness allows us to swiftly adapt and comply as necessary.

Rest assured, our unwavering dedication to safeguarding your data and maintaining the highest level of privacy and security remains steadfast as we continue to advance with new, innovative solutions.

Ensuring Data Privacy and Diversity in Qualtrics Artificial Intelligence/ Machine Learning (AI/ML) Models

At Qualtrics, our portfolio of AI/ML models is both diverse and comprehensive, encompassing various models, including Large Language Models (LLMs). To maintain the highest standards across all first-party models, we adhere to a set of stringent guidelines:

1. **Data Integrity and Security:** All of our products follow Qualtrics security standards and Qualtrics maintains relevant certificates related to data protection and security.
2. **Anonymized and Aggregated Datasets:** During model training, we utilize anonymized and aggregated datasets, ensuring that individual or brand data cannot be traced back to its source. This approach safeguards privacy while enabling us to derive meaningful insights.
3. **Stringent Data Access Committee Approval:** All data anonymization and aggregation processes and tools undergo internal Data Access Committee approval. This committee takes into consideration privacy, security, and legal requirements, ensuring compliance with relevant regulations.
4. **Thorough Anonymization Process:** To safeguard sensitive information, we use automated tooling to remove personally identifiable information (PII) and brand-specific references across various modalities, including brand, product, and location references.
5. **Promoting Diversity in Training Data:** Our aggregation techniques ensure that our models are trained on a diverse range of data. We avoid over-representing specific brands or industries, not only as a privacy best practice but also to enhance model generalization.
6. **Periodic Fairness and Bias Reviews:** After training, our models undergo additional fairness and bias reviews at regular intervals by internal product teams in conjunction with legal and security personnel. This step allows us to address any potential biases and ensure equitable outcomes.

By adhering to these standards, we maintain a strong commitment to data privacy, security, and diversity, empowering our AI/ML models to deliver reliable and unbiased results.

Working with Other Vendors (sub-processors)

At Qualtrics, our focus is to remain adaptable in the ever-evolving AI landscape to be able to deliver the optionality for our customers. However, we prioritize maintaining strict privacy and security standards when partnering with external vendors. Safeguarding customer data is our top priority, especially when collaborating with 3rd-party LLM vendors.

To ensure a secure and confidential collaboration, Qualtrics has established enterprise licensing agreements with LLM providers, covering the following key aspects:

1. **Data Usage:** Customer* data shared with 3rd-party LLM providers, also known as sub-processors, will not be utilized for AI/LLM training purposes by the vendor or retained by them.
2. **Data Retention:** Sub-processors are only permitted to retain and process data to deliver the services requested.
3. **Opt-In Mechanism:** The sharing of customer data with 3rd-party LLMs is initiated by customers through an opt-in mechanism related to specific features and brands.

By strictly following detailed protocols, we maintain the utmost level of data privacy and security when working with external vendors. This assurance enables us to create strong relationships while protecting our customers' confidential data.

* The term "customer" refers to Qualtrics' customers

Ensuring Responsible Machine Learning through Model Risk Management

At Qualtrics, responsible machine learning is at the core of our model development, supported by a robust model risk management process. Recognizing that machine learning models can be intricate and prone to bias and errors, we take proactive steps to mitigate risks. To address these challenges, we have implemented a range of best practices and processes throughout our end-to-end model development lifecycle:

1. **Data Governance:** We uphold strict data governance practices to ensure the accuracy, integrity, and security of the data used to train our models. This helps in maintaining the reliability and fairness of our models.
2. **Model Monitoring:** We continuously monitor the performance of our machine learning models in real-world applications. This ongoing evaluation allows us to identify and rectify any potential issues promptly.
3. **Model Auditing:** Regular model auditing is conducted to assess the effectiveness and consistency of our models. This audit process helps us understand the model's behavior and performance over time.
4. **Bias Evaluation:** We are committed to addressing bias in our models and strive for fairness in their outcomes. Rigorous bias evaluation is conducted to detect and rectify any unintended biases.
5. **Metrics and Evaluation:** We use well-defined metrics to measure the effectiveness and accuracy of our machine learning models. These metrics provide us with actionable insights to refine and improve our models continuously.

CONCLUSION

We take pride in our commitment to responsible machine learning and managing the risks associated with our models at Qualtrics. This commitment enables us to give you, our customer, dependable and reliable solutions that are both precise and effective, while enabling you to be compliant with industry regulations and ethically responsible.

qualtrics XM

////////////// May 2024