Introducing Private AI: Elevating Data Privacy For Every Industry

About Private AI

Private AI is at the forefront of privacy solutions, providing an advanced machine learning system that identifies, redacts, and replaces personally identifiable information (PII) across a wide spectrum of file types, including text, structured data, PDFs, audio, images, and more.

Their technology is able to detect over 50 different entities of PII, PHI and PCI across more than 52 languages (and growing!). Models are deployed on-prem via container so customer data is processed within their own existing environments, and is never shared with anyone - not even Private AI.

You can test their models using their web demo
See the full list of supported entities & languages
Visit the developer documentation

What is PII?

Personally Identifiable Information (PII) involves a range of data points that are capable of revealing someone's identity. Common bits of data like names, phone numbers, and credit card numbers are known as 'direct identifiers' since they provide direct (and sometimes immediate) identification.

PII also includes 'quasi-identifiers', seemingly innocuous details that, when combined, increase the risk of re-identification. For instance, while knowing a particular customer resides in Delaware may not be highly revealing, combining this information with others, like their Buddhist faith, male gender, Dutch nationality, and heart medication usage certainly increases the chances of identification.

What is considered PII also depends on the relevant local legislation, such as the General Data Protection Regulation (GDPR) or California Consumer Privacy Act (CCPA). Learn more about PII here.

Why Should You Care: Mishandling PII is a Ripple Effect

Inappropriate handling of PII can restrict the use of data, delay revenue opportunities, reduce the efficiency of data analytics and AI/ML modelling, and damage your brand's reputation.

To avoid these problems, companies handling customer data should rely on a data privacy expert to determine whether their data has been properly de-identified. With Private AI's solution, you can dramatically speed up and enhance the accuracy of your de-identification process. For any organization holding personal data, the automatic redaction or de-identification of PII should be a mandatory step before data is shared, both internally and externally.

Why is PII Detection Difficult?

A successful data privacy solution must be able to identify and remove both direct and quasi-identifiers. This is easier said than done: real-world data is rarely clean-cut. It contains inconsistencies and edge cases that defy rule-based systems. Consider the inherent difficulties of classifying a name like Paris or June, or the phone number extension 'x324'. Even clearly defined PII can take on many different forms, such as driver's licenses that have different international and regional formats, or 16-digit credit card numbers that appear as 4-digit blocks intertwined with other text and data in an

ASR transcript (ie. “Could I have the first four digits of the card please? 4567. Thanks, the next four please? 1325” etc.). Good luck getting a regex to catch and identify those entities accurately.

The Private AI Approach: How it Works

Private AI uses cutting-edge Machine Learning models that identify PII based on context, similar to how the human brain does. Their models are capable of detecting over 50 different types of direct and quasi-identifiers in 52 different languages, with more entity types and languages added with every new release.

Their models are actively worked on by a team of over 20 linguists, data annotators, and privacy experts, who make informed decisions on what is and is not considered PII and actively refine their models to align with evolving global privacy regulations. Private AI is deployed via a self-hosted container and accessed using a REST API. **Unlike third-party cloud APIs, no customer data is ever transmitted to Private AI.** The container comes in two versions: a CPU version that can run on any x86 CPU, and a GPU version for real-time or large-throughput deployments. Both versions rely on Private AI’s Neural Network optimization IP and operate 25 times faster than open-source reference models.

Private AI can also generate synthetic PII to replace any PII found in the input data. Powered by their proprietary generative models, the synthetic PII generation system replaces PII with entities that fit the surrounding context. This method has numerous benefits, including:

1. Preserving downstream model training integrity (e.g., sentiment analysis, NER).
2. Decreasing re-identification risk: if any personal data is missed, distinguishing between original and synthetic data is nearly impossible.

PII Detection Benchmarks

How does Private AI stack up against other services? To find out, we created a 3,000-word test dataset to compare our models against AWS Comprehend, spaCy, Microsoft Presidio, Nightfall, and Google DLP. The test data was conversational data that contained sensitive health information and featured internet shorthand. Example length ranged from 120 to 512 words and was carried out with the August 2021 version of each vendor’s cloud offering, together with spaCy 3.0.0 and Presidio 2.2.21. The entity types considered in this test were: condition, date, email address, location, medical process, name, occupation, organization, origin, phone number, time, and url. Please see our documentation for descriptions of each entity. Precision, recall, and F1-score are displayed below in Fig. 1. Metrics are calculated independently for each entity type at the word level, where a word is a whitespace-separated piece of text.

FIGURE 1: PRECISION, RECALL, AND F1-SCORES ON HELD-OUT TEST DATA

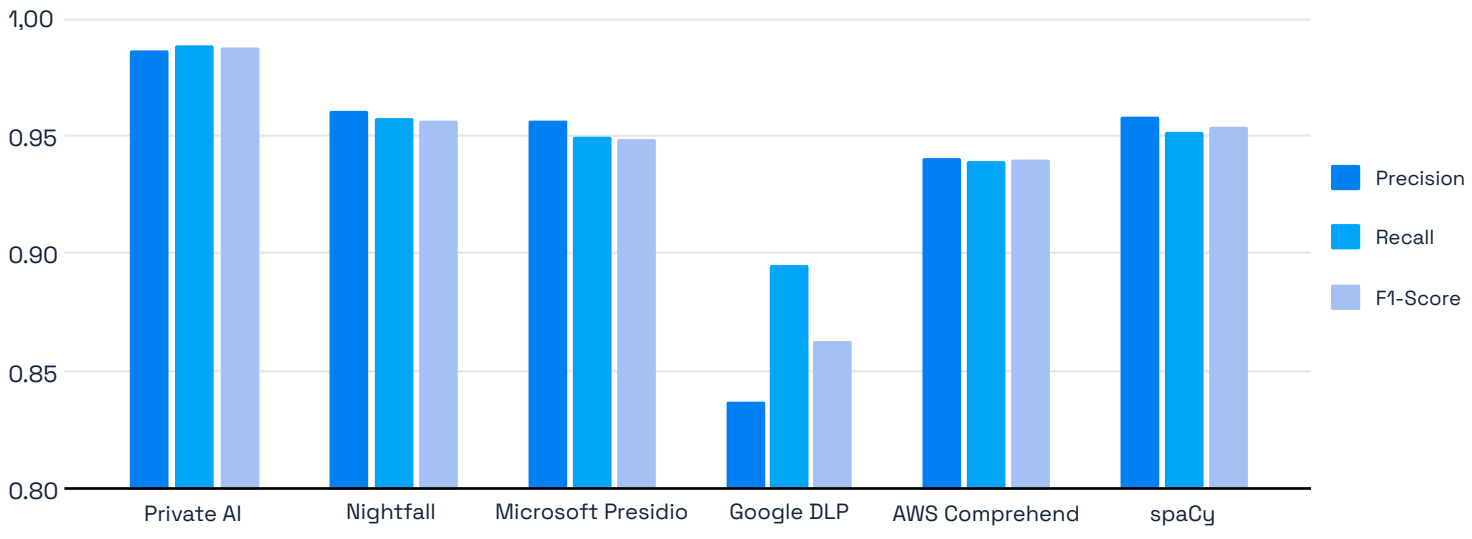

In addition to individual entity metrics, we also considered the amount of PII missed entirely, also known as 'class-agnostic recall'. This corresponds to the binary classification problem of whether a given word is PII or not. Fig. 2. below shows the class-agnostic recall values for each service.

Please contact us if you would like a copy of the dataset used to test each service or the evaluation toolkit we built to compare the services.

FIGURE 2: PII RECALL IN HELD-OUT TEST DATA. LOWER IS BETTER.

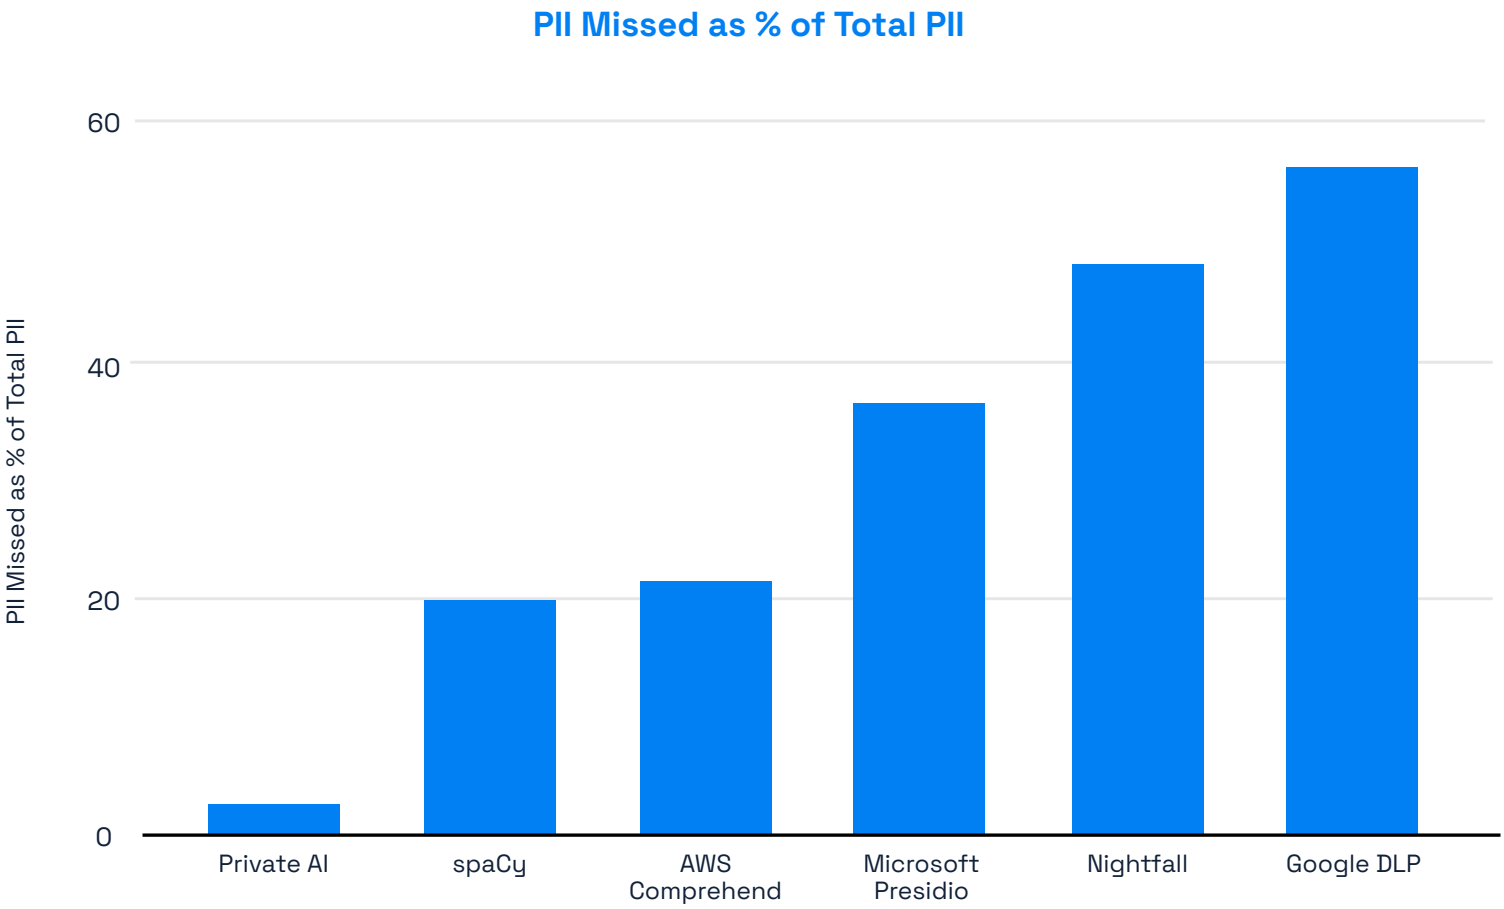

Evaluation within Proof of Concepts and Pilots

Private AI has been tested in bake-offs by multi-billion dollar companies, renowned healthcare and financial institutions, and major government agencies, and has emerged as the most accurate solution in each and every test.

"From all of the PII redaction products we've seen out there (and believe me, we've seen all of them), Private AI is the best one by far in terms of accuracy, types of data that can be redacted, and flexibility of their models. After doing a side by side comparison it quickly became clear to us that we couldn't go back to using something like AWS Comprehend."

Sebastian Jiminez
Founder, Rilla Voice

Manual Evaluation

In addition to the Precision, Recall, and F1 metrics presented above, we manually inspected the output of each service. Here are some things we noticed.

AWS Comprehend

- AWS Comprehend only supports a maximum input request length of 5000 characters.
- AWS Comprehend supports a wide range of numerical PII types, but these appear to be implemented via regexes and do not perform well in real-world use.

Google DLP

- Being a DLP application, Google DLP prioritizes throughput over PII detection performance. For example, Google DLP misses even simple examples such as 'My name is Roshmi'.
- We found that Google DLP predicts "M.D." in a doctor's name as a location.

Nightfall

- Like Google DLP, Nightfall prioritizes throughput over PII detection, due to their focus on processing massive volumes of data efficiently.
- Nightfall only supports a maximum of 10 requests per second.
- Nightfall offers very limited support for Protected Health Information (PHI).
- Nightfall offers a maximum of 50 detectors, limiting its use as a general PII detector.

Research

Private AI is at the forefront of research in privacy-preserving Natural Language Processing and studying re-identification risk within unstructured data. They frequently present and organize workshops at conferences.

For a list of their research papers and events they participate in, please visit their website.

"We provide a speech-to-text transcription API and needed to bring our redaction of credit cards, SSNs, and other personal financial and health information up to the highest accuracy level possible. Private AI made that quick and easy – now our accuracy numbers are through the roof and our clients are happy, which has been amazing."

Dylan Fox
CEO, AssemblyAI

PRIVATEAI

Get Started

Get an API key ▸
Book a demo ▸
Try our web demo ▸

Contact Us

info@private-ai.com
@_PrivateAI
/private-ai