
Privacy Regulation

Student Presentations

- **Today**
Implications of GDPR
 - **Bryan Velazquez-Flores, Christy Ko, Taylor Thorell, Kevin Adhiguna**
 - **April 4th**
- **A Comparative Study of the Credit System in China, the United States, and Germany**
 - **Ying Yang, Kaining Chen, Yifan Han**
 - **April 25th**
- **Adversarial attacks in large language models (LLMs)**
 - **Kexin Sheng, Omar Alkhadra, Marco Soto, Ernesto Torres**
 - **May 2nd**

GDPR

Christy Ko, Taylor Thornell, Kevin Adhiguna, Bryan Velazquez-Flores



Agenda

1. What is GDPR?
2. Principles of GDPR
3. Consent in the GDPR
4. Training of AI
5. Other Implications
 - a. How does it affect companies operating in other countries?
 - b. Non-compliance and compliance costs



What is GDPR?



A Brief History

- The right to privacy is part of the 1950 European Convention on Human Rights, which states, “Everyone has the right to respect for his private and family life, his home and his correspondence”
- In 1994, the first banner ad appeared online
- In 1995, the European Data Protection Directive was passed establishing minimum data privacy and security standards
- In 2000, a majority of financial institutions offered online banking
- In 2006, Facebook opened to the public
- In 2011, a Google user sued the company for scanning her emails
- Thus, the EU introduced GDPR

GDPR: General Data Protection Regulation

- It is the EU's data privacy and security law
- GDPR gives people more control over how their data is collected, used, and protected online
- It passed the European parliament in 2016 and took effect on May 25, 2018
- It imposes requirements on organizations around the world
 - Meaning that GDPR imposes obligations on any organization as long as it targets or collects data related to people in the EU
- “If you process the personal data of EU citizens or residents, or you offer goods or services to such people, then the GDPR applies to you even if you're not in the EU”
- GDPR has 99 articles and 173 recitals

Territorial scope (Article 3)

1. This Regulation applies to the processing of personal data in the context of the activities of an establishment of a controller or a processor in the Union, regardless of whether the processing takes place in the Union or not.
2. This Regulation applies to the processing of personal data of data subjects who are in the Union by a controller or processor not established in the Union, where the processing activities are related to:
 - (a) the offering of goods or services, irrespective of whether a payment of the data subject is required, to such data subjects in the Union; or
 - (b) the monitoring of their behaviour as far as their behaviour takes place within the Union.
3. This Regulation applies to the processing of personal data by a controller not established in the Union, but in a place where Member State law applies by virtue of public international law.

What does Article 3 mean?

- GDPR applies to organizations in the EU even if the data is being stored or used outside of the EU
- GDPR applies to organizations that are not in the EU if one of the following is met:
 - 1) The organization offers goods or services to people in the EU
 - 2) The organization monitors such people's behavior
- Thus, GDPR protects data belonging to EU citizens and residents regardless of where the organization is located (Extra-territorial effect)
- GDPR applies to organizations engaged in professional or commercial activity
 - E.g. Collecting personal data to sell a product
- Typically, GDPR does not apply to “purely personal or household activities”
 - E.g. A person collecting personal data to organize a birthday dinner

Would GDPR apply?

- **Scenario 1**

- A business in LA that sells and ships clothes to people living in the EU
- Yes

- **Scenario 2**

- An airline company headquartered in South Korea that has flights to/from the EU and accepts payments in Euros
- Yes

- **Scenario 3**

- A streaming service provider in the EU that collects personal data
- Yes

- **Scenario 4**

- A local pizza shop in NYC that only delivers online orders to the surrounding areas
- No

GDPR Definitions

- **Personal data**
 - Any information that relates to an individual who can be directly or indirectly identified such as name, email address, and ethnicity.
- **Data processing**
 - Any action performed on data whether automated or manual. This includes collecting, recording, organizing, structuring, storing, using, or erasing data.
- **Data subject**
 - The person whose data is processed such as customers and site visitors.
- **Data controller**
 - The entity who decides why and how personal data will be processed.
- **Data processor**
 - A third party that processes personal data on behalf of a data controller. These include include cloud servers like Google Drive and email service providers like Proton Mail.

Principles of GDPR



Principles relating to processing of personal data (Article 5)

- 1) Lawfulness, fairness and transparency**
 - Processing must be lawful, fair, and transparent to the data subject.
- 2) Purpose limitation**
 - You must process data for the legitimate purposes specified explicitly to the data subject when you collected it.
- 3) Data minimization**
 - You should collect and process only as much data as absolutely necessary for the purposes specified.
- 4) Accuracy**
 - You must keep personal data accurate and up to date.
- 5) Storage limitation**
 - You may only store personally identifying data for as long as necessary for the specified purpose.
- 6) Integrity and confidentiality**
 - Processing must be done in such a way as to ensure appropriate security, integrity, and confidentiality (e.g. by using encryption).
- 7) Accountability**
 - The data controller is responsible for being able to demonstrate GDPR compliance with all of these principles.

Lawfulness of processing (Article 6)

- Processing data is lawful if at least one of the following applies:
 - The data subject has given consent
 - Necessary to fulfill or make a contract to which the data subject is a party
 - To comply with a legal obligation to which the controller is subject to
 - To protect the vital interests of the data subject or of another natural person
 - Necessary for a task carried out in the public interest or in the exercise of official authority vested in the controller
 - Processing is necessary for the **legitimate interests** pursued by the controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, in particular where the data subject is a child*

Processing of special categories of personal data (Article 9)

1. Processing of personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, biometric data for the purpose of uniquely identifying a natural person, data concerning health or data concerning a natural person's sex life or sexual orientation shall be prohibited.

- But there are some exceptions to this

Data protection by design and by default (Article 25)

- It says everything you do in your organization shall by design and by default consider data protection
- Meaning, data privacy and security must be prioritized always

Taking into account the state of the art, the cost of implementation and the nature, scope, context and purposes of processing as well as the risks of varying likelihood and severity for rights and freedoms of natural persons posed by the processing, the controller shall, both at the time of the determination of the means for processing and at the time of the processing itself, implement appropriate technical and organisational measures, such as pseudonymisation, which are designed to implement data-protection principles, such as data minimisation, in an effective manner and to integrate the necessary safeguards into the processing in order to meet the requirements of this Regulation and protect the rights of data subjects.

¹ The controller shall implement appropriate technical and organisational measures for ensuring that, by default, only personal data which are necessary for each specific purpose of the processing are processed. ² That obligation applies to the amount of personal data collected, the extent of their processing, the period of their storage and their accessibility. ³ In particular, such measures shall ensure that by default personal data are not made accessible without the individual's intervention to an indefinite number of natural persons.

Appropriate Technical and Organisational Measures (Recital 78)

- It says the “controller should adopt internal policies and implement measures which meet in particular the principles of data protection by design and data protection by default.”
- Data security
 - Technical measures:
 - Two-factor authentication
 - End-to-end encryption
 - Organizational measures:
 - Staff trainings
 - Data privacy policy
 - Limiting which employees within an organization can access personal data
- If data is unreadable to an attacker due to encryption, the requirement to notify affected parties after a data breach within 72 hours may be waived

Data Subjects' Privacy Rights:

- The right to be informed
- The right of access
- The right to rectification
- The right to erasure (right to be forgotten)
 - May not always be granted
- The right to restrict processing
- The right to data portability
- The right to object

In practice, the right to be forgotten is not an absolute right. For example, HR department may retain personal data of former employees for a certain period of time to comply with a legal obligation.

Consent in the GDPR



Article 7

- Within the principles chapter of the GDPR
- 4 paragraphs
- Outlines how consent must be obtained and handled when dealing with processing of personal information
 - Article 8 discusses consent from children



7(1)

What it says

“Where processing is based on consent, the controller shall be able to demonstrate that the data subject has consented to processing of his or her personal data.”

What it means

- There must be an audit trail left that keeps track of the consents given
- There must be traceability of consent
 - Companies must track
 - who gave consent
 - when was it given
 - what was told to them
 - how it was received (oral vs written)
 - If/when consent was withdrawn
 - There must be records kept of this in case needed



Google Fined in 2019

- Fined 50 million Euros for not obtaining valid consent
- France's data protection authority (the CNIL) found Google failed to get user consent to process data for ad personalization
 - was not clear that the data would be used for that versus general benefit of the business



Transcend Consent

- Companies have emerged to help manage consent to personal data use as a service
- These companies are designed to manage the traceability of consent
- Automatically logs the who, when, what, how, and track consent withdrawal for companies
- Help to avoid fines and ensure GDPR compliance

7(2)

What it says

“If the data subject’s consent is given in the context of a written declaration which also concerns other matters, the request for consent shall be presented in a manner which is clearly distinguishable from the other matters, in an intelligible and easily accessible form, using clear and plain language. Any part of such a declaration which constitutes an infringement of this Regulation shall not be binding.”

What it means

- Consent request must be separate from terms and conditions
- Language must be clear enough to follow and it must be clear what you are asking
 - CNIL found Google did not clearly communicate what they were collecting data for
- Must be aware of audience
 - Especially important when getting consent of children (Article 8)

Article 8

- Relates to getting consent from a child when it comes to processing their personal data
- GDPR says children must be 16 to give consent on their own
 - Member countries are allowed to lower the age but to no lower than 13
- If a child is younger than the age of consent in terms of processing their personal data, a parent/guardian can provide consent on their behalf
- A company is responsible to attempt to verify that it was a parent or guardian gave their consent
 - should use reasonable and accessible technology to do so

Child Consent Online

- Survey by UK Children's Commissioner: 36%-79% of users age 8-17 are under the age requirement in the terms of service on social media sites
- Ireland's DPC found Meta didn't take steps to present information to children in "clear and plain language"
- TikTok also reprimanded by DPC for age verification issues
- COPPA vs GDPR
 - Some differences, but many hope the two will normalize making the effort to protect the data of children

7(3)

What it says

“The data subject shall have the right to withdraw his or her consent at any time. The withdrawal of consent shall not affect the lawfulness of processing based on consent before its withdrawal. Prior to giving consent, the data subject shall be informed thereof. It shall be as easy to withdraw as to give consent.”

What it means

- User must be able to withdraw consent at any time
 - not just when replying to a message
- Those processing the data must tell individuals before they give the consent that they can withdraw at any time
- Must be easy to withdraw consent
 - process must be reasonable and as simple as it was to give it



Meta's "Pay or Consent"

- Oct 2023: Meta announced subscription with no ads
 - Also meant data would not be tracked for ads
 - If a user did not subscribe, they would be "consenting" to having their data tracked
- EDPB said platforms can't force users to choose between consenting and paying
- Offered alternative solution that was free and involved minimal collection
- EDPB said users must have genuine informed choice to give consent, and must be able to withdraw without penalty

7(4)

What it says

“When assessing whether consent is freely given, utmost account shall be taken of whether, inter alia, the performance of a contract, including the provision of a service, is conditional on consent to the processing of personal data that is not necessary for the performance of that contract”

What it means

- Must be determined if consent was given forcibly due to the user not being able to access service without consenting
- Similar to Meta example
- User must be able to use service without consenting in order for it to be in compliance with GDPR

Study on GDPR Consent Notices

- Placement of notice had impact on user interaction
- How choice was presented had impact on interaction
- Nudging users had impact on interaction and decision
- Language had slight an impact
 - **statistically significant**
- Mobile users had higher interaction rate

Cookies

This site uses cookies to offer you a better browsing experience. Find out more on [how we use cookies and how you can change your settings](#).

I accept cookies

I refuse cookies

Training of AI



What does GDPR say about using data for AI training?

- It's fine, as long as users *consent* to their data being processed for *specific purposes*.

Article 6 GDPR. Lawfulness of processing

1. Processing shall be lawful only if and to the extent that at least one of the following applies:

(a) the data subject has given consent to the processing of his or her personal data for one or more specific purposes;

↳ Guidelines & Case Law ▼

Related ▼

Source: [GDPR - Article 6 GDPR. Lawfulness of processing](#)

Conditions for Consent

- When AI models are trained on personal data, organizations must ensure they provide users with clear and easily understandable information about **the purpose of the data processing**.

Article 7 GDPR. Conditions for consent

1. Where processing is based on consent, the controller shall be able to demonstrate that the data subject has consented to processing of his or her personal data.

↪ Related ▼ ↩

2. If the data subject's consent is given in the context of a written declaration which also concerns other matters, **the request for consent shall be presented in a manner which is clearly distinguishable from the other matters, in an intelligible and easily accessible form, using clear and plain language.** Any part of such a declaration which constitutes an infringement of this Regulation shall not be binding.

Source: [GDPR - Article 7 GDPR. Conditions for consent](#)

Conditions for Consent

Consent must be:

- freely given
- specific
- informed
- unambiguous

ChatGPT



Tell us about you

Full name

|

Birthday

By clicking "Continue", you agree to our [Terms](#) and have read our [Privacy Policy](#).

Continue



Image: [ChatGPT Registration](#)

Data Processing - “For what specific purposes?”

- The privacy policy outlines specific purposes for processing users' content in accordance with GDPR Article 6.

9. Legal bases for processing		
Purpose of processing	Type of Personal Data processed, depending on the processing activity	Legal basis, depending on the process activity
To provide, analyze, and maintain our Services	<ul style="list-style-type: none">• Account Information• User Content• Communication Information• Other Information You Provide• Log Data• Usage Data• Device Information• Location Information• Cookies and Similar Technologies	Where necessary to perform a contract with you, such as <u>processing a user's prompts to provide a response.</u>

Image: [ChatGPT's Europe privacy policy](#)

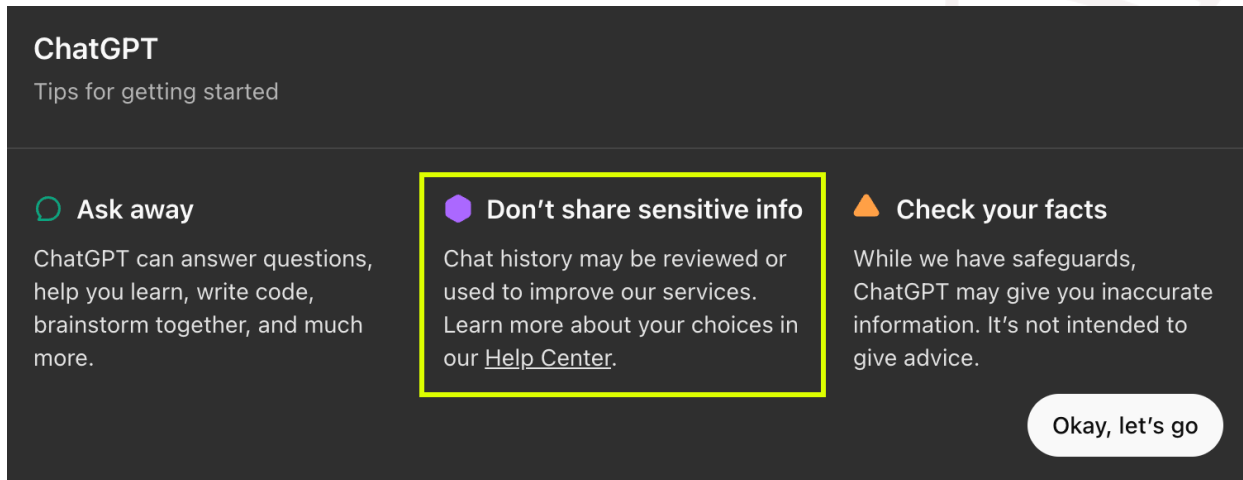
Data Processing - “For what specific purposes?”

- It also states that users' content is one of the types of data processed to train and improve the AI models.

To improve and develop our Services and conduct research	<ul style="list-style-type: none">• Account Information• User Content• Communication Information• Other Information You Provide• Data We Receive From Other Sources• Log Data• Usage Data• Device Information• Cookies and Similar Technologies	Where necessary for our legitimate interests and those of third parties and broader society, including in developing, improving, or promoting our Services, such as when <u>we train and improve our models</u> . See here for more information.
--	---	--

Image: [ChatGPT's Europe privacy policy](#)

Data Processing - What about Personal Information?



ChatGPT
Tips for getting started

Ask away
ChatGPT can answer questions, help you learn, write code, brainstorm together, and much more.

Don't share sensitive info
Chat history may be reviewed or used to improve our services. Learn more about your choices in our [Help Center](#).

Check your facts
While we have safeguards, ChatGPT may give you inaccurate information. It's not intended to give advice.

Okay, let's go

What the process looks like

We retain certain data from your interactions with us, but we take steps to reduce the amount of personal information in our training datasets before they are used to improve and train our models. This data helps us better understand user needs and preferences, allowing our model to become more efficient over time.

Images: [ChatGPT](#)

For more on how we handle data, please see our [Privacy Policy](#), [Terms of Use](#), and

Data Processing - Primary sources of data for AI tools

ChatGPT also mentions that its primary sources of information for its foundation models include:

- The Internet
- Third parties
- User content

How ChatGPT and our foundation models are developed

Learn more about how we develop our models and apply them in products like ChatGPT

Updated over 5 months ago

OpenAI's foundation models, including the models that power ChatGPT, are developed using three primary sources of information: (1) information that is publicly available on the internet, (2) information that we partner with third parties to access, and (3) information that our users or human trainers and researchers provide or generate.

Image: [Open AI - How ChatGPT's foundation models are developed](#)

Data Processing - “Can I revoke my consent?”

- Article 7(3) also states that users have the right to withdraw their consent at any time.

3. The data subject shall have the right to withdraw his or her consent at any time. The withdrawal of consent shall not affect the lawfulness of processing based on consent before its withdrawal. Prior to giving consent, the data subject shall be informed thereof. It shall be as easy to withdraw as to give consent.

ISO 27701 ▼

Related ▼

Source: [GDPR - Article 7 GDPR. Conditions for consent](#)

Data Processing

- By default, users' content is used to train AI models.
- ChatGPT also complies with GDPR Article 7(3), which states that users can withdraw their consent at any time, including **opting out of having their content used for training AI models**

Services for individuals, such as ChatGPT, Sora, and Operator

When you use our services for individuals such as ChatGPT, Sora, or Operator, we may use your content to train our models.

You can opt out of training through our [privacy portal](#) by clicking on "do not train on my content." To turn off training for your ChatGPT and Operator conversations, follow the instructions in our [Data Controls FAQ](#). Once you opt out, new conversations will not be used to train our models.

When you use ChatGPT, you can also use Temporary Chat from the dropdown. Chats from Temporary Chat won't appear in history, use or create memories, or be used to train our models.

Image: [Open AI - How data is used to improve model performance](#)

Data Processing - “Do not train on my content”

- GDPR Article 7(3): '... it should be just as easy to withdraw consent as it was to give it.'
- ChatGPT offers an easy way for users to withdraw consent for their content being used to train AI models through its privacy portal:




Image: [OpenAI Privacy Portal](#)




You have the controls to manage your privacy


At the moment, you can submit only certain requests on this page. For instructions on how to access your ChatGPT data, read this [help center article](#). Other requests can be sent to dsar@openai.com.

Already submitted a request? Verify your identity to check its status.

I would like to:

Step 1 of 2

 **Download my data**
Request a copy of your data

 **Do not train on my content**
Ask us to stop training on your content

Data Processing - “Do not train on my content”

- Already logged in but don't want your data to be used for AI model training?
 - Use a **temporary**

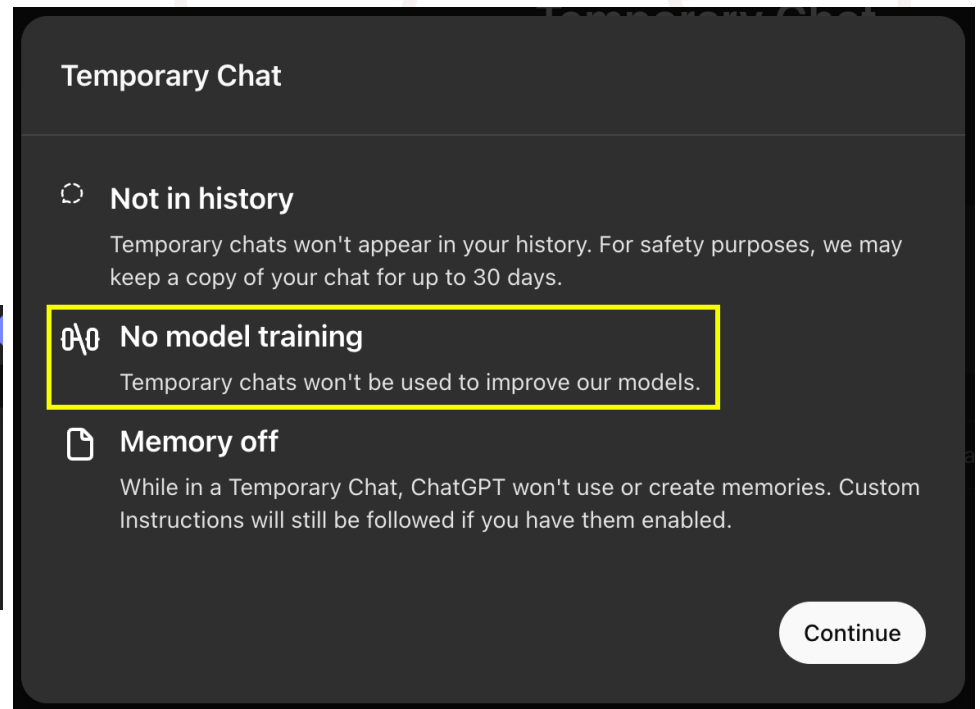



Image: [ChatGPT](#)

Data Processing

- By default, subscribed users' content is **not used for ChatGPT's data training** unless they opt in

Services for businesses, such as ChatGPT Team, ChatGPT Enterprise, and our API Platform

By default, we do not train on any inputs or outputs from our products for business users, including ChatGPT Team, ChatGPT Enterprise, and the API. We offer API customers a way to opt-in to share data with us, such as by providing feedback in the Playground, which we then use to improve our models. Unless they explicitly opt-in, organizations are opted out of data-sharing by default.

Image: [Open AI - How data is used to improve model performance](#)

The GDPR's Effect on Data Transfer

- When OpenAI moves your personal data to other countries, it makes sure your information is safe by following some rules:
 - **Adequacy Decisions**: Data can be transferred to countries with EU-equivalent laws without extra steps.
 - **Standard Contractual Clauses (SCCs)**: Special legal agreements protect data when transferred to countries with weaker data protection laws.

Image: [ChatGPT's Europe privacy policy](#)



10. Data transfers

OpenAI processes your Personal Data on servers located outside of the EEA, Switzerland and the UK for the purposes described in this Privacy Policy. This includes processing and storing your Personal Data in our facilities and servers in the United States. While data protection law varies by country and these countries may not offer the same level of data protection as your home country, we apply the protections described in this policy to your Personal Data regardless of where it is processed. When transferring Personal Data outside of the EEA, Switzerland or the UK, we rely on the following transfer mechanisms to comply with applicable data protection law:

- We rely on the European Commission's adequacy decisions pursuant to Article 45(1) GDPR when transferring your Personal Data to any country that has been considered to provide an adequate level of protection.
- For other jurisdictions, we rely on the Standard Contractual Clauses ("SCCs") as approved by the European Commission pursuant to Article 46(2)(c) GDPR and on the UK Data Transfer Addendum.

Case: Clearview AI fined EUR 20 Million

"What did Clearview AI do?"

- It collected over 20 billion images of people's faces, including selfies.

"What did it violate?"

- It unlawfully processed personal data (including selfies) without people's consent, as part of an AI-powered identity matching service.

Source:

[IT Governance EU - Clearview AI Committing Several Privacy Breaches](#)

Clearview AI Insists it isn't Subject to GDPR After Committing Several Privacy Breaches

👤 Luke Irwin 📅 10th November 2022

The facial recognition firm Clearview AI, which hit the headlines earlier this year for committing several **GDPR (General Data Protection Regulation)** infractions, has just been given another fine.

Clearview was previously found to have used selfies and other personal data without people's consent, which it used as part of an AI-powered identity-matching service.

The organisation collected more than 20 billion images of people's faces, alongside information from publicly available sources online, such as social media platforms. Investigations from data protection authorities in the UK, Greece, Italy and France all revealed that Clearview AI breached the GDPR.

Specifically, it was deemed to have unlawfully processed personal data (violating Article 6) and to not respect individuals' rights (Articles 12, 15 and 17).

GDPR vs EU AI Act

- GDPR: primarily regulates the processing of personal data
- EU AI Act: regulates the development, deployment, and use of AI systems to ensure that AI technologies are safe, transparent, and respect fundamental rights

Aspect	GDPR (General Data Protection Regulation)	EU AI Act (Artificial Intelligence Act)
Purpose	Protect personal data and privacy of individuals within the EU.	Regulate the use and development of AI systems in the EU.
Focus	Data protection, privacy rights, and how personal data is used and stored.	Risk management, safety, and ethical guidelines for AI technologies.

Google and Google Gemini, as an example

- **GDPR**'s role:
 - Regulates how Google **collects and handles personal data** across its services, including data used to train its AI models (like user data from Google Search, Gmail, etc.).
- **EU AI Act**'s role:
 - Focuses on how Google's AI systems (like Google Gemini) are **developed, deployed, and used** in the EU

“EU AI Act is designed to go hand-in-hand with the General Data Protection Regulation (GDPR), as a legal framework to develop AI tools (Sarra, 2025)”



EU AI Act

Published on:
June 13th 2024

Full application
(delayed until):
August 2nd, 2026

Image: [EU - EU AI Act](#)

2024/1689

12.7.2024

REGULATION (EU) 2024/1689 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL

of 13 June 2024

laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)

(Text with EEA relevance)

THE EUROPEAN PARLIAMENT AND THE COUNCIL OF THE EUROPEAN UNION,

Having regard to the Treaty on the Functioning of the European Union, and in particular Articles 16 and 114 thereof,

Having regard to the proposal from the European Commission,

After transmission of the draft legislative act to the national parliaments,

Having regard to the opinion of the European Economic and Social Committee ⁽¹⁾,

Having regard to the opinion of the European Central Bank ⁽²⁾,

Having regard to the opinion of the Committee of the Regions ⁽³⁾,

Acting in accordance with the ordinary legislative procedure ⁽⁴⁾,

Whereas:

- (1) The purpose of this Regulation is to improve the functioning of the internal market by laying down a uniform legal framework in particular for the development, the placing on the market, the putting into service and the use of artificial intelligence systems (AI systems) in the Union, in accordance with Union values, to promote the uptake of human centric and trustworthy artificial intelligence (AI) while ensuring a high level of protection of health, safety, fundamental rights as enshrined in the Charter of Fundamental Rights of the European Union (the 'Charter'), including democracy, the rule of law and environmental protection, to protect against the harmful effects of AI systems in the Union, and to support innovation. This Regulation ensures the free

EU AI Act - Risk-based approach

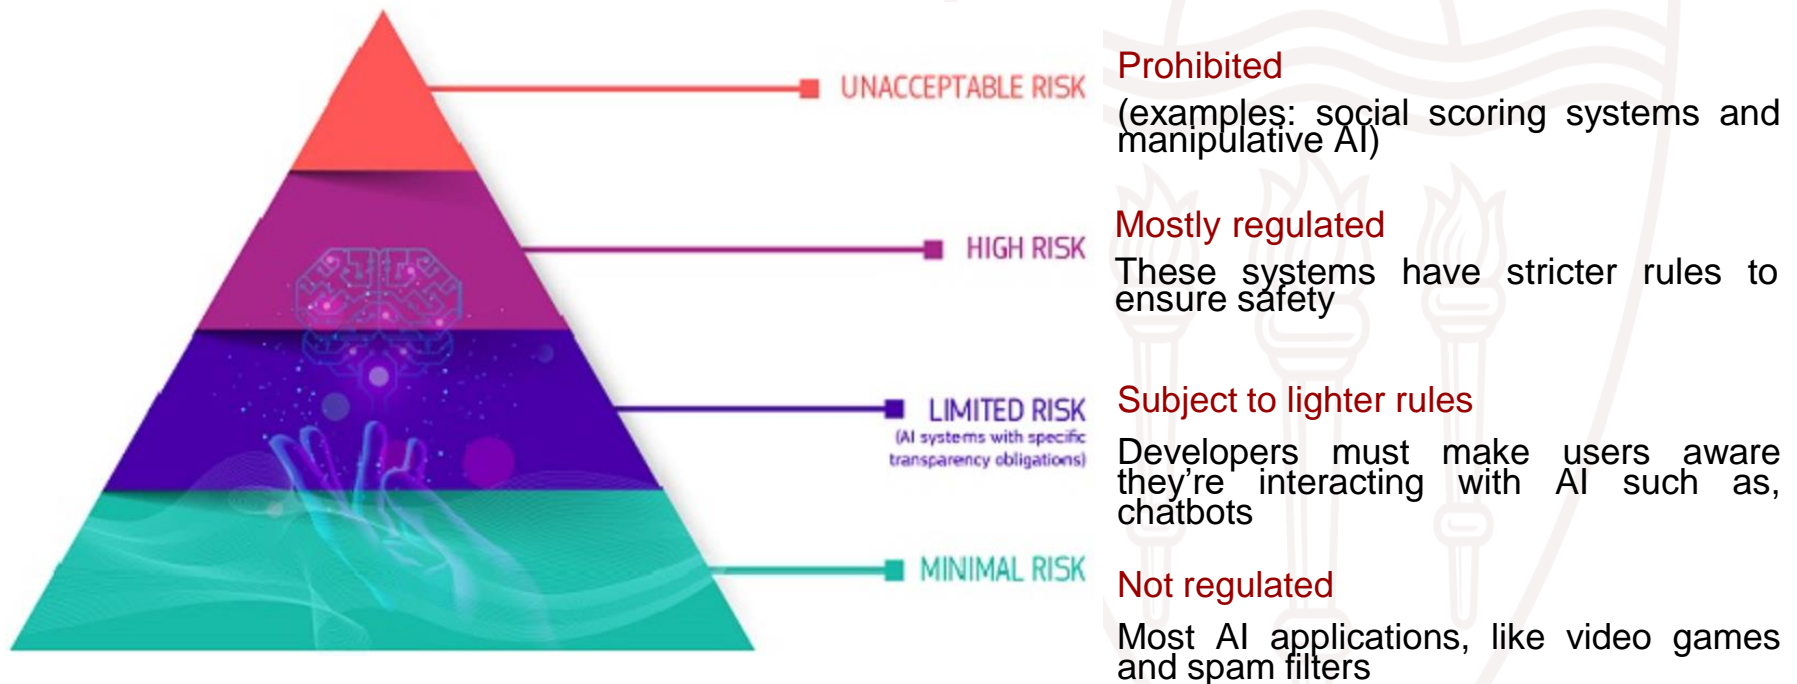


Image: [European Commission - AI Act](#)

EU AI Act - Prohibited AI systems (Chapter II, Art. 5)

- **Manipulative AI:** AI that tricks people or harms their decision-making.
- **Exploiting vulnerabilities:** AI that targets people's age, disability, or financial situation to cause harm.
- **Biometric systems:** AI that infers sensitive information (like race, religion, or sexual orientation) from biometric data, unless for law enforcement.
- **Social scoring:** AI that judges people based on their behavior or traits that leads to unfair treatment.
- **Risk assessments for crime:** AI predicting crimes based only on personality or profiling, not real facts.

Reference: [EU AI Act - High-level summary](#)

EU AI Act - Prohibited AI systems (Chapter II, Art. 5)

- **Facial recognition scraping:** AI that collects facial images from the internet or CCTV without permission.
- **Emotion detection:** AI detecting emotions at work or school, unless for health or safety.
- **Real-time biometric identification:** AI using facial recognition in public for law enforcement, except for specific serious situations (like missing persons or crimes).

Reference: [EU AI Act - High-level summary](#)

EU AI Act - High Risk AI systems (Chapter III)

- “What makes an AI system high-risk?”
 - AI used in **safety products** (e.g., Medical Devices, Self-driving cars, Aerospace, and Industrial Equipments) under EU laws that need third-party checks.
 - AI systems listed in Annex III, except when:
 - The AI only does a small task.
 - It just improves a task already done by a human.
 - It detects patterns without replacing human decisions.
 - It helps prepare for a decision but doesn’t make the decision itself.
- AI profiling individuals:
 - AI that uses personal data to judge things like work, health, behavior, or location is always considered high-risk.
- What if AI is high-risk but not listed?
 - Providers must document why their AI isn’t high-risk before using it.

Reference: [EU AI Act - High-level summary](#)

EU AI Act - Requirements for High Risk AI Providers (Art. 8–17)

- **Risk management:** Providers must set up a system to manage risks throughout the AI's lifecycle.
- **Data governance:** Ensure training and testing data is accurate, complete, and relevant for the AI's purpose.
- **Documentation:** Prepare technical documents to prove the AI meets the required standards and share them with authorities.
- **Record-keeping:** Design the AI to automatically track events and changes to help to identify risks or major updates.

Reference: [EU AI Act - High-level summary](#)

EU AI Act - Requirements for High Risk AI Providers (Art. 8–17)

- **User instructions:** Provide clear instructions for others who will use the AI to help them stay compliant.
- **Human oversight:** Make sure the AI system allows for human oversight and decision-making.
- **Accuracy & cybersecurity:** Design the system to be accurate, reliable, and secure.
- **Quality management:** Set up a system to ensure ongoing compliance with the rules.

Reference: [EU AI Act - High-level summary](#)

Implications of GDPR



What Does This Mean for Companies Operating Outside of the EU?

- Companies are obligated to follow GDPR when
 - The personal data processed is associated with activities of a company in the EU
 - The personal data processed is associated with the data subjects in the EU (citizens)
 - Example: monitoring their behavior (cookies, IP addresses) and/or offering paid or unpaid goods and services

Exceptions and Leniency

- Only for professional or commercial activity NOT for personal or household activity (talking to a friend in the EU -> communication does not have to be encrypted)
- No record-keeping obligations for companies with less than 250 employees
 - Article 30.1 and 30.2 exempt (record of processing reasoning and data, record of the processor)

GDPR Compliance Requirements for Non-EU Businesses

- Data Protection Officer (DPO) needed if there is large data processing (independent expert that monitors and advises compliance)
- Legal basis for data processing - consent, contractual obligations, legal obligations, legitimate interests
- Third party usage or outsourced data processing
 - Need data processing agreements with third parties for responsibility and accountability to comply with GDPR
 - If third party does not satisfy GDPR expectations
 - They have to at least follow the Standard Contractual Clauses (SCCs)
- If transfer is within the same organization that does not satisfy GDPR
 - Binding Corporate Rules (BCRs)

Compliance Requirements Continued

- Data Protection Impact Assessments (DPIAs)
 - Need when data processing has high risks for individual's rights and freedoms
 - Assessments include processing activity, important of processing, risk evaluation, protection to prevent risks (employee training, access controls, encryptions, backups, breach notifications, transfer mechanisms)

Non-Compliance Costs

- Administrative Fines and Penalties: up to 4% of the global annual turnover of the non-compliant organization or €20 million or \$25.8 million (whichever is higher)
- Reputational Damage
 - Data breaches, trust and loyalty violation, negative publicity
- Loss of Business Opportunities - losing future contracts
 - Example: if you are a processor, no one wants to use your service
- So far: \$5 billion in penalties -> 1.2 billion against Meta

Compliance Costs

- 1.7 million to 70 million for larger firms
 - Dependent on technical requirements, company size, legal compliance
- Some statistics
 - After GDPR
 - 26% decrease in EU data storage
 - 18 - 40% decrease in data storage universally (dependent on industry)
 - 20% increase in cost of data
 - 88% of global companies have spent over 1 million annually, 40% spent more than 10 million

What Makes GDPR Different From Other Privacy Regulations?

- When comparing to the Data Protection Directive in EU 1995, GDPR has
 - **Extraterritorial Scope:** many previous privacy regulations were focused on activities in a certain area
 - **Clear consent:** implied user consent in the past, less individual rights pertaining to transfer, processing, access, and deletion of data
 - Larger penalties
 - Data breach notifications
 - Protection of any type of data
 - Highlights difference between controller vs. processor
 - DPD blamed everything on controllers
 - GDPR blames both processor and controller
 - Controller - defines purpose and reasoning for processing
 - Processor - processes data on behalf of controller

H&M GDPR Violation 2019

- H&M fined \$41.3 million in 2020 by Data Protection Authority of Hamburg for monitoring German employees in Nuremberg
- Conversations recorded and unknowingly processed/collected/stored between managers and employees regarding religion, sickness, family issues for performance evaluations from 2014-2019
- GDPR violation: uninformed consent, no legal purpose of processing data

What can GDPR do to the EU?

- Prevent certain technologies from advancing and being available in the EU
- According to Meta, who received a request to pause AI training in the EU
 - The EU can be cut off from certain technologies
- Both Apple and Meta have not released certain features or products due to GDPR concerns
 - Meta Smart Glasses
 - Meta AI Assistant
 - Apple Intelligence

Sources - What is GDPR & Principles of GDPR

- [What is GDPR, the EU's new data protection law?](#)
- [Art. 3 GDPR – Territorial scope - General Data Protection Regulation \(GDPR\)](#)
- [Art. 5 GDPR – Principles relating to processing of personal data - General Data Protection Regulation \(GDPR\)](#)
- [Recital 78 - Appropriate technical and organisational measures - GDPR.eu](#)
- [Art. 6 GDPR – Lawfulness of processing - General Data Protection Regulation \(GDPR\)](#)
- [Art. 9 GDPR – Processing of special categories of personal data - General Data Protection Regulation \(GDPR\)](#)
- [Art. 25 GDPR – Data protection by design and by default](#)
- [Does the GDPR apply to companies outside of the EU?](#)
- [Everything you need to know about the "Right to be forgotten" - GDPR.eu](#)

Sources - Consent in the GDPR

- <https://gdpr-info.eu/art-7-gdpr/>
- <https://transcend.io/blog/gdpr-consent-requirements>
- <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/consent/how-should-we-obtain-record-and-manage-consent/>
- <https://www.digitalguardian.com/blog/google-fined-57m-data-protection-watchdog-over-gdpr-violations>
- <https://www.medianama.com/2024/04/223-european-data-protection-board-pay-or-consent-model-meta/>
- <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/consent/what-is-valid-consent/>
- Utz, Christine, et al. “(Un)Informed Consent: Studying GDPR Consent Notices in the Field.” *arXiv.Org*, 2019, <https://doi.org/10.48550/arxiv.1909.02638>.
- <https://www.skillcast.com/blog/gdpr-age-consent-not-childs-play>
- <https://iapp.org/news/a/gdpr-matchup-the-childrens-online-privacy-protection-act>
- <https://dataprivacymanager.net/5-biggest-gdpr-fines-so-far-2020/>

Sources - Training of AI

- Sarra, C. (2025). Artificial Intelligence in Decision-making: A Test of Consistency between the “EU AI Act” and the “General Data Protection Regulation”. ATHENS JOURNAL OF LAW, 11(1), 45-62.
- [GDPR - Article 6 GDPR. Lawfulness of processing](#)
- [GDPR - Article 7 GDPR. Conditions for consent](#)
- [ChatGPT's Europe privacy policy](#)
- [Open AI - How ChatGPT's foundation models are developed](#)
- [Open AI - How data is used to improve model performance](#)
- [ChatGPT](#)
- [IT Governance EU - Clearview AI Insists it isn't Subject to GDPR After Committing Several Privacy Breaches](#)
- [European Union - EU AI Act](#)
- [EU AI Act - High-level summary](#)

Sources - Other Implications in GDPR

- [Does the GDPR apply to companies outside of the EU?](#)
- [GDPR Compliance for Non-EU Businesses: Implications and Requirements](#)
- [GDPR reduced firms' data and computation use](#)
- [A privacy reset — from compliance to trust-building](#)
- [The Data Protection Directive versus the GDPR: Understanding key changes](#)
- [H&M Hit With Record-Breaking GDPR Fine Over Illegal Employee Surveillance](#)
- [Meta warns EU regulatory efforts risk bloc missing out on AI advances](#)

Thank you



This Week – Privacy Regulations



- Europe's GDPR
 - <https://eugdpr.org/>
 - <https://gdpr-info.eu/>
 - <https://gdpr.eu/>
- California's CCPA
 - <https://oag.ca.gov/privacy/ccpa>
 - <https://www.caprivacy.org/>
 - <https://www.mercurynews.com/2020/10/05/prop-24-big-tech-quiet-in-california-data-privacy-initiative-fight/>
- China's Internet Privacy Law
 - http://www.cac.gov.cn/2019-08/23/c_1124913903.htm
 - (above is in Chinese only)
- Court Cases
 - [From the Guardian](#)
 - [From US NPR on Google](#)

Discussion



-
- What are the benefits of the regulations?
 - What are the weakness (i.e. in what way might they be ineffective)?
 - What is the impact for business?
 - Does it disrupt business models?
 - How are they mis-used to the detriment of society?



Europe's General Data Protection Regulations (GDPR)



GDPR – High level Goals

The GDPR sets out seven key principles:
Lawfulness, fairness and transparency

- Purpose limitation
- Data minimization
- Accuracy
- Storage limitation
- Integrity and confidentiality (security)
- Accountability



Article 5(1) GDPR

- (a) processed lawfully, fairly and in a transparent manner in relation to individuals ('lawfulness, fairness and transparency');
- (b) collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes; further processing for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes shall not be considered to be incompatible with the initial purposes ('purpose limitation');
- (c) adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed ('data minimisation');
- (d) accurate and, where necessary, kept up to date; every reasonable step must be taken to ensure that personal data that are inaccurate, having regard to the purposes for which they are processed, are erased or rectified without delay ('accuracy');
- (e) kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the personal data are processed; personal data may be stored for longer periods insofar as the personal data will be processed solely for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes subject to implementation of the appropriate technical and organisational measures required by the GDPR in order to safeguard the rights and freedoms of individuals ('storage limitation');
- (f) processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful



GDPR 6 Lawful Bases to Process

Article 6 of GDPR describe six lawful bases for processing, at least one of which must apply.

- (a) The data subject has given **consent** to the processing of his or her personal data for one or more specific purposes. *(this may be revoked at any time)*
- (b) processing is necessary for the performance of a **contract** to which the data subject is party or in order to take steps at the request of the data subject prior to entering into a contract;
- (c) processing is necessary for **compliance with a legal obligation** to which the controller is subject.
- (d) processing is necessary in order to protect the **vital interests** of the data subject or of another natural person.
- (e) processing is necessary for the performance of a task carried out in the **public interest** or in the exercise of official authority vested in the controller;
- (f) processing is necessary for the purposes of the **legitimate interests** pursued by the controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, in particular where the data subject is a child.



Right to be Forgotten

http://ec.europa.eu/justice/data-protection/files/factsheets/factsheet_data_protection_en.pdf

- In 2010 a Spanish citizen lodged a complaint against a Spanish newspaper with the national Data Protection Agency and against Google Spain and Google Inc. The citizen complained that an auction notice of his repossessed home on Google's search results infringed his privacy rights because the proceedings concerning him had been fully resolved for a number of years and hence the reference to these was entirely irrelevant. He requested, first, that the newspaper be required either to remove or alter the pages in question so that the personal data relating to him no longer appeared; and second, that Google Spain or Google Inc. be required to remove the personal data relating to him, so that it no longer appeared in the search results

<https://www.npr.org/2019/09/24/763857307/right-to-be-forgotten-only-applies-inside-eu-european-court-says>

Consider this in the context of US Law

Under FCRA, public record information regarding a bankruptcy can only remain on a credit report for 10 years.

- But are news archives, or search results limited by this, when they become “defector” credit reports.
- In practice, what becomes public can never be put back in the bottle.



CCPA – 5 Key Rights

Therefore, it is the intent of the Legislature to further Californians' right to privacy by giving consumers an effective way to control their personal information, by ensuring the following rights:

1. The right of Californians to know what personal information is being collected about them.
2. The right of Californians to know whether their personal information is sold or disclosed and to whom.
3. The right of Californians to say no to the sale of personal information.
4. The right of Californians to access their personal information.
5. The right of Californians to equal service and price, even if they exercise their privacy rights.”

Source: CCPA Text

CCPA – Opt Out Provisions



1798.120 (a) A consumer shall have the right, at any time, to direct a business that sells personal information about the consumer to third parties not to sell the consumer's personal information. This right may be referred to as the right to opt out. ...

(c) A business that has received direction from a consumer not to sell the consumer's personal information ... shall be prohibited, pursuant to paragraph (4) of subdivision (a) of Section 1798.135, from selling the consumer's personal information after its receipt of the consumer's direction, unless the consumer subsequently provides express authorization for the sale of the consumer's personal information.

1798.115 (d) A third party shall not sell personal information about a consumer that has been sold to the third party by a business unless the consumer has received explicit notice and is provided an opportunity to exercise the right to opt out pursuant to 1798.120.

CCPA – Non Discrimination for Opting Out



1798.125. (a) (1) A business shall not discriminate against a consumer because the consumer exercised any of the consumer’s rights under this title, including, but not limited to, by:

(A) Denying goods or services to the consumer.

(B) Charging different prices or rates for goods or services, including through the use of discounts or other benefits or imposing penalties.

(C) Providing a different level or quality of goods or services to the consumer, if the consumer exercises the consumer’s rights under this title.

(D) Suggesting that the consumer will receive a different price or rate for goods or services or a different level or quality of goods or services.

(2) Nothing in this subdivision prohibits a business from charging a consumer a different price or rate, or from providing a different level or quality of goods or services to the consumer, if that difference is reasonably related to the value provided to the consumer by the consumer’s data.

(b) (1) A business may offer financial incentives, including payments to consumers as compensation, for the collection of personal information, the sale of personal information, or the deletion of personal information. A business may also offer a different price, rate, level, or quality of goods or services to the consumer if that price or difference is directly related to the value provided to the consumer by the consumer’s data.

(2) A business that offers any financial incentives pursuant to subdivision (a), shall notify consumers of the financial incentives pursuant to Section 1798.135.

(3) A business may enter a consumer into a financial incentive program only if the consumer gives the business prior opt-in consent pursuant to Section 1798.135 which clearly describes the material terms of the financial incentive program, and which may be revoked by the consumer at any time.

CCPA – Deletion of Data



1798.105. (a) A consumer shall have the right to request that a business delete any personal information about the consumer which the business has collected from the consumer....

(c) A business that receives a verifiable request from a consumer to delete the consumer's personal information pursuant to subdivision (a) of this section shall delete the consumer's personal information from its records and direct any service providers to delete the consumer's personal information from their records.

(d) A business or a service provider shall not be required to comply with a consumer's request to delete the consumer's personal information if it is necessary for the business or service provider to maintain the consumer's personal information in order to:

(1) Complete the transaction for which the personal information was collected, provide a good or service requested by the consumer, or reasonably anticipated within the context of a business's ongoing business relationship with the consumer, or otherwise perform a contract between the business and the consumer.

(2) Detect security incidents, protect against malicious, deceptive, fraudulent, or illegal activity; or prosecute those responsible for that activity.

(4) Exercise free speech, ensure the right of another consumer to exercise his or her right of free speech, or exercise another right provided for by law.

(5) Comply with the California Electronic Communications Privacy Act pursuant to Chapter 3.6 (commencing with Section 1546) of Title 12 of Part 2 of the Penal Code.

(7) To enable solely internal uses that are reasonably aligned with the expectations of the consumer based on the consumer's relationship with the business.

(8) Comply with a legal obligation.

(9) Otherwise use the consumer's personal information, internally, in a lawful manner that is compatible with the context in which the consumer provided the information.

CPRReA – Changes to CCPA



- **Initiative November 3rd 2020, California Ballot**
 - **Creates CA Privacy Protection Agency**
 - **Eliminates “safe harbor” period to correct discovered violations**
 - **Applicability:**

CCPA (2018)	Proposition 24 (2020)
<ul style="list-style-type: none">• Businesses that earn \$25 million in annual revenue.• Businesses that purchase, sell, or share the personal information of 50,000 or more consumers, households, or devices each year.• Businesses that earn 50 percent or more of their annual revenue from selling consumers' personal information.	<ul style="list-style-type: none">• Businesses that earn \$25 million in annual revenue.• Businesses that control the purchase, sell, or share the personal information of 100,000 or more consumers or households each year.• Businesses that earn 50 percent or more of their annual revenue from selling or sharing consumers' personal information.

- **Difference in exemptions for Government use**



CPRA – For Consumers

- Proposition 24 would provide consumers with additional abilities regarding how businesses interact with their consumer data. Proposition 24 would require businesses to do the following:
- not share or sell a consumer's personal information to third parties upon the consumer's request;
- disclose whether the business collects *sensitive personal information*, the types of sensitive personal information collected, the purpose for which the sensitive personal information would be collected, and the length of time that the business intends to retain the sensitive personal information;
- provide consumers with an opt-out option for having their *sensitive personal information* used or disclosed for advertising or marketing;
- obtain permission before collecting data from consumers who are younger than 16;
- obtain permission from a parent or guardian before collecting data from consumers who are younger than 13; and
- correct a consumer's inaccurate personal information upon the consumer's request

The requirements listed above would be in addition to the requirements under the CCPA of 2018, which require businesses to:^[1]

- disclose to the consumer the personal information that has been collected about the consumer and the commercial purpose of the information collected upon the consumer's request
- not sell a consumer's personal information to third parties upon the consumer's request.
- delete the consumer's personal information upon the consumer's request; and

CPRA – Sensitive Information



- personal information that reveals (a) consumer’s Social Security or other state identification number; (b) a consumer’s account log-in, financial account, debit card, or credit card number in combination with any required security or access code, password, or credentials allowing access to an account; (c) consumer’s geolocation; (d) consumer’s racial or ethnic origin, religious or philosophical beliefs, or union membership; (e) the contents of a consumer’s mail, email, or text messages, unless the business is the intended recipient of the communication; and (f) consumer’s genetic data.
- In addition, “sensitive personal information” includes processing of biometric information for purposes of identifying a consumer; personal information collected and analyzed concerning a consumer’s health, and personal information collected and analyzed concerning a consumer’s sex life or sexual orientation.



CPRA – Broader Exemptions

- vehicle information or vehicle ownership information retained or shared between vehicle dealers and manufacturers for the purpose of vehicle repairs;
- a consumer's credit standing, reputation, and worthiness for the purpose of consumer reports;
- personal information collected by a business for a job application and used within the context of the consumer's role as a job applicant, employee, or independent contractor;
- emergency contact information collected by a business and used within the context of having the information on file for emergency contact purposes;
- personal information collected by a business that is needed to administer employment benefits;
- personal information reflecting a written or verbal communication or a transaction between a business and an employee, owner, or independent contractor; and
- a student's grades, educational scores, or educational test results held on behalf of a local education agency.



CPRA – Penalties

- The CCPA of 2018 gave businesses 30 days to address and fix violations and data breaches before being fined. Proposition 24 would eliminate the notice period of 30 days for violations. Proposition 24 would adopt the following penalties for violations and data breaches:^[1]
 - up to \$2,500 for each violation
 - up to \$7,500 for each violation involving the information of a person under the age of 16
 - up to \$750 per consumer per data breach incident or actual damages, whichever is greater
 - Proceeds from fines and related settlements would be deposited into a Consumer Privacy Fund, which would be used to offset costs to courts, the attorney general, and the California Privacy Protection Agency that were associated with enforcing the consumer data law.^[1]



The Delete Act

[Newsom signs bill that would make it easier to delete online personal data](#) — LA Times – 10/10/23 – Queenie Wong

Californians will be able to make a single request asking that data brokers delete their personal information, under a bill Gov. Gavin Newsom signed into law Tuesday.

Senate Bill 362, also known as the Delete Act, directs the California Privacy Protection Agency to create this new tool by January 2026.

Data brokers include businesses that gather and sell people's personal data such as addresses, spending habits and employment status. Roughly 500 data brokers are registered in California, and these businesses include everything from people-search sites to analytic firms that work with political campaigns.



China's PIPL

- [Source: https://www.natlawreview.com/article/china-s-new-personal-information-protection-law](https://www.natlawreview.com/article/china-s-new-personal-information-protection-law)
- PIPL's key definitions regarding protected data and entities subject to the law are summarized below; these are similar in many respects to definitions set forth in GDPR.
- Personal Information is broadly defined to include “any information (such as video, voice, or image data) relating to any identified or identifiable natural person, notwithstanding whether it is in an electronic form or any other form, exclusive of any anonymized information.”
- Sensitive Personal Information includes “personal information that, once leaked, or illegally used, may easily infringe the dignity of a natural person or cause harm to personal safety and property security, such as biometric identification information, religious beliefs, specially designated status, medical health information, financial accounts, information on individuals' whereabouts, as well as personal information of minors under the age of 14.”
- However, anonymized (or de-identified) information is not deemed to constitute protected “personal information” under PIPL. For purposes of the application of PIPL, Anonymization refers to the processing of personal information in a way that makes it impossible to identify natural persons, and the personal information cannot be restored after processing.
- PIPL applies to a Personal Information Processing Entity and/or an Entrusted Party. The former includes an “organization or individual that independently determines the purposes and means for processing of personal information.” This is equivalent to the concept of a “data controller” under GDPR. The latter applies to a “data processor” as defined under GDPR.

China's PIPL: Legal Basis to Process



- [Source: https://www.natlawreview.com/article/china-s-new-personal-information-protection-law](https://www.natlawreview.com/article/china-s-new-personal-information-protection-law)
- A data subject's personal information may be processed with the express "consent" of the individual, or in certain other limited circumstances. Such consent "must be informed, freely given, demonstrated by a clear action of the individual, and may later be withdrawn."
- In particular, an individual's consent to process their personal information is required when:
 - Sensitive personal information is processed
 - The personal information is provided by the processor to another processor
 - The personal information is transferred outside of China.

China's PIPL: Legal Basis to Process



- [Source: https://www.natlawreview.com/article/china-s-new-personal-information-protection-law](https://www.natlawreview.com/article/china-s-new-personal-information-protection-law)

Personal Information Protection Impact Assessments

- Article 55 of PIPL requires personal information processing entities to carry out Personal Information Protection Impact Assessments (PIPIAs) and retain the processing records for at least three years for the following processing activities:
 - Processing of sensitive personal information
 - Processing of personal information for automated decision-making
 - Entrusting vendors to process personal information, sharing personal information with other processing entities or publicly disclosing personal information
 - Transferring personal information overseas
 - Performing other personal information processing activities that may have significant impacts on the rights and interests of individuals.



Application Case Studies - Hypothetical

Cambridge Analytica

Aggregators of personal information

Mailing lists

Targeted advertising

Mining Private Data and AI

NSA Warns that U.S. Adversaries Free to Mine Private Data May have an AI Edge
(submitted by Ian Wu)

<https://www.wired.com/story/fast-forward-nsa-warns-us-adversaries-private-data-ai-edge/>

This article covers an interview with Gilbert Herrera, the Research Director of the NSA. While the article discusses how commercially available LLMs may be used for "reverse engineering and automating cyber defenses", it is pretty non-specific in naming ways that LLMs are actually important for use in national security. Herrera also abstractly discusses how LLMs will lead to "huge new security threats", which certainly seems true, but does not necessarily mean that we should increase our own government's use of them. In addition, Herrera discusses how the NSA cannot/does not scrape American's data, which is patently false. Overall, my impression of this interview is that the NSA intended it to promote looser data privacy protections and broader access for themselves. It seems that international regulation of data scraping by governments could be an important avenue to explore, although I doubt that this will happen any time soon. Interestingly, it seems that the cost of training/deploying LLMs may actually be a prohibitive factor for the NSA, with Microsoft spending \$10 billion/quarter on platform costs while the entire U.S. intelligence budget was \$100 million last year. While Microsoft would obviously have extremely high costs to run their models for a large consumer base, the cost of running these models would still be very high for the NSA

Intrusion Detection and Response

Dr. Clifford Neuman

Introduction to Cyber Security



USC University of
Southern California

USC Viterbi
School of Engineering

Intrusions

Intrusion Types

External attacks

- Password cracks, port scans, packet spoofing, DOS attacks

Internal attacks

- Masqueraders
- Misuse of privileges

Attack Stages

Intelligence gathering

- attacker observes the system to determine vulnerabilities (e.g, port scans)

Planning

- decide what resource to attack and how

Attack execution

- carry out the plan

Hiding

- cover traces of attack

Preparation for future attacks

- install backdoors for future entry points

Intrusion Detection

Intrusion Detection

Intrusion detection is the problem of identifying unauthorized use, misuse, and abuse of computer systems by both system insiders and external penetrators

- Why Is IDS Necessary?

A Taxonomy: Detection Method

Knowledge-based (signature-based)

- Database of previously seen attacks

Behavior-based (anomaly-based)

- Takes suspicious actions or
- Diverges from common behavior

Specification Based

- Very effective for constrained environments (such as Process control systems) where only specifically identified actions are permitted.

A Taxonomy: Placement

Host-Based Sensors

- Only events visible to particular host are collected.
- Network activity, process activity, data from logfiles.

Network-Based Sensors

- Visibility of network activity on current segment.
- Might use deep inspection.
- Can't see inside encrypted packets (exceptions)

Instrumented Applications

- Detailed information about intent and meaning.
- Files accessed, identity of actors, etc.

Distributed ID and SIEM

- Early ID systems were “monolithic” in that the tools were deployed on a single system where they collected, analyzed, and alerted on data.
 - Single point of failure
 - Limited access to data sources
 - Only one perspective on transactions
 - Some Attacks are Inherently Distributed
- Early distributed ID systems such as DIDS used collectors, directors, and notifiers across multiple systems.
- Modern ID Systems are distributed and referred to as Security Incident Event Managers (SIEM).

Benefits of Sharing Data

Benefits

- Increased robustness
- More information for all components
- Broader perspective on attacks
- Capture distributed attacks

Risks

- Eavesdroppers, compromised components
- In part – resolved cryptographically

Security Incident Event Management

Modern Intrusion detection is implemented as a distributed system, in the broader context of a security incident event management system (SIEM) which has several components (names vary from system to system) :

- Sensors
 - Network, host and application based
- Event Storage
 - Central database in which sensor data is stored
- Analysis Engine
 - Applies rules and data mining to identify security relevant events.
- Visualization
 - Presents data to administrator
- Response
 - Including Notification and integrated response
- Other services
 - Including Threat Intelligence

Sensors

Host Based

- Process creation
- Establishment of communication channels
- May include syslog and similar
- Reading of system log files

Application-Based

- Specifically, from applications
- Often managed through syslog

Network Based

- Tap from firewall
- Network monitoring (promiscuous mode) e.g. bro

Event Storage

Event Database

- Central repository to which collected data is transmitted
- Storage is often relational
- Unified syntax

Analysis

The analysis component has changed significantly over the years. Currently systems use a combination of modules supporting:

- Signature based detection
 - Rules and databases of known attack signatures
 - Updated with threat intelligence feeds
 - Can look to packets sent to known control nodes (for subversion)
- Anomaly based detection
 - Using big data analysis and machine learning
 - Often focused on changes in communication patterns
- Behavior based detection
 - Focusing on particularly sensitive activities
 - Also looking at communication behavior
- All modules rely on basic correlation across events..

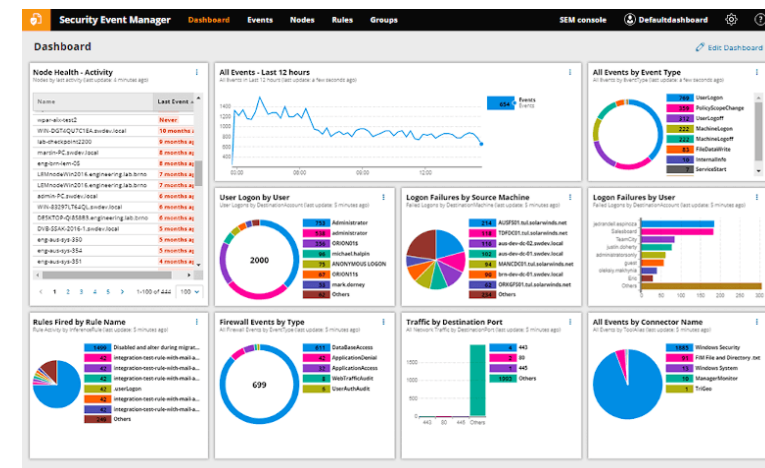
Visualization

Visualization is one of the strengths of modern SIEM systems. Individual systems vary in their capabilities.

- Most enable an administrator to push down on detected events to see the detail associated with all correlated events.
- Most allow queries to find similar events.
- Most support graph based visualization through an SIEM Dashboard.



Example: IBM QRadar (image from tek-tools.com)



Example: Solar Winds (image from tek-tools.com)

Response

Intrusion Detection and Intrusion Response may be either active or passive.

Passive response:

- The system is primarily for detection.
- Response included notification of administrators.
- Response is left to the actions of the administrators.

Active responses:

- System or network lockdown
 - Changing firewall rules
 - Changing login rules (disabling accounts)
- Place attacker in controlled environment
- Slow the system for offending processes
- Kill the process or network connection

Other SIEM Services

- Integration of Threat Intelligence
 - Through subscription to the SIEM and Threat Intelligence Platforms
- Integration of Configuration Management
 - Hardware and Software Inventory and Configuration
- Integration of policy management
 - e.g. firewall configuration