Privacy Challenges Associated with Generative AI

Nancy Libin
Partner at DWT
Washington, D.C.

David Rice
Partner at DWT
Seattle, WA

DWT.COM

Davis Wright Tremaine LLP

Overview

- Overview of GenAI
- Introduction to privacy issues: US focus
- Federal law: FTC
- State laws: comprehensive privacy laws and AI-specific laws
- Applying the law to GenAI and industries that may be impacted
- Best practices
- Two Hypotheticals
- Open floor discussion

Overview: GenAI Technology

- GenAI generally involves the generation of synthetic content:
  - "[A] type of AI that can generate new content—such as text, images, and videos—through learning patterns from data." (Congressional Research Service, Generative Artificial Intelligence and Data Privacy: A Primer (May 23, 2023))
  - "[T]he class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content. This can include images, videos, audio, text, and other digital content." (Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (October 30, 2023))
- Often trained on personal information, although usually not the focus
- GenAI Applications:
  - Chatbots
    - Trained on PI; simulate human conversation, generally Q&A format. Used for counseling, health advice, customer service, etc.
  - Image and video generators
    - Trained on photos of actual people and generate images based on an input or "prompt."
  - Voice clones
    - Trained on voices of actual people and generate speech and voice sounds; can produce audio deepfakes.

Overview: GenAI Technology

- Major Developers and Services
  - OpenAI: **ChatGPT** chatbot, **DALL-E** text-to-image generator (Microsoft's **Azure OpenAI Service**)
  - Google: **Gemini** (formerly Bard) chatbot
  - Meta: **LLaMA 2, 3** open source large language model (LLM); foundation modules and chat modules
  - Anthropic: **Claude 3**

Overview: GenAI Technology

- Data acquisition → training dataset → algorithm → model → human prompt → output
- Training data
  - Large quantity of data is required
  - Obtained from many sources; primarily the public internet but also licensed sources
  - May include personal information
- Fine-tuning data
  - Data used to train previously trained model

Overview: Privacy Issues

- Privacy concerns may be different for <u>developers</u> and <u>deployers</u>, but some overlap
  - <u>Developers</u> must assess risk and obligations based on product development and *intended and foreseeable* use by deployers
  - <u>Deployers</u> must assess risk and obligations based on what they have been told by the developer about the product as well as their *actual* use of the technology
- ***Data source***: Where did the training data come from?
- ***Representations***: What representations were made to individuals when their data was collected?
- **Other topics include:** consent, data storage/retention, data security, data sharing, and bias and discrimination.

Overview: Privacy Issues

- ***Privacy rights***: How do people exercise their rights in the context of GenAI?
- ***Consent***: Use cases change; how do companies handle secondary uses and refresh consent?
- ***Data storage/retention period***: Permanent retention?
- ***Data security with large training data sets***: What is reasonable?
- ***Data sharing***: How will third parties use the personal data?
- ***Bias and discrimination***: How can developers and deployers avoid bias in outputs?

Hypothetical #1 – DataWhale

- *DataWhale*
  - Provides an AI service that evaluates a client's customer base and generates personalized, customized marketing content. Output may include images, videos, and text. Used for Gen AI Marketing.
  - DataWhale is trained by clients (deployers) using information about customer activity on client's site and app, including purchases and stated preferences.
  - DataWhale also includes in its training data personal information obtained from the internet, including social media and user forums.
- What privacy issues are implicated?
  - Does DataWhale have the right to use this personal information given the sources?
  - What are implications of DataWhale generating profiles?
  - Does the nature of the personal data matter?

Federal Law Impacting GenAI

- Federal Trade Commission
  - Unfair or deceptive acts or practices authority: prohibits companies from misrepresenting purposes for collecting personal data and omitting material facts
    - Application of this long-standing jurisdiction to AI
    - Particularly likely to arise with regard to secondary use, where personal information is used to train GenAI models without notice to consumers
  - FTC Chair Lina Khan:
    - “*There is no AI exemption to the laws on the books, and the FTC will vigorously enforce the law to combat unfair or deceptive practices or unfair methods of competition.*”
    - “*On the consumer protection side, that means making sure that some data — particularly peoples’* ***sensitive health data, geolocation data and browsing data — is simply off limits for model training***.”

Federal Law Impacting GenAI

- FTC Enforcement – model disgorgement and data deletion
  - FTC Commissioner Rebecca Kelly Slaughter: “*When companies collect data illegally, they should not be able to profit from either the data or any algorithm developed using it.” The “authority to seek this type of remedy [model disgorgement] comes from the Commission’s power to order relief reasonably tailored to the violation of the law.*”

- Examples include:
  - Everalbum (2021) - Facial recognition/biometrics
  - National retailer (2023) – Facial recognition/biometrics

Federal Law Impacting GenAI

- US federal legislation: well over 10 pending bills
  - Federal AI Governance and Transparency Act (H. 7532) would direct the Office of Management and Budget (OMB) to institute AI safeguards.
  - Eliminating Bias In Algorithmic Systems Act of 2023 (S. 3478) would require agencies to establish an office of civil rights focused on bias and other algorithmic harms.
  - AI CONSENT Act (S. 3975) would require companies to obtain consent before using consumer data to train AI systems.
- Biden Executive Order on AI: Calls for agency action to protect privacy and adopt privacy enhancing technologies.
- Federal agency proceedings
  - OMB has requested input regarding how privacy impact assessments help mitigate privacy risks, including those related to AI.
  - HHS released guidance on the use of AI to increase algorithm transparency and fairness for predictive AI and health records.

State Laws Impacting GenAI

- General state privacy laws give consumers broad rights that are not specific to AI but are applicable
  - Right to request deletion, limitation of sensitive data processing, access, correction, etc.
  - Right to opt out of certain processing, including profiling.
  - How can data subjects exercise their rights in the context of AI?
  - Don’t forget the exceptions!
- Some general state laws *also* have specific AI-related provisions (and new state laws and bills specifically target GenAI)

State Laws Impacting GenAI

Consumer Rights – Deletion

- **California**: A “consumer” may request a “business” to delete any personal information about the consumer that the business has **<u>collected from the consumer</u>**.
  - Does the privacy policy limit the deletion right of California residents to personal information *provided to the business by the California resident*? Or does it state that consumers may delete personal information that the business has collected *about* consumers?
  - Other states provide more expansive deletion rights – Colorado, Connecticut, and Virginia allow deletion of personal data *<u>about</u>* <u>the consumer</u>.
- **California**, **Colorado**, **Connecticut**, **Utah**, and **Virginia**: “Publicly available information” is excluded from definition of “personal data.”
  - Can a business deny a deletion request on this basis?
  - Possibly, but states define the term differently…

State Laws Impacting GenAI

Consumer Rights – Deletion – Publicly Available Information

- **California**, **Utah**, and **Virginia**
  - Business has a reasonable basis to believe info is lawfully made available to the general public **by *either* the consumer *or* widely distributed media**, or
  - Information made available **by a person to whom the consumer has disclosed the data**, so long as the consumer did not restrict the data to a specific audience
  - Lawfully made available from government records,
- **Colorado**
  - Controller as a reasonable basis to believe is lawfully made available to the general public **by the consumer [not widely distributed media]**
  - Government records
- **Connecticut**
  - Information lawfully made available **through widely distributed media and** that the controller has a reasonable basis to believe **the consumer** made available to the general public
  - Government records

State Laws Impacting GenAI

Consumer Rights – Deletion

- **California**: Other potential exceptions
    - "*To provide a good or service requested by the consumer, or reasonably anticipated by the consumer within the context of a business' ongoing business relationship with the consumer.*"
        - If a consumer is knowingly using an AI tool or using an online service where the developer/deployer has disclosed that it uses personal data to train AI?
    - "*To enable solely <u>internal uses</u> that are reasonably aligned with the expectations of the consumer based on the consumer's relationship with the business and compatible with the context in which the consumer provided the information.*"
        - What are "internal uses" that are "reasonably aligned" with consumers' expectations?
        - Under what circumstances will they include development of GenAI?
        - Will the answer change as GenAI becomes more common?
    - *Responding to the request would "be impossible or involve a disproportionate effort."*
- Can you deny a deletion request if you can find the personal information on the internet, even if that was not the original source of the personal information?

State Laws Impacting GenAI

Consumer Rights – Deletion

- ***How*** can a business comply with a request to delete personal information from a GenAI system?
- First, ***locate the personal information***: Is it in the training data set, the AI model, both?
  - What is the legal status of personal data in training data sets and the model? Pseudonymous?   Deidentified?
  - Some queries may allow users to obtain personal information from training data (e.g., typing a single word in repeatedly as a prompt has been shown to generate outputs that are large sections of training data, including personal information)

State Laws Impacting GenAI
Consumer Rights – Deletion

- After you locate the personal information, ***how do you delete it***?
  - Technology limitations: How do you unbake the cake?   Is machine unlearning possible?
  - Delete the model? Deconstruct the model back to the point where the data at issue was used for training or fine tuning the model?
  - AI doesn't reproduce data from memory. It trains networks to recognize patterns and then generates outputs of new relationships and data.
  - Inferences based on connections?

State Laws Impacting GenAI

Consumer Rights – Deletion

- Archiving/deidentification?
  - Delete personal data except for backed up or archived data?
  - Deidentify data so that it cannot be linked to an identifiable individual?
  - How do you keep the technology from re-establishing the relationships that were severed to de-identify?
- Another escape hatch? Exceptions for unstructured data
  - **California**: Businesses are *not required to reidentify or otherwise link information that, in the ordinary course of business, is not maintained in a manner that would be considered "personal information," or maintain information in identifiable, linkable, or associable form, or collect, obtain, retain, or access any data or technology, in order to be capable of linking or associating a verifiable consumer request with personal information***.**
  - Other states have variations of this exception.

Hypothetical #2 - ZenCorp

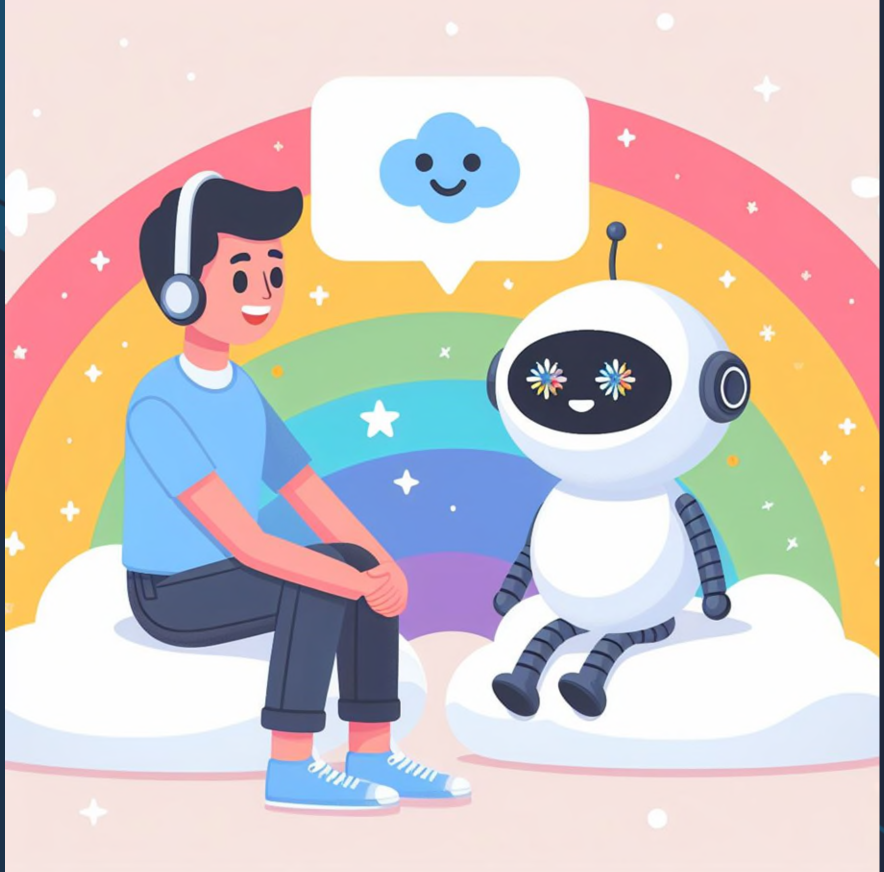

- *ZenCorp*
  - Provides a wellness chatbot that answers consumer questions about fitness, nutrition, workouts. Likely not covered by HIPAA.
  - ZenCorp's model provides the best answer based on the information provided by the consumer. The AI model is trained on personal information obtained across the internet and from licensed sources.
- Questions
  - What are the privacy issues and concerns?
  - How do you handle a request to delete personal information that was used to train the underlying model?
  - How do you handle a request from a current user to delete personal information?

State Laws Impacting GenAI

Consumer Rights – Correction and Access

- *Correction* of personal information
  - Same technical challenges as deletion: Where is it located? How do you correct it?
  - Erroneous outputs: Who is responsible – the user who prompted the error, the deployer, or the developer? How do you prevent erroneous AI outputs? Erroneous inputs?
- *Access* to personal information
  - Again, where is it located? How can you collect it so that you can provide access?
  - Keep exceptions in mind.

State Laws Impacting GenAI

Consumer Rights – Profiling Opt-Out Right

- Right to opt out of profiling that involves automated processing and has legal or other significant impact
- **Colorado** regulations:
  - "'*Profiling' … means any form of <u>automated processing</u> of personal data to evaluate, analyze, or predict personal aspects concerning an identified or identifiable individual's economic situation, health, personal preferences, interests, reliability, behavior, location, or movements.*"
  - "*Consumers have the right to <u>opt out of Profiling</u>…when the Profiling is done in furtherance of a decision that results in the provision or denial of financial or lending services, housing, insurance, education enrollment or opportunity, criminal justice, employment opportunities, health-care services, or access to <u>essential</u> goods or services …*"

State Laws Impacting GenAI

Developer and Deployer – Transparency

- Developer and Deployer Obligations: Transparency
  - Generally, businesses subject to state privacy laws must disclose the collection of personal information, categories of personal information collected, purposes for collection, and whether that personal information is sold or shared, among other things.
  - Disclosed in privacy policies. What does that look like with regard to GenAI? AI model card info?
- **Colorado**: For profiling using AI, "*[c]lear, understandable, and transparent information to Consumers in the required privacy notice.*" 6 specific factors including:
  - What decisions are subject to profiling.
  - "*Non-technical, plain language explanation of the logic used in the Profiling process.*"
  - "*A non-technical, plain language explanation of how Profiling is used in the decisionmaking process, including the role of human involvement, if any.*"

State Laws Impacting GenAI

Developer and Deployer – Transparency

- Other states require less detailed disclosures regarding AI for profiling:
  - **Oregon**: Provide "*a clear and <u>conspicuous description of any processing of personal data</u> in which the controller engages for the purpose of targeted advertising or for the purpose of profiling the consumer in furtherance of decisions that produce legal effects or effects of similar significance, and a procedure by which the consumer may opt out of this type of processing.*"

Other State Laws Impacting GenAI

- Washington My Health My Data Act
  - Consumer Health Data - broadly defined
  - To “collect,” need consent or necessity to provide a requested service
  - “’Personal information" does not include publicly available information.’”
  - Private right of action
- Utah Artificial Intelligence Policy Act
  - Amends consumer protection statute to require disclosure when consumer interacts with “generative artificial intelligence,” if “asked or prompted by that consumer,” that the consumer is interacting with GenAI and not a human
- Pending state legislation and regs: California, Colorado, Illinois, Maryland, New Jersey, New York, Tennessee, Indiana, and Alaska.
  - Over 50 pending bills across the states

Another Look at Zencorp and DataWhale…

- ZenCorp
  - How much of the training data is regulated under an AI-specific state law? Does the training constitute profiling?
  - What about the My Health My Data Act?
  - Does this affect whether ZenCorp can use the personal information to train the chatbot?
  - Logic disclosures
    - How realistic is it to expect ZenCorp to explain the logic of the chatbot in plain language when due to fine-tuning, the weights of the model may have changed?
- DataWhale
  - What are DataWhale's responsibilities with respect to training data that constitutes personal information? What are DataWhale's clients' responsibilities?
  - How is the Colorado law implicated?

For Developers and Deployers of GenAI Technology: *Action Items and Best Practices*

- Developers
  - Understand the source of the data
  - Know what is contained in the training dataset
  - Understand what the algorithm is doing and impact of machine learning
- Deployers
  - Implement processes for evaluating and honoring consumer rights
  - Ensure transparency and conduct risk assessments

QUESTIONS?