California Releases Draft Comprehensive Regs Governing Artificial Intelligence and Automated Decisionmaking Technology

Published: November 28, 2023

By: Daniel M. Goldberg

On November 27, the California Privacy Protection Agency (CPPA) released its long-anticipated draft regulations governing the use of automated decisionmaking technology (ADMT). The CPPA is set to discuss the draft ADMT regs at its next board meeting on December 8. As predicted, the draft ADMT regs are extremely comprehensive, and arguably create the first U.S. state framework governing the use of artificial intelligence (AI). Below are some quick takeaways from my initial read of the draft ADMT regs. This list is not comprehensive, and you should speak with legal counsel regarding potential implications. If you have any questions or thoughts, please send me a message.

An Important Read

The draft ADMT regs provide the first in-depth insight into how the CPPA intends to regulate ADMT and AI. At 17 pages, the draft ADMT regs lay out comprehensive obligations regarding ADMT. Along with the draft ADMT regs, the CPPA published a presentation regarding ADMT, which is worth a read. The CPPA also published revised drafts of its cybersecurity audit regs and risk assessment regs, both of which were updated to further address ADMT and better correspond with the ADMT regs. Notably, the draft ADMT regs only reference AI once; however, the risk assessment regs continue to include a definition and express obligations around AI. The interplay between these documents strongly suggests that the ADMT regs are intended to cover AI - as we all know, AI is on the minds of regulators and consumers.

As with the draft cybersecurity audit regs and risk assessment regs, the draft ADMT regs come with a disclaimer that they were prepared by subcommittee and the CPPA has not yet started its formal rulemaking process. Despite the disclaimer, I expect the draft regs to closely resemble the eventual final version. The CPPA has historically stuck with most of its proposed language. Further, the updated versions of the draft cybersecurity audit regs and risk assessment regs carry over most of the proposed language from the initial drafts. For takeaways on the revised draft cybersecurity audit regs, see our prior post here.

ADMT Framework: Three Requirements

The ADMT regs primarily focus on a new proposed framework governing the use of ADMT. There are three main requirements:

(1) Pre-Use Notice Requirement: Under this requirement, any business that uses a covered ADMT must provide consumers with certain disclosures regarding use of the ADMT.

(2) Opt-Out Requirement: Under this requirement, any business that uses a covered ADMT must provide consumers with the ability to opt-out of their personal information being processed using the ADMT.

(3) Access Right Requirement: Under this requirement, any business that uses a covered ADMT must provide consumers with the ability to request details about the business's use of the ADMT to process their personal information.

I expect the CPPA to spend significant time at its next board meeting discussing the ADMT framework. As California often leads US regulatory compliance, the final version of this ADMT framework may serve as a model for other states as they implement their own ADMT and AI laws.

Threshold for Requirements

The good news from a business compliance perspective is that the ADMT framework only applies to certain covered processing operations:

First, there must be an ADMT involved. This is not a high bar to meet. Pursuant to the ADMT regs, an ADMT is "any system, software, or process—including one derived from machine-learning, statistics, or other data-processing or artificial intelligence— that processes personal information and uses computation as whole or part of a system to make or execute a decision or facilitate human decisionmaking."

Second, the ADMT must be involved in at least one of the following uses (pursuant to Section 7030(b)):

(1) A decision that produces legal or similarly significant effects concerning a consumer (such as decisions to provide or deny employment opportunities)

(2) Profiling a consumer who is acting in their capacity as an employee, independent contractor, job applicant, or student

(3) Profiling a consumer while they are in a publicly available space

(4) Profiling a consumer for behavioral advertising

(5) Profiling a consumer that the business has actual knowledge is under the age of 16

(6) Processing the personal information of consumers to train ADMT

If you have been following privacy law over the past decade, none of these should be surprising. All six of these processing activities are expressly identified in the draft risk assessment regs as “present[ing] significant risk to consumers’ privacy”, which requires the conducting of a risk assessment. (6) feels like a catch-all for AI, in particular.

Notably, the ADMT regs specify that (4), (5), and (6) are options for board discussion. Again, given that the CPPA has historically stuck with most of its proposed language, I expect this language (or similar language) to be in the final version.

Detailed Requirements

Below are some relevant details for each of the requirements under the ADMT framework.

<u>Pre-Use Notice Requirement</u>
The pre-use notice requirement is similar to the “notice at collection” requirement under the original CPRA regs. Under the pre-notice requirement, any business that uses a covered ADMT must inform consumers of the following:

- that the business uses the ADMT,
- the purpose of the use of the ADMT (which must not be in generic terms),
- that consumers have the right to opt-out of the use of the ADMT, and
- that consumers have the right to access information about the use of the ADMT

While the pre-notice requirement may not be difficult to comply with, businesses may have difficulty complying with a related obligation that any business that uses a covered ADMT must provide a resource (such as through a layered notice or hyperlink) where consumers can obtain additional information about the business’s general use of ADMT, including an explanation of:

- the logic used by ADMT,
- the intended output of ADMT,
- how the business intends to use the output of ADMT (including any human involvement), and
- whether the business’s use of ADMT has been evaluated (and the outcome of any such evaluation)

This looks like a public bias audit or risk assessment, which is similar to transparency obligations found in other laws (such as New York City’s AI Bias law).

<u>Opt-Out Requirement</u>
The ADMT opt-out requirement shares much in common with the Do Not Sell/Share opt-out requirement under the original CPRA regs. Under the ADMT opt-out requirement, any business that uses a covered ADMT must provide consumers with the ability to opt-out of the ADMT.

Where a consumer opt-outs, the business must cease processing the consumer's personal information using that ADMT within 15 business days, and notify all downstream recipients of the personal information to comply with the opt-out with respect to the ADMT.

The method for submitting ADMT opt-outs appears to be a combination of Do Not Sell/Share and verifiable consumer request methods. Some notable aspects:

- A business must offer an interactive form as well as at least one other method for the opt out
- A business may require verification if the business determines and documents that consumers are more likely than not to be negatively impacted absent verification; however a business may not require verification for opt-outs of profiling for behavioral advertising
- A business must provide a means by which consumers can confirm the business processed their requests
- A business must respond to authorized agent requests if the authorized agent provides written permission signed by the consumer
- A business must offer an ADMT opt-out specific to ADMT requests; however, relying on cookie banners or cookie controls is not sufficient to address this right
- There is no express obligation to respond to preference signals for ADMT use, such as GPC signals

Access Right Requirement

The access right requirement is similar to the "right to know" requirement under the original CPRA regs. Under the access right requirement, any business that uses a covered ADMT must provide consumers with the ability to request information about the business's use of ADMT with respect to their personal information. Consumers must verify their identities, and businesses must address verifiable consumer requests within 45 days.

Where a consumer exercises their right, the business shall provide the following:

- The purpose for which the business used the ADMT,
- The output of the ADMT with respect to the consumer,
- How the business used the output with respect to the consumer,
- If the business plans to use the output to make a decision, a specific explanation regarding that decision,
- How the ADMT worked with respect to the consumer,
- A method by which the consumer can obtain a range of possible outputs, which may include aggregate output statistics,
- Instructions for how the consumer can exercise their other CPRA rights, and
- Instructions regarding methods by which the consumer can submit a complaint to the business, the CPPA, and the AG's Office regarding ADMT

Similar to the ADMT explanation in the pre-use notice requirement, this access right requirement may be difficult for businesses to address. Also, some may argue that this level of required detail goes far beyond the CPRA statutory text and privacy law. I anticipate that there may be legal challenges to some of these requirements.

Exceptions – Section 7030(m)

Section 7030(m) is one of the most important provisions of the new ADMT regs. This section sets out exceptions to the pre-use notice, opt-out, and access rights requirements. Pursuant to this section, a business is not required to provide consumers with pre-use notice, opt-out, or access rights where the business's ADMT use is necessary to achieve, and is solely for, the following purposes:

(1) Security: To prevent, detect, and investigate security incidents
(2) Fraud prevention: To resist malicious, deceptive, fraudulent, or illegal actions
(3) Safety: To protect the life and physical safety of consumers
(4) Requested Good or Service: To provide the good or perform the service specifically requested by the consumer, provided that the business has no reasonable alternative method of processing. There is a rebuttable presumption that the business has a reasonable alternative method of processing.

Notably, there is some ambiguity around the exceptions. For example, Section 7030 states that the opt-out right exception only applies where the ADMT complies with Section 7002 (the reasonable expectation test) while Section 70301 states that the access right exception only applies where the response would compromise the processing for purposes (1)-(3). It is not clear why (or if) these rights and corresponding exceptions should be treated differently. I expect the exceptions to be a topic of discussion during the board meeting.

Limited Obligations for Service Providers

There is only a single line imposing obligations on service providers. Under the draft ADMT regs, a service provider must provide assistance to the business in responding to verifiable consumer access requests. Of course, other parts of the CPRA regs impose specific obligations on service providers, but it is interesting we did not see more here.

Special Rules for Children Under 16

The draft ADMT regs also set out some specific rules for children under 16. Where a business has actual knowledge that it profiles a consumer less than 16, it must obtain opt-in consent. For under 13, that consent must be from the parent, and must be separate from the verifiable parental consent required under COPPA. I am surprised that the ADMT regs allow for any profiling of children under 16 (given the robust obligations under the pending Age Appropriate Design Code), and I could see that changing in the final version.

Submission of Risk Assessments to the CPPA

While not part of the ADMT regs, the revised draft risk assessment regs include new language regarding the process for submitting risk assessments to the CPPA. Per the regs, the first submission is due 24 months from the effective date of the regs, and subsequent risk assessments are due on an annual basis every calendar year. Risk assessments must be submitted through the CPPA’s website. From a practical perspective, I can’t see how the CPPA will be able (or want) to review all these risk assessments. Rather than proactively submit risk assessments, it seems that risk assessments should be provided upon request.

Cybersecurity Audit Regulations Under CCPA

Published: August 31, 2023

By: Rick Borden

The California Privacy Protection Agency (the "Agency") released draft Cybersecurity Audit Regulations ("Draft Regulations") for consideration by the Board of the Agency at a meeting schedule for September 8. The Draft Regulations provide that every business whose processing of personal information presents a significant risk to consumers' security will be required to perform a cybersecurity audit. Interestingly, the threshold for "significant risk" is proposed to be size-based instead of based on other risk factors, such as processing of sensitive personal information. This means that the audit requirement will be broad. As discussed below, there are a number of significant challenges for businesses subject to this requirement.

And, audit will mean audit, not assessment. That means that the auditor must use procedures and standards generally accepted in the auditing profession. These include (i) impartiality, and (ii) a reporting structure outside of the management chain that oversees the cybersecurity function. This also includes direct reporting to the board of directors, unless the company does not have a board. For companies that have an internal; audit function, the audit may be performed internally. Those that do not will have to hire an audit/accounting firm.

The audit requirement also has a thoroughness component, which includes the scope, criteria, and specific evidence observed and assessed. The audit will be required to be presented to the board, and a board member will be required to certify that they have reviewed the audit and understand the findings. This is not a trivial matter.

The scope of the required audit goes beyond a SOC2, which is currently the most widely offered cybersecurity audit. The cybersecurity audit will be required to "assess and document the business's cybersecurity program that is appropriate to the business's size and complexity and the nature and scope of its processing activities, taking into account the state of the art and cost of implementation." Additionally, the Agency Board has been provided options concerning additional scope requirements. Some of these potential scope requirements go beyond cybersecurity into core privacy requirements. These include:

- Impairing consumers' control over their personal information associated with the unauthorized access, destruction, use, modification, or disclosure of personal information; or unauthorized activity resulting in the loss of availability of personal information.
- Economic harm to consumers associated with the unauthorized access, destruction, use, modification, or disclosure of personal information; or unauthorized activity resulting in the loss of availability of personal information. This includes, for example, the direct and indirect costs associated with identity theft.

- Physical harm to consumers or to property associated with the unauthorized access, destruction, use, modification, or disclosure of personal information; or unauthorized activity resulting in the loss of availability of personal information.
- Psychological harm to consumers, including emotional distress, anxiety, embarrassment, fear, frustration, shame, and feelings of violation associated with the unauthorized access, destruction, use, modification, or disclosure of personal information; or unauthorized activity resulting in the loss of availability of personal information.
- Reputational harm to consumers, including stigmatization associated with the unauthorized access, destruction, use, modification, or disclosure of personal information; or unauthorized activity resulting in the loss of availability of personal information.

The scope includes a list of 18 specific technical and administrative safeguards (plus numerous sub-items) that must be assessed, and if they are not in place, the audit (not the business) must provide an explanation as to why they are not necessary, and how the safeguards in place have at least equivalent security. One particular requirement required in the audit scope is zero trust architecture ("ZTA"). This is very new in cybersecurity, and is not widely used throughout existing networks. For those who would like to understand some of the technical complexities, the NIST Special Publication 800-207 may be found here. For those who are not as technical, this statement from NIST 800-207 will give you a flavor: "Transitioning to ZTA is a journey concerning how an organization evaluates risk in its mission and cannot simply be accomplished with a wholesale replacement of technology." When NIST says that something is a journey, it will be very, very hard to audit!

Service providers and contractors will be required to "assist" the company in completing its cybersecurity audit. What this means is not clear. Presumably, they will have to provide information to the company. But, the requirements could be read much more broadly. The 18 technical and administrative controls described above may, or may not, apply to service providers or contractors. If they do, the amount of information required to be collected from service providers and contractors may be overwhelming, except for the largest enterprises. Currently, large enterprises struggle with service provider cybersecurity assessments. Making this an audit requirement will significantly increase the costs for many service providers and businesses.

Perhaps most concerning is that the Agency appears to have borrowed the concept of certification from the New York Department of Financial Services. Businesses required to conduct audits will be required to "submit to the Agency either: (1) A written certification that the business complied with the requirements set forth in this Article during the 12 months that the audit covers; or (2) A written acknowledgment that the business did not fully comply with the requirements set forth in this Article during the 12 months that the audit covers. The written acknowledgement shall: (A) Identify all sections and subsections of this Article that the business has not complied with and describe the nature and extent of such noncompliance; and (B) Provide a remediation timeline or confirmation that remediation has been completed." A member of the company's board, or a specified officer if there is no board, will be required to sign the certification or acknowledgment. That means that the Agency will have evidence every year of the companies that have conducted audits, and those that have not. It makes enforcement much easier.

Businesses will have 24 months from the date the regulations are completed to finalize cybersecurity audits, plus an updated audit annually thereafter. As a SOC2 generally takes at least a year, as it includes 6 months of control effectiveness testing, in-scope businesses should call their counsel and auditor **today** to get on their calendar and begin the process. It is likely that the number of businesses required to complete these audits will far outstrip the supply of auditors. The next 24 months will be a wild ride!

FloSports Settles "Meta Pixel" Litigation Claim for $2.6 Million

Published: August 2, 2023
By: Daniel M. Goldberg and Zach Lewis

In a development highlighting the recent wave of litigation around tracking technologies, live event streaming platform FloSports Inc. has agreed to a $2.6 million class action settlement to resolve claims related to its alleged use of the “Meta Pixel.” The Meta Pixel (or “Facebook Pixel”) “is a piece of code” on a company’s website that helps them understand “the actions people take on [their] site, like visiting a page or adding an item to their cart.” The Meta Pixel also allows companies to “see when customers took an action after seeing [their] ad on Facebook and Instagram[.]”

FloSports joins the ranks of hundreds of other companies across various industries that have recently faced lawsuits alleging violations of the Federal Wiretap Act, Video Privacy Protection Act (VPPA), Health Insurance Portability and Accountability Act (HIPAA), and state wiretapping laws in connection with their use of the Meta Pixel and similar tracking technologies. In these disputes, plaintiffs claim that by using such tracking technologies, website operators share users’ private information with third parties without their consent. Here, the plaintiffs claimed that when users watch a video on FloSports’ website, the Meta Pixel communicates to Meta the title of the video, its URL, and the viewer’s Facebook ID (which can be used to view and access their Facebook profile), which constitutes a violation of the VPPA. The plaintiffs sought $2,500 per violation of the VPPA in addition to attorneys’ fees and expenses.

According to an unopposed motion for preliminary approval of the class action settlement, FloSports revealed information during mediation proceedings that suggested it could not satisfy an adverse judgment and that “a verdict in favor of the class in the full amount of claimed damages likely would force the company into bankruptcy.” Pursuant to the settlement agreement, FloSports will be required to limit its use of the Meta Pixel on its websites to comply with the VPPA in addition to making the $2.6 million payment.

Key Takeaways:

- This case demonstrates that deploying tracking technologies on a website carries inherent risk of litigation claims. While the Meta Pixel is widely used across the internet, plaintiff's attorneys are targeting companies for quick pay days for alleged violations of privacy laws.
- While this lawsuit resulted in a settlement, we believe these claims likely would have failed at trial. However, most companies are afraid to risk an unfavorable judgment, and therefore decide to settle such claims.
- Due diligence is key. Companies must regularly review their data privacy policies and practices, especially as they relate to tracking technologies. Be careful when rolling out tracking technologies that are more intrusive, such as those that capture key strokes or record screens.

Our team is very familiar with these types of cases. Please contact us if you need any help.

> "Under the proposed Settlement, FloSports will create a $2.625 million non-reversionary cash fund for the benefit of the Settlement Class. FloSports also will agree to suspend operation of the Facebook Pixel on portions of its website relevant to VPPA compliance—i.e., webpages that both include video content and have a URL that identifies the video content viewed."

The Generative AI Journey Continues

Published: November 13, 2023

By: Brian Murphy

When it comes to tech powered by generative AI, we are on the verge of a Cambrian Explosion, a period during which we can expect an unprecedented variety of innovative and spectacularly useful tools to emerge from the pre-AI primordial muck. And, at the dawn of this era, courts already are issuing preliminary rulings in cases where plaintiffs are challenging, and defendants are championing, these tools, cases that could determine how and to what extent these tools are allowed to flourish. Jeremy Goldman wrote about a recent case (see this post), and today I write about another.

Andersen v. Stability AI, LTD is a class action lawsuit brought by three visual artists against Stability AI, the developer of Stable Diffusion, a latent, text-to-image model capable (in its words) of “generating photo-realistic images given any text input.” The plaintiffs also sued DeviantArt and Midjourney, whose generative AI tools (respectively, DreamUp and the eponymous Midjourney) are alleged to incorporate Stable Diffusion technology. The plaintiffs contend that these tools infringe upon their copyrights, rights of publicity and other rights.

Last month, Judge William H. Orrick dismissed the majority of the plaintiffs' claims. But this is not the end … not by a long stretch. The court let stand arguably the most significant of the plaintiffs' claims (the claim for direct copyright infringement arising from the use of billions of images to train Stable Diffusion), and the court gave the plaintiffs leave to amend “to provide clarity regarding their theories" underlying their other claims.

Below I summarize the court's ruling on the motion to dismiss the plaintiffs' direct copyright infringement and right of publicity claims.

Direct Copyright Infringement (Input) Against Stability

The plaintiffs alleged that defendant Stability was liable for direct copyright infringement because it had used their works (along with billions of others) as training images for its Stable Diffusion product. Specifically, the plaintiffs alleged that Stability paid Large-Scale Artificial Intelligence Open Network (LAION) - a non-profit that (in its own words) aims “to make large-scale machine learning models, datasets and related code available to the general public" - to scrape from the internet over five billion images (among them the plaintiffs') for training.

Because two of the plaintiffs (McKernan and Ortiz) had not registered their works with the Copyright Office (a prerequisite for suing, 17 U.S.C. § 411), the court dismissed their claims with prejudice. (It's not clear why the court dismissed these claims with prejudice - perhaps because counsel stated during oral argument that they were not pursuing copyright claims on behalf of these plaintiffs?) The defendants' argued that the claims by the third plaintiff (Andersen) - who had properly registered sixteen collections of her works - should be dismissed because she had failed to specify which of the works in those collections had actually been used by Stability in training. The court disagreed, finding that Andersen's reliance on the search results from the website “ihavebeentrained.com” (which indicated that many of her works had, in fact, been used in training), combined with the complaint's allegation that LAION had scraped over five billion images for its training datasets, supported “the plausibility and reasonableness of her belief" that those of her registered works that were posted online had, in fact, been scraped into the training datasets.

Finally, the court held that the plaintiffs' allegations that Stability had, without permission, “downloaded or otherwise acquired copies of billions of copyrighted images without permission" and caused those “images to be stored at and incorporated into Stable Diffusion as compressed copies" were sufficient to plead direct copyright infringement by Stability. According, the court refused to dismiss the plaintiffs' “primary theory” of direct copyright infringement.

Direct Copyright Infringement (Input) Against DeviantArt

Defendant DeviantArt hosts an online community where digital artists can share their works; it also offers its own AI image generator, DreamUp, which is powered by Stable Diffusion. Plaintiffs alleged that one of the LAION datasets used to train Stable Diffusion was created by scraping DeviantArt’s site. At the outset, the court held that merely “being a primary source” for training images did not support a claim for direct infringement.

However, the plaintiffs also alleged that (1) “compressed copies” of training images are embedded within Stable Diffusion, and (2) DeviantArt was liable for direct infringement because it distributed Stable Diffusion (and, therefore, the “compressed copies" of training images embodied therein) as part its own DreamUp product. The plaintiffs did not appear to be alleging that actual copies of all of the training images were incorporated within Stable Diffusion. Indeed, as the defendants noted, that would be impossible since no active application could compress five billion images. Instead, the complaint described Stable Diffusion as providing “an alternative way of storing a copy of [training] images” that used “statistical and mathematical methods to store these images in an even more efficient and compressed manner.”

Ultimately, the court agreed with the defendants that plaintiffs' allegations were unclear and contradictory and invited the plaintiffs to amend the complaint to allege with greater clarity precisely how the training images were “embedded” within Stable Diffusion:

> “If plaintiffs contend Stable Diffusion contains “compressed copies” of the Training Images, they need to define “compressed copies” and explain plausible facts in support. And if plaintiffs’ compressed copies theory is based on a contention that Stable Diffusion contains mathematical or statistical methods that can be carried out through algorithms or instructions in order to reconstruct the Training Images

> in whole or in part to create the new Output Images, they need to clarify that and provide plausible facts in support."

Direct Copyright Infringement (Output) Against DeviantArt

Plaintiffs also alleged that DeviantArt's DreamUp program produces and distributes output images that are infringing derivative works of the training images. The defendants urged the court to dismiss the claim because the plaintiffs had failed to allege that the output images were substantially similar to the plaintiffs' copyrighted works: on the contrary, the plaintiffs had admitted in the complaint that "none of the Stable Diffusion output images provided in response to a particular Text Prompt is likely to be a close match for any specific image in the training data." In response, the plaintiffs contended that *all* elements of plaintiff Anderson's copyrighted works (and the copyrighted works of all others in the purported class) "were copied wholesale as Training Images and therefore the Output Images are necessarily derivative."

Once again, the court found there were numerous defects in the plaintiffs' complaint. As with the direct infringement (input) claims, the plaintiffs' "theory regarding compressed copies and DeviantArt's copying needs to be clarified and adequately supported by plausible facts." Moreover, the court found it "simply not plausible" that all of the images used to train Stable Diffusion were copyrighted (as opposed to copyrightable), or that all the output images were derivative of copyrighted training images. And, perhaps most important, the court was "not convinced that copyright claims based on a derivative theory can survive absent 'substantial similarity' type allegations." (*See* this post.) Accordingly, the court dismissed the claim, with leave to amend.

Direct Copyright Infringement (Input and Output) Against Midjourney

The court also dismissed (with leave to amend) the plaintiffs' direct infringement claims against Midjourney, which were nearly identical to those it brought against DeviantArt. However, the court called out that the plaintiffs had failed to allege facts regarding what training, if any, Midjourney had conducted for its Midjourney product, and that the plaintiffs needed to clarify if their theory of liability "is it based on Midjourney's use of Stable Diffusion, on Midjourney's own independent use of Training Images to train the Midjourney product, or both?"

Plaintiffs' Right of Publicity Claims

In their complaint, the plaintiffs asserted that the defendants misappropriated their names and their "artistic identities," in violation of their statutory and common law right of publicity, because the defendants' AI tools allow users to request art "in the style of their" names. In their brief and at the hearing, the plaintiffs "clarified" that their claims were based on the defendants' use of their names to advertise and promote their DreamStudio, DreamUp, and Midjourney products.

The court dismissed the claims, once again with leave to amend:

> "The problem for plaintiffs is that nowhere in the Complaint have they provided any facts specific to the *three named plaintiffs* to plausibly allege that any defendant has used a named plaintiff's name to advertise, sell, or solicit purchase of DreamStudio, DreamUp or the Midjourney product. Nor are there any allegations regarding how use of these plaintiffs' names in the products' text prompts would produce an "AI-generated image similar enough that people familiar with Plaintiffs' artistic style could believe that Plaintiffs created the image," and result in plausible harm to their goodwill associated with their names, in light of the arguably contradictory allegation that none of the Output Images are likely to be a "close match" for any of the Training Images. Plaintiffs need to clarify their right of publicity theories as well as allege plausible facts in support regarding each defendants' use of each plaintiffs' name in connection with advertising specifically and any other commercial interests of defendants." (Cleaned up.)

Since it had dismissed the right of publicity claims with leave to amend, the court refused at this juncture to consider the defendants' First Amendment defense - i.e., that the output of these tools was "transformative" under *Comedy III Productions, Inc. v. Gary Saderup, Inc*., 25 Cal.4th 387 (2001). The court invited the defendants to raise this defense again after the plaintiffs have amended their complaint and clarified their theories of liability for the right to publicity claims.

Andersen v. Stability AI LTD, No. 23-cv-00201-WHO (N.D. Cal. Oct. 30, 2023)