Spring 2025 **Blind ID Number: _____**

UNIVERSITY OF CALIFORNIA
LOS ANGELES
SCHOOL OF LAW

Prof. Selbst | Law 495. §1 Artificial Intelligence Law | May 7, 2025

Final Examination
OPEN BOOK
FULLY SECURE
(3 hours, 30 minutes)

Be sure to avoid making any remarks that will jeopardize your (or anyone else's) anonymity either on the exam or in any other context, before the exam grades are posted. Any questions during or after the exam should be addressed to the Records Office or Dean Martí.

Exam Instructions

1. Please write your blind ID in the top right corner of this page.

2. This is an open book, fully secured examination. You may use your casebook, assigned course materials, class PowerPoints, and any printed notes and outlines you have prepared. You may not use electronic materials outside of the exam software.

3. The word limit for the entire exam is **4000** words. I will not read past that number of words. Do not waste words restating the facts or the exam question!

4. The exam has three questions—two issue spotters and one essay question. I recommend you allocate your time according to the percentage indicated in each question/part.

5. Your answers will be evaluated according to responsiveness to the question, clear expression of your analysis, reliance on appropriate authority, and discussion of policy considerations where appropriate. You should discuss relevant cases and statues explicitly and draw the relevant comparisons to the fact patterns. I want to see that you understand the relevant steps and considerations in answering each question, and that you have the ability to work through a problem by identifying and citing the correct sources of law.

6. You may leave the room to use the restroom or for an emergency, but you may not take anything out with you, or bring anything in with you, or consult anyone or anything while you are outside the room.

7. If you finish the examination early (up until the 5 minute warning is given), you may bring your exam materials to the front of the room and quietly leave the room. All of these materials must be turned in, or you will receive no credit for the examination.

8. Exam answers must be entered in the exam software, not on this question booklet.

9. If you feel that there has been a typographical error or other mistake on the exam, identify the error and answer the question as you think it should be corrected to read.

* * *

Examplify Features

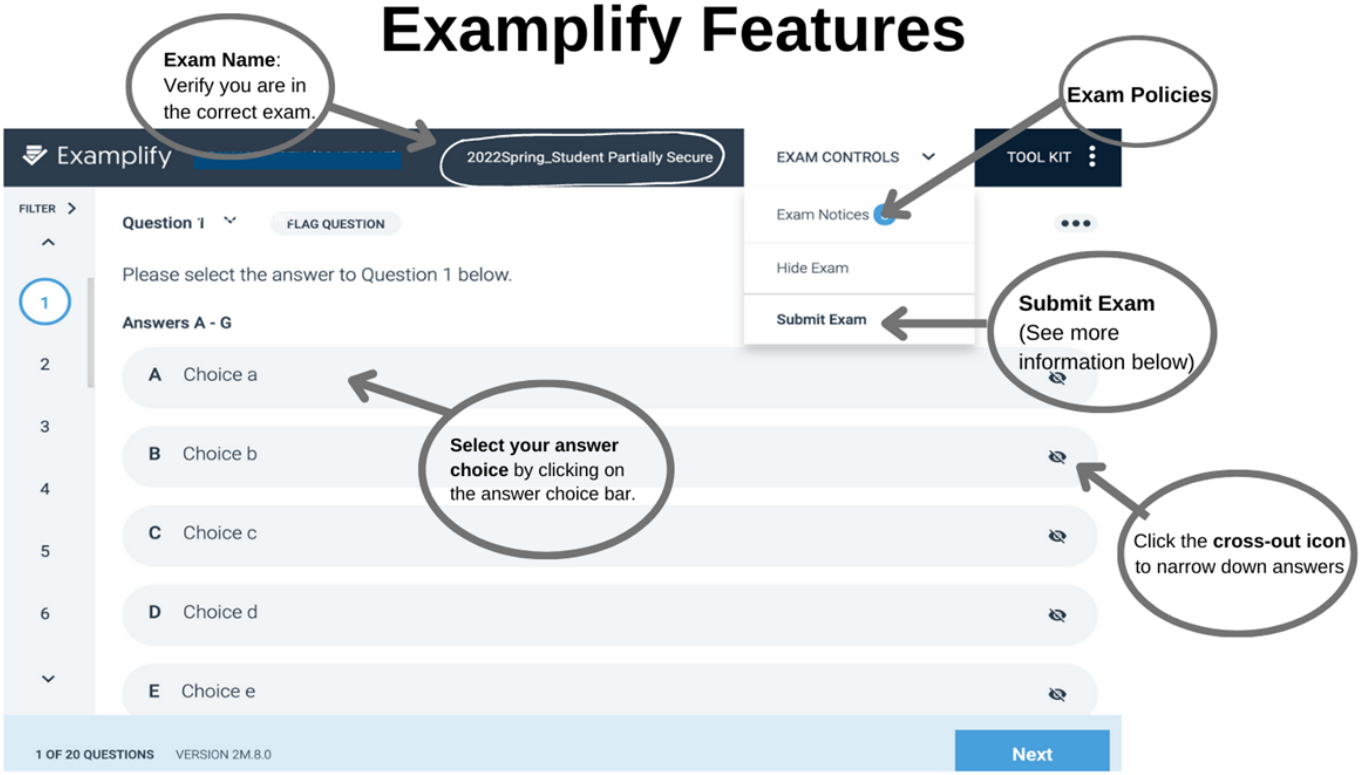

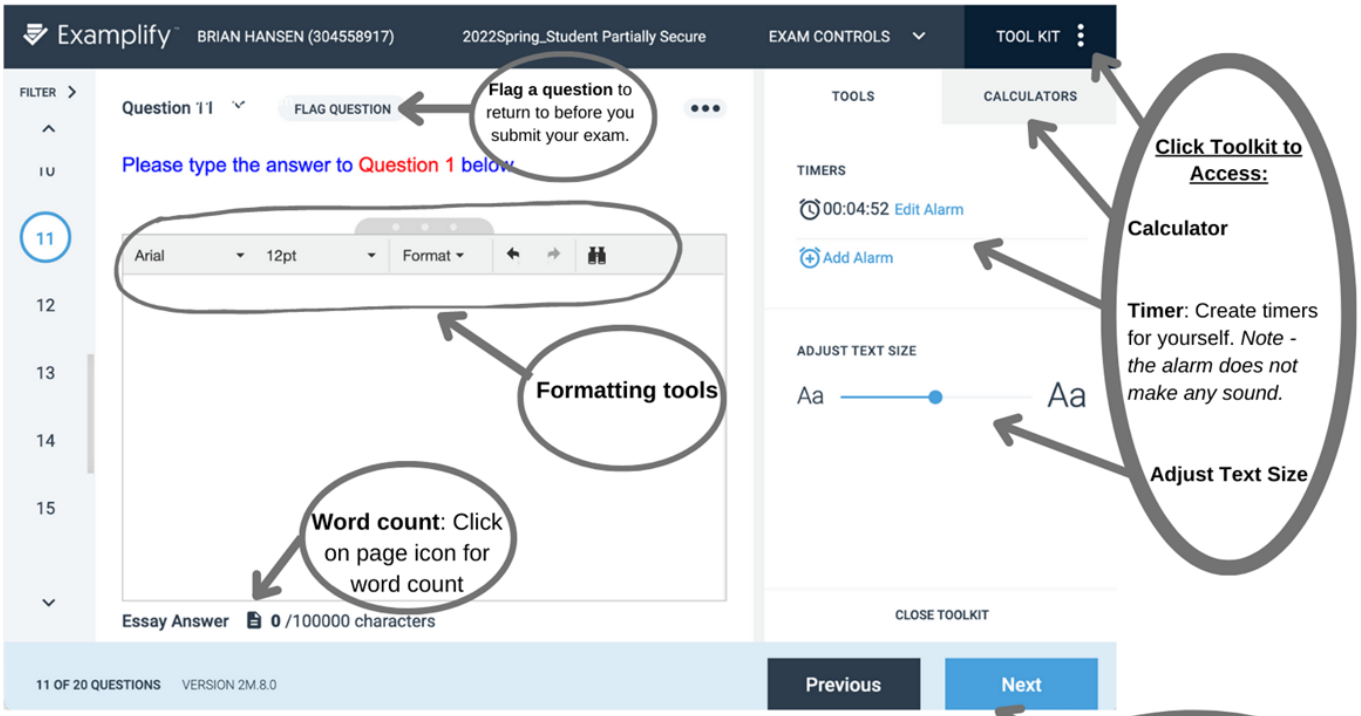

Submitting Your Exam:

- Click "Exam Controls" then "Submit Exam"
- Your computer will reconnect to the internet. *Note: If you computer has trouble reconnecting, try manually connecting to the UCLA_WEB wireless network.*
- Make sure you receive a green confirmation screen before you leave the exam room.

UCLA School of Law

Question 1. (90 minutes)

SocioTech Systems is a technology company that contracts with governments across the United States and Europe to provide automated eligibility determinations for public assistance programs (such as Medicaid, SNAP benefits, unemployment insurance, and housing subsidies). Its product, BenefitMax, is an AI system that analyzes large volumes of data from various government databases—including income, employment history, residency, eviction records, criminal records, and even "fraud risk indicators" (such as frequent address changes or gaps in employment) to determine eligibility for various benefit programs. Applicants are assigned eligibility scores between 0 and 100. The scores are based on complex, proprietary machine learning model trained on years of administrative data (e.g., past eligibility determinations, interactions with social services, employment history, housing instability).

SocioTech's government clients have only limited access to the systems. Depending on the particular system and jurisdicton, the eligibility thresholds and proceudres for review vary; each are determined in consulation with the government agency in question. Ultimately, agency workers receive eligibility decisions and high-level "reason codes" (e.g., "insufficient employment history" or "unverified residence") but cannot audit or independently replicate the underlying scores. SocioTech claims trade secrecy over its algorithm, and refuses to disclose its internal logic, weighting of factors, or training datasets, even under protective order. On its website, SocioTech states that it conducts regular validation and fairness assessments of BenefitMax.

Part A

The Missouri Department of Social Services (MDSS) uses BenefitMax to automate decisions about eligibility for Medicaid and SNAP benefits. When a person applies for Medicaid or SNAP benefits through the State Department of Assistance, they submit their data through an online portal. The BenefitMax system automatically takes in and processes the data. Applicants then receive a simple eligibility notification by email: either "Approved," "Denied," or "Pending Review." If denied, the notification includes only a pre-populated reason code, as well as a description of the appeal procedure.

Caseworkers are supposed to review BenefitMax denials before finalization. The review is limited to checking that data fields were properly filled (e.g., no missing social security number)—not a substantive reassessment of eligibility. Caseworkers have five minutes allocated per case and are instructed to defer to the system output unless a clerical error is obvious. Upon receipt of the denial, applicants have 15 days to file an appeal by filling out a web form or mailing a letter. Upon appeal, a second BenefitMax review occurs. The appeal does not address the algorithm's underlying analysis. The appeals adjudicators cannot access the underlying model; they are allowed to overturn denials only if the applicant provides new documentary evidence (e.g., updated pay stubs). Applicants are not provided any additional information about how their eligibility was assessed unless they formally request their "administrative file," which contains only input data and the system's "reason code," not the logic or scoring model. No opportunity for a live hearing is provided at the appeal stage; adjudication is based entirely on written submissions. Appeals decisions are considered final, allowing denied applicants to challenge the decision in state court.

After six months of use, an internal audit revealed that Black and Latino applicants are disproportionately denied benefits compared to white applicants, even when income levels are comparable. An investigation by the nonprofit Techtonic Justice revealed that many of the "fraud risk indicators" used by BenefitMax—like unstable housing or frequent job changes—are themselves correlated with race and poverty due to historical and systemic inequalities.

You are an attorney with Missouri Legal Aid. Patrice Davis, a Black applicant denied Medicaid benefits, has come to you asking if you could challenge the decision on her behalf. Please write a memo evaluating the strength of potential claims under Title VI and due process. (Title VI bars discrimination by programs receiving federal funding, which includes Missouri's Medicaid program. Assume Title VI includes disparate

impact provisions that parallel Title VII's). Please discuss with particularity where each claim is strong and what challenges you expect to encounter. If there are facts you will need to find in discovery to evaluate your case, please explain their signficance to the case, whether you will face challenges in finding out what you need, and what is likely to happen depending on what you do or do not find.

Part B

In Bruges, Belgium, the Centre Public d'Action Sociale (CPAS)—the local public assistance agency—implements BenefitMax to automate eligibility for unemployment assistance. Applicants' information (income history, job searches, address history, prior social assistance use) is processed by the proprietary BenefitMax system. The system assigns a risk score, then: If the score falls below 50, the system recommends: "deny." If the score is above 50, the system recommends "manual review."

Before finalizing denials, an agency worker reviews the BenefitMax output, checking for data-entry errors and verifying that an automated reason code is populated. In the last six months of operation, 98% of the "deny" recommendations are processed without modification.

If a case is marked for manual review, the application is forwarded to a human caseworker along with a "recommended outcome"—either approval or denial. Caseworkers are instructed to independently assess eligibility based on available applicant documentation. Caseworkers can also request additional evidence from the applicant. Internal policy guidelines state that the system's recommendation is "advisory but important." If a caseworker deviates from the BenefitMax recommendation, they must write a detailed jsutificaiton memo explaining the deviation. Caseworkers' decisions, whether following or deviating from the BenefitMax recommendation, are subject to a second-level supervisory check for consistency. Agency audits reveal that in 85% of "manual review" cases, caseworkers align with the system's recommendation.

A group of applicants who have been denied benefits have filed a challenge to this system under the GDPR with the Beligan data protection authority (DPA). You are now at attorney in the DPA who is determining whether to begin an investigation into the Bruges CPAS and SocioTech. In addition to GDPR concerns, the DPA is considering whether SocioTech's BenefitMax system falls within the scope of the EU AI Act, which recently entered into force.

Please write a memo evaluating the applicants' case under Article 22 and Articles 13–15 of the GDPR. In your memo, please also include an analysis of the proper risk classification for BenefitMax under the AI Act and a discussion of what SocioTech needed to do to obtain a conformity assessment. Be specific about your reasons.

EXAM CONTINUES ON NEXT PAGE

Question 2. (75 minutes)

CraftGen is a generative AI model developed by a San Francisco startup, Craft.ai. The system is a version of Meta's Llama3.1 open weights model, fine-tuned on how-to guides, blogs, books, and videos sourced from the internet. Many of the sources were protected by copyright, including paywalled e-books and instructional manuals.

CraftGen is sold as a downloadable software product to small businesses, hobbyists, and freelance artisans. Users input a desired craft or building project—for example, "Make a child's treehouse out of reclaimed wood"—and CraftGen generates a detailed step-by-step plan with text, diagrams, material lists, and safety warnings. In the CraftGen Pro version, a user can also submit pictures of a work site, so that the instructions can be more specifically tailored to the project.

Quinn Hernandez is a freelance carpenter in northern California. His clients, Marc and Emily Romero, have hired him to build for a large treehouse for their two children, Lila and Jaime, aged 8 and 6. Hernandez has 15 years' experience as a carpenter, but has never built a large treehouse before. Hernandez is active on the Freelance Carpenters subreddit, and for months everyone there has been raving about CraftGen Pro. It seems like in just the last year, nearly every freelancer on the internet has used it to great effect, and there are constant discussions about how it's going to revolutionize the industry (coupled with worries that they will all be out of jobs). He's been waiting for a chance to give it a shot, so he figures this is the perfect opportunity.

Hernandez purchases CraftGen Pro. He uploads pictures of the treehouse site, and then instructs CraftGen to "design a structurally sound, safe treehouse using recycled wood and local materials." CraftGen outputs a 7-page instruction set that includes:

- A list of materials,
- Step-by-step building instructions,
- Load-bearing assumptions for platform supports,
- A general safety disclaimer ("consult your local codes and safety standards").

Hernandez reviews the plans, and they seem sensible enough to him. Relying on CraftGen's plans, he builds the treehouse. One month later, the structure collapses, seriously injuring Lila.

The Romeros sue Hernandez, alleging negligence in the construction. They attach an affidavit from structural engineer who reviewed the AI-generated plans, and opined that the load-bearing assumptions in the plan were flawed and dangerous. Hernandez then impleads Craft.ai, alleging that the liability belongs with the company because the software was defective.

Meanwhile, Practical House Books, a large publisher of how-to books, sues Craft.ai in a separate case for copyright infringement, arguing that large portions of its how-to books were used to train CraftGen without permission, that some instructions generated by CraftGen closely resemble paragraphs from its manuals or transcriptions of demonstration videos it has produced for their YouTube channel, which is well-known in the amateur DIY community.

You are clerking for a federal district judge handling both sets of cases. Your judge asks you to draft a memo evaluating:

- The negligence claims against Hernandez in *Romero v. Hernandez and Craft.ai*
- The products liability claims against Craft.ai in *Romero v. Hernandez and Craft.ai*.
- Copyright claims in *Practical House Books v. Craft.ai*

For the tort claims, your judge tells you that she used to teach torts, so she knows the basics. Accordingly, while it's important to mention the elements for completeness, she would like you to focus the memo on the aspects that are particularly relevant to the new AI developments.

Question 3. (45 minutes)

Throughout this course, we have frequently discussed the concept of "humans in the loop" (HITL) as a regulatory strategy. Often, we have criticized this approach for overestimating a human's ability to meaningfully intervene in or monitor algorithmic systems—especially under real-world conditions of complexity, opacity, and time pressure. Still, regulators and legislators repeatedly return to the idea. Why?

Please write an essay defending at least one reason for maintaining a human-in-the-loop requirement in a specific applied context—even if the human performs worse than an AI alternative in terms of accuracy or efficiency. In other words, identify and defend a non-accuracy-based justification for requiring human involvement. Your answer should:

- Identify a concrete domain or regulatory context (e.g., parole decisions, benefits administration, hiring, medical triage, etc.)
- Clearly state your proposed justification
- Use readings, cases, and/or discussions from class to support your argument

* * * *

END

* * * *

Good luck with the rest of your exams.

* * *

APPENDICES ON NEXT PAGE

Appendix A: Selected Text of the General Data Protection Regulation

Art. 5: Principles relating to processing of personal data

1. Personal data shall be:
    (a) processed lawfully, fairly and in a transparent manner in relation to the data subject (‘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, in accordance with Article 89 (1), 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 in accordance with Article 89 (1) subject to implementation of the appropriate technical and organisational measures required by this Regulation in order to safeguard the rights and freedoms of the data subject (‘storage limitation’);
    (f) processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures (‘integrity and confidentiality’).

2. The controller shall be responsible for, and be able to demonstrate compliance with, paragraph 1 (‘accountability’).

Art. 6: 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;
    (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.

Point (f) of the first subparagraph shall not apply to processing carried out by public authorities in the performance of their tasks.

. . .

4. Where the processing for a purpose other than that for which the personal data have been collected is not based on the data subject's consent or on a Union or Member State law which constitutes a necessary and proportionate measure in a democratic society to safeguard the objectives referred to in Article 23(1), the controller shall, in order to ascertain whether processing for another purpose is compatible with the purpose for which the personal data are initially collected, take into account, inter alia:
    (a) any link between the purposes for which the personal data have been collected and the purposes of the intended further processing;
    (b) the context in which the personal data have been collected, in particular regarding the relationship between data subjects and the controller;
    (c) the nature of the personal data, in particular whether special categories of personal data are processed, pursuant to Article 9, or whether personal data related to criminal convictions and offences are processed, pursuant to Article 10;
    (d) the possible consequences of the intended further processing for data subjects;
    (e) the existence of appropriate safeguards, which may include encryption or 8seudonymization.

Art. 9: Processing of special categories of personal data

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.

[2–4 omitted]

Art. 13: Information to be provided where personal data are collected from the data subject

1. Where personal data relating to a data subject are collected from the data subject, the controller shall, at the time when personal data are obtained, provide the data subject with all of the following information:
    (a) the identity and the contact details of the controller and, where applicable, of the controller's representative;
    (b) the contact details of the data protection officer, where applicable;
    (c) the purposes of the processing for which the personal data are intended as well as the legal basis for the processing;
    (d) where the processing is based on point (f) of Article 6(1), the legitimate interests pursued by the controller or by a third party;
    (e) the recipients or categories of recipients of the personal data, if any;
    (f) where applicable, the fact that the controller intends to transfer personal data to a third country or international organisation and the existence or absence of an adequacy decision by the

Commission, or in the case of transfers referred to in Article 46 or 47, or the second subparagraph of Article 49(1), reference to the appropriate or suitable safeguards and the means by which to obtain a copy of them or where they have been made available.

2. In addition to the information referred to in paragraph 1, the controller shall, at the time when personal data are obtained, provide the data subject with the following further information necessary to ensure fair and transparent processing:
    (a) the period for which the personal data will be stored, or if that is not possible, the criteria used to determine that period;
    (b) the existence of the right to request from the controller access to and rectification or erasure of personal data or restriction of processing concerning the data subject or to object to processing as well as the right to data portability;
    (c) where the processing is based on point (a) of Article 6(1) or point (a) of Article 9(2),the existence of the right to withdraw consent at any time, without affecting the lawfulness of processing based on consent before its withdrawal;
    (d) the right to lodge a complaint with a supervisory authority;
    (e) whether the provision of personal data is a statutory or contractual requirement, or a requirement necessary to enter into a contract, as well as whether the data subject is obliged to provide the personal data and of the possible consequences of failure to provide such data;
    (f) the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject.

[3–4 omitted]

Art. 14: Information to be provided where personal data have not been obtained from the data subject

1. Where personal data have not been obtained from the data subject, the controller shall provide the data subject with the following information:
    (a) the identity and the contact details of the controller and, where applicable, of the controller's representative;
    (b) the contact details of the data protection officer, where applicable;
    (c) the purposes of the processing for which the personal data are intended as well as the legal basis for the processing;
    (d) the categories of personal data concerned;
    (e) the recipients or categories of recipients of the personal data, if any;
    (f) where applicable, that the controller intends to transfer personal data to a recipient in a third country or international 9rganization and the existence or absence of an adequacy decision by the Commission, or in the case of transfers referred to in Article 46 or 47, or the second subparagraph of Article 49(1), reference to the appropriate or suitable safeguards and the means to obtain a copy of them or where they have been made available.

2. In addition to the information referred to in paragraph 1, the controller shall provide the data subject with the following information necessary to ensure fair and transparent processing in respect of the data subject:

(a) the period for which the personal data will be stored, or if that is not possible, the criteria used to determine that period;
(b) where the processing is based on point (f) of Article 6 (1), the legitimate interests pursued by the controller or by a third party;
(c) the existence of the right to request from the controller access to and rectification or erasure of personal data or restriction of processing concerning the data subject and to object to processing as well as the right to data portability;
(d) where processing is based on point (a) of Article 6 (1) or point (a) of Article 9 (2), the existence of the right to withdraw consent at any time, without affecting the lawfulness of processing based on consent before its withdrawal;
(e) the right to lodge a complaint with a supervisory authority;
(f) from which source the personal data originate, and if applicable, whether it came from publicly accessible sources;
(g) the existence of automated decision-making, including profiling, referred to in Article 22 (1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject.

3. The controller shall provide the information referred to in paragraphs 1 and 2:

(a) within a reasonable period after obtaining the personal data, but at the latest within one month, having regard to the specific circumstances in which the personal data are processed;
(b) if the personal data are to be used for communication with the data subject, at the latest at the time of the first communication to that data subject; or
(c) if a disclosure to another recipient is envisaged, at the latest when the personal data are first disclosed.

[4–5 omitted]

Art. 15: Right of access by the data subject

1. The data subject shall have the right to obtain from the controller confirmation as to whether or not personal data concerning him or her are being processed, and, where that is the case, access to the personal data and the following information:

(a) the purposes of the processing;
(b) the categories of personal data concerned;
(c) the recipients or categories of recipient to whom the personal data have been or will be disclosed, in particular recipients in third countries or international organisations;
(d) where possible, the envisaged period for which the personal data will be stored, or, if not possible, the criteria used to determine that period;
(e) the existence of the right to request from the controller rectification or erasure of personal data or restriction of processing of personal data concerning the data subject or to object to such processing;
(f) the right to lodge a complaint with a supervisory authority;
(g) where the personal data are not collected from the data subject, any available information as to their source;

(h) the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject.

[2–4 omitted]

Art. 22: Automated individual decision-making, including profiling

1. The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.

2. Paragraph 1 shall not apply if the decision:

   (a) is necessary for entering into, or performance of, a contract between the data subject and a data controller;
   (b) is authorised by Union or Member State law to which the controller is subject and which also lays down suitable measures to safeguard the data subject's rights and freedoms and legitimate interests; or
   (c) is based on the data subject's explicit consent.

3. In the cases referred to in points (a) and (c) of paragraph 2, the data controller shall implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision.

4. Decisions referred to in paragraph 2 shall not be based on special categories of personal data referred to in Article 9(1), unless point (a) or (g) of Article 9(2) applies and suitable measures to safeguard the data subject's rights and freedoms and legitimate interests are in place.

APPENDIX B ON NEXT PAGE

Appendix B: Selected Text of the EU Artificial Intelligence Act

Article 5: Prohibited AI Practices

1. The following AI practices shall be prohibited:
    (a) the placing on the market, the putting into service or the use of an AI system that deploys subliminal techniques beyond a person's consciousness or purposefully manipulative or deceptive techniques, with the objective, or the effect of materially distorting the behaviour of a person or a group of persons by appreciably impairing their ability to make an informed decision, thereby causing them to take a decision that they would not have otherwise taken in a manner that causes or is reasonably likely to cause that person, another person or group of persons significant harm
    (b) the placing on the market, the putting into service or the use of an AI system that exploits any of the vulnerabilities of a natural person or a specific group of persons due to their age, disability or a specific social or economic situation, with the objective, or the effect, of materially distorting the behaviour of that person or a person belonging to that group in a manner that causes or is reasonably likely to cause that person or another person significant harm;
    (c) the placing on the market, the putting into service or the use of AI systems for the evaluation or classification of natural persons or groups of persons over a certain period of time based on their social behaviour or known, inferred or predicted personal or personality characteristics, with the social score leading to either or both of the following:
        (i) detrimental or unfavourable treatment of certain natural persons or groups of persons in social contexts that are unrelated to the contexts in which the data was originally generated or collected;
        (ii) detrimental or unfavourable treatment of certain natural persons or groups of persons that is unjustified or disproportionate to their social behaviour or its gravity;
    (d) the placing on the market, the putting into service for this specific purpose, or the use of an AI system for making risk assessments of natural persons in order to assess or predict the risk of a natural person committing a criminal offence, based solely on the profiling of a natural person or on assessing their personality traits and characteristics; this prohibition shall not apply to AI systems used to support the human assessment of the involvement of a person in a criminal activity, which is already based on objective and verifiable facts directly linked to a criminal activity;
    (e) the placing on the market, the putting into service for this specific purpose, or the use of AI systems that create or expand facial recognition databases through the untargeted scraping of facial images from the internet or CCTV footage;
    (f) the placing on the market, the putting into service for this specific purpose, or the use of AI systems to infer emotions of a natural person in the areas of workplace and education institutions, except where the use of the AI system is intended to be put in place or into the market for medical or safety reasons;
    (g) the placing on the market, the putting into service for this specific purpose, or the use of biometric categorisation systems that categorise individually natural persons based on their biometric data to deduce or infer their race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation; this prohibition does not cover any

labelling or filtering of lawfully acquired biometric datasets, such as images, based on biometric data or categorizing of biometric data in the area of law enforcement;

(h) the use of 'real-time' remote biometric identification systems in publicly accessible spaces for the purposes of law enforcement, unless and in so far as such use is strictly necessary for one of the following objectives:

(i) the targeted search for specific victims of abduction, trafficking in human beings or sexual exploitation of human beings, as well as the search for missing persons;

(ii) the prevention of a specific, substantial and imminent threat to the life or physical safety of natural persons or a genuine and present or genuine and foreseeable threat of a terrorist attack;

(iii) the localisation or identification of a person suspected of having committed a criminal offence, for the purpose of conducting a criminal investigation or prosecution or executing a criminal penalty for offences referred to in Annex II and punishable in the Member State concerned by a custodial sentence or a detention order for a maximum period of at least four years. Point (h) of the first subparagraph is without prejudice to Article 9 of Regulation (EU) 2016/679 for the processing of biometric data for purposes other than law enforcement.

2. The use of 'real-time' remote biometric identification systems in publicly accessible spaces for the purposes of law enforcement for any of the objectives referred to in paragraph 1, first subparagraph, point (h), shall be deployed for the purposes set out in that point only to confirm the identity of the specifically targeted individual, and it shall take into account the following elements:

(a) the nature of the situation giving rise to the possible use, in particular the seriousness, probability and scale of the harm that would be caused if the system were not used;

(b) the consequences of the use of the system for the rights and freedoms of all persons concerned, in particular the seriousness, probability and scale of those consequences. In addition, the use of 'real-time' remote biometric identification systems in publicly accessible spaces for the purposes of law enforcement for any of the objectives referred to in paragraph 1, first subparagraph, point (h), of this Article shall comply with necessary and proportionate safeguards and conditions in relation to the use in accordance with the national law authorising the use thereof, in particular as regards the temporal, geographic and personal limitations. The use of the 'real-time' remote biometric identification system in publicly accessible spaces shall be authorised only if the law enforcement authority has completed a fundamental rights impact assessment as provided for in Article 27 and has registered the system in the EU database according to Article 49. However, in duly justified cases of urgency, the use of such systems may be commenced without the registration in the EU database, provided that such registration is completed without undue delay.

[3–8 omitted]

Article 6: Classification Rules for High-Risk AI Systems

1. Irrespective of whether an AI system is placed on the market or put into service independently of the products referred to in points (a) and (b), that AI system shall be considered to be high-risk where both of the following conditions are fulfilled:

(a) the AI system is intended to be used as a safety component of a product, or the AI system is itself a product, covered by the Union harmonisation legislation listed in Annex I;

(b) the product whose safety component pursuant to point (a) is the AI system, or the AI system itself as a product, is required to undergo a third-party conformity assessment, with a view to the placing on the market or the putting into service of that product pursuant to the Union harmonisation legislation listed in Annex I.

2. In addition to the high-risk AI systems referred to in paragraph 1, AI systems referred to in Annex III be considered to be high-risk.

3. By derogation from paragraph 2, an AI system referred to in Annex III shall not be considered to be high-risk where it does not pose a significant risk of harm to the health, safety or fundamental rights of natural persons, including by not materially influencing the outcome of decision making. The first subparagraph shall apply where any of the following conditions is fulfilled:

(a) the AI system is intended to perform a narrow procedural task;
(b) the AI system is intended to improve the result of a previously completed human activity;
(c) the AI system is intended to detect decision-making patterns or deviations from prior decision-making patterns and is not meant to replace or influence the previously completed human assessment, without proper human review; or
(d) the AI system is intended to perform a preparatory task to an assessment relevant for the purposes of the use cases listed in Annex III.

Notwithstanding the first subparagraph, an AI system referred to in Annex III shall always be considered to be high-risk where the AI system performs profiling of natural persons.

4. A provider who considers that an AI system referred to in Annex III is not high-risk shall document its assessment before that system is placed on the market or put into service. Such provider shall be subject to the registration obligation set out in Article 49(2). Upon request of national competent authorities, the provider shall provide the documentation of the assessment.

[5–8 omitted]

Article 17: Quality Management System

1. Providers of high-risk AI systems shall put a quality management system in place that ensures compliance with this Regulation. That system shall be documented in a systematic and orderly manner in the form of written policies, procedures and instructions, and shall include at least the following aspects:

(a) a strategy for regulatory compliance, including compliance with conformity assessment procedures and procedures for the management of modifications to the high-risk AI system;
(b) techniques, procedures and systematic actions to be used for the design, design control and design verification of the high-risk AI system;
(c) techniques, procedures and systematic actions to be used for the development, quality control and quality assurance of the high-risk AI system;
(d) examination, test and validation procedures to be carried out before, during and after the development of the high-risk AI system, and the frequency with which they have to be carried out;
(e) technical specifications, including standards, to be applied and, where the relevant harmonised standards are not applied in full or do not cover all of the relevant requirements set out in Section 2, the means to be used to ensure that the high-risk AI system complies with those requirements;

(f) systems and procedures for data management, including data acquisition, data collection, data analysis, data labelling, data storage, data filtration, data mining, data aggregation, data retention and any other operation regarding the data that is performed before and for the purpose of the placing on the market or the putting into service of high-risk AI systems;
(g) the risk management system referred to in Article 9;
(h) the setting-up, implementation and maintenance of a post-market monitoring system, in accordance with Article 72;
(i) procedures related to the reporting of a serious incident in accordance with Article 73;
(j) the handling of communication with national competent authorities, other relevant authorities, including those providing or supporting the access to data, notified bodies, other operators, customers or other interested parties;
(k) systems and procedures for record-keeping of all relevant documentation and information;
(l) resource management, including security-of-supply related measures;
(m) an accountability framework setting out the responsibilities of the management and other staff with regard to all the aspects listed in this paragraph.

Article 40: Harmonised Standards and Standardisation Deliverables

1. High-risk AI systems or general-purpose AI models which are in conformity with 15standardiz standards or parts thereof the references of which have been published in the Official Journal of the European Union . . . shall be presumed to be in conformity with the requirements set out in Section 2 of this Chapter . . . to the extent that those standards cover those requirements or obligations.

. . .

3. The participants in the standardisation process shall seek to promote investment and innovation in AI, including through increasing legal certainty, as well as the competitiveness and growth of the Union market, to contribute to strengthening global cooperation on standardisation and taking into account existing international standards in the field of AI that are consistent with Union values, fundamental rights and interests, and to enhance multi-stakeholder governance ensuring a balanced representation of interests and the effective participation of all relevant stakeholders in accordance with Articles 5, 6, and 7 of Regulation (EU) No 1025/2012.

Article 43: Conformity assessment

1. For high-risk AI systems listed in point 1 of Annex III, where, in demonstrating the compliance of a high-risk AI system with the requirements set out in Section 2, the provider has applied harmonised standards referred to in Article 40 . . . the provider shall opt for one of the following conformity assessment procedures based on:

(a) the internal control referred to in Annex VI; or
(b) the assessment of the quality management system and the assessment of the technical documentation, with the involvement of a notified body, referred to in Annex VII.

In demonstrating the compliance of a high-risk AI system with the requirements set out in Section 2, the provider shall follow the conformity assessment procedure set out in Annex VII where:

(a) harmonised standards referred to in Article 40 do not exist,
(b) the provider has not applied, or has applied only part of, the harmonised standard;

[(c)–(d) omitted]

2. For high-risk AI systems referred to in points 2 to 8 of Annex III, providers shall follow the conformity assessment procedure based on internal control as referred to in Annex VI, which does not provide for the involvement of a notified body.

3. For high-risk AI systems covered by the Union harmonisation legislation listed in Section A of Annex I, the provider shall follow the relevant conformity assessment procedure as required under those legal acts. The requirements set out in Section 2 of this Chapter shall apply to those high-risk AI systems and shall be part of that assessment.

Annex I: List of Union Harmonisation Legislation

Section A – List of Union harmonisation legislation based on the New Legislative Framework

1. Directive 2006/42/EC of the European Parliament and of the Council of 17 May 2006 on machinery, and amending Directive 95/16/EC (OJ L 157, 9.6.2006, p. 24);

2. Directive 2009/48/EC of the European Parliament and of the Council of 18 June 2009 on the safety of toys (OJ L 170, 30.6.2009, p. 1);

3. Directive 2013/53/EU of the European Parliament and of the Council of 20 November 2013 on recreational craft and personal watercraft and repealing Directive 94/25/EC (OJ L 354, 28.12.2013, p. 90);

4. Directive 2014/33/EU of the European Parliament and of the Council of 26 February 2014 on the harmonisation of the laws of the Member States relating to lifts and safety components for lifts (OJ L 96, 29.3.2014, p. 251);

[5–12 omitted]

Section B – List of other Union harmonisation legislation

[13–20 omitted]

Annex III: High-Risk AI Systems Referred to in Article 6 (2)

High-risk AI systems pursuant to Article 6(2) are the AI systems listed in any of the following areas:

1. Biometrics, in so far as their use is permitted under relevant Union or national law:
    (a) remote biometric identification systems.
    (b) This shall not include AI systems intended to be used for biometric verification the sole purpose of which is to confirm that a specific natural person is the person he or she claims to be;
    (c) AI systems intended to be used for biometric categorisation, according to sensitive or protected attributes or characteristics based on the inference of those attributes or characteristics;
    (d) AI systems intended to be used for emotion recognition.

2. Critical infrastructure: AI systems intended to be used as safety components in the management and operation of critical digital infrastructure, road traffic, or in the supply of water, gas, heating or electricity.

3. Education and vocational training:

    (a) AI systems intended to be used to determine access or admission or to assign natural persons to educational and vocational training institutions at all levels;
    (b) AI systems intended to be used to evaluate learning outcomes, including when those outcomes are used to steer the learning process of natural persons in educational and vocational training institutions at all levels;
    (c) AI systems intended to be used for the purpose of assessing the appropriate level of education that an individual will receive or will be able to access, in the context of or within educational and vocational training institutions at all levels;
    (d) AI systems intended to be used for monitoring and detecting prohibited behaviour of students during tests in the context of or within educational and vocational training institutions at all levels.

4. Employment, workers' management and access to self-employment:

    (a) AI systems intended to be used for the recruitment or selection of natural persons, in particular to place targeted job advertisements, to analyse and filter job applications, and to evaluate candidates;
    (b) AI systems intended to be used to make decisions affecting terms of work-related relationships, the promotion or termination of work-related contractual relationships, to allocate tasks based on individual behaviour or personal traits or characteristics or to monitor and evaluate the performance and behaviour of persons in such relationships.

5. Access to and enjoyment of essential private services and essential public services and benefits:

    (a) AI systems intended to be used by public authorities or on behalf of public authorities to evaluate the eligibility of natural persons for essential public assistance benefits and services, including healthcare services, as well as to grant, reduce, revoke, or reclaim such benefits and services;
    (b) AI systems intended to be used to evaluate the creditworthiness of natural persons or establish their credit score, with the exception of AI systems used for the purpose of detecting financial fraud;
    (c) AI systems intended to be used for risk assessment and pricing in relation to natural persons in the case of life and health insurance;
    (d) AI systems intended to evaluate and classify emergency calls by natural persons or to be used to dispatch . . . emergency first response services, including by police, firefighters and medical aid, as well as of emergency healthcare patient triage systems.

6. Law enforcement, in so far as their use is permitted under relevant Union or national law:

    (a) AI systems intended to be used by or on behalf of law enforcement authorities . . . to assess the risk of a natural person becoming the victim of criminal offences;
    (b) AI systems intended to be used by or on behalf of law enforcement authorities . . . as polygraphs or similar tools;
    (c) AI systems intended to be used by or on behalf of law enforcement authorities . . . to evaluate the reliability of evidence in the course of the investigation or prosecution of criminal offences;

(d) AI systems intended to be used by law enforcement authorities or on their behalf . . . for assessing the risk of a natural person offending or re-offending not solely on the basis of the profiling of natural persons . . . or to assess personality traits and characteristics or past criminal behaviour of natural persons or groups;
(e) AI systems intended to be used by or on behalf of law enforcement authorities . . . for the profiling of natural persons . . . in the course of the detection, investigation or prosecution of criminal offences.

7. Migration, asylum and border control management, in so far as their use is permitted under relevant Union or national law:

(a) AI systems intended to be used by or on behalf of competent public authorities or by Union institutions, bodies, offices or agencies as polygraphs or similar tools;
(b) AI systems intended to be used by or on behalf of competent public authorities or by Union institutions, bodies, offices or agencies to assess a risk, including a security risk, a risk of irregular migration, or a health risk, posed by a natural person who intends to enter or who has entered into the territory of a Member State;
(c) AI systems intended to be used by or on behalf of competent public authorities or by Union institutions, bodies, offices or agencies to assist competent public authorities for the examination of applications for asylum, visa or residence permits and for associated complaints with regard to the eligibility of the natural persons applying for a status, including related assessments of the reliability of evidence;
(d) AI systems intended to be used by or on behalf of competent public authorities, or by Union institutions, bodies, offices or agencies, in the context of migration, asylum or border control management, for the purpose of detecting, recognising or identifying natural persons, with the exception of the verification of travel documents.

8. Administration of justice and democratic processes:

(a) AI systems intended to be used by a judicial authority or on their behalf to assist a judicial authority in researching and interpreting facts and the law and in applying the law to a concrete set of facts, or to be used in a similar way in alternative dispute resolution;
(b) AI systems intended to be used for influencing the outcome of an election or referendum or the voting behaviour of natural persons in the exercise of their vote in elections or referenda. This does not include AI systems to the output of which natural persons are not directly exposed, such as tools used to organise, optimise or structure political campaigns from an administrative or logistical point of view.

Annex VI: Conformity Assessment Procedure Based on Internal Control

1. The conformity assessment procedure based on internal control is the conformity assessment procedure based on points 2, 3 and 4.

2. The provider verifies that the established quality management system is in compliance with the requirements of Article 17.

3. The provider examines the information contained in the technical documentation in order to assess the compliance of the AI system with the relevant essential requirements set out in Chapter III, Section 2.

4. The provider also verifies that the design and development process of the AI system and its post-market monitoring as referred to in Article 72 is consistent with the technical documentation.

Annex VII: Conformity Based on an Assessment of the Quality Management System and an assessment of the technical documentation

2. Overview

The approved quality management system for the design, development and testing of AI systems pursuant to Article 17 shall be examined in accordance with point 3 and shall be subject to surveillance as specified in point 5. The technical documentation of the AI system shall be examined in accordance with point 4.

3. Quality management system

3.2 The quality management system shall be assessed by the notified body, which shall determine whether it satisfies the requirements referred to in Article 17.

The decision shall be notified to the provider or its authorised representative.

The notification shall contain the conclusions of the assessment of the quality management system and the reasoned assessment decision.

3.3 The quality management system as approved shall continue to be implemented and maintained by the provider so that it remains adequate and efficient.

3.4 Any intended change to the approved quality management system or the list of AI systems covered by the latter shall be brought to the attention of the notified body by the provider.

The proposed changes shall be examined by the notified body, which shall decide whether the modified quality management system continues to satisfy the requirements referred to in point 3.2 or whether a reassessment is necessary.

The notified body shall notify the provider of its decision. The notification shall contain the conclusions of the examination of the changes and the reasoned assessment decision.

4. Control of the technical documentation.

4.1 In addition to the application referred to in point 3, an application with a notified body of their choice shall be lodged by the provider for the assessment of the technical documentation relating to the AI system which the provider intends to place on the market or put into service and which is covered by the quality management system referred to under point 3.

4.4 In examining the technical documentation, the notified body may require that the provider supply further evidence or carry out further tests so as to enable a proper assessment of the conformity of the AI system with the requirements set out in Chapter III, Section 2. Where the notified body is not satisfied with the tests carried out by the provider, the notified body shall itself directly carry out adequate tests, as appropriate.

4.6 The decision of the notified body shall be notified to the provider or its authorised representative. The notification shall contain the conclusions of the assessment of the technical documentation and the reasoned assessment decision.

Where the AI system is in conformity with the requirements set out in Chapter III, Section 2, the notified body shall issue a Union technical documentation assessment certificate.

Where the AI system is not in conformity with the requirements set out in Chapter III, Section 2, the notified body shall refuse to issue a Union technical documentation assessment certificate and shall inform the applicant accordingly, giving detailed reasons for its refusal.

4.7 Any change to the AI system that could affect the compliance of the AI system with the requirements or its intended purpose shall be assessed by the notified body which issued the Union technical documentation assessment certificate.

5. Surveillance of the approved quality management system.

5.3 The notified body shall carry out periodic audits to make sure that the provider maintains and applies the quality management system and shall provide the provider with an audit report. In the context of those audits, the notified body may carry out additional tests of the AI systems for which a Union technical documentation assessment certificate was issued.

Blind ID Number: _____

SCRATCHPAPER

Blind ID Number: _____

SCRATCHPAPER