

OECD AI Transparency Report

Organization: Ubie (JPN) Reporting Period: Q2 2026 Published: October 2, 2026

Section 0 – Your role in the AI system lifecycle

Responding as (please select all that apply):

- Application developer (/provider)

If any of the fields above are selected, please indicate your organization size:

Medium-sized enterprise: 50-249 employees

Section 1 – Risk Identification and Evaluation

1. How does your organization define and/or classify different types of risks related to AI?¹

- Maintain an internal taxonomy, classification system, and/or set of risk area definitions
- Align with specific external frameworks, standards, reports, and/or guidance documents such as taxonomies, classification systems, and/or sets of definitions for risk areas (e.g. the NIST AI Risk Management Framework, the ISO/IEC 42001, the OECD AI Principles, the EU General-Purpose AI Code of Practice, the International AI Safety Report). (please specify): METI/MIC AI Guidelines for Business (Ver. 1.1) and its Appendix; Hiroshima AI Process (HAIP) International Code of Conduct; OECD AI Principles; NIST AI Risk Management Framework; Japan AI Safety Institute (AIS) Guide to Evaluation Perspectives on AI Safety, and its healthcare edition, the Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector (published April 2026), which Ubie co-developed as leader of the AIS Healthcare Sub-Working Group on behalf of The Japan Digital Health Alliance (JaDHA); JaDHA Guide to the Utilization of Generative AI for Healthcare Providers; the Two Guidelines from Three Ministries (the MHLW Guidelines for the Safety Management of Medical Information Systems and the METI/MIC Safety Management Guidelines for Information System and Service Providers Handling Medical Information); ISO/IEC 27001:2022 (certified). Alignment with ISO/IEC 42001 is under consideration.
- Consult with external stakeholders (e.g., from academia, industry groups, civil society, or standard development organisations) when developing or updating risk taxonomies

How does your organization define thresholds at which severe risks, unless adequately mitigated, may be considered "unreasonable"? Please describe the framework(s) most relevant to your organization's role and context. If your organisation does not use formal thresholds, you may describe instead the policies, decision criteria, or governance processes you apply when assessing severe AI risks. You may also describe:

- Specific examples of contexts in which your models or systems could, in your assessment, pose unreasonable or intolerable risk;
- How mitigations are graduated by capability or

propensity level; • How your organisation considers the specificities of different AI technologies, such as generative AI and agentic AI, including any specific risk assessment or governance mechanisms applied to agentic AI capabilities.

Ubie does not use numeric capability thresholds, as we are not a foundation model developer. We apply decision criteria and governance processes instead.

Decision criteria. Severity is determined on two axes: imminence and severity of harm. Inputs are the use case, the users, the sensitivity of the data, the influence of outputs on clinical or health decisions, the degree of AI autonomy, whether human confirmation is in place, and whether an external foundation model is used. Because health and medical information can influence user behaviour, we treat three points as non-negotiable: our services provide general information; they do not constitute diagnosis or treatment; and they complement rather than replace human judgment.

How mitigations are graduated. Mitigations escalate with the assessed level of risk: (1) baseline - prompt and guardrail design, input and output filtering, disclaimers, terms of use; (2) elevated - mandatory medical review by physicians before release, sample evaluation of generative outputs, adversarial testing, and restriction of the intended purpose; (3) high - limitation of the scope of provision, access rights restricted so that highly confidential data is accessible only to the medical institution concerned, mandatory human confirmation with the original source always retrievable, and additional operational procedures and customer training; (4) unacceptable - postponement or non-release, or suspension of a feature already in service.

Generative AI and agentic AI. For generative AI features we assume that outputs of external foundation models may contain bias or hallucination, and we design for verification rather than trust: the original text is displayed by default with the AI-generated summary shown alongside (our "safety net structure"), sources are cited, and inputs and creation history are logged for ex post verification. Agentic AI systems that plan or act to a partial extent are under evaluation. Because an error in the medical domain can be irreversible, our governance for agentic capabilities requires an explicit boundary of autonomous action, a human decision point before any action affecting a patient or a medical record, transparency of the decision logic, and evaluation against the AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector before any phased introduction. Evaluation perspectives for agentic AI are being organised through our research organization Ubie Lab and the AISI Healthcare Sub-Working Group. Governance. Risk assessment is performed by dedicated specialists independent of the development and business units in order to ensure objectivity. A two-layer structure - the Generative AI Promotion Team on the promotion side and the Risk and Compliance Committee (legal, security, public relations and other functions) on the risk side - shares information and deliberates trade-offs, and the Trust & Safety Team performs cross-product quality management. This structure is featured as an advanced case (Column 11) in the Appendix to the METI/MIC AI Guidelines for Business (Ver. 1.1).

2. Both before and after deployment, what practices does your organization use to identify and evaluate vulnerabilities, incidents, emerging risks and misuse throughout the AI system lifecycle?

- Adversarial testing to detect technical vulnerabilities
- Internal benchmarks, standard metrics (e.g., frequency, severity scores)
- Qualitative assessments (e.g., expert judgment, narrative scoring)
- Specific risk identification practices for vulnerable groups

- Impact assessments such as personal data protection assessment and fundamental rights impact assessments
- Internal model, system, or application-independent research initiatives
- External model, system, or application-independent research initiatives
- Vulnerability reports
- Incident reports, including reports shared by other organizations
- Automated monitoring / remote status reporting (e.g., telemetry pipelines, automated alerts when risk thresholds are breached)
- Multilingual and cross-cultural evaluation, including collaboration with local ecosystems (e.g., national AI institutes, language-specific evaluation consortia) for under-represented languages and cultural contexts (please specify): Japanese is an under-represented language in medical LLM evaluation. Ubie collaborates with the Japan AI Safety Institute (AIS), Japan's national AI safety institute, as leader of the Healthcare Sub-Working Group under its Business Demonstration Working Group, on behalf of The Japan Digital Health Alliance (JaDHA). In that capacity we co-developed the Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector (published April 2026), which includes robustness evaluation for Japanese dialects, abbreviations and non-standard medical terminology. In connection with our U.S. expansion, we conducted performance evaluation using large-scale patient data in the Mayo Clinic Platform_Accelerate program, covering a different language and clinical-cultural context.
- Internal red teaming³

Please briefly describe your organization's decision-making process for deploying an AI model or system including: who has decision authority; what evidence and criteria are required (e.g., performance, safety/security, robustness, fairness, privacy, compliance, independent external testing, evaluation, validation and verification (TEVV) or assurance); contextual caveats or limitations required in interpreting risk evaluation metric results; what approvals/sign-offs are needed; and what safeguards (e.g., mitigations, monitoring, or phased release) must be in place to move into deployment.

Decision authority. Release of an AI feature is decided by the responsible product team on the basis of a cross-product review by the Trust & Safety Team. Risk assessment is performed by dedicated specialists independent of the development and business units in order to ensure objectivity. Where a material concern remains, the decision escalates to the Risk and Compliance Committee (legal, security, public relations and other functions), which deliberates the trade-off between AI utilization and risk together with the Generative AI Promotion Team. This two-layer structure is featured as an advanced case (Column 11) in the Appendix to the METI/MIC AI Guidelines for Business (Ver. 1.1). Evidence and criteria required. Review and approval gates are applied at each stage. At the planning stage we assess the use case, the users, data sensitivity, and impacts on medical safety, privacy and security. Before release, the evidence required against defined test criteria comprises: functional testing; medical review by physicians of clinical accuracy; sample evaluation of generative AI outputs, including hallucination and source attribution; security assessment, including prompt injection; privacy review covering the scope of data sent to external foundation models and retention in logs; legal and regulatory compliance review (Pharmaceuticals and Medical Devices Act, Medical Practitioners Act, Act on the Protection of Personal Information, and the Two Guidelines from Three Ministries); and fairness review across disease, age, sex, expression of symptoms, and region or access to healthcare. Independent external assurance is used for information security, through ISMS (ISO/IEC 27001:2022)

audit and security assessment. Contextual caveats in interpreting risk evaluation metric results. Evaluation results are read as conditional rather than absolute. Sample evaluation of generative outputs is not exhaustive, so a passing result is not evidence of the absence of hallucination. Benchmark and test-set performance does not transfer directly to real clinical language, which includes dialects, abbreviations and non-standard medical terminology. Because our products are built on external foundation models, results are valid only for the model and version tested, and are re-examined when the model, version or data integration method changes. For these reasons we do not treat any single metric as a sole release criterion, and human confirmation is retained as the controlling safeguard. Approvals and sign-offs. Confirmation of conformity with our quality standards by the Trust & Safety Team is required before release. Deployment at a medical institution additionally requires confirmation of that customer's security and operational requirements under contract. A fresh security review is triggered by material changes to external foundation models, cloud services or data integration methods, and by amendments to applicable laws or guidelines. Safeguards required to move into deployment. Prompt and guardrail design with input and output filtering; a human confirmation step for any output that materially affects a care decision; the "safety net structure", in which the original text is displayed by default and can always be referenced; access rights restricted so that highly confidential data of a medical institution is accessible only to that institution; cautionary notices in wording and UI; operating procedures and customer training; and continuous log and error monitoring after release. Where risk warrants it, we limit the scope of provision or release in phases. Where material risk remains, we postpone or do not release.

Following deployment, how do you use incident reports — including those shared by other organizations — to identify emerging issues, update risk assessments, and trigger corrective actions.

Yes. Internally, incidents that occur or are reported, inquiries, and feedback from medical institutions are managed centrally in an incident management tool. For each we record the type, scope of impact, root cause, detection channel and recurrence-prevention measures, and we then check whether the same type of risk exists in other products or other use cases. Root cause analysis and recurrence-prevention measures are documented as post-mortems.

Externally, we continuously refer to the checklists and case examples in the AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector and the JaDHA Guide to the Utilization of Generative AI for Healthcare Providers, to publicly available information on medical information systems, personal data protection and cybersecurity, and to incident cases reported by other companies and other industries - including cases outside healthcare where the failure mode is transferable.

Updating risk assessments. Findings feed back into our internal risk taxonomy and its categories, into the test criteria applied before release (new test items are added), and into internal education. Corrective actions triggered include revision of prompts, guardrails, UI and explanatory text; additional output filters; updates to our medical information database; security patches; review of permission settings; restriction or suspension of a feature; and revision of operational guides, templates and customer training. Where notification is required by law or contract, we notify affected customer medical institutions and users, and report to public authorities such as the Personal Information Protection Commission. Knowledge is shared with the industry through the AISI Healthcare Sub-Working Group, JaDHA, and the AI Governance Association (AIGA).

3. Does your organization have mechanisms to receive and/or share reports of risks, incidents or vulnerabilities by/with third parties?

Yes, both receiving and sharing

If yes, please indicate which stakeholders reporting mechanisms may be accessible to, recognising specifics will depend on the nature of the report and risk level. For each, indicate whether access is on a standing basis (Always), activated based on context (Ad hoc/as relevant), or not currently accessible.

	Always	Ad hoc / as relevant	Not currently accessible
Relevant employees (security, legal/compliance, etc.)	☐		
Public authorities (e.g., public reporting offices for security incidents or data breaches)		☐	
Supervisory authorities (e.g., data protection authorities)		☐	
Industry partners or industry fora	☐		
Users (end users, enterprise customers, developers using APIs)	☐		
Business customers / deploying organizations (if different from “users”)	☐		
Supply-chain partners (cloud, data providers, etc.)	☐		
Academia / researchers		☐	
Independent auditors		☐	
Standards bodies / multi-stakeholder initiatives (OECD.AI, ISO/IEC, IEEE, etc.)		☐	
Specific civil society organizations		☐	
Media / general public (people or communities indirectly impacted)		☐	
Other: <i>Any member of the public, through our publicly open channels (the Vulnerability Coordination and Disclosure Policy and our consumer inquiry desk)</i>		☐	

Please describe how your organisation determines when and how reporting mechanisms are made accessible to different stakeholders (e.g. based on legal or regulatory requirements, internal policies, or risk level). Please also provide public URLs or details on how to submit reports. If none, write “Not applicable”

Accessibility of our reporting mechanisms is determined by three factors: legal and regulatory requirements, internal policy, and the level of risk involved.

Publicly open channels, available to anyone without a prior relationship. (1) The Ubie Vulnerability Coordination and Disclosure Policy (VCDP) sets out how external security researchers and users may submit vulnerability reports, and the coordinated disclosure process we follow: <https://ubie.life/safety/vcdp> (2) A general inquiry channel for consumer users of our services. (3) The contact point stated in our privacy policies, for questions and requests concerning personal data: <https://ubie.app/policy> (consumers) and <https://intro.dr-ubie.com/policy> (medical institutions). Our overall approach to safety and security, including these channels, is published at <https://ubie.life/safety>

Channels restricted to contracted parties. Customer medical institutions report through a dedicated support desk and through the operational communication channels agreed in their contract. Access is restricted because these channels may carry patient-related and other confidential information, and are therefore governed by the security and confidentiality terms of the contract rather than being publicly open.

Internal channels. Employees and contractors report information security incidents to the Information Security Manager without delay under our Information Security Policy, and incidents are recorded and tracked in an incident management tool used across all products.

Statutory reporting. Reporting to public authorities, including the Personal Information Protection Commission, follows the requirements of applicable law rather than our discretion, and is triggered by the nature and scope of the incident.

Industry sharing. Risk and vulnerability knowledge is shared with the industry through the AISI Healthcare Sub-Working Group, The Japan Digital Health Alliance (JaDHA) and the AI Governance Association (AIGA), within the scope permitted by confidentiality obligations.

How does your organization share relevant information as appropriate?

- Publicly
- Information-sharing partnerships (with confidentiality agreements as appropriate)
- Reporting to regulatory authorities as relevant

Please describe what type of information is shared.

Publicly shared. The intended uses and limitations of our services, including that they are not medical devices, are not intended for diagnosis or treatment, and should not be used in emergencies; our data sources and physician supervision process (<https://ubie.app/about>); our privacy policies and our approach to data protection and privacy management (https://ubie.life/safety/data_protection, https://ubie.life/safety/privacy_policy_management); our Information Security Policy; our Vulnerability Coordination and Disclosure Policy and the coordinated disclosure process it sets out (<https://ubie.life/safety/vcdp>); and case studies and measured outcomes of deployment at medical institutions.

Shared individually with customer medical institutions and pharmaceutical companies. Deployment materials and operational guidance; the premise that AI outputs are reviewed and edited by healthcare professionals, and which tasks the technology is appropriate for; security requirements and conformity with the Two Guidelines from Three Ministries; confirmation that input data is not used for open model training; and notification of incidents that affect them, including scope of impact and corrective action.

Shared with industry bodies and the Japan AI Safety Institute. Evaluation perspectives and practical checklists; risk categories specific to the healthcare sector, such as hallucination, off-label use of non-medical-device products, privacy and security; and practical mitigation measures and operational

lessons learned. This is the substance of the Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector and the JaDHA Guide to the Utilization of Generative AI for Healthcare Providers. Reported to regulatory authorities. Incident information to the extent required by applicable law, including reports to the Personal Information Protection Commission.

4. If your organization has programs to incentivize risk(s) disclosure, incidents and systemic vulnerabilities, please share what type of programs you use?

- Internal incentives and culture

Please provide public URLs and/or details on your incentive program(s). If none, write “Not applicable.”

Our incentive program is internal and is not published, so there is no public URL.

Ubic promotes a culture in which employees and contractors are encouraged to report risks, potential vulnerabilities and incidents responsibly. Reports are used for organisational learning rather than for attributing fault: root cause analysis and recurrence-prevention measures are documented as post-mortems and fed back into the test criteria applied before release, into internal education, and into policies and operating procedures. Our Information Security Policy requires employees and contractors to report information security incidents to the Information Security Manager without delay, and incidents are recorded and tracked in an incident management tool used across all products. We also operate an internal award program that recognises appropriate and responsible use of AI (the Generative AI Utilization Award).

For reports originating outside the company, our Vulnerability Coordination and Disclosure Policy is published at <https://ubic.life/safety/vcdp>. This is a disclosure channel rather than an incentive program. No monetary rewards (bug bounty) are currently offered.

5. Does your organization use or contribute to the development of international technical standards, best practices or tools for the identification, assessment, and evaluation of risks?

Contributes and uses

If contributes or uses, please select all that apply:

- ISO/IEC standards (e.g., ISO/IEC 27001, ISO/IEC 42001, ISO/IEC 23894, ISO/IEC 42005)
- NIST AI Risk Management Framework (AI RMF)
- Industry-led best practices or codes of conduct
- Other international or regional standards (please specify): OECD AI Principles; the METI/MIC AI Guidelines for Business (Ver. 1.1), in whose Appendix our AI governance structure is featured as an advanced case (Column 11); the Japan AI Safety Institute (AIS) Guide to Evaluation Perspectives on AI Safety and its healthcare edition, the Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector, which Ubie led the development of as leader of the AISI Healthcare Sub-Working Group on behalf of The Japan Digital Health Alliance (JaDHA); and the Two Guidelines from Three Ministries (the MHLW Guidelines for the Safety Management of Medical Information Systems and the METI/MIC Safety Management Guidelines for Information System and Service Providers Handling Medical Information), which govern the handling of medical information in Japan.

6. Does your organization collaborate with relevant stakeholders across sectors to assess and adopt risk mitigation measures to address risks that could have broad and potentially cascading socio-economic impacts? Please indicate which stakeholders you collaborate with, and whether collaboration is on a standing basis (Always-on), activated based on context (Ad hoc/as relevant), or there is no current collaboration.

	Always-on	Ad hoc / as relevant	No current collaboration
Relevant employees (security, legal/compliance, etc.)	☐		
Public authorities / regulators / AI safety institutes / other public sector bodies and inter-governmental coordination structures	☐		
Industry partners or industry fora	☐		
Users (end users, enterprise customers, developers using APIs)	☐		
Business customers / deploying organizations (if different from "users")	☐		
Supply-chain partners (cloud, data providers, etc.)	☐		
Academia / researchers		☐	
Independent auditors		☐	
Standards bodies / multi-stakeholder initiatives (OECD.AI; ISO/IEC; IEEE; AIID etc.)		☐	
Independent civil society organizations (e.g., for evaluations, audits, or policy review)		☐	
Media / general public (people or communities indirectly impacted)		☐	
Other(s) (please add any other relevant stakeholder groups):: <i>Healthcare professionals and supervising physicians (in-house and external medical reviewers), and medical societies</i>	☐		
Not applicable			

Please describe examples of how your organization collaborates. Examples may include: participating in multi-stakeholder or industry initiatives; partnering with academia, civil society, or the public sector; establishing internal or joint governance structures with external experts or in partnership with other efforts to share knowledge and jointly address risks.

Ubic collaborates to address risks in medical AI - medical safety, privacy, security and public trust - which can have broad and cascading effects on access to care and on public confidence in healthcare. Multi-stakeholder and industry initiatives. As leader of the Healthcare Sub-Working Group under the AISI Business Demonstration Working Group, on behalf of The Japan Digital Health Alliance (JaDHA),

Ubie led the development of the Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector, published in April 2026 jointly with the Japan AI Safety Institute. Ubie also led the development of the JaDHA Guide to the Utilization of Generative AI for Healthcare Providers, is a regular member of the AI Governance Association (AIGA), a director of the Japan Medical Venture Association (JMVA), and a member of the Japan Trust & Safety Association (JTSA).

Partnership with the public sector and academia. We have presented at IPA/AISI Business Demonstration Working Group briefing sessions and spoken at the Hiroshima Global Forum for Trustworthy AI. With university hospitals and acute care hospitals we run demonstration projects and clinical research that evaluate how humans should verify generative AI outputs, which tasks the technology should be applied to, and what cautionary notices and training are required.

Joint governance and technical partnerships. Our internal governance combines the Generative AI Promotion Team, the Risk and Compliance Committee, and the Trust & Safety Team, with medical review by in-house and external supervising physicians built into the content pipeline. With technology partners we conduct safety and performance evaluation - our work with Google Cloud/Gemini, and performance evaluation using large-scale patient data in the Mayo Clinic Platform_Accelerate program.

Scope of our systemic risk management. Because Ubie is not a large-scale developer of foundation models, we manage broad socio-economic risk primarily from the position of a user of, and a service provider built on, external foundation models. We focus on sector-specific risks: misuse in the medical field, over-reliance on AI, privacy infringement, and disparities in access to healthcare.

Any further comments and implementation documentation.

Role and scope of this response. Ubie responds as a company that provides AI-based services in the medical and healthcare domain (an AI Provider), not as a developer of foundation models. The scope covers Ubie (consumer-facing), and AI Health Assistant and Ubie Generative AI for medical institutions, among others. Our products are non-medical-device software (Non-SaMD).

Standards contributed to and used. Contributed: as leader of the AISI Healthcare Sub-Working Group on behalf of The Japan Digital Health Alliance (JaDHA), Ubie led the development of the Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector (published April 2026, jointly with the Japan AI Safety Institute), which provides practical checklists and evaluation perspectives to the industry and is published in part in Markdown so that it can be loaded into AI services. Ubie also led the development of the JaDHA Guide to the Utilization of Generative AI for Healthcare Providers. Our AI governance structure is featured as an advanced case (Column 11) in the Appendix to the METI/MIC AI Guidelines for Business (Ver. 1.1). Used: the AI Guidelines for Business; the Hiroshima AI Process International Code of Conduct; the OECD AI Principles; the NIST AI Risk Management Framework; the Two Guidelines from Three Ministries; and ISO/IEC 27001:2022, for which we hold and operate certification. Alignment with ISO/IEC 42001 is under consideration.

Implementation documentation.

- Commitment to safety and security: <https://ubie.life/safety>
- Data protection: https://ubie.life/safety/data_protection
- Privacy management structure: https://ubie.life/safety/privacy_policy_management
- Vulnerability Coordination and Disclosure Policy (VCDP): <https://ubie.life/safety/vcdp>
- AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector (Japanese and English editions): https://aisi.go.jp/output/output_information/260402/

- Appendix to the AI Guidelines for Business (Ver. 1.1), Column 11:
https://www.meti.go.jp/shingikai/mono_info_service/ai_shakai_jisso/pdf/20250328_3.pdf
- The Japan Digital Health Alliance (JaDHA): <https://jadha.jp/>
- Data sources and physician supervision process for the consumer-facing service: <https://ubie.app/about>

Section 2 – Risk Management and Information Security

7. Across the AI system lifecycle, what steps does your organization take to manage risks and vulnerabilities? Please indicate which specific practices, tools or metrics your organization uses at each lifecycle stage. You may refer to tools or metrics present in the OECD.AI Catalogue of Tools and Metrics or suggest additions. Respondents may complete only the lifecycle stages relevant to their organisation's role. Stages that do not apply may be left blank.

Plan and design

Collect and process data

Build, integrate and interpret model

Verify and validate

Deploy

Operate and monitor

Please provide a consolidated description of your organisation's approach to risk management across the AI system lifecycle, including how the practices identified above operate together and any cross-cutting governance, review or escalation processes. Please also indicate how often these practices are reviewed and what types of events trigger reviews and updates. Where relevant, please distinguish between different types of AI systems (e.g. generative or agentic AI).

Ubie integrates risk management across the AI lifecycle through three layers that operate together. (1) Two-layer governance covering promotion and risk. The Generative AI Promotion Team pursues value from AI utilization, while the Risk and Compliance Committee (legal, security, public relations and other functions) examines risk. The two share information and deliberate trade-offs, so that promotion and restraint are weighed in the same forum rather than in sequence. Risk assessment itself is performed by dedicated specialists independent of the development and business units, which separates the assessment from the incentive to ship. This structure is featured as an advanced case (Column 11) in the Appendix to the METI/MIC AI Guidelines for Business (Ver. 1.1). (2) Cross-product quality management by the Trust & Safety Team, which reviews releases across all products, collects and classifies incidents, and distinguishes issues common to several products from issues specific to one. (3) Policies and guidelines that bind the above into repeatable practice: the Information Security Policy, the Privacy Policy, our quality guidelines, and the Ubie Trustworthy AI Policy. Escalation. Where a material concern remains after assessment, the specialists escalate to the Risk and Compliance Committee, which may decide to postpone release, limit the scope of provision, restrict the intended purpose, or require additional safeguards. Information security incidents are escalated to the

Information Security Manager without delay under the Information Security Policy. Material incidents are documented as post-mortems and fed back into the test criteria applied before release. Review frequency and triggers. Practices are reviewed on a regular cycle, and additionally as needed when: a new AI or generative AI feature is introduced; an external foundation model, cloud service or data integration method changes materially; a customer medical institution's requirements are confirmed for a new deployment; a material vulnerability, incident or item of external threat intelligence is confirmed; laws or guidelines are amended, for example the AI Guidelines for Business, the AI Promotion Act, the Act on the Protection of Personal Information, the Two Guidelines from Three Ministries, the Pharmaceuticals and Medical Devices Act or the Medical Practitioners Act; or an issue is raised in an audit or by a customer. How test results feed back. Defects and risks detected are prioritised by severity and reflected in prompt revisions, additional output filters, and review of the intended scope of use. After remediation we re-test and check whether the same weakness exists in other features or other products. Distinguishing types of AI system. Generative AI features are managed on the assumption that outputs of external foundation models may contain bias or hallucination, so controls centre on verification: retrieval grounding with source citation, output filtering, and a human confirmation step, with the original text displayed by default and always retrievable (our "safety net structure"). Agentic AI systems that plan or act to a partial extent are under evaluation. Both categories are managed on the premise of a human-in-the-loop, in which a healthcare professional or a user makes the final decision. Because an error in the medical domain can be irreversible, any future agentic capability would additionally require an explicit boundary of autonomous action and a human decision point before any action affecting a patient or a medical record.

Please describe any measures used to ensure data quality and fairness across the data collection and curation stage (e.g., dataset audits, filtering training data to remove harmful content or personal information, fine-tuning on curated data, use of moderation tools, assessment of limitations to model generalisability).

Ubie is not in a position to directly manage the pre-training data of foundation models. What we manage is the quality of the medical information, evaluation data and retrieval targets that our products use. Data quality. In our consumer-facing service we restrict sources to publicly known, reliable material - more than 50,000 medical papers and various clinical practice guidelines - and we have built a medical knowledge database covering more than 1,100 disease names. Medical information is created by in-house physicians and checked by supervising physicians, and is updated as needed in light of the latest clinical guidelines and standards of care. Each item is traceable to a cited source.

Fairness. To prevent inappropriate guidance arising from differences in disease, age, sex, expression of symptoms, or region and access to healthcare, we improve our services continuously through medical review, user feedback, feedback from customer medical institutions, and checks of evaluation data.

Assessment of limitations to generalisability. We treat instability of output quality across diverse inputs as an explicit robustness concern, in particular Japanese dialects, abbreviations and non-standard medical terminology. Benchmark or test-set performance is not read as transferring directly to real clinical language, and evaluation results are treated as valid only for the model and version tested.

Personal information. Our medical knowledge database is built from publicly known sources and does not contain patient personal data. Testing and evaluation use dummy data or anonymized and de-identified data, and personally identifiable information sent to external large language models is minimised.

Generative AI features. On the premise that outputs of external foundation models may contain bias or hallucination, we apply confirmation by healthcare professionals, prompt design, limitation of the intended purpose, cautionary notices on outputs, and feedback loops.

8. What security measures are used in testing AI systems?

- Isolated and secure test environments
- Isolation of staging environments (from production / operational environments)
- Controlled environments for red-teaming and adversarial testing
- Containment measures for prompt-injection and jailbreak attempts
- Controlled security vulnerability testing
- Human-in-the-loop controls for higher-risk test cases
- Access controls, including roles-based permissions
- Audit logging and monitoring
- Use of “AI-safe” test-data, as appropriate
- Use of synthetic, anonymized, or dummy data, as appropriate
- Other(s) (please specify): Testing is conducted on cloud infrastructure compliant with the Two Guidelines from Three Ministries, which govern the handling of medical information in Japan, and under our ISMS (ISO/IEC 27001:2022). For generative AI features we verify, by configuration and by contract, that input data is not used for open model training, and we check the scope of data transmitted to external APIs, unnecessary retention in logs, and the boundaries of integrations with external services. Test data is restricted to dummy data, anonymized or de-identified data, or data whose use is permitted under contract, consent and access controls.

8.A. Please indicate, if relevant, the team or role responsible for overseeing secure testing and whether this responsibility is shared across multiple teams and/or functions.

Responsibility is shared across several teams and functions rather than held by a single team.

Product teams. Each product team performs QA and functional testing for its own features, and designs the prompts, guardrails and output filtering that are then tested.

Trust & Safety Team. Performs cross-product review before release and confirms conformity with our quality standards. It also collects and classifies incidents across all products, so that findings from testing on one product are checked against the others.

Dedicated risk assessment specialists. Independent of the development and business units, they carry out risk assessment so that the assessment is separate from the incentive to ship, and they escalate to the Risk and Compliance Committee where a material concern remains.

Information security function. Under our Information Security Policy and ISMS (ISO/IEC 27001:2022), security reviews and vulnerability assessment are conducted when AI features are introduced or materially changed, and the Information Security Manager receives escalation of information security incidents.

Supervising physicians. In-house and external physicians perform medical review of clinical accuracy. This is treated as part of pre-release testing rather than as a separate content step.

Risk and Compliance Committee. Legal, security, public relations and other functions deliberate residual risk and may require additional safeguards or postpone release.

Because our products handle medical information, oversight of testing deliberately combines a security function and a clinical function. Neither alone is sufficient to clear a release.

9. How does your organization protect intellectual property, including copyright-protected content?

9.A. Safeguards that promote lawful data use and transparency in AI training:

- Publishing “model cards” or contextual information on training data and model limitations
- Apply provenance controls to guard against training data being sourced from websites or repositories that have been determined to engage in large-scale piracy
- Refrain from the unlawful circumvention of technological protection measures in data sourcing
- Other(s) (please specify): Ubie does not perform pre-training of foundation models. Our products are built on pre-trained models accessed via API, so most safeguards concerning training data rest with the foundation model providers, for example OpenAI and Google, and are governed by their published policies. At our application layer: the medical information in our own knowledge database is authored in-house by physicians and checked by supervising physicians, on the basis of publicly known, reliable sources such as clinical practice guidelines, clinical protocols and medical papers, with each item traceable to a cited source. We publish our approach to data sources and our supervision process at <https://ubie.app/about> When adopting an external foundation model or cloud service, we verify the contract terms, the terms of service, the scope of permitted data use, whether input data is used for model training, and the treatment of generated outputs. We confirm by configuration and by contract that customer and user input is not used for open model training.

9.B. Safeguards that reduce the risk of dissemination of outputs that may infringe copyright:

- Restricting outputs that raise ethical, human rights, or legal concerns—such as non-consensual intimate imagery, unauthorized deepfakes, or outputs that are substantially similar to copyrighted works.
- Designating clear points of contact so creators or rights holders can raise concerns about infringing outputs
- Other(s) (specify): The outputs of our services consist primarily of medical information content created in-house under physician supervision, and summaries and drafts of business documents. We do not perform general-purpose content generation such as images or music, so filtering against a corpus of known infringing works is not a principal area of control for us. For our services for medical institutions, summaries and drafts are generated from documents that the institution itself has entered, rather than from third-party works, and outputs are produced on the premise that a healthcare professional reviews and edits them. We have review processes for output content so that outputs do not improperly reproduce third-party copyrighted works. Our terms of service prohibit unauthorized copying or reproduction of content, automated bulk collection, and reverse engineering, and we publish a contact point for inquiries concerning rights.

10. How do you protect privacy and prevent your systems from disclosing confidential or sensitive data across the AI lifecycle?

- Use security controls, including encryption at rest and in transit

- Make information about data privacy, practices, and controls accessible to a diverse set of stakeholders
- Other(s) (specify): Ubie handles sensitive personal information, including medical and health information, and positions privacy protection as one of its highest priorities. With respect to sensitive personal information, our policy is not to collect it or to provide it to third parties without the individual's consent, except where permitted by law. Medical information is managed in accordance with the Two Guidelines from Three Ministries, the Japanese guidelines governing the safety management of medical information. Detailed measures are described in the following field.

Describe key privacy measures and data-protection standards applied.

Technical measures. Personally identifiable information is stored separately in a dedicated store accessible only to a limited number of staff. We apply encryption at rest and in transit, least-privilege access control, and log monitoring (https://ubie.life/safety/data_protection).

Management of medical information. Safety management is compliant with the Two Guidelines from Three Ministries. For highly confidential data of a medical institution, we adopt a design in which only that institution holds access rights.

When using generative AI. We verify settings and contracts ensuring that input data is not used for open model training; minimise personally identifiable information sent to external large language models; prevent unnecessary retention in logs; and provide explanation and consent handling where patient information is involved.

Organisational measures. Access management, restrictions on taking data outside the company, supervision of subcontractors, information security education, and incident reporting with root cause analysis and recurrence prevention. Our privacy management structure is published at https://ubie.life/safety/privacy_policy_management

Protection of users' rights. Clear statement of purposes of use; provision of a path to withdraw consent; deletion of data through account deletion; and procedures for disclosure, correction and suspension of use. Transfers of data overseas, to the United States, are carried out with appropriate protective measures in accordance with applicable laws.

Standards applied. ISO/IEC 27001:2022 (ISMS certification held and operated); the Two Guidelines from Three Ministries for the safety management of medical information; and the Act on the Protection of Personal Information, under which medical and health information constitutes sensitive personal information.

11. How does your organization assess cybersecurity risks and implement policies to enhance the cybersecurity of advanced AI systems across the AI system lifecycle? Respondents may complete the checkbox subsections in 11.A–11.D or describe their overall approach under 11.E.

11.A. Governance, policy and assessments

- Conduct regular cybersecurity risk and vulnerability assessments tailored to AI models or systems
- Perform threat modeling that includes AI-specific risks (e.g., data poisoning, prompt injection)
- Maintain or update security policies aligned with recognized standards (e.g., NIST CSF, ISO/IEC 27001, SOC 2, CIS Critical Security Controls), and review them regularly in response to emerging risks or regulatory guidance

- Maintain vulnerability management processes (scanning, channels for external reporting, patching, remediation)
- Implement incident response, disaster recovery, or business continuity plans, including escalation procedures for high-severity incidents
- Train staff on cybersecurity best practices and AI-specific threats, including how to recognize and respond to misuse or unintended outcomes of advanced AI systems
- Integrate cybersecurity AI risks into broader cybersecurity risk management activities
- Regularly review security measures and update them in response to emerging risks, incidents, or major releases
- Other(s) (specify): As a service handling medical information, our security governance additionally covers two areas. First, cloud usage compliant with the Two Guidelines from Three Ministries, the Japanese guidelines governing the safety management of medical information, and conformity with the security requirements of our customer medical institutions, which are confirmed at the time of each deployment under contract. Second, where we use external foundation models and cloud services, the scope of data use and the allocation of responsibility are clarified by contract. Responsibility is not delegated to those providers; we retain accountability for the security of our own products and verify provider terms and configurations as part of our own security review.

11.B. Technical and operational controls

- Apply multi-factor authentication (MFA) and role-based access controls for advanced AI systems
- Encrypt AI assets (datasets, models, configurations) in storage and in transit as appropriate
- Monitor system activity and log access to detect anomalies or misuse
- Conduct secure AI development and/or deployment practices
- Conduct security testing focused on AI-specific risks (e.g., red-teaming, adversarial testing, stress tests)
- Other(s) (please specify): Our products run on cloud infrastructure rather than on servers we own, so physical access control to servers and accelerators is provided by the cloud provider. We select cloud environments that hold third-party certifications such as ISO and SOC 2, primarily Google Cloud, and we store data there in encrypted form. Where we use external foundation models and cloud services, the scope of data use and the allocation of responsibility are clarified by contract, and we verify provider terms and configurations as part of our own security review rather than delegating responsibility.

11.C. Protection of sensitive or high value AI assets

- Limit access to proprietary AI assets based on role or need-to-know
- Store model weights, datasets, and algorithms in secure, access-controlled environments
- Track and log access to sensitive AI resources
- Require NDAs or confidentiality agreements for staff and contractors with access

11.D. Insider risk protection detections

- Provide staff training on cybersecurity best practices and AI-specific threats
- Monitor for unauthorized internal access or suspicious activity
- Implement insider-risk escalation or response procedures

- Other (please specify): Ubie treats insider threats as part of information security management rather than as a separate programme. Our Information Security Policy and internal rules apply to employees and contractors alike, and cover measures in the event of violations, access management, log management and monitoring, restrictions on taking information assets outside the company, and supervision of subcontractors. For work involving medical and health information or information belonging to customer medical institutions, we combine minimum necessary access privileges, segregation of duties, auditable logs, confidentiality obligations, and communication of internal rules. For highly confidential data of a medical institution, we adopt a design in which only that institution holds access rights, which limits what any insider at Ubie can reach.
- Conduct pre-employment screening or role-sensitive vetting

11.E. Please describe the main security measures your organization applies and indicate which are specific to particularly sensitive AI assets, where relevant.

General information security management. We have obtained and operate ISMS certification (ISO/IEC 27001:2022) ourselves, comply with the Two Guidelines from Three Ministries for the handling of medical information, and publish our Information Security Policy.

AI-specific risks we focus on. Prompt injection; unintended output of personal data; excessive transmission of data to external foundation models and APIs; retention of sensitive information in logs; inadequate permissions on RAG and search targets; and over-reliance on generative AI outputs in business operations.

Countermeasures. Management of the scope of input data; permission management; audit logs; masking and anonymization; contractual and security checks on external services; settings ensuring that input data is not used for open model training; secure usage environments; human review of outputs; limitation of the intended purpose of use; and operational support for customer medical institutions.

Measures specific to particularly sensitive assets. For features handling particularly sensitive medical information, we adopt healthcare-specific safety designs: a structure in which only the medical institution concerned holds access rights to its highly confidential data, and a "safety net structure" in which the original text is displayed by default and can always be referenced, so that an AI-generated summary is never the only available record. Personally identifiable information is held in a dedicated store accessible only to a limited number of staff.

12. What is your organization's vulnerability management process? Does your organization take actions to address identified risks and vulnerabilities, including in collaboration with other stakeholders?

- Internal vulnerability disclosure receipt and tracking processes
- Investigation of reported or identified risks and vulnerabilities
- Development of patches or other mitigations for confirmed risks or vulnerabilities
- Testing of patches or other mitigations
- Scanning for and mitigation of similar risks or vulnerabilities
- Collaborative vulnerability sharing (e.g., with partners, industry groups)
- Reporting to regulators / authorities when relevant

- Defined escalation criteria and follow-on actions for identified vulnerabilities (e.g., for external reporting or information sharing)
- Other(s) (please specify): Detection channels feeding our vulnerability management process are: internal development and operational monitoring; security reviews; log and error monitoring; reports from customer medical institutions; inquiries from consumer users; notifications from cloud vendors and other suppliers; publicly disclosed vulnerability information; and reports from external security researchers through our Vulnerability Coordination and Disclosure Policy (VCDP, <https://ubie.life/safety/vcdp>), under which we respond on the basis of a coordinated disclosure process. The process runs from detection, intake and prioritisation through remediation, verification, recurrence prevention and any required notification. Each confirmed vulnerability is assessed for scope of impact, exploitability, impact on medical and personal information, and whether notification is required by law or contract. Depending on severity we carry out remediation, a workaround, suspension of the feature, or notification of affected customers. After remediation we check whether similar vulnerabilities exist in other features or other products.

13. What post-deployment practices does your organization use to monitor and respond to vulnerabilities, incidents, emerging risks and misuse of advanced AI systems?

- Other(s) (please specify): Post-deployment we identify issues early through inquiries and support from consumer users and customer medical institutions, hands-on post-deployment support, log and error monitoring, external vulnerability information, and updates to medical and AI-related guidelines. Depending on severity and scope of impact, our responses include emergency response; suspension or restriction of a feature; revision of prompts, guardrails, UI and explanatory text; updates to our medical information database; security patches; notification of customers and users; and reporting as required by law or contract. For generative AI features we maintain the premise that healthcare professionals review the outputs, and where misuse is confirmed we review training, operational guides, templates and permission settings. Root cause analysis and recurrence-prevention measures are documented as post-mortems, and we check whether the same risk exists in other features or other products.
- Perform root-cause analysis and apply lessons learned to improve future development and/or deployment practices
- Conduct post-deployment monitoring

14. Do you develop or deploy agentic AI systems (systems that can autonomously plan/act or use tools)?⁴

Yes

14.A. If yes, please indicate which categories of controls your organization applies. Please select all.

- Human oversight or interruptibility mechanisms (approval requirements for sensitive actions, escalation paths, interruption interfaces)
- Logging of agent actions
- Other (specify): Agentic capability is under evaluation and has not been rolled out to external-facing services. The controls indicated here are those applied across our existing products and within the evaluation environment. Human-in-the-loop is a precondition across all our AI features: a healthcare professional or a user makes the final decision, and AI output is positioned as support for that

judgment. In our patient intake summary feature the original text is displayed by default and can always be referenced, so that an AI-generated summary is never the only available record. Inputs, outputs and creation history are logged in our business systems, which preserves ex post verifiability. Evaluation perspectives for autonomous action-taking, tool use and multi-step planning are being organised through our research organization Ubie Lab and the AISI Healthcare Sub-Working Group. Any future introduction would follow the evaluation perspectives in the AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector, which Ubie led the development of.

Please provide any further comments and implementation documentation.

We recognise the safety of agentic AI as an important emerging risk area, and we are organising and researching evaluation perspectives through our research organization Ubie Lab and our participation in the AISI Healthcare Sub-Working Group. The Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector, which Ubie led the development of, is published in part in Markdown format so that it can be loaded into AI services, and a practical guide containing examples of prompts and agent skills is also planned for publication.

In using generative AI and agent technologies for internal operations, we rigorously apply the principle that final decisions are made by humans, together with guardrail design and transparency of decision logic. If we consider introducing agent technologies into external-facing services in the future, we will do so on the premise of a phased and safe rollout based on the evaluation perspectives in that Guide.

Implementation documentation.

- Commitment to safety and security: <https://ubie.life/safety>
- Data protection: https://ubie.life/safety/data_protection
- Privacy management structure: https://ubie.life/safety/privacy_policy_management
- Vulnerability Coordination and Disclosure Policy (VCDP): <https://ubie.life/safety/vcdp>
- Privacy policies: <https://ubie.app/policy> (consumers), <https://intro.dr-ubie.com/policy> (medical institutions)
- AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector (Japanese and English editions): https://aisi.go.jp/output/output_information/260402/

Section 3 – Transparency Reporting on Advanced AI Systems

15. Does your organization publish public reports and/or technical documentation related to the capabilities, limitations, and domains of appropriate and inappropriate use of advanced AI systems? Are these clear and understandable by non-technical audiences?

Yes

15.A. If yes, which of the following information is included in your organization's publicly available documentation for any of your AI products?

- Capabilities, behaviours, and limitations, including domains of appropriate and inappropriate use or indicative thresholds associated with systemic risks.
- Technical documentation and instructions for use if relevant.
- Plain-language or non-technical summaries

Please provide links or references showing where this information is made public.

Capabilities, behaviours and limitations, including appropriate and inappropriate use:

- Ubie (consumer-facing service), including statements that the service is not a medical device, is not intended for diagnosis or treatment, and should not be used in emergencies: <https://ubie.app/about>

- Terms of service and in-service notices, presented within the service and on the app store listings

Technical documentation and instructions for use:

- Commitment to safety and security, covering privacy, security and content reliability:

<https://ubie.life/safety>

- Data protection: https://ubie.life/safety/data_protection

- Privacy management structure: https://ubie.life/safety/privacy_policy_management

- Vulnerability Coordination and Disclosure Policy (VCDP): <https://ubie.life/safety/vcdp>

- For medical institutions, further documentation on features, security, conformity with the Two Guidelines from Three Ministries, and confirmation that input data is not used for open model training is provided through deployment materials and training under contract

Plain-language and non-technical summaries:

- <https://ubie.life/safety> and the in-service explanations for consumer users, written in plain Japanese

- Help Center for consumer users

- Deployment training and materials for medical institutions, aimed also at non-technical staff such as medical clerks

Related transparency documentation:

- Data sources and physician supervision process: <https://ubie.app/about>

- Privacy policies: <https://ubie.app/policy> (consumers), <https://intro.dr-ubie.com/policy> (medical institutions)

- AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector, which Ubie led the development of: https://aisi.go.jp/output/output_information/260402/

15.B. If yes, how often are such reports usually updated?

Other (e.g., triggered by material changes, incidents, or new risks)

16. How does your organization share information with a diverse set of stakeholders (other organizations, governments, civil society and academia, etc.) regarding the outcome of evaluations of risks and impacts related to an advanced AI system?

- Public documentation
- Protected information sharing
- Regulatory reporting as relevant
- Direct engagement with AI safety institutes, civil society or academic partners

Describe how evaluation results are shared and with which stakeholders (e.g. other organizations, governments, civil society and academia, etc). You may also indicate whether non-technical or plain-language summaries are provided for broader audiences.

Sharing methods and audiences.

Public disclosure. Policies, case studies and press releases, together with our commitment to safety and security page (<https://ubie.life/safety>), set out the intended uses and limitations of our services, our

handling of data, security, and the results of deployment.

Individual sharing with customers. Medical institutions and pharmaceutical companies receive deployment materials, operational explanations and incident notifications under contract. Highly confidential risk assessment results, vulnerabilities and customer-specific information are shared only with relevant parties, within the scope required by contract, law and security considerations.

Government and public bodies. We have presented at briefing sessions held by the Japan AI Safety Institute (AISi) and the Information-technology Promotion Agency (IPA), and we make reports required by law.

Industry associations. We share knowledge through The Japan Digital Health Alliance (JaDHA), the AI Governance Association (AIGA), the Japan Medical Venture Association (JMVA) and the Japan Trust & Safety Association (JTSA).

Academia. We run demonstration projects with university hospitals and speak at academic conferences and external meetings.

Direct engagement, in detail. As leader of the Healthcare Sub-Working Group under the AISi Business Demonstration Working Group, on behalf of JaDHA, we co-developed and published the Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector in April 2026. Through demonstration projects with medical institutions and university hospitals we share findings on effectiveness, limitations and operational precautions. We have also spoken at the Hiroshima Global Forum for Trustworthy AI.

Non-technical summaries. Yes. Our commitment to safety and security page (<https://ubie.life/safety>) and the explanations presented within our services are written in plain Japanese for general consumers. For medical institutions we provide training and deployment materials aimed also at non-technical staff such as medical clerks and nurses.

17. Does your organization publish public privacy policies addressing the use of personal data, user prompts, and/or the outputs of advanced AI systems?

Yes

If yes, provide links to privacy policies or notices covering personal data, user inputs, and AI-generated outputs.

Privacy policies.

- Consumer-facing service (Ubie): <https://ubie.app/policy> - sets out the information we collect, purposes of use, provision to third parties, cookies and external transmission, safety management measures, and contact points.

- Services for medical institutions: <https://intro.dr-ubie.com/policy> - sets out the methods of collecting personal data including sensitive personal information, purposes of use, provision to third parties, use of statistical data, safety management measures, provision overseas, and procedures for disclosure, correction and suspension of use.

User inputs. That user input, such as symptoms and the content of a consultation, is recorded and used for health management is explained within the service, and deletion requests are handled. For services for medical institutions, we publish that input data is not used for open model training.

AI-generated outputs. Our terms of service explain that we provide general reference information based on the information entered by the user, and that service outputs do not constitute diagnosis, medical examination or treatment. The same statements are presented within the service interface.

Related published documentation.

- Commitment to safety and security: <https://ubie.life/safety>
- Data protection: https://ubie.life/safety/data_protection
- Privacy management structure: https://ubie.life/safety/privacy_policy_management

18. How does your organization provide information about the sources of data used for the training of advanced AI systems, as appropriate, including information related to the sourcing of data annotation and enrichment?

- Public documentation
- Provide information or tools to downstream actors on data quality
- Other (specify): bie is not in a position to manage or disclose the pre-training datasets of foundation models. Information on foundation model training data depends on the public disclosures and contract terms of each foundation model provider. At our application layer, the medical information used by our services is based on publicly known, reliable sources - clinical practice guidelines, clinical protocols and medical papers - and each item is traceable to a cited source. Content is authored by in-house physicians and checked by supervising physicians, and we publish our approach to data sources and this supervision process. Part of the list of data used is also published. Data held within medical institutions and handled by our generative AI for medical institutions - electronic medical records, audio and documents - is handled in accordance with the contract, access permissions, consent and security requirements agreed with the customer medical institution. We publish that input data is not used for open model training.

Please provide links to any publicly available information.

- Data sources and physician supervision process for the consumer-facing service (Ubie): <https://ubie.app/about>
- Privacy policy, consumers: <https://ubie.app/policy>
- Privacy policy, medical institutions: <https://intro.dr-ubie.com/policy>
- Commitment to safety and security: <https://ubie.life/safety>
- Commitment to data protection: https://ubie.life/safety/data_protection

19. Does your organization demonstrate transparency related to advanced AI systems through any other methods?

Yes

If yes, describe additional method(s), including where relevant transparency measures related to AI agents

In-service disclosure. Our service screens state that the service is not a medical device, that it is not intended for diagnosis or treatment, and that it should not be used in emergencies. Disclaimers are displayed on AI-generated responses. In the patient intake summary feature for medical institutions, the UI clearly distinguishes an AI-generated summary from the original text entered by a human, with the original text displayed by default and always retrievable. Publication of supervision and sourcing. We publish our physician supervision process and our data sources, and we provide contact points for inquiries. Explanation at the point of deployment. When deploying at a medical institution we provide

hands-on support and training that explain the premise that AI outputs are reviewed and edited by healthcare professionals, the intended scope of use, and the applicable security and privacy requirements. Publication of how the technology is actually used. Case studies and press releases set out which tasks the technology is used for and how humans review its use, rather than presenting capability claims alone. Institutional transparency. Publication of this transparency report under the Hiroshima AI Process reporting framework; inclusion as an advanced case (Column 11) in the Appendix to the METI/MIC AI Guidelines for Business (Ver. 1.1); co-development and publication of the AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector; and activity as a regular member of the AI Governance Association (AIGA). Transparency measures related to AI agents. Agentic capability is under evaluation, so user-facing transparency measures specific to agents are not applicable at present. Evaluation perspectives for autonomous action-taking, tool use and multi-step planning are being organised through our research organization Ubie Lab and the AISI Healthcare Sub-Working Group, and the results of that work are made public through the guide we co-developed. Inputs, outputs and creation history are logged in our business systems, which preserves ex post verifiability for any future agentic feature.

20. Does your organization publish statistical insights derived from anonymised, aggregated usage data to advance understanding of AI adoption and its economic impacts in areas including occupations and tasks by region, sector and country? To this end, do you contribute to international cooperation efforts?

Yes

If yes, please provide links or describe the scope of insights published, including where relevant contributions to understanding AI adoption in specific sectors (e.g. public administration, education systems).

Scope of insights published.

Consumer symptom and care-seeking behaviour. Our consumer-facing service Ubie is used by more than 13 million people per month. From anonymised, aggregated usage data we publish statistical insights on symptoms and care-seeking behaviour, in survey reports and press releases. These describe how people move from a symptom to a decision about seeking care, which is difficult to observe at population scale by other means.

AI adoption in the healthcare sector, at the level of specific occupations and tasks. For our services for medical institutions we publish measured effects of AI adoption on clinical documentation work: a 42.5% reduction in the time required to create discharge summaries, and a 54% improvement in the efficiency of creating admission summaries. These are task-level figures for physicians, nurses and medical clerks in the healthcare sector, which is the sector in which we operate.

Links.

- Improved AI performance in the Mayo Clinic Platform_Accelerate program:

<https://prtimes.jp/main/html/rd/p/000000175.000048083.html>

- Establishment of our research organization Ubie Lab:

<https://prtimes.jp/main/html/rd/p/000000060.000048083.html>

- Ubie's work with Gemini, on the official Google blog: <https://blog.google/intl/ja-jp/company-news/technology/ubie/>

- About Ubie, including data sources and the physician supervision process: <https://ubie.app/about>

- AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector:

https://aisi.go.jp/output/output_information/260402/

Contributions to international cooperation. Publication of this transparency report under the Hiroshima AI Process reporting framework. Co-development of the Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector, published in April 2026, as leader of the AISI Healthcare Sub-Working Group on behalf of The Japan Digital Health Alliance. Speaking at the Hiroshima Global Forum for Trustworthy AI. Participation in international programmes including Mayo Clinic Platform_Accelerate, through which we evaluated performance on large-scale

Please estimate the total staff time and annual expenses associated with conducting compliance activities and include any further comments and implementation documentation (optional).

No answer provided

Section 4 – Organizational Governance, Incident Management, and Transparency

21. How has AI risk management been embedded in your organization's governance framework?

- Governance council or committee
- Other (e.g., distributed / multi-institutional governance arrangement) (please specify): In addition to the oversight bodies indicated above, AI risk management at Ubie is embedded through a two-layer structure. The Generative AI Promotion Team pursues value from AI utilization - product value, internal productivity and brand value - while the Risk and Compliance Committee, comprising legal, security, public relations and other functions, examines risk. The two share information and deliberate trade-offs with each other, so that promotion and restraint are weighed together rather than in sequence. Risk assessment itself is performed by dedicated specialists who are independent of the development and business units, which ensures objectivity, and concerns are escalated to the Committee for deliberation. Since January 2026 the Trust & Safety Team has been responsible for cross-product quality management across all products. This structure is featured as an advanced case (Column 11) in the Appendix to the METI/MIC AI Guidelines for Business (Ver. 1.1):
https://www.meti.go.jp/shingikai/mono_info_service/ai_shakai_jisso/pdf/20250328_3.pdf
- Board of Directors oversight
- Executive oversight
- Senior management oversight

Please describe how AI risk management is integrated into governance structures, including: which roles or bodies have advisory versus decision-making authority; which roles or bodies have final accountability for AI risk-related decisions; and how often policies are reviewed or updated (e.g., following major incidents, regulatory changes, or material updates to systems).

Advisory versus decision-making authority.

Advisory. Dedicated specialists, independent of the development and business units, perform risk assessment and raise concerns. They assess rather than decide, which keeps the assessment

separate from the incentive to ship.

Decision-making. Each product team decides the release of its own features, but only after the Trust & Safety Team has confirmed conformity with our quality standards; that confirmation is a precondition, not an opinion. Where a material concern remains, the matter is escalated to the Risk and Compliance Committee, which comprises legal, security, public relations and other functions and which can decide to postpone release, limit the scope of provision, restrict the intended purpose, or require additional safeguards.

Final accountability. Final accountability for AI risk-related decisions rests with the Risk and Compliance Committee, under executive-level oversight.

How AI risk management is integrated. It is embedded in each of our processes for product development, medical information quality management, information security, privacy, compliance, and customer deployment support. It operates on the basis of our Information Security Policy, Privacy Policy, terms of service, medical information supervision process, internal rules on the use of AI and generative AI, and incident response procedures. Since January 2026 the Trust & Safety Team has carried out cross-product quality management, which allows a finding in one product to be checked against the others.

How often policies are reviewed or updated. In addition to review on a regular cycle, we update policies as needed when: laws or guidelines change, including the AI Guidelines for Business, the AI Promotion Act, the Act on the Protection of Personal Information, the Two Guidelines from Three Ministries, the Pharmaceuticals and Medical Devices Act and the Medical Practitioners Act; service specifications or the external foundation models we use change materially; a material incident occurs; a significant issue is raised by a customer medical institution or a user; or new knowledge on AI safety becomes available.

22. Are relevant staff trained on your organization's governance policies and risk management practices?

Yes

If yes, describe training methods (e.g., onboarding modules, regular sessions, and confidential channels to report concerns).

Channels to report concerns. Ubie operates an internal whistleblowing channel through which employees and contractors can report concerns confidentially. Separately, under our Information Security Policy, employees and contractors are required to report information security incidents to the Information Security Manager without delay, and incidents are recorded and tracked in an incident management tool used across all products. Internally we promote a culture that encourages the responsible reporting of risks, potential vulnerabilities and incidents, and reports are used for learning through post-mortems rather than for attributing fault.

23. Are steps taken to address reported incidents documented and maintained internally?

Yes

If yes, describe how incident records are kept and used for follow-up (e.g., incident logs, root-cause analysis reports, internal tracking systems). You may indicate whether and how you define escalation thresholds for when incidents must be surfaced to senior leadership or external parties.

How records are kept. Incidents across all products are recorded and tracked in an incident management tool. Each record captures the date of occurrence, the detection channel, the scope of impact, the products involved, the root cause, the interim response, the permanent response, recurrence-prevention measures, and whether notification or reporting is required. Root cause analysis and recurrence-prevention measures are documented as post-mortems.

How records are used for follow-up. Documented findings are used to take stock of similar risks, to add test items to the criteria applied before release, for internal education, and to review policies and operating procedures. After remediation we check whether the same weakness exists in other features or other products. The Trust & Safety Team collects and classifies incidents across the organisation, which allows common issues to be distinguished from issues specific to a single product.

Escalation. For information security incidents, the person who discovers the incident reports to the Information Security Manager without delay under our Information Security Policy, and emergency measures, root cause analysis and recurrence-prevention measures follow; this is managed within our ISMS (ISO/IEC 27001) framework. Beyond that, escalation is determined by assessment rather than by a fixed numeric threshold: for each incident or vulnerability we assess the scope of impact, exploitability, the impact on medical and personal information, and whether notification is required by law or contract. Where a material concern remains, the matter is escalated to the Risk and Compliance Committee, which comprises legal, security, public relations and other functions.

External surfacing. External notification is driven by legal and contractual obligation rather than discretion. Affected customer medical institutions and pharmaceutical companies are notified, and reports are made to public authorities including the Personal Information Protection Commission where required by law. Depending on severity, the response may also include a workaround, suspension of the feature, or customer notification.

24. Does your organization share research and best practices on addressing or managing risk?

Yes

If yes, describe how research or good practices are shared (e.g., publications, conferences, collaborations with other sectors).

Published guides. As leader of the Healthcare Sub-Working Group under the AISI Business Demonstration Working Group, on behalf of The Japan Digital Health Alliance (JaDHA), Ubie led the development of the Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector, published jointly with the Japan AI Safety Institute in April 2026. The guide provides practical checklists and evaluation perspectives to the industry, and part of it is published in Markdown format so that it can be loaded into AI services. Ubie also led the development of the JaDHA Guide to the Utilization of Generative AI for Healthcare Providers. Both are freely available:

https://aisi.go.jp/output/output_information/260402/ and <https://jadha.jp/>

Conferences and public briefings. We have spoken at the Hiroshima Global Forum for Trustworthy AI and at IPA/AISI Business Demonstration Working Group briefing sessions. We also host the Japan Medical Generative AI Forum, which brings healthcare professionals and vendors together on the safe use of generative AI in clinical settings.

Research. Our research organization Ubie Lab conducts safety verification of generative AI and research on policy developments, and makes policy recommendations.

Case studies and press releases. We publish deployment cases and measured outcomes, including our work with Google Cloud/Gemini, performance evaluation in the Mayo Clinic Platform_Accelerate program, and demonstration projects with university hospitals. These describe not only results but which tasks the technology is used for and how humans review its use.

Collaboration across sectors. We participate in study group activities as a regular member of the AI Governance Association (AIGA), and are a director of the Japan Medical Venture Association (JMVA) and a member of the Japan Trust & Safety Association (JTSA). Through these bodies our healthcare-specific findings reach practitioners outside healthcare, and we take in failure modes reported from other industries.

25. Does your organization use international technical standards or best practices for AI risk management and governance policies?

Yes

If yes, please indicate which standards or frameworks apply:

- ISO/IEC standards (e.g., ISO/IEC 27001, ISO/IEC 42001)
- NIST AI Risk Management Framework (AI RMF)
- OECD AI Principles
- Industry-led best practices or codes of conduct
- Other international or regional standards (please specify): The METI/MIC AI Guidelines for Business (Ver. 1.1), in whose Appendix our AI governance structure is featured as an advanced case (Column 11); the Japan AI Safety Institute (AISI) Guide to Evaluation Perspectives on AI Safety and its healthcare edition, the Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector, which Ubie led the development of; the JaDHA Guide to the Utilization of Generative AI for Healthcare Providers; and the Two Guidelines from Three Ministries, the Japanese guidelines governing the safety management of medical information.

If yes, describe which international technical standards or best practices your organization uses.

AI governance. We refer to and adopt the Hiroshima AI Process International Code of Conduct, the OECD AI Principles, the NIST AI Risk Management Framework, the METI/MIC AI Guidelines for Business, and the AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector. These inform our internal risk taxonomy and the countermeasure policies applied at each lifecycle stage.

Information security. We have obtained and operate ISO/IEC 27001:2022 (ISMS) certification ourselves, and run a PDCA cycle for risk assessment and countermeasures under it. Medical information is managed in compliance with the Two Guidelines from Three Ministries. Cloud services are selected from those holding third-party certifications such as ISO and SOC 2.

Healthcare-specific. We comply with domestic law including the Pharmaceuticals and Medical Devices Act and the Medical Practitioners Act, and we put the JaDHA Guide to the Utilization of Generative AI for Healthcare Providers into practice within our own organisation, not only as its lead developer.

Under consideration. We are examining alignment with and adoption of ISO/IEC 42001, the AI management system standard.

Any further comments and implementation documentation.

No answer provided

Section 5 – Content Authentication & Provenance Mechanisms

26. What mechanisms, if any, does your organization put in place to allow end-users, where possible and appropriate, to know when they are interacting with an advanced AI system developed by your organization?

- User disclosure (e.g., text notice, verbal statement, or system label)
- Distinct branding or labeling for AI-generated content (e.g., “AI-generated” tag, metadata label)
- Documentation or FAQs explaining when and how AI is used in the system
- Other (specify): For services provided to medical institutions, disclosure is made not only in the interface but through the contract, the deployment materials and the training we provide, so that healthcare professionals understand that AI performs branching of questions, summarisation, document drafting and search support. In the patient intake summary feature the UI clearly distinguishes an AI-generated summary from the original text entered by a human, with the original text displayed by default and always retrievable. We also use service naming and branding that make the use of AI explicit.

Provide any additional details or context on how or when users are informed they are interacting with advanced AI systems (e.g. if such mechanisms are planned, in testing, or applied only to certain products).

At the start of service use. Users are informed that the service uses AI at the point they agree to the terms of service.

During use. On-screen indications and disclaimers allow users to recognise this continuously. In conversational services the interface makes clear that the user is interacting with AI.

At the point of deployment at a medical institution. Disclosure is made through the contract, the deployment materials and the training we provide, so that healthcare professionals understand which tasks AI performs - branching of questions, summarisation, document drafting and search support.

Scope of application. These mechanisms are already applied in production across our main products in service: Ubie for consumers, and AI Health Assistant and Ubie Generative AI for medical institutions.

They are not planned or in testing.

Not applicable at present. Agentic capability is under evaluation, so disclosure relating to an AI system acting autonomously or executing actions on a user's behalf does not currently arise.

27. Does your organization use content provenance detection, labeling or watermarking mechanisms that enable users to identify content generated by advanced AI systems?

Yes

27.A. If yes, how?

- Manifest disclosure (e.g., visible label)

27.B. If yes, does your organization use international technical standards or best practices when developing or implementing content provenance?

No

27.C. Does your organization collaborate on or invest in research to advance content authentication and provenance?

Yes

Describe how provenance mechanisms are implemented and which standards or tools are used, and any collaborations or investments to improve content authentication (e.g., C2PA credentials, joint R&D projects, open-source initiatives, standard-setting groups). You may also indicate whether these mechanisms are tested for robustness against removal, tampering, or circumvention.

Scope. Our services centre on the provision of medical information and support for business operations, primarily in text form. We are not a generative AI platform that broadly distributes content such as images, video or audio. Accordingly, watermarking and metadata embedding such as C2PA are not currently a principal area of control for us, and we do not implement latent provenance signals or detection tools for them.

How disclosure is implemented. We rely on manifest disclosure. First, the UI states explicitly that a response or output is AI-generated. Second, in the patient intake summary feature for medical institutions, an AI-generated summary is clearly distinguished from the original text entered by a human: the original text is displayed by default and can always be referenced, a design we refer to as a safety net structure, so that an AI-generated summary is never the only available record. Third, the use of AI is stated explicitly in contracts and terms of service.

Verifiability rather than watermarking. Documents, summaries and drafts generated by AI are treated as supporting information on the premise that a healthcare professional reviews and edits them. Input data, creation history and logs are managed in our business systems, which preserves ex post verifiability of what was generated and by whom it was reviewed. In a clinical setting this traceability serves the purpose that provenance signals serve elsewhere.

Collaboration and investment in research. The AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector, which Ubie led the development of as leader of the AISI Healthcare Sub-Working Group on behalf of The Japan Digital Health Alliance, includes evaluation perspectives on ensuring the reliability and verifiability of AI-generated content: https://aisi.go.jp/output/output_information/260402/ Future scope. If the scope of our handling of patient-facing generated content or externally published content expands, we will consider mechanisms suited to the characteristics of our services in light of developments in content authentication, provenance and labelling standards such as C2PA.

Any further comments and implementation documentation.

No answer provided

Section 6 – Research & Investment to Advance AI Safety & Mitigate Societal Risks

28. Does your organization advance research and investment in the following areas?

- Security
- Safety
- Data quality and representativeness
- Information integrity
- System performance and outcomes (e.g. accuracy, consistency)
- Explainability or interpretability
- Transparency
- Robustness or reliability

If yes, please describe key research or investment initiatives in the areas selected above (e.g., dedicated safety teams, funding of research programs, contributions to open-source safety tools, publishing research, partnerships and collaborative research initiatives). You may also indicate whether related outputs are publicly shared, shared with partners or within consortia or not publicly disclosed.

Dedicated research organization. Ubie Lab, established in July 2023, advances verification of the underlying technologies and safety of generative AI, covering countermeasures against hallucination, output stability and medical accuracy. It also conducts research on policy developments and makes policy recommendations. Ubie Lab operates under the principle of achieving both innovation and safety.

Publicly shared: <https://prtimes.jp/main/html/rd/p/000000060.000048083.html>

Dedicated safety function. Since January 2026 the Trust & Safety Team has carried out cross-product quality management, incident collection and classification, and pre-release review across all products. Risk assessment is performed by specialists independent of the development and business units. Not publicly disclosed as a research output, but described in this report and in the Appendix to the METI/MIC AI Guidelines for Business (Ver. 1.1), Column 11.

Leading a sector-wide safety standard. As leader of the Healthcare Sub-Working Group under the AISI Business Demonstration Working Group, on behalf of The Japan Digital Health Alliance (JaDHA), Ubie led the development of the Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector, published in April 2026 jointly with the Japan AI Safety Institute. It sets out ten evaluation perspectives across five lifecycle phases, covering security, safety, data quality, information integrity, explainability, robustness and verifiability. Publicly shared, free of charge, and published in part in Markdown so that it can be loaded into AI services: https://aisi.go.jp/output/output_information/260402/ Ubie also led the development of the JaDHA Guide to the Utilization of Generative AI for Healthcare Providers, publicly shared at <https://jadha.jp/>

Data quality and representativeness. We invest continuously in a medical knowledge database built from more than 50,000 medical papers and clinical practice guidelines, covering more than 1,100 disease names, authored in-house by physicians and supervised by more than 50 physicians, with scheduled review and updating against the latest clinical guidelines. The approach and the supervision process are publicly shared: <https://ubie.app/about>

System performance and outcomes. In the Mayo Clinic Platform_Accelerate program we conducted performance evaluation using large-scale patient data, covering a different language and clinical-cultural context from our domestic services. Publicly shared:

<https://prtimes.jp/main/html/rd/p/000000175.000048083.html> We also measure deployment outcomes at customer medical institutions, and publish them as case studies.

Explainability and transparency. Rather than pursuing post-hoc explanation of model internals, we invest in designs that make the basis of an output inspectable in clinical use: retrieval grounding with source citation, and the safety net structure in which the original text is displayed by default and remains retrievable alongside any AI-generated summary. Publicly described at <https://ubie.life/safety> Partnerships. Safety and performance evaluation is conducted with technology partners, including our work with Google Cloud/Gemini, publicly shared on the official Google blog: <https://blog.google/intl/ja-jp/company-news/technology/ubie/> We also participate in study groups as a regular member of the AI Governance Association (AIGA), where findings are shared within the consortium.

Disclosure posture. Guides, evaluation perspectives and deployment outcomes are publicly shared. Internal test criteria, prompts, guardrail configurations and specific vulnerability details are not publicly disclosed, since disclosure would reduce their effectiveness or enable exploitation.

29. Does your organization pursue research or investment related to socio-economic or environmental considerations in the development and deployment of AI?

- Socio-economic considerations (e.g., AI adoption and productivity, workforce transitions, skills development, access to AI tools)
- Environmental considerations (e.g., energy use, computational efficiency, resource use)

If yes, please provide examples of relevant research or investment initiatives for the areas selected above.

Socio-economic considerations.

We treat four issues as priorities: the excessive workload of healthcare professionals, uneven access to healthcare, incorrect self-assessment by patients, and over-reliance on AI. Our research and investment is directed at three responses.

First, guidance to appropriate care. Our consumer-facing symptom checker, used by more than 13 million people per month, guides people from symptoms to an appropriate medical institution and clinical department. This addresses uneven access to care and supports earlier consultation and earlier detection.

Second, operational efficiency at medical institutions. We invest in patient intake, document creation, summarisation and DPC support - DPC being Japan's case-mix based inpatient payment system - which reduces the documentation burden on physicians, nurses and medical clerks. Measured outcomes include a 42.5% reduction in the time required to create discharge summaries and a 54% improvement in the efficiency of creating admission summaries. We also promote productivity improvement through the use of generative AI within our own organisation.

Third, skills development for healthcare professionals. We invest in training, hands-on deployment support and the sharing of prompt templates, and we host the Japan Medical Generative AI Forum, so that healthcare professionals can understand the nature, capabilities and limitations of generative AI and use it safely.

Environmental considerations.

Because Ubie does not perform pre-training of foundation models, our direct consumption of large-scale computing resources is limited. Within that scope we work to curb unnecessary consumption of computing resources through cost and quality optimisation of the large language models we use,

selection of cloud services and foundation models, and reduction of wasteful processing. Operational efficiency gains at medical institutions also reduce paper, travel and time costs.

Publicly shared outputs.

- Ubie's work with Gemini, on the official Google blog: <https://blog.google/intl/ja-jp/company-news/technology/ubie/>

- Improved AI performance in the Mayo Clinic Platform_Accelerate program:

<https://prtimes.jp/main/html/rd/p/000000175.000048083.html>

- Establishment of our research organization Ubie Lab:

<https://prtimes.jp/main/html/rd/p/000000060.000048083.html>

Any further comments and implementation documentation.

No answer provided

Section 7 – Advancing Human and Global Interests

30. Does your organization support any digital literacy, education or training initiatives to improve user awareness and/or help people understand the nature, capabilities and limitations of advanced AI systems?

Yes

If yes, describe any digital literacy or training initiatives your organization supports (e.g., user education, public AI-education programs, school or university partnerships, community workshops). You may also indicate whether these efforts include or prioritise underserved or underrepresented communities.

For consumer users. Through the service interface, the terms of service, our Help Center and the publication of our data sources and physician supervision process, we explain in plain Japanese that the information AI provides is general reference information, that it does not constitute diagnosis or treatment, and that the service should not be used in emergencies. The explanation is written for people with no technical or medical background: <https://ubie.app/about> and <https://ubie.life/safety>

For healthcare professionals. At the point of deployment we provide explanation, hands-on support, training and seminars, publish case studies, and share prompt templates. We host the Japan Medical Generative AI Forum, which brings healthcare professionals and vendors together on the safe use of generative AI in clinical settings. Training is aimed not only at physicians but also at nurses and medical clerks, who are often excluded from AI training programmes despite handling the resulting documents.

For the industry and for policy. We led the development of the JaDHA Guide to the Utilization of Generative AI for Healthcare Providers and the AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector, both of which are designed to raise the AI literacy of healthcare businesses that do not have AI specialists in-house. We have presented at IPA/AISI Business Demonstration Working Group briefing sessions, and our practices are provided as an advanced case (Column 11) in the Appendix to the METI/MIC AI Guidelines for Business (Ver. 1.1).

Internally. Onboarding training, generative AI study sessions for leadership, ideathons and hackathons, and an internal award programme raise AI literacy across the company.

31. Does your organization pursue AI projects that aim to deliver societal or public benefits, including through partnerships with civil society and community groups?

Yes

If yes, please provide examples, including any collaborations with civil society, community groups, or efforts to address broader societal challenges.

Ubie prioritises AI projects that connect directly to the Sustainable Development Goals, in particular Good Health and Well-being (SDG 3).

Consumer-facing service. Ubie provides more than 13 million people per month with guidance from symptoms to an appropriate medical institution and clinical department. This addresses two societal challenges: disparities in access to healthcare, and delayed consultation and detection.

Services for medical institutions. By reducing the documentation and administrative burden on healthcare professionals, our services for medical institutions create time for clinicians to spend with patients and support improvements in the quality of care. This responds to the excessive workload of healthcare professionals, which is a structural issue in Japan's healthcare system.

Collaboration with medical institutions and academia. We work with local medical institutions, university hospitals, acute care hospitals and clinics to identify and develop AI solutions to problems that arise in clinical practice - clinical research, discharge summaries, summarisation of referral letters and DPC coding work - always on the premise that healthcare professionals make the final decision.

Multi-stakeholder collaboration. The deliberation structure for the AISI Guide to Evaluation Perspectives on AI Safety in the Healthcare Sector brought together multiple companies, academic and public-interest experts, and a research institute acting as secretariat, and we led it as Healthcare Sub-Working Group leader on behalf of The Japan Digital Health Alliance. We also participate in industry alliances: JaDHA, the AI Governance Association (AIGA), the Japan Medical Venture Association (JMVA) and the Japan Trust & Safety Association (JTSA).

Principles applied to these projects. Physician supervision, reliance on trustworthy medical information sources, final human judgment, privacy protection, and explanation of the limitations of AI outputs. We treat these as conditions of the work rather than as additions to it, because in healthcare a project that delivers public benefit unsafely delivers no benefit.

Any further comments and implementation documentation.

No answer provided