

# Research and Development of AI Trust and Governance

Innovations brought about by AI and digitalization are significantly transforming the concepts of trust and governance. To ascertain more clearly what these two concepts mean in a digital society, Hitachi has worked with the World Economic Forum and Japan's Ministry of Economy, Trade and Industry to develop the Trust Governance Framework. Hitachi is also moving ahead with work on AI governance, including assessing the impacts and risks of AI on the economy and wider society through participation in the Ministry of Internal Affairs and Communications' Conference toward AI Network Society. This article describes studies into how to build trust in ways that will foster wellbeing in digital societies and the research and development of technologies that underpin AI governance.

**Naokazu Uchida**  
**Tadashi Kaji, Ph.D.**  
**Nick Blake**  
**Masayoshi Mase, Ph.D.**  
**Hiroki Ohashi**

**Dipanjan Ghosh, Ph.D.**  
**Chetan Gupta, Ph.D.**  
**Ken Naono, Ph.D.**  
**Mika Takata**

## 1. Introduction

Artificial intelligence (AI) has achieved a high level of performance in recent years, with applications growing in fields such as railways, energy, healthcare, and finance. At the same time, the innovations arising out of AI and digitalization are forcing a re-evaluation of how people think about trust and governance. For example, while numerous different systems interoperate with one another to create new value in a digital society, this also brings a lack of clarity over who is responsible when problems occur. Such arrangements raise issues that are difficult to address under existing law, such as responsibility for decisions made by AI and how to ensure the reliability of algorithms and data.

In a digital society, trust is built by stakeholders working together. Technologies and services come to be trusted by society when they provide evidence that such trust is warranted (a basis of trust), and when this evidence is accepted by users and the general public. It is on the basis of this philosophy that Hitachi has been undertaking research and

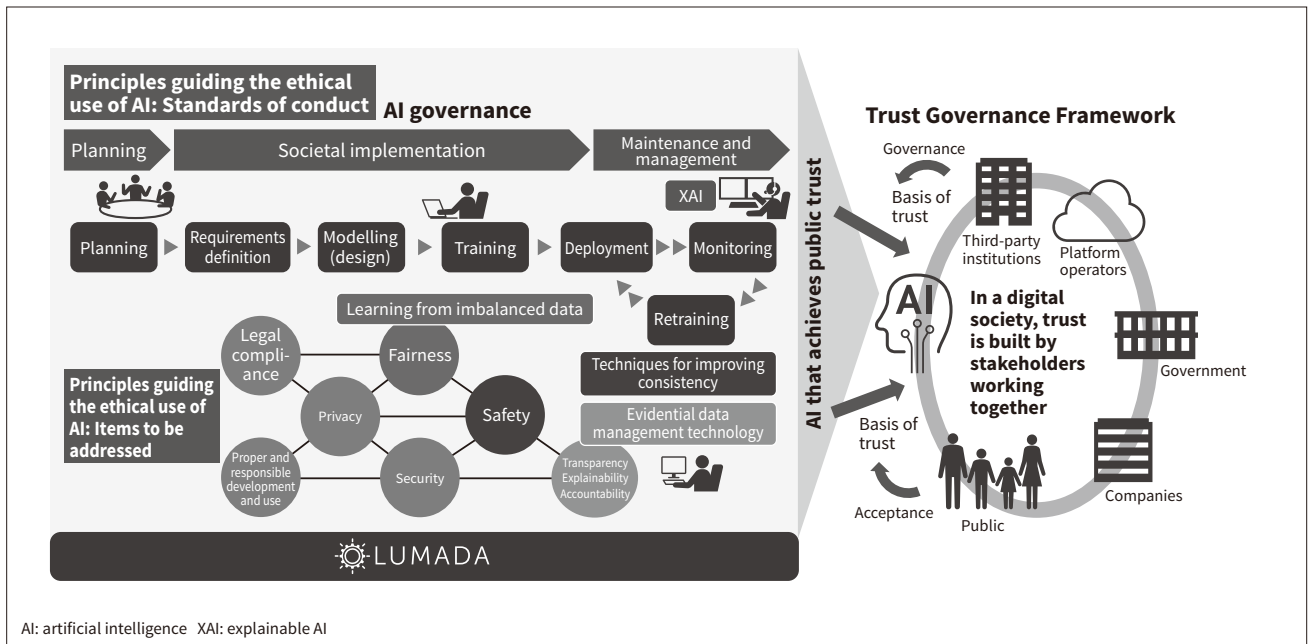
development aimed at achieving public trust in AI. Working in partnership with the World Economic Forum and Japan's Ministry of Economy, Trade and Industry, this has included the publication of a white paper on the nature of trust in a digital society entitled, "Rebuilding Trust and Governance: Towards Data Free Flow with Trust (DFFT)"<sup>(2)</sup> and the development of the Trust Governance Framework that describes the trust and governance relationships needed by such a society. Progress is also being made on developing a variety of different technologies, such as explainable AI (XAI), that will underpin AI governance, and on deploying these in Lumada, the engine behind Hitachi's Social Innovation Business (see **Figure 1**).

AI governance is a vital part of winning public trust in the technology. Along with safety, security, and privacy, it also requires that consideration be given to notions such as fairness, transparency, explainability and accountability.

One example is the opaque nature of decision-making by AI. Here, XAI can enhance the transparency of systems that use AI by shedding light on why particular decisions were made, such as by highlighting which factors influenced the outcome. In addition to its use with the aid of expert

**Figure 1 — AI Governance and Trust Governance Framework**

The diagram presents an overview of the Trust Governance Framework showing the trust and governance relationships needed by a digital society and how AI governance can achieve public trust in AI.



interpretation to improve AI models during the societal implementation phase of AI deployment, it can also foster trust among users during the maintenance and management phase by making the analysis process more transparent to them. Unfortunately, because the outputs of XAI do not always align with what an expert would anticipate, there is also a need for further improvements through research and development aimed at incorporating such expert knowledge more effectively.

Learning from data is a key feature of AI, and as such it is subject to the various biases that might be present in that data. One example is how fairness can be compromised by poor classification accuracy for categories with small sample sizes relative to those with a lot of data. In response, Hitachi is using bias minimization techniques to overcome this problem and prevent AI decision making from accentuating discrimination and prejudice. Data quantity is just one of the factors involved in the problem of classification accuracy being different for different categories, with things like the nature of the data also playing a part. Accordingly, Hitachi is trying a number of different approaches in its research and development work on reducing bias.

Once in use, AI models sometimes require retraining to deal with changes in the environment in which they operate. As AI behavior is a statistical outcome determined by its training data, guaranteeing that an AI will continue to behave in the same way after retraining is difficult. Retraining runs the risk of seriously compromising prediction accuracy. To prevent this, Hitachi uses techniques for assessing whether prediction performance after retraining is consistent with past operation. Moreover, the

safety requirements of AI systems call for consistency to be maintained not only for accuracy, but also across the broader aspects of their behavior.

Public concerns about the use of data are likely to arise if there is no way of tracking which data is used in the development and operation of AI, and how it is manipulated. This calls for the employment of techniques for tracking this “data lineage” to provide greater transparency over the use of data. Data lineage management enhances the trustworthiness of AI training data and enables appropriate management across all processes up to the practical deployment of the trained model.

Research and development at Hitachi takes all of these different considerations into account. The following section provides an overview of research into the building of trust, and this is followed by a section describing research and development intended to underpin AI governance.

## 2. Trust and Governance in Digital Society

This section describes a governance framework for ensuring trust in digital societies together with digital trust practices that give stakeholders the confidence to trust one another and generate new value through collaborative creation.

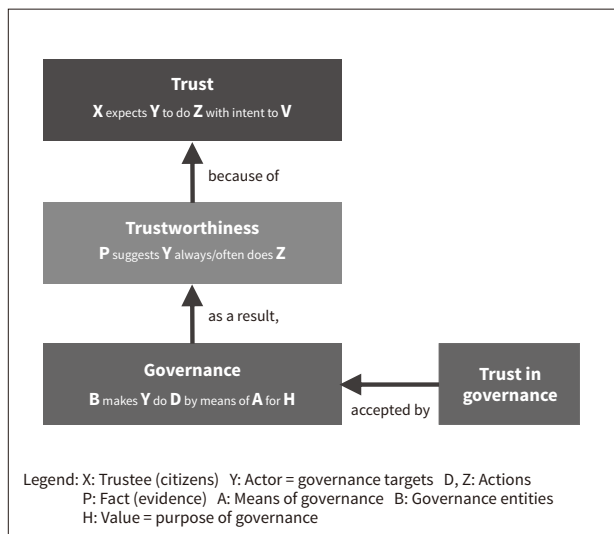
### 2.1

#### Trust Governance Framework

Digital societies have the potential to deliver ongoing improvements in human wellbeing through the interconnection of different services to generate new value. For this

**Figure 2 — Trust Governance Framework**

In addition to formalizing the concepts of trust, trustworthiness, and governance, the framework uses the relationships between them to define a structure for trust that can be utilized in the formulation of trust-building policies.



new value to reach the public, however, the public’s trust must first be earned. As trust is intrinsically subjective, there is no one single way of achieving this. What is needed, rather, is to combine a variety of different trust-building measures to suit different target groups, circumstances, and types of service.

The Trust Governance Framework collates the best ways of engaging with stakeholders to help win public trust. In particular, it is seen as a providing a common standpoint from which to engage in consultation and debate aimed at gaining the trust of multiple stakeholders.

The framework provides one of models to shows how trust, trustworthiness, and governance are interrelated, modeling trust as something that is built up through a process of: (1) implementing governance in ways that facilitate trust building, and, as an outcome of this governance, (2) accumulating evidence that demonstrates trustworthiness, and (3) having stakeholders accept that this evidence does in fact demonstrate trustworthiness (see **Figure 2**).

The Trust Governance Framework also formalizes trust, trustworthiness, and governance, such that the process of formulating measures for winning trust can be made more efficient by substituting the relevant entities into the formalism’s variables (the actors, means of governance, etc.) based on the particular circumstances (who is subject to governance, etc.). One example might be how the *public’s* trust in the *safety of autonomous driving* is achieved through the *government* imposing governance *rules* on *autonomous vehicle manufacturers*, leading to an accumulation of evidence that suggests the *safety of autonomous driving* in the form of *data on traffic accidents* that is in turn accepted by the *public*. Here, the terms in italics represent formalism variables, the values of which are based on the target

groups, circumstances, and type of service involved in this particular example.

The accelerating speed of change experienced by digital societies means that past governance practices struggle to maintain trust once it has been acquired. The above example included governance of autonomous vehicle manufacturers by means of government-imposed rules. In a digital society, however, the flexible and timely updating of rules is needed to keep up with technological innovation and changes in the social environment. Accordingly, the framework also models an agile trust-building process for keeping pace with such change.

## 2.2

### Building Digital Trust

Digital Trust that enables stakeholders in digital society to collaborate has two sides: “Trust of Digital” and “Trust by Digital.”

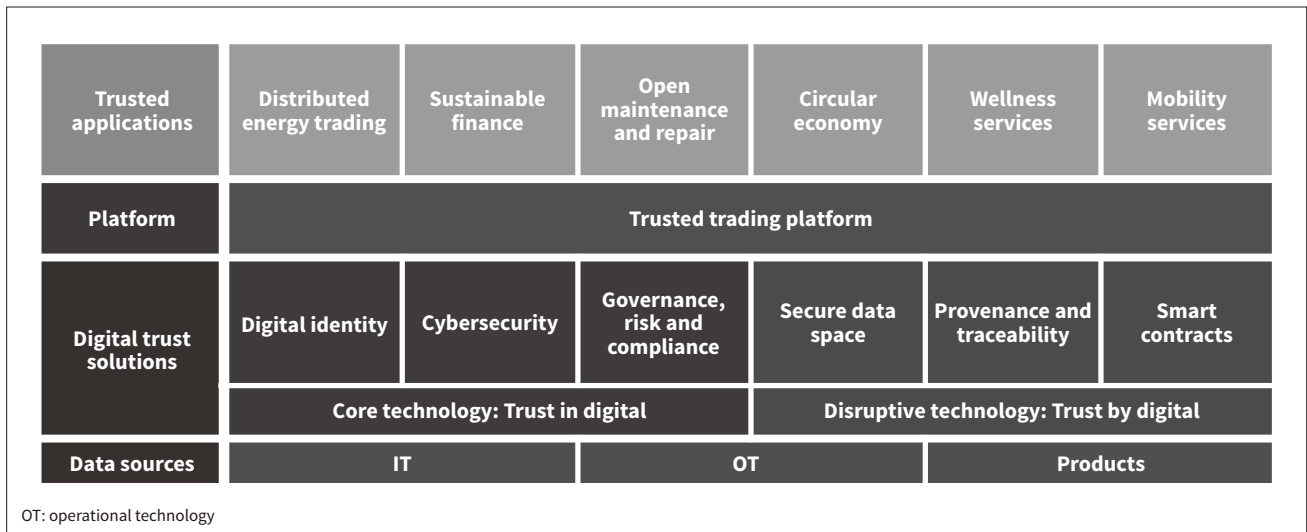
“Trust of Digital” refers to the trust that stakeholders need in order to fully accept and adopt digital systems. This means ensuring that systems and data are safe from attack, abuse, or interference from cyber criminals and hostile agents; that digital rights and privileges are protected through secure biometric authentication; and that all parties comply with local, regional, and global rules and regulations. Two examples of such rules and regulations are the General Data Protection Regulation (GDPR)<sup>(3)</sup>, which was passed by the European Commission and came into force in 2018, and the Artificial Intelligence Act (AI Act)<sup>(4)</sup>, which was drafted by the Commission in April 2021 and is expected to come into force by 2023 at the earliest. The AI Act is designed to promote the adoption of trustworthy and ethical AI that provides effective protection of people’s safety and fundamental rights by establishing boundaries around the purpose and application of AI.

“Trust by Digital” refers to how digital systems are themselves central to strengthening and reinforcing corporate, social, and governmental trust. In order to deliver this trust, digital systems that increasingly span traditional organizational boundaries, will be required to provide full traceability, transparency, and digital notarization, so that parties are able to trade and share data with full confidence that all other parties are respecting mutually agreed standards relating to sustainability, ethicality, and safety (see **Figure 3**).

Hitachi Europe Ltd. developed autonomous control software using state-of-the-art AI in the UK autonomous vehicle project HumanDrive<sup>(5)</sup>. One of the key innovations of this technology was the creation of an intelligent data management tool, namely DRIVBAS, an acronym of “driving behaviour analysis software.” DRIVBAS extracts balanced training data from terabytes of human driving data that can then be used in AI learning. This avoids bias in the resulting AI model, which is able to interpret road environments and generate a safe path for the vehicle.

**Figure 3 — Building Digital Trust**

Digital trust is a form of trust that enables the collaborative creation by all stakeholders in a digital society of agreed rules and standards for safety, security, privacy, and ethics on the basis of confidence that compliance will be mutual and collective and achieved with transparency.



### 3. Technologies that Support AI Governance

Advances in AI are prompting calls for its use in mission-critical social infrastructure applications. Given the difficulty of understanding the complex behavior of AI models, however, making the actions of AI easier to understand is one of the issues to be addressed. A number of effective ways exist for achieving fair AI governance so that companies and individuals will accept the technology and use it on an ongoing basis. XAI, for example, sheds light on the AI decision-making process, making it amenable to review by experts and, when necessary, revision. Others include techniques for learning from imbalanced data in ways that prevent these imbalances from causing bias in the AI model’s outputs, techniques for improving consistency over long-term AI operation, and evidential data management to ensure transparency in the AI analysis process. The following sections describe the current state of this research and development.

#### 3.1

##### XAI and Measures for Achieving Fairness

Hitachi has developed multifaceted diagnostic techniques for the models and training data used by XAI that it supplies through its support services for AI installation and operation. In home loan screening, for example, XAI can be applied to the applicant scores output by a black-box AI, indicating which factors influenced the decision or explaining the reasons in terms of precedents.

The basic approach to understanding the behavior of a black-box AI is to quantify which of the model’s input variables played a part in a prediction. This is a long-standing topic in the study and application of statistics,

with numerous approaches having been adopted to the question of what makes particular variables important. In many methods, a variable is deemed important if varying its value changes the output of the model. When considering fairness, however, the problem is that correlations in the data can result in outcomes that are disadvantageous to a particular group even when variables that are indicative of that group are excluded from the input of the AI model. An analysis of this phenomenon calls for an approach based on the idea that a particular variable is important if knowledge of its value is informative for estimating the model’s output. The cohort Shapley value was developed to provide a way of doing this (see **Figure 4**).

Along with the problem of model output being different for different groups, there is also a need when addressing the issue of AI fairness to consider the tradeoffs between different indicators, such as the negative consequences of incorrect predictions<sup>(6)</sup>. To this end, analysis considers the various different definitions of fairness defined in terms of both black-box AI predictions and data on the correct answers. In this instance, which considers the disparity between groups, in terms of prediction results and false positive predictions, using the cohort Shapley value enables the analysis not only of conventional group-wide bias, but also of which particular instances contribute to bias, which is done from the histogram of values<sup>(7)</sup>. In the future, Hitachi intends to support the development and operation of trustworthy AI models through the use of analysis techniques like this in consultation with domain experts.

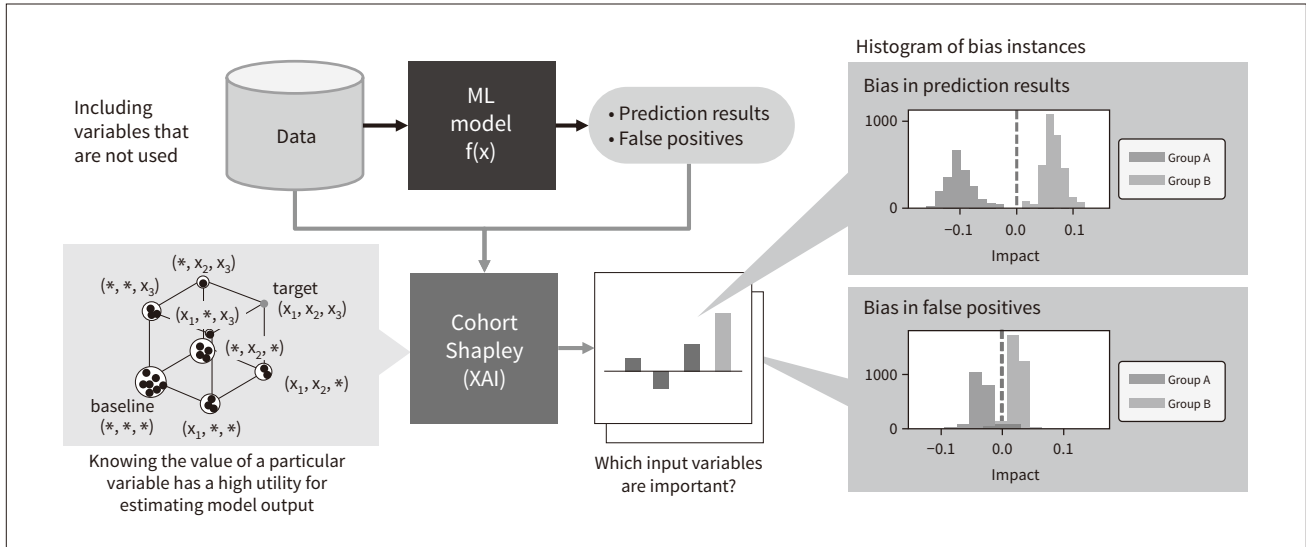
#### 3.2

##### Mitigating Bias when Learning from Imbalanced Data for Achieving Fairness

AI has seen significant progress over recent years, not only due to advances in algorithms and computing hardware, but

**Figure 4 — Cohort Shapley Value**


When applied to the deviation between prediction results and false positive classes, the histogram of cohort Shapley values provides a way to analyze which particular instances contribute to bias, in addition to conventional group-wide bias.



also thanks to the availability of large open datasets that can be accessed by anyone. Unfortunately, these large datasets may contain unintended biases, giving rise to unexpected

**Figure 5 — Distribution of Data Quantities in Action Recognition Dataset and Recognition Accuracy**

The upper graph shows the number of data points for each of the action categories in an action recognition dataset (the EGTEA dataset<sup>(11)</sup>) and the lower graph shows corresponding recognition accuracies for an action recognition model trained using the dataset. Despite the presence of a large imbalance in the number of data points, this demonstrates that categories with a low number of data points do not necessarily suffer from poorer accuracy.



inequities when the models are deployed in practice<sup>(8)</sup>. Bias in the distribution of data is a typical example. A team led by US computer scientist Joy Buolamwini looked at a dataset that had been created by the US government for use as a facial recognition benchmark with explicit consideration given to ensuring geographic diversity in the photographs it contained. What they found, however, was that the dataset contained unambiguous bias, with a very low proportion of photographs showing women with dark skin compared to men with light skin<sup>(9)</sup>. It was already widely understood that, if an AI is trained using biased datasets such as this, the resulting model will perform particularly poorly on minority groups<sup>(10)</sup>.

Hitachi has been researching and developing AI techniques that can be used, even when models are trained using imbalanced datasets, to deliver fairer outcomes by minimizing the loss of AI model accuracy for data that belongs to categories present in small numbers. The traditional way of addressing this issue when using imbalanced datasets has been to seek to overcome the imbalance in the training process by considering the number of data points in each category, increasing the weighting of under-represented categories and reducing that of well-represented categories. Behind this lies the hypothesis that the fewer data points are available, the more difficult it is to train the model using data for that category. Hitachi, however, has demonstrated experimentally that this hypothesis is not always true (see Figure 5).

Based on this discovery, Hitachi has proposed a method for correcting imbalance in model training by continuously monitoring the difficulty of learning for each category and adjusting the category weightings accordingly. At the time of publishing a paper, it had already been demonstrated that this method outperformed past methods when learning from imbalanced datasets<sup>(12)</sup>.



## 3.3

### Techniques for Improving Consistency over Long-term AI Operation

Even as greater use is made of AI in industrial systems, human operators still have a key role to play. If an AI-driven system is not trusted, it can lead to a myriad of issues ranging from low productivity to safety and financial impact. For example, an un-trustworthy robotic arm in a manufacturing shop can be hazardous to operators and an untrustworthy driverless vehicle can cause accidents.

To increase trustworthiness in AI, Hitachi America, Ltd. has observed that for current systems it involves two techniques: (1) human-in-the-loop and rule-based safeguarding mechanisms, and (2) modifications to model training methodologies. While explainable models represent an innovative new approach, development is still in its early days. Human-in-the-loop and rule-based safeguards are straightforward, but they are also very domain- and use-case-specific. The more robust approach would be to modify the model training methodologies as these are fundamental in nature and can be applied to a wide variety of use-cases and domains. Accordingly, they are the focus of research.

Trust can be considered as a function of consistent behavior. From an AI perspective, it means that, given the same input, the user would expect the same output, especially for correct outputs, or in other words consistently correct outputs. In the case of industrial systems and mission critical systems, inconsistency in model outputs can have adverse effects on business processes along with safety issues. So, Hitachi America has studied model behavior in the context of periodic retraining of deployed models, where the outputs from successive generations of the models might not agree on the correct labels assigned to the same input (see **Figure 6**).

Model consistency is defined as the ability to make consistent predictions across successive model generations for the same input. While consistency is applicable to both correct and incorrect outputs, producing consistently correct outputs for the same inputs is the more desirable case. Hitachi defines “correct consistency” as the ability to make consistent correct predictions across successive model

generations for the same input. To quantify (correct) consistency, metrics have also been developed to measure these. For deep learning models, this has involved the development of an efficient ensemble learning technique called the dynamic snapshot ensemble method<sup>(13)</sup> and obtaining theoretical proof of the conditions under which the consistency of deep learning models is improved. With this research, Hitachi recommends that these proposed metrics be used in addition to existing aggregate metrics like accuracy to evaluate the AI systems before deployment, and that they be part of AI-driven system governance.

## 3.4

### Evidential Data Management for Transparency in AI Analysis Process

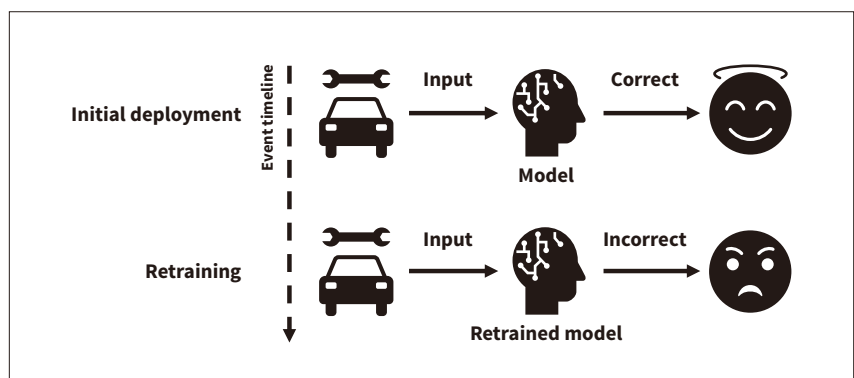
With the increasingly integral role of AI in society, a growing proportion of the general public are becoming concerned about how their personal data is being used. One example might be if an AI predicts that a person has a 45% probability of needing nursing care in 10 years’ time. In the absence of any evidence as to what the 45% figure actually means, and what the person can do to bring the probability down, all this will do, unfortunately, is make the person feel anxious. Alleviating this anxiety requires transparency as to what data the AI used as the basis for its prediction, and how that data was processed. This will provide experts with a deeper understanding of the AI prediction process so that they can ascertain what outputs actually mean and explain this to the public. Doing so will make it easier for the public to find more appropriate ways of avoiding the need to go into nursing care. Specifically, this means being able to trace back to the sources of data that formed the basis of the prediction to determine which factors went into the predicted 45% probability of needing nursing care and which factors have the potential to bring the probability down.

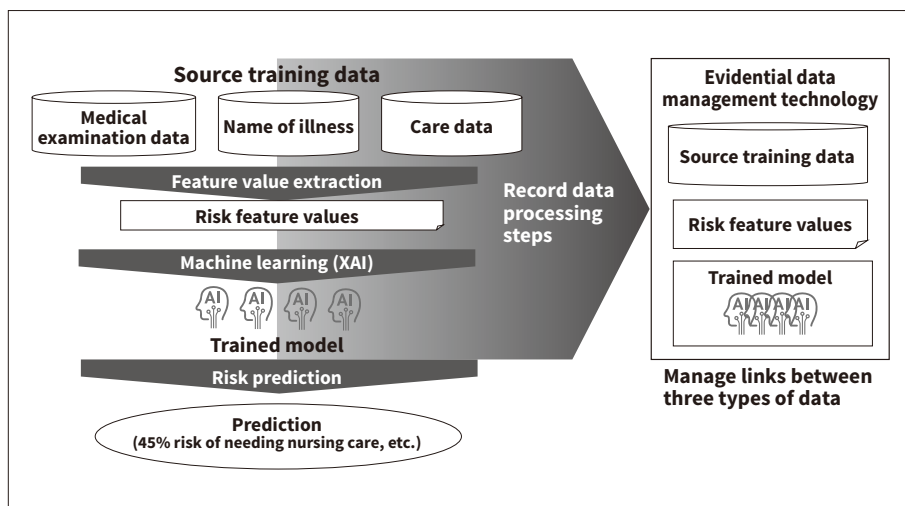
This approach of managing the evidential data associated with AI predictions should provide trust and confidence in the technology.

Hitachi has developed this “evidential data management” through collaborative creation with Partners HealthCare, a US healthcare provider. The objective of the project was to

**Figure 6 — Inconsistency in AI Model Prediction after Retraining**

While AI models require ongoing retraining to maintain performance after entering service, this retraining sometimes makes their predictions inconsistent (changes prediction outcomes).





**Figure 7 — Overview of Evidential Data Management Technology (Example for Medical and Elderly Care Sector)**

Evidential data management technology works by recording the data processing steps, from the source training data to determining the risk feature values and creating the trained model, and then managing this data in a way that links these three types of data together. This enables rapid analysis to determine which source data explains a prediction outcome, such as a person having a 45% risk of needing nursing care.

avoid the re-admission of hospital patients within 30 days of discharge by establishing practices for providing appropriate follow-up targeted at those discharged patients who exhibit more than a predetermined level of risk. Along with predicting the probability of re-admission, this also requires the provision of information about which data served as the reason for the prediction. This was done by recording all steps in the AI analytics process, covering the creation of the trained model on the basis of risk feature values that were generated from the training data, and developing a technique for managing data in a way that links these three types of data together (training data, risk feature values, and trained model) (see **Figure 7**). Named “evidential data management technology,” information about this technique was announced in December 2018<sup>(14)</sup>.

The technique has since been used in a variety of projects and it is being further enhanced to turn it into a generally applicable technique for data management. Its potential is increasingly being recognized, especially given the widespread acknowledgement in recent years of the importance of data quality control in AI. Examples of where evidential data management technology is being called for include the requirement to use complete data in the European Union (EU), where an AI Act is currently under consideration, and a policy paper by the Japan Electronics and Information Technology Industries Association (JEITA)<sup>(15)</sup> that offers a realistic counterproposal whereby the use of data lineage to demonstrate data trustworthiness, consent management techniques (evidential data management), and other such practices will provide transparency as to the source of data and facilitate the more appropriate management of processes from data processing to machine learning and use of the outcomes of machine learning. As it works to create a world in which the general public is able to enjoy the benefits of AI with confidence, Hitachi intends to continue pursuing the wider adoption of evidential data management technology as a form of data management that is compatible with AI ethics.

## 4. Conclusions

This article has described a framework for trust and governance in digital societies together with techniques that facilitate AI governance.

Achieving public trust in AI is a prerequisite for Hitachi’s Social Innovation Business, and to this end it is establishing AI governance on the basis of Lumada and its AI Ethical Principles. Hitachi is also pursuing a wide range of research and development, encompassing topics such as AI quality assurance and AI-specific privacy protection and security measures not covered in this article. By combining these different approaches, it is helping to create a safe, secure, and resilient digital society and to foster human wellbeing.

## Acknowledgements

Hitachi would like to acknowledge the significant assistance received in the development of the Trust Governance Framework described in this article, notably from Keita Nishiyama, Visiting Professor, the University of Tokyo. Similarly, the cohort Shapley value described in this article is the result of joint work with Professor Art B. Owen and Benjamin B. Seiler of Stanford University. The authors would like to express their deep gratitude.

## References

- 1) Ministry of Internal Affairs and Communications, “Conference toward AI Network Society” in Japanese, [https://www.soumu.go.jp/main\\_sosiki/kenkyu/ai\\_network/index.html](https://www.soumu.go.jp/main_sosiki/kenkyu/ai_network/index.html)
- 2) World Economic Forum, “Rebuilding Trust and Governance: Towards Data Free Flow with Trust (DFFT)” (Mar. 2021), <https://www.weforum.org/whitepapers/rebuilding-trust-and-governance-towards-data-free-flow-with-trust-dfft>
- 3) European Union, “General Data Protection Regulation (GDPR),” <https://gdpr.eu/>

- 4) European Union, "Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts," <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206>
- 5) Hitachi Europe Ltd., "HumanDrive Project Achieves UK's Longest and Most Complex Autonomous Journey," <https://www.hitachi.eu/en/press/humandrive-project>
- 6) A. Chouldechova, "Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments," *Big Data*, 5 (2), pp. 153–163 (Jun. 2017).
- 7) M. Mase et al., "Cohort Shapley Values for Algorithmic Fairness," Technical Report, arXiv: 2105.07168 (May 2021).
- 8) A. Caliskan et al., "Semantics Derived Automatically from Language Corpora Contain Human-like Biases," *Science*, 356 (6334), pp. 183–186 (Apr. 2017).
- 9) J. Buolamwin et al., "Intersectional Accuracy Disparities in Commercial Gender Classification," *Proceedings of Machine Learning Research*, 81, pp. 77–91 (Feb. 2018).
- 10) N. Japkowicz et al., "The Class Imbalance Problem: A Systematic Study," *Intelligent Data Analysis*, 6, pp. 429–449 (Oct. 2002).
- 11) Y. Li et al., "In the Eye of Beholder: Joint Learning of Gaze and Actions in First Person Video," *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 619–635 (Sep. 2018).
- 12) S. Sinha et al., "Class-Wise Difficulty-Balanced Loss for Solving Class-Imbalance," *Proceedings of the Asian Conference on Computer Vision (ACCV) 2020*, pp. 549–565 (Feb. 2021).
- 13) D. Ghosh, "Wisdom of the Ensemble: Improving Consistency of Deep Learning Models" (Dec. 2020), <https://www.hitachi.com/rd/sc/aiblog/035/index.html>
- 14) Hitachi, Ltd., "Development of Information Dashboard that Can Present AI-based Predictions of Patient Re-admission Risk and the Data on which They are Based" (Dec. 2018) in Japanese, <https://www.hitachi.co.jp/rd/news/topics/2018/1212.html>
- 15) European Commission, "Artificial Intelligence – Ethical and Legal Requirements: Feedback from: Japan Electronics and Information Technology Industries Association (JEITA)" (Aug. 2021), [https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12527-Artificial-intelligence-ethical-and-legal-requirements/details/F2665415\\_ja](https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12527-Artificial-intelligence-ethical-and-legal-requirements/details/F2665415_ja)

## Authors



### Naokazu Uchida

Media Intelligent Processing Research Department, Center for Technology Innovation – Advanced Artificial Intelligence, Research & Development Group, Hitachi, Ltd. *Current work and research:* Research and development of language models and dialogue systems. *Society memberships:* The Japanese Society for Artificial Intelligence (JSIAI).



### Tadashi Kaji, Ph.D.

Center for Technology Innovation – Societal Systems Engineering, Research & Development Group, Hitachi, Ltd. *Current work and research:* Research and development of cybersecurity and digital trust. *Society memberships:* IEEE.



### Nick Blake

European Research and Development Centre, Big Data Laboratory, Hitachi Europe Ltd. *Current work and research:* Leading Hitachi's social innovation in smart spaces and digital trust.



### Masayoshi Mase, Ph.D.

Media Intelligent Processing Research Department, Center for Technology Innovation – Advanced Artificial Intelligence, Research & Development Group, Hitachi, Ltd. *Current work and research:* Research and development on interpretability and explainability of machine learning. *Society memberships:* The Information Processing Society of Japan (IPSJ), the IEEE Computer Society, the Association for Computing Machinery (ACM), and the American Statistical Association (ASA).



### Hiroki Ohashi

Intelligent Vision Research Department, Center for Technology Innovation – Advanced Artificial Intelligence, Research & Development Group, Hitachi, Ltd. *Current work and research:* Research on machine learning and computer vision for supporting skills transfer, human error reduction, workers' safety, and productivity improvement in industry. *Society memberships:* IPSJ and JSAI.



### Dipanjan Ghosh, Ph.D.

Industrial AI Laboratory, Research and Development, Hitachi America, Ltd. *Current work and research:* Development of prognostics solutions with artificial intelligence.



### Chetan Gupta, Ph.D.

Industrial AI Laboratory, Research and Development, Hitachi America, Ltd. *Current work and research:* Development of industrial solutions with artificial intelligence.



### Ken Naono, Ph.D.

Data Management Research Department, Center for Technology Innovation – Digital Technology, Research & Development Group, Hitachi, Ltd. *Current work and research:* Research and development of data processing technology in the field of medical and nursing care. *Society memberships:* Director of the Japan Society for Industrial and Applied Mathematics (JSIAM) and a member of IPSJ.



### Mika Takata

Data Management Research Department, Center for Technology Innovation – Digital Technology, Research & Development Group, Hitachi, Ltd. *Current work and research:* Research and development of data management and machine learning model management technologies. *Society memberships:* JSIAI and IPSJ.