


What Is “Trustworthy AI”?

—Young Fujitsu AI engineers discuss AI issues and AI's future—



Now that we have entered the third AI boom with the evolution of machine learning technologies such as deep learning, research and development for practical applications of AI are progressing in various fields, but at the same time questions as to whether AI can be trusted are increasingly arising. For example, research on explainable AI (XAI), which clarifies the way in which AI draws conclusions, is actively being carried out around the world. Further, at the G20 Ministerial Meeting on Trade and Digital Economy held prior to the G20 Osaka Summit, “AI Principles” summarizing what AI should be were drawn up. For its part, Fujitsu formulated the [“Fujitsu Group AI Commitment”](#). Against this backdrop, Fujitsu’s AI engineers continuously ask themselves what they can do to make the AI they develop trustworthy for customers, partners, and even end users. As part of this quest, a round-table discussion with four young AI engineers who are active at the forefront of AI application projects was organized to clarify what trustworthy AI is through their efforts.

[Participants]

Masaki Kakisako (joined Fujitsu in 2015)
Fujitsu Limited
Data Analytics Department, Professional Services
Division, Data and AI Business Unit

Yuichi Ike (joined Fujitsu laboratories in 2018)
Fujitsu Laboratories Ltd.
Artificial Intelligence Laboratory, Autonomous
Machine Learning Project

Takamori Kanao (joined Fujitsu in 2016)
Fujitsu Limited
NLP Business Department, AI Software
Development Division, Software Business Unit

[Facilitator]
Yuzuru Yamakage
Fujitsu Limited
Data Analytics Department, Professional Services
Division, Data and AI Business Unit

Toru Fujita (joined Fujitsu in 2017)
Fujitsu Limited
AI Service Department, Professional Services
Division, Data and AI Business Unit

The department, division and business unit names that appear in this article are current as of the time of the round-table discussion.

[Part 1] Reasons Why These AI Engineers Chose Fujitsu

—Today, we have gathered here young AI engineers who are active in various divisions of Fujitsu. I think that this is our first meeting although we all work in the same field, so let's start by introducing each other first.

Ike: This is my second year with the company, I work at the Artificial Intelligence Laboratory at Fujitsu Laboratories. Utilizing my experience as a graduate school math major, I am working with [Inria](#), a French national research institute, on the study of new data analysis methods using topology.

Kakisako: Speaking of Inria, it is one of the leading research centers in Europe. It is amazing that you have been carrying out joint research with such a big name since first joining the company.

Ike: No, not at all (laughs). When you belong to a research lab, it's hard to get a sense of the business scene, so today I'm eager to learn from you guys from the Business Unit. Kakisako-san and Fujita-san, you both belong to the Data and AI Business Unit. What kind of work do you do?

Kakisako: Our Data and AI Business Unit provides consulting services for customers who want to use AI in their business. Starting with interviews about their business content, we perform total solutions that cover everything, from the analysis of issues to be solved by AI, to data learning and evaluation, and demonstration and verification. I belong to the Data Analytics Department and mainly promote co-creation projects using AI with customers.

Fujita: I belong to the AI Service Department in the Data and AI Business Unit. I am mainly in charge of verification support and introduction services for image-based deep learning using the Zinrai Utilization Support Service, which provides [FUJITSU Human Centric AI Zinrai](#) to corporate customers. During my student days, I was doing research using GPUs, so I wanted to work in a business where various image analysis AIs could be applied to actual tasks and businesses, so I sought to be assigned to an AI-related department. This is my third year in the company, and I've been in this department since the beginning, but five years

ago, when Kakisako-san joined the company, the Business Unit had not been launched yet.

Kakisako: Yes, that's right. When I joined the company, big data was popular. I joined the company as a data scientist and was assigned to a department that provided marketing support using big data analysis. I transferred to my current department in my third year with the company.



Masaki Kakisako, Fujitsu Limited

Kanao: I too transferred like Kakisako-san. This is my fourth year with the company, but since my second year, I've been with the Software Business Unit, where I'm developing a FAQ search system that uses AI.

Fujita: I've been in the same department since I joined the company, so I'm jealous of you guys for having had a wide range of experience in multiple departments. Kanao-san, is the work you are doing now different from the work you did at your first department when you joined the company?

Kanao: Yes. When I joined the company, I was assigned to the Incubation Center, an organization completely separate from the Business Units, and I was able to work freely in a unique workplace with just 20 workers, all of them newcomers (laughs). In that department, I was working on Proof of Value (PoV) through agile development, and specifically I was in charge of developing screens for managers that would "visualize" the operation status of data centers.

—Now that we have heard your profiles, I would like to ask you one more question before we get started on the main topic, and that is: What motivated you to join Fujitsu? I believe that each of you had a number of choices of workplaces, so what led you to choose Fujitsu?

Fujita: The reason I became interested in Fujitsu is that it struck me as an interesting company at the career fair I attended. Especially the company's attitude: It's ambitious, but also frank and honest, as symbolized by the 3Gs (Gemba, Genbutsu, and Genjitsu, meaning actual place, actual things, and actual activities), and I felt that this closely matched my own views.

Kakisako: Do you feel that there is a gap between your impression of the company before and after you joined?

Fujita: I think that there are some things that could be done in a smarter way, but the company turned out to be frank and honest exactly as I had imagined. My current job gives me many opportunities to interact with customers, so I particularly feel it to be customer-oriented. I truly appreciate the attitude of providing valuable things to our customers.

Kanao: What attracted me to Fujitsu is how it works with customers in many different fields. I think there is strength in specializing in one area, but the issues of customers are multifaceted and there is not always a single solution. I thought that having a variety of skills and being able to bring them together to solve problems was a great strength of the company.



Takamori Kanao, Fujitsu Limited

Ike: I agree, but there are other companies besides Fujitsu that also have a wide range of technologies.

What was the decisive factor that led you to choose Fujitsu?

Kanao: It may sound strange, but color was the decisive factor. I like the color red.

Everyone: That was it?!

Kanao: Yup, color is important (laughs). When looking for employment, how you get along with the Human Resources people also counts for a lot. In that sense, I thought that liking the company's color was important, and I was interested in the spirit of challenge, the humanity, and the exciting corporate culture, all expressed by the company's corporate color, "FUJITSU Red".

Kakisako: I see. Kanao-san, was there a gap between your impression of the company before and after joining?

Kanao: It being a large company, I expected some stuffiness and rigidity, but actually it was quite different. I think that it may have to do with the fact that my initial assignment was the Incubation Center as I mentioned earlier, but I was able to do what I wanted in an at-home atmosphere. After transferring to the Business Unit, I found out that I couldn't do just as I liked (laughs), but I want to make the most of the merits of agile development, including the flexibility and nimbleness that I had cultivated at my first workplace.

Fujita: And you, Kakisako-san, did you feel there was a gap between your impression of the company before and after joining?

Kakisako: I felt a gap, but in a good way. I joined the company as a data scientist, and I always had in mind a career analyzing data. However, my job actually gives me lots of opportunities to communicate with customers and I am getting increasingly involved in business, and I feel that this suits me fine.

Kanao: Communicating with customers has its challenges, but it can also be a tremendous motivation to work, right?

Kakisako: Absolutely. In the case of a development-oriented company, there is this image of the roles being divided between the department that actually carries out the development work and the department that does the business planning with customers, but my current work is interesting because it allows me to experience both sides.

Ike: When I listen to you all, it looks like you have contact with customers and you get direct feedback of your

output, which is something that I'm envious about. When you work in a lab like me, such opportunities are quite rare.

Kakisako: Ike-san, what motivated you to join the company?

Ike: From a researcher's point of view, what I found attractive was the focus not just on the immediately useful, but also on basic research. I think both are important, but I'm rather interested in the latter, and since joining the company, I've been able to do what I want from the beginning.

Fujita: So you mean to say that joint research with Inria was what you aimed for even before joining the company?

Ike: Originally, my collaborative research with Inria started before I joined the company, and I was told in an interview that "If you join, you can go to France," so I was looking forward to this and joined the company (laughs). In actual fact, I had the opportunity to stay at Inria for two months and concentrate on my research.

[Part 2] AI Development in Various Fields

—Now that you have learned each other's profiles in Part 1, could you tell us about the AI projects you are currently working on?

Ike: The theme of my joint research with Inria is TDA (Topological Data Analysis), which applies topology to the analysis of time-series data. More specifically, we try to read features from the shapes shown by the time-series data to allow AI to perform tasks that until now required the knowledge of experts.

Kakisako: Specifically, what kinds of data do you handle?

Ike: We handle a wide range of data, such as data of sensors embedded in bridges, from which anomalies can be detected, and data from ECGs and the like, from which arrhythmia can be detected, for example.

Kakisako: That sounds interesting. But in the case of joint research, isn't how to share the results with the other parties problematic? I too have some experience in joint research, and divvying up things can be difficult, right?

Ike: That's right. Right now, I am writing a research paper, and I am struggling with that. I'd love to hear your own experience and use it as reference.

Kakisako: I am in charge of a co-creation project with a major measuring equipment manufacturer, in which we are working to equip a mass spectrometer, which examines the ingredients in foods, beverages, and pharmaceuticals, with a function that automates the task of peak picking for data analysis with AI.

Kanao: Peak picking means reading the height and width of waveforms (peaks) from obtained data graphs. I've heard that it takes a lot of work.

Kakisako: At laboratories that handle large amounts of data, skilled technicians spend hours on this each day, so we are trying to automate this with AI. However, when reading waveforms, the reading range varies from person to person and the criteria are ambiguous, so we had to create criteria together with the customer. Also, when installing the finally completed AI, there were hardware limitations. That is, having to make do with the processing speed of a normal PC was another difficulty.

Ike: Just because you develop a new technology doesn't mean that it can be commercialized, which is something that one doesn't easily see when working in a laboratory as I do.

Kakisako: Yes, exactly. There is also the difficulty of divvying up the results, as mentioned earlier. In this particular case, the product would be sold as the customer's product, but the problem was how to share the profit. We solved this through many negotiations, thanks in part to our relationship of trust.

Ike: What is needed to make such co-creation successful in business terms?

Kakisako: In my opinion, what's needed is people who can apply the technology generated in the laboratory to actual business. Although such human resources were scarce until a few years ago, we have been enhancing training for the entire Business Unit, so I think that there will be more and more opportunities to use new technologies for business in the future.

—Now that it has expanded into a variety of topics ranging from the difficulty of AI development to business issues, this has become a pretty interesting discussion. Kanao-san, what kind of project are you in charge of?

Kanao: It's a project to develop an AI-based FAQ search system to boost the competitiveness of call center companies. Call centers often create FAQs based on past

inquiry data so that operators can respond quickly to inquiries from users. Operators search the FAQs based on the content of the question and find the best answer, but it takes skill to master the huge number of FAQs, so the aim is to use AI to support operators.

Kakisako: Specifically, how is support provided?

Kanao: AI for FAQs has been evolving rapidly in recent years. In the past, AI was capable only of determining which were the more appropriate QAs from the searched QAs and displaying them at the top. What we are developing now makes it possible to perform searches while displaying the verbal exchanges between the user and the operator in real time using voice recognition.

Kakisako: Does the UI display conversations as alternating text messages, like chat apps?

Kanao: Yes, that's right. Ideally, we would like to develop this technology to the point where AI analyzes the information converted to text and finds the appropriate answer on its own; but this is still difficult, so the current arrangement is to have the operator perform searches by clicking the displayed text while conversing.

Fujita: Isn't it difficult to perform searches while interacting with the user?

Kanao: Yes, it is; so we needed to make the search process as simple as possible. Moreover, the operator's dialogue includes a lot of conversational gap fillers and interjections, and the text-converted conversation flows rapidly, making searches difficult. We sought to deal flexibly with such problems.

Ike: I see. We need not only to apply AI to business, but also to develop a UI that is tailored to the actual use site to ensure ease of use.

Kanao: Yes. Through this project, I realized that for customers, improving the usability of AI tools is actually more important than the accuracy of the AI.

Fujita: I think that it is very important not to make AI an accuracy battle. The engineers who develop AI tend to focus on accuracy no matter what. Of course, high accuracy is a selling point in business terms, but high accuracy does not necessarily translate into high ratings by customers.

Kanao: That's true. Accuracy and ease of use are completely different things. The criteria by which the customer evaluates AI depend on what the customer wants from AI, so I have recently become keenly aware that if you do not control customer expectations, your business

will not go well. How about you, Fujita-san? Have you felt the same kind of thing in your own projects?

Fujita: I am currently in charge of manufacturer customers to whom we provide verification support services for image analysis AI for safety measures in factories. The verification support service started at the end of last year, and we have run a number of data collection and learning model trials. Right now, using image data captured in the factory, we are currently verifying the validity of our prototype AI model. The person in charge at the customer side is eager to introduce AI, but of course they are not an AI expert, so they have no concept of specific evaluation methods and indicators. The latest issue is how to evaluate the verification results, and I feel that customers need to understand the basics of AI.



Toru Fujita, Fujitsu Limited

Kanao: Certainly, the person in charge, as well as the management that makes the final decision on whether to introduce AI into the actual business, need a minimum of knowledge to make their decision.

Kakisako: "Teaching" may not exactly be the right term, but an important part of our job is to deepen our customers' understanding of AI while interacting with them on a daily basis. If they don't understand AI, they will not get to know the value and convenience of AI, so it's important that we communicate with customers on a daily basis and answer all their questions. Conversely, it is very important for us that our customers tell us about their technologies and know-how, as we can learn a lot from this. That's how important communication is.

Ike: For someone like me who doesn't have the opportunity to interact with customers, this is quite informative.

–You all seem to have thought about AI in depth through the various issues you have run into while working on distinctive projects. Next, based on these experiences, let's finally discuss the main topic, "Trustworthy AI".

[Part 3] What Is "Trustworthy AI?"

–Our discussions thus far contained a number of hints as to our main topic, which we will finally start discussing now. However, the simple expression "AI that Can Be Trusted" actually involves various aspects. First is the reliability of AI as a technology. Another is reliability in terms of the ease of use that the AI being introduced into a business brings the customer. Additionally, there is also the reliability of the company that offers the AI. What are your respective thoughts about these various aspects?

Ike: Within the context of my joint research with Inria, my research topic is directly related to the theme of this round-table discussion: it is explainable AI, which has become a hot topic in recent years. For example, if TDA is used to detect abnormalities in time-series data but the reason why something has been determined to be abnormal cannot be explained, that technology will not be trusted by the customer. I am researching how to create such a mechanism.



Yuichi Ike, Fujitsu Laboratories Ltd.

Kanao: Explainable AI is something I've been interested in for quite some time, but in reality, what kind of

output are we talking about?

Ike: That is a very sharp question (laughs). The difficulty lies in the fact that there is no definition as to what is explainable. One approach is to explain from the results that come out of the black box, while another is to clarify the content of AI proper, and given the large variety of choices, we are just at the stage of trying to figure out what we should do.

Fujita: When we say "explainable AI", of course, it is not the AI that does the explaining, but humans. In other words, humans translate the results of AI, but, even there, there is already a gap between us developers and society at large, including our customers.

Ike: In regard to this point, I think how to explain the answers provided by AI in an easy-to-understand manner will be an important theme for AI going forward. Particularly when deep learning is used, AI tends to become a black box, so this is a question that will come up.

–Kakisako-san, what do you think about the reliability of AI in light of your experience of actually developing products equipped with AI?

Kakisako: I think that there are two phases to getting people to trust AI: understanding AI and then using AI. As mentioned earlier, we first explain to customers what they can do with AI and make them adjust the level of their expectations. It is important to discuss with customers what AI can and cannot do.

Fujita: Definitely. Unless we make customers understand AI, we will be asked, "AI can do anything, right?" In reality, using AI for business is not easy. This is because we need to first understand the customer's business, then identify the problem spots, and then make proposals on how to solve these problems with AI, what data to collect, and how to process it using deep learning.

Kakisako: And to get from there to the phase where AI is actually used, I think that the conventional approach is not enough. The conventional way is to submit a report to the customer and then be done with the project, but even if the customer understands what can be done with AI, it is not easy for the customer to decide to actually use it.

Kanao: Yes, indeed. In order to apply AI to business, it is necessary not only to understand the theory, but also to have the trust of the actual user.

Kakisako: And the means of overcoming that hurdle is the prototype. It's only by getting the customer to actually experience the results of AI and see how usable it is that you truly get the customer's trust in the true sense of the word. It was terribly difficult when I actually tried it (laughs), but I felt it was important to carry out these two phases carefully.

Ike: Yes, but even if the decision to introduce AI is made this way, the next hurdle is how to actually incorporate AI into products.

Kakisako: Yes. The most important aspect is the quality assurance of AI, but it is very difficult to delineate how far the quality can be guaranteed. In particular, in my project, the customer sells their products all around the world, so I needed to make the customer understand how they could explain the products' quality to the end user.

—Both Fujitsu and Kakisako-san's customer were satisfied with the level of quality assurance achieved in the project and said "Neither company tolerates any compromises, so the fact that we were able to bring this product to a level that satisfied both sides makes this project truly epoch-making." Kanao-san, what about you?

Kanao: When customers apply AI, they don't use it on a standalone basis but rather as a system where AI is combined with other functions. From the perspective of customers and end users, the AI and other functions are not used separately, so in my opinion improving the reliability of the entire system is what eventually leads to the reliability of the AI.

Ike: Certainly, AI is one of the functions that make products easier to use, but AI itself is not the main thing.

Kanao: From that viewpoint, I think it is also necessary to have the flexibility of not seeking to do all things with AI. In my own project, there was a Q&A item that exactly matched the text searched for, but sometimes that search result did not come out on top, so the customer could only ask why. Judging that it was not a matter of AI accuracy, we provided a function to display any Q&A item that was a perfect match at the top of the list, and the customer was pleased.

Fujita: That result was because AI learns not only FAQs but also the content of answers. Still, from their

viewpoint, customers may well tell you that AI can't be trusted with such results.

Kakisako: Even in the measurement equipment project I was involved in, although the data was clear, AI sometimes made a judgment different from the correct answer, and the customer would ask us why. How you explain this greatly affects the level of customer trust.

Kanao: This experience made me realize that just because a system uses AI, it is not necessary to rely on AI for everything, but rather there are situations where it is better to deal with conventional programs. Given the variety of customers, thinking about wherein lies the value that the customer wants is likely to lead to improvements in the reliability of AI.

—Lastly, I would like to ask Fujita-san, who has experience in AI consulting, to give us his views from the perspective of consulting.



Yuzuru Yamakage, Fujitsu Limited

Fujita: In the case of AI consulting, customers may have excessive expectations for AI. This being business, I want to say "Of course, it's okay," but one can't let customers entertain excessive expectations. For that reason, I try to make proposals that reflect actual conditions.

Kanao: By excessive expectations, do you mean expectations regarding functional aspects of AI? Or do you mean expectations regarding accuracy?

Fujita: Both, but the most difficult is accuracy assurance. Naturally, we try to improve accuracy, but I think that our job is to propose ways to use AI even if the accuracy is not 100%.

Kakisako: We want to meet customer expectations, and also contribute to sales, so it is only natural to tend to just say “Yes, that can be done”, but readily giving customers expectations produces negative outcomes.

Fujita: If results later fall short of expectations, we will be in trouble ourselves and also cause inconvenience to our customers. Fujitsu has a culture of not overcommitting, precisely because we care for our customers. Having the honesty of promising only what we deliver results in trust, right?

—We were able to hear many interesting stories and opinions about “AI that Can Be Trusted,” and I’ve come away with the feeling that besides offering technology that further raises trust in AI, it is also important that we increase the customer’s trust in us who offer AI. Lastly, could each of you give us your impressions of today’s discussion?

Fujita: This coming together of engineers who work on AI from various positions, from the most advanced laboratories to the forefront of business, to talk from various perspectives, was very helpful. I plan on applying the stories I heard here today to my future work.

Kanao: I really appreciate this rare opportunity. It made me aware of the importance of exchanging various opinions about the way AI should be, so I think it would be nice to have more opportunities to communicate like this.

Ike: Since I belong to a lab, I rarely get to hear from the business units, so I was honestly worried about what kind of topics we’d discuss today (laughs). Being able to hear helpful stories from a variety of perspectives has been very stimulating.

Kakisako: Even though we deal with the same AI at the same company, there are not many opportunities for communication, so let’s get together again!

—Looks like we had a meaningful round-table discussion, which is great. I look forward to seeing what kind of activities all of you who are creating Fujitsu’s AI future engage in from here on out.

All company and product names mentioned herein are trademarks or registered trademarks of their respective owners.



This article first appeared in Fujitsu Technical Review, one of Fujitsu’s technical information media. Please check out the other articles.

Fujitsu Technical Review

<https://www.fujitsu.com/global/technicalreview/>

