

ESOMAR 20

20 questions to help buyers of
AI-based services for market research
and insights

toluna 





Table of contents

Q1 What experience and know-how does your company have in providing AI-based solutions for research?

Q2 Where do you think AI-based services can have a positive impact for research? What features and benefits does AI bring, and what problems does it address?

Q3 What practical problems and issues have you encountered in the use and deployment of AI? What has worked well and how, and what has worked less well and why?

Q4 Can you explain the role of AI in your service offer in simple, non-technical terms in a way that can be easily understood by researchers and stakeholders? What are the key Functionalities?

Q5 What is the AI model used? Are your company's AI solutions primarily developed internally, or do they integrate an existing AI system and/or involve a third party and if so, which?

Q6 How do the algorithms deployed deliver the desired results? Can you summarize the underlying data and the way in which it interacts with the model to train your AI service?

Q7 What are the processes to verify and validate the output for accuracy, and are they documented? How do you measure and assess validity? Is there a process to identify and handle cases where the system yields unreliable, skewed or biased results? Do you use any specific techniques to fine-tune the output? How do you ensure that the results generated are 'fit for purpose'?

Q8 What are the limitations of your AI models and how do you mitigate them?

Q9 What considerations, if any, have you taken into account, to design your service with a duty of care to humans in mind?

Q10 **Transparency:** How do you ensure that it is clear when AI technologies are being used in any part of the service?




Table of contents

- Q11** Do you have ethical principles explicitly defined for your AI-driven solution, and how in practice does that help to determine the AI's behaviour? How do you ensure that human-defined ethical principles are the governing force behind AI-driven solutions?
- Q12** **Responsible Innovation:** How does your AI solution integrate human oversight to ensure ethical compliance?
- Q13** **Data quality:** How do you assess if the training data used for AI models is accurate, complete, and relevant to the research objectives in the interests of reliable results and as required by some data privacy laws?
- Q14** **Data lineage:** Do you document the origin and processing of training or input data, and are these sources made available?
- Q15** Please provide the link to your privacy notice (sometimes referred to as a privacy policy). If your company uses different privacy notices for different products or services, please provide an example relevant to the products or services covered in your response to this question.
- Q16** What steps do you take to comply with data protection laws and implement measures to protect the privacy of research participants? Have you evaluated any risks to the individual as required by privacy legislation and ensured you have obtained consent for data processing where necessary or have another legal basis?
- Q17** What steps do you follow to ensure AI systems are resilient to adversarial attacks, noise and other potential disruptions? Which information security frameworks and standards do you use?
- Q18** **Data ownership:** Do you clearly define and communicate the ownership of data, including Intellectual property rights and usage permissions?
- Q19** **Data sovereignty:** Do you restrict what can be done with the data?
- Q20** **Ownership:** Are you clear about who owns the output?

Q1

What experience and know-how does your company have in providing AI-based solutions for research?

In 2018 – earlier than most – we set up a dedicated AI/ML lab of software engineers and data scientists. **We're also the first company in the industry** to launch an insights solution that uses only AI and human intelligence (HI) without surveys: a copy testing solution called ACT Instant that we launched back in 2019. Today, we have a vast team of engineers, data scientists, and product and research experts across the Toluna Group who work on our AI-based capabilities and solutions.



Q2

Where do you think AI-based services can have a positive impact for research? What features and benefits does AI bring, and what problems does it address?

We expect – and have already observed – AI models, tools, and solutions to bring at least two major positive benefits for our clients: increased depth of insights and increased speed to action. Developing AI applications to achieve those benefits requires high-quality data inputs, human-in-the-loop review and feedback, an in-depth understanding of technology, AI algorithms and research design, and a digital-native test and learn mindset – all of which are readily available across the Toluna Group.



Q3

What practical problems and issues have you encountered in the use and deployment of AI? What has worked well and how, and what has worked less well and why?

One thing that has worked well at the Toluna Group is the human review and feedback on the intermediate and final outputs from our AI models. Assessing the accuracy of classifications relative to human coding, reviewing the rationale behind generated outputs, client beta tests, and researchers providing feedback (not only during user acceptance testing but also post-launch) have contributed greatly to the robustness of our AI-based solutions and capabilities.

Another thing that benefits Toluna and our clients is our global footprint and the fact that the teams working on our AI-based solutions reside across the globe. This diversity enables us to stay on top of language and cultural differences as we build our capabilities and solutions.



Q4



Can you explain the role of AI in your service offer in simple, non-technical terms in a way that can be easily understood by researchers and stakeholders? What are the key Functionalities?

At Toluna we use AI across a wide range of areas: detecting low quality or fraudulent survey responses, open-end coding, translations, follow-up question probing, survey routing, generating synthetic data, panel member services, transcript coding, first draft report writing, script or code writing, synthetic persona development, data analysis, ads and package feature coding, and so much more.

For each specific AI-powered solution or capability that we release, we follow the tenets of Explainable AI and provide explanations from simple to complex, depending on the intended audience.

To create an audience-appropriate, accurate methodology statement for each of our AI-based capabilities and solutions, we first create what we call an AI Model Trust Card for each AI model underneath the hood of each capability or solution. The intended audience for our AI Model Trust Cards is AI subject matter experts/champions within client organizations – experts who often do the initial and more detailed vetting of new capabilities or solutions. From these, we create a slimmed down AI Methodology Statement – consistent with industry guidelines for methodology statements – for more general use in client proposals and reports.

Q5

What is the AI model used? Are your company's AI solutions primarily developed internally, or do they integrate an existing AI system and/or involve a third party and if so, which?

Sometimes Toluna builds its own proprietary AI models using publicly available data or our own first-party data. Examples of where we've built our own proprietary models include our Luna chatbot for panel member services, our ACT Instant model for predicting ad success, and our predictive model for sample feasibility and routing.

At other times, we leverage third-party large language models (LLMs) as a base model within our AI-infused capabilities or solutions to perform tasks such as classification, summarization, auto-prompting, or synthetic persona generation. The LLMs we use most frequently are versions of ChatGPT from OpenAI. Our licenses ensure data is kept confidential and not used to train OpenAI models.

And finally, some machine learning algorithms we apply were originally outlined in publicly available academic articles.




Q6

How do the algorithms deployed deliver the desired results? Can you summarize the underlying data and the way in which it interacts with the model to train your AI service?

Adding to the examples shared above, we do not routinely train AI models in that we do not change model weights in the final layer of third-party algorithms. We use RAG (retrieval augmented generation) to ensure that our models are (a) grounded in the data and behaviour of real human beings, and (b) guided by research and market information that makes them relevant and accurate for gaining insights.

Depending on the solution, we use different proprietary sources of data that we own or have consent to use, such as stored panelist information, concept and ad testing benchmark databases, internal surveys, feasibility check logs, and census-type population data, to name a few.



Q7

What are the processes to verify and validate the output for accuracy, and are they documented? How do you measure and assess validity? Is there a process to identify and handle cases where the system yields unreliable, skewed or biased results? Do you use any specific techniques to fine-tune the output? How do you ensure that the results generated are ‘fit for purpose’?

Q8

What are the limitations of your AI models and how do you mitigate them?

Q9

What considerations, if any, have you taken into account, to design your service with a duty of care to humans in mind?

We assess AI model performance and validation not only in terms of model fit or accuracy, but also in terms of the guardrails implemented to mitigate bias or harm.

Accuracy. Human review of AI model outputs plays an important role in our validation process. We are transparent about our accuracy rates.

Bias mitigation. We test our models in multiple languages and countries to ensure that if there are any skews in the training data used by a third-party algorithm those skews are not contributing to any biasing effects in the results we deliver. We also build guardrails into our model applications to mitigate hallucinations.


Duty of care. This is especially important in situations where any output of an AI model is part of an autonomous action towards our fellow human beings. For example, our AI-based probing tool, QProbe – that identifies vague or incomplete open-ended responses and then probes for greater detail through tailored follow-up questions – was designed to ensure that survey participants are not probed on sensitive topics they may mention in their response, unless that topic is crucial to the survey topic (of which the participant has already been briefed). During testing, we also ran analyses to ensure that this prompt was not being served up in a way that would indicate a skew toward any gender, age, ethnicity, or education level.



Q10

Transparency: How do you ensure that it is clear when AI technologies are being used in any part of the service?

We indicate the use of AI in our DIY survey platform and in client reports using this symbol:



In data files, we create a separate field that indicates whether the response is synthetic or not.



Q11

Do you have ethical principles explicitly defined for your AI-driven solution, and how in practice does that help to determine the AI's behaviour? How do you ensure that human-defined ethical principles are the governing force behind AI-driven solutions?

Q12

**Responsible Innovation:
How does your AI solution integrate human oversight to ensure ethical compliance?**



We believe in trustworthy augmented intelligence. As such, two key ethical principles for the Toluna Group are explainability (see our answer to Question #4) and human review/human-in-the-loop.

For Toluna, 'augmented intelligence' – artificial intelligence combined with human intelligence – is the way of the future. Our global footprint and globally distributed engineers, data scientists, and researchers help us keep cultural differences and sensitivities top of mind.

We also take Toluna's company values seriously. One value is embracing empathy, defined in terms of caring and respecting others, inclusivity, and elevating others in an egoless way.

In terms of organizational design, we have a 3-person AI Review team that reviews planned AI initiatives and a team of Product Value managers who ensure our AI Model Trust Cards and AI Methodology Statements are accurate and up to date.

Q13



Data quality: How do you assess if the training data used for AI models is accurate, complete, and relevant to the research objectives in the interests of reliable results and as required by some data privacy laws?

We assess the overall quality of our training data in four ways:

1. Accuracy checks

We verify the accuracy of our data by checking whether the data correctly represents real-world scenarios as they are supposed to. We do this through a combination of automated tools and manual reviews of large samples to catch and correct errors and edge cases. Accuracy is checked both on the input and output of the AI model.

2. Completeness evaluation

An incomplete dataset can lead to biased AI decisions, so we look for gaps in our data and try to fill them. The way this is done is by defining initial expected clusters, such as cultural or demographics based that our model aims to support, and balancing the data between these clusters, which often leads to gathering more data or enhancing existing data sets to cover missing areas. The latter is achieved through a combination of automated tools and manual sampling.

3. Compliance with data privacy laws

We ensure that all data used complies with applicable privacy laws and regulations, protecting personal information and respecting user privacy, as well as being transparent on the use of individual personal data.

We use anonymization techniques when training models. If third-party models are involved, we ensure they don't use information we submit to train their models.

4. Continuous monitoring

We continuously monitor our models and update them as necessary to adapt to new insights or changes in the real world. This ongoing process helps us maintain a high standard of data quality throughout the lifecycle of our AI models. Monitoring is done by automated anomalies detection, client/user feedback, and frequent manual checks of random samples.

Q14

Data lineage: Do you document the origin and processing of training or input data, and are these sources made available?

Toluna's approach is as follows:

1. Tracking data origins

We keep records of where each piece of data comes from. Whether it's gathered from our own data collection, public datasets, collected through partnerships, or generated by users, we document the source. This helps us ensure the data's initial quality and legality.

2. Preprocessing details

We also document how the data is processed. This includes the tools and techniques we use to clean, sort, fill gaps, and prepare data for use. By keeping a clear record, we can fix issues by adjusting processing in the pipeline if necessary.

3. Transparency

For publicly available data, we often provide references or links so that others can access the same datasets. For Toluna proprietary data, we provide as much information as possible within the bounds of confidentiality agreements.



Q15

Please provide the link to your privacy notice (sometimes referred to as a privacy policy). If your company uses different privacy notices for different products or services, please provide an example relevant to the products or services covered in your response to this question.

Our Privacy notice can be found at www.toluna.com



Q16



What steps do you take to comply with data protection laws and implement measures to protect the privacy of research participants? Have you evaluated any risks to the individual as required by privacy legislation and ensured you have obtained consent for data processing where necessary or have another legal basis?

Toluna applies GDPR as its base for compliance with privacy laws around the world and increases the level of compliance to adhere to any higher obligation under local law.

Toluna has a dedicated privacy management program, which was implemented as part of our GDPR/Privacy and ISO27001 certification programs. The teams in the program are co-managed by operating in two layers: at an executive level (Group General Counsel and the CTO) and at an IT/Information Security and Operations level. In each country in which it operates, Toluna has a privacy champion, who is generally a member of the Client Services teams. The legal team includes a dedicated Privacy & Compliance legal Counsel, who has International Association of Privacy Professionals (IAPP) qualifications (CIPP/E and CIPM) and is a member of the ESOMAR Legal Affairs Committee.

Toluna monitors compliance with legal/regulatory standards by leveraging compliance tools like OneTrust and has processes in place for analyzing and implementing changes required by applicable legal/regulatory standards. Additionally, members of the legal team receive information

about updates to legal/regulatory standards from a variety of legal resources.

Toluna is transparent about the processing of individuals' personal data. All panelists join the Toluna panel and/or participate in research programs by providing their consent to such participation. Toluna also uses personal data for other legal bases e.g. legitimate interest, where the respondent understands how their personal data will be used for other purposes by being provided with fair processing notices.

Toluna undertakes data protection impact assessments (DPIA's), where the processing of personal data may pose greater risks and employs a data privacy by design approach in relation to its approach to new processing, so that privacy is embedded within the design of all our technologies.

Toluna's approach to the use of AI ensures transparency of the use of individual's data and it ensures human oversight is integrated in its systems to ensure the use of AI is safe and ethical.

Q17



What steps do you follow to ensure AI systems are resilient to adversarial attacks, noise and other potential disruptions? Which information security frameworks and standards do you use?

Toluna builds resilience and security in multiple ways:

1. Guardrails

This includes restrictions on what the AI can generate or decide, with monitoring to detect and alert on unusual behavior. It also includes mechanisms that allow human intervention if the AI behaves unexpectedly. This helps us maintain control over the AI's actions and ensures it operates within the ethical and operational boundaries we've set.

2. Noise and disruption management

We use techniques like data augmentation and regularization to make our models less sensitive to noise and more capable of handling real-world variability.

3. Fallback plans

In case of a significant problem, we have recovery strategies and backups in place to restore functionality. This ensures minimal disruption to our services and protects data integrity.

AI products are subject to secure software review, architecture, data governance, and third-party assessments. They are checked for prompt injection, OWASP top 10 for LLM models.

We also ensure AI safety and security assurance by training our teams on these topics.

Finally, we also check for existing vulnerabilities in LLM models.

Q18

Data ownership:

Do you clearly define and communicate the ownership of data, including Intellectual property rights and usage permissions?

Q19

Data sovereignty:

Do you restrict what can be done with the data?

Q20

Ownership:

Are you clear about who owns the output?



Toluna owns the intellectual property in anything that arises or is obtained/developed by Toluna and/or third parties and which may be developed independently of the Services. Examples include but are not limited to any methodologies, systems, computer programs and/or software used by Toluna to perform the services, including AI tools and methodologies.

The client owns the results of the services, which are created exclusively for them. However, to stay in line with market research best practices, the client should use Deliverables for their own internal purposes and maintain confidentiality.