


ESOMAR

20 Questions
to Help Buyers of
AI-Based Services
for Market Research
and Insights



© 2024 ESOMAR. Issued March 2024.

These Questions were drafted in English and the English text (available at www.esomar.org) is the definitive version. The text may be copied, distributed and transmitted under the condition that appropriate attribution is made and the following notice is included.

ISBN: 978-90-8336785-9

Disclaimer:

This document should not be considered as legal advice. You should seek appropriate legal counsel for your own situation in the appropriate jurisdiction. In spite of careful preparation and editing, this publication may contain errors and imperfections. ESOMAR and contributors do not accept any responsibility for the consequences that may arise as a result thereof.

The answers that companies provide to the AI Questions are their responsibility and do not reflect the views of ESOMAR or its members. The companies that have answered the questions are always in full control of the statements that they provide and therefore ESOMAR cannot guarantee or be held responsible for their veracity and compliance or the consequences that may arise as a result thereof.

Published by ESOMAR, Amsterdam, The Netherlands.

For further information, please email professional.standards@esomar.org

ESOMAR

20 Questions to Help Buyers of AI-Based Services for Market Research and Insights

About ESOMAR

Since 1947, ESOMAR has been the global hub for research, insights and analytics. Reaching 50,000+ individuals, 750+ companies and 130+ countries, we are a worldwide membership organisation that empowers insights professionals and businesses to unlock their potential on both the global and local stage, fostering connections, collaboration, growth and knowledge.

Driven by our core values of inclusivity, caring, innovation and trust, we have led the industry through a rapidly evolving landscape for more than 75 years. We continue our commitment to raising professional and ethical standards, facilitating education, advocating with legislators, sharing best practices and promoting evidence-based solutions for decision-makers. www.esomar.org

Contents

Purpose and scope	5
A. Company profile	6
B. Is the AI capability/service explainable and fit for purpose?	7
C. Is the AI capability/service trustworthy, ethical and transparent?	8
D. How do you provide Human Oversight of your AI system	9
E. What are the Data Governance protocols?	11
Glossary	14
Project Team	16

Purpose and scope

The objective of this Checklist is to serve as guidance for market, opinion and social researchers and data analysts seeking to commission AI-based services. The aim is to help AI service buyers and suppliers ensure there is an appropriate level of transparency, trust and confidence in the field and that applicable intellectual, privacy and AI laws are respected. The intention is to promote fair, transparent, responsible and ethical usage of AI in market, opinion and social research and data analytics (hereafter referred to as “research”).

In terms of the current scope of AI activity within research, we see three areas of consideration.

1. The use of AI tools in relation to the business itself, developed outside of the field of research, e.g. generalised business processes.
2. The use of AI whilst undertaking research activities, for instance, a commercial research offering introduced by a supplier or a bespoke offer made to the buyer. This could include, for example, the development of specific research products and services, ways in which research organisations interact with research participants, methods used in research operations activities, post-analysis and preparation of research conclusions, and methods used in the development and analysis of proprietary syndicated datasets.
3. The use of the research data and findings in the AI programmes of third parties, including the use of research data as AI “training data”. These third parties would usually be clients and commissioners of research programmes.

The scope of this guidance covers the second area and excludes the first and third.

The research sector is not alone in aiming to apply ethical and professional principles to this fast-developing area. The project team has therefore referred to a range of existing ethical codes, together with information about the likely legislative landscape, whilst creating this guidance.

The definition of AI is itself a question that many legislators are still discussing. In this context, the material is designed to be used with any organisation claiming to offer AI-based research solutions or in any situation where the user believes that AI is being utilised and where this framework may be helpful. In terms of terminology, we have referred to the existing IAPP¹ AI definition set. The IAPP terms that are relevant to this material are published in the appendix, together with the link to the full definition set and IAPP copyright obligations.

The working group on this project recognises the dynamic nature of the topic. At the time of drafting, nearly all areas of research are being discussed as possible candidates for AI-based activity, meaning that buyers and suppliers are seeking to keep abreast of the rapid changes and understand their implications. Many new and existing organisations are offering AI-powered solutions, some of which have the capability to transform the industry.

It is understood that this guidance may be out of date within a short time window, it may be more relevant to some organisations than others and more relevant to certain AI applications by virtue of the fact that the field of AI is accelerating and diversifying at an astonishing rate. Nonetheless, it is necessary to start somewhere. Recognising all these points, this material will be updated on a frequent basis considering the dynamic business and legislative environment.

¹IAPP = International Association of Privacy Professionals

A. Company profile

What this section is about

This section will help to gain an initial understanding of the credentials of the supplier organisation.

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

Context

The answer will help the buyer to form an opinion about the experience and knowledge of the supplier offering AI-based solutions, both in the fields of research and AI; for example, does the supplier have a good mix of research skills and AI-relevant skills, such as data science, deep learning or natural language processing? This discussion will help the buyer to understand the level and depth of expertise on offer.

2. 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?

Context

This question will give an indication of how well the supplier has thought about the way in which AI-based solutions can solve known research issues and improve the research process and contribution to decision making and society. It may be helpful to understand the supplier's general views on the use of AI tools before going into a more detailed discussion about the services on offer.

3. 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?

Context

This answer will help the buyer to gauge the level of activity within the supplier organisation and what it has learned. It will also provide an indication of how experienced the supplier is in managing challenges around emerging technologies.

B. Is the AI capability/service explainable and fit for purpose?

What this section is about

This section will help the buyer to evaluate AI services from a practical standpoint. It will enable the buyer to determine whether the capability on offer aligns with their business purpose and is likely to provide a clear benefit.

4. 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?

Context

Explainability is key to understanding how the methodology works in broad terms, how AI is used to power the method, and the benefits the approach delivers.

5. 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?

Context

Buyers may find it useful to distinguish between a custom-built and an open-source solution to be able to assess capabilities and possible risks. It should be clear to buyers whether publicly available solutions are integrated into the suppliers' offering so they can assess potential consequences, such as the handling of proprietary information.

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

Context

Information regarding what data is used, how the data is used, and the algorithmic processes applied to that data can help determine if the AI is fit for the required purpose.

It could be helpful to understand whether client data and/or synthetic data is used or required in the training data and if buyers can opt out of this. Many AI systems continually learn from the data they ingest, and this can benefit users, but there should be consideration of data provenance and veracity.

It may be helpful to probe if the model is trained using non-English content depending on your use case or service of interest.

C. Is the AI capability/service trustworthy, ethical and transparent?

What this section is about

This section will help to clarify whether the buyer and supplier are aligned on ethical principles, and whether the supplier has considered other important topics such as potential biases, data security and resilience.

7. 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'?

Context

AI can pose challenges related to bias and representativeness and risk generating inconsistent or incorrect outcomes, depending upon the objectives. This question aims to understand the processes that the supplier uses to ensure reliable results. For instance, the answer should indicate how the training datasets are selected, what they cover, and if they are sufficiently up-to-date, and screened to be representative of the environment in which the system will be deployed.

Supplementary questions and context for generative AI and synthetic data are shown below:

- Tools using Generative AI: AI may have limited knowledge of the real world and events after a particular date and can sometimes produce biased, incorrect or otherwise misleading answers. To what extent have such tools been incorporated into the AI model, and what measures are taken to identify "hallucinations", incorrect answers and minimise bias?
- Tools using synthetic data: what validation has been done to compare synthetic data outputs to primary research outputs or real-world results? Are there criteria for data recency to be able to assess if the data is appropriate for use? How does the supplier flag and distinguish, in their systems and approach, between data derived directly from natural persons, and data which is derived synthetically?

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

Context

Understanding service limitations is crucial to determine the potential risk of inconsistent or inaccurate output. This question aims to evaluate if the supplier is transparent about the limitations of their offering. Buyers should be provided with sufficient technical information to enable them to assess the validity of the results and any conclusions drawn. Mitigation strategies might include quality metrics, testing and approval processes, plus regular reviews against agreed criteria to identify and manage emerging issues.

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

Context

Ethics play a vital role in market research. This question will help buyers to evaluate whether the supplier has considered, within the design of their service, whether there might be any potential negative consequences of using their AI model for people (e.g. increased prejudice, financial harm, misinformation).

D. How do you provide Human Oversight of your AI system?

What this section is about

This section will help buyers understand how human involvement and oversight have been considered in both the development and the operation of the AI applications on offer. Buyers should expect the supplier to be able to discuss human oversight in their process and/or how the user of the method is able to stress test the outputs.

Answers to questions in this section will help determine how to identify what role the human plays when building solutions driven with AI and working with data that is processed/analysed with AI in an ethical and responsible way.

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

Context

Identifying AI-generated images/text will be helpful to indicate the original source of outputs to buyers. Suppliers should be able to communicate the type of AI technologies used in their service. For instance, if relevant, the supplier may flag where images/text are 'generated by AI'.

11. 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?

Context

This question is designed to ascertain if the supplier has defined ethical principles and probes the extent to which humans have been in the loop when testing software modules.

For instance, training AI systems via reinforcement learning of the models from using human feedback is the main technique for using or aligning AI models and can help to avoid creating or reinforcing bias. Buyers may need to know that there has been human validation and checking where outputs produced by AI are being used in decision making or when it potentially has an impact on individuals.

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

Context

This question will allow buyers to understand whether the supplier has integrated any tools, techniques or processes involving humans in the design and development of an AI solution for research. The extent of these interventions will depend on the size and nature of the AI supplier organisation and the type of service or AI on offer. They might include, for example:

- Human-in-the-loop² : this approach calls for humans to be involved throughout the process of building and updating machine learning algorithms.

²Human-in-the-loop (HITL) - in the context of utilising AI-based systems, HITL refers to a framework for integrating human judgement into the process. This can include, but is not limited to, ongoing supervision, human overrides, final decisions being made by humans, and exception handling. Source: ESOMAR.

- Ethical review boards: setting up independent ethical review boards comprising experts from diverse backgrounds.
- Participatory design: involving diverse stakeholders, including non-western communities, in the design process of AI solutions to help ensure that the solutions are tailored to the needs of different communities.
- Cultural sensitivity training: providing cultural sensitivity training to developers and data scientists working on AI solutions to ensure that they are aware of cultural differences and can develop solutions that are more inclusive.
- Human-guided data curation: involving creating detailed networks called knowledge graphs. Each point in these graphs holds data and every connection represents a meaningful relationship. People can add carefully selected information to these graphs with curated data, offering invaluable context for queries powered by large language model (LLM)-driven queries. Any additional data created by machines should be checked by human experts.
- Human-engineered ontologies: Using these organised models that let people influence how machines think to help guide LLMs and set limits on their actions. Any changes to these models need to be checked by humans before being used in a solution.

E. What are the Data Governance protocols?

This section can be used in combination with the other questions in this guidance, or as a standalone checklist. To this extent, there may be some elements of overlap with other sections in this material.

What this section is about

This section will help buyers understand whether the supplier is appropriately aware of the legal frameworks that govern AI based activities.

AI suppliers and their clients are subject to data protection and related information security requirements imposed by data protection laws and regulations. These laws and regulations vary by jurisdiction with different laws and regulations applying in different countries or states within countries and are generally interpreted based on where the data were collected or the location of the provider.³

In addition, they may be subject to laws and regulations relating to intellectual property and copyright.

Answers to the questions in this section can help buyers understand the data protection, information security and compliance policies, procedures and practices that a supplier has implemented.

13. 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?

Context

Inaccurate, incomplete, or irrelevant data may affect the outcomes from the AI. Data quality for AI is a key step for ensuring that the AI works as desired. Low quality data is likely to lead to an output, insights and conclusions that are not valid. For example, even within LLMs, certain nationalities or demographics within nationalities are likely to be under-represented. How has this been accounted for when determining insights relating to a wider demographic?

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

Context

It is important to understand where the data was sourced from, how and who processes it and whether proprietary or publicly available data is used. Buyers will want to understand that data is responsibly and ethically sourced and processed and is fit for the purpose of the AI system.

³Applicable data protection laws and regulations include, but are not limited to: the Act on the Protection of Personal Information or APPI (Japan); the Australian Privacy Act (Australia); the California Consumer Protection Act or CCPA (state of California in the US); the Children's Online Privacy Protection Act or COPPA (US); the Personal Information Protection Law of PIPL (China); the Data Protection Act (UK); amendments regarding data localisation requirements to the Data Protection Act (Russian Federation); the General Data Protection Law (Brazil); the EU General Data Protection Regulation or EU-GDPR (EU/ EEA); the Health Insurance Portability and Accountability Act or HIPAA (US); the Graham-Leach Bliley Act or GLBA (US); and PIPEDA (Canada).

15. 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.

Context

A privacy notice is required by data protection laws and regulations, as well as other laws and regulations, and many market research industry codes. It should disclose information about the personal data that a supplier collects and processes, the purpose the personal data is used for and the way that personal data is used, disclosed and managed. A review of the supplier's privacy notice can help buyers understand their procedures and practices related to personal data and the degree to which they comply with applicable laws, regulations and industry codes.

16. 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?

Context

Understanding the supplier's compliance position is essential to ensure adherence to laws and help protect the personal data of individuals.

Global privacy laws require researchers to assess the risks that the AI may present to people and, in many cases, obtain their consent to process their data for a particular purpose. This is typically done through a data protection impact assessment (DPIA), and guidance is available from a country's data protection regulator.

Where personal data relating to people is being processed, privacy laws typically require researchers to be transparent with individuals about how their data will be processed, who it will be shared with, the purpose it will be put to and how it is kept secure when used with AI. Buyers and suppliers are subject to data protection and related information security requirements imposed by data protection laws and regulations.

17. What steps do you follow to ensure AI systems are resilient to adversarial attacks, noise and other potential disruptions? Which information security frameworks and standards⁴ do you use?

Context

The answer should include a process to identify and address vulnerabilities and possible threats to the AI system, design, or technical faults such as manipulation of training data, or possible misuse or inappropriate use of the system. The response should also cover measures to detect and tackle cyber-attacks, cybersecurity systems and whether the model is compliant with specific security standards. Are there fall-back plans that enable a back-up plan in case of problems? It may be useful to understand where the data is located and whether it is stored on cloud or on-device.

⁴ Examples include, but are not limited to COBIT, HITRUST, ISO 27001, the NIST Cybersecurity Framework and SOC 2.

18. Data ownership: Do you clearly define and communicate the ownership of data, including intellectual property rights and usage permissions?

Context

Is it clear who owns the data being input into the model? The data being input may have been information that you have created, or information that you have sourced from a client or another supplier. If a supplier or a client own the data, they may have retained ownership of the data and place controls on how you can use the data (for example to not upload it into Chat GPT and/or other equivalent models). It is helpful to clarify at the outset what you can or cannot do with the data, and to ensure this is appropriately documented.

Data owners may want their information kept confidential or secret and may require that it is not published in the public domain or into AI system cloud-based tools. It is important to be clear about which data are deemed confidential and therefore should not be uploaded into publicly available technologies.

19. Data sovereignty: Do you restrict what can be done with the data?

Context

Data owners may not be comfortable with their data being processed by AI. A number of AI technologies are new, and the data owner may not want their data uploaded into a system of which they are unaware, have not reviewed, or mistrust. Data owners may not permit their data to be co-mingled with other datasets. There are a range of emerging AI laws, and data owners will not want their data loaded into potentially non-compliant AI technologies. It is advisable to include in the contract whether a client allows or forbids any of their data to be processed by the AI in either data learning and/or in the production of outputs using AI tools. Data owners may also need to stipulate which country or region the data or models should be held in.

Data owners should also confirm which party owns the IP in the inputs (prompts) to the AI model. Data owners may also need to stipulate which country or region the data or models should be held in.

20. Ownership: Are you clear about who owns the output?

Context

Data owners may be concerned about which entity owns the outputs, i.e. who owns the intellectual property (IP) of the information produced by the AI technology. For example, a third-party AI technology owner may state in terms and conditions that they own the IP of any output from the AI technology, which may restrict how you can use the information, particularly if the contemplated use is of a commercial nature.

Glossary

The glossary used in the preparation of this material has been provided by the IAPP (International Association of Privacy Professionals). The “Key Terms for AI Governance” produced by the International Association of Privacy Professionals originally appeared in the IAPP Resource Center. It is reprinted with permission. This version was published June 2023.

The following terms are a sub-set of the glossary relevant to this ESOMAR guidance:

Artificial Intelligence⁵

Artificial intelligence is a broad term used to describe an engineered system that uses various computational techniques to perform or automate tasks. This may include techniques, such as machine learning, where machines learn from experience, adjusting to new input data and potentially performing tasks previously done by humans. More specifically, it is a field of computer science dedicated to simulating intelligent behavior in computers. It may include automated decision-making. Acronym: AI

Bias

There are several types of bias within the AI field. Computational bias is a systematic error or deviation from the true value of a prediction that originates from a model’s assumptions or the input data itself. Cognitive bias refers to inaccurate individual judgment or distorted thinking, while societal bias leads to systemic prejudice, favoritism and/or discrimination in favour of or against an individual or group. Bias can impact outcomes and pose a risk to individual rights and liberties.

Explainability

The ability to describe or provide sufficient information about how an AI system generates a specific output or arrives at a decision in a specific context to a predetermined addressee. XAI is important in maintaining transparency and trust in AI.

Acronym: XAI

Generative AI

A field of AI that uses deep learning trained on large datasets to create new content, such as written text, code, images, music, simulations and videos. Unlike discriminative models, Generative AI makes predictions on existing data rather than new data. These models are capable of generating novel outputs based on input data or user prompts.

Hallucinations

Instances where a generative AI model creates content that either contradicts the source or creates factually incorrect output under the appearance of fact.

Input data

Data provided to or directly acquired by a learning algorithm or machine learning model for the purpose of producing an output. It forms the basis upon which the machine learning model will learn, make predictions and/or carry out tasks.

⁵The Organisation for Economic Cooperation and Development (OECD) defines AI as: An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.

The definition of AI in the EU AI Act is identical but excludes the final sentence.

Large Language Model

A form of AI that utilises deep learning algorithms to create models (see machine learning model) pre-trained on massive text datasets for the general purpose of language learning to analyse and learn patterns and relationships among characters, words and phrases. There are generally two types of LLMs: generative models that make text predictions based on the probabilities of word sequences learned from its training data (see generative AI) and discriminative models that make classification predictions based on probabilities of data features and weights learned from its training data (see discriminative model). The term “large” generally refers to the model’s capacity measured by the number of parameters and to the enormous datasets that it is trained on.

Acronym: LLM

Machine Learning

A subfield of AI involving algorithms that enable computer systems to iteratively learn from and then make decisions, inferences or predictions based on input data. These algorithms build a model from training data to perform a specific task on new data without being explicitly programmed to do so. Machine learning implements various algorithms that learn and improve by experience in a problem-solving process that includes data cleansing, feature selection, training, testing and validation. Companies and government agencies deploy machine learning algorithms for tasks such as fraud detection, recommender systems, customer inquiries, health care, or transport and logistics.

Acronym: ML

Reliability

An attribute of an AI system that ensures it behaves as expected and performs its intended function consistently and accurately, even with new data that it has not been trained on

Synthetic data

Data generated by a system or model that can mimic and resemble the structure and statistical properties of real data. It is often used for testing or training machine learning models, particularly in cases where real-world data is limited, unavailable or too sensitive to use.

Training data

A subset of the dataset that is used to train a machine learning model until it can accurately predict outcomes, find patterns or identify structures within the training data.

Transparency

The extent to which information regarding an AI system is made available to stakeholders, including disclosing whether AI is used and explaining how the model works. It implies openness, comprehensibility and accountability in the way AI algorithms function and make decisions.

Trustworthy AI

In most cases used interchangeably with the terms responsible AI and ethical AI, which all refer to principle-based AI governance and development, including the principles of security, safety, transparency, explainability, accountability, privacy, non-discrimination/nonbias (see bias), among others.

Validation data

A subset of the dataset used to assess the performance of the machine learning model during the training phase. Validation data is used to fine-tune the parameters of a model and prevent overfitting before the final evaluation using the test dataset.

Permission should be sought directly from the IAPP to publish or restate any part of this glossary.

Project Team

Barry Jennings, Director, Cloud + Commercial Business Planning Insights at Microsoft

Kathy Joe, Consultant to ESOMAR Professional Standards Committee

Judith Passingham, Chair of ESOMAR Professional Standards Committee

Ravinder Roopra, Head of Privacy at Kantar

Phil Sutcliffe, Managing Partner, Nexxt Intelligence

Facilitated by:

Paula Fernandez, Standards Programmes Coordinator at ESOMAR

The bottom half of the page features a dark grey background with a repeating pattern of hexagons. In the lower right corner, the word "ESOMAR" is written in a large, bold, white sans-serif font. Below it, the website address "www.esomar.org" is written in a smaller, white sans-serif font.

ESOMAR
www.esomar.org