

Research Question Examples: AI & Machine Learning

About this document

This PDF contains 30 **AI & machine learning research question examples**, each quoted word for word from one of 15 completed dissertations and theses (2022–2026), grouped into 5 sections. Each question links to the thesis that asked it, so you can read how it was answered. Use them to see what a well-formed AI & machine learning research question looks like, then adapt the shape to your own study.

By [Derek Jansen](#) (MBA) · Reviewed by [Eunice Rautenbach](#) (DTech) · Grad Coach · October 2026

Online version: gradcoach.com/research-question-examples/ai-machine-learning/

© 2026 Grad Coach. Free to share and adapt under [CC BY 4.0](#) with credit to Grad Coach and a link to gradcoach.com/research-question-examples/ai-machine-learning/. Research questions and thesis titles quoted here remain their authors'; each dissertation keeps its own license.

Contents

Descriptive questions · 10 questions	2
Exploratory questions · 8 questions	3
Comparative questions · 7 questions	4
Causal questions · 3 questions	5
Relational questions · 2 questions	5
How to adapt a question to your own study	6

Descriptive questions

These questions ask what is happening: how much, how often, or what kinds. They describe a situation rather than explain it.

1. How well are these AI-enhanced project management tools integrated with their AI feature?

Asked in: [Analysis and comparison of generative AI chatbot applications](#) · Ahuja, Maghav · Unitec Institute of Technology · 2024

2. How do patients and family supporters perceive the trustworthiness and social appropriateness of an AI companion that enforces structural role boundaries and manages the flow of information between them?

Asked in: [Designing, Prototyping and Evaluating a Family-Aware AI Coaching Agent for Type 2 Diabetes Patients](#) · Leon Samir Saidavi · Uppsala University · 2026

3. How effective is a chatbot using retrieval-augmented generation at retrieving and providing contextually relevant information from domain-specific manuals?

Asked in: [DOMAIN-SPECIFIC INFORMATION RETRIEVAL FROM A LARGE LANGUAGE MODEL CHATBOT](#) · Edis Sabic · Mälardalen University · 2025

4. How do practitioners perceive the chatbot's usefulness, accuracy, and effectiveness in assisting their tasks?

Asked in: [DOMAIN-SPECIFIC INFORMATION RETRIEVAL FROM A LARGE LANGUAGE MODEL CHATBOT](#) · Edis Sabic · Mälardalen University · 2025

5. Does the application of a lightweight object detection model on raw video data create a feasible solution to fish monitoring?

Asked in: [Enhancing Fish Detection Using Synthetic Datasets](#) · Igor Górski · Mälardalen University · 2024

6. Are synthetic datasets a viable method for enhancing the model's scope and thus improving the possible results given a limited amount of available data?

Asked in: [Enhancing Fish Detection Using Synthetic Datasets](#) · Igor Górski · Mälardalen University · 2024

7. What are the critical research areas, user level and design level insights discussed in the previous literature?

Asked in: [Explainable artificial intelligence \(XAI\) : making AI understandable for end users](#) · AKM Baharul Haque · LUT University · 2025

8. What are the limitations and benefits when widely-used deep reinforcement learning (DRL) approaches are used to address constrained and combinatorial optimization problems in wireless networks, and are there tailored solutions to overcome the inherent drawbacks?

Asked in: [Machine Learning-Based Efficient Resource Scheduling for Future Wireless Communication Networks](#) · Yaxiong Yuan · University of Luxembourg · 2022

9. How to predict the performance of [Deep Neural Networks] under continuous distribution shifts?

Asked in: [Towards safe and efficient application of deep neural networks in resource-constrained real-time embedded systems](#) · Luan, Siyu · Umeå University · 2023

10. To what extent are open-source LLMs suitable for open-vocabulary tasks?

Asked in: [Avoimen lähdekoodin suurten kielimallien hienosäätö avointen kommentojen tulkitsemiseen robottinavigoinnissa](#) · Palmulaakso, Joonas · Aalto University · 2025

Exploratory questions

These questions ask how or why something happens, usually where little research exists yet. They are common in qualitative studies built on interviews or observation.

11. In what ways does a role-separated, on-demand conversational AI companion buffer or alleviate the emotional burden and interpersonal conflict (e.g., "diabetes nagging") that can arise in family-based self-management?

Asked in: [Designing, Prototyping and Evaluating a Family-Aware AI Coaching Agent for Type 2 Diabetes Patients](#) · Leon Samir Saidavi · Uppsala University · 2026

12. How do users want the explanations to be presented/ delivered, and how will access to explanations impact their usage?

Asked in: [Explainable artificial intelligence \(XAI\) : making AI understandable for end users](#) · AKM Baharul Haque · LUT University · 2025

13. How can a structured fairness evaluation framework be adapted to the Aotearoa New Zealand context?

Asked in: [Improving fairness in AI systems: A framework for bias mitigation](#) · Loganathan, Manochitra · Unitec Institute of Technology · 2025

14. How to overcome the shortcomings of extensively adopted end-to-end learning in addressing resource management problems, and which types of features are suited to be learned if supervised learning is applied?

Asked in: [Machine Learning-Based Efficient Resource Scheduling for Future Wireless Communication Networks](#) · Yaxiong Yuan · University of Luxembourg · 2022

15. How to enable ML-based approaches to timely adapt to dynamic and complex wireless environments?

Asked in: [Machine Learning-Based Efficient Resource Scheduling for Future Wireless Communication Networks](#) · Yaxiong Yuan · University of Luxembourg · 2022

16. How to design a data pipeline which enables self configuration (i.e. how to design effective Machine Learning (ML) pipelines)?

Asked in: [SUPPORTING COMPANIES IN THEIR DIGITAL TRANSITION TO SMART MANUFACTURING SYSTEMS](#) · Giuseppe De Martino · University of Milan · 2025

17. How can the 'Items' define in the WPS be systematically mapped to quantifiable website elements?

Asked in: [Utilizing artificial intelligence for website personality detection](#) · Shafquat Ali Chishti · Unitec Institute of Technology · 2026

18. How can ML modules be designed to classify website personality across multiple traits and dimensions?

Asked in: [Utilizing artificial intelligence for website personality detection](#) · Chishti, Shafquat Ali ·
Unitec Institute of Technology · 2026

Comparative questions

These questions compare two or more groups or settings, such as two schools or two teaching methods, to see whether they differ.

19. What is the level of reliability and agreement between deterministic approaches based on keyword matching and semantic approaches based on advanced language models in determining the maturity of a startup?

Asked in: [Automating startup evaluation: from qualitative descriptions to a multidimensional, text-based measure of maturity and investment readiness](#) · Mariannelli, Veronica · University of Genoa · 2026

20. How efficient are fairness metrics in effectively detecting and measuring bias across single and intersectional demographic groups?

Asked in: [Improving fairness in AI systems: A framework for bias mitigation](#) · Loganathan, Manochitra ·
Unitec Institute of Technology · 2025

21. How different are patterns of bias in Aotearoa New Zealand from global benchmark datasets?

Asked in: [Improving fairness in AI systems: A framework for bias mitigation](#) · Loganathan, Manochitra ·
Unitec Institute of Technology · 2025

22. How do different LLM models perform in terms of precision and analytical level in producing competitor analysis?

Asked in: [Market Intelligence in a New Era : a case study on AI in market intelligence](#) · Battal, Nour ·
Mälardalen University · 2026

23. How do adversarial patch defence algorithms perform on different hardware platforms with varying computing capabilities?

Asked in: [Performance Measurement and Analysis of Defences against Adversarial Patch Attacks](#) · Gao, Zeyu ·
Umeå University · 2024

24. Which of the selected contemporary approaches achieves the best performance for survival prediction from multichannel mIF microscopy images?

Asked in: [Self-Supervised Representation Learning for Survival Prediction from Multichannel Microscopy Images of Cancer Tissue](#) · Dilmi Nawanjana Wijesinghe Muhandiramlage · Uppsala University · 2026

25. How can the developed modules be validated against human perception of website personality?

Asked in: [Utilizing artificial intelligence for website personality detection](#) · Chishti, Shafquat Ali ·
Unitec Institute of Technology · 2026

Causal questions

These questions ask whether one thing causes a change in another. Answering them usually takes an experiment or a careful design that rules out other explanations.

26. How do bias mitigation techniques affect the balance between fairness and accuracy when applied at different stages of the AI development lifecycle?

Asked in: [Improving fairness in AI systems: A framework for bias mitigation](#) · Loganathan, Manochitra · Unitec Institute of Technology · 2025

27. To what extent do data-level augmentation strategies, such as [Synthetic Minority Oversampling Technique] variants and [Generative Adversarial Networks], affect group fairness outcomes and predictive performance?

Asked in: [Improving fairness in AI systems: A framework for bias mitigation](#) · Loganathan, Manochitra · Eastern Institute of Technology · 2025

28. How can iterative methods improve the quality of AI-generated market intelligence?

Asked in: [Market Intelligence in a New Era : a case study on AI in market intelligence](#) · Battal, Nour · Mälardalen University · 2026

Relational questions

These questions ask whether two things tend to go together, such as study time and exam results, without claiming that one causes the other.

29. How do heuristics-based adversarial defence algorithms perform with increasing patch sizes?

Asked in: [Performance Measurement and Analysis of Defences against Adversarial Patch Attacks](#) · Gao, Zeyu · Umeå University · 2024

30. Can image representations learned in a self-supervised manner from multichannel Multiplexed immunofluorescence (mIF) microscopy images of cancer tissue provide prognostic information for lung cancer survival prediction?

Asked in: [Self-Supervised Representation Learning for Survival Prediction from Multichannel Microscopy Images of Cancer Tissue](#) · Dilmi Nawanjana Wijesinghe Muhandiramlage · Uppsala University · 2026

How to adapt a question to your own study

Every question here was asked by someone else, about their own setting. These four steps turn one into a question that is yours.

- 1. Pick the shape, not the wording.** Notice whether the question compares, explores, explains or measures. That shape is what transfers to your study.
Recommended reading: [How to write a research question](#)
- 2. Swap in your own population and setting.** Change who, where and when, and check the question still makes sense.
Recommended reading: [How to find a research gap](#)
- 3. Check it fits your method.** A "how" or "why" question usually points to qualitative work; "how much" or "is there a relationship" to quantitative.
Recommended reading: [Qualitative vs quantitative research](#)
- 4. Line up your aims, objectives and questions.** Your committee will check that they agree.
Recommended reading: [Research aims, objectives and questions](#)

Do this with two or three questions and compare them: the one you can actually answer with the time and access you have is usually the right one.