

University of Essex Online

MSc Artificial Intelligence

Transcript

Research Proposal Presentation

Submission date: 12.01.2026

Contents

Transcript	2
Slide 1: Title Slide	2
Slide 2: Agenda.....	2
Slide 3: Introduction	2
Slide 4: Research Problem & Research Question	3
Slide 5: Key Literature.....	5
Slide 6: Research Design.....	6
Slide 7: Ethical considerations and risk assessment.....	7
Slide 8: Description of artefacts	8
Slide 9: Timeline	9
Slide 10: Thank you	9

Transcript

Slide 1: Title Slide

Hello and welcome to my presentation. My name is Elias Medig. Today I'm going to present you my research proposal outline for the Master of Science in Artificial Intelligence program at the University of Essex online.

Slide 2: Agenda

We are going to start with an agenda. I will start with a short introduction followed by looking into the research problem and research questions. Then we'll have a brief discussion about the key literature, research design and the ethical considerations and risk assessments that I did. Then to end the presentation, I will give a description of the artifacts that I expect to get because of my thesis. And then finally, we're going to look at the whole timeline that I propose. Let's jump right directly into this by looking into the introduction.

Slide 3: Introduction

First, the project or project title of my thesis would be design and evaluate an AI enabled agent for enterprise architecture compliance. Now, how does that look like or what is it about? The significance or let's say the contribution to the discipline and research problem will mainly be about the three, let's say key problems, where in EA, compliance relies heavily on manual review of metadata and documentation. That means a lot of things that could be done automatically or maybe even with AI agents. That's what I suppose with this thesis are currently done by humans, by manual

interacting with data, by manual work. Then on top of that, portfolio data is often incomplete, inconsistent and/or outdated, which means that a lot of the information that is tracked in general IT portfolio management often is not up to date, might have some errors etc. And that's like one of the big problems in enterprise architecture. As a result of all of that, so that's the last problem, we have slow decision taking, we have hidden redundancies, and we have a general governance risk. Now, the significance of the work that I want to do by designing and evaluating AI enabled agents for the enterprise architecture discipline would especially come from applying machine learning and large language models to EA governance. It's under my research area. Then it would be a proposal of an agent-based decision support approach, which means that I would like to propose an agent-based solution that helps enterprise architects or say the enterprise architecture discipline to come up with automated ways of coping with the problems that I mentioned earlier. And then in the end, the contribution would also be that what I would be doing here would provide measurable evaluation criteria, so accuracy and time to decision taking would especially be the key factors that I would look at. So how accurate, what is the, let's say data quality that we have, and what is the time to decide that is usually involved in enterprise architecture decisions after they would employ AI agents. And of course, I would use reproducible synthetic data sets so that the whole work that I would be doing, the whole agents that I would be building are reproducible.

Slide 4: Research Problem & Research Question

Then the next topic is the research problem and the research question. So maybe mainly I would focus myself on three major research questions. The first one would be which enterprise architecture governance tasks are suitable for AI support? It's

not everything that enterprise architects do can be done or substituted with artificial intelligence. The second one is how accurately can AI agents detect portfolio anomalies, redundancies, and non-compliance? Here that's really to cope with the problem that I mentioned earlier regarding data quality when it comes to IT asset portfolios used in the enterprise architecture domain to do their work. Maybe one good example would be the application portfolio management that is done in nearly all large companies to a certain degree, but also all large companies from the know-how that I have regarding my work lack certain capabilities there, especially when it comes to data accuracy, data quality, completeness of the data sets that they have, and their AI agents could detect such anomalies and then either correct them or try to get them filled out. Then the last research question would be what measurable impact does AI support have on decision quality and time to decision? Here again it goes a little bit into both other research questions because if AI agents can support in decision taking, in data quality, increasing in let's say certain operative tasks that the enterprise architects would have to do, then of course it can also help to reduce the time to decision, but also increase the quality of decisions because of a better foundation, better data that the enterprise architects have to take their decisions on. Now aims and objectives, the aim is really to design and evaluate an AI-enabled agent for enterprise architecture compliance support. That means I want to design and build and evaluate different agent models to see how they could theoretically respond to my three research questions and how they would perform there. The objectives would really be to analyze AI governance and AI literature to get a better understanding of the whole enterprise architecture and artificial intelligence environment, then also of course going through certain interviews. I would have to build up different interview questions. I would focus on qualitative interviews with

architects, with CTOs or other let's say architecture-related subject matter experts. Then one of the most important parts, I would build a POCO proof of concept machine learning and large language model agent to really deliver the value that I want to test against and then also of course evaluate it and see hey did those agents really provide the value expected, did they help, did they make the enterprise architecture domain let's say easier or more efficient.

Slide 5: Key Literature

Then next we are talking about the key literature related to the project. Most of all it's about three major areas. We have the enterprise architecture and governance area, the AI techniques for decision support and then the intelligent agents. For all of those I would look into different literature sources and for me it's not mainly about just designing the topics but rather to look into the most important parts which means to critically engage with literature, identify gaps and limitations but also try to choose up-to-date literature not coming up with any literature from the 1990s that is not up-to-date anymore. I really want to be ahead of the herd here and investigate what's new, what's upcoming and what could really help the enterprise architecture domain. Then of course important is the sourcing of the literature. We want to use credible sources. We want to do it in a way that it really says anything about the whole market, the whole areas that I'm looking in and not just some paid media sources. Here I think especially in the AI area it's very important to keep that in mind because there might be a lot of marketing sources in terms of what's the best model, why it's good, how it forms etc. Just to sell their AI models. That's something that I will be looking into. Then last but not least of course I want to include various science and relevant industry sources so it means I don't only want to look into science or into

industry, I want to really combine both, take the best out of both worlds, looking at what do let's say large companies say about it, what do think tanks say about it, what do researchers say about it and then come into a position where I can really take an opinion and add it to my work.

Slide 6: Research Design

Then we have the research design which is a very important part. For me the whole thesis would be a three-phase model let's say. I would start with building the foundations, that's really the whole structured literature review for the different areas that I'm looking into. Then we have the identification of an AI, a EI governance task. What can even be done with artificial intelligence in the EA area. That's phase one, building the foundations. Then we have phase two, the qualitative study. Here I would use this semi-structured interviews with EA subject matter experts where I can jump into, talk to them, understand their needs, understand their concerns, understand their challenges and getting their perspective into my work. Then of course, thematic analysis to understand the system requirements that I will have against my AI agent systems. Then the last phase that builds up on phase one and two will be the real development and the evaluation of the AI agent. I would there try to focus on an iterative prototyping way of building the AI enabled agent by having several iterations, by testing it, by coming to conclusions, optimizing it and then grow from there. I would of course, for the evaluation to evaluate the machine learning components, which just standard metrics, like precision, recall, the F1 score, etc. Really understand how while my machine learning components work, how they predict. Then in the end, of course, I would to do an end-to-end evaluation to then

understand whether an AI enabled agent could really be effective in the enterprise architecture domain.

Slide 7: Ethical considerations and risk assessment

Then of course, an important part about all these is the ethical considerations and risk assessment. Here I would mainly look into four, let's say, key ethical considerations that would have to take into consideration. I would need the informed consent for my interview participants. I don't want to just record them without their knowing. I don't want to cite them in the wrong way, etc. Then of course, it could be that I need to take care of anonymization of participants or organizations. Because what I want to do is to talk to architects or SMEs or key CTOs or even a CIO from different companies and maybe they don't want to be affiliated with my work or can generally not be cited due to their internal policies, whatever. I really need to take care of that. Then of course, it goes into the next point. No use of real enterprise or sensitive data. It's hand in hand. Not only do I want to anonymize the data of the people, the companies that I would be working with, but of course also not use actual enterprise data when it comes to building the agents. Then AI is used for decision support, not for autonomous decisions. That's very important. Decisions need to be taken by humans. It is about supporting this decision-taking process with an AI, not by replacing humans with AI when it comes to decision-taking. Then of course, we have the risk assessment. There is like limited external validity. Large language models come with hallucinations. We tend to over-trust AI outputs, especially if things go well at the beginning. We might just put more of our own decision-taking process onto the AI. Then of course, synthetic data could also lead to behaviors that we don't want to have in our AI agents. Those are, let's say, the ethical

considerations and for risk assessment points that I want to take care of and look into.

Slide 8: Description of artefacts

Then of course, we have the description of the artifacts that I would be building. Here we have the primary artifacts, which will be a proof of concept AI-enabled agent for the EA compliance support. Let's call it the product that I will be building. Then the machine learning-based portfolio analysis. Here that's really about the assets that we have as companies and where we lack quality and let's say correctness of data, where we may find anomalies, redundancies, etc., which can be pointed out by a machine learning-based agent. Then of course, large language model-based policy and documentation checking. Here that's really about using large language models to dig deeper, to understand, let's say, different different documentations, different different portfolios that we have and to analyze that not only with machine learning and taking predictions, but also with large language models. Then we have several supporting artifacts, like the synthetic EA data sets that I will have to build. As mentioned before, I won't take any real company data. Then we may have a formalized policy and exception rule sets, which I need to come up with for the agents. Then also of course the evaluation script and metrics that I need to have to evaluate the final agents. Here it's important scope. The scope for all of the artifacts for all of the work is decision support only. Then again, no real enterprise data and of course, no production deployment. It's about a POC running with synthetic data to cope EA problems identified via interviews and literature review.

Slide 9: Timeline

Then last but not least, I took a look into a possible timeline. I started it first of March, just so that we have a good date. I too laid the gun chart on and yeah, I created the gun chart. Looks like that. It might be a bit small depending on the monitor you have to look at, but basically I try to really break it down into four different phases. We have research and design, development, evaluation, and then the last phase of writing and submission. Each of those comes with its own milestone, which we have here. First, the research foundation should be completed end of March. The POC should be ready. End of April, evaluation completed end of May. Then one month later, end of June, the thesis could be submitted. Now of course, this timeline might be changed in the future, but that's more or less how I expected.

Slide 10: Thank you

Now that was it already. Thank you for listening. I'm looking forward to the feedback I receive, and I wish you a nice rest of the day.