



Agile Approaches to Policy Development

HÅKAN BURDEN, RISE Research Institutes of Sweden

SUSANNE STENBERG, RISE Research Institutes of Sweden

To both understand the implications of EU's proposed AI Act and involve relevant stakeholders in the process we have adopted an agile approach to policy development. We, a former judge and a PhD in computer science, reason on our experiences from the perspective of the Agile Manifesto. A main difference from traditional agile processes is that we have deliberately produced presentations and not software.

1. THE AI ACT

The growing usage of Artificial Intelligence (AI) poses both opportunities and risks. To ensure future trust in AI as technology, the EU commission has therefore proposed a regulation for AI as a technology to guarantee EU citizens' their safety, health and fundamental rights. The initiative is part of the EU strategy for the Digital Decade (2021) which aims for digital sovereignty and taking a leading role internationally towards trustworthy digitalization. The proposed regulation is referred to as the AI Act (2021) and will be a binding legislative act across the harmonized market, e.g. across all the EU's member states.

The final wording of the AI Act is currently negotiated by the trilogue - the European Parliament, the European Council and the European Commission. That means that the wording of the AI Act is changed iteratively as negotiations adds, deletes or edits the content. Slovenia held the presidency during the autumn of 2021 and the changes made to the proposed AI Act during this period are collectively known as the Slovenian compromise proposal. During the spring of 2022 France holds the presidency which in turn will result in a French compromise proposal. The ambition at the current point of time is that the negotiations will be finished by spring 2023 during the Swedish presidency. The AI Act will come into force two years after the final version has been agreed upon. Until the trilogue has agreed upon a final version the proposal is a living document.

The current version of the AI Act defines AI as a system that receives input from a machine and/or a human, infers how to achieve a given set of human-defined objectives using learning, reasoning, or modelling, and generates output in the form of content, predictions, recommendations, or decisions, which influence the environments it interacts with. Inference can be done by three main technologies or approaches - machine learning approaches (including supervised, unsupervised and reinforcement learning as well as deep learning); logic- and knowledge-based approaches (including knowledge representation, inductive programming, knowledge bases, inference and deductive engines, reasoning and expert systems); statistical approaches (including Bayesian estimation, search and optimization methods).

The AI Act also proposes to differentiate the risk an AI system poses *ex ante*, e.g. the assessment of the risk connected to AI systems is done before the actual AI system is taken into use. The risk levels range from:

Low: Typically, chatbots and similar. It should be obvious for a user that they are interacting with an AI system and over time there should be codes of conduct that guide developers and users for trustworthy AI systems.

High: AI systems that are either safety components or constitute a part of a safety system in certain products (Annex II.A) as well as stand-alone AI systems that can have a negative effect on citizens' health and fundamental rights (Annex III). High-risk AI systems must comply with the proposed certification process, affix the CE-mark to their system and be registered in the official database governed by the Commission.

Prohibited: Using subliminal techniques with the aim to distort someone's behaviour, taking advantage of vulnerable groups, socially ranking citizens based on their behaviour and biometric identification of individuals without their consent.

In order to assess the impact of the proposed regulation we have adopted an agile approach. Before we give an overview of our experiences of the approach in section three, we will present ourselves in the next section. Section four will then recount a significant anecdote that fleshes out some of our findings in further detail. The last section will discuss our main findings in terms of product and process and highlight future work.

2. PERSONAL BACKGROUND

We come from different disciplines and form a (very) small cross-functional and inter-disciplinary team. I, Susanne Stenberg, LL.M (master of law), am a practitioner of the Swedish judiciary and within the legislators. I have worked as a judge in District Courts and the Courts of Appeal for over a decade and been assigned to of the Governments office as Secretary of Inquiries on several legislative proposals, ranging from the rules of democracy (Wästberg et al., 2014) to the legal basis of digitalization within the public sector (Jungstedt and Stenberg, 2018). When I entered the area of research in 2020, I had not professionally come in contact with the words “iterative” and “incremental”. At RISE I joined a team of researchers on mobility and electrified vehicles, adding legal perspectives and policy development into R&D projects. In short, my expertise includes to investigate the needs for and suggesting new laws as well as assessing cases based on existing law.

I, Håkan Burden, have a background in computer science with a PhD on Model-Driven Engineering (Burden, 2014). After I defended my thesis, I left the university to be a researcher at a research institute. I kept my affiliation with my alma mater and have been teaching agile project management for software engineers (Steghöfer, et al., 2016, Steghöfer et al., 2017) and sustainable development for computer engineers (Burden and Sprei, 2020). While my teaching has remained close to my PhD topic, my research has shifted towards policy and how to bridge the perceived gap between regulatory, technological and business development (Burden et al., 2021). The latter is a team effort where I can use my skills to understand and manage complex systems. But until now my colleagues have been responsible for reading and analysing the regulatory space.

Together we have previously worked on policy development for autonomous vehicles (Burden et al., 2021) and smart ships (Burden et al., 2022).

3. WAY OF WORKING

This section introduces our way of working and reasoning in terms of an agile approach. We have chosen to structure the section based on the core values of the Agile Manifesto (2001).

3.1 Individuals and Interactions

During our work with the analysis we have opted to not have set routines in terms of when we meet whom our do what. Rather, we have interacted with external stakeholders as they reach out to us. Those meetings have often followed a plan-act-reflect-cycle as we have first discussed the purpose of the meeting and who will participate in order to understand what kind of message is fit for purpose. We have then put our analysis in place, presented it and in the subsequent discussion both had an opportunity to validate our analysis and get insights in needs and wishes we so far had not encountered. After the meeting we have had our own short retrospective in order to see what we want to change and what we want to continue with.

Interacting with external stakeholders has been facilitated by our dual experience of software and justice systems. We can adapt our way of communication to the intended audience and anchor our insights in examples relevant for their ways of working. A key perspective here has been to explore how the mandate of the stakeholders will change if the AI Act comes into effect.

3.2 Working Software

The nature of our work has meant that we never had any software. Instead, we have ensured to come up with a first analysis that represents value for someone outside of the team but requires as little effort as possible to define. After a first reading of the AI Act as it was presented in spring 2021, we summarized our first analysis in an e-mail to a representative for the automotive industry to see if what we found interesting resonated with their analysis. This was done primarily to see if we understood our potential customers’ take on the act and after positive response, we refined the text into a presentation that could be shared with other actors. The initial text and the subsequent presentation followed the idea of slicing the assignment thinly into MVPs, following the allegory of elephant carpaccio (Cockburn, 2009). Since then our product has been a deck of slides representing different views of the AI act. The slides have been updated or deleted as our analysis continuously evolves.

When preparing for presenting our analysis for new stakeholders we have taken a subset of the deck, added new slides as required and configured our examples with the ambition of making the analysis concrete to the current target audience. A guiding principle has been the notion of fit for purpose. As our internal product owner has not had the network or stakeholder analysis needed to assess what that means for different contexts it has been up to us to make our own judgements. This may be somewhat of an oddity in a software development context, but coming from the judicial arena it is plain obvious and ordinary business as usual.

This has given us a range of freedom that most likely is unusual within agile organisations. But then we have been fortunate to have no dependency on other teams in terms of deliveries.

User responsibility

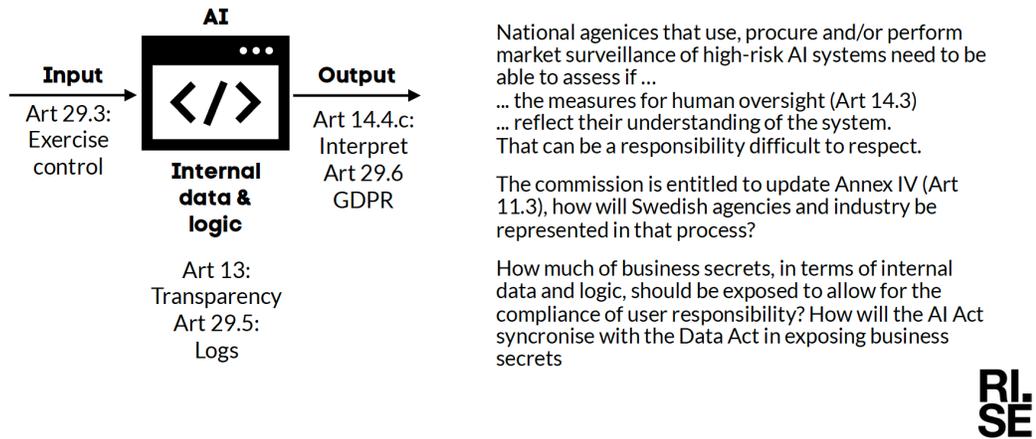


Figure 1. An example of a slide representing our analysis. Traces to the AI Act are shown as article numbers.

In terms of documentation, we have kept it as simple as possible by focusing on traceability. For each partial analysis we have made sure to list which articles that the analysis refers to. In this way it is possible to go back and cross-read the current (or previous) wording of the articles and see how they reflect our analysis. If there is need to update the analysis due to new understandings or changes to the articles this is done in the corresponding slides. If the article numbering is the same as before the trace remains untouched, if new articles are included in the article we add the corresponding numbers to the slide. An example of a slide is given in Figure 1.

3.3 Customer Collaboration

When we started out we did not have any external customers. Instead, we relied on internal sponsors that wanted to explore if our organization could find new opportunities in relation to EU’s ambitions to regulate digitalisation. From there we have had an open mandate to find external stakeholders to interact with and build long-term relationships with. We then got involved in emerging networks designated to discussing possible interpretations and consequences of the proposed regulation. This gave us an opportunity to collect new perspectives of what would be valuable for a larger set of stakeholders. The networks also gave us an opportunity to demonstrate incremental additions and changes to our initial analysis to the stakeholders. Some of the contacts were also willing to share their experiences of working with the act which gave us a community for conducting retrospectives.

We have now reached a tipping point where we have our first paying customers. The customer represents a private company that wants a general description of the AI Act and ensure that we are available in the future for analysing their business in relation to the regulation. A similar set up is being negotiated with an authority with the ambition of finding funding for a mutual research project. It seems in this case that focusing on building relationships with stakeholders has led to funding opportunities, an approach made possible due to the internal funding.

3.4 Responding to change

Late in the spring of 2021 we were given the task to write the remiss for RISE. The remiss is thus our first public assignment and official product launch. While our own organization gave few opportunities for reviews and reflections the remiss has opened new doors for interaction. Susanne was invited to a national center for policy studies which gave new perspectives on stakeholder value and opportunities for reviewing our own analysis. We were now ready to contact employees of the EU Commission to present our key insights and hear

their view. The interaction with various stakeholders has in turn given us new contacts so that we now find us to be in a situation where we respond to requests from national agencies and business networks wanting to explore what the act could mean for them. We are now striving for continuous delivery of incremental changes to our analysis, tailored for specific stakeholder needs.

From our perspective the shift resonates to transitioning into a Kanban-inspired flow as the product matures and the focus shifts from developing a marketable product to maintaining and tuning the product for various stakeholder segments. As the act is still being negotiated our work will continue in close collaboration with invested stakeholders. In what direction is difficult to foresee, managing change and uncertainty is part of the process.

4. CLASH OF CULTURES

The interaction within a team of two is in many ways simple as there is only one relationship to consider, a reflection on our part is that maybe the setup resembles a programming pair more than a team (see e.g. Andersson and Bendix, 2006, for how XP has been used in other contexts than software development). That said, we have had some differences in how to approach the regulation and share our insights based on our different backgrounds.

4.1 Håkan's Story

We were working at another project when Susanne said that there was a new EU proposal for regulating AI. I can remember how I immediately thought that it would be interesting to hear what she finds out from reading the proposal. That I would read the proposal myself to have an opinion was beyond my imagination, mostly because I had no picture of what an EU proposal looks like and what purpose it would serve if I had an opinion. But then Susanne said the reading was both fascinating and difficult since the regulation touched upon technological definitions she was unsure about. That got me thinking that I know enough about AI and software to have a look at the proposed definitions.

Now, an EU proposal is a bit like reading a file of source code. First you have the high-level comments describing the purpose of the code and various calls to other packages detailing the context of the code under consideration. In this case at least 20 pages of legal references to existing regulations and the political ambitions of the EU Commission. I enjoyed it! So much that I postponed checking out the technical definitions. And I could make sense of what I read and formulate opinions. That was a revelation! It was such a buzz to be able to discuss an EU regulation under development with a former judge and feel that I could contribute to the discussion. All I needed now was someone in our business networks to discuss with to see if they had made the same analysis as we had.

As I read more of the AI Act, I eventually got to the articles that are the binding statements. Susanne had of course already read them, and we could start our analysis for the remiss. As the work went on and we found new things that seemed counter-intuitive it really made a difference to have Susanne to reflect with. She could so easily assess if it were normal procedure or an oddity. Having that expertise next to me made it so much easier to try out different approaches since her feedback was immediate. It felt like I got a crash course in the law-making process while we were formulating our own analysis of someone else's ambitions to regulate innovative technology.

From my perspective it really made sense to use my knowledge and skills from agile development to take on a proposed EU regulation. There was no way I would be able to manage the task in a more waterfall way, working iteratively and taking on new increments as I got a grasp of old ones made it possible to define actionable tasks. I also found that engaging other stakeholders to understand their unknowns and insights in relation to the AI Act was beneficial for my own understanding, it gave a possibility to assess what I took for known and gave me new perspectives to integrate into the emerging analysis. Taking on the proposal from a use case perspective gave new ways of analysing the proposal, compared to taking on the proposal as it is structured by the Commission. An example is how the perspective of procuring came up in talks with local and national authorities and that perspective is not in the given structure, but it can be dug out from reading different sections of the proposal and combining that into an analysis by itself. Having written the remiss it felt obvious to engage a broader public by writing a discussion piece for a national newspaper.

4.2 Susan's story

There's no experience or pre-understanding of the agile method on my part before entering into the research arena and collaboration with Håkan and especially not in assessing propositions of legal instruments, which I

have done earlier as the lead on several State Inquiries leading up to new legislation. My previous juridical work has an enormous emphasis on the written word, the training into the profession includes weighing every nuance of optional words, for instance in a judgement. To me, understanding and thereafter clarity is essential, and preferably unambiguous phrasing. Though there are similarities to agile thinking within mediation, where the next step is vital to visualize to the parties in order to make that next step possible to take, coming from the judicial area and several years of training into the profession, the importance of the written word can hardly be exaggerated.

The proposal for a new regulation of AI was made public in April. We had internal funding to work on an analysis of the proposal at the end of May and did a two-week-sprint to understand parts of it and work out an opinion on the significance of the proposal. After the response to the proposal was anchored in our organization's current management team, we formally answered the Swedish governments request of an opinion and sent in our response, as a basis for the governments work with the Swedish position on the proposal. From my horizon we had delivered at a good level. We then discussed whether we should write a debate article based on the texts we already had, we talked to each other and Håkan produced a draft and wanted to submit it. I was very hesitant. I had not yet thoroughly read through all the articles of the AI Act and was far from ready to go public with results to the daily press. For me, it is important that the analysis is correct, clear and the wording of a text - the written word - is the main working tool. Here we worked with a proposal for regulation, and a complicated and for many areas comprehensive one, of great importance to many Swedish companies, authorities and people as citizens.

A clash between our way of thinking came to be, very much apparent, and we did not send the text to the press. The analysis was for me at a too immature stage, and I was also unsure of the role I had, as a researcher, to go public with a debate article. The clash prompted reflection: what was the purpose of our analysis, is writing a debate article a step in reaching a deeper understanding of the significance of the proposed regulation for various activities in society? How is policy development perceived, and done? Negotiation was still ongoing, the proposal still under development. There's no doubt our way of working with policy development, however successful, is not easily done. To approach policy development with agile thinking and as work in progress in practice takes courage coming from a profession of always being right, but is so rewarding since our analysis gets to be robustly rooted in concrete examples, hence possible to relate to and easier to understand for stakeholders. To me, Håkan's know-how of the agile process and ability to visualize relevant aspects within complexity of a system is vital in order for me to trust our approach to policy development.

4.3 Shared Retrospective on Going Public

Perhaps the most important aspect of an agile approach was to accept that we will not necessarily reach a point where we are finished or know what the analysis will be. We just have to accept that there are possible unknown unknowns and be prepared to address them when they show up. By keeping a dialogue running with various stakeholders we are more likely to find them sooner than later. And since we do not have a date by which we must have a verdict we can afford ourselves to release partial analyses continuously. At the same time it is important to be clear about our own uncertainty and when we have changed our analysis so that external stakeholders can make their own judgements on our way of working and its implications for the overall outcome. There has been a strong sense of learning from each other, taking it in turns to be the expert and the novice (Collins, Brown and Newman, 1987; Burden and Steghöfer, 2019).

5. DISCUSSIONS AND FUTURE WORK

In terms of product it seems that our analysis is an appreciated contribution. We are today in a situation where companies and trade associations are contacting us with commercial offers for them to take part of our analysis in relation to their business. A similar stance is taken by several national authorities that want to collaborate to better understand what the impact will be on their current digitalisation efforts. A third consequence is that the governments office contacts us before the negotiation meetings to get our feedback on the articles under scope. This is not an honour exclusively given to us, other actors are also contacted to get different points of view before formalising the Swedish stance. Taken together, we foresee that we will continue working with the implications with the AI Act but with external funding from research and innovation funds instead of internal development funds. We have already started to share our analysis to a new audience by having an experience report accepted for publication at a research arena for transportation (Burden and Stenberg, 2022). We foresee more publications coming as the scope of our analysis grows.

In terms of process we would like to emphasise the impact of working in a cross-functional team. While two persons make up a small team we have had the mandate to plan, act and assess our own way of working. That entails how we have been free to find our own external stakeholders to interact with, what we want to achieve together with them and the subsequent evaluation. Our endeavour had been more difficult if someone else had been in charge of setting up the collaboration or we had been dependent on other teams to deliver. Another benefit of the cross-functional team has been the multi-disciplinary background and experiences we bring into the collaboration. Less senior or less diverse team members might have find it more difficult to take on the tasks we set ourselves. Finally, the teamwork built on a strong sense of trust which needed time to grow at the same time as we conducted our agile way of working. The trust was not only required as we conducted our analysis and interacted with external stakeholders, it was also essential during our retrospectives as we wanted to find out what works for us and what needs to change. It has not always been easy to find a common ground.

6. ACKNOWLEDGEMENTS

We would want to thank all those who have collaborated with us during the process. You are too numerous to recount. We would also like to thank the sponsors within RISE who have enabled us to conduct this work while exploring the possibilities for new collaborations. Finally, we really appreciated our discussions with our shepherd Avraham Poupko!

REFERENCES

- AI Act. 2021. Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial intelligence act) and amending certain union legislative acts. Brussels, Belgium.
- Roy Andersson and Lars Bendix. 2006. eXtreme teaching: A framework for continuous improvement, *Computer Science Education*, 16, 3, 175-184.
- Kent Beck, Mike Beedle, Arie van Bennekum, Alistair Cockburn, Ward Cunningham, Martin Fowler et al. 2001. Manifesto for Agile Software Development. Available from: <http://www.agilemanifesto.org/>.
- Håkan Burden. 2014. A scholarship approach to model-driven engineering. PhD Thesis, Chalmers University of Technology and University of Gothenburg. Gothenburg, Sweden.
- Håkan Burden, Cilli Sobiech, Kristina Andersson, Martin Skoglund och Susanne Stenberg. 2021. The role of policy labs for introducing autonomous vehicles. Presented at the ITS World Congress 2021. 11-15 Oct, 2021. Hamburg, Germany.
- Håkan Burden and Frances Sprei. 2020. Teaching sustainable development through entrepreneurial experiences. *International Journal of Sustainability in Higher Education*, 22, 1, 142–156.
- Håkan Burden and Jan-Philipp Steghöfer. 2019. Teaching and Fostering Reflection in Software Engineering Project Courses. In *Agile and Lean Concepts for Teaching and Learning: Bringing Methodologies from Industry to the Classroom*, 231-262.
- Håkan Burden and Susanne Stenberg. 2022. Implications of the AI Act in relation to mobility. Transport Research Arena. Lisbon, Portugal.
- Håkan Burden, Lisa Carlgren, Ted Sjöblom and Susanne Stenberg. 2022. The Swedish policy lab for maritime autonomous surface ships. Lisbon, Portugal.
- Alistair Cockburn. 2009. Software Carpaccio. Available from: <https://alistair.cockburn.us/wp-content/uploads/2018/02/Elephant-Carpaccio-exercise-instructions.pdf>.
- Allan Collins, John Seely Brown and Susan E. Newman. 1987. Cognitive apprenticeship: Teaching the craft of reading, writing and mathematics (Technical Report No. 403). BBN Laboratories, Cambridge, MA. Centre for the Study of Reading, University of Illinois.
- Digital Decade. 2021. European Commission, Communication on a Europe’s digital decade: 2030 digital targets. Brussels, Belgium.
- Sofia Jungstedt and Susanne Stenberg. 2018. Nya regler om faderskap och föräldraskap (SOU 2018:68). Stockholm, Sweden
- Jan-Philipp Steghöfer, Håkan Burden, Hiva Alahyari and Dominik Haneberg. 2017. No silver brick : Opportunities and limitations of teaching Scrum with Lego workshops. *Journal of Systems and Software*, 131, 230–247.
- Jan-Philipp Steghöfer, Eric Knauss, Emil ALegroth, Imed Hammouda, Håkan Burden and Morgan Ericsson. 2016. Teaching Agile : Addressing the conflict between project delivery and application of Agile methods. *Proceedings - International Conference on Software Engineering*, 303–312
- Olle Wästberg et al. 2014. *Demokratiutredningen (Ju 2014:19)*. Stockholm, Sweden.