



## UvA-DARE (Digital Academic Repository)

### An Interdisciplinary Toolbox for Researching the AI-Act

Metikoš, L.

**DOI**

[10.17176/20230908-062850-0](https://doi.org/10.17176/20230908-062850-0)

**Publication date**

2023

**Document Version**

Other version

**License**

CC BY-SA

[Link to publication](#)

**Citation for published version (APA):**

Metikoš, L. (2023). An Interdisciplinary Toolbox for Researching the AI-Act. Web publication or website, Verfassungsblog. <https://doi.org/10.17176/20230908-062850-0>

**General rights**

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

**Disclaimer/Complaints regulations**

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

# An Interdisciplinary Toolbox for Researching the AI-Act

---

Ljubisa Metikos

2023-09-08T03:10:59

[The proposed AI-act](#) (AIA) will fundamentally transform the production, distribution, and use of AI-systems across the EU. As a result of its extremely wide scope, the AIA contains many broad and vaguely formulated provisions. It is therefore clear to see that legal research has an important role to play in both clarifying and evaluating the AIA. To this end, legal researchers may employ a *legal-doctrinal* method, and focus on the AIA's provisions and recitals to describe or evaluate its obligations. However, legal-doctrinal research is not a panacea that can fully operationalize or evaluate the AIA on its own. Rather, with the support of *interdisciplinary* research, we can better understand the purpose of the AIA's vague provisions, test its real-life application, and create practical design requirements for the developers of AI-systems. To achieve these goals, legal-doctrinal research should be supplemented by normative, empirical, and computer-science based research. Additionally, such research should be specified for each of the different sectors that the AIA aims to regulate, in order to provide an in-depth and practical understanding of this new law.

Nevertheless, interdisciplinary research can be [a daunting endeavour](#), particularly for starting legal researchers. In this blogpost, I will give a short glimpse into the methodological toolbox. I will first discuss the issue of the AIA's wide scope, as well as the need to specify what its provisions entail for each of the individual sectors that it aims to regulate. I then zoom in on a specific case study – the requirement of human oversight in AI-assisted adjudication – to give a short and practical illustration of how interdisciplinary research can be conducted on the AIA. To this end, this blogpost is subdivided into four sections. These detail how one can perform internal and external normative research, empirical research, as well as computer-science based research. I finally conclude with a short reflection on the relationship between these different disciplines, and how they can strengthen one another.

## One-size fits all

The AIA aims to regulate AI in the broadest sense of the word. It does this by differentiating between different kinds of risk-levels. In annex III of the AIA, we can see a list of various sectors where the use of AI will be considered '*high risk*'. This list includes the use of AI in healthcare, migration, and judicial adjudicatory proceedings, to just name a few. Consequently, all these sectors will be subject to the same set of regulatory requirements. Unfortunately, these provisions are often very vague and are not specified for each of the widely different areas that the AIA aims to regulate.

To give an example, it is unclear what the requirement for '*human oversight*' under art. 14 AIA practically entail. Art. 14 (1) AIA requires high-risk AI-systems to be designed in such a way that '*they can be effectively overseen by natural persons*'.

But how can we design AI-systems in such a manner? What measures need to be taken? To operationalize this requirement, we first need to take a step back and define what the *goal* of human oversight actually is. Defining this goal is crucial, as it will influence what measures need to be taken to best meet it.

Let us first take a legal-doctrinal approach and see what the AIA itself has to say about the goals of human oversight. Art. 14 (2) AIA states that human oversight is aimed at preventing risks to '*health, safety or fundamental rights*' posed by high-risk AI-systems. Unfortunately, this wording is so incredibly broad that it does not give us much guidance. Solely focusing on the text of the AIA itself, is therefore insufficient to formulate the goal of art. 14 AIA.

Alternatively, we can try to use a *normative* method to formulate the goal of art. 14 AIA. We can start this endeavour by first focusing on a specific high-risk sector that the AIA aims to regulate. Different high-risk sectors that employ AI-systems are subject to different needs, goals, and obstacles, after all. Let us therefore now zoom in on the context of AI used in adjudication as a test-case for all the different research methods we shall discuss.

## Internal normative research

Having chosen our sector, we now need to ascertain what the goal is of human oversight in adjudication. We need to ask questions such as: '*what is the role that a human overseer fulfils in adjudication?*' '*What do we wish to achieve with human oversight requirements for judges?*' To this end, [normative legal research](#) can be quite useful. Normative legal research, in short, is concerned with what the law *should* be, and can help formulate policy goals for legislators. Typically, it is used to [evaluate](#) the law. By applying a certain normative perspective, we can say that law X is bad, because it conflicts with a specific human rights framework or a particular moral theory. However, in the following paragraphs, I will demonstrate how normative research can also help us to not only *evaluate*, but also *operationalize* the AIA.

Let us first look at the *internal* normative approach. Here, we still use a legal-doctrinal method. But, rather than focusing on the text of the AIA itself, we focus primarily on external legal sources. It does deserve mention, however, that legal scholars can differently delineate and conceptualize normative and legal-doctrinal research. This is because legal methodological concepts can be rather fluid, and context-dependent. For the purposes of this blogpost, I partially include legal-doctrinal research under the internal normative method.

To evaluate the AIA using an internal normative approach, we can look, for example, at human rights frameworks such as the European Convention on Human Rights (ECHR). We can use the ECHR to *critique* the AIA, by arguing that the AIA violates certain human rights. However, the ECHR can also help *clarify* our interpretation of the provisions of the AIA. Because art. 14 AIA aims to safeguard fundamental rights, it also aims to safeguard the right to a fair trial, which is laid down in art. 6 ECHR. We might therefore try to analyse how human oversight measures in

adjudication might help safeguard this right. Art. 6 ECHR provides, for example, a right to an impartial and independent judge. On the other hand, art. 14 (4) (b) AIA requires that human overseers should be aware of *'automation bias'*. By reading these provisions together, we can argue that one possible goal of human oversight requirements for judges, is to safeguard the independence of the judiciary against uncritical adherence to the output of an AI-system.

## External normative research

Normative research does not only have to rely on (positive) legal texts, however. Researchers can also include external non-legal sources to study the AIA. To define what a judge should do when they exercise *'human oversight'*, we can look at what others have written about the preferable ways in which judges should interact with AI-systems. To this end, we could for example look at international guidelines, such as the [Ethical Charter on the use of artificial intelligence in judicial systems](#). Moreover, it can also be useful to re-examine older notions of how judges should behave *in general*. We can for example look at philosophical theories that discuss [the virtues that a judge](#) might need to possess. We might then say that the goal of human oversight in adjudication is to enable a judge to exercise these virtues. These could for example be, again, independence and impartiality. Based on this, we could argue that one of the goals of human oversight is to protect the judge from undue interference from outside actors through AI-systems.

These external perspectives can help us to [evaluate](#) the AIA and make recommendations for legal reform. Here, it has to be noted that evaluative legal researchers cannot *implicitly* assume the validity of their normative perspectives, however. Rather, a methodologically responsible legal scholar should explicitly mention the external perspective that they are employing.

Besides evaluation, external normative research can also help us to *operationalize* the AIA's provisions. Relying on philosophical frameworks can help us to explain legal concepts when positive law sources are insufficient to provide sufficient clarity. The previous external normative sources that we have discussed, could therefore not only serve as an aspirational goal for legal reforms, but also as a definition of what the goal of human oversight actually entails.

In short, I laid down two possible avenues of using internal and external normative methods for research on the AIA: *operationalization* and *evaluation*. However, normative research might also be useful in many other endeavours. For those interested in exploring the other ways that external normative research can support legal-doctrinal research, I suggest reading up on the work of [Taekema and Van der Burg](#). They describe how philosophical perspectives can, for example, help explain why certain fundamental principles exist in a legal order, generate fresh new perspectives on old legal concepts, or structurally critique the legal order.

For those interested in these structural critiques of the law, there are a myriad of [Critical Theory](#)-based perspectives one might be able to use in legal research. For the sake of brevity, I will not be discussing them in this blogpost. However, for

starting researchers who wish to learn more about this topic, I recommend watching George Tsouris' series of video essays over on YouTube. There, he discusses topics such as [the different feminist legal theories](#), as well as [Critical Race Theory](#), in an easy to understand format.

## Empirical research

In the previous sections, we used normative research to both operationalize and evaluate the goals of human oversight. Now that we have defined these goals, the next step is to understand what measures would be *effective* to achieve them. To this end, we might use *empirical* methods to help us evaluate the AIA.

In the simplest of terms, empirical legal research focuses on analysing data to understand legal phenomena. For empirical research on the AIA, a variety of different methods might be applicable, depending on the topic of research. These can include the use of case studies, interviews, content analysis, surveys, ethnography, and various kinds of experiments. Fundamentally, we can delineate between [quantitative and qualitative](#) empirical research methods. We can describe quantitative methods to be focused on statistically analysing *numerical* data pertaining to legal phenomena. Qualitative research, on the other hand, is interested in *non-numerical* data to understand the underlying meaning, motivation, and context of legal phenomena.

Less experienced researchers who are interested in pursuing empirical legal research might find the [Leiden Law Methods Portal](#) to be a helpful first resource. This website is open-access and can serve as a valuable alternative for individuals without access to legal libraries or those unable to afford empirical research handbooks. It provides a step-by-step guide to all kinds of empirical methods, describing how to formulate a research question, gather and analyse datasets, as well as how to write an empirical paper.

Empirical research on the AIA can, however, be obstructed by two fundamental issues. The first obstacle is the fact that the AIA is still a *proposed* piece of legislation. It is therefore not yet possible to research any effects that have occurred as a result of its implementation. However, some of the requirements of the AIA have already been introduced in the past, either in other pieces of legislation, or as measures that particular organizations have taken. To better understand our test-case, human oversight measures in adjudication, we might turn to the myriad of human oversight measures that have already been implemented in other contexts. We can for example look at the research of [Ben Green](#), who has researched and criticized the real-life effectiveness of human oversight in other kinds of decision-making processes.

A second obstacle to empirical legal research on the AIA, may arise from the fact that AI might not yet be used much in the sector that you are researching. This is for example the case in adjudication. A lack of adoption does not mean, however, that research is not warranted in such fields. The societal adoption of AI is occurring

at a break-neck speed. Pre-emptive research that tries to inform and future-proof legislation can therefore definitely be warranted.

Nevertheless, to evaluate art. 14 AIA's requirements on human oversight in judiciaries, it is important to first ascertain how users would actually interact with AI-systems. A pilot program might therefore be developed for this purpose, as Manuela van der Put did for [the Dutch court of 's Hertogenbosch](#) during her PhD research project. As part of her research, Van der Put performed qualitative research by interviewing the paralegals that used her AI-system. This led to the discovery that the paralegals distrusted the system, often preferring not to use it. However, art. 14 (4) (b) AIA heavily emphasizes that users of high-risk systems must remain aware of '*automation bias*' and not overly rely on the AI-system. Considering this empirical data, we can now, for example, criticize the need for such a provision in the context of adjudication.

## Computer science-based research

Computer science-based methods also play an important role in legal research on the AIA. This is because the AIA contains a number of provisions that require us to design AI-systems in a certain manner. Art. 14 (1) AIA, for example, requires that high-risk AI-systems '*be designed and developed in such a way to ensure that their operation is sufficiently transparent to enable users to interpret the system's output and use it appropriately*'. But what does this mean in practice? How can an AI-developer create a transparent and interpretable system? Before computer scientists can engage with the requirements of the AIA, we need to operationalize these provisions further.

Here, the previously discussed normative and empirical perspectives, cross paths with the field of computer science. This is because normative and empirical research methods can help kick-off computer science-based legal research. First of all, normative research can practically help operationalize the AIA for AI developers. A good example of this can be found in the paper '[News exposure diversity as a design principle for recommender systems](#)'. Here, Helberger et al. operationalized the regulatory concept of '*news exposure diversity*' for programmers. By defining this term in light of a number of different philosophical theories, they developed concrete parameters and benchmarks for the design of recommender systems. This example shows us that normative methods can help conceptualize practical guidelines for AI developers, making compliance with the AIA easier.

Empirical research can also help us to operationalize the AIA for AI developers. Art. 14 AIA, for example, requires that users of high-risk AI have easy access to a variety of control mechanisms to oversee the functioning of the AI-system. Empirical methods can show us what user interfaces could best enable this goal by testing how users actually interact with AI-systems. This in turn, can help AI developers design control mechanisms that best suit the goals of art. 14 AIA. An example of this can be found in the work of [Harambam et al.](#) on control mechanisms for algorithmic news recommender systems. Through the use of '*moderated think-aloud sessions*'



Harambam et al. tried to see what kind of user interface could best promote certain normative goals.

These examples of normative and empirical research interacting with computer science-based research, show us that interdisciplinarity in law really is a *connected* and *constructive* endeavour. Different kinds of research methodologies can build on one another. Research from one field could give rise to new research questions in another, leading to knowledge that was unattainable if one had stayed within their methodological comfort zone.

## Conclusion

The AIA needs interdisciplinary research to address its vagueness, evaluate its effectiveness, and operationalize its provisions. Moreover, such research should also be specified for each of the different high-risk sectors that the AIA aims to regulate. This way, we can do justice to the different needs and obstacles that the use of AI faces in these areas. To this end, we can use normative methods to uncover what the AIA aims to achieve in a given sector, and whether those aims are desirable in the first place. Empirical research, in turn, can test what measures can help achieve these normative goals in practice. Lastly, by using both normative and empirical methods in collaboration with computer science-based research, we can better operationalize the abstract design requirements of the AIA.

To study the AIA, legal researchers need to dare to step outside of the limits of purely legal-doctrinal analysis. So, I encourage you, the researcher reading this blog post, to start working. Reach out, cross fields, and explore new ways to do legal research!

*I wish to thank prof. dr. Natali Helberger and Isabella Banks from the Institute for Information Law, University of Amsterdam and mr. Marleen Kappé from the Institute for European Law, KU Leuven for their helpful thoughts and comments.*

