FAIR AI SCRUM HANDBOOK

Eliza Hobo
Nina van Liebergen
Quirine Smit
Cor Veenman

TNO

Co-funded by
the European Union

CONTENTS

Introduction	4
Pillar A: Share the responsibility	10
Pillar B: Examine your own position	12
Pillar C: Let communities co-own participation	13
Pillar D: Position the AI in its social context	15
Pillar E: Find the nuance in data, models and metrics	16
Case Study: University Admissions System	19
Appendices	23

This first edition of the Fair AI Scrum Handbook (2026) was created within DIVERSIFAIR. DIVERSIFAIR is an Erasmus+ project (2023–2026) that brings together eight partners from six European countries: CorTexter (NL), Eticas (ES), Sciences Po (FR), TNO (NL), Turing College (LT), University College Dublin (IE), Women4Cyber (BE) and Women in AI (FR). The project aims to support a new generation of AI experts who not only have strong technical skills, but also understand how to identify and address intersectional biases. More information is available at: https://diversifair-project.eu

This project has received funding from the European Education and Culture Executive Agency (EACEA) in the framework of Erasmus+, EU solidarity Corps A.2 - Skills and Innovation under grant agreement 101107969.

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Culture Executive Agency (EACEA). Neither the European Union nor the EACEA can be held responsible for them.

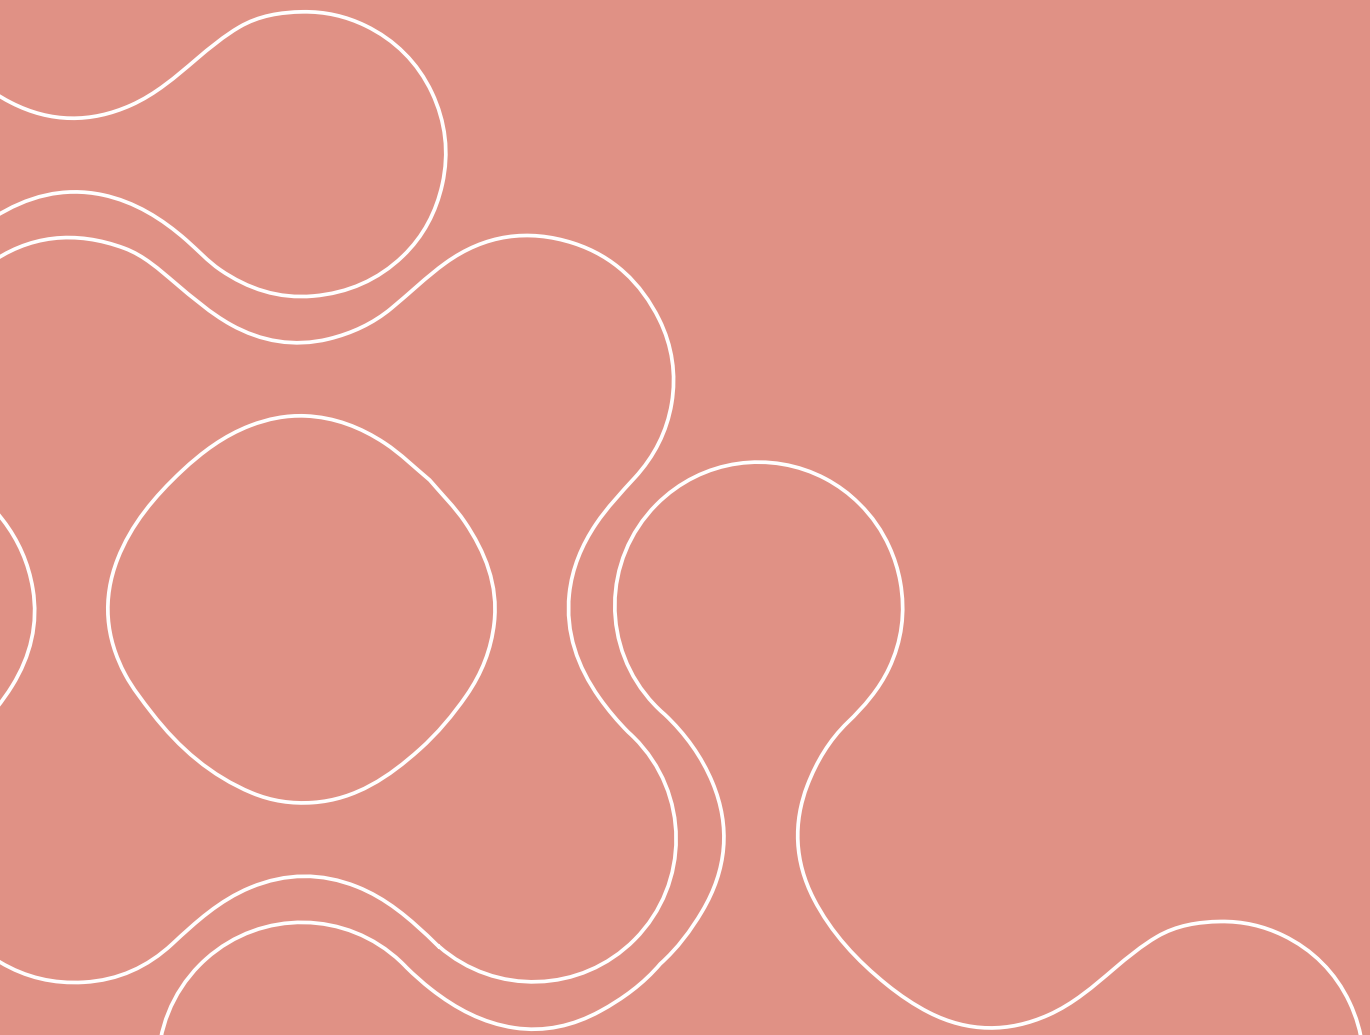

INTRODUCTION

As AI systems become increasingly embedded in the world around us, from health care to hiring and law enforcement, the question of fairness in their design and deployment has grown in urgency. This handbook serves as a guide for teams developing fair AI solutions, specifically those developing high-risk and limited-risk AI applications, drawing upon the Agile Scrum framework. Agile Scrum is widely adopted by teams to develop technology and AI solutions and is grounded in the principles of respect, openness and courage, all which are essential to developing fair AI solutions.

While many fair AI efforts focus on technical fixes within the data and models, this handbook goes further. Achieving fair solutions requires attention not just to algorithms, but also to the social context in which AI operates. This handbook offers practical tips and recommendations to help you to embed fairness throughout your development processes, leading to fairer solutions with real-world impact.

For each Scrum Role, Event and Artifact this handbook offers various recommendations designed to incorporate fairness at every level. By making the development process itself fairer, we believe the resulting solution will be fairer too.

In the introduction we will further explain what we mean with fairness and how it relates to Scrum. Then we explain how this handbook can be used and give some tips on how to implement this new process.

Disclaimer:
Fair AI Cannot be Achieved with a Checklist

This handbook is designed to support Scrum teams in integrating fairness into their AI solutions by providing starting points and tools. However, it is important to emphasize that fairness is not a fixed outcome, nor something that can be achieved simply by following a checklist. Instead, it requires ongoing dialogue, reflection, and adaptation. This process can be difficult and uncomfortable, it will take time and not everything will go right the first time. Adhering to the guidelines in this handbook does not guarantee a fair AI solution, but it does help teams move actively in that direction.

Fairness

Throughout this handbook, we understand fairness through the lens of intersectionality. This is a broader, more nuanced view than is commonly found in AI, where fairness often focuses on bias mitigation methods and outcome metrics, such as achieving equal performance across different demographic groups. In contrast, intersectionality focuses on the social and societal effects of AI solutions on people's lived experiences. As we cannot explore every nuance of intersectionality here, we assume readers have a foundational understanding of the concept. If you're new to intersectionality in AI, see the appendix for a selection of resources, ranging from videos and blog posts and scientific papers.

Adopting an intersectional definition of fairness means accepting that achieving a "fair" or "unbiased" AI solution in an absolute sense is impossible. Instead, it invites working towards fairer AI systems, that work towards social justice, "the idea that all people should have the same rights and opportunities (...)"[1]. This extends the goal from equality (treating everyone the same) toward equity (providing what each person needs to support), as is illustrated in Figure 1. As a consequence, the Fair AI Scrum Handbook is not a checklist or a step-by-step guide to achieve a "fair" AI solution. Instead, it assists in creating fair development process.

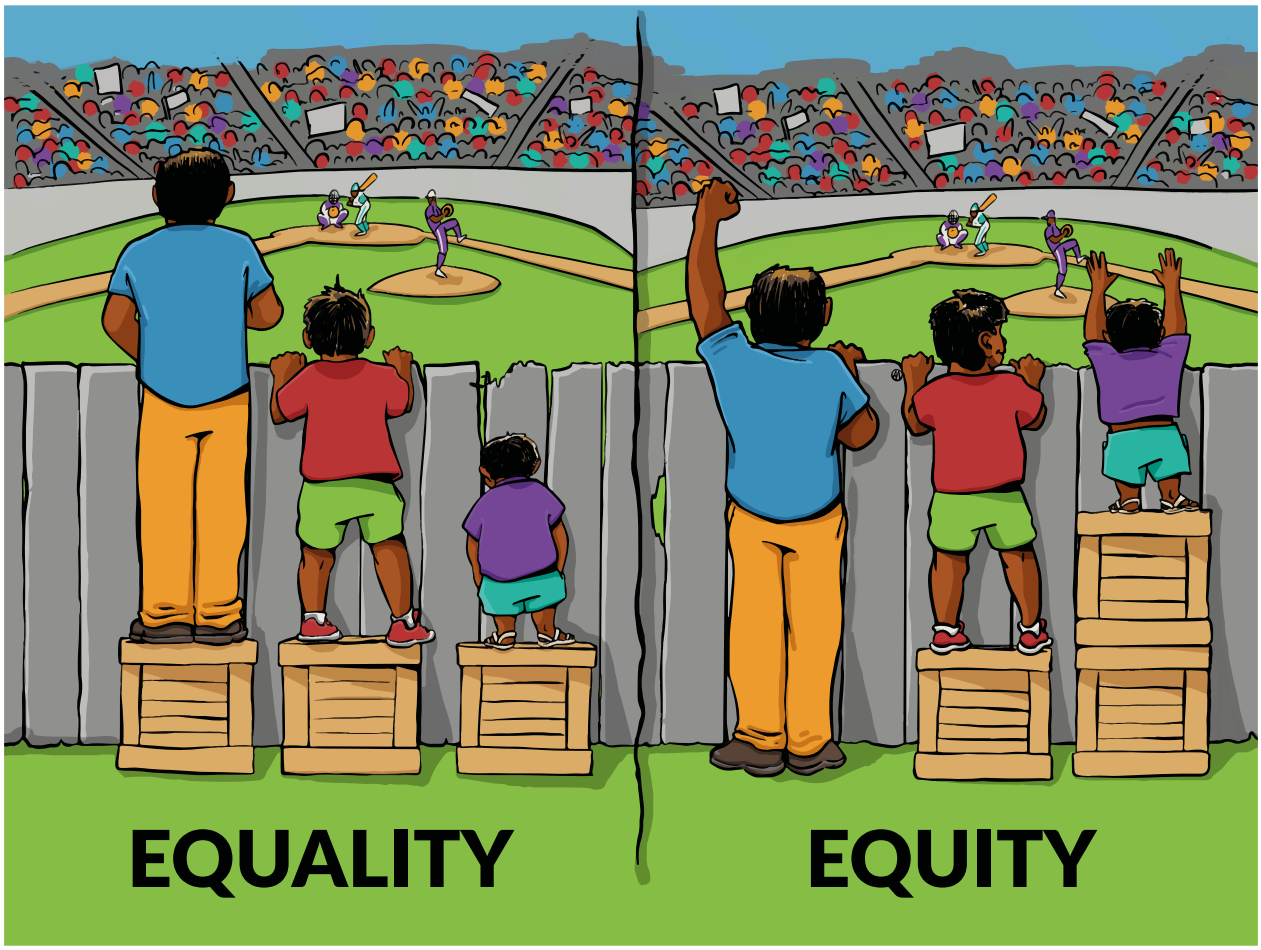

Figure 1: The difference between equality and equity[2]

1 https://dictionary.cambridge.org/dictionary/english/social-justice

2 https://interactioninstitute.org/illustrating-equality-vs-equity/, Interaction Institute for Social Change | Artist: Angus Maguire

Taking an intersectional approach to fairness involves, for example, carefully considering who develops, uses and is affected by AI systems. A diverse and inclusive development team is key to intersectional fairness. Research[3] consistently shows that such teams are more likely to create products that address the needs of a wide range of users, spot potential risks early, and deliver better results overall. Team members also report higher satisfaction, creativity, and resilience. In short, embracing intersectionality helps foster both fairness and stronger, more innovative teams.

Agile and Scrum

Agile is a philosophy and set of guiding principles for software development. It works iteratively and emphasises flexibility and collaboration. Rather than following a strict, linear process, Agile teams work in short cycles, constantly evaluating and adapting their work to meet changing requirements and customer needs.

Scrum is one of the most popular frameworks that teams use to apply the Agile principles. It provides a structured set of Roles, Events and Artifacts that guide teams in their work. Scrum helps teams break complex work into manageable pieces, fosters regular feedback and communication, and encourages continuous improvement.

Consequently, Scrum is well-suited method for fair AI development:

- It is iterative and adaptive, so plans can be revised as teams continue to learn about the solution and its social context. As the project progresses and more information becomes available, additional risks may be identified which may lead to a different solution.

- It builds in time for reflection and discussion, allowing all perspectives to be heard. Transparency, one of the pillars of Scrum, contributes to this.

- Openness and courage are two of the Scrum values. This is important as it encourages everyone to share their expertise and what they consider important.

3 Dixon-Fyle, S., Dolan, K., Hunt, V., & Prince, S. (2020, May 19). Diversity wins: How inclusion matters. McKinsey & Company. https://www.mckinsey.com/featured-insights/diversity-and-inclusion/diversity-wins-how-inclusion-matters Galinsky, A. D., Todd, A. R., Homan, A. C., Phillips, K. W., Apfelbaum, E. P., Sasaki, S. J., ... & Maddux, W. W. (2015). Maximizing the gains and minimizing the pains of diversity: A policy perspective. Perspectives on Psychological Science, 10(6), 742-748., ISO 690, Herring, C. (2009). Does diversity pay? Race, gender, and the business case for diversity. American sociological review, 74(2), 208-224.

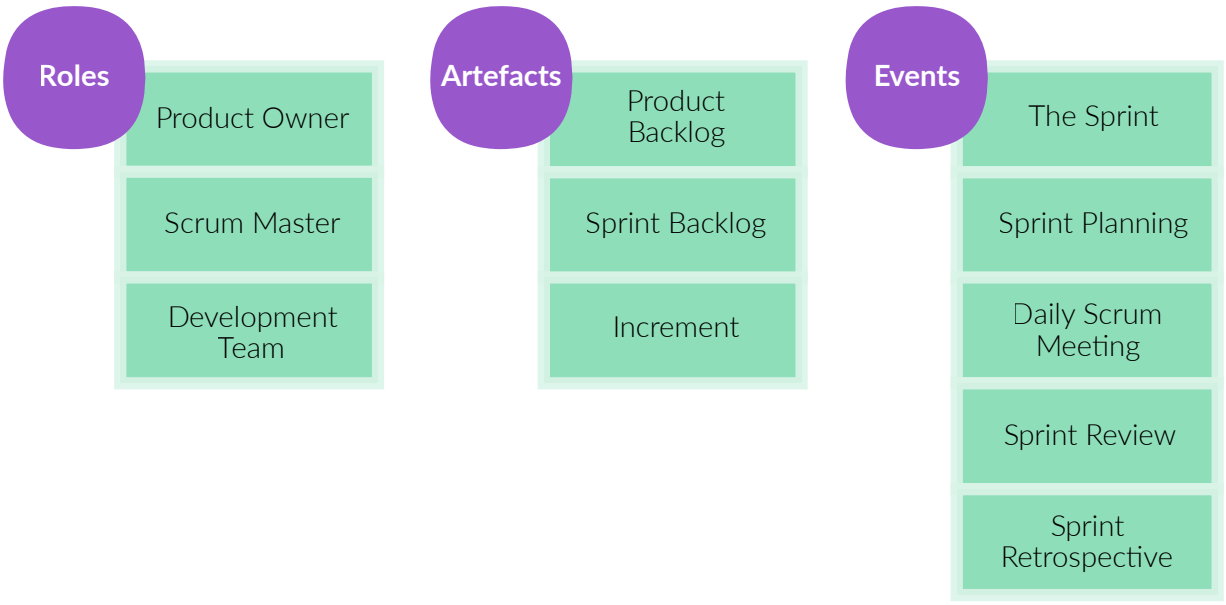

Figure 2: The Roles, Artefacts and Events in Scrum

However, some aspects of the Scrum methodology can challenge the intersectional interpretation of fairness outlined in this handbook. Scrum proposes non-hierarchical teams where everyone is considered equally important and all opinions are valued. While this is a great ambition, perceived hierarchies often persist in practice. These may arise from company dynamics such as seniority, or from social discriminating factors outside the workplace, including sexism, racism, ableism, homophobia, transphobia, and other systems of oppression. Even with the best intentions, some team members may still experience or perceive these hierarchies.

Additionally, Scrum is inherently product focussed. It starts with defining a Product Goal, and while this may be adapted during the process the focus is very much on developing a product. In contrast, Fair AI Scrum is aimed at solving an issue instead of the developing tool: with the end goal of social justice. A solution may take the form of an AI tool or application but can also be a non-AI or even non-technical solution. What ever best contributes to fairness in the real world.

Applying the Handbook

The Fair AI Scrum methodology embeds fairness in AI development by dividing it into five pillars and aligning these with Scrum Roles, Events, and Artefacts. These pillars are based on the five key recommendations from the scientific article *Fairness beyond the Algorithmic Frame* (Vethman et al., 2025). Each chapter of the Fair AI Scrum handbook introduces one pillar and explains its rationale. Practical tips for integrating within Scrum are indicated by the green boxes.

Figure 3: The five pillars of Fair AI Scrum

The tips and guidelines in the handbook are not meant as a checklist nor as a comprehensive guide. Instead, they serve as starting points for discussion and creative thinking, and we invite you to adapt them to your specific context. We understand that factors such as time, budget and organisational structures can naturally shape what is possible, and that not everything is relevant in every situation. We encourage you to thoughtfully consider the reasons behind the recommendations and to adopt what feels possible for your team. Please approach this work with openness and courage, embracing learning opportunities and recognising that occasional missteps are part of growth. Finally, we emphasize that progress is more important than perfection.

To integrate Fair AI Scrum into your existing Scrum process, we recommend you begin with reviewing and documenting your current Scrum practices to identify opportunities for introducing fairness considerations. Then you can start to modify Scrum Roles, Events and Artifacts based on the recommendations in this handbook. Pilot these modifications in a single sprint and gather feedback. Finally, you can continuously refine your approach by gradually expanding implementation as the team becomes more comfortable.

This handbook is intended for everyone in the Scrum team, from Product Owners and Scrum Masters to Developers. While the Fair AI Scrum methodology is especially relevant to AI projects, it can help in other technological or non-technological projects too. We specifically encourage projects that classify as “high-risk” or “limited-risk” as defined by the EU AI Act to adopt Fair AI Scrum.

PILLAR A: SHARE THE RESPONSIBILITY

The first pillar of Fair AI Scrum states that fairness is a shared responsibility. Often, one person handles fairness as a technical task, overlooking the broader impact of AI solutions. Instead, teams with diverse disciplines should address fairness together. This section offers ways to distribute responsibility for fairness across diverse perspectives.

Example: A team of software developers builds a CV filtering algorithm, to select candidates for job interviews. They find out that the algorithm is biased towards women and they try to improve it by masking the names on the CVs. After discussing this with a multidisciplinary team, they realised their intervention was not effective for diverse hiring. Together with social scientists and labour market experts, they changed the vacancy and updated the algorithm.

1. Multiple disciplines participate

It is beneficial if teams consist of people with diverse professional backgrounds. This allows for incorporating diverse perspectives into the development process by design, not as an afterthought. These perspectives can come from other professional disciplines, such as social scientists and domain experts, or from people who know the social context of the AI solution.

First Sprint(s): Reflect on the diversity of disciplines in your team and find ways to invite more diverse backgrounds: by joining as a member or through regular check ins.

Sprint Planning/Sprint Review: Invite stakeholders and experts from various disciplines to provide feedback on your ideas and plans. Ensure that everyone can give feedback and contribute, regardless of their technical expertise.

Definition of Done: Define the DoD with multiple disciplines, so fairness is integrated in each increment.

2. Take time to develop a common language between disciplines

Working with multiple disciplines can be time intensive, as each discipline has its own system of reference and vocabulary. For example, equity can mean something different to a technical developer (the same numbers as outcome) than a social scientist (having the same opportunities in society). Similarly, the word model can refer to various concepts in different disciplines. Take time for this exploration, since it can be frustrating to not understand each other immediately.

First Sprint(s): Take time to find a common language.

Scrum Master: Facilitate sessions that promote finding a common language, for example by using games or designing visual diagrams of different scenario's.

3. Don't start with the technical details

Before the technical work starts, take time to reflect on the desired impact of the AI solution in the real-world. Keep the desired impact in mind throughout the process and keep monitoring it. If you discuss the potential positive and negative outcomes before development, the team will spend their time and energy on a worthy goal.

First Sprint(s): Focus on really understanding each other and the various disciplines. The aim of these Sprints is not to technically develop something, but to learn about the question, the team and the various disciplines. Be curious about how the disciplines differ and complement each other.

Sprint Goal: Set the first Sprint Goal to a non-technical output, but rather related to creating a mutual understanding of each other, their disciplines and the goal of the project.

4. Dedicate time and effort to create a psychologically safe environment

To facilitate meaningful conversations with different disciplines and stakeholders, it is of the utmost importance to ensure a psychologically safe environment. The Scrum Handbook states that the Scrum team is non-hierarchical. However, hierarchies will exist nonetheless (for example, because of seniority in the company or societal hierarchies). One needs to be aware of the possible power inequalities, for example due to knowledge differences, privileges or seniority in the work environment or different lived experiences in society. Opinions can be conflicting, but to understand all implications of the AI solution, all opinions are valuable.

Scrum Master: Work towards a psychologically safe environment. If available, ask for the support of professionals within the organization. Even if you have the feeling that every team member feels safe, it could be that some signals are invisible to you. Therefore, it is always good to ask for a second (independent) check. Ensure dedication to psychological safety of everyone involved in all scrum events: Sprint, Retrospective, Review, Daily Scrum.

Scrum Master: If possible, try to keep your Scrum team small. This improves the psychological safety, since people are more easily seen and heard and aligns with the Scrum handbook.

First Sprint: Refrain from setting a specific goal in the first discussions, as this can create a frame where critical voices are not heard.

PILLAR B: EXAMINE YOUR OWN POSITION

The second pillar of Fair AI Scrum is examining your own position. Your perspectives are shaped by your academic and professional training, as well as your lived experiences such as age, gender, cultural background, and upbringing. This influences how you approach questions of fairness.

The recommendations in this section describe how to get started with cultivating awareness about one's own background in relation to the project, individually and as a multidisciplinary team. It is important to be aware of the perspectives represented in the team, and which perspectives might be missing, specific to the context of the project.

Example: a team developing an AI tool for job candidate screening consisted entirely of engineers with similar demographic backgrounds. They overlooked how their assumptions shaped the definition of a "qualified" candidate. After reflecting on their own positions, they invited HR staff and community advocates to join the process. This broadened the team's perspective and led to a fairer, more inclusive screening model.

1. Each team member writes a positionality statement and reflects on it.[4]

A positionality statement describes one's position in society, their beliefs and what shapes them and how they come into play in the project. Writing a positionality statement can open your eyes to the uniqueness of your position, and that of everyone else.

Retrospective: Invite all team members to write a positionality statement in the first retrospective. In subsequent retrospectives, they can reflect on it and identify potential new relevant positions.

2. Document perspectives and decisions throughout the project

By being open to new insights and perspectives, you learn a lot about the impact of your project. It is essential to thoroughly document significant decisions and identified risks as the project progresses. This practice helps guard against negligence by creating a transparent record of why certain choices were made and how potential risks were considered.

Sprint Planning: Document important perspectives and decisions about sprint goals, priorities and risk assessment.

Backlog: In the Product and Sprint backlog, keep track of all the identified risks of the AI application.

4 Example: The Equality Institute | Why positioning identity matters in decolonising research and knowledge production: How to write a 'positionality statement'

PILLAR C: LET COMMUNITIES CO-OWN PARTICIPATION

The third Fair AI Scrum pillar is co-owning the process with communities. Solutions are fairer stakeholders are actively involved throughout the development process. Think of stakeholders here as including the communities who will use the AI solution as well as other groups who may be impacted by it. By inviting participation, teams enable these stakeholders to speak up, learn about the process, and assess whether the system works well for them. This co-ownership fosters real-world fairness. Note that this requires active research and reaching out. Involving community members may be easier in some contexts than others.

Example: A city launched an AI tool to manage public housing applications without consulting residents. The system overlooked local needs and created barriers for vulnerable groups. After reaching out to affected communities, the city redesigned the process with their input. The AI now reflects real-world priorities and improves access fairly.

1. Invite communities to co-own the participation process

   Including communities who will use the solution and those who may be impacted by the solution as part of the stakeholder group can both improve the quality of the product and make it fairer.

First Sprint(s): Include stakeholders who may be impacted by the AI from the very first Sprints, when the project's goals are defined.

Product Owner: As the guardian of the Product Goal, the Product Owner must consider not only the client's objectives but also the broader impact of the product. What risks does it pose? Who could be affected? Take into account the needs, values, and concerns of impacted stakeholders, and integrate them into the vision and development priorities.

Scrum Master: Organize inclusive sessions with stakeholders, where all voices are welcomed and psychological safety is promoted. Encourage discussions on topics such as long-term sustainability and ethical implications.

Developers: Invite people who may be impacted by the product to be part of the team of Developers. Try to include as many perspectives as possible, reflect on who is missing and document the risks.

Sprint Planning: Invite community stakeholders to participate in Sprint Planning. Ensure they understand the process and are empowered to influence decisions. Sprint Planning, as outlined in the Scrum Guide, is already open to external input—use this as a recurring opportunity for feedback rather than limiting it to the project's start or conclusion.

Sprint Review: Include affected stakeholders in the Sprint Review process. Use this opportunity to reflect on outcomes, surface challenges, and assess potential risks collaboratively.

Sprint Planning/Sprint Review: Recognize that the goals of stakeholders from impacted communities may diverge from those of the Product Owner and Scrum Master. In some cases, holding separate sessions for community members can foster more open and honest dialogue, especially when psychological safety is a concern.

2. Make sure the involved communities have the knowledge and power to make a difference.

When involving a wider group of stakeholders and seeking their input, it is important that they are can contribute effectively. The discussion should be conducted in a language accessible to all participants, and technical information should be presented clearly to accommodate those without specialized training

First Sprint(s): Spend time understanding and educating each other. Ensure everyone involved is able to understand the discussions, to give them a voice to participate.

Sprint Retrospective: Plan an additional Sprint Retrospective with the Scrum Master, Product Owner and select stakeholder groups, without the Developers, to ensure open communication.

Product Owner: Reflect on the conversations with stakeholders: are they able to openly discuss their worries and priorities?
Product Backlog/ Sprint Backlog: Ensure documentation can be found and understood by all stakeholders.

3. Make participation financially sustainable for the involved communities

Insights from community members contribute to improving the solution, so providing compensation can help ensure that collaboration is mutually beneficial.

Scrum Master: Make the participation of community representatives mutually beneficial and financially sustainable. You can involve a third party to determine a fair compensation rate.

4. Design a mechanism where impacted communities can safely voice concerns

There can be a perceived hierarchy when involving a diverse group of stakeholders. There are some actions that can be taken to ensure that the stakeholders can feel empowered to give honest feedback and input.

Scrum Master: Ensure there is a mechanism for impacted communities to voice their worries, both during development and in deployment. This process should combine passive channels for feedback with proactive outreach to engage and hear from affected individuals or groups.

Sprint Review: Take time to make concrete what the potential effects of mistakes or malfunctioning are. Document these considerations.

Sprint Review: Consult with stakeholders at regular intervals, to identify any unforeseen risks or ethical concerns. Their perspective can help catch blind spots before deployment.

PILLAR D: POSITION THE AI IN ITS SOCIAL CONTEXT

The fourth pillar of Fair AI Scrum is positioning the AI solution in its social context. It is necessary to understand the broader social context in which the project and the AI solution exist. For this, analyse the underlying power relations. Who benefits, who might be harmed, who makes the decisions? Recognising the uneven distribution of power in the social context of the project, helps teams to ensure their solutions contribute positively to fairness, rather than reinforcing existing inequalities.

Example: A school district used an AI tool to allocate funding based on standardized test scores. This favored wealthier schools and disadvantaged under-resourced communities. After analyzing who benefits and who is harmed, the district revised the model to include socioeconomic factors and community needs. The AI now supports a fairer distribution of resources.

1. Position the AI solution within social context and define the present power relations.
   An AI solution does not exist in a vacuum, but it interacts with people's lives. Make sure you position the AI solution in the historical, social and cultural context and the existing power relations. Who are the people affected by the AI system, what is their history, how do they view the problem and solution? Engage with these communities to accurately and fairly position the solution in the real world.

   First Sprint(s): Analyze the problem that the Product Goal is a solution to with the interdisciplinary Scrum team. Try to understand the problem through the lens of power relations and hierarchies.

   Scrum Master: Do an impact assessment of the AI solution to get a systematic view on the effects of the AI.

   Product Goal and Sprint Goal: Define the Product Goal and Sprint Goals such that they contribute towards social justice.

2. Redefine concepts with power and social context in mind.
   In awareness of power relations and social context, the priorities and goals may change. Continually discuss with the team and stakeholders how you understand and operationalize terms such as fairness, accountability, responsibility and transparency.

   Sprint Retrospective: Reflect on the understanding of concepts with power and social context in mind. A common understanding of fairness, accountability, responsibility and transparency in the social context yields more effective AI solutions.

   Product Goal: Carefully examine who is the client and who is formulating the Product Goal. Explore what their power relations are in the social context of the project. Consider what their goal is and how that relates to the goals of social justice. Take this into account throughout the process.

   Product Goal and Sprint Goal: Be open to redefining the Product Goal and Sprint Goals based on what is learnt during the process.

PILLAR E: FIND THE NUANCE IN DATA, MODELS AND METRICS

The fourth and final pillar of Fair AI Scrum is aimed at finding the nuances: a careful consideration of data, models, and (fairness) metrics. These are very helpful tools, do not capture the world in its full complexity, but rather as an abstract representation.

In some cases, reflecting on the use of data, models and metrics in the social context may reveal that a fair solution does not require AI or technology at all. Perhaps a non-technical solution is fairer in the real-world situation.

Where an AI solution is required to contribute towards fairness, it is still important to critically discuss the strengths, weaknesses and representativeness of the data, models and metrics used. Re-examine throughout the process how these abstractions reflect or fail to reflect society and the real world and how that affects the fairness of the solution.

Example: A city used historical crime data to train an AI model for to predict where crimes might happen. Those records were already bias: they showed more police activity in certain neighbourhoods, not necessarily more crime. After re-examining the data and metrics, the city paused the AI rollout and invested in community-led safety programs. Fairness was better served without technology.

1. Be critical on why you use data, models and metrics.
   Consider carefully why you want to use specific data, models and metrics: are they truly addressing the core issue? Why do you believe this approach is the right solution? Remain open to the possibility that the use of certain data, models, metrics, or even technology might not align with the overarching goal of social justice. Be willing to adjust your approach as needed.

First Sprint(s): Take the first Sprint(s) to analyze the problem that you are trying to solve. The goal should go beyond the technically solvable issues.

Product goal: Evaluate whether a technical or AI solution is necessary. Would a non-tech solution be better for the problem at hand?

Sprint review: Don't be afraid to redefine the Sprint Goal or Product Goal as a non-AI/ non-technical solution, so that it works more towards social justice.

Product/Sprint Backlog: Be critical with assigning high priority to using data and metrics.

Product Owner: Be open to the possibility that the solution will not be AI.

2. Be critical about the intended use and limitations of the data, models and metrics.

The use of data, models and metrics can support in developing an AI solution that supports social justice. However, remember that they are a limited representation of the world. It is crucial to explore and question the data and models that are used. What types of biases exist in the data? Who is represented in the data, who is not? How can the model be used? Which use is needed in the project?

Sprint Planning: Ensure there is enough time to critically test and evaluate the intended data, models and metrics from multidisciplinary perspectives.

Sprint Review: Engage stakeholders in finding or evaluating intended data, models and metrics.

3. Augment quantitative approaches with qualitative research and participatory design.

Even when the reason and goal for using data and metrics are clearly defined, data and metrics alone cannot fully capture the complexities and subtleties of discrimination and social justice issues. To address this, it is important to complement quantitative approaches with qualitative approaches, such as interviews and focus groups. They can reveal perspectives and experiences that numbers may overlook. By integrating both quantitative and qualitative methods, you can be more effective towards fairness and social justice.

Sprint Planning: Consider what to explore with qualitative measures. This is relevant for the Product Owner when discussing with the client and stakeholders, and for the Developers when working towards the Sprint Goals.

Developer Team: Aim for a diverse Developer team that is able to perform both qualitative and quantitative analyses.

4. Document clearly on the intended use and limitations of data, model and metrics.

Document clearly how you establish the added value and limitations of your data and metrics.

Definition of Done: Part of the Definition of Done is a documentation of the strengths and limitations and conditions of use. Clearly document the intended use and limitations of the data, model and metrics. Document the potential harms, if and how they were mitigated.

Sprint Review: Go over all increments and their documentation and discuss the impact of these considerations combined.

CASE STUDY: UNIVERSITY ADMISSIONS SYSTEM

The tips and examples in this handbook can be applied to many different use cases. In this chapter, we illustrate one of such examples: building an AI-based university admission system. The team assess their fairness challenges and uses fair AI Scrum to tackle them.

Background:

A university data science team set out to build an AI-based admissions system. The system would analyse application materials, predict student success, and generate initial rankings for the admissions committee to review. They aim to predict student success potential using application data, test scores, essays, and extracurricular activities.

Fairness Challenges:

Early testing showed socioeconomic and racial disparities in acceptance rates. The system favoured applicants from well-resourced high schools and penalized non-traditional academic paths. It also exhibited language biases when analysing personal statements from non-native English speakers. The team's initial attempt at integrating fairness failed despite good intentions. Their standard Scrum process created several gaps:

1. Artifact Gap: User stories focused solely on functional requirements. *"As an admissions officer, I want to rank applicants by predicted GPA"* lacked fairness dimensions, pushing developers to optimize only for correlation with historical GPA data.
2. Event Gap: Sprint reviews showed only aggregate accuracy metrics. This approach masked significant disparities across socioeconomic groups until late testing.
3. Role Confusion: No one owned fairness outcomes. The Product Owner prioritized efficiency features over fairness fixes. Developers assumed someone else would handle bias issues "later."

The social context made these gaps especially problematic. University admissions directly impact educational access and life opportunities. Historical discrimination patterns in education amplified the risk of perpetuating existing inequities through automated systems.

Solution Implementation

The team applied Fair AI Scrum to address these gaps. The team expanded beyond technical solutions to address the broader social context of university admissions, bringing in perspectives from student advocacy groups, faculty, and education researchers. The approach balanced fairness with other objectives by integrating fairness work into existing events. It emphasized disaggregated metrics alongside aggregate ones, allowing the team to address fairness without sacrificing overall system performance.

Update user stories

Before: "As an admissions officer, I want to rank applicants by predicted success."

After: "As an admissions officer, I want to rank applicants by predicted success, ensuring equivalent accuracy across socioeconomic backgrounds, racial groups, and geographic origins."

Update Workflow

Process Integration: embed fairness in workflow:

Add fairness tests to CI/CD pipeline.

Create minimum fairness thresholds as merge blockers.

Required fairness documentation for model changes.

Add to Definition of Done:

Disaggregated performance metrics across intersectional demographic groups.

Documentation of fairness interventions with results.

Review by a diverse panel including student advocates.

Update Events

Sprint Planning: explicit fairness risk analysis for each feature.

Daily Standups: dedicated time for fairness blockers.

Sprint Reviews: Required presentation of disaggregated metrics.

Retrospectives: Added "equity impact" as a standing topic.

Role Accountabilities

Product Owner: Prioritized fairness tasks alongside features.

Scrum Master: Tracked fairness blockers and facilitated discussions with stakeholders.

Developers: Implemented fairness tests within feature code.

Data Scientists: Analysed bias patterns and developed mitigation strategies.

Outcomes and Lessons

Their Fair AI Scrum implementation yielded significant results. This case demonstrates how Fair AI Scrum bridges fairness principles with practical implementation. By modifying standard Scrum elements rather than creating separate fairness activities, the team made fairness an integral part of their development process rather than a separate consideration

1. Process Improvements:
- All user stories now include explicit fairness dimensions.
- Fairness issues surface 2-3 sprints earlier in development.
- Fairness test coverage jumped from 30% to 92% of features.

2. Fairness Metrics Improvements:
- Socioeconomic disparity in acceptance recommendations dropped from 18% to 4%.
- Racial performance differences decreased from 15% to 6%.
- Geographic origin biases reduced by 78%.

3. Business Outcomes:
- Reduced compliance risk and potential legal challenges.
- Stronger alignment with the university's equity mission.
- More diverse admitted student body with equivalent academic success rates.

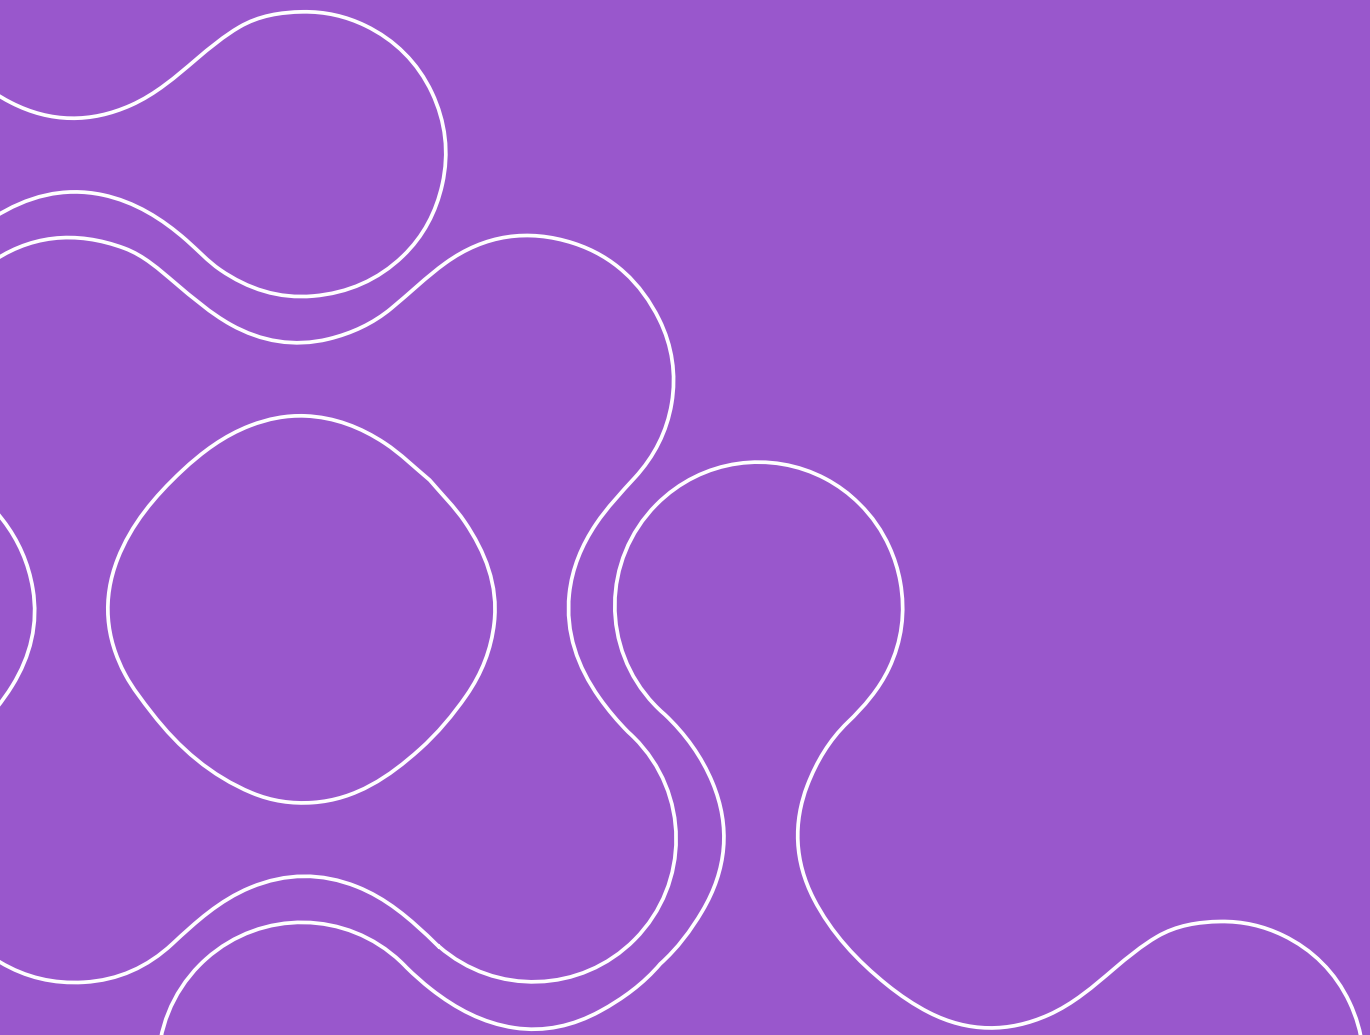

APPENDICES

Appendix A – References and Further reading

References

Steven Vethman, Quirine T. S. Smit, Nina M. van Liebergen, and Cor J. Veenman. 2025. Fairness Beyond the Algorithmic Frame: Actionable Recommendations for an Intersectional Approach. In Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25). Association for Computing Machinery, New York, NY, USA, 3276–3290. https://doi.org/10.1145/3715275.3732210

Buolamwini, J., & Gebru, T. (2018, January). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency* (pp. 77-91). PMLR.

What is intersectionality?

YouTube video by Peter Hopkins, Newcastle University
https://www.youtube.com/watch?v=O1islM0ytkE

Diversifair Columns

A selection of columns, delving into different aspects of intersectionality:

Introduction to Intersectionality and AI Fairness
https://diversifair-project.eu/courses/div-column-introduction-article/

Structures of Oppression
https://diversifair-project.eu/courses/div-column-structures-of-oppression/

Power and Intersectionality
https://diversifair-project.eu/courses/div-column-power-and-intersectionality/

AI Fairness
https://diversifair-project.eu/courses/div-column-ai-fairness/

Recommendations for Intersectional Fairness
https://diversifair-project.eu/courses/div-column-recommendations-ai-fairness/

Business case for diversity

Dixon-Fyle, S., Dolan, K., Hunt, V., & Prince, S. (2020, May 19). Diversity wins: How inclusion matters. McKinsey & Company. https://www.mckinsey.com/featured-insights/diversity-and-inclusion/diversity-wins-how-inclusion-matters

Galinsky, A. D., Todd, A. R., Homan, A. C., Phillips, K. W., Apfelbaum, E. P., Sasaki, S. J., ... & Maddux, W. W. (2015). Maximizing the gains and minimizing the pains of diversity: A policy perspective. *Perspectives on Psychological Science*, 10(6), 742-748., ISO 690,

Herring, C. (2009). Does diversity pay? Race, gender, and the business case for diversity. *American sociological review*, 74(2), 208-224.

Appendix B -
Tips per Role, Event & Artefact

The chapters in the handbook follow the five pillars of Fair AI Scrum and list tips for integrating the pillars in the different roles, events and artefacts. In this appendix, we list said tips, categorized by Scrum component:

Roles

Developers
Invite people who may be impacted by the product to be part of the team of Developers. Try to include as many perspectives as possible, reflect on who is missing and document the risks.

Aim for a diverse Developer team that is able to perform both qualitative and quantitative analyses.

Scrum Master
Facilitate sessions that promote finding a common language, for example by using games or designing visual diagrams of different scenarios.

Work towards a psychologically safe environment. If available, ask for the support of professionals within the organization. Even if you have the feeling that every team member feels safe, it could be that some signals are invisible to you. Therefore, it is always good to ask for a second (independent) check. Ensure dedication to psychological safety of everyone involved in all scrum events: Sprint, Retrospective, Review, Daily Scrum.

If possible, try to keep your Scrum team small. This improves the psychological safety, since people are more easily seen and heard and aligns with the Scrum handbook.

Organize inclusive sessions with stakeholders, where all voices are welcomed and psychological safety is promoted. Encourage discussions on topics such as long-term sustainability and ethical implications.

Make the participation of community representatives mutually beneficial and financially sustainable. You can involve a third party to determine a fair compensation rate.

Ensure there is a mechanism for impacted communities to voice their worries, both during development and in deployment. This process should combine passive channels for feedback with proactive outreach to engage and hear from affected individuals or groups.

Do an impact assessment of the AI solution to get a systematic view on the effects of the AI.

Product Owner
As the guardian of the Product Goal, the Product Owner must consider not only the client's objectives but also the broader impact of the product. What risks does it pose? Who could be affected? Take into account the needs, values, and concerns of impacted stakeholders, and integrate them into the vision and development priorities.

Reflect on the conversations with stakeholders: are they able to openly discuss their worries and priorities?

Be open to the possibility that the solution will not be AI.

Events

First Sprint(s)
Reflect on the diversity of disciplines in your team and find ways to invite more diverse backgrounds: by joining as a member or through regular check ins.

Take time to find a common language.

Focus on really understanding each other and the various disciplines. The aim of these Sprints is not to technically develop something, but to learn about the question, the team and the various disciplines. Be curious about how the disciplines differ and complement each other.

Refrain from setting a specific goal in the first discussions, as this can create a frame where critical voices are not heard.

Include stakeholders who may be impacted by the AI from the very first Sprints, when the project's goals are defined.

Spend time understanding and educating each other. Ensure everyone involved is able to understand the discussions, to give them a voice to participate.

Analyze the problem that the Product Goal is a solution to with the interdisciplinary Scrum team. Try to understand the problem through the lens of power relations and hierarchies.

Take the first Sprint(s) to analyze the problem that you are trying to solve. The goal should go beyond the technically solvable issues.

Sprint Planning
Invite stakeholders and experts from various disciplines to provide feedback on your ideas and plans. Ensure that everyone can give feedback and contribute, regardless of their technical expertise.

Document important perspectives and decisions about sprint goals, priorities and risk assessment.

Invite community stakeholders to participate in Sprint Planning. Ensure they understand the process and are empowered to influence decisions. Sprint Planning, as outlined in the Scrum Guide, is already open to external input—use this as a recurring opportunity for feedback rather than limiting it to the project's start or conclusion.

Recognize that the goals of stakeholders from impacted communities may diverge from those of the Product Owner and Scrum Master. In some cases, holding separate sessions for community members can foster more open and honest dialogue, especially when psychological safety is a concern.

Ensure there is enough time to critically test and evaluate the intended data, models and metrics from multidisciplinary perspectives.

Consider what to explore with qualitative measures. This is relevant for the Product Owner when discussing with the client and stakeholders, and for the Developers when working towards the Sprint Goals.

Sprint Review
Invite stakeholders and experts from various disciplines to provide feedback on your ideas and plans. Ensure that everyone can give feedback and contribute, regardless of their technical expertise.

Include affected stakeholders in the Sprint Review process. Use this opportunity to reflect on outcomes, surface challenges, and assess potential risks collaboratively.

Take time to make concrete what the potential effects of mistakes or malfunctioning are. Document these considerations.

Consult with stakeholders at regular intervals, to identify any unforeseen risks or ethical concerns. Their perspective can help catch blind spots before deployment.

Don't be afraid to redefine the Sprint Goal or Product Goal as a non-AI/non-technical solution, so that it works more towards social justice.

Engage stakeholders in finding or evaluating intended data, models and metrics.

Go over all increments and their documentation and discuss the impact of these considerations combined.

Recognize that the goals of stakeholders from impacted communities may diverge from those of the Product Owner and Scrum Master. In some cases, holding separate sessions for community members can foster more open and honest dialogue, especially when psychological safety is a concern.

Sprint Retrospective
Invite all team members to write a positionality statement in the first retrospective. In subsequent retrospectives, they can reflect on it and identify potential new relevant positions.

Plan an additional Sprint Retrospective with the Scrum Master, Product Owner and select stakeholder groups, without the Developers, to ensure open communication.

Reflect on the understanding of concepts with power and social context in mind. A common understanding of fairness, accountability, responsibility and transparency in the social context yields more effective AI solutions.

Artefacts

Product Backlog/Goal
Define the Product Goal and Sprint Goals such that they contribute towards social justice.

Carefully examine who is the client and who is formulating the Product Goal. Explore what their power relations are in the social context of the project. Consider what their goal is and how that relates to the goals of social justice. Take this into account throughout the process.

Be open to redefining the Product Goal and Sprint Goals based on what is learnt during the process. Evaluate whether a technical or AI solution is necessary. Would a non-tech solution be better for the problem at hand?

Be critical with assigning high priority to using data and metrics.

In the Product and Sprint backlog, keep track of all the identified risks of the AI application.

Sprint Backlog/Goal
Set the first Sprint Goal to a non-technical output, but rather related to creating a mutual understanding of each other, their disciplines and the goal of the project.

In the Product and Sprint backlog, keep track of all the identified risks of the AI application.

Ensure documentation can be found and understood by all stakeholders.

Define the Product Goal and Sprint Goals such that they contribute towards social justice.

Be open to redefining the Product Goal and Sprint Goals based on what is learnt during the process.

Be critical with assigning high priority to using data and metrics.

Increment/Definition of Done
Define the DoD with multiple disciplines, so fairness is integrated in each increment.

Part of the Definition of Done is a documentation of the strengths and limitations and conditions of use. Clearly document the intended use and limitations of the data, model and metrics. Document the potential harms, if and how they were mitigated.

DIVERSIFAIR

TNO

Co-funded by
the European Union