

How the BIAS project is re-coding fairness into Europe's AI future

The rapid integration of artificial intelligence (AI) into professional settings—from recruitment and training to engagement and disciplinary monitoring—promises speed and efficiency. Yet, this digital transformation presents a fundamental ethical challenge: how do we ensure that the algorithms governing access to opportunity are fair?

The European labour market, despite striving for equality, does show discrimination across factors such as gender and nationality. The central mission of the BIAS project is to address this very concern: ensuring AI either contributes to or actively mitigates discrimination.

The challenge: when code copies prejudice

AI in human resources (HR) is increasingly used to generate job postings, sift and rank large volumes of applications, and automatically extract information from

CVs and cover letters. It is also used to manage employees by generating shifts and monitoring schedules, working hours, efficiency, and other activities. The information gathered by AI can feed into automated decisions about who to invite for an interview, who to hire, or who to promote.

The power of AI, particularly natural language processing (NLP), lies in its ability to analyse vast amounts of text and make complex inferences. But NLP systems are not neutral: they learn from the data they are trained on. What happens if these systems are biased and reproduce societal stereotypes, leading

to unfair outcomes? Algorithmic bias is often encoded in training data that may reflect historical inequalities and can then be reproduced in models and automated decisions. These dynamics risk undermining the European Pillar of Social Rights on work and employment, as well as several United Nations Sustainable Development Goals.

The BIAS project, a 4-year initiative funded by the European Union's Horizon Europe research and innovation programme, was created to help the AI and human resource management (HRM) communities recognise, address, and mitigate these harmful algorithmic biases.

The four pillars: an interdisciplinary methodology

The BIAS project adopts a unique interdisciplinary research and impact methodology, ensuring that technological development is grounded in real-world human experience and policy requirements. This approach rests on 4 distinct pillars:

National Labs and stakeholder engagement

The project established [National Labs](#) in each participating country. These labs function as communities composed of practitioners, employees, HRM specialists, and AI experts, with a specific focus on underrepresented communities. Members of these labs are integral to the project's evolution, participating in needs analysis and stakeholder involvement via surveys, interviews, and essential co-creation workshops.

Objectives

The project is structured around 6 mutually reinforcing objectives designed to create both robust technological solutions and profound social change:

- 1. Develop novel tools.** Create reliable tools for identifying and mitigating bias in existing AI and NLP systems, forming the project's core technological capability.
- 2. Empower the community.** Raise awareness and build capacity in AI and HRM so professionals can design better technology and adopt practices that reduce algorithmic bias.
- 3. Enhance understanding.** Deepen knowledge of biases in recruitment to improve HRM practices, advance worker studies, and support effective capacity-building across Europe.
- 4. Broaden participation.** Increase the involvement of underrepresented individuals in AI research, bringing in diverse perspectives and revealing overlooked forms of bias.
- 5. Reduce bias.** Decrease biases in hiring by promoting fairer recruitment practices and supporting more transparent, non-discriminatory hiring decisions.
- 6. Co-create the Debiaser.** Engage stakeholders in co-creating the Debiaser, the BIAS proof-of-concept technology to identify and mitigate biases and unfairness in decision-making.

AI research and development (R&D)

The R&D pillar focuses on advanced AI methodologies, in particular NLP and case-based reasoning (CBR). CBR addresses new problems by drawing on past cases, reusing solutions that proved successful in similar situations.

The creation of the Debiaser, the BIAS proof-of-concept technology

The Debiaser aims to provide a toolkit for identifying and mitigating biases in language models, making them significantly safer for application in HRM. During the recruitment phase, the Debiaser highlights areas of application that may be subject to discrimination and offers suggestions for mitigation. Crucially, the tool provides human-understandable explanations for the automated decisions presented to recruiters, thereby building essential

trust in the process.

The fundamental challenge in automated recruitment is defining fairness. The Debiaser addresses this by creating a use-case-specific definition of fairness that ensures similar candidates are treated similarly, moving beyond generic, one-size-fits-all definitions.

Ethnographic fieldwork

Research integrity is ensured through comprehensive ethnographic fieldwork, with PhD candidates and researchers studying employers, employees, and AI developers across Europe to understand how AI shapes work now and in the future. Some partners have shared these insights in the [Bias On The Move](#) vlog series. This work underpins an impact strategy focused on raising awareness of gender and intersectional biases in AI, building capacity in the AI and HRM communities, and laying the ground for a product that helps companies reduce bias in their HR practices.



https://www.youtube.com/watch?v=Oy_BdBzGaiM



<https://www.youtube.com/watch?v=tCk8i5na09Q>

The infrastructure: consortium and advisory board

The success of the BIAS project relies on a balanced and expert infrastructure. Our consortium is robust, comprising 4 universities, 3 communication partners, one large industrial organisation, and one small and medium-sized enterprise (SME). This mix ensures a breadth of expertise spanning academic research, industrial application, and effective public communication.

The BIAS Advisory Board is guiding the consortium. Each member of the Advisory Board was selected based on their profound experience and knowledge, ensuring balanced expertise and geographical coverage. They provide essential advice on project implementation and dissemination throughout the project's entire lifespan.

Milestones: progress towards true fairness

The project has already achieved several significant milestones over the past 3 years, demonstrating the commitment to practical outputs and collaborative development.

Expert interviews

In an intensive period of 6 months, the project conducted 71 expert interviews, capturing detailed insights from 35 AI

developers and 36 HRM managers. This expansive dataset provides the necessary foundation for understanding current perspectives and informing the Debiaser's development.

The main conclusions of this action are available in a [factsheet](#).

Co-creation workshops: developing the Debiaser

Throughout 2023, the BIAS project ran a series of co-creation workshops to support the early development of the Debiaser tool and CBR systems: **national workshops (June–July 2023)** focused on identifying categories of wordlists for bias detection, producing terms and expressions that help AI experts build more robust detection models; **national workshops (August–October 2023)** explored how to define and operationalise fairness in early recruitment stages, from principles and process features to how candidate attributes should be prioritised; and an **international workshop (December 2023)** simulated AI-based recruitment using the Candidate Ranker and Mitigation Tool, aligning system requirements with ALTAI ethical guidelines and gathering inputs for future training packages.

Watch: [Co-creation workshops playlist](#)

National Labs

The core foundation was established by building national communities

of stakeholders across 2 main ecosystems: those involved in human resources and recruitment policies, and those composed of AI experts and practitioners. These stakeholders are continuously engaged in BIAS activities and initiatives, and the National Labs are constantly open to new members.

Training and dissemination

The project has initiated a dynamic **Capacity-Building programme** composed of 2 rounds of sessions. This programme explores the complex role of AI within HRM, ensuring participants gain essential knowledge and skills to recognise and address bias, explore how AI can support fairness, and reflect on the human impact of the technology in everyday work environments. The main outcomes of the training will be made available in an e-learning course very soon.

Dissemination efforts are comprehensive, including the formation of the **AI Fairness Cluster**—a network composed of the European projects AEQUITAS, BIAS, FINDHR, and MAMMOth. This cluster is a crucial component of the European Commission's strategy to ensure the trustworthiness of AI, working to raise awareness, coordinate forums, and collectively contribute input to the European Commission regarding emerging challenges.

Furthermore, the project has produced 5 awareness videos covering essential topics such as what bias is, how it's experienced, and the legal, social, and technical perspectives on fairness.



BIAS Spring School on Trustworthy and Fair AI

A major event was the **BIAS Spring School on Trustworthy and Fair AI** in Tallinn in April 2025. This event brought together academics, industry leaders, and policymakers for keynotes, capacity-building workshops, and hands-on coding sessions. A key highlight was the **AI Fairness Connect: Networking Reception on Trustworthy AI**, which fostered crucial connections between participants, business leaders, government representatives, and researchers.

Trustworthy AI Helix

The BIAS virtual helix brings together experts, researchers, and innovators from academia and industry in a specialised community focused on Trustworthy AI. It provides a structure for sharing BIAS activities and results, and for facilitating collaboration on technical and methodological issues.

Final thoughts

The BIAS project is not just developing technology; it is systematically building a multi-layered defence against inequality in the labour market. By uniting technological rigour, sociological insight, and stakeholder collaboration, the project is actively ensuring that the future of work is not just fast and efficient, but fundamentally fair.

BIAS

Mitigating biases of AI in the labour market

PROJECT SUMMARY

The BIAS project is a European Union-funded initiative that empowers the artificial intelligence and human resources management communities. Through interdisciplinary research, a capacity-building programme, and the creation of a proof-of-concept technology to mitigate algorithmic bias in recruitment, the project aims to ensure fairer practices across Europe.

PROJECT PARTNERS

The BIAS consortium brings together 9 partners from 9 countries—4 universities, 3 communication organisations, 1 large company, and 1 SME—covering AI/NLP, social sciences, diversity and inclusion in HR, communication, and industrial uptake. The project is coordinated by the Norwegian University of Science and Technology (NTNU).

PROJECT LEAD PROFILE

Roger A. Søråa is the BIAS Principal Investigator and Professor in NTNU's Department of Interdisciplinary Studies of Culture. He leads DigiKULT, a research group on digital technologies, social robots, and the cultural practices they shape, providing strong interdisciplinary leadership for this diverse consortium.

PROJECT CONTACTS

BIAS Principal Investigator: Roger A. Søråa

✉ roger.soraa@ntnu.no

Project Coordination Team

✉ info@biasproject.eu

🌐 biasproject.eu

📘 [@BIASProjectEU](https://www.facebook.com/BIASProjectEU)

📧 [@BIASProjectEU](https://twitter.com/BIASProjectEU)

📷 [@biasprojecteu](https://www.instagram.com/biasprojecteu)

🌐 [/company/biasprojecteu](https://www.linkedin.com/company/biasprojecteu)

📺 [@BIASProjectEU](https://www.youtube.com/@BIASProjectEU)

FUNDING



Funded by the European Union

Funded by the European Union. The Associated Partner Bern University of Applied Sciences has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI).

Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the Swiss State Secretariat for Education, Research and Innovation can be held responsible for them.

Project funded by



Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs,
Education and Research EAER
State Secretariat for Education,
Research and Innovation SERI



Adobe Stock, generated with AI © Emanuel