



**Global Dialogue
on AI Governance**



Requirement Specification – Equity Lens v0.0

This document is based on workshop feedback gathered through chat, oral Q&A, polls, and presentations. It consolidates the main points into an annotated requirement specification to support internal review and the next implementation phase.

Document Status: Draft (Post-Workshop)

Date: July 2026

Version: 0

Next Steps: Internal review → Pilot phase → Public release

1. Purpose

Equity Lens is a transparency tool that helps selection committees detect potential bias patterns before final decisions are made. It does not rank candidates, recommend outcomes, or record votes. Its purpose is to make invisible bias patterns visible so that committees can apply informed human judgment with greater confidence. In this way, the tool supports decision-making rather than replacing it, and it is designed to strengthen fairness review without taking ownership of the decision itself.

2. Core Requirements (from Workshop Feedback)

R1. Context-sensitive configuration

ID	Requirement	Source
R1.1	The tool shall allow committees to configure which bias flags are active for their specific selection context.	Chat Qc4, Qo5
R1.2	The tool shall support different selection types (professorships, junior hires, speaker invitations, awards) with context-appropriate criteria.	Qo5
R1.3	Committees shall be able to adjust weights for different metrics (e.g., h-index importance for certain research groups like particle physics or gender bias in citations).	Qc5b, Qo4
R1.4	The tool shall allow multiple selection rounds (e.g., shortlisting).	Qo1
R1.5	The tool shall allow committees to include soft skills criteria in candidate profiles, if relevant to the selection context.	Qo4

R2: Bias Detection (What to Flag)

ID	Requirement	Source
R2.2	Flag institutional concentration (when $\geq 40\%$ of nominees from one institution).	Core design
R2.3	Flag familiarity bias (candidates with ≥ 2 prior invitations in 3 years).	Core design, polls
R2.4	Flag career-stage mismatch (comparing junior to senior unfairly).	Core design
R2.5	Flag network effects (candidates from same institute, advisor or research group).	Polls
R2.6	Flag citation bias – note that h-index itself may carry gender bias.	Chat Qc5a
R2.7	Field-adjustable metrics – allow different h-index interpretations for different fields.	Chat Qc5b
R2.8	The tool shall offer double-blind mode as an optional setting (masking name and gender) in preventive mode, not as a default or mandate.	Qe1, Polls

R3: Privacy & Data Protection (DSGVO/GDPR Compliance)

ID	Requirement	Relevant GDPR Article	Source
R3.1	Never store individual votes – only aggregated totals.	Art. 5(1)(c) – Data minimisation	Core design
R3.2	Candidate names shall be hashed after 30 days.	Art. 32 – Security of processing	DPIA*
R3.3	Users shall be able to delete all their data (Right to erasure).	Art. 17 – Right to erasure ('right to be forgotten')	DPIA
R3.4	Data shall be encrypted at rest (AES-256) and in transit (TLS 1.3).	Art. 32 – Security of processing	DPIA
R3.5	Differential privacy ($\epsilon=0.1$) for all aggregated outputs – where applicable.	Art. 25 – Data protection by design and by default	DPIA
R3.6	Minimum cell size (<3 individuals = data suppressed).		Core design
R3.7	Users shall be able to opt out of data being used for AI training.	Art. 7 – Conditions for consent (and Recital 42, 43)	DPIA

*DPIA: Data Protection Impact Assessment

R4: Usability & Workflow

ID	Requirement	Source
R4.1	5-minute rule – total committee time from login to vote ≤5 minutes.	Core design
R4.2	Automated candidate upload – support CSV or similar bulk upload.	Chat Qc6
R4.3	No training required – intuitive interface with clear labels.	Core design
R4.4	Plain-language flags – each flag explains what it means and suggests an action.	Core design
R4.5	Multi-device support – works on desktop, tablet, and mobile.	Core design

R5: Verification & Trust

ID	Requirement	Source
R5.1	The tool itself shall be auditable – code open-source and inspectable.	Qo3, Polls
R5.2	The tool shall support third-party verification of its bias detection logic.	Polls
R5.3	The tool shall provide transparency on its own limitations – what it does not detect.	Polls
R5.4	Bias flags shall include confidence indicators (e.g., "strong signal" vs. "weak signal").	Polls
R5.5	The tool shall allow committees to dismiss flags with justification (optional).	Qc5, Qe1

R6: Transparency & Accountability

ID	Requirement	Source
R6.1	The tool shall provide a dashboard showing aggregated, anonymised patterns.	Core design
R6.2	The tool shall never influence committee choices – it only informs.	Core design
R6.3	The tool shall support export of anonymised session data for institutional learning.	Core design
R6.4	The tool shall include a privacy notice explaining what data is collected and why.	DPIA

R7: Committee Self-Awareness & Institutional Learning

ID	Requirement	Source
R7.1	The tool shall provide committees with aggregated, anonymised feedback on their own decision patterns over time (e.g., "In the last 3 years, your committee selected women 20% of the time from a pool that was 40% women").	Qc2
R7.2	The tool shall not identify or track individual committee members' votes – only aggregate patterns.	Core design, Qc2
R7.3	The tool shall provide comparison to anonymised benchmarks (e.g., "Compared to similar committees, yours selected fewer women from the same pool").	Qc2