

Session Outcome Document

Beyond Ethics and Regulation : Compassionate AI for a Just and Sustainable World

Satyarthi Movement for Global Compassion

8 July, 16:00 – 16:45 CEST

<https://www.itu.int/net4/wsis/forum/2026/Agenda/Session/574>

Key Issues discussed:

- **Translating principles into practice:** Responsible AI principles are broadly accepted but hard to operationalise. Standards were discussed as the "middle layer" between policy intent and technical implementation, and measurable indicators — as the SDGs gave the Millennium goals — as what turn aspiration into something technologists can build.
- **Whether values can be standardised:** Algorithms inevitably carry a philosophy reflecting their creators and training data, so the question is not whether AI has values but whether they are visible. The aim should be to standardise not values themselves — which differ across cultures — but the transparency of values and decision-making, so societies can choose the systems they prefer.
- **Upstream design versus downstream governance:** There are limits of governing AI only after harm occurs. In every other sector, products are assessed for impact before release, raising the case for impact assessment before deployment.
- **Incentives and the market's blind spot:** The market does not currently value compassionate AI or voluntary standards, and that mandates are essential and can be supplemented with well-designed incentives (procurement conditions, funding alignment) to shift behaviour. Incentives must be aligned with goals the market will not solve on its own, and financial reward alone is not sustainable without a shared, well-defined vision of a desirable future state.
- **Inclusion, language and access:** Multilingualism and local-language data was highlighted as bridging the risk-focused and access-focused camps in AI governance: it reduces bias and extends access to those otherwise excluded at once.
- **Bringing stakeholders together:** Progress requires closing a long-standing gap between those who build technology and those who understand its human stakes — technologists historically dismissive of human-rights considerations, civil society often lacking technical fluency.
- **Compassion as a design input:** Compassion (understood as mindful problem-solving, contextual awareness, fairness, and moral reasoning) must be built into AI systems at the training stage, and this requires reframing the underlying narrative of technological

progress away from pure efficiency and profit maximisation toward solidarity and shared human values.

Key Outcomes of the session

- Compassion was advanced as a design principle for AI — embedded at the training stage so harm is prevented upstream rather than remedied after deployment — and engaged seriously across industry, multilateral and civil-society perspectives.
- The main constraint on responsible AI is now operational, not aspirational: the principles are broadly agreed; what is missing are practical, measurable ways to implement them.
- Governing AI only after deployment is widely seen as insufficient, with the emphasis shifting upstream — toward transparency in how systems are built and decide, rather than the harder task of standardising values across societies.
- Market forces alone will not produce compassionate AI; mandates, reinforced by well-designed incentives, are needed to change how systems are built.

Key Recommendations and Forward-Looking Actions

Embed compassion at the design stage. To build Compassionate AI, systems that recognise harm, understand context, and weigh human consequences before optimising for outcomes, compassion must be designed into how systems are trained and built, not added after release.

- SMGC offers its Satyarthi Compassion Quotient (SCQ) as the framework for doing this: a structured way to make deliberate choices at the training stage about whose realities the data represents, what harms are designed against, and what values are treated as non-negotiable. Embedding these upstream also makes them measurable as a defined standard a system can be held to before it is deployed.

Train AI teams to see what they build. Developers cannot design for realities they have never encountered. Standardised training should bring technologists into direct contact with the communities that bear the costs of these systems — especially children and vulnerable groups — so those realities shape design from the outset.

Move scrutiny upstream. Introduce impact assessment during design and before deployment, drawing on the pre-market scrutiny that exists in every other sector.

Standardise transparency and align incentives. Require disclosure of how systems are trained, what data they use, and how they decide — and pair enforceable requirements with incentives such as procurement conditions and funding alignment.

Build inclusion in by design. Treat multilingualism, local-language data, and the participation of affected and under-represented communities as core inputs rather than afterthoughts, advancing both risk reduction and equitable access, in line with WSIS's people-centred vision and its 2035 horizon.