

The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems

Key Information, Milestones, and FAQs about The Initiative

Introduction

This document is designed to provide an introduction to the work of The IEEE Global Initiative.

Please provide a brief description of the initiative, including key objectives, milestones, scope and partners.

The mission of the IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems mission is to ensure every stakeholder involved in the design and development of autonomous and intelligent systems is educated, trained, and empowered to prioritize ethical considerations so that these technologies are advanced for the benefit of humanity.

Key milestones include:

- The creation of a seminal document created by hundreds of global experts over the course of three years called, [Ethically Aligned Design \(EAD\): A Vision for Prioritizing Human Wellbeing with Autonomous and Intelligent Systems](#). The original version came out in 2016, received over 500 pages of feedback, and was released again in 2017(EAD/v2). The most recent version of EAD was released in 2019.
- The creation of over twelve standards working groups inspired by *Ethically Aligned Design (EAD)*. The [IEEE P7000 standards series](#) addresses key socio-technical issues identified by EAD in pragmatic and actionable ways to put principles into practice for Artificial Intelligence Systems (AIS).

What are the initiative's key success factors?

During the open consultation period for the first two versions of *Ethically Aligned Design* there was great interest in the work and we received over 500 pages of feedback. This garnered significant interest in participating in the IEEE work from various stakeholders from around the globe, resulting in over 500 additional members engaging in various activities within the IEEE Global Initiative community.

The second iteration of EAD was utilized by the OECD to create the OECD Principles on Artificial Intelligence. It also informed the work of the Future of Life Institute, the EU High Level Experts Group, and multiple companies, including IBM. In aggregate, all three versions of EAD have been mentioned in more than three dozen global policy documents, highly cited academic journals and articles, and in the media.

The work of the Global Initiative and the IEEE 7000 series of standards inspired the IEEE SA to lead the formation of the [Open Community for Ethics in Autonomous and Intelligent Systems](#)

(OCEANIS), a global forum that brings together organizations interested in the development and use of standards as a means to address ethical matters in autonomous and intelligent systems.

What has contributed to the success of the initiative most in the past, and what will support it in the future?

Participatory Design: The IEEE Standards Association cultivates collaboration via participatory models that help to mobilize people from diverse backgrounds and sectors from around the world. These models are embedded in the structure of the IEEE standards development process that is governed by a proven set of policies and procedures and guiding principles, including direct participation, openness, transparency, public comment and broad consensus. These same principles are embraced and adhered to by the Global Initiative and its participants.

Consensus: In the Global Initiative all viewpoints are considered and addressed with decisions made by a majority of participants and where no individual or organization has undue influence or power.

Multiple Discipline Focus: The Global Initiative's Executive Committee recognized the need to have multi-disciplinary participation so as to best identify and address the socio-technical and design oriented issues related to AIS and ethics. To that end, we invited philosophers, marketers, policy makers, designers, and multiple technical experts within the AIS realm (data scientists, machine learning experts) to partake in the development and promotion of *Ethically Aligned Design* and related work. What was evident from early on is that different verticals or disciplines have their own vocabularies and epistemologies, which may lead to a sort of bias and affect their worldview. Addressing this becomes critically important and this resulted in the [IEEE AIS Glossary](#).

Ethically Aligned Design's first edition received feedback from non-Western regions indicating that the document, while deeply helpful, may have been more Western in nature. To that end, we invited individuals and organizations from China, South Korea, Japan, Mexico, and multiple countries in other regions to participate. Among greater regional diversity overall, the [Classical Ethics in A/IS Chapter](#) in EAD1e, which incorporates ethical traditions from around the world, came directly from this feedback.

Recognition of the Socio-Technical: It is critical to remember that values and the output of a system are dependent upon the perspectives of the user. A Western roboticist, well intentioned and highly skilled in avoiding risk, may still not realize that featuring eyes in the head of an android that look directly into a user's face could be considered rude in some Eastern cultures. This is why ethically aligned, human-centric AI is focused on deeply understanding and honoring end users. This demands an understanding of cultural and societal issues and employing methodologies into one's work at the outset of design to recognize and address these issues.

A Focus on Design: While there are moral and ethical issues that must be addressed in AIS design, such as bias in data sets affecting algorithms, AIS and how these systems affect individuals and society at large are still widely untested. But this does not mean we should focus on fear as dystopian attitudes may cause more harm. This is why prioritizing ethically aligned design methodologies and examining the issues laid out in EAD and the IEEE 7000 standards series provides users with frameworks to examine their AIS in light of how it will be released to the world and to support the responsible design and use of AIS.

What does success look like for your initiative? What metrics do you use to measure success? What are your next steps towards it?

Our success continues to reside in our open community that provides diverse perspectives and information, the community's strong ability to address evolving issues and needs, and the Global Initiative's ability to initiate and support practical programs. This includes the creation of [The Ethics Certification Program for Autonomous and Intelligent Systems](#) (ECPAIS). We were one of the first organizations to address certification process needs in 2018, recognizing that organizations are ready to provide evidence for the due diligence they are doing to ensure the trustworthiness of the AIS products and services they build. We also continue to launch new committees like the EAD for Business Committee, whose report, [A Call to Action for Business Using AI](#), provides an AI Ethics Readiness Framework providing users with timely information and guidance for addressing how they can move toward being more deeply responsible and innovative in creating AIS.

Success also is measured by the feedback we receive from the interviews, keynotes, and panel discussions our members and leadership engaged since 2015. Members of the Global Initiative have been a part of over 500 global events in the past five years.

What impact has your initiative achieved so far?

IEEE was the first in the larger socio-technical arena and as an organization of our size to boldly declare the need to prioritize applied ethical decisions at the outset of design. We even adapted our [Code of Ethics](#) to consider these emerging matters (see clauses 1 and 2, explicitly mentioning "ethical design" and "intelligent systems").

As mentioned previously, organizations like the OECD, UNESCO, UNICEF, the EU High Level Experts Group, and policy makers from the European Commission, UAE, India, Australia, United States and Canada, among others, seek guidance from the IEEE Global Initiative and its members on how to instantiate the ethical principles and standards for AI they wish to build and implement and cite and use the outputs of the Global Initiative.

Due to the open and inclusive nature of IEEE, we have been able to highlight similar work as much as it has been the need to prioritize our own efforts. The processes IEEE SA has in place enable global, ongoing collaborations. In this regard, we are uniquely positioned to help advance ethical socio-technical agendas both with the private and the public sector, as a

convener of experts in this work space and through our processes and expertise, enabling us to advance principles into practical applications.

What are the main socio-economic, technical, operational, commercial and/or scientific challenges you have faced in developing or implementing your initiative?

A key challenge for any AIS-related work is in recognizing that it is human data driving many, if not most of the broadly used algorithmic systems. A foundational aspect to our work is in creating a mental model for sovereign data honoring the need for all people to be able to access and share their data in parity with how it is already tracked.

For human agency and identity to flourish people need to be able to identify and create their own algorithmic level terms and conditions. These will serve as key ways for governments and businesses, among others, to communicate with trust to all users. The [Personal Data and Agency](#) chapter of EAD provides a blueprint for how this can work, along with a number of our [IEEE P7000 Standards Working Groups](#), especially IEEE P7012, [Standard for Machine Readable Personal Privacy Terms](#). We are also honored to have recently started the IEEE P2890 Working Group, which is a [Recommended Practice for the Provenance of Indigenous Peoples' Data](#).

Are you proactively seeking cross regional and cross sectoral collaboration? If so, how? If not, why not?

As IEEE is a global organization, it is an ongoing priority for us to have as much regional and cross sectoral collaboration and participation as possible. This includes members from countries such as India, Australia, Brazil, Eastern Europe, China, Japan, The Russian Federation and Thailand, along with North America, Western Europe and the United Kingdom. We recently partnered with C-Minds based in Mexico City that is helping drive our efforts in Latin America. Our Committees feature members from sectors including academia, policy, and business.

Do you collaborate with relevant stakeholders in charge of implementing and monitoring Ethical or policy based practices relating to Artificial Intelligence System?

We collaborate with the OECD, UNICEF, UNESCO, the European Commission, the European Union High Level Expert Group, and agencies of the U.S. Government. This is in addition to the coordination with multiple Standards Development Organizations and global academic, engineering and AIS-driven bodies. Moreover, IEEE was the driving force behind the creation of the [OCEANIS](#) open forum, dedicated to addressing the impact of trust and ethical concerns through standards in ICT platforms, in particular in the domain of autonomous and intelligent systems.

Which types of AI principles is your initiative helping to implement?

A sample of some generally accepted AI Principles that Ethically Aligned Design supports include:

- Sustainable development
- well-being
- Human-centered values
- fairness
- Transparency and explainability
- Robustness, security and safety
- Accountability

Ethically Aligned Design, First Edition, covers the above issues in great depth in the various chapters. For example, the [Wellbeing chapter](#) directly addresses how AIS and economic indicators focused on sustainability and human flourishing should drive AIS design.

Many of the IEEE P7000 Working Groups also focus on these issues. For example, IEEE [P7001 is focused on Transparency](#). Transparency, Accountability and issues of Bias are the three areas of focus for the criteria developed for our certification program, [ECPAIS](#).

How does the IEEE Initiative help to implement principles like these?

The principles are largely implemented within the growing series of AIS-related standards that promote innovation, foster interoperability and honor human values; through certification programs, such as IEEE ECPAIS, developing metrics and processes towards the implementation of a certification methodology addressing transparency, accountability and algorithmic bias; and through [workshops](#) and [education](#) to empower technologists, professionals and users to practically address and implement ethical considerations when developing and using autonomous and intelligent products and services.

What are the lessons learned and good practices you would like to share with other people launching initiatives advancing the adoption of responsible AI?

Foremost is to be cross-disciplinary from the start: always be asking, “who should be in this room to make sure the AIS we create is as safe and responsible as possible?” For example, if you created an AIS toy with affective computing sensors, it would be vital to have psychologists and children’s experts on your team. Equally important is to recognize that you must include multiple global perspectives including ethical frameworks from around the world and from various stakeholder communities.

What tools, if any, did you use to implement AI principles and responsible AI?

There are multiple frameworks and checklists within *Ethically Aligned Design*, the IEEE 7000 standards series, and papers like *EAD for Business* as mentioned previously.

Could you briefly describe your team’s structure and background?

The Global Initiative Executive Committee is about 50% male / female and includes experts from academia (including from humanities sciences), engineering and business. We are continuously working to increase global diversity. We work to ensure that our committees and members consider gender and geo-location inclusion as well as multidisciplinary expertise as standing principles.

In the rest of our past and present committees we have over 2,000 volunteers from multiple global regions involved in our work.

Did you include perspectives from stakeholders outside the AI ecosystem, including those marginalized in the production and design of AI systems, and the sectors or groups of population that are supposed to benefit from your project? If so, how?

In addition to what is noted above, we are also now focusing our efforts on including indigenous and First Nations perspectives in updates of *Ethically Aligned Design* and our standards.

What is your theory of change or theory of impact?

You can see our [EAD Conceptual Framework Here](#).

You can download [Ethically Aligned Design, 1e here](#).

You can download [EAD for Business here](#).

Where are the website(s) focused on GIEAIS?

[The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems](#).

Secondary website: <https://ethicsinaction.ieee.org/>