

A Call to Action for Businesses Using AI

Ethically Aligned Design for Business



Introduction



Hi there. If you are reading this, you might be initiating your first Artificial Intelligence (AI) project. Maybe you are uneasy about what an AI-infused future means for humankind. You might be in the middle of an AI effort and need some guidance. No matter what your specific situation is, the need for AI ethics is something that should be on everyone's mind, and we've provided a solid starting point for anyone interested in AI.

In this paper, you will find a straightforward overview of the value and necessity of utilizing AI ethics for businesses, recommendations for creating a sustainable culture of AI ethics, the skills required to do this, and how to hire and staff for such an effort. We have also included an AI Ethics Readiness Framework to help you determine how prepared you are to introduce and grow AI ethics in your business.

This is the first paper in the *Ethically Aligned Design for Business* series created by [The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems](#) (A/IS). More than a dozen industry experts have provided insights in the following pages based on their frontline experiences as AI ethics thought leaders and practitioners in their organizations. By relating the recommendations within *Ethically Aligned Design* to their industry efforts, our

committee members provide, for inspiration and mutual learning, pragmatic applications of AI ethics inspired by IEEE's global community of academics, data scientists, engineers, and tech entrepreneurs.

Future papers from this committee will focus further on some of the themes and ideas introduced here such as skill building, review boards, and creating risk taxonomies. This paper is meant to establish the general perspective that will inform future efforts.

*We are also launching this paper as a Request for Input (RFI) so we can hear from you. We want your **questions, stories, and insights** as we move forward. Find details on how to submit your feedback on page 16.*

Overview

Ethical decision-making is not just another form of technical problem solving. A tech-centric focus that solely revolves around improving the capabilities of an autonomous or intelligent system doesn't sufficiently consider human needs and the long-term impact on the future of society. Simply put, if you are working in the world of AI, you are, in fact, working for the future of humanity, so you need to embed ethics practices across all teams responsible for these types of systems.

An ethical, human-centric AI must be designed and developed in a manner that is aligned with the values and ethical principles of the society or community it affects.

To be human-centered, businesses must first establish a culture of trust, transparency, and accountability internally in order to effectively spread these values externally. An understanding of end-user values in light of AI's use in the wild allows companies to establish deep trust and relevance for and with users. To be successful, the ability of a business to cross-pollinate this expertise among employees through every stage of design is critical. Yet, while AI ethics skills are increasing in demand for businesses, [these skills are not yet clearly defined](#).¹



It is easy for anyone to say, "I am not an ethicist, so it is not my responsibility," but the duty to implement AI ethically belongs to everybody who works with it. This is not solely about efficiency of business processes but about holding the human at the center of all design. Designers, developers, and other stakeholders should never work in a vacuum and must always stay in tune with users' needs and concerns for the entire lifespan of an AI-infused product or service.

The time for action is now. While the growing [lists of AI Principles](#)² provide general guidance for businesses, the specific opportunity for today's businesses is to pragmatically utilize AI ethics methodologies to improve the design and building of their AI systems.

In 2011, the United Nations endorsed the [Guiding Principles for Business and Human Rights](#), which defines the responsibilities that businesses and states have to protect the rights and liberties afforded to all individuals. Protection of human rights should be the first foundational principle³

that technology is built upon; however, it is not sufficient on its own. Additional guidelines or principles are needed to protect society from harms still worth addressing that are not necessarily a violation of human rights.

It might be tempting to delegate decisions around AI ethics to the academics, technologists, and policymakers releasing multiple AI principles around the world. Now that AI technologies are so ubiquitous, however, they require first-hand understanding from everyone involved in their design, development, and use.

¹ "AI & Global Governance: Human Rights and AI Ethics—Why Ethics Cannot be Replaced by the UDHR," United Nations University Centre for Policy Research.

² [An Updated Round Up of Ethical Principles of Robotics and AI](#), Alan Winfield.

³ Human Rights is the first General Principle of [Ethically Aligned Design, First Edition](#). See pages 19-20 for details.

The value and necessity of AI ethics



By adopting and practicing AI ethics skills, businesses can become more ethically aware, increase levels of transparency, mitigate biases within their AI, and instill responsibility and accountability in those who work with AI while developing AI technologies that will be of greater benefit to humanity.

By doing so, AI ethics also becomes a positive differentiator. It can help build trust with stakeholders, employees, and customers. It increases employee morale when members of the workforce know they work for an ethical company. [Most customers believe companies have a responsibility to positively impact society.](#)⁴

Additionally, [investors may begin to see AI ethics as a critical factor.](#)⁵ Larry Fink, CEO of global investment management firm BlackRock, stated in his 2019 letter to CEOs,⁶ “Purpose unifies management, employees, and communities. It drives ethical behavior and creates an essential check on actions that go against the best interests of stakeholders...Purpose is not a mere tagline or marketing campaign; it is a company’s fundamental reason for being—what it does every day to create value for its stakeholders. Purpose is not the sole pursuit of profits but the animating force for achieving them. Profits are in no way inconsistent with purpose—in fact, profits and purpose are inextricably linked.”

⁴ Ethical Leadership and Business,” Salesforce Research.

⁵ “Google and Microsoft Warn that AI May Do Dumb Things,” *Wired*.

⁶ https://www.blackrock.com/corporate/investor-relations/larry-fink-ceo-letter?cid=ppc:BlackRock_USWA.

⁷ Excerpt from the [Methods to Guide Ethical Research and Design Chapter](#) of EAD1e.



“Autonomous and intelligent systems (A/IS) research and design must be developed against the backdrop that technology is not neutral. A/IS embody values and biases that can influence important social processes like voting, policing, and banking. To ensure that A/IS benefit humanity, A/IS research and design must be underpinned by ethical and legal norms. These should be instantiated through values-based research and design methods. Such methods put human well-being at the core of A/IS development.”

—*Ethically Aligned Design*,
First Edition⁷

The financial potential of AI is not in question. A recent [Accenture report](#)⁸ predicts that by 2035 AI could double annual economic growth rates for the twelve developed economies covered in the report. New AI businesses joining an unregulated market and utilizing large sets of data will present a challenge to those companies that have not yet modernized.

Whether dealing with the principles of fairness, accountability, or explainability in relationship to AI, one point is clear — without being able to demonstrate these attributes to stakeholders and customers, an organization’s AI won’t inspire trust. According to PwC’s 22nd Global CEO Survey, 85% of CEOs admitted that AI will change their business practices in the next five years, and 84% said AI decisions must be explainable to be trusted.⁹ Forty one

percent of senior executives report to have abandoned an AI system altogether when an ethical issue had been raised.¹⁰

At the same time, the risks of ignoring ethical AI issues are growing. A 2018 Deloitte survey states that 32% of AI-aware executives ranked the ethical risks of AI as one of their top three AI-related concerns.¹¹ According to a July, 2019 report from Capgemini Research Institute, (surveying 1580 executives from large organizations across ten countries and over 4400 consumers across six countries), “when consumers’ AI interactions result in ethical issues, it threatens both reputation and the bottom line: 41% said they would complain in case an AI interaction resulted in ethical issues, 36% would demand an explanation, and 34% would stop interacting with the company.”¹²

“Organizations must create ethical systems and practices for the use of AI if they are to gain people’s trust. This is not just a compliance issue, but one that can create a significant benefit in terms of loyalty, endorsement, and engagement. To achieve this, organizations...must not only define a code of conduct based on their own values, but also implement it as an ‘ethics-by-design’ approach.”

—Anne-Laure Thieullent,
*AI and Analytics Group
Offer Leader at Capgemini*

It is not enough to simply utilize AI to transform your business. AI ethics and its practitioners are required to describe the benefits and potential harms of a company’s products in the algorithmic age. By identifying principles that frame a company’s core values, they can begin the values-based design analysis that benefits their customers.

32%

of AI-aware executives ranked the ethical risks of AI as one of their top three AI-related concerns.¹¹

⁸ Why Artificial Intelligence is The Future of Growth,” Mark Purdy and Paul Daugherty, Accenture, 2016.

⁹ “CEOs’ curbed confidence spells caution,” PwC’s 22nd Annual Global CEO Survey.

¹⁰ “Organizations must address ethics in AI to gain public’s trust and loyalty,” Capgemini, 2019.

¹¹ “State of AI in the Enterprise,” 2nd Edition, *Deloitte Insights*, 2018.

¹² “Why addressing ethical questions in AI will benefit organizations,” Capgemini Research Institute, July, 2019.

Creating a sustainable culture of AI ethics

Defining and embedding AI ethics skills is about behavioral change within a wide group of people in a company. It will take time for businesses to consider AI ethics as a core competency across roles instead of only considering it for issues of compliance or in response to negative PR.

We see the inculcation of a culture imbued with AI ethics happening in two waves. For the first wave, people are introduced to new concepts regarding AI design and ethics at scale. Employee skills to make this happen require people who can offer direct support and/or who can start up programs that focus largely on general educational and awareness aspects.



“The C-suite needs to actively drive and engage in the end-to-end integration of a responsible and ethically led strategy for the development of AI in order to balance the economic potential gains with the once-in-a-generation transformation it can make on business and society. One without the other represents fundamental reputational, operational, and financial risks.”

—Anand Rao, *PwC US Global AI Leader*, [Interview Of The Week: Anand Rao](#)¹³

The second wave involves helping people understand these new ideas in the context of where they are working in an organization. This involves localizing different areas of a business and identifying a point person or evangelist who can help determine how AI ethics work can progress until everyone is aligned. The goal is to have these specific employees strategically spread out first, then grow the message from there.

There may be employees that are already working on ethical AI practices in the absence of a wider strategy. It’s crucial to elevate and empower those voices when creating any programs or uniting efforts under a wider framework.

In terms of [values-sensitive design](#), both waves focus on emotional intelligence as a key skill. All employees need to be able to think about their users or customers holistically and not just through the lens of technology, offerings, or short-term outcomes.



“In order to develop a product that is ethically aligned, stakeholders’ feedback is crucial to design a system that takes ethical and social issues into account. To ensure representation of stakeholders, organizations should enact a planned and controlled set of activities to account for the interests of the full range of stakeholders or practitioners who will be working alongside A/IS and incorporating their insights to build upon, rather than circumvent or ignore, the social and practical wisdom of involved practitioners and other stakeholders.”

—*Ethically Aligned Design*,
First Edition¹⁴



“AI ETHICS:

All employees need to be able to think about their users or customers holistically”

¹³ Interview of the Week: Anand Rao by Jennifer L. Schenker, *The Innovator*. July 7, 2019.

¹⁴ Except from the [Methods to Guide Ethical Research and Design Chapter](#) of EAD1e.

Our recommendations for creating a culture of ethics

Create awareness of the need for AI ethics. This is a critical first step. Once this is accomplished, your colleagues and leaders can start to mirror and embed these practices into their teams' daily processes.

- Consider offering a webinar utilizing this paper with a Q&A session afterward with members of your organization.
- Give this paper to your senior leadership as well as key product or design leads and any AI, data science, or ML leads.
- Make sure to quote from this paper regarding the statistics and research denoting the need for AI ethics as a business differentiator and a way to establish greater trust with stakeholders and customers.
- Assess where your team is according to an AI Ethics Readiness framework (see page 13).

Identify, recruit, and train a core team.

- Identify a multidisciplinary and diverse team of individuals that may already be doing ethical AI work, and support them with additional resources and visibility within the organization. Identify key stakeholders and those who are natural evangelizers for specific teams. Product managers are key—make sure you know what gets them excited and how to speak their language.
- Include key decision makers and team leaders or you may unintentionally exclude them from the process.
- Make the value of AI ethics clear from the beginning and as a core function of each person's role. This will move more people from the awareness stage to the action phase for this type of work.

Modify existing design and development practices to include ethical AI considerations.

- Work with executives to identify ethical principles based on the company's existing core values. Identify critical inflection points to see which principles can make the biggest impact when identified and addressed. Some teams may start with a list of [commonly agreed-upon AI principles](#).
- Ensure stakeholder alignment on the potential risks and benefits of AI implementation at the beginning of every design phase. Create a taxonomy of risks in order to appropriately [prioritize and address](#) those you have identified.
- Look for the most relevant subject matter experts who can provide deep and nuanced recommendations for each new design opportunity.
- Make sure you understand your audience and can demonstrate value in that context. Underrepresented users (e.g., people of color, women, etc.) can be disproportionately affected when proactive prioritization of equitable AI is not built into the core strategy. Ensure that you create actionable goals that help avoid unintended bias and discrimination, and bring in subject matter experts from the affected communities as necessary.

Talk with the people you are building products for and with (participatory design) to avoid negative unintended consequences and to encourage trust and transparency in regard to any newly created product or system.

Document your work and progress. Recommendations include:

- Names of any advisory board members, champions, evangelists, or SMEs that can be contacted for counsel.
- Discussions with any internal or external corporate governance boards or ethical advisory councils.
- Product goals, audience, and intended uses.
- Team members, past and present.
- Key decisions made and milestones met.
- Model cards for models.
- Datasheets for data sets.
- Document non-goals (i.e., who should *not* use the product, and identifying use cases the product should not be used for).
- Intended and unintended consequences with mitigation plans.
- Risks categorized with taxonomy rating.
- Notes from ethical reviews/discussions.
- Customer feedback.
- Known incidents.

Build structures to embed Ethically Aligned Design/AI ethics across the organization:

- Launch communications campaigns to enhance visibility of AI ethics work, encourage dialogue around *Ethically Aligned Design* and this paper to share lessons learned.
- Develop and cascade ethical skill-building programs.
- Define the cultural norms and behaviors that might enhance or undermine AI ethics practices, and develop change management programs to transform culture and bridge the behavioral gaps.
- Align reward and recognition structures to incentivize ethical decision making and development practices.
- Work with executives to ensure ethics is part of their board, governance, and leadership agendas so they can unlock resources (funds, people, technology) dedicated to ethical AI systems.
- Create (or adapt) governance structures to minimize ethical failures.
- Define metrics and KPIs.
- Adapt agile and human-centered design processes to ensure ethics are embedded across the development, design, and delivery chain.
- Highlight and celebrate successes.

AI ethics skills and hiring



It may be easy to assume the skills required to apply ethically aligned design methodologies may belong solely to a small cadre of experts that can attend to discrete problems or opportunities. This is not the case. As noted in IBM’s paper, *Everyday Ethics for Artificial Intelligence*:

“Every person involved in the creation of AI at any step is accountable for considering the system’s impact on the world, as are the companies invested in its development and deployment. Human judgment plays a role throughout a seemingly objective system of logical decisions. It is humans who write algorithms, who define success or failure, who make decisions about the uses of systems and who may be affected by a system’s outcomes.”¹⁵

To guide the implementation of AI in its multiple manifestations, design in the algorithmic age requires systemic adoption in all parts of a business supported from the top down, bottom up, outside in, and inside out.

According to Gartner Research, by 2023 more than 75% of large organizations will hire specialists in AI behavior, forensics, privacy, and customer trust to reduce brand and reputation risk.¹⁷

¹⁵ “Everyday Ethics for Artificial Intelligence,” IBM (<https://www.ibm.com/watson/assets/duo/pdf/everydayethics.pdf>).

¹⁶ Excerpt from the [Methods to Guide Ethical Research and Design Chapter](#) of *EAD1e*.

¹⁷ “[Predicts 2019: Digital Ethics, Policy and Governance Are Key to Success With Artificial Intelligence](#),” Gartner Research, December 2018.



“Companies should create roles for senior-level marketers, engineers, and lawyers who can collectively and pragmatically implement ethically aligned design. There is also a need for more in-house ethicists, or positions that fulfill similar roles.”

—*Ethically Aligned Design*,
First Edition¹⁶

75%

of large organizations will hire specialists in AI behavior, forensics, privacy, and customer trust to reduce brand and reputation risk.¹⁷

Recognizing ethical risk beyond issues of compliance means understanding the nature of algorithmic bias or how AI may be interpreted by users based on their cultural or personal values. Utilizing a form of empathetic introspection, AI ethics experts may have skills ranging in areas from anthropology, psychology, human rights, policy, philosophy, and/or design. However, it's when the ethics experts have working knowledge of machine learning and the ability to communicate potential outputs of AI that they bring the greatest benefit to their colleagues and companies.

A core challenge to make this happen involves communication and education. While it may be easy to recommend that engineers should be provided

applied ethics training to identify issues like bias, it is just as critical for non-tech employees, including the C-suite, to understand the basics of machine learning to the point where they can provide contextual ethical recommendations that can be codified as code. The same goes for sales and marketing employees. Incentive structures must be in place to encourage and reward ethical behavior.

It is important to identify what specific skills are needed in designing ethically for AI, for individual hires as well as augmenting existing team members' skills across and throughout a company. The particular combination of desired skills will be different based on the culture and needs of the company itself.


SPECIFIC AI ETHICS SKILLS INCLUDE BUT ARE NOT LIMITED TO:

- Strong communication, facilitation, and negotiation skills
- Ability to implement and drive an AI ethics code of conduct, e.g., non-discrimination, human rights, diversity and inclusion
- Program management skills
- Understands the conditions and situations in which the AI system works and doesn't work
- Identifies the edge cases/out of scope cases where an AI system is out of its depth
- Uses vocabulary and models from philosophy, including ethical theories, logic, and aesthetics to modify and facilitate the development of ethical decision making models
- Uses vocabulary and models of risk assessment to explain ethical dilemmas
- Understands the flaws in the current state (in which data is used as the baseline)
- Understands the concepts of bias and fairness, accountability and remediability, transparency, interpretability, and explainability
- Understands the ways in which an AI system can perpetuate existing biases
- Evaluates issues of safety in human-AI interactions, cybersecurity safety, etc.
- Demonstrates emotional intelligence—in particular, self-awareness, introspection, and self-management skills to encourage values-identification and alignment
- Uses concepts and leading research (psychology, sociology, anthropology, etc.) to ensure ethical standards ultimately advance human well-being and societal flourishing
- Understands the value judgments that are implicit and explicit in a system
- Understands the impact of AI on individual privacy, control, dignity, sense of agency, human well-being, and human rights
- Demonstrates cross-functional influence and fosters relationships to engage Corporate Social Responsibility, Governance, Diversity & Inclusion, etc. functions in the development of ethical standards
- Understands the impact of AI on second and third-tier consequences on human-to-human behavior, environment, sustainable development goals, civil rights, democracy, societal well-being, automation, labor trends, and job loss

If you are not hiring separately for AI ethics roles, you should equip technologists with the necessary skills to be able to conduct AI audits, which cover ethical and regulatory considerations. Analogous to independent audits of a company's financial statements, AI audits assess AI models for things such as fitness for purpose, potential for bias (i.e., because of design or data issues whether algorithms used in the model are unintentionally excluding certain groups or favoring some other groups), and that AI systems are working as intended. Ethical red teams could also be created in which an internal group looks for ethical holes (i.e., ways to abuse your product or ways your product might be used in alarmingly negative but not malicious ways).

There is also a timing and training aspect to AI ethics skills that is relevant to consider. Currently, it is key to identify technologists with the depth to understand the under-the-hood workings of machine learning. This means identifying people who have a grasp of math, probability, and statistics, and then partnering them with non-technical stakeholders as needed.

While it can be easy to find specialists who are machine learning engineers or ethicists, the twain don't often seem to meet in any individual. You cannot easily find a machine learning engineer who grasps potential ethical concerns nor an ethics graduate who understands the process of how technology products are built or the nuances of AI techniques. We expect this to change in the next five to ten years in terms of people's training or dual-purpose experience.



AI ETHICS READINESS FRAMEWORK

	Lagging	Basic	Advanced	Leading
Internal training, support, and people resources	<ul style="list-style-type: none"> • Employees on their own to find appropriate AI ethics resources • May be encouraged but no official support • More focused on compliance 	<ul style="list-style-type: none"> • Workshops and certifications required for team members • Access to and support from evangelizers • Review board of experts 	<ul style="list-style-type: none"> • Advisory boards • Key stakeholders and evangelizers for each product/solution • Employees add to existing processes for their specific use cases 	<ul style="list-style-type: none"> • Ethical AI practices are embedded into decisions rather than inserted as modules • Infused into all roles and onboarding
Leadership buy-in	<ul style="list-style-type: none"> • Leadership recognizes but does not prioritize AI ethics 	<ul style="list-style-type: none"> • Completed introductory level of training • Compliance focused knowledge 	<ul style="list-style-type: none"> • Including AI ethics standards in new projects • AI ethics is included in collective agreements • Leadership is updated/aware of team efforts 	<ul style="list-style-type: none"> • Ethical practices and perspectives are infused into product strategy • Incentive strategy to reward ethical behavior and create consequences for unethical behavior • Leadership champions AI ethics efforts
Metrics and KPIs	<ul style="list-style-type: none"> • No clear qualitative metrics beyond generally speaking to any AI ethics principles 	<ul style="list-style-type: none"> • Basic qualitative metrics are defined (human rights assessments, societal well-being indicators) • Some metrics implemented in user research process 	<ul style="list-style-type: none"> • Metrics on user trust and understanding are further developed and maintained • Value is tied to health of product • Taxonomy of risks to prioritize mitigation • Differential privacy 	<ul style="list-style-type: none"> • Sprints and goals to reach healthy levels of trust based on research • Constant user feedback for adjustments • Tied to revenue • Common vulnerabilities and exposures (CVE)
Organizational impact	<ul style="list-style-type: none"> • Does not overtly change organization, siloed effects 	<ul style="list-style-type: none"> • Principles and accountability organization wide are tied to practices between teams 	<ul style="list-style-type: none"> • Teams working closely with each other to improve processes • Understanding and accountability from top to bottom of organization • Creating and embracing a corporate culture of dialogue, discussion, and constructive critique 	<ul style="list-style-type: none"> • Changes direction of products/relationship with users • Mindset change

Our committee members on the importance of AI ethics¹⁸

“Salesforce’s core values drive everything we do. Trust, customer success, innovation, and equality mean that developing AI our customers and society can trust is critical to our mission.”

—Kathy Baxter,
Salesforce

“We have an ethical foundation which we continuously work to strengthen and expand. In the past this has included adopting a progressive code of conduct, developing and adhering to global human rights principles, and assessing our supply chain to enable sustainable development including pursuing conflict-free minerals. As we become more data and AI-centric, we see the need to expand our principles in such a way to mitigate any future harm our hardware, software, or data science contributions could make in the world.”

— Anna Bethke,
Intel

“We see opportunities for AI and automation to be embedded into our day-to-day technology design and operational processes, so it’s vital that our people are aware of and understand AI ethics in the context of their work. This will enable us to proactively manage potential impacts on employees, customers, and society.”

—Amy Oding,
Vodafone

“We pursue AI ethics because we’re committed to building products for everyone, with everyone. From our Product Inclusion initiatives to our AI principles, we are committed to building responsibly, being accountable to people, and avoiding unfair bias.”

—Annie Jean-Baptiste,
Google

“We see a lack of awareness about the unintended consequences of AI among organizations, be they in the private sector or in the public sector, and we’re trying to raise awareness about the various issues involved. For example, people intuitively understand there are moral dilemmas involved when a self-driving car has to choose who gets hurt in an accident scenario. However, when it comes to not-so-vivid examples such as automated decisions by algorithms as part of a business process, we find that people put AI on a pedestal. This is why educating all stakeholders about what gives rise to algorithmic bias and what are the steps they can take to minimize unwanted behaviors of AI systems is part of our mission.”

—Kashyap Kompella,
rpa2ai Research

¹⁸ Quotes not attributed to an outside source come directly from Committee members in development of this paper. Quotes are based on member’s own opinions and do not formally reflect the positions of their organizations.

“Our industry has started to embrace and apply key technologies such as AI and robotics, which could potentially disrupt the whole energy value chain. We should be thoughtful about delivering this change in the right way, and avoiding unintended consequences, before the technology becomes widespread.”

—Bob Flint, *BP*

“We want to design healthy relationships with our users. The potential of AI is wrapped up in its longevity as a solution—meaning everything we design must address current and future needs for users. To truly understand those needs, we need an inclusive and ethical approach to the entire process. Globally, we are starting to see the repercussions that come when companies do not prioritize AI ethics in their solutions. We want to make sure that ethical practices are ingrained on our teams so they can then be embedded into the products themselves.”

—Milena Pribić,
IBM

“The LEGO Group is a values-driven company who puts our responsibilities toward kids and parents at the center of our decision making. We aim to apply the same rigor to our implementation of technical experiences/ solutions as we do to production, quality, and safety.”

—Sam Coates, *LEGO, Creative Play Lab*

“Fashion, like any other creative industry, is currently facing the challenge of adapting AI for the right level of usage, with the goal of still remaining a human centric and ethical business. AI-driven design, if adapted freely without any regulation and lack of transparency, can become a destructive tool for unfair competition and radically disrupt the market in a unique way. Ethical guidelines for both technology and business shall become, and remain probably for quite a long time, the only mechanism of regulating the levels of adaptation of AI across the creative industries.”

—Andrey Golub,
ELSE Corp

“All Tech Is Human is committed to making AI ethics more inclusive, multidisciplinary, and participatory. We have a tech problem and need a societal solution. We believe it is important for society to offer insight and guidance to inform and influence how technologies are developed and deployed. Together, we can all co-create a more thoughtful future toward technology.”

—Renee Farris, *All Tech Is Human*

Conclusion

The proliferation of A/IS presents a profoundly human moment. Collectively, we are standing in the nexus of history.

While it's always been essential to know your customer and their needs, the specific nuances of AI when interacting with people demands a higher level of awareness around things like bias, identity, emotion, and cultural relevance. It also means recognizing the very real possibility that, outside of anyone's positive intentions for what they build, an end-user's experience is not fully up to the designer—it is up to the end user. This is why ethically aligned design is so focused on the end-user and how they and their values can be a part of AI design from the beginning.

We want to encourage a more participatory, collaborative relationship with our users in order to build long-term, meaningful products and solutions. Our goal is to build trust with our customers in order to provide real value while honoring their culture and context.

With that said, it is not easy to suggest one-size-fits-all solutions or recommendations to achieve this. What works for a 20-person company may not work for a 200,000-person enterprise. What we have provided here is an introduction to the need, general concepts, and first steps to implementing what is really an evolution of design practices in the algorithmic age.

This paper is meant for generating awareness. While there are plenty of sources linked here, this is about starting the conversation. Future papers will be focused on specific and deeper ways for you to take action.

Establishing a culture of ethical practices is definitely a journey. It's not a "one and done" process. Ethics is a perpetually evolving practice, and we thank you for joining us on this journey.

Request for input

As we've already pointed out, this is our first paper in a series designed to help you out on every step of our collective journey. We want your **questions, stories, and insights** as we move forward. That's why we're launching this paper as a Request for Input (RFI) and growing our committee with other AI ethics practitioners, no matter what size your business or organization.

So shoot us an email and tell us:

- Your thoughts about the paper (and remember to make it actionable critique versus opinion alone).
- Your stories from the AI ethics trenches (and let us know if we can reprint in future articles).
- If you'd like to join our committee (and kindly send your bio and AI ethics experience).
- Send all your thoughts to aiopps@ieee.org with the subject line, "EAD for Business," and thanks for contributing.

About the Ethically Aligned Design (EAD) for Business Series



THE IEEE GLOBAL INITIATIVE'S COMMITMENT TO PRIORITIZE ETHICAL A/IS DESIGN

Ethically Aligned Design, First Edition (EAD1e) was created by hundreds of global experts from academia, industry, and policy. [The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems](#) (A/IS), the program that drove the iteration of *Ethically Aligned Design* since 2015, currently has more than 3000 members and has inspired the creation of fourteen Standards Working Groups, a certification program, and multiple education modules.

The *EAD for Business* committee represents the beginning of the next phase of work for The IEEE Global Initiative. This continuing series of work focused on highlighting *Ethically Aligned Design* focuses on specific topical areas such as health, artists, and parenting along with business related issues.

The Ethics Certification Program for Autonomous and Intelligent Systems (ECPAIS)

The goal of The Ethics Certification Program for Autonomous and Intelligent Systems (ECPAIS) is to create specifications for certification and marking processes that advance transparency, accountability and reduction in algorithmic bias in Autonomous and Intelligent Systems (A/IS).

More information: <https://standards.ieee.org/industry-connections/ecpais.html>

Other Projects Inspired by EAD Influencing the Pragmatic Adoption of A/IS Ethical Design:

P7000™ Approved Standardization Projects

The IEEE P7000 series of standards projects (inspired by the work of The IEEE Global Initiative on Ethics of

Autonomous and Intelligent Systems and their landmark document, *Ethically Aligned Design*) under development represents a unique addition to the collection of over 1900 global IEEE standards and projects. Whereas more traditional standards have a focus on technology interoperability, functionality, safety, and trade facilitation, the IEEE P7000 series addresses specific issues at the intersection of technological and ethical considerations. The IEEE P7000 series also represents the largest suite of projects of its kind on the planet, with over five hundred working group members from dozens of countries around the world.

Joining a Standards Working Group is an easy and pragmatic way to genuinely shape the future of the algorithmic age. All groups are free to join.

More information: <https://ethicsinaction.ieee.org/>

The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems

The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems (The IEEE Global Initiative) is a program of the IEEE Standards Association with the status of an Operating Unit of The Institute of Electrical and Electronics Engineers, Incorporated (IEEE), the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity with over 420,000 members in more than 160 countries.

More information: <https://ethicsinaction.ieee.org/>

Ethically Aligned Design—First Edition Glossary

The Glossary features more than two hundred pages of terms that help to define the context of A/IS ethics for multiple stakeholder groups, specifically: engineers, policy makers, philosophers, standards developers, and computational disciplines experts. It is currently in its second iteration and has also been informed by the IEEE P7000™ standards working groups.

Download the Glossary at:

https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e_glossary.pdf



EAD for Business Committee Members

Members:

- Kathy Baxter Salesforce, Architect, Ethical AI Practice
- Anna Bethke Intel, Head of AI for Social Good, Intel Artificial Intelligence Products Group
- Hermann Brand IEEE, European Standards Affairs Director
- Caroline Chubb Calderon HelloHumanity, Founder and CEO
- Raja Chatila IEEE, Chair, IEEE Global A/IS Ethics Initiative
- Sam Coates LEGO, Senior Innovation Director, LEGO Creative Play Lab
- Adam Cutler (Co-Chair) IBM, Distinguished Designer
- Renee Farris All Tech Is Human, Advisor
- Bob Flint BP, Technology Director, Digital Innovation Organization
- Andrey Golub ELSE Corp, co-Founder and CEO
- Daniela Hanauer PwC, Partner, Integrity-driven Governance Risk and Compliance Solutions
- John C. Havens IEEE, E.D., IEEE Global A/IS Ethics Initiative
- Annie Jean-Baptiste Google, Global Head of Product Inclusion, Research and Activation
- David Joa Bank of America, SVP, Senior Data Scientist
- Kashyap Kompella rpa2ai Research, CEO and Chief Analyst
- Karen McCabe Senior Director, Public Affairs and Marketing, IEEE SA
- Eva Muller PwC, Senior Associate
- Clara Neppel Senior Director, European Business Operations, IEEE
- Amy Oding Vodafone NZ, Technology Planning Manager
- Ben Olsen Microsoft, Principal PM, Responsible AI
- Milena Pribic (Co-Chair) IBM, Advisory Designer

Committee members represent themselves only. Participation is not an indication of the formal views of the organizations in which they work.

Visit us at standards.ieee.org and **connect** with us on:



Facebook: facebook.com/ieeesa



Twitter: [@ieeesa](https://twitter.com/ieeesa)



LinkedIn: linkedin.com/company/ieee-sa-ieee-standards-association



YouTube: youtube.com/ieeesa



Instagram: instagram.com/ieeesa



Beyond Standards blog: beyondstandards.ieee.org

Powered by

IEEE SA

