ETHICALINTELLIGENCE.CO

ISSUE NO.4

EQUATION

THE TECH ETHIC

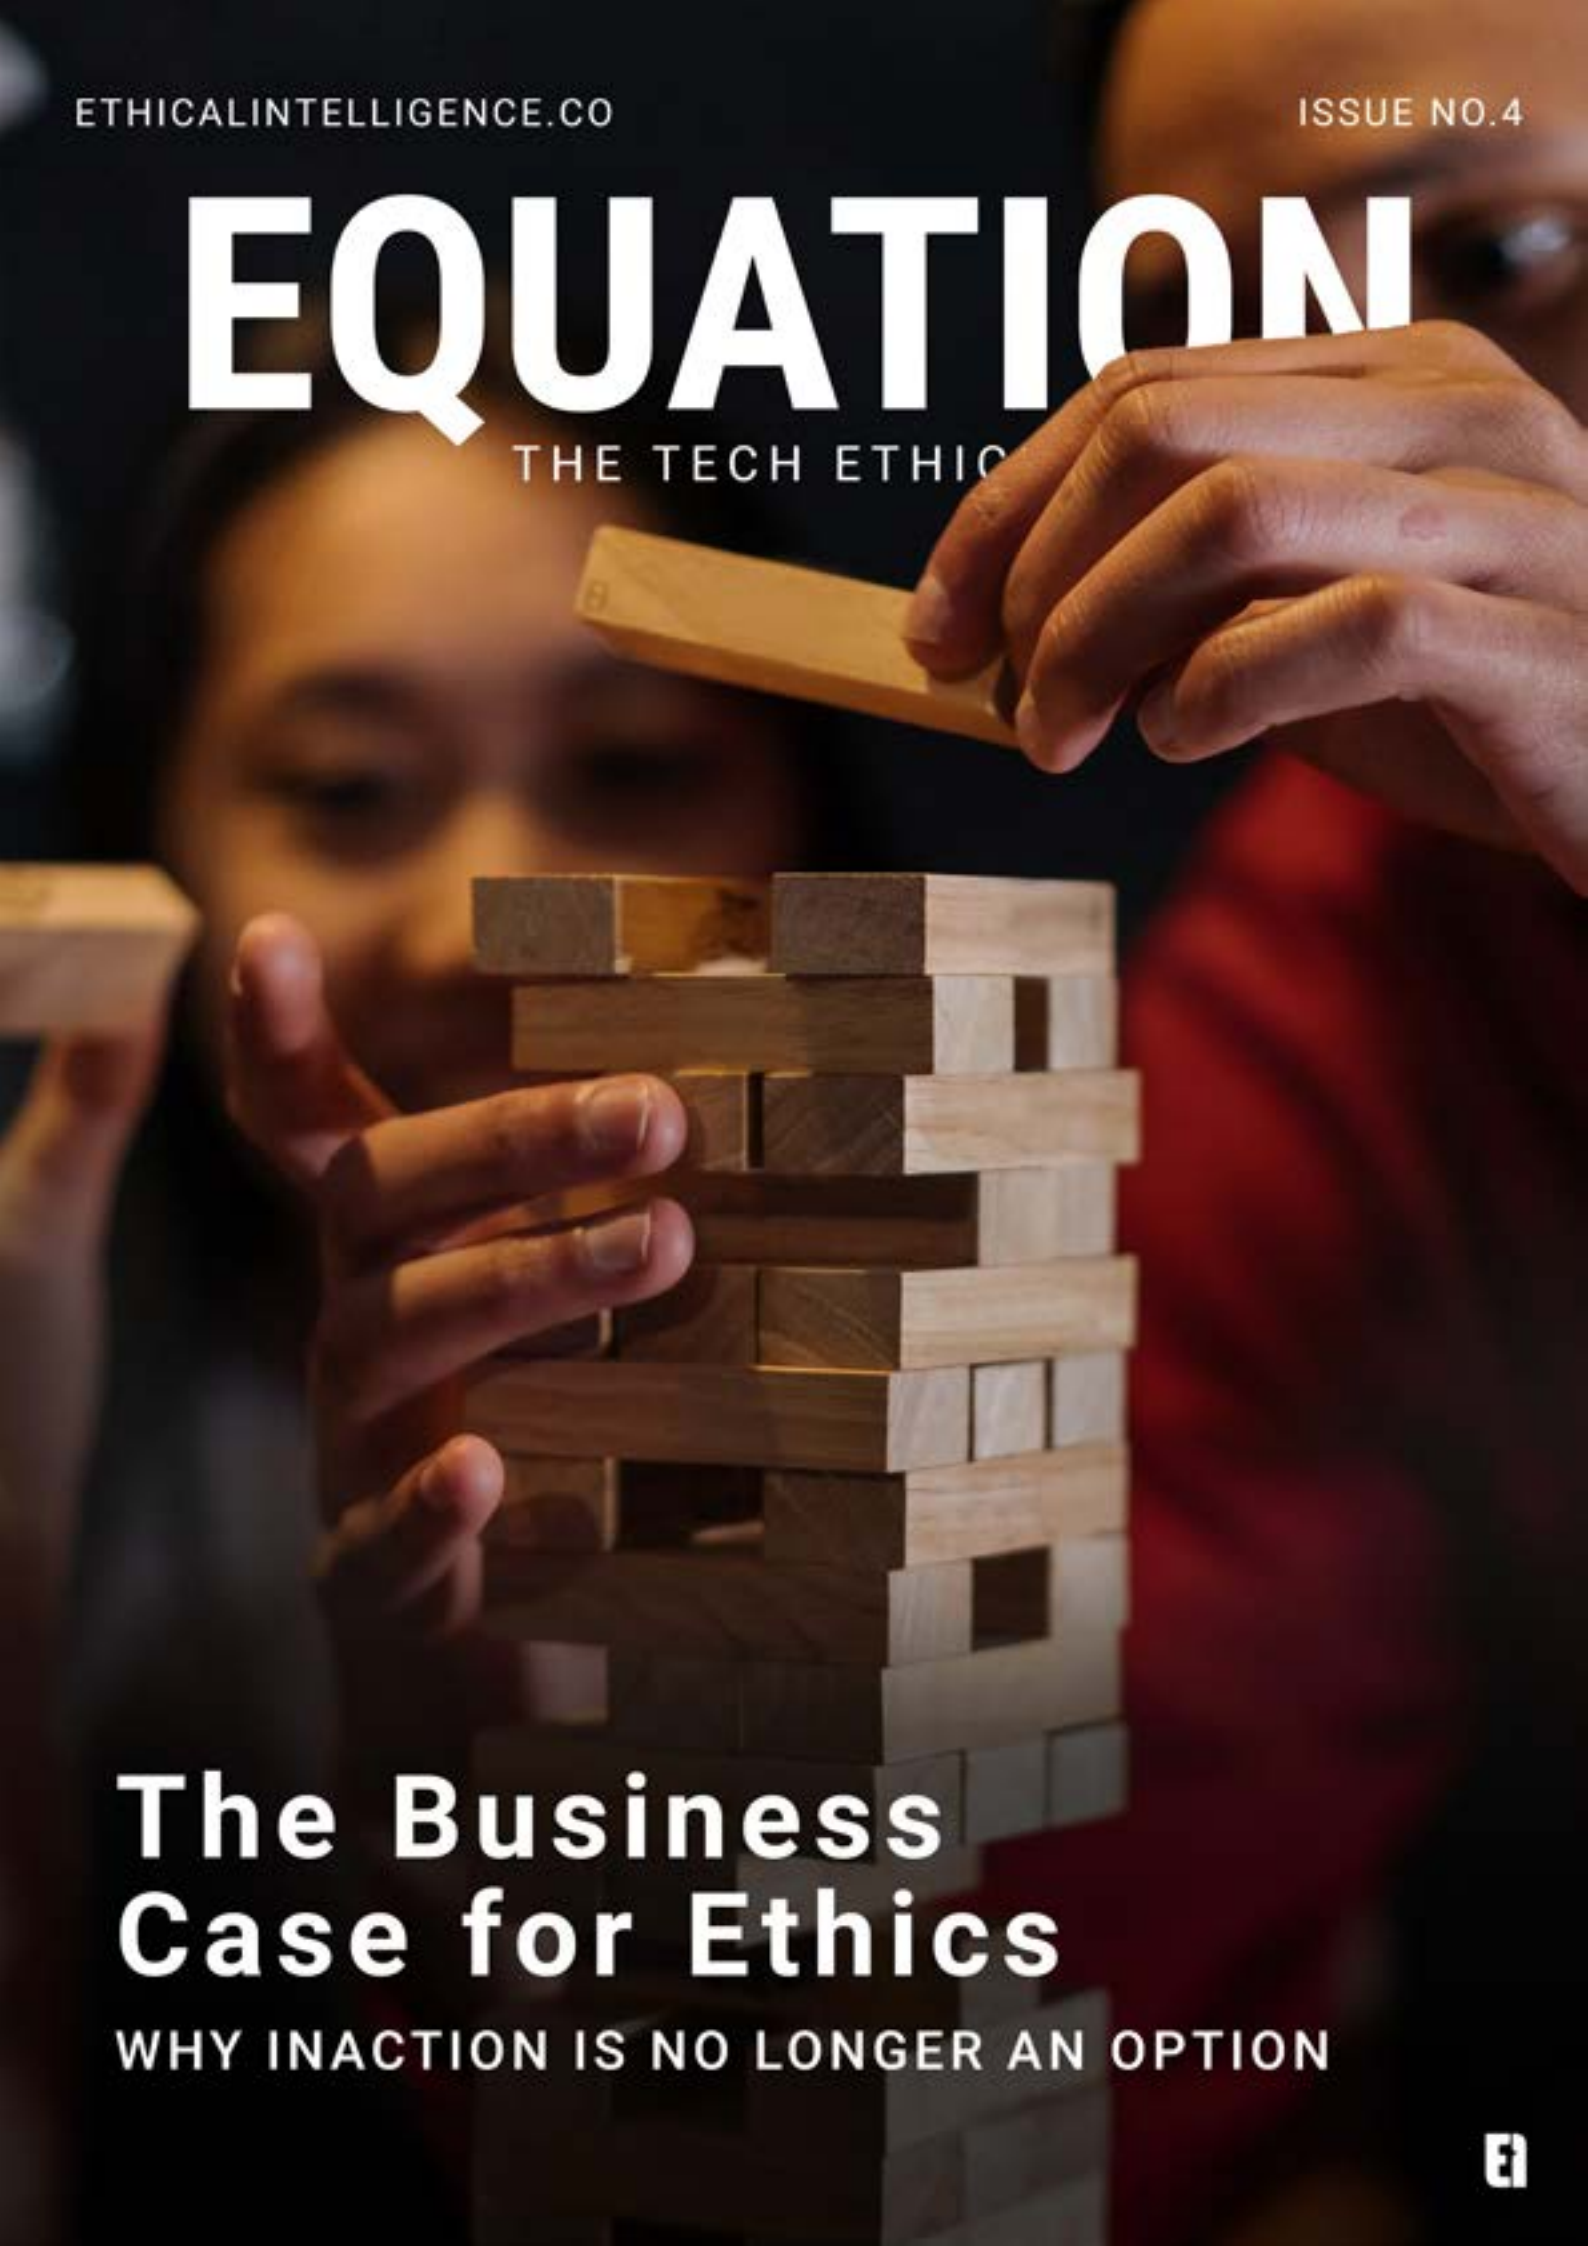

The Business Case for Ethics

WHY INACTION IS NO LONGER AN OPTION

TABLE OF CONTENTS

Letter from the Editor

Dear Reader,

There are a lot of lessons I've had to learn over the years as an AI Ethicist. Some have been hard, others quite unexpected, and a rare few could be classified as comical. However, one lesson in particular stood out. It was especially difficult to master, came with a hard pill to swallow, but has since proven to be one of the most important lessons I've learned to date.

What I had to learn early on was that if AI Ethics doesn't make sense in terms of business growth and development, nothing is ever going to change in the AI industry, no matter how well intentioned or moti- vated the people behind the technology are.

This realization would have been a great time to give up and count AI Ethics as just another lost cause. But instead of wallowing in dis- illusionment and frustration, I decided to set out and prove my hunch that ethics not only makes sense in terms of business growth and development, but is ultimately a source of strategic innovation in AI.

Turns out, that hunch was spot on.

Over the years, thanks to thousands of brilliant researchers, leaders, and responsible tech workers, we have an abundance of proof that AI Ethics is one of the most strategically beneficial tools for long-term success in AI. It can no longer be ignored, ethics is simply good business.

So now I'd like to offer you a new lesson to learn: instead of asking whether or not ethics is a good business decision, ask yourself ***just how much good business you will create by making the decision to embed ethics*****?**

Happy reading,

Responsible AI is More Than a Technical Problem:

Why an AI Ethics strategy is an essential part of the solution

Written by Simone Larsson & Olivia Gambelin

Artificial Intelligence is creating exponential opportunities for business and is predicted to reach $1394.30 billion by 2029 in market size. However, despite the promising developments, there lurk serious concerns on the horizon for the future of AI.

Gone are the days of moving fast and breaking things, as the countless scandals of technology abuse are not so subtly hinting at. As with any industry, our approach and best practices to AI must adapt and mature to align with market needs if we are ever to reach the full potential we have promised ourselves.

As we say good-bye to the honeymoon phase of AI, we must now usher in a new era that more accurately reflects the needs of the market. And never has it been more clear - the market demands Responsible & Ethical AI.

In a recent study by Salesforce it was found that 93% of consumers think companies have a responsibility to look beyond profit to positively impact society, while 79% of the workforce would consider leaving an employer that demonstrates poor ethics. In other words, the push for ethics is coming from both sides. On the one hand, consumers are voting for Ethical AI with their checkbooks, exposing companies without any Responsible AI (RAI) initiatives to extreme risks of losing a large percentage of business. On the other hand, employees are voting for Ethical AI with their time, again, exposing companies without RAI to extreme risks of losing top talent.

And it's not merely that customers, employees, and shareholders are expecting companies to take a public stance on Responsible and Ethical AI. They also expect companies to follow through with a considered approach, one that leads to outcomes that align to their company values and principles.

This means that the creators, architects, & product visionaries - both on an operational and a strategic level - are accountable for how these AI products are designed and built from an ethical and responsible perspective. Meeting accountability demands means incorporating ethics during the design and development process of AI powered products, rather than just at the operationalisation stage. We will delve into why this is important later on.

–

Barriers to developing RAI

Now, let's not pretend that the responsibility is negligible. For those in the driving seat, the challenges and complexity of embedding responsibility and ethics into product development and design can feel like an enormous mountain to climb. This may explain why 84% of executives state RAI as a top priority going into 2023 but only 25% have any kind of active program in place - and only 21% of those programs have reached maturity. Executives are likely to face a range of different challenges specific to their context, but there are four main challenges companies consistently face when it comes to implementing RAI:

1. Foreseeing the future implication of AI products 1, 3, 5 years from now,
2. Recognising and addressing ethical AI challenges specific to their organizational context and disciplines in AI they focus on,
3. Cutting through the noise and keeping track of impending and recently implemented AI regulatory mandates and the associated codes and requirements they need to adhere to, and lastly
4. Reducing or altogether eliminating the technical team's moral burden of navigating ethical questions and considerations, as often they lack the expertise and credentials to lead the discussion on ethics in a meaningful way.

Dealing with these challenges and the resulting complexities frequently creates noise, leads to inertia, and an unsettling disparity between intentions and implementation. However, these challenges are surmountable, this inertia can be overcome, and the gap between intention and action can be closed all with one simple step: seeking support from ethical experts to help unpack where and how to start. An AI Ethics strategy, the who, what, why, and how of proactively mitigating unintended consequences with AI is a good place to start.

–

What does an AI Ethics strategy look like

Coming from a technology strategy background, I can sense the silent sigh of 'yet another strategy', but hear me out on this one. RAI is not only a technical challenge. Instead, it includes operational and sociocultural aspects which, if not given proper attention, will either make or break the success of an RAI initiative. An AI Ethics strategy takes this into account, making it not only important, but also different in its approach. It's different because without it your AI Governance Strategy potentially only focuses on operational governance: MLOps, DevOps and DataOps. A one sided approach is not comprehensive and does not factor in ethical considerations in the upfront design and development.

An AI Ethics strategy allows companies to be deliberate about how and why their business engages with AI. It acts as a pathway to align a company's values with AI product development. AI Ethics strategies prioritize people first, then the process, and finally, the technology. AI Ethics strategies examine the purpose of an AI powered product, and the process of building it, asking: what is being built?, who is building the technology?, and how it is being built? Deliberating on such questions allows a company to better understand the intricacies of their product development process. An AI Ethics strategy also presents itself as a call to action for the management team and employees. The best approach is a top-down bottom-up approach whereby both leadership and delivery teams are consulted to ensure comprehensiveness and clarity on what to prioritize first.

Why we should bother with ethics

Now why care about any of this in the first place? Building AI is hard enough as it is, why add yet another layer to the mix?

Because, in these uncertain times you need to be doing everything you can to ensure the survival of your business, and Ethical AI has been proven time and again to yield immense benefits to those who embrace it.

- Findings from Blackrock and Morningstar show that the most ethical companies outperform Large Cap Index companies over five years by 14.4 pct., and over three years by 10.5 pct.

- Capgemini reports that 55% of surveyed consumers would purchase more products if a company's AI was perceived to be ethical, while 34% would stop interacting with a company altogether if its AI outcomes resulted in ethical issues.

- Finally, Accenture reports that, in Europe and North America, 73% of consumers are willing to share more personal information if brands are transparent about how it is being used. By sharing more data, users contribute towards higher quality training data, enabling more accurate and tailored outcomes, such as recommendations; and creating a virtuous cycle of positive outcomes for both the company and the end user.

Why we should bother with ethics is clear. It's much more than a 'feel good' exercise. Embedding ethics and having an Ethical AI strategy will not only relieve companies of the moral burden of navigating ethical considerations, it's good for business too.

–

Conclusion

I challenge you to be the champion of operationalising AI ethics within your organization. From the start, not as an afterthought when your AI product has already been operationalised. It's too late then. Eat the frog. Get the AI Experts in to contextualize Ethics and Responsibility for the industry, AI disciplines, and AI regulatory context you operate in. Allow them to advise you on your strategy and along the prioritization and implementation journey. Then sit back and reap the benefits and ROI.

Making the case for ethics

Talking points:

- Building Responsible and Ethical AI starts with strong company culture and clear governance processes to support the AI development lifecycle

- Creating an AI Ethics strategy addresses the operational and sociocultural aspects of Responsible AI, effectively increasing the adoption and impact of any RAI initiative

- Working with an ethicist to establish an AI Ethics strategy ensures effective and efficient use of time, effort, and budget

By the numbers:

- Ethical companies outperform Large Cap Index companies over five years by **14.4** pct. and over three years by **10.5** pct. (Blackrock and Morningstar)

- **55%** of consumers would purchase more products if a company's AI was perceived to be ethical, while **34%** would stop interacting with a company altogether if its AI outcomes resulted in ethical issues (Capgemini)

- **73%** of consumers are willing to share more personal information if brands are transparent about how it is being used (Accenture)

Growing Your Confidence in AI:

an evolving landscape and the role of advisors

Written by Patricia Shaw & Jacob Beswick

Leveraging AI within products and services across business functions has long been associated with a promissory note guaranteeing improved productivity, innovation, and ultimately bottom lines. This promissory note has been accompanied by concerns that AI's benefits are tied to risks. Regulation, standards, and tools are evolving and being designed to provide us with confidence that our AI products and services can indeed be "trustworthy". No one jurisdiction or sector has the monopoly but there are "first movers".

The evolving landscape

Regulation:
There is a lot going on in the policy, legislative and regulatory space in the EU, US, UK, & APAC. Legislators and regulators are on a journey, trying to find ways of making new and emerging technologies safe for the public. Not all legal requirements are homogenous or even synergistic, so it can be hard to find the overlap, and obligations can be open to interpretation, especially when current and proposed laws do not neatly fit into the context of your AI operations. This space can be really hard to navigate and can feel a little overwhelming. Whilst lawyers may get you so far in achieving legal "compliance", we appreciate that there is still often that nagging feeling, "have I missed something?"

Standards:
Standards are vital to the good and proper functioning of our technologies in general. Standards draw on best practice processes, methodologies and techniques to build in safety, security, reliability, robustness from the beginning, but there are now standards, certifications, and AI audit rules emerging that help businesses put not just principles but values like accountability, algorithmic bias, ethical privacy, fairness, transparency, sustainability and wellbeing into practice. In doing so, standards will be a vehicle for regulatory compliance once those requirements are formally established. Like regulation, standards are an important means to assuring that AI-based products, services, and systems are trustworthy.

Prospective regulation and standards have in common desired impacts, or outcomes that shape the world in a way that we collectively (as some baseline) wish to exist. Today, these impacts speak to *trustworthiness, responsible AI, safety, risk mitigation* (to individuals, groups, and society).

With the desired impacts of regulation and standards clear, we do not yet know the prescribed course of action to achieve those impacts. Given this, how is any AI business (or business utilising AI as just one small part of its business offering or back office function) expected to navigate this landscape in a way which is not only empowering them to be "regulation ready" but helps them to be prepared in the face of fears that AI (when it is put into service or launched on the market) won't result in harm and liability, be a PR headache and lose trust, and/or worse still, not adopted at all?

–

Taking first steps

While we know the desired impacts of regulation and standards, we do not yet know the practices or means by which those impacts can be realised. And so before regulation and standards are formalised, organisations can take steps to prepare for the anticipated requirements and the related change management necessary.

Some helpful questions to ask and answer:

- What are your organisational priorities (e.g. improving revenue; fairness, transparency, accountability; etc.) and what role does AI play in relation to them?

- Is the way your organisation builds, buys, and leverages AI living up to your priorities?

Depending on your organisation's maturity, asking and answering the above may prove challenging in and of itself. One route is to work with Responsible AI or AI Ethics Advisors. Generally, such advisors can assist organisations to clarify their priorities with respect to AI ethics and principles, risk management, and AI governance; contextualise these priorities in the wider world of regulation and standards; and ultimately support organisations with implementation.

Afterall, you are building for your organisation a positive ecosystem of care. If all that can be encapsulated in one word: TRUST (Transparency, Responsibility, Understanding, Safeguarding, and Truth). And why is responsible AI so important? Because it is about producing FAIR (Fairness, Auditable where inputs and process can be traceable, Insight which is proportionate and necessary, and Reliable) outcomes.

–

Practical considerations

You may be asking yourself now, where do I begin? Here are some hints and tips for starters:

First things first

If you are going to implement a Responsible AI approach within your organisation, you are going to need to get C-suite buy-in.

Any organisation-wide successful roll out of Responsible AI requires both top-down and bottom-up engagement. The top-down effort helps to align a Responsible AI approach with overarching strategic priorities and a bottom-up approach helps with understanding current and enforcing future practices. These might not happen at the same time; indeed, C-suite may launch this as an executive initiative at first, bringing in teams over time. Otherwise, data science teams may see the value in this and generate their own approaches, making efforts to widen the impact of their work over time. As you can imagine, the latter approach can have many challenges.

Finding your principled approach
To get started there's a real need to understand what principles or values your organisation should prioritise. Priorities can be informed by the current values of the business, by existing regulations, or by the adoption of external frameworks. Part of making your organisation's principled approach stick is making sure it is relevant. Having 'fairness' as a principle for a company that leverages AI in manufacturing steering columns for cars may not make sense, although 'robustness' may resonate more meaningfully. Spend time profiling your AI systems to understand the depth and breadth of the values (and virtues) that apply to you and your stakeholders.

Setting firm foundations
Design a Responsible AI approach that suits your business, values, and culture. Understanding where you are on your Responsible AI journey (whether you're a newcomer, novice, maturing or a seasoned veteran) will help your business gauge what practices, processes/procedures, and policies are needed for now and to be prioritised, and what will be needed over a Responsible AI Roadmap as your business gets regulation, standards or even certification ready.

This is where a Responsible AI advisor can really help. In setting firm foundations (in general), you need to understand:

- Your AI systems
- Your data flows
- Your people and culture
- Your value chain
- Your user journey
- Your risks based on your intended and unintended/undesirable outcomes

Mind the gaps
For more established businesses, where there may already be some AI and/or data governance in place, there may be some need for a gap analysis so as not to reinvent the wheel, to find meaningful ways of adapting good practices or ditching poor practices and behaviours.

From tolerable beginnings to mature endings
Responsible AI governance cannot be considered as a "one-stop-shop" or as a "set-it-and-forget-it" approach. It likely won't start perfectly and will require maturation and refinement. An organisation's Responsible AI approach does not (and cannot) operate in a vacuum and will be informed by new products and services developed; new and evolving standards; and new customer or wider stakeholder expectations.

As Responsible AI advisors, we are here to help your business:
- Design and build AI governance strategy
- Build competence, capability and capacity at an executive and operational level
- Recognise co-existing frameworks and regimes for data, platform services, and AI
- Navigate the hazards of data, model and system/service overlaps and dependencies
- Build in dynamic algorithmic risk and impact assessment into the process
- Draw on your most powerful asset: people
- Innovate and create realisable positive societal / environmental impact
- Provide expertise for ongoing AI and data specific governance functions such as AI Ethics Advisory Boards, Data Ethics Committees, etc.

Having the assurance that an independent expert advisor has assisted you on your Responsible AI journey can provide that important support that builds confidence across your organisation. In a world where AI is being leveraged in all manner of use cases, the benefits of expert advisory extend from executives, to data science and business teams, teams covering risk, compliance and legal, to those covering HR, sales, and contracts and procurement.

Preparedness is empowering. Responsible AI is about preparing for your business' future in a way which is ethically aligned and societally impactful.

Making the case for ethics

Talking points:

- Compliance and ethics are two separate things; just because you are compliant does not mean you are ethical
- Lawyers ensure your product is compliant with regulation; ethicists ensure your product meets industry standards
- A holistic approach to AI preparedness includes both lawyers and ethicists working together, ideally on an ethics board, to provide the necessary guidance and advice to companies

By the numbers:

- **193** countries have adopted UNESCO's global agreement on the Ethics of Artificial Intelligence (United Nations)
- There are over **800** AI policy initiatives from **69** countries worldwide (OECD)
- The EU AI Act (the first-ever legal framework on AI) is estimated to come into effect later half of **2024** with as wide of an impact as the GDPR in 2018 (European Union)

EInsight
Transforming tech scandals for the better, subscribe for a monthly dose of ethics in action.

Subscribe here.

EImpact
Listen in as ethicists and responsible technologists tell their success stories.

Listen here.

EI Blog
Hear straight from the experts on developing topics in AI, Ethics, and Responsible Tech.

Read here.

Supporting Tech Teams in Ethical AI Development

Written by Divyansh Agarwal & Anna Danés

Tech teams in organizations around the globe are increasingly incorporating AI based tools and methods into their systems. This has resulted in the proliferation of human-AI interaction in the real world, which is projected to grow exponentially as new AI tools are developed and applied in myriad ways. Ethical development of software engineering systems has been around for a while, from the perspective of ethical data usage, privacy, security and the well-being of humans using these systems.

However, the increase of human-AI interaction has introduced an important turn towards the ethical development of AI, which poses scenarios previously unexplored, and challenging to address.

The inherent black-box nature of AI methods introduces challenges for ethical experts, as well as for researchers who study different ethical principles such as bias, fairness, transparency etc., in AI systems. Keeping the pace of innovation in perspective, more tech teams are realizing the need to wrap their heads around concepts in ethical AI. Tech teams are realizing that AI applications may have ethical blind spots if the development process doesn't involve individuals who are skilled in identifying them. Turning to an ethics board, or using Ethics as a Service (EaaS) is an effective strategy many organizations are opting for. An ethics board essentially guides the digital ethics transformation of the organization, working closely with tech teams that develop and build systems. The ethics board brings in experts from different backgrounds, who can oversee the development of AI applications and may be involved at various stages from the design to the deployment lifecycle.

However, this begs the question - How can tech teams working on building AI systems feel more supported in the context of developing within ethical boundaries? It is equally important to understand how the ethical principles themselves are changing with the complexity of human-AI applications. The needs of different organizations and their tech teams vary, so how can the ethics board identify ways to be flexible enough for supporting different use cases?

–

What does the AI development process look like?

The AI development process at organizations varies depending on the different stakeholders and objectives in the equation. Tech teams that build and deploy AI systems operate within varying metrics such as expected outcome, intended users, available resources and time restrictions, which in turn determines the ethical support they need. These factors may be more clearly defined for bigger organizations with dedicated AI teams, as they can draw on their own archive of experience in developing and deploying AI systems. Such organizations may also incur a high cost associated with fixing and redeploying the system in the event of an ethical oversight. Hence, they might prefer to employ an in-house team of ethical AI experts that work together with their tech teams early on in the process.

In contrast, smaller organizations may still be using AI in an experimental manner, to probe the overall efficacy it brings to their software systems, and possibly for specific clients/use-cases. Their tech teams may be transitioning from traditional software systems to developing AI systems, and thus their needs would be constantly evolving. This evolving nature, however, isn't necessarily a set back. It may allow for the flexibility tech teams require for ethical AI reviews and involving experts in a more ad-hoc nature. It provides an opportunity to use Ethics as a Service (EaaS), where the ethics board can jump in and provide the necessary support the tech team might require.

Depending on these scenarios, tech teams themselves might feel the need to be supported

by an ethics board in different, more specific ways. The composition of the ethics board itself should be organization-specific, to support the varying nature of these tech teams, which ties back to the idea of incorporating EaaS. The recent trend towards democratization of AI has opened the door for smaller organizations to incorporate AI in their software systems, for satisfying specific business needs or customer requirements - with smaller AI teams that are still in the incubation phase. Such an AI tech team may initially be interested in developing around a specific ethical principle. For instance, they might ensure that their AI models are trained on sample data diverse enough to capture the representative users, and therefore would require ethicists and experts who are experienced in the model fairness domain. For AI teams involving researchers who publish and present their system/findings, the involvement of an ethical expert may be required more end to end.

In most AI projects, ethical experts can (and should) be involved simultaneously with the development process, with an ethics board acting as a sounding board at different stages.

As tech organizations make this transition to developing AI systems, it's important to understand the specific ways in which the development process and the needs of tech teams have evolved.

–

What are the needs of traditional tech teams & how do they change compared to teams that develop AI systems?

Tech teams developing human-usable software applications and products have traditionally used (and refined) agile engineering practices over time. This software development methodology takes into consideration the need for flexibility, in an evolving landscape. An effective development process is also required to be human-centered, taking into account ethical guidelines in decision making for managing user privacy, security and well-being in interacting with the software application. With the advent of AI-powered software applications, the development process, as well as the ethical outlook has changed dramatically.

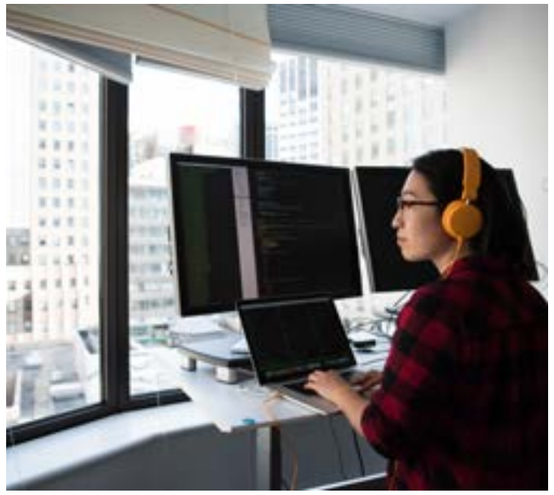

Owing to the fundamental differences in the complexities of AI-based software applications, the needs of the tech teams have evolved substantially. Large AI models that power applications can be inherently non-deterministic due to their black-box nature. This has brought greater emphasis on the volume, quality and the overall management of the data used in building AI models, including sourcing, cleaning, tuning, dissemination etc., more so with the recent push towards data-centric AI. AI-based software applications also require a different kind of post-deployment oversight, making it essential for tech teams to constantly evaluate, update and redeploy their models. This even creates new challenges for the security and the robustness of AI models behind software applications. These advancements have nudged tech teams to re-evaluate their needs, and build expertise in acquiring deep knowledge in the field of AI. Big tech organizations like Microsoft are incorporating AI workflows as an essential part of their software development process. Simultaneously, this has induced a progression in the field of ethical AI, with principles of fairness, access, transparency and security being constantly updated in the context of AI systems.

While tech teams may find it challenging to align themselves with the evolving ethical guidelines, ethical AI experts also need to adapt to the different ways in which AI technologies affect the humans using them, while keeping up with the pace of innovation. Consequently, organizations have also begun to realize the pertinent need for experience in ethical development of AI systems amongst their ranks.

Practical takeaways on the effectivenes of an ethics board

In recent past, the media has highlighted many instances of ethical oversight in big tech organizations. However, most tech teams nowadays are aware of the ethical implications of their work. Team members are aware of privacy issues with their own personal use of technology and have a rich ethical landscape, even if some still deny any accountability. But knowing that there might be issues in the work that you do is not the same as identifying them, understanding those issues and knowing how to fix them. This is where the support of an Ethics Board can be helpful.

When the Ethical Intelligence team first started working with Anyone, they found a very ethically mature C-level and dev team. The entire company was aligned with the idea of creating a solution that not only did no harm but would also do good to the society. In such a mature environment the Ethics Board takes the issues that they have raised and helps them get perspective and identify new ones. Due to the nature of the project, the Ethics Board was composed of individuals with a strong development background, a philosophical background and expertise in bringing a business perspective to the team. This “balanced approach” to composing the team allowed for a level communication between the “ethical advisor” and the development team that increased blind spot detection. Having an engineer in the Ethics Board allowed for smooth communication and the depth that the project required.

The engineering team felt understood, supported and guided towards finding the right answers.

The team could focus on their work, knowing that there are colleagues on the Ethics Board looking after the ethical aspects of their work. This allowed them to pick up speed on the project and meet their deadlines without the fear of having to go to the drawing board again after deploying their work. This was not only the most efficient process to develop the system, but the team members felt supported and empowered in dealing with ethical dilemmas. This clearly increased the connection with the company and their motivation, leading to a decrease in attrition.

Ethics boards are an easy piece to integrate in small and large companies. In a start-up or medium size company, ethical experts will seamlessly connect to the affected teams and provide solutions. For large companies, the ethics board acts as an integral piece in the creation of successful products and services. Most organizations that understand their social responsibility invest in an ethics board.

–

Conclusion

It is important to consider if an ethics board can serve AI and data-centric tech teams in developing ethical AI systems. AI engineering teams have specific needs and operate in different ways depending on market needs and technological developments. The onus largely lies on organizations to identify the flavor of ethical expertise that would be suitable for their tech teams.

A conscious organization would make sure that teams have the support of ethical experts to navigate the nuances in AI software development and accomplish their vision.

Making the case for ethics

Talking points:

- Tech teams face a diverse set of ethical challenges and need a variety of ethical solutions tailored to fit their specific context
- Having ethical AI expertise on-hand helps technical teams responsibly keep pace with innovation, feel supported in making important ethical decisions, and avoid ethical fatigue
- Ethics boards should be designed to adapt to the changing horizontal and vertical needs of a tech team

By the numbers:

- **90%** of organizations are aware of at least one instance where an AI system had resulted in ethical issues for business
- **60%** of organizations have attracted legal scrutiny because of poor ethical decisions made by an AI system
- **22%** of organizations have faced customer backlash as a result of AI-enabled decisions

(Enterprisers Project)

StartUp Awards Scotland

2022 Winner

EI ETHICS BOARD

Create, connect & manage your customized ethics board all in one place

Access your EI Ethics Board anytime, anywhere. Designed with simplicity in mind, your EI portal powers all communication, scheduling, strategy building, and resource development with your EI Ethics Board.

MATCH EXPERTISE

Match expertise to your needs with the selection of your personal EI Experts. Each board is composed of a technologist, ethicist, and industry expert specific to your company

MATCH SCHEDULES

Customize your board by setting frequency of meetings that best compliments your team workflow, and metrics of success that support your company objectives

MATCH METRICS

Select your metrics of success and improve on your ethical decision-making by building out your database of decisions and ethics resources, all hosted securely in platform

For more information, email **info@ethicalintelligence.co**

Missing the Forest for the Trees:

A More Balanced Approach to AI Ethics

Written by Geoffrey M Schaefer

The field of AI Ethics is in a paradoxical state. On the one hand, it has never been healthier. It is one of the fastest - if not the fastest - growing subfields of AI. There has been a profusion of ethical guidance in the form of frameworks, principles, and toolkits from organizations big and small, public and private. The issues of bias, fairness, transparency, and other focal points have been cemented as domains in their own right.

And it is now considered impolitic to establish an AI practice without an accompanying body of work in Responsible AI.

But on the other hand, this work has resulted in a kind-of developmental stasis. Despite noble intentions, the field has committed to an overly narrow focus on the risks and harms of AI - treating questions of risk mitigation and avoidance of harm as paramount - with too little attention being paid to how such systems can produce better societal outcomes. In fact, the current state of AI Ethics would be entirely foreign to the classical philosophers, like Aristotle, who focused on a different question: What does it mean to live a good life? Aristotle's conception of this topic is termed Eudaimonia. This crucial, albeit missing, half of the equation has turned AI Ethics into a discipline of compliance, not progress. A discipline concerned with the sins of commission, not omission. This is an unbalanced moral calculus - and the opportunity cost to society is huge.

–

A Brief History of Eudaimonia

Aristotle's principal focus was on the deceptively complex concept of *Eudaimonia*, a sort of fusion of "happiness" and "flourishing," and his measure of a life well lived. Achieving eudaimonia requires a lifetime pursuit of one's goals with virtue and excellence. It's a lifetime pursuit precisely because one doesn't know how these goals (e.g., rearing kids, building a career, etc.) will turn out until the twilight of their years. Life contains many twists and turns, peaks and troughs; but, according to Aristotle, it's the net result that's important, not one's degree of happiness at any discrete moment. A successful life can be filled with hardship but still work out according to one's goals in the end.

The twin pillars of Eudaimonia, virtue and excellence, are complex in themselves. For example, excellence is not always synonymous with virtue; something can be excellent without being virtuous. An excellent soldier may act in contravention to human rights and other accepted norms of warfare. Excellence must have its own moral character, then, to fulfill its role in our collective human flourishing. Aristotle also stressed that a virtuous person is not someone who simply acts in a virtuous way. Rather, they must understand what the right thing is, actually do it, and do it for the right reason. In other words, a virtuous person will do something precisely because it's the right thing to do, and will feel no temptation to do otherwise. As a contemporary example, imagine an auditor who chooses to conduct a secondary analysis of an account that's showing red flags while being pressured to finish the assessment on schedule, no matter the results. While this may be a hopelessly high bar, it demonstrates the challenge of evaluating ethical actions and has important implications for modern AI Ethics.

–

Rebalancing the Equation

Let's shift our focus back to AI Ethics. As we discussed, the most prominent questions in AI Ethics are focused on mitigating risk and avoiding harm. Examples include: ensuring the privacy of users; reducing bias in training data; and maximizing the explainability of a system's outputs. But if we shift our attention from questions of harm, to concerns of human flourishing, we completely change our ethical calculus. This also constitutes a more uplifting and actionable approach to designing and deploying AI systems.

This is not to say that issues of bias, fairness, transparency, and other concerns are not without merit. To be sure, these are important considerations in the design and use of any AI system. But they are merely half of the equation. Introducing the concept of Eudaimonia into our assessment leads us to ask a different set of questions, resulting in a more balanced assessment of an AI system's ethical dimensions. It also enhances our capacity to build AI systems that produce or facilitate good outcomes in themselves. Let's demonstrate this approach by looking at two case studies.

Healthcare

Some of the most promising applications of AI are centered in the healthcare space, including precision medicine and early-stage diagnostics. Yet, two of the most common AI Ethics questions in this space concern *How we can protect patient privacy?* and, *How can we ensure our training data isn't biased?* While both of these questions are important, neither of them address one of the principal goals of healthcare: reducing human suffering. To better align with this goal, what if we asked questions such as *How can we better treat rare diseases?* or, *How can we prevent a patient from getting sick in the first place?* These questions drive much of the science behind our most promising medical advances, but they rarely factor into our ethical calculus of AI's role in these advancements. This is not to say that the previous questions are unimportant. Indeed, it is the combination of these questions that should guide our analysis. Instead, the field of AI Ethics is missing the forest for the trees.

DeepMind's AlphaFold system is a perfect example. This algorithm has been widely touted as having cracked biology's 50-year challenge of protein folding. Proteins are the building blocks of life, "underlying every biological pro- cess in every living thing." But deciphering their structure – a process called "folding" – is no- toriously difficult, taking months or years to achieve, if at all. Understanding protein folding is essential for everything from developing medicine to understanding antibiotic resistance. And while there are hundreds of millions of proteins known to science, AlphaFold has already folded over 200 million of them. Eric Topol, Founder and Director of the Scripps Research Translational Institute, said that with AlphaFold mapping "nearly the entire protein universe, we can expect more biological mysteries to be solved each day."

AlphaFold has rightly garnered the attention and admiration of the AI community and beyond. But we don't celebrate AlphaFold as an *ethical* algorithm. Why not? Given the criticality of protein folding to modern medicine, it stands to become one of the biggest drivers of human flourishing in history. And it will do so in manifold ways, from the elimination of debilitating disease to the strengthening of our immune systems. A more holistic approach to AI Ethics, then, would encourage the development of similar types of AI systems across a range of critical sectors and applications that impact human lives. In the case of AlphaFold, specifically, we might encourage the system's generalizability to ensure it impacts the broadest cross section of society possible.

The Justice System

The justice system is replete with controversial applications of AI. One of the more infamous examples is the COMPAS system, which has proven demonstrably biased against minorities in predicting their risk of recidivism. However, COMPAS is just one of many AI systems designed to augment, if not automate entirely, sentencing recommendations, criminal risk assessments, and the rendering of judicial decisions altogether. This latter type of AI system is referred to as a "robot judge" and the core ethical questions include: *What information is the AI system basing its judgements on?* and, *How do I appeal a decision it got wrong?* On the surface, we might think the central focus of these questions is justified, given that these types of systems consistently fail to deliver excellence in transparency, explainability, and reduction of bias. So, electing a different ethical perspective seems dubious, if not irresponsible. But if we focus on a broader set of ethical concerns - balancing risk and *Eudaimonia* - a new and potentially more serious issue becomes clear.

A case that's perpetually delayed fails to deliver justice just as much as an AI system's sub-optimal ruling. Justice delayed, after all, is justice denied. This problem is increasing dramatically, essentially cutting off vast populations from their own justice system. Richard Susskind, author of the recent book, *Online Courts and the Future of Justice*, describes the stunning backlog of 100 million cases in Brazil and 30 million in India as instances of "manifest injustice." In the U.S., Milwaukee County *alone* had a backlog of 1,700 felony cases and 3,100 misdemeanor cases as of December 2021. These are scenarios where neither party to a case can attain closure, accountability, or redress based on a definitive ruling; their plight is frozen in time, subject to no ruling at all. The examples of Brazil and India show that, in sufficiently large numbers, case backlogs can render an entire society's justice system moot.

In modern society, justice often facilitates flourishing – or at least removes impediments to it. This begs a different set of questions, such as: *How might we increase access to the justice system for everyone?* and, *How do we ensure fair outcomes with a non-traditional judicial process?* An approach that balances both risk and Eudaimonia would focus on the specific ways in which we can make these AI systems more effective, not whether we should build them at all. Practical examples include adjusting a model's weights to reduce the impact of historical bias in the data; ensuring a legal right to interrogate the system's decision-making process; and developing a special class of appeals for algorithmic rulings. With this perspective, we might argue that AI-driven judicial systems, while currently imperfect, are moral necessities. And that a more Eudaimonia-driven AI Ethics would seek to improve their efficacy instead of letting the demonstrably unjust status quo rule the day. Instead, their risks too often crowd out discussion of their potential to help us flourish. We must rebalance that equation.

–

Conclusion

What does it mean to live a good life? How can AI help us flourish? These are questions that AI Ethicists should make central to their work. No longer should an AI system's potential benefits be separated from its potential risks. In fact, a more robust – and historically accurate – ethical calculus will focus on the net good that an AI system will generate over its operational life. To start, the field should focus its analytical approach on three questions:

1. **What is the maximal good an AI system can do**
2. **What are the potential risks in its design;**
3. **How can we mitigate these risks so that the maximal good can be achieved? The order of these questions is intentional, shifting the locus of concern from harms to happiness, from risks to flourishing.**

This will help break AI Ethics out of its stasis, opening up new avenues for its practitioners to advance into. Ethics was never about compliance. Nor was it simply about the difference between right and wrong. Rather, it provided the overriding question of philosophy in ancient times: How can we be happy and flourish? Let us recenter Eudaimonia in our work. Downside risks must be put into context. AI's potential for societal good must be prioritized. AI Ethics must not miss the forest for the trees.

Making the case for ethics

Talking points:

- Ethics is more than just compliance; it is also a tool that is meant to be used to determine the maximal good a system can achieve
- Ethicists are trained to identify potential risks as well as potential opportunities within a system
- Having an ethics board unlocks the capacity for responsible innovation at scale

By the numbers:

- **93**% of consumers think that companies have a responsibility to positively impact society (Salesforce)
- **40%** of managers in companies with at least $100 million in annual revenues reported that they are already experiencing accelerated innovation as a result of their Responsible & Ethical AI efforts (MIT Sloan)
- **43%** of responsible AI leaders report accelerated innovation as a result of their Responsible & Ethical AI efforts (BCG)

How Diversity Leads to Better Business Outcomes

Written by Yattish Ramhorry

The game of basketball is played with teams of 15, with five on the basketball court at any time. Teams consist of two forwards, two guards, and a center. The object of the game is for one team to score more points than the opposing team.

Every player on a basketball team brings unique abilities and skillsets that they employ for the duration of the game. To succeed, each member of the team must work together as a cohesive unit and defeat the opposing team.

But what does the game of basketball have to do with AI Ethics or even business?

Sports and business share similar characteristics. In both basketball, and business, every member of the organization, or team, brings a unique set of skills, knowledge, and experience based on their academic and cultural backgrounds. A company is usually composed of two or more individuals from different backgrounds and levels of experience.

-

What is cognitive and demographic diversity?

When a wide range of people with different backgrounds and experiences are involved in the development and governance of AI, we say that there is “cognitive and demographic diversity.”

Cognitive diversity refers to the different ways that people think and process information. We all have different biases and ways of looking at the world. This can be a good thing, as it allows us to see things from different perspectives and develop new and innovative solutions.

On the other hand, demographic diversity refers to the variety of backgrounds and experiences people have. It includes gender, race, ethnicity, and socioeconomic status in decision-making.

Demographic diversity is essential because it ensures that a wide range of voices are heard in the AI debate.

Cognitive Diversity
Variation of thought, skillset, and manner of processing information

Demographic Diversity
Variation of gender, race, ethnicity, and socioeconomic status

Why should we create AI systems with cognitive and demographically diverse teams?

Cognitive and demographic diversity are vital to creating fair and inclusive AI. It can help us to identify potential ethical problems with AI systems. For example, people from different cultures may have different views on what is considered acceptable or ethical behavior.

AI designed by people with different perspectives and backgrounds will likely be more resilient to potential ethical problems. This is because they will have been designed with a broader range of inputs and perspectives.

It’s essential to have diverse opinions regarding ethical decision-making in AI to avoid conflicts and misunderstandings.

“If everyone thinks the same way, we are more likely to make the same mistakes.”

When it comes to AI, we need to be careful that we’re not creating echo chambers. Echo chambers refer to a situation where AI is only exposed to one type of data or opinion. This leads to AI that makes decisions that are not reflective of real-world situations or not in line with ethical or moral standards.

Ensuring that diverse perspectives are included in the development and regulation of AI will help us to avoid the risk of perpetuating harmful biases.

Cognitive and demographic diversity can help to ensure that the AI we build is designed for the benefit of all people. If AI systems are designed by people from a narrow range of backgrounds, they may inadvertently exclude or discriminate against certain groups of people.

For example, if the only people designing and building AI systems are white males, it becomes increasingly likely that those systems will reflect the biases and values of that group. This inevitably leads to systems that discriminate against women and minority groups.

Big tech companies and smaller startups are only now beginning to realize the true value and potential of having a diverse team. However, it wasn't always this way, as several big tech companies like Facebook, Apple, and a few others had to learn these lessons of diversity inclusion the hard way.

An article published online for ProPublica in October 2016 reported that Facebook's ad delivery process allowed advertisers to exclude black, Hispanic, and other "ethnic affinities" from seeing ads posted to their platforms. Facebook was allowing advertisers to target users not only by their interests or background, but it also gave advertisers the ability to exclude specific groups that it calls "Ethnic affinity."

In another article published by Wired magazine in November 2019, it was reported that the Apple Card (credit card) algorithm offered smaller lines of credit to women than to men. The scandal broke on Twitter when several disgruntled customers of Apple's credit card lashed out, saying that the Apple card was "sexist." No one at Apple could explain how the algorithm worked or even justify its output. Goldman Sachs, the issuing bank for the Apple Card, insisted that there was no gender bias in their algorithm; however, it failed to provide any proof.

By considering a diversity of perspectives, we can ensure that AI systems are designed fairly and equitably for all.

–

How can cognitive diversity improve business outcomes?

Cognitive and demographic diversity is an invaluable asset to businesses. It allows companies to benefit from a range of perspectives and skill sets.

Organizations that embrace cognitive diversity and demographic diversity can create an environment where everyone's ideas and opinions are valued. This environment encourages collaboration and the sharing of ideas, which can help identify new opportunities and solutions to existing problems.

In addition to leading to better business outcomes, cognitive diversity can also help increase employee engagement and motivation. When employees feel that their opinions and perspectives are heard and respected, they are more likely to feel motivated and engaged in their work.

This can lead to increased productivity, as employees are more likely to go the extra mile when their ideas are taken seriously.

Cognitive diversity is an essential tool for businesses looking to gain a competitive edge and increase their profitability. By recognizing and leveraging the unique perspectives of their employees, organizations can benefit from new ideas and creative solutions. This can lead to improved business outcomes and increased engagement among employees.

A diverse workplace with different backgrounds, life experiences, and perspectives can help broaden the scope of ideas and solutions. For example, a team consisting of both men and women can bring different outlooks and opinions to the table, leading to a more creative and successful outcome.

In basketball, success depends on each member bringing their unique talents, working together towards the same goal. The same is true of businesses. To succeed, they must have a balance in decision-making, experience, and skills which can only come with cognitive and demographic diversity.

Having a dynamic and experienced team of Ethics Experts on your (development) team will help you to use ethics to guide your design and development decisions and align them with your values to reach your end goal of sound technology.

Making the case for ethics

Talking points:

- Teams with cognitive and demographic diversity are better positioned to unlock innovation that drives market growth
- Ethicists are trained to identify when and where cognitive and/or demographic diversity is needed
- Establishing an ethics board is an effective solution to bringing in necessary diversity to the design and development processes

By the numbers:

- Diverse companies are **70%** more likely to capture new markets (HBR)
- Diverse management teams lead to **19%** higher revenue (BCG)
- Diverse teams are **87%** better at making decisions (People Management)
- Corporations identified as more diverse and inclusive are **35%** more likely to outperform their competitors (McKinsey)

THANK YOU TO OUR CONTRIBUTORS

Anna Danés is an AI Ethicist who works with organizations to create frameworks and practical methods to ensure ethical developments of AI.

Divyansh Agarwal is a Senior Research Engineer with the interactive AI team at Salesforce Research.

Geoffrey M Schaefer heads the AI Ethics and Safety Team at Booz Allen Hamilton.

Yattish Ramhorry is a Data Strategist, Software Engineer, and the Founder of 4IR Technologies.

Simone Larsson is Head of Strategic Development at Ethical Intelligence.

Jacob Beswick is the Director for AI Governance Solutions at Dataiku.

Patricia Shaw LLB (Hons), LLM, FRSA, CEO of Beyond Reach Consulting.

THANK YOU FOR READING

Subscribe for future issues delivered straight to your inbox

Let us know what you think of the EQUATION

Twitter | **@ethicalai_co**
LinkedIn | **Ethical Intelligence**
Email | **info@ethicalintelligence.co**
Website | **www.ethicalintelligence.co**