



How To Build Responsible AI

Step 1: Accountability



How To Build Responsible AI

Step 1: Accountability

The development, deployment and operation of irresponsible AI has done, and will continue to do, significant damage to individuals, business, markets, societies and economies of every scale. Now is the time to be explicit in the processes and systems that we create.

I recently outlined the six essential elements of a responsible AI framework that can be adopted by any organization in any industry:

1. Accountability
2. Impartiality
3. Resilience
4. Transparency
5. Security
6. Governance

In a series of articles, I will explore each one of these elements and its crucial role in building the responsible AI of the future. The first component of responsible AI that I will address in this second article in the series is accountability, which is especially important in areas such as supply chain, finance, national security and intelligence, cyberbalkanization, data protection, data destruction and data/algorithm aggregations.

Rather than assume we all mean the same thing when we use the term “accountability,” I’ll now suggest three critical features for how we distill the term and understand it beyond its etymology. These three features provide expanded ways to measure or assess performance in order to improve accountability in AI; they are: the audit, the ledger and the reasoning.

1. Audit

When you have accountability in AI, you are able to show the provenance of your data. In an audit — which I often call a “gnarly audit” because it gives you unvarnished and sometimes uncomfortable facts — you trace the trail of data back to its source. Where did it come from? How was it aggregated? How was it trained? Who was responsible for it?

Have we done a meaningful review of first-order data use, as well as all secondary and subsequent uses? Are all of these applications clearly defined and communicated?

Follow the same process for auditing algorithms, models and ensembles. As data is aggregated into more complex forms, rendered in such ways as neural networks, some people use its complexity as an excuse not to explain its origin. But that’s not acceptable. An audit ensures you understand each link in the data chain. If you’ve taken humans out of data preparation, how have you automated processes, personas, decisions and actions? When an action occurs, when and how is it recorded? What consequences, good or bad, does it lead to?

For high-profile examples of potential negative consequences, just look at the news. Major tech companies, including Google and Microsoft, have come under fire for racial, ethnic and gender biases in their AI face detection services. Facebook recently issued an apology after its AI software asked users watching a video featuring Black men if they wanted to keep watching “videos about primates,” causing the company to investigate and disable the AI-powered feature pushing the message. For many errors in image analysis systems like this one, experts have pointed to flaws in the datasets used to train algorithms. This is precisely the need for an audit.

When flawed data and biases are built into influential AI programs, like those used by law enforcement agencies, financial institutions or HR departments, they can do serious and lasting harm. An audit is an honest risk assessment, revealing problems in your data before they get out of control — and can otherwise be the means of exacting accountability when needed.

2. Ledger

In tech, we often focus on technical debt, or the time, money and resources your business will have to expend in the future to make up for the limited solution or shortcut you opted for today. We would do far better to talk less about technical debt and talk more about technical margin — what you are doing now to be ahead of your competition in the future. What steps are you taking to create a margin instead of always operating in the red?



How To Build Responsible AI

Step 1: Accountability

To develop greater accountability, create a ledger to measure your technical debt vs. technical margin, treating data as a strategic asset. Get started with a simple ledger that collects your accounts; for each, note a starting balance and record transactions, either debit or credit, and a closing balance. This can include risk, data health, decisioning speed, readiness or otherwise. You've already tracked the data, processes, algorithms, personas, decisions and consequences in the audit. The ledger helps you understand the interactions among those systems, recording what is happening and what the outcomes are in real time.

3. Reasoning

A third feature for enhancing accountability is to provide controls, standards and processes for documenting the logic for how conclusions are drawn — what I call “reasoned reasons,” or simply “reasoning.” In a responsible AI framework, you need strong checks, counterbalances, consequences and rewards. With the complexity and velocity of AI, we can't (or shouldn't) understand reasoning purely by outcome. It is best to evaluate our reasoning constructs in AI by Expected Value (EV) and how those are deduced, considered and acted upon.

Because of this, I don't believe we're likely to be successful by “striving for balance.” A ballet dancer never balances herself; she is always counter-balancing herself, and the same is true with AI. We're constantly counter-balancing against risks and other environmental effects. How we reason and what we do to understand that reasoning is what will make us more or less accountable.

If you identify an issue in your business, what can you do to counteract it? What checks or guardrails do you have in place? For example, what kind of board or bias team are you establishing to prevent racial or gender biases in your AI systems? Who is ultimately accountable for what happens with your AI? Clearly define the consequences and rewards for responsible parties within your company.

Accountability is an integral component of responsible AI. We will be held accountable for our actions and non-actions in relation to AI, data, algorithms and their operationalization. Let's not be slaves to process or insist on perfect conditions to start developing accountable AI. We won't be able to anticipate every problem or contingency plan — but it's far

better to have an imperfect plan that can be changed and updated to meet the circumstances than no plan at all.

We can be dutiful in concentrating on these three key features of accountability in order to develop, measure and deliver ever-more accountable AI for our individual, business, market, social and economic best interests.



Aaron Burciaga, CAP, ACE

Chair of the Analytics Certification Board. Roles have included VP, CTO, and Practice Director of Fortune 500s and co-founder of data and technology startups.

Aaron is a Forbes contributor, frequently invited keynote and speaker, and Certified Analytics Professional (CAP). He is an appointed member of the U.S. Department of Commerce's National Technology Information Service (NTIS) advisory board. Aaron received his M.S. Operations Research from the Naval Postgraduate School and his B.S. from the US Naval Academy.