## ADP:  Ethics in Artificial Intelligence

ADP is a leading global technology company that supports tens of millions of people in 140 countries with its comprehensive human capital management solutions. With over 70 years of experience responding to changing workforce needs with insights and solutions, ADP has built a legacy of innovation and service. This legacy has yielded ADP a profound understanding of the lives of people and their success in the world of work and has embedded within the organization a great respect for the trust placed in ADP.

ADP holds the trust of its clients and their employees at the heart of its mission, vision and values. First and foremost, we are committed to upholding strong ethics as part of our core business approach—not just because we believe it gives us a competitive advantage, but because it is the right thing to do. Our associates reflect this every day in their actions and commitments, including placing a great focus on Privacy.

As technology evolves, new opportunities arise; with them, however, can come the potential for unintended effects. Automation fueled by artificial intelligence (AI) is helping transform organizations by reimagining the way we work, and even the nature of the work itself. Machine learning (ML), a branch of AI, allows machines to make decisions by feeding them data to learn on their own, without constant human supervision. ML is helping accelerate innovation by reducing the time required for process-driven work and providing opportunities to perform tasks more effectively, quicker, and at a scale that exceeds human capacity.

Aspects of this automation are at work today: identifying discrepancies in payroll registries, automating the completion of forms, or providing recommendations for financial programs or retirement offerings. While powering positive advancements, ML applications can potentially have negative implications as well, such as arriving at incorrect recommendations or amplifying factors of bias. To counter such impacts, ADP has adopted a rigorous set of principles and processes to govern its use of newer technologies, including real-time, operational monitoring of automated decisions, much in the same way we monitor transaction and payroll processing. This awareness ensures that our models are relevant and functional, and do not "drift" from their intended goal unknowingly or inappropriately.
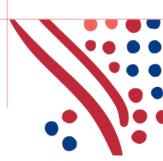
Included in this approach is the isolation of unintended bias. Bias in the world of work—that is, favoring or disfavoring one group compared with another—is caused by a variety of factors. Ethically, anyone developing ML technologies should be vigilant not to reproduce such bias in any ML-enabled product or service. Even when accounting for potential unintentional bias in the source data, coding, or use of an AI-enabled product or service, there can be unexpected or unforeseen bias that come into play. ADP's goal is to continually strive to identify new and unexpected sources of bias and then refresh and enhance the design of our client offerings to address them.

We at ADP believe it is important to abide by the following core set of guiding principles, which put people at the center of everything we do:

1. *Human Oversight.*  ADP believes that human oversight is essential to the reliable operation of ML models and making proper use of their results. Our solutions provide recommendations to human decision-makers, which they can then decide how to act upon.

2. *Governance.*  While AI holds the potential to mitigate human bias, without proper oversight it can incorporate bias as well.  We have implemented audit and risk assessments to test our ML models as the baseline of our oversight methodologies. We continue to actively monitor and improve our models and systems to ensure that changes in the underlying data or model conditions do not inappropriately affect the desired results. And we apply our existing compliance, business ethics, and risk management governance structures to our ML development activities.

3. *Privacy-By-Design.*  We have implemented an enterprise-wide Privacy-By-Design approach that incorporates privacy and data security into our ML model development and data processing systems more generally. ADP provides information about how we handle personal data in privacy statements made available to our clients' employees, consumers, and job applicants.  Our ML models seek to minimize access to identifiable information to ensure we use only the personal data we need to generate insights. We also maintain a robust security program for our ML models, including designing them in line with our security standards and protecting them against misuse or compromise.

4. *Explainability and Transparency.* We strive to develop ML models that are explainable and direct, with clear purposes. Our ML models are designed with understanding as a key attribute, measured against an expressed desired outcome. We test and evaluate our ML models accordingly, adjusting as needed to maintain accuracy in line with the models' purposes. We provide our clients with information about how our ML models operate, their proper use, and their limitations, so that clients can implement those models in accordance with their design and purpose, operate them effectively, and use their outputs as intended.

5. *Data Quality.* Understanding how we use data, and the sources from which we obtain it, is key to our ML model development. We maintain processes and systems to track and manage our data usage and retention from across ADP systems or processes. If we use external information in our models, such as government reports or industry terminologies, we understand the processes and impact of that information in our models. All data included in our ML models is monitored for its quality and accuracy, as well as for changes that could alter the desired outcomes.

6. *Culture of Responsible AI.* We have an active AI & Data Ethics Committee, comprised of both industry leaders and ADP experts across our business, which meets on a regular cadence and reviews our design principles. The Committee advises on emerging industry trends and concerns and provides guidance with respect to the principles ADP should follow while developing products, systems and applications that involve AI and data.

7. *Inclusion and Training.* We are committed to having diverse teams design and develop our ML models, to ensure a wide variety of perspectives and experience are considered. After all, ML models impact humans, and human experience should inform that impact. In addition, we support skills development to accelerate the growth of a diverse workforce that can develop and deploy the ML solutions of the future.