



The AI Manifesto

Nine guiding principles for responsible & trustworthy application of AI with impact

AUTHORED BY: PETER VAN DER PUTTEN



Is AI a cool technology? A business strategy?

How do you keep control while reaping benefits for your business and your customers?

With public discourse bouncing between techno-utopianism and doomerism, we've created the definitive manifesto on AI for the enterprise. Armed with these principles, you'll be able to cut through the AI hype and chart your journey to drive real, business-defining impact.

PETER VAN DER PUTTEN



TABLE OF CONTENTS

01. **Business goals** | The why

Intelligence put to action is the key to unlocking value in AI. Put it to work in processes and interactions to drive meaningful results as you head on the road to an autonomous enterprise.

03. **AI strategy** | The what & how

Strong AI strategy needs to cover all forms of AI to be complete, actionable, and successful. Build your strategy by working side by side with AI, exploiting AI's rational and creative powers, and using both public AI and private data.

07. **Ethical, responsible & trustworthy AI** | The right way

When it comes to balancing business goals and customer needs, what is the right way? Learn how we're operationalizing ethical principles into intelligence assets and best practice to do what is right, not just what is allowed.

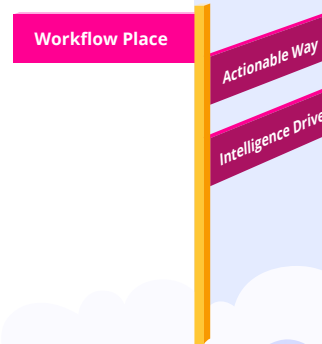
01.

AI without action is a science experiment with no path to value

Intelligence is useless if not put into action in business processes or customer interactions. Start with the real business outcomes you want to achieve to understand which AI-powered decisions and generative models drive the most value. Then identify the specific workflows and customer interactions you want to impact. Having well-structured interactions, workflows, and case management creates the data and feedback loop required for additional optimization through artificial intelligence.



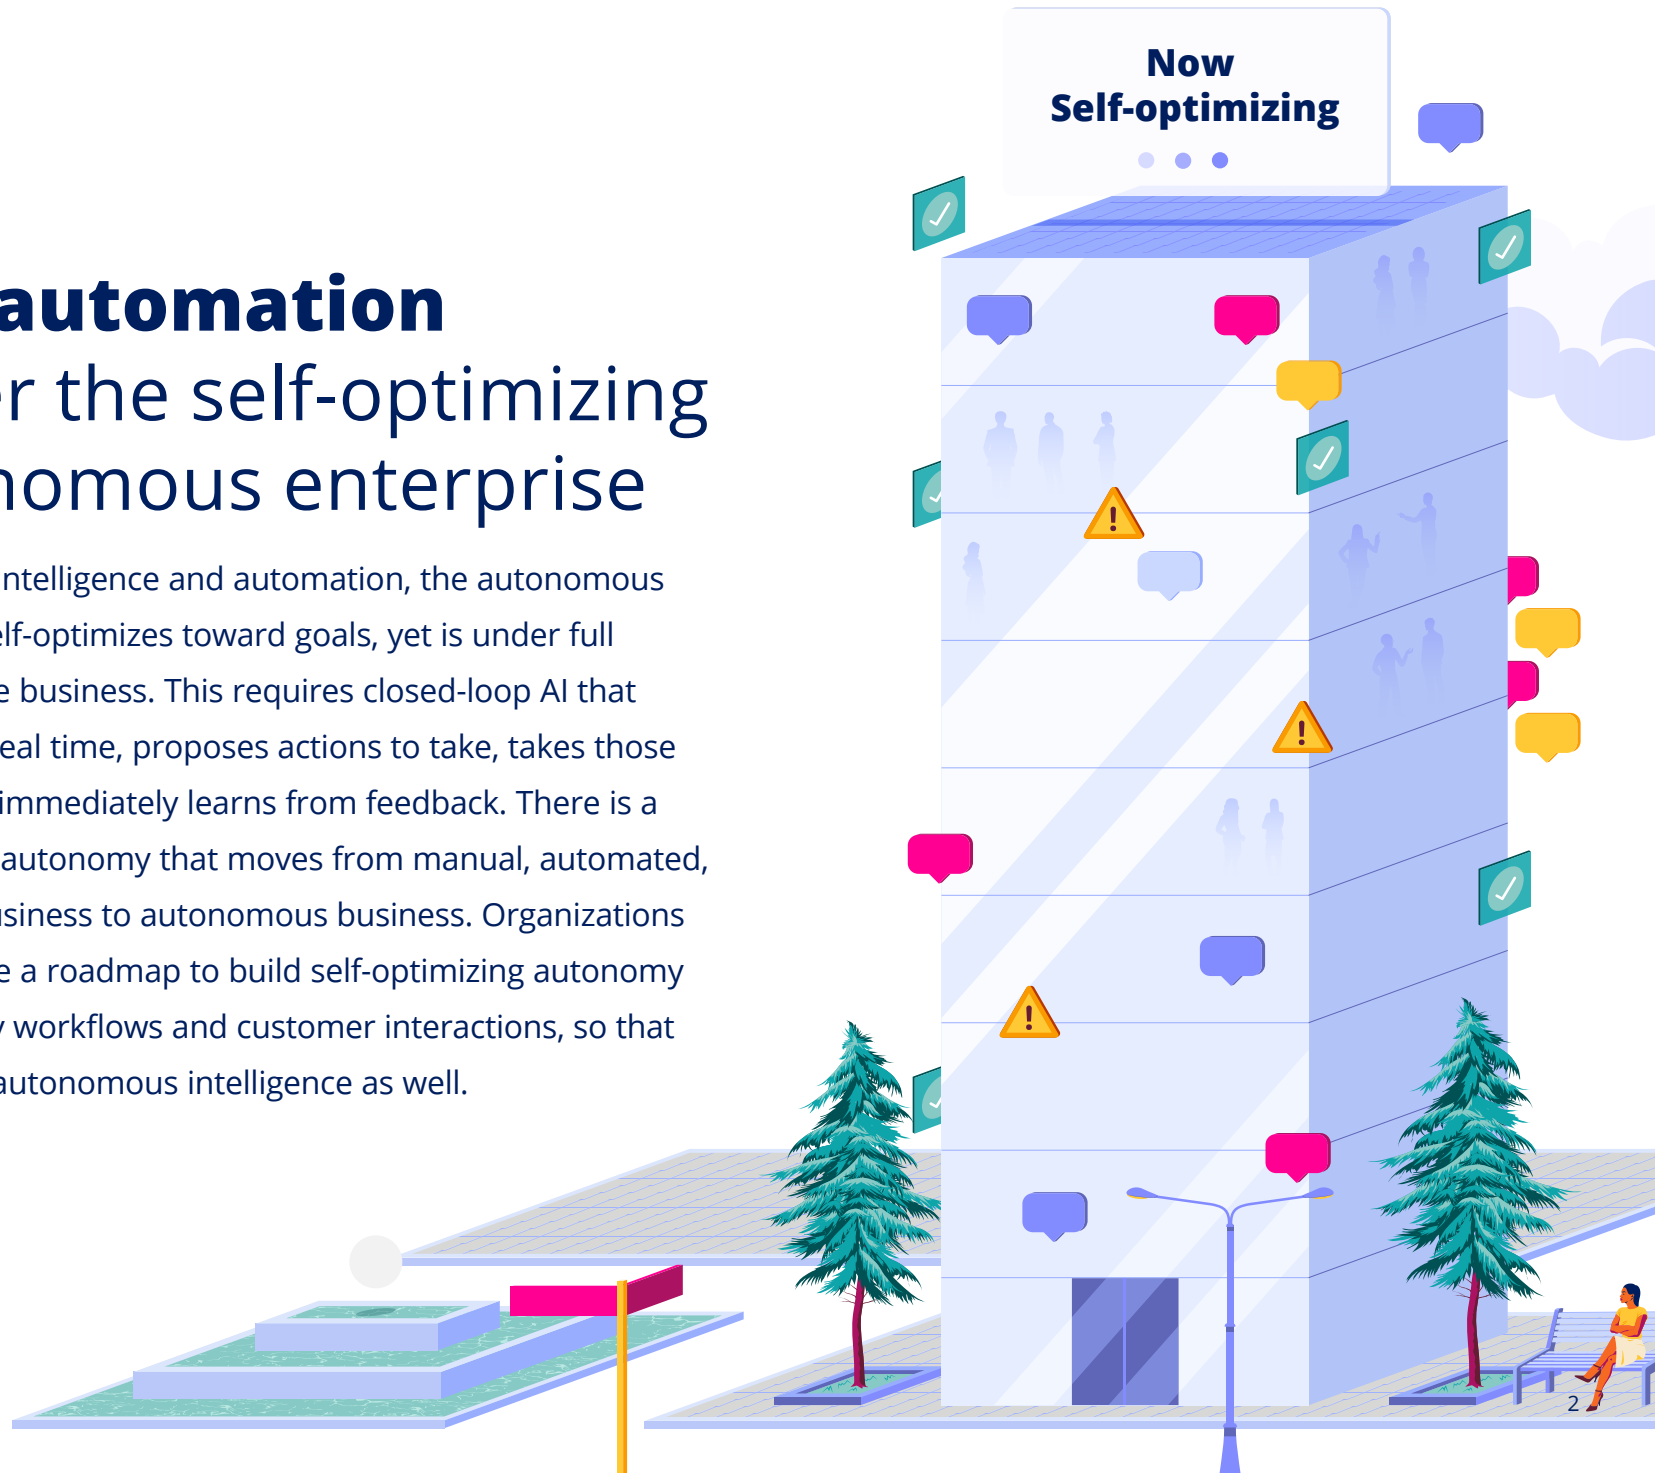
Intelligence is useless if not put into action in business processes or customer interactions.



02.


AI & automation power the self-optimizing autonomous enterprise

Powered by intelligence and automation, the autonomous enterprise self-optimizes toward goals, yet is under full control of the business. This requires closed-loop AI that executes in real time, proposes actions to take, takes those actions, and immediately learns from feedback. There is a path toward autonomy that moves from manual, automated, intelligent business to autonomous business. Organizations should create a roadmap to build self-optimizing autonomy into their key workflows and customer interactions, so that AI becomes autonomous intelligence as well.





...AI is a tool designed by humans, and it is most valuable in instances where it is assisting...



AI STRATEGY

The what & how

03.


AI is augmented intelligence – **it's best with humans in control**

Hollywood movies and “AI doomer” blogs want to make you believe AI is all about taking control away from humans. The reality is that AI is a tool designed by humans, and it is most valuable in instances where it is assisting the agent, employee, developer, marketer, or customer by providing guidance and taking on tasks. To minimize risk, AI strategy should prioritize human control by considering “human-in-the-loop” use cases or ensuring humans are monitoring and steering autonomous AI in play.

04.

There is more to AI than just gen AI – **you need left & right brain AI**

Generative AI is injecting creativity – right-brain thinking – into the way we do business and design applications. But many automated decisions in workflows and gen AI driven interactions require “left-brain AI,” the analytical and rational AI capabilities that drive real-time automated decision-making, recommend the next best action across all channels, and proactively spot process inefficiencies before they become challenges. Generative AI may have gotten your C-suite excited about AI-powered transformation, but your AI strategy needs to cover all forms of AI to be complete, actionable, and successful.





05.

Start with outcomes & decisions, not with data and models

It is easy to get lost in data swamps, or marvel at all kinds of machine learning models that could be created. A solid and secure data strategy is important, but you must think top-down, not just bottom-up. What outcomes in the business do you want AI to optimize? What customer experiences do you want to improve? What are the automated decisions and generated intelligence that can drive these outcomes? Which models can automate the decisions? Answering these questions will help prioritize what data matters most to an organization.





06.

You'll need an open, **best-of-breed model** ecosystem

The debate between “public” and “private” AI is a distraction. Organizations will need to run private models – trained on their own data – for differentiated use cases such as customer engagement and process optimization. In other cases, publicly available models can provide a secure, fast, and scalable path to value for commoditized use cases, such as optimal character recognition (OCR) and image recognition. Many generative AI foundation models and services are available to build upon. Establishing an AI architecture that protects private data while supporting private models, fine-tuning public and open-source models, and using public AI services provides the flexibility and security to adopt and scale the right AI, for the right use cases, in the timeframes markets will demand.

07.

AI should be empathic to all stakeholders

Empathy is about putting yourself in others' shoes and doing what's right not just for you, but for everyone else. This is key for ensuring AI is trusted, responsible, accepted, wanted, and compliant. Employees and customers are fearful that organizations may use AI to drive profits at their expense. To ensure buy-in, your AI initiatives should deliver value to customers and employees, and striking the right balance should be ingrained at all levels of automated decisioning and AI.




To ensure buy-in, your AI initiatives should deliver value to customers and employees.

08.

Build ethical principles into your tools & processes

Everyone will state that AI needs to be fair, transparent, safe, accurate, robust, privacy-friendly, auditable, managed, and accountable – promoting organizational benefits as well as societal and environmental well-being. You must bake these principles into your tools and best practices. Understand and govern which decisions and use cases require explainable AI, and where more opaque models are acceptable. Explanations are not just for data scientists but should also target end users such as agents and customers. Create processes that continuously check for hidden bias in models and automated decisions. Ensure all automated decisions and gen AI interactions are audited – both for you and your customers.





The real question is,

within what is allowed,
what should be done?

09.

AI ethics goes beyond **regulatory compliance**

The U.S., EU, and other governments globally are moving rapidly to create guidelines and laws for the use of AI, but it's a trap to think about AI ethics and responsible AI solely in terms of what is allowed and what is not. The real question is, within what is allowed, what should be done? For instance, using AI to promote conversations that benefit customers and are relevant to their needs will deliver far greater payoff in the long run than pushing sales opportunities in the short term. Ensure that the customer is at the center of your AI strategy by driving outcomes that deliver both a great customer experience and improvements to your bottom line – turning an ethical and empathetic approach into a competitive advantage.



Pega delivers a powerful low-code platform that helps the world's leading brands adapt to rapid change. Clients use our AI-powered decisioning and workflow automation to solve their most pressing business challenges – from personalizing engagement to automating service to streamlining operations. Since 1983, we've built our scalable and flexible architecture to help people focus on what matters most, so they can meet today's customer demands while continuously transforming for tomorrow.