

The AI Product Manager

Beyond Prompts and Output

A Taste of What's Inside

05

The current state of AI in Product Teams

06

From prompt engineering to conversational orchestration

07

The new failure mode: context erosion

08

What we learned while building our own product

10

The capability gap is not technical, it is behavioral

11

Experimentation and trust: fintech and consumer brands

12

Efficiency is not the point, effectiveness is

13

Software as clay and what it might mean for product roles



Summary

In our earlier work on AI-powered product discovery, we described a shift that many product teams are starting to feel: the rhythm of discovery is changing. AI reshapes what becomes easy, what becomes optional, and what becomes fragile.

This paper is a follow-up based on what we have observed in the market and what we have learned while building our own product using AI, developer expertise, and agent-based workflows. Our conclusion is simple: the role of the Product Manager and Product Owner is evolving in ways that require new thinking.

Most organizations currently use AI to automate what they already do: writing PRDs, generating user stories, summarizing meetings, expanding backlog items. This creates a risk worth taking seriously: organizations that focus on automating existing work may find themselves producing more while choosing, learning, and steering less effectively.

We believe the next era of product management requires both imagination and rigor. AI makes building and exploration easier, raising the premium on problem selection and creative thinking. At the same time, AI increases the risk of synthetic progress, context erosion, and the illusion of certainty and stability.

In high-trust environments, whether regulatory-driven like fintech or reputation-driven like global consumer brands, that rigor must include safety, quality consistency, and trust as a strategic asset.

This paper proposes a practical view of the AI PM as a guardian of effectiveness in a system that can now accelerate itself.

Why this is a **follow-up** and not a restart

The public conversation about AI in product management has become predictable. AI tends to be framed either as a productivity tool or as an existential threat to the role. Both framings are reductive and miss what is actually changing.

AI changes the economics of product work. It changes what is scarce. When building becomes cheaper and iteration becomes easier, the bottleneck shifts toward clarity, coherence, judgment, and trust. This shift is already visible in the way teams work with AI today, and in the kind of mistakes they are likely to make next.

This paper focuses on the person and the role: what the Product Professional and Product Manager must become when AI can produce vast amounts of content, analysis, and prototypes with minimal friction.



The current **state of AI** in product teams

Across industries, we see only a small number of product people using AI in a way that materially changes their work. Most usage remains shallow, even in organizations that talk about AI frequently.

The typical pattern looks like this

- PRDs are generated **faster**, but the underlying assumptions are not tested any earlier.
- User stories are expanded **automatically**, but the user problem is not better understood.
- Backlog items contain more “**context**,” but no one reads it because it is too long or too generic.
- Documentation is produced in greater **volume**, but decision quality does not improve.

AI naturally fits into existing habits, which explains why most usage remains shallow.

Many Product Professionals are trained to operate inside delivery systems that reward output, predictability theater, and solution mode.

When AI is introduced into such a system, it tends to optimize the current workflow rather than improve outcomes.

In practice, AI becomes an accelerator for the organization's current maturity level, amplifying whatever way of working is already in place.



From prompt engineering to conversational orchestration


A lot of AI training for product roles still revolves around prompt engineering. We understand why. It is tangible, teachable, and it produces visible artifacts. But we increasingly believe prompt engineering is the wrong mental model for the Product Professional and Product Manager.

Conversations over one mega prompt

The more accurate model is conversational orchestration: the ability to think in dialogue, to structure reasoning through interaction, and to iterate with AI as a companion that can surface patterns, gaps, contradictions, and alternative frames.

This matters because AI can become a cognitive partner in the messy part of product work: framing problems, testing logic, exploring options, and exposing assumptions.

In that world, the job is maintaining a disciplined conversation that leads to better choices.



This is also where many teams fail. A conversation with AI can easily become output theater. People can generate convincing narratives and extensive documentation without strengthening their understanding of the problem. The activity looks productive, but it does not necessarily improve direction.

The new failure mode: context erosion

One of the most under-discussed risks of AI in product organizations is what we call context erosion.

AI enables teams to generate richer artifacts: more background, more stakeholder summaries, more “analysis.” In response, teams stop reading. They rely on AI to summarize what AI produced, or they paste the artifacts back into AI to extract the “essence.”

The result is progressive compression of meaning.

Strategic intent becomes diluted. Trade-offs disappear. Constraints get lost. What remains is a simplified narrative that is less accurate, less nuanced, and less useful for making hard decisions. AI can create a situation where the organization produces more context while losing collective understanding. Adding more context in this situation deepens the problem.

This is why the future AI PM is a curator and guardian of context. Their responsibility is to decide what must remain explicit, what can be compressed, and what must never be delegated to automated summarization.

*This means
preserving
things like*

- Strategic intent
- Non-decisions and explicit “no”
- Trade-offs and consequences
- Safety boundaries and compliance constraints
- What would make an initiative unacceptable, not merely unsuccessful

Those elements should not be treated as optional metadata. They are the core of coherent product work.

What we learned while **building** our own product

Our view is not only based on observation. We have been building our own product and using AI actively in the process. We started in a familiar place: validating assumptions with prototypes. That early validation gave us clarity on what was valuable and what was not. We did not begin with AI. We began with hypothesis and evidence.

Once value was validated, we used conversational AI to shape the logic and functionality. Because we had sufficient technical fluency, we could translate conversations into working behavior. We could build, test, adjust, and build again without waiting for a full team to interpret our intent through layers of documentation.

We then involved a developer-expert to evaluate quality. This became a recurring lesson: AI can accelerate creation, but it does not guarantee quality. A human expert remains essential for judgment. AI generates what you say it should generate, a human expert understands whether it is put together in such a way that in the long-term it still works, and not just at this moment.

Next, we introduced a group of agents, effectively a small development team, to take on tasks we had already described. Remarkably, they could reverse engineer and create documentation based on earlier prompts and prior decisions, essentially reconstructing logic and intent after the fact.



This experience reinforced two critical points:

First, AI can collapse the distance between intent and implementation if the person driving it has clarity and enough technical fluency to verify what is being produced.

Second, AI can also collapse responsibility if the organization is not careful. The more the system generates and reverse engineers, the easier it becomes for people to assume coherence exists because documentation exists. But coherence is not a property of documents. It is a property of disciplined decision making.

The future Product Professional needs the capacity to stay in learning mode even when solutions are cheap.

This is where Product Professionals must evolve: away from writing artifacts and toward designing and protecting the integrity of the decision process.

The **capability gap** is behavioral

A common assumption is that the gap in Product Professionals is primarily technical: that they lack development knowledge or data fluency. Technical fluency matters. But the bigger gap we see is behavioral.

Many Product Professionals are not trained, rewarded, or supported to be creative and curious. They are rewarded for being responsive, for managing stakeholders, and for keeping output flowing. They enter solution mode quickly because that is what the system demands. AI makes this worse by making solution generation effortless.

This is why we agree with the observation we recently heard from a senior Product Professional: this may become the era of the creatives, where small groups can move extremely fast. But speed without curiosity is just faster repetition. Creativity in product work is not artistic expression. It is the ability to reframe the problem, to imagine alternatives, and to resist premature convergence on “the obvious solution.”

The future Product Professional needs the capacity to stay in learning mode even when solutions are cheap.

Experimentation and trust: FinTech and consumer brands

In FinTech and similar environments, we see a consistent pattern: organizations are curious but also deeply afraid. Fear is not irrational. In regulated domains, trust is institutional. Mistakes carry direct legal and financial consequences. Experimentation is constrained, and the path to production requires governance and auditability.

In global consumer brands, the constraints can look different. Compliance may be lighter in many areas, but reputational risk can be brutal. Trust is emotional and brand-based. One AI-driven mistake can scale globally through social channels and permanently damage perception.

Consumer brands can often design experimentation differently, with more options to scale experiments in controlled ways if they choose the right use cases and the appropriate boundaries. This changes how risk is managed.

Conditions for AI experiments

- PRDs are generated faster, but the underlying assumptions are not tested any earlier.
- User stories are expanded automatically, but the user problem is not better understood.
- Backlog items contain more “context,” but no one reads it because it is too long or too generic.
- Documentation is produced in greater volume, but decision quality does not improve.

These are product requirements in an AI world.

Effectiveness is the point

Many AI narratives frame the future of product management as a story of efficiency. The assumption is that work remains the same, only faster. We believe this frame is strategically limiting.

AI may indeed make teams faster, but that is not where the real advantage lies.

The advantage lies in effectiveness: improving the quality of decisions, reducing wrong bets, strengthening learning loops, and increasing the ability to deliver impact rather than output.

Effectiveness also requires patience, which is becoming rarer. AI increases the pressure to move quickly because it removes many frictions. Product work requires deliberate slowing at the right moments: to reframe, to validate, to check trust implications, to clarify trade-offs.

This is where the future Product Professional becomes a leader, by protecting the organization from mistaking speed for progress.

Software as clay and what it might mean for product roles

A final theme we have been exploring is the idea that software may become more like clay: less defined by fixed features and more defined by modelability. Rather than shipping personalized and customized solutions through heavy engineering, organizations may offer systems that can be shaped and configured more deeply, more dynamically, and closer to user preference.

If that trajectory holds, product management shifts toward capability design. The core question becomes "what should be possible" within a safe, coherent, and trustable set of constraints.

In that world, the Product Professional becomes even more responsible for:

- Coherence across many possible configurations
- Simplicity despite flexibility
- Guardrails that protect users and the brand
- Trust as an experience, not a compliance checkbox

This increases the need for product leadership.



The **AI PM**, defined is there a future for the AI PM?

Yes, and it looks different from what many expect. The role calls for genuine transformation, well beyond automating PRDs or optimizing for efficiency.

The AI PM of the future is someone who can:

- Use conversational AI to structure thinking, not to replace it
- Maintain hypothesis discipline when solutions are cheap
- Design experiments at the right scale and with the right boundaries
- Prevent context erosion by curating what must remain explicit
- Treat trust, safety, and quality consistency as strategic assets
- Protect patience and coherence in a system that wants to accelerate itself

AI tests the AI PM. And in that test, the winners are the most effective.

For Heads of Product, the question is "what kind of product behavior do we want to amplify." AI will amplify your system, making a learning-driven organization faster at learning, and an output-driven organization faster at producing features.

For Product Professionals and Product Managers, the question is "how do I remain a disciplined, curious, and coherent product thinker in an environment that produces convincing output on demand."

The future is effectiveness under acceleration.



About Stephan van Rooden

With over 15 years of experience in product management, Stephan van Rooden is a recognized authority on product strategy, stakeholder communication, and leadership. He has trained and mentored thousands of professionals, helping them unlock their full potential as product managers and product owners.

s.vanrooden@expandior.nl
+31(0)6 38757584

Expandior Product Academy

Expandior Product Academy helps product professionals and organizations raise their level of product management through structured learning journeys, assessments, and real-world practice.

We believe that great product organizations are built by developing capabilities, not just teaching frameworks. Our programs combine training, coaching, and tools to support continuous professional growth and measurable business impact.

Academy Learning

- Structured learning journeys for product professionals
- Practical training based on real product challenges
- Continuous capability development

GO Product Navigator (GPN)

- Capability assessment for product professionals
- Personalized learning paths
- Insight into strengths and growth areas

Product Management Curriculum

Training modules covering the core capabilities of product management:

- Product Strategy
- Product Discovery
- Data & Decisions
- Customer Centricity
- Go-to-Market
- Leadership,
- Collaboration
- Communication

About the Saasmore Group



Digital product development

GlobalOrange develops (SaaS) products for companies that want to distinguish themselves with their product. Software specifically developed for the requirements of your organisation and your end users.



Development teams and IT Staffing

Amsterdam Standard provides individual nearshore software developers in a variety of languages, nearshore development teams, and extended teams. Flexible and on demand.



Product Management Academy

Expandior helps organisations to raise their level of Product Management. Through a combination of training and coaching we help you to create winning products.



Truly exceptional user experience

YouSir is a SaaS focussed UX Design Studio in Amsterdam. YouSir makes your product feel 'just 4right'. By doing so, they help you retain your users and thus optimise your bottom line.