


State of the art

AI applied to agile work

Reference map of current professional knowledge



Updated October 2026

About this document

This document is a curated map of the knowledge, practices and models in professional use as of its publication date for bringing artificial intelligence into agile work with sound judgment, whatever the role of the person using it. For each entry, the map places it in context, indicates its degree of adoption in current professional practice and points to where to find reference information.

Use it as an observatory and reference point to check whether the professional knowledge you rely on is aligned with current practice and at the leading edge.

It is complemented by an in-depth companion guide that develops each concept and by the training and assessment platform in Skill Arena. In the [AI applied to agile work](#) area you can test your level of knowledge and, if you pass, obtain a diploma that provides curricular accreditation of professional competence and currency in this area.



State of knowledge

Knowledge about AI applied to agile work evolves extremely fast: tools, standards, and legal obligations change within months. The fundamentals (giving the model the right context, asking precisely, verifying what it returns, and protecting information) have settled, and what keeps moving is work with agents and the rules each team uses to govern its use. Throughout the document, the labels (ESTABLISHED, CONSOLIDATING, EMERGING) help identify the maturity of each concept:

ESTABLISHED settled consensus; knowledge taken as required.

CONSOLIDATING gaining adoption fast; not yet universal but already relevant.

EMERGING recent frontier; high relevance and high volatility.

Block A — Fundamentals of interacting with AI

Cross-cutting competencies, valid for any role that works with AI, whatever its discipline: giving the model the right context, asking precisely, verifying what it returns, protecting information, and agreeing as a team on how AI is used. They are the ground on which everything else rests.




Figure 1. The context window: everything that goes in competes for the model's attention, and managing context means deciding what is kept outside, what is brought in, what is summarized, and what is kept apart.

Context engineering · CONSOLIDATING · new since 2026

Deliberately designing what information the model receives at each moment, in what order, and what is left out, beyond how a single instruction is worded. The model works with everything in its context window (system instructions, history, retrieved documents, tool results, and examples), and that space is finite: whatever goes in competes for its attention. The core competency has shifted from writing the perfect instruction (the prompt) to governing the context around it.

Why it is here now. As work with AI has moved from one-off queries to sustained tasks and agents, what sets practitioners apart is deciding what goes into the model's window. The evidence supports this: Chroma's study on context rot (July 2025) measured across 18 models that their performance becomes less reliable as the input grows, even on simple tasks. The techniques are still moving fast, with conversation compaction and memory already turned into product features, which is why the card remains consolidating.

Where to look. [Anthropic — Effective context engineering for AI agents](#) · [Chroma — Context rot](#) · [Scrum Manager BoK — Context engineering \(in Spanish\)](#)

Structured, reusable prompting · ESTABLISHED

Specifying a request with the components that make its result predictable and verifiable: a goal with observable success criteria, the minimum sufficient context, an output format, and explicit constraints. Structural instructions, the ones that set the shape of the response, are what deliver consistency; qualitative ones (“be more thorough”) can

improve the result, but they do not replace the former. A request that works is saved, versioned, and shared as a team asset.

Why it is here now. It is the foundation for everything else: without a well-framed request, verification and work with agents inherit inconsistent results. Vendor guidance has been refined with recent models, which follow instructions more precisely. Anthropic recommends asking for exactly what you want, placing long documents before the question, and toning down emphatic language (“CRITICAL,” “you MUST”), which with current models leads to overcompliance. Today it is understood as the base case of context engineering.

Where to look. [Anthropic — Prompting best practices](#) · [OpenAI — Prompt engineering](#) · [Scrum Manager BoK — Prompt engineering \(in Spanish\)](#)

Verification and reliability of responses · ESTABLISHED

Critically evaluating what a model produces before acting on it: separating facts from inferences and from recommendations, asking for the assumptions and evidence behind a claim, and distrusting specific figures or certainties that arrive without support. Fluent text does not guarantee accuracy, and the verification effort is calibrated to the cost of being wrong. Verifying also requires keeping your own judgment, because whoever delegates the entire reasoning stops learning what they need in order to judge.




Figure 2. Verify before acting: each part of a response calls for a different check, and a figure without a source is treated as suspect.

Why it is here now. It is the most timeless competency in the area, and it gains weight with agents: the more work is delegated, the more critical it is to check what the system delivers. Why it matters has been studied. Kalai, Nachum, Vempala, and Zhang (2025) show that models hallucinate partly because their evaluations reward guessing over admitting they do not know. So has its cost: in an Anthropic experiment (January 2026), those who learned a programming library with AI assistance scored 50% on the comprehension quiz, compared with 67% for those who learned it without it.

Where to look. [Kalai et al. — Why language models hallucinate](#) · [Anthropic — How AI assistance impacts the formation of coding skills](#) · [Scrum Manager BoK — AI hallucination \(in Spanish\)](#)

Trust boundary: instructions versus data, privacy, and secrets ·

ESTABLISHED

Treating all external content as untrusted data by default, because a model does not reliably distinguish the instructions it should follow from the text it should only process, and a command hidden in a document or a web page can hijack it (prompt injection). And applying minimization and effective anonymization before sharing information with a model, without relying on “confidential” labels or on the system's promises.

Why it is here now. Prompt injection remains in first place in the 2026 edition of the OWASP Top 10 for large language model applications (August 2026), which moves excessive agency up to third. With agents acting on real systems, an injected instruction no longer produces a bad answer but an unwanted action, so the defense lies in design: Simon Willison calls the combination of private data, untrusted content, and the ability to communicate externally in a single agent the lethal trifecta, and the prudent course is not to bring all three together.

Where to look. [OWASP — GenAI LLM Top 10 2026](#) · [Simon Willison — The lethal trifecta](#) · [EDPB — Guidelines 02/2026 on anonymisation \(draft under consultation\)](#)

The team's AI usage policy · CONSOLIDATING · new since 2026

Agreeing as a team, explicitly and in a way everyone knows, which AI tools are used, for which tasks, with what data, and with what controls: what is prohibited, what is allowed with guardrails, and what is encouraged. It is one more working agreement, reviewed in the retrospective like any other, and it includes how the team learns to use AI with sound judgment. Without it, each person decides on their own and use goes underground.




Figure 3. Three levels of use: the team's policy states what is prohibited, what is allowed with guardrails, and what is encouraged.

Why it is here now. Shadow AI, the undeclared use of tools nobody has evaluated, grows when it is unclear what is allowed. DORA counts a clear and communicated AI stance among the capabilities that amplify the effect of AI, and IBM's *Cost of a Data Breach Report* (2025) found that 63% of breached organizations had no AI governance policy or were still developing one. The EU AI Act also requires those who deploy AI to take measures to ensure AI literacy among their staff (Article 4).

Where to look. [DORA — Clear and communicated AI stance](#) · [EU AI Act — Article 4: AI literacy](#) · [Scrum Manager BoK — AI governance \(in Spanish\)](#)

Block B — Working with AI agents

The recent paradigm shift: moving from using AI to think better, where it answers, to delegating work to it, where it executes. It is the most fast-moving area of this map and, therefore, the one that is revised most often.




Figure 4. From chat to agent: the assistant answers and the person acts, while the agent executes within set limits and against acceptance criteria.

From chat to agent: delegating verifiable work · CONSOLIDATING · new since 2026

Understanding the difference between an assistant that answers questions and an agent that completes multistep tasks on its own (it searches, decides, uses tools, and delivers a result), and knowing when to delegate and when not to. A good candidate is a task whose result can be checked against a clear criterion and whose errors are acceptable and reversible. The agile professional's role shifts from verifying answers to designing what work is delegated, with what limits, and with what acceptance criteria.

Why it is here now. Agents have moved beyond software development and are now offered for office work in general: in July 2026 OpenAI launched ChatGPT Work, with agents that run complete workflows and produce finished reports and documents. This reshapes agile work, because the unit of work now includes the task that is orchestrated and supervised, and it forces teams to revisit planning, tracking, and the Definition of Done (DoD).

Where to look. [Anthropic — Building effective agents](#) · [Scrum Manager BoK — AI agent \(in Spanish\)](#)

Designing human oversight (human-in-the-loop) · CONSOLIDATING · new since 2026

Deciding where an AI-assisted flow stops so that a person can review, approve, or redirect it before it continues, according to the risk of each action: with the human in the loop (approving beforehand), on the loop (supervising and able to intervene), or out of the loop (reviewing afterward, in aggregate). Irreversible actions or those with access to sensitive systems sit behind an approval gate, and every decision is logged. And whoever supervises needs the information, authority, and time to genuinely influence the outcome.




Figure 5. Three positions of human oversight: the level of oversight is set by the risk of each action, and anything irreversible goes through an approval gate.

Why it is here now. It is the necessary counterpart of delegation, and real cases show what its absence costs: in August 2026 METR investigated how some 1,200 agents that OpenAI was evaluating, and that were supposed to be isolated, coordinated through an unauthorized message board, and how some 700 of them attacked Hugging Face. It also has a legal dimension. The EU AI Act requires effective human oversight of high-risk systems, with obligations that Regulation (EU) 2026/1744 postponed to December 2, 2027, and, for systems embedded in regulated products, to August 2, 2028.

Where to look. [European Commission — AI Omnibus enters into force](#) · [EU AI Act — Article 26: Obligations of deployers of high-risk AI systems](#) · [Scrum Manager BoK — Human-in-the-loop \(in Spanish\)](#)

Agentic ecosystem standards · CONSOLIDATING · new since 2026

Knowing the open standards that are stabilizing work with agents and what each one solves. The Model Context Protocol (MCP) connects models with tools and data; AGENTS.md gives agents reusable instructions from a file in the project; A2A (Agent2Agent) lets agents from different vendors communicate; and Agent Skills packages capabilities that the agent loads only when it needs them. There is no need to implement them, but there is a need to understand where they are taking the ecosystem.

Why it is here now. They have gone from promise to infrastructure under neutral governance. MCP, AGENTS.md and, since August 2026, A2A are part of the Linux Foundation's Agentic AI Foundation, which already brings together almost 250 organizations. The July 2026 MCP specification came with a deprecation schedule, and

its TypeScript and Python SDKs have each passed one billion downloads; AGENTS.md is used in more than 60,000 open-source projects. They are still changing fast, which is why they are consolidating rather than established.

Where to look. [Model Context Protocol — The 2026-07-28 specification](#) · [AGENTS.md](#) · [A2A — A new chapter: joining the Agentic AI Foundation](#) · [Agent Skills](#)

Block C — Application to the agile work cycle

Use cases for AI in the workflow. The first two serve any role on an agile team; the third, marked for development roles, is especially relevant for those who build software.




Figure 6. From a fuzzy story to a refined item: the AI reviews and proposes, and the team decides in the refinement conversation.

AI in the product and backlog flow · ESTABLISHED

Using AI to detect ambiguity in user stories and generate refinement questions, propose slicing options that keep demonstrable value, draft verifiable acceptance criteria with negative scenarios and edge cases, and formulate hypotheses and experiments with metrics and guardrails. It performs better reviewing than drafting: asking it to critique a story written by the team is more reliable than asking it to write one. Its proposals are the starting point for the team's conversation, not requirements.

Why it is here now. It is the most everyday and cross-cutting use of AI in agile, though still a shallow one. In Stefan Wolpers's *AI4Agile Practitioners Report 2026*, with 289 practitioners, 83% use AI tools, but 55% devote 10% of their time or less to them, because they do not yet know where the tools fit into their work. What is settled is the competency, which is taken as necessary for any role that works with the backlog.

Where to look. [Skill Arena — User stories](#) · [Stefan Wolpers — The AI4Agile Practitioners Report 2026](#) · [Scrum Manager BoK — Acceptance criteria \(in Spanish\)](#)

AI in facilitation and team dynamics · CONSOLIDATING · new since 2026

Relying on AI to prepare facilitation, synthesize the material from discovery sessions, and qualitatively analyze the output of retrospectives, without letting it replace human interaction or the reading of the team's mood, which remain the facilitator's job. AI touches facilitation at its edges, the preparation beforehand and the analysis afterward, and leaves its center intact. Recording or transcribing a session also requires informing participants and respecting their data.




Figure 7. AI at the edges of facilitation: it prepares beforehand and synthesizes afterward, while the center of the session remains the interaction between people.

Why it is here now. It is the most useful application for Scrum Masters and agile coaches who do not work with code, and the evidence places it where this map had it. In Wolpers's report, the uses that work are preparing retrospectives, clustering comments, and drafting first versions, and the main concern is losing people-centered collaboration. The boundary also has a legal basis: since February 2025 the EU AI Act has prohibited systems that infer people's emotions in the workplace, except for medical or safety reasons.

Where to look. [AEPD — AI voice transcription \(in Spanish\)](#) · [EU AI Act — Article 5: Prohibited AI practices](#) · [Scrum Manager BoK — Retrospective \(in Spanish\)](#)

AI in quality and delivery (development roles) · ESTABLISHED

Generating and reviewing test cases aligned with acceptance criteria and risks, covering negative scenarios and edge cases beyond the happy path, and using AI as support in code review with a focus on risks: input validation, error handling, permissions, and side effects. Generated code is reviewed with the same rigor as handwritten code, or more, and the final decision on a change still belongs to whoever reviews it.

Why it is here now. It remains a high-value use for technical teams, with two measured caveats. Veracode's *GenAI Code Security Report* (July 2026) found that 44% of generation tasks introduced a vulnerability, almost the same share as a year earlier. And METR observed in 2025 that experienced developers took 19% longer with AI while believing they were going faster; in 2026 it warned that its new design no longer measures the difference reliably, but the gap between perceived and measured speed remains the lesson.

Where to look. [Veracode — 2026 GenAI Code Security Report](#) · [OWASP — Secure Coding with AI Cheat Sheet](#) · [Scrum Manager BoK — AI-assisted QA \(in Spanish\)](#)

A note on the references

The links included were available and verified as of the update date. Given the pace of change in this area, some may change; each revision of the map also updates its references.

© 2026 Scrum Manager®. This work is published under a Creative Commons Attribution-NonCommercial 4.0 International licence (CC BY-NC 4.0). Official Scrum Manager trainers and centres are licensed under the terms of CC BY 4.0 for their training activity.