

The Tandem Manifesto

A working agreement between software engineers and AI agents

Software generation has become a commodity. It is fast, easy, and cheap. Control, quality, reliability, and accountability are the new hard problems.

Software is not progress if no individual can accept accountability for its outcome — regardless of how quickly it was produced.

Solving these problems requires a clear distribution of roles and a clear mode of collaboration between humans and AI agents.

This is what we recognize — and what we commit to.

● WHAT WE RECOGNIZE

1. **The primary bottleneck of software development has shifted:** from creating software to organizing and managing its development.
2. **AI is neither a replacement for developers nor a traditional development tool.** It is a new type of teammate with an exceptional profile — its strengths and its limitations both shape the process. Human and AI form a tandem: they move together, and the human steers.
3. **Result = f(Context).** The quality of the result is a function of the quality of the context. The primary lever is not better prompting. It is better context.

● WHAT WE COMMIT TO

4. **We fill two contexts — and overfill neither.** The machine context determines how closely the output will match the expectations. The human context determines whether the human can understand, explain, and accept accountability for the result. For both, too much context is as harmful as too little.
5. **We treat the human context window as the scarcest resource.** We deliberately organize work around this boundary.
6. **Our agreement binds both sides.** The human provides curated context, clear boundaries, and timely decisions. The AI separates autonomous execution from conscious decision-making: when a relevant gap appears, it stops and asks — it never fills the gap with a silent assumption.
7. **We make important decisions explicit.** Assumptions, specifications, and changes are documented before execution continues. Nothing changes silently.

8. **We synchronize expected outcome, strategic goals, and target architecture** within the tandem before implementation begins — and re-synchronize as the implementation evolves.
9. **We re-synchronize when it is needed** — not because the calendar requires it.
10. **We shape the solution approach together — human and AI** — in a tactical specification before implementation begins. This is where the mental model of the solution takes root in the human mind. It is this model that makes review fast and accountability real.
11. **We do not ship what no human has reviewed and approved** — but human review does not mean reading every line. An independent peer AI inspects the implementation in depth, and the implementing AI resolves or justifies every finding. The human verifies that the agreed approach was followed.
12. **We scale review depth to criticality.** Where software is mission-critical — or regulation demands it — human review deepens, down to every single line. The required depth is an explicit decision, made during specification — never an afterthought.
13. **We specify governance and security before implementation begins** and validate them continuously as the solution evolves. “Not required” is an explicit decision, not an omission.
14. **We commit to principles, not products.** No AI model, tool, vendor, or jurisdiction is prescribed — the mode of collaboration is what endures.

**We move together.
The human steers.
Nothing changes silently.**

● SIGN THE MANIFESTO

If you believe that software built with AI must remain software that humans can understand, control, and stand behind — say so publicly.

There is no form to fill in. Post that you have signed the Tandem Manifesto, under your own name, and tag it **#TandemManifesto**.

Initiated by Anastasia Galani & Ingo Rube · 2026