AI News

Foreword

2026-05-02 - 14:44

Compiled with Claude AI

1 Introduction

By now, almost anyone paying attention has noticed AI's explosive evolution and its rapid integration into daily life. What was experimental two years ago is now part of how email gets answered, how documents get drafted, how code gets written, and how decisions get framed. It is, on any reasonable reading, a reshaping of how people organise their work — and a reshaping that organisations are forced to address whether they are ready or not.

That is not a calm thing to live through. Companies are racing to adapt. So are individuals, governments, and professions. The question is no longer whether AI changes how work gets done. The question is how that change is being managed, and at what cost.

These pages collect monthly observations of the AI ecosystem — the models, the products, the funding, the regulation. This Foreword exists for a different purpose: to step back from the data and reflect on where the change actually feels uneasy. There are three places where the unease keeps returning.

2 The reality on the ground

Large organisations are caught between two pressures. On one side, the fear of missing the departure — of becoming disconnected from how customers actually want to work and what they will increasingly expect. On the other side, the practical difficulty of bringing AI inside the organisation responsibly. Adapting AI across an entire company — its data, its tooling, its workflows, its compliance posture — is not a quick exercise. It carries real risks.

The most significant risk is one that few executive presentations highlight clearly. Employees who use AI well can produce work of noticeably higher quality, and at noticeably higher speed, than they could a year ago. That is the upside, and it is real. The downside is that producing higher-quality output and verifying higher-quality output are different skills. Many employees are not yet trained on how to use AI for routine tasks, and even fewer are trained on how to cross-check, validate, or sanity-test what AI produces before it leaves the building.

A polished document that says something subtly wrong is harder to catch than an obviously rough draft that says something wrong. The risk is not that employees produce worse work — they will not. The risk is that they produce smoother work that requires the same scrutiny as before, and they have not yet built the muscle to apply that scrutiny at speed. The output looks finished. The reading required to be confident in it has not got any shorter.

3 Software, where the dynamic is sharper

In software development the same dynamic plays out in concentrated form. In most organisations, building and shipping software is a slow, careful process. The cycle from specification to production passes through reviews, tests, validations, and approvals that exist for good reasons. AI-assisted coding can compress that cycle by several gears at once.

The benefit is real. Engineers who learn to work with current coding agents can move much faster on routine work, and the work is often genuinely better than it would otherwise have been. The cost is that the surrounding controls — review queues, test coverage, staging environments, security audits — are calibrated for the old velocity. Pull requests arrive at code review faster than reviewers can keep up. Tests cover the cases the AI was trained on and miss the edge cases that humans used to catch by hand. Security review becomes a bottleneck where it used to be a step in the flow.

The teams responsible for testing and validating output are rarely serious enough about the new failure modes to catch them in time. This is not a criticism of those teams. It is a description of what happens when one part of the pipeline accelerates while the rest does not. The result is software that ships faster and that breaks in less expected ways. The mistakes are not bigger than before. They are differently shaped, and the existing controls were not designed to find that shape.

4 The path forward

The direction of travel is clear. AI will reshape the world of work. The pace is what is open to debate.

The popular discourse oscillates between two extremes. One says everything changes within five years; the other says AI is overhyped and will plateau soon. The data in the monthly Trends sits between those poles. Capability is advancing meaningfully every quarter. So is enterprise adoption. So is the gap between what AI can demonstrate on a *benchmark* and what it can deliver reliably in production.

The honest framing is that AI today still makes too many mistakes for the kind of unsupervised deployment that some product roadmaps assume. We are not yet at a point where fully correct results are routine on hard tasks. We are at a point where mostly correct results are routine on tasks that previously required full human effort — which is a different change, and an equally consequential one.

The path forward, for organisations and for individuals, is not to bet on either extreme. It is to build the working practice that lets AI’s strengths show up in production while keeping its mistakes contained. That practice is harder to build than to describe — and it is what the next few years will be about. These pages will track the change month by month.