The Next Era of Knowledge Work

How Codex is helping people navigate the complexity of modern work

June 2, 2026

5M	**>6x**	**>3x**
Codex weekly active users	Growth in Codex weekly active users since the desktop app launched	Knowledge workers are now adopting Codex >3x faster than developers

Data Analysis, Research, Knowledge Artifacts

Codex's fastest-growth task types for knowledge workers

Knowledge Work in the Modern World

Knowledge work dominates advanced economies. More than 40 percent of US labor, roughly 72 million people, works primarily with information: analysis, code, documents, designs, systems, decisions, and communication. This is a recent historical development. In 1850, approximately 60 percent of Americans worked in agriculture. By 1970, that had fallen to about 4 percent even as output soared. Manufacturing absorbed much of the displaced labor and peaked around 26 percent of US employment in 1960, then receded with globalization and robotics.

Knowledge work filled the space. Peter Drucker coined the term in 1959 to describe occupations whose primary inputs and outputs are symbols rather than crops or manufactured goods. The category now includes software, engineering, science, management, education, healthcare, law, finance, design, marketing, journalism, consulting, and much of public administration.

The result is a strange abundance. Modern workers can produce documents, messages, dashboards, models, and presentations faster than ever. Yet they spend a remarkable share of their time looking for context, reconciling conflicting versions, waiting for responses, and moving information across systems. A McKinsey Global Institute study found that the average knowledge worker spends roughly 28 percent of the workweek managing email and nearly 20 percent of it looking for internal information or tracking down colleagues who can help with specific tasks.

This is the inheritance of piecemeal digitization. Email made correspondence instantaneous. Documents became collaborative. Spreadsheets, dashboards, ticketing systems, CRMs, knowledge bases, chat tools, and SaaS applications captured more work. Each solved a local problem but gave us fragmented tools. The information we need now hides in inboxes, folders, chats, comments, permissions, dashboards, tickets, meetings, and our own memory. The office became digital without becoming integrated.

Three frictions now define the daily cost of knowledge work:

- First, search is the cost of finding the relevant inputs across sprawling, untransparent systems: the right file, clause, file path, precedent, dataset, message, or expert are all needles in obscure haystacks.
- Second, coordination is the cost of moving information and decisions through teams, tools and formats, while navigating an organization's divergent and shifting incentives.
- And third, approval and verification are the costs of getting work accepted and ensuring that it survives contact with reality. In engineering, that means tests, reviews, deployments, and ongoing monitoring. In law or consulting, it means partner review, client acceptance, and defensible reasoning. In science, it requires experiments, replication, and evidence.

These frictions explain why computers invaded the office long before they increased its productivity. Robert Solow captured this when he observed that the computer age was visible everywhere except in the productivity statistics. Erik Brynjolfsson later gave that puzzle a name, the productivity paradox, and developed the deeper theory that information technology produces large gains only when organizations redesign processes, skills, management structures, and workflows around the tasks that the technology has made cheap. Electricity followed the same pattern. Large gains arrived only after factories stopped treating electric power as a substitute for steam engines and rebuilt floorplans around distributed electric motors at each machine. That took decades.

Knowledge work is still waiting for its factory redesign. Previous generations of workplace software lowered the cost of producing intermediate artifacts, but did not reduce the attention required to consume them. Email made correspondence cheap, then multiplied correspondence. Docs made drafting cheap, then multiplied drafts and review cycles. The result is an excess of documents and tools, and ever scarcer time and attention.

Codex is that redesign. Factories put electric motors next to each machine, and Codex places AI closer to each problem to be solved. Large organizations were built around the heavy cost of producing and moving artifacts: secretary pools, cross-functional teams, pyramids of clerks, long chains of review. Codex puts more agency in the hands of the people with the greatest need. It dissolves the bottlenecks before, during and after artifact production. It can find the inputs, coordinate the workflow, produce the deliverables, check their quality and chase down the necessary approvals. Codex brings a new way of work to teams in fragmented digital organizations. Companies can now build faster, with fewer delays and administrative burden, so their people can focus on the most impactful tasks.

GroundVue – founded by Travis Hoppe, Ann Lewis, and Shannon Arvizu – offers a glimpse of how Codex is expanding beyond software into knowledge work. The company helps governments learn from one another by making public meetings searchable and comparable at scale. Across roughly 90,000 government bodies, critical information is scattered across videos, websites, and local platforms. GroundVue uses Codex to find hard-to-reach public sources and build systems that continuously collect and organize them. Tasks that once took days or weeks now take minutes, allowing a small team to perform work that previously would have required large groups of technologists and researchers. In effect, Codex helps turn fragmented public information into usable institutional knowledge.

Beyond Software: How Agentic Tools Untangle Knowledge Work

Codex now has more than 5 million weekly active users, up more than 6x since the launch of the desktop app in February. While it began as a tool for software development, faster growth is increasingly coming from a broader category: knowledge work.

Data scientists use Codex to clean datasets, build models, and automate analysis. Finance teams use it to reconcile information, build forecasts, and generate reports. Designers use it to prototype products and create assets. Marketers use it to analyze campaigns, produce content, and synthesize customer feedback.

Knowledge workers now represent about 20 percent of Codex users and are adopting it more than 3 times as fast as developers. This includes roles that span product and project management, design, research and academia. Personal users represent more than 5 percent of Codex users and are growing more than 4 times as fast as developers, with substantial use in hobbies and creative work, education and self-learning, personal finance, and entertainment.

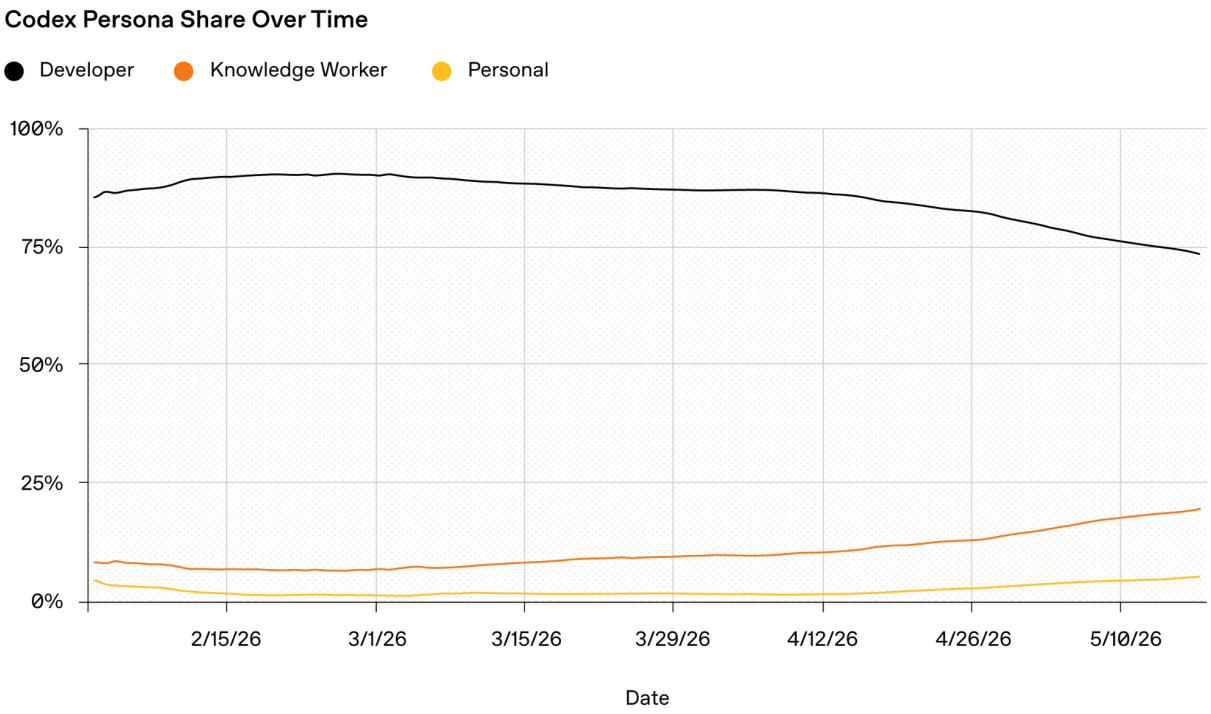

The task mix among knowledge workers is the most revealing piece of the data. Each week, 72 percent of these users produce artifacts: text documents such as reports, memos, and contracts; multimedia assets such as images, audio, and video; and, increasingly, PDFs and spreadsheets. The next categories are less obvious: engineering operations at 47 percent, code implementation at 46 percent, and research at 41 percent. The boundary between software work and knowledge work has blurred, because AI enables people to reach beyond their formal role and build what their goal requires. Developers use Codex for knowledge artifacts; knowledge workers use it for code and engineering operations. A product manager builds the dashboard rather than commissioning it from another team. A researcher writes the dataset-cleaning script herself instead of routing it through engineering. A designer ships the prototype without a developer in the loop. An executive constructs an internal tool that reconciles files and produces a weekly report. Codex gives the person closest to the work a way to build the missing tool without waiting for a formal software roadmap to authorize and deliver it.

Knowledge Worker Weekly Task Penetration

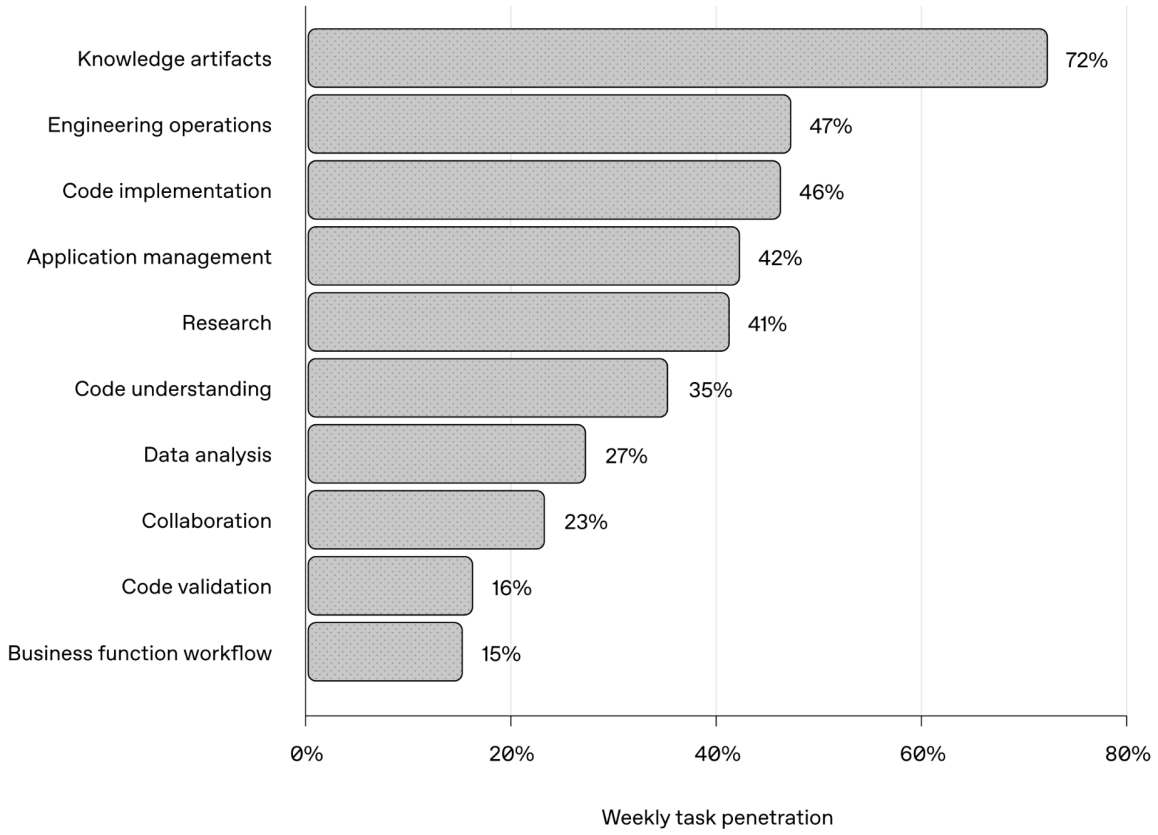

The most consequential shift in behavior that our data show is toward parallel tasks. Roughly 50 percent of users now have more than one Codex task running simultaneously at some point during the day, up from less than one third in mid-April. The shift from sequential to parallel use is what lets a single knowledge worker operate at the scale of a small team: one turn to inspect a dataset, another to draft a script, another to assemble a report, another to check an application. The user becomes the orchestrator of workstreams rather than executing a single task at a time.

Modern knowledge work is beset by imperfect tools, scattered context, time-consuming handoffs, and artifacts that must be assembled, revised, and approved. Codex gives people a way to operate efficiently inside that fragmentation, to turn scattered, local knowledge into shared, working systems, and converting individual effort into parallel, reviewable streams of work.

Proaction shows how Codex helps small teams perform work that once required larger organizations. The company helps fleets manage vehicles and equipment whose data is scattered across telematics systems, maintenance platforms, spreadsheets, and institutional memory. Using Codex, co-founder Colin Knudsen turns customer conversations into customized proposals, workflow prototypes, and working demos tailored to each prospect's operations. Instead of relying on generic sales pitches, Proaction can quickly build and validate solutions before a contract is signed. Codex effectively connects customer discovery, sales, and product development, helping a five-person startup move faster and compete well above its size.

Where Codex Is Growing Fastest

The fastest growing task types for knowledge workers are: Data Analysis (110% growth week over week); Research (+37%), and Knowledge Artifacts (+36%):

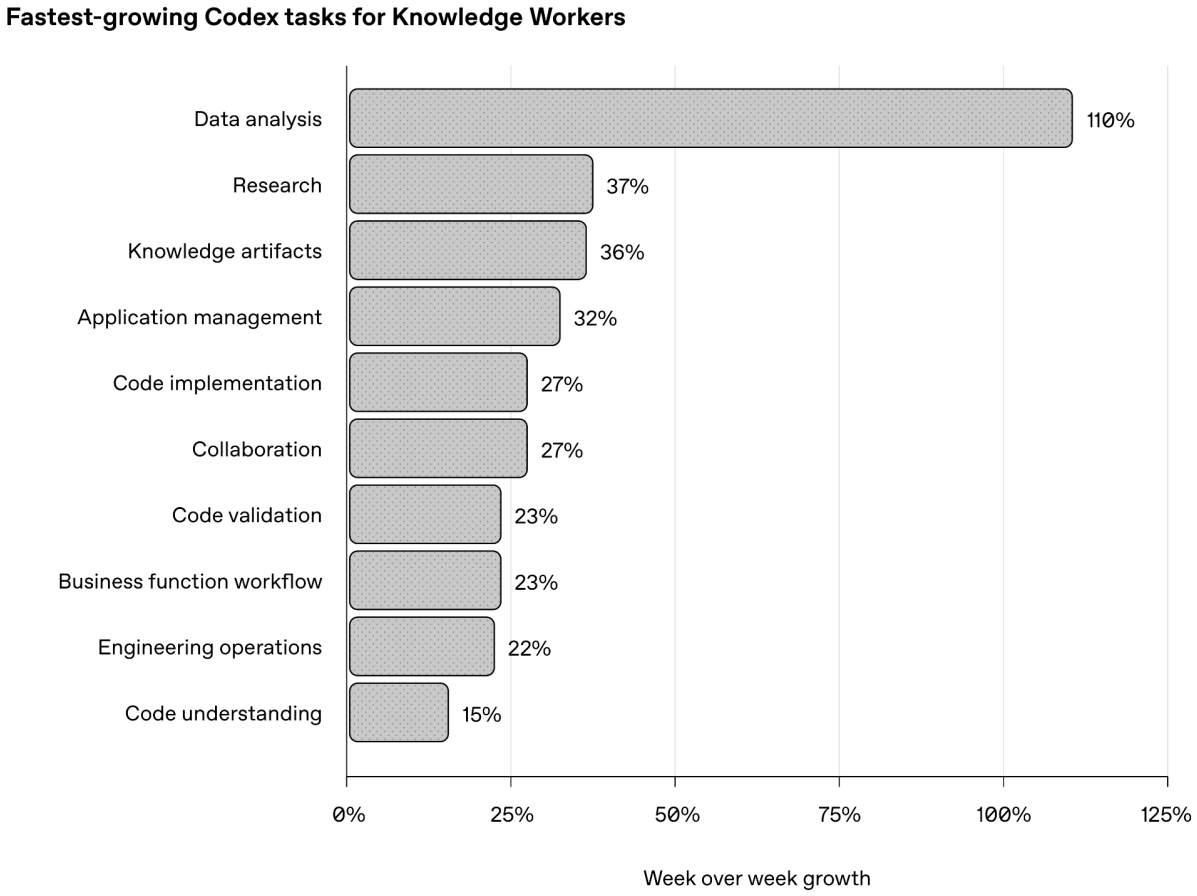

Among the fastest growing task types, a few are notable: within Knowledge Artifacts, users are most commonly working with text documents (Google Docs, Word docs, reports, memos, contracts) and multimedia assets (images, audio, and video). We see more than 50 percent growth in users working with PDFs and spreadsheets. In Research, we commonly see searches of the web and internal knowledge, with substantial growth in market research: into companies, industries, competitors, market size and position. Data labeling dominates Data Analysis, with the majority of usage and the fastest growth rate among tasks. Finally, we see more than 40 percent growth in a variety of other tasks: drafting messages, building and designing products, understanding contracts, regulations, and policies, and hiring and interviewing.

Mathematics professor Taiyo Inoue uses Codex to automate one of the least rewarding and most time-consuming parts of teaching: maintaining course information in a learning management system. By helping him generate scripts that update assignments, calendars, materials, and announcements in Canvas, Codex handles work that previously consumed hours of manual effort each week. Inoue estimates the workflow saves him four to five hours weekly. Instead of spending that time on administrative upkeep, he has redesigned his classes around collaborative problem-solving, giving students at California State University more opportunities to work through mathematics together and in person. Codex helps turn routine academic administration into time and energy that Inoue reinvests in teaching.

Policy for the Agentic Era

Countries and organizations that give people access to these tools and teach them to build and delegate responsibly will see the most productivity gains. Codex is an early example of that future: not a replacement for knowledge workers, but a way for people closest to the work to operate with more agency and capacity to solve hard problems. The points below are intended to start a conversation with policymakers: as agentic tools move from software into the broader work of public agencies, schools, startups, nonprofits, research labs, and business.

1. Modernize workflows and measure what works

Public agencies should use agents to reduce administrative backlogs, improve software systems, search and reconcile records, support scientific research, and deliver public services faster. Success should be measured in outcomes that serve people and they easily understand: shorter wait times, fewer forms, faster permitting, better benefits delivery, lower administrative costs, and more responsive government. The benefits of agentic AI should not be limited to the largest firms and policy should support broad access for schools, startups, small businesses, nonprofits, researchers, and government.

2. Make AI fluency a core workforce skill

Governments should treat AI fluency as basic economic infrastructure. That means funding hands-on AI training through schools, community colleges, public agencies, libraries, and employer partnerships. Employers should be encouraged to support training through grants, tax incentives, procurement, and partnerships to give workers paid time to learn these tools on real workflows.

3. Put workers at the center of AI adoption

The people closest to the work often know best where AI can make a job better, safer, and more productive. A nurse knows which forms end up slowing care, a caseworker knows where benefits delivery breaks down, a machinist knows which handoffs lead to errors, and a teacher knows how much administrative work reduces their time with students. Governments should support worker-led AI adoption through small-business grants, public-sector innovation funds, technical assistance, and formal channels for workers and managers to shape how AI is used. The highest-return uses will often be local and specific: a small tool for a clinic, a lab, a local government office, a small manufacturer, or a nonprofit that helps people remove dangerous, repetitive, administrative, or exhausting work and focus on higher-value tasks.

4. Update public procurement around jobs to be done

Governments should update procurement so agencies can buy AI tools that solve operational problems, not merely purchase software licenses. Public-sector pilots should require privacy, security, auditability, and human oversight. Agencies should also be empowered to test tools in secure sandboxes and scale what measurably reduces backlogs, improves quality, or lowers administrative burden.

Luke Xing used Codex to build a desktop app that helps compensate for major and variable hearing loss in his left ear. By describing the problem in plain language to Codex, he created a tool that tests hearing across frequencies and adjusts audio output for different devices, helping restore balance to music, calls, and everyday listening. The app is not a medical device, but a personal solution to a highly specific challenge that commercial software has not addressed. Luke's experience illustrates a broader shift: people no longer need to wait for someone else to build the tools they need. Increasingly, they can build those tools themselves.