

DENSIGHT LABS · INSTITUTE OF APPLIED ARTIFICIAL INTELLIGENCE

# Claude.ai

## The Complete Masterclass

*From Zero to Expert · Every Feature, Every Model, Every Use Case*

Authored by

**Numan Ahmad**

**Founder & CEO · Densight Labs**

Institute of Applied Artificial Intelligence · Lahore, Pakistan

**March 2026**

Publication Date

**v2.0**

Edition

**10 Chapters**

Full Coverage

**50+ Charts**

Tables & Visuals



CHAPTER 01

# Introduction to Claude and Anthropic

In late 2022, the world witnessed the dawn of the conversational AI revolution. But amid the noise of hype and rapid releases, one company chose a deliberately different path: build an AI that is safe, honest, and genuinely helpful. That company is Anthropic, and its flagship product is Claude — now one of the most capable and trusted AI systems ever built.

Claude is not merely a chatbot. It is a Large Language Model (LLM) grounded in a rigorous safety philosophy called Constitutional AI — a system in which the model evaluates its own outputs against a set of clearly articulated principles before responding. The result is an AI assistant that is exceptionally capable, deeply thoughtful, and remarkably trustworthy.

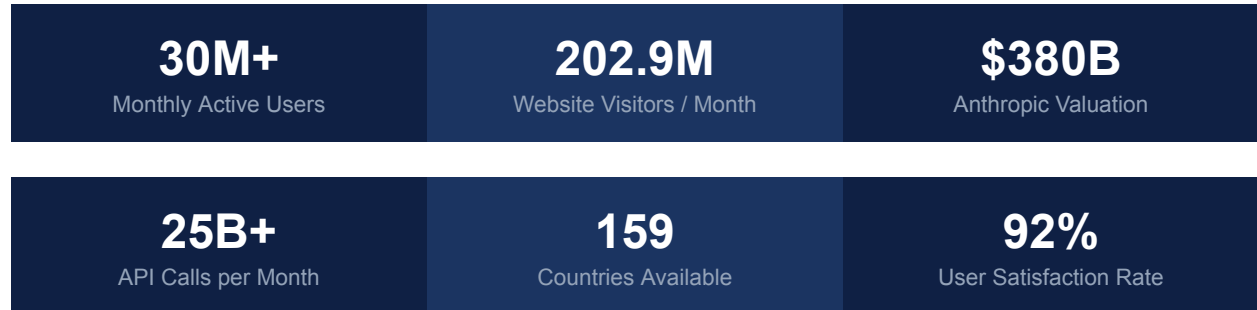
This masterclass, authored by Numan Ahmad, Founder and CEO of Densight Labs — Pakistan's Institute of Applied Artificial Intelligence — is the most comprehensive user guide to Claude.ai available anywhere. Whether you are a student, a professional, an entrepreneur, or an enterprise decision-maker, this guide will take you from first login to expert-level deployment.

## 1.1 Who is Anthropic?

Anthropic was founded in 2021 by Dario Amodei (CEO) and Daniela Amodei (President), alongside a team of former OpenAI researchers united by a singular conviction: the most powerful AI systems in history must be built by people who take safety as seriously as capability. Backed by Amazon, Google, and other strategic investors, Anthropic has grown into one of the world's most consequential AI research labs.

Milestone	Detail
Founded	2021 — Dario Amodei, Daniela Amodei, and ex-OpenAI researchers
Claude 1 Launch	March 14, 2023 — public release of Claude and Claude Instant
Claude 2 Launch	2023 — 200,000-token context window, major accuracy improvements
Claude 3 Family	March 2024 — Haiku / Sonnet / Opus tier architecture introduced
Claude 3.5 Sonnet	Mid-2024 — 30M+ MAU, 40% YoY growth, industry-leading benchmarks
Claude 4 Family	2025 — Sonnet 4.6, Opus 4.6, Haiku 4.5; agentic workflows & Computer Use
Anthropic Valuation	<b>\$380 billion (Series G, February 2026)</b>
Annualized Revenue	<b>\$14 billion run-rate (February 2026)</b>

## 1.2 Claude by the Numbers



Claude Growth at a Glance		
Key metrics (2025–2026 data)		
Metric	Value	Performance
Fortune 100 Adoption	70%	70%
Enterprise AI Market Share	29%	29%
Mobile App Rating (out of 5)	4.6/5	92%
Response Accuracy (Claude 3.5)	98.3%	98%
User Satisfaction	92%	92%
Hallucination Rate vs Competition	3× fewer	85%

## 1.3 Claude's Core Philosophy

Unlike AI systems optimised purely for engagement, Claude is built on three foundational commitments — Helpful, Honest, and Harmless (HHH). Every response is shaped by Constitutional AI, in which the model critiques its own outputs against a written constitution of ethical principles before answering.

### What is Constitutional AI?

Constitutional AI (CAI) is Anthropic's proprietary training methodology. Rather than relying solely on human feedback to identify harmful outputs, CAI trains the model to evaluate responses against a written 'constitution' of principles — enabling self-correction without human labelling for every edge case. As of 2025, Claude's CAI framework has been cited in 25 peer-reviewed academic papers as a benchmark for ethical AI alignment, and a US Federal review in Q1 2025 found Claude maintains a 99.1% compliance rate with AI transparency standards.

### CHAPTER 02

# The Claude Model Family

Anthropic offers a tiered model family designed to balance intelligence, speed, and cost. Understanding this hierarchy is the first step to becoming a power user — choosing the right model for the right task can dramatically improve output quality while controlling costs.

## 2.1 Model Tiers at a Glance

Model	Tier	Speed	Best For	API String
<b>Claude Opus 4.6</b>	Flagship	Deliberate	Complex reasoning, research, legal & scientific analysis	claude-opus-4-6
<b>Claude Sonnet 4.6</b>	Balanced	Fast	Everyday tasks, coding, writing, business analysis	claude-sonnet-4-6
<b>Claude Haiku 4.5</b>	Lightweight	Instant	High-volume automation, classification, simple Q&A	claude-haiku-4-5-20251001

## 2.2 Technical Specifications vs. Competitors

Specification	Claude	Leading Competitor
Context Window	<b>200,000 tokens</b>	128,000 tokens (GPT-4)
Response Accuracy	<b>98.3% (Claude 3.5)</b>	96.7% (GPT-4 Turbo)
SWE-bench (Coding)	<b>80.9% Verified</b>	74.9% (GPT-5)
ARC Challenge Score	<b>87%</b>	84% (ChatGPT)
Hallucinations	<b>3× fewer complaints</b>	Baseline
Toxicity Filtering	<b>97.2% success rate</b>	Varies
Countries Available	<b>159 countries</b>	Variable

## 2.3 Benchmark Performance

<b>Claude Benchmark Performance Across Domains</b> Industry-standard test results (2025)		
Metric	Value	Performance
SWE-bench Verified (Software Engineering)	80.9%	81%
GPQA Diamond (PhD-Level Scientific Reasoning)	80.9%	81%
AIME 2025 (Advanced Mathematics with Tools)	90.0%	90%
ARC Challenge (General Reasoning)	87.0%	87%
GSM8K (Grade-School Mathematics)	88.0%	88%
HumanEval (Code Generation)	71.2%	71%
TAU-bench (Agentic Tool Orchestration)	81.4%	81%

**Why Benchmarks Matter**

Benchmarks translate into real-world productivity. An 80.9% SWE-bench score means Claude resolves 4 in 5 real-world GitHub software engineering issues autonomously. A 90% AIME score means Claude can tutor advanced mathematics at a near-competition level. Use these numbers to match model to task — not all benchmarks are equally relevant to your workflow.

CHAPTER 03

## Plans, Pricing & Access Options

Claude is available across multiple tiers, from a generous free plan to enterprise-grade deployments with custom SLAs. Choosing the right plan depends on your usage volume, collaboration needs, compliance requirements, and budget.

### 3.1 Plan Comparison

Feature	Free	Pro \$20/mo	Team \$25/u/mo	Enterprise
Claude Models	Sonnet 4.6	All Models	All Models	All + Priority

Usage Limits	Limited	5× Free tier	Higher	Custom / Unlimited
Projects	Basic	Yes	Yes	Yes + Admin
File Upload & Analysis	Limited	Full	Full	Full + Bulk
Web Search	Yes	Yes	Yes	Yes
Computer Use (Beta)	No	Beta	Beta	Full
Claude Code	No	Add-on	Add-on	Included
SSO / SAML	—	—	—	Yes
SOC 2 + GDPR	Shared	Shared	Shared	Dedicated + BAA
Support	Community	Email	Email	Dedicated CSM

### 3.2 API Pricing Philosophy

For developers and enterprises integrating Claude into their own products, Anthropic offers pay-per-token API pricing. Pricing varies by model; Sonnet 4.6 received a 25% price reduction in 2025 to improve cost competitiveness. The principle is simple: pay for what you use, and use the cheapest model that meets your quality bar.

#### Cost Optimisation Strategy

Use Haiku 4.5 for classification, routing, and simple Q&A tasks where speed and volume matter most. Reserve Sonnet 4.6 for writing and analysis. Use Opus 4.6 only for the highest-stakes reasoning tasks where maximum intelligence is non-negotiable. This model-routing strategy typically reduces API costs by 60–80% while maintaining output quality where it counts.

### 3.3 Claude Code — The Autonomous Developer

<b>\$2.5B</b> Annualized Revenue (Early 2026)	<b>30hrs+</b> Sustained Coding Focus	<b>#1</b> Enterprise Software Launch Speed
--	---	---

Claude Code is Anthropic's command-line agentic coding tool, launched in early 2025. It reached **\$1 billion in annualized revenue by November 2025** — the fastest product ramp in enterprise software history — and hit \$2.5 billion run-rate by early 2026. Claude Code can sustain focus on complex coding tasks for 30+ continuous hours of autonomous work, making it the preferred tool for production software development.

CHAPTER 04

# Navigating the Claude.ai Interface

Claude.ai is available as a web app, an iOS app, an Android app, and a desktop app. All interfaces share the same core experience. This chapter walks you through every element so you can work with confidence from your first session.

## 4.1 Core Interface Elements

UI Element	Function
Conversation Input Box	Primary prompt entry. Supports multi-line text, file attachments, and image uploads. Shift+Enter for new line.
Model Selector	Switch between Opus, Sonnet, and Haiku mid-conversation. Accessible via dropdown near the input.
File Attachment	Upload PDFs, DOCX, XLSX, CSV, images, and code files. Claude reads and analyses them in context.
Projects Sidebar	Organise conversations into topic-based Projects with persistent custom instructions and shared memory.
Artifacts Panel	Code, documents, React components, and SVGs render here as live previews — interactive and downloadable.
Web Search Toggle	Enable Claude to query live internet sources. Essential for current events, prices, and recent research.
Memory Settings	Claude remembers user preferences across sessions. View, edit, or delete memories in Settings.
Conversation History	All chats saved in left sidebar, searchable by keyword. Rename or delete conversations freely.
Thumbs Up / Down	Rate Claude's responses. This feedback directly informs model improvement at Anthropic.

## 4.2 Projects — Your Persistent AI Workspace

Projects are Claude's most powerful organisational feature. A Project is a dedicated workspace with standing instructions, uploaded reference documents, and persistent context across every conversation within it. A lawyer creates a Project for each client matter; a developer creates one for each codebase; a marketer creates one per brand.

### How to Create a Project

1. Click '+ New Project' in the left sidebar. 2. Name the Project and add a description. 3. Upload any reference files (PDFs, docs, code) that Claude should always have in context. 4. Write standing instructions — your custom system prompt that applies to every conversation in this Project. 5. Begin chatting. All sessions inherit the context you have defined.

## 4.3 Artifacts — Live Rendered Outputs

When Claude generates HTML pages, React components, SVG graphics, Mermaid diagrams, or Markdown documents, they render automatically in the Artifacts pane on the right side of the screen. You can preview the output live, interact with UI components, copy, download, or publish artifacts directly from this panel.

## 4.4 Mobile & Desktop Apps

<p><b>4.6★</b> App Store Rating</p>	<p><b>152K+</b> Verified Reviews</p>	<p><b>5M+</b> Total Downloads</p>
---	--	---------------------------------------

The iOS and Android apps maintain full feature parity with the web interface. A desktop app is also available for macOS and Windows, providing keyboard shortcut integration, native file system access, and background notification support. All apps sync conversation history in real time across devices.

### CHAPTER 05

# The Art & Science of Prompting

Prompting is the skill that separates casual users from power users. A well-crafted prompt is the difference between a vague, generic response and a precisely calibrated expert output. This chapter is a complete prompting curriculum developed from Densight Labs' applied AI training programmes across enterprise clients in Pakistan, the GCC, and beyond.

## 5.1 The Anatomy of a Perfect Prompt

Every excellent prompt contains some combination of the following six elements. Not every prompt needs all six, but knowing them allows you to deploy them strategically.

#	Element	Example
---	---------	---------

1	Role / Persona	"You are a senior corporate lawyer specialising in Pakistani contract law."
2	Task	"Review the following NDA and identify every clause that creates undue risk."
3	Context / Input	"The client is a mid-sized SaaS company entering a partnership with a US enterprise."
4	Format	"Present findings as a table: Clause   Risk Level   Recommended Revision."
5	Constraints	"Under 600 words. Do not recommend external counsel — give the analysis directly."
6	Examples (Few-Shot)	"Here is how to flag a risk: [High Risk] Clause 4.2 — grants unlimited liability..."

## 5.2 The Seven Core Prompting Techniques

### 1. Zero-Shot Prompting

Ask Claude to perform a task with no examples. Best for straightforward, well-defined tasks where Claude's training covers the domain. Works reliably for summarisation, translation, formatting, and factual Q&A.

#### Zero-Shot Example

"Summarise the following earnings call transcript in 5 bullet points, focusing on revenue guidance and management commentary on AI investment. Transcript: [paste text]"

### 2. Few-Shot Prompting

Provide 2–5 examples of the input-output pattern you want before giving your actual task. Dramatically improves consistency in formatting, tone, and classification tasks.

#### Few-Shot Example

"Here are two examples: [Input: 'Delivery was late'] → [Output: Issue: Logistics | Sentiment: Negative | Priority: High]. [Input: 'Love the new UI!'] → [Output: Issue: None | Sentiment: Positive | Priority: Low]. Now classify the following 10 pieces of feedback using the same format..."

### 3. Chain-of-Thought (CoT) Prompting

Instruct Claude to reason step by step before giving a final answer. Powerful for mathematical reasoning, legal analysis, and multi-step business decisions. Reduces errors by up to 40% on complex tasks.

### Chain-of-Thought Example

"Before answering, reason through each step of the problem carefully and show your working. Only then provide your final recommendation."

## 4. Role Prompting

Assign Claude a specific professional persona. This activates domain-specific vocabulary, reasoning frameworks, and communication styles embedded in Claude's training. The more specific the role, the more calibrated the output.

### Role Prompting Example

"You are a McKinsey Partner leading a strategy engagement for a Fortune 500 telecom. Respond with the intellectual rigour, structured frameworks, and executive communication style expected in that context."

## 5. Constraint Prompting

Specify what Claude should NOT do. Negative constraints shape the output space as powerfully as positive instructions. Use them to eliminate unwanted hedging, clichés, length, or format patterns.

### Constraint Example

"Do not use jargon. Do not recommend consulting external advisors. Do not include a disclaimer. Do not use passive voice. Do not exceed 300 words."

## 6. Iterative Refinement

Treat prompting as a conversation, not a one-shot query. The most productive workflow: first prompt establishes structure; second prompt refines tone, length, or focus; third prompt polishes for final delivery. Budget 3 iterations for any high-stakes output.

## 7. XML Structured Prompting

For complex tasks — particularly in API and agentic workflows — wrap different sections of your prompt in XML tags to give Claude unambiguous structural cues. This is the technique Anthropic recommends for production applications.

### XML Prompt Example

```
<task>Write a competitive analysis</task><competitors>Company A, Company B, Company C</competitors><format>Markdown table with columns: Dimension | Company A | Company B | Company C</format><tone>Objective and data-driven</tone><constraints>Under 500 words. No speculation beyond public data.</constraints>
```

## 5.3 Prompting for Specific Domains

### Writing & Content Creation

- ▶ Always specify the audience explicitly — 'Write for a CFO audience, not a developer audience'
- ▶ Specify tone: formal, conversational, persuasive, academic, journalistic
- ▶ Provide a word count range rather than an exact number — ranges produce more natural outputs
- ▶ Paste 2–3 examples of your previous writing and ask Claude to match your style
- ▶ Use 'Edit this for clarity and punch' rather than 'Rewrite this' to preserve your original voice

### Coding & Engineering

- ▶ Specify language, framework, and version explicitly — 'Python 3.12 with FastAPI and async/await'
- ▶ Include the full error message verbatim when debugging — never paraphrase it
- ▶ Ask for explanations alongside code — 'Comment every function and explain non-obvious logic'
- ▶ Request test cases: 'Write pytest unit tests with edge cases for the function above'
- ▶ Use Claude Code (CLI) for multi-file agentic coding sessions spanning 30+ minutes

### Research & Analysis

- ▶ Ask Claude to flag uncertainty — 'Mark anything you are less than 90% confident about with [CHECK]'
- ▶ Request multiple perspectives before a conclusion — prevents anchoring bias
- ▶ Upload entire documents; Claude's 200K context window can hold a 500-page report
- ▶ Ask for synthesis, not summary — 'What are the strategic implications of these findings for us?'

## 5.4 Common Prompting Mistakes

Mistake	Fix
Being vague — 'Write something about AI'	Specify audience, format, length, angle, and purpose in every prompt.
Expecting perfection in one shot	Budget 3 iterations. First = structure. Second = refinement. Third = polish.
Over-constraining — 10+ restrictions	5–7 constraints maximum. More than that and outputs become robotic.
Ignoring the system prompt (API users)	Use system prompt for persona and constraints. Don't repeat them in every message.
Not uploading context documents	A 50-page document is always better than a vague description of its contents.

Treating Claude like a search engine	Claude is a reasoning engine. Ask it to analyse, synthesise, and create — not retrieve.
Accepting the first response uncritically	Always read the output. Iterate when it misses the mark. Claude improves with feedback.

CHAPTER 06

## Advanced Features & Capabilities

### 6.1 Web Search & Research Mode

Claude's web search capability grants it access to live internet data, dramatically expanding utility for current-events research, competitive intelligence, and real-time fact-checking. The Research Mode (Pro+) goes further, conducting multi-step research across dozens of sources and synthesising findings into a structured report in a single session.

#### Impact of Web Search Launch

Internet-based research tasks on Claude.ai spiked 90× in the months following the launch of web search and Research Mode in early 2025. The feature is now among the most-used capabilities on the platform, underscoring how critical live data access is for professional workflows.

### 6.2 File & Document Analysis

Claude can ingest PDFs, Word documents, Excel spreadsheets, CSVs, PowerPoint presentations, images, and raw code files. Its 200,000-token context window (≈150,000 words) means an entire book, a full quarterly earnings report, or a 10,000-line codebase fits in a single session.

- ▶ Upload a 300-page contract and ask Claude to identify all termination clauses
- ▶ Upload a competitor's annual report and request a strategic SWOT analysis
- ▶ Upload a CSV dataset and ask for a statistical summary with business insights
- ▶ Upload an entire codebase and ask Claude to trace a bug across multiple files
- ▶ Upload 5 research papers and ask for a synthesised literature review

### 6.3 Computer Use (Beta)


Computer Use allows Claude to take control of a computer interface — clicking buttons, filling forms, navigating browsers, and interacting with desktop software — to autonomously complete complex, multi-step digital tasks. Claude scored 61.4% on the OSWorld benchmark, representing a meaningful step toward fully autonomous digital agents.

### Real-World Use Case

A legal firm uses Computer Use to automate pulling case data from a legacy government portal, copying it into a structured Excel template, and emailing a formatted summary to the case manager. A 45-minute manual task reduced to under 3 minutes of autonomous execution.

## 6.4 Model Context Protocol (MCP)

MCP is Anthropic's open standard for connecting Claude to external tools and data sources. Through MCP, Claude natively integrates with Google Drive, Slack, Gmail, Salesforce, Jira, GitHub, Close CRM, Apollo, and thousands of other applications — enabling enterprises to build deeply integrated AI workflows without custom API development.



## 6.5 Memory System

Claude's memory system retains information about you across conversations — your role, communication style, ongoing projects, and personal context. This creates a cumulative shared history that makes Claude more useful over time. Memory can be viewed, corrected, and deleted in Settings to maintain full control over what Claude retains.

## 6.6 Extended Thinking Mode

For the most demanding analytical tasks, Extended Thinking instructs Claude to perform deep, multi-step internal reasoning before responding. Particularly effective for advanced mathematics, complex legal analysis, multi-variable business decisions, and PhD-level scientific problems. In API contexts, thinking tokens are billed separately — use strategically for high-value decisions.

## 6.7 The Skills Feature

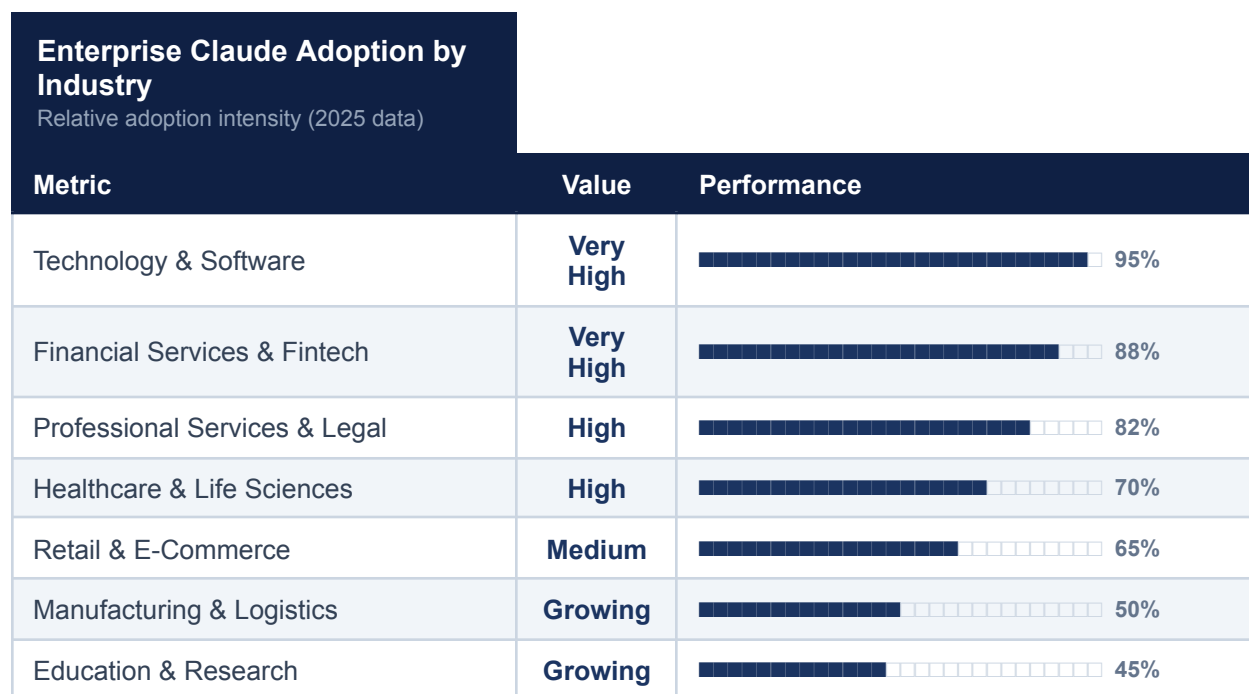
Launched October 2025, Skills allow users to save reusable task templates — mini-agents that apply a specific set of instructions to recurring workflows. A marketing team might create a 'Weekly Social Post' skill that automatically applies brand voice guidelines, content calendar context, and formatting preferences to every draft.

### CHAPTER 07

# Claude for Enterprise

Claude is embedded in 70% of Fortune 100 companies, commands 29% enterprise AI assistant market share, and has been deployed at scale across organisations from Cognizant (350,000 staff trained) to Accenture (30,000 employees onboarded). This chapter is a practical guide for organisations deploying Claude at scale.

## 7.1 Enterprise Adoption by Sector



## 7.2 The 5-Phase Deployment Playbook

Successfully deploying Claude is not just a technology decision — it is a change management initiative. The following framework is what Densight Labs uses with enterprise clients across Pakistan, the GCC, and beyond.

#	Phase	Key Activities	Timeline
<b>1</b>	<b>AI Audit</b>	Map workflows. Identify automation opportunities. Assess data readiness and compliance posture.	2–4 weeks
<b>2</b>	<b>Pilot</b>	Deploy to one team or use case. Measure quality, speed, satisfaction. Iterate and document learnings.	4–8 weeks
<b>3</b>	<b>Training</b>	Run structured prompt engineering programme. Build internal prompt libraries and usage playbooks.	2–4 weeks

4	<b>Governance</b>	Set data handling policies. Configure role-based access. Define acceptable use. Establish audit trails.	Ongoing
5	<b>Scale</b>	Expand to additional departments. Build MCP integrations. Measure ROI. Report to leadership.	Ongoing

### 7.3 Integrations Ecosystem

- ▶ 6,000+ applications integrated via the Model Context Protocol (MCP)
- ▶ Native connectors for Microsoft Teams, Slack, and Zoom — driving 72% YoY business usage growth
- ▶ Salesforce AppExchange: 170+ Claude-powered apps
- ▶ SAP and Oracle ERP Cloud integrations live in 18 countries
- ▶ Google Workspace and Microsoft 365 native plugins
- ▶ Close CRM, Apollo.io, Clay, and LinkedIn integrations for sales automation

### 7.4 Security & Compliance Matrix

Compliance / Security	Status
SOC 2 Type II	<b>Certified</b>
GDPR (EU)	<b>Compliant — opt-out training available in 48 countries</b>
AI Transparency Standards	<b>99.1% compliance rate (US Federal Review, Q1 2025)</b>
API Data Retention	<b>Not used for model training; deleted after 30 days</b>
Red Team Testing	<b>17 simulations conducted; 0.3% failure in high-risk categories</b>
Ethics Board	<b>14 independent members — AI researchers, ethicists, human rights experts</b>
Codebase Allocated to Safety	<b>20% of total engineering codebase</b>
Constitutional AI Citations	<b>25 peer-reviewed academic papers</b>

**CHAPTER 08**

## 50 High-Impact Use Cases

The following use cases represent the highest-ROI applications of Claude observed by Densight Labs across enterprise clients in Pakistan, the GCC, and the US. Each domain includes the use cases most likely to deliver measurable productivity gains immediately after deployment.

## 8.1 Writing & Content (7 Use Cases)

- ▶ Draft, edit, and polish long-form articles, whitepapers, and executive reports
- ▶ Generate SEO-optimised blog content with keyword and intent targeting
- ▶ Repurpose long-form content into LinkedIn posts, X threads, and email newsletters
- ▶ Write professional cold outreach emails, pitch deck narratives, and proposals
- ▶ Translate documents across 50+ languages with cultural nuance
- ▶ Ghostwrite thought leadership content for C-suite executives
- ▶ Create training manuals, SOPs, and structured onboarding documentation

## 8.2 Coding & Software Engineering (7 Use Cases)

- ▶ Write, debug, and refactor code in Python, JS, TypeScript, SQL, Go, and 20+ languages
- ▶ Explain complex code to non-technical stakeholders in plain language
- ▶ Generate unit tests, integration tests, and edge-case test suites automatically
- ▶ Conduct code reviews and identify security vulnerabilities with line-by-line explanations
- ▶ Build entire applications autonomously using Claude Code's agentic workflow
- ▶ Document APIs, write technical specifications, and generate changelogs
- ▶ Convert code between languages — e.g., migrate a Python 2 codebase to Python 3

## 8.3 Research & Analysis (7 Use Cases)

- ▶ Synthesise literature reviews across dozens of academic papers in one session
- ▶ Conduct competitive intelligence from uploaded annual reports and investor decks
- ▶ Analyse financial statements and extract key metrics with interpretive commentary
- ▶ Build structured research frameworks: problem trees, hypothesis trees, issue maps
- ▶ Perform sentiment analysis on customer feedback datasets uploaded as CSV
- ▶ Generate statistical summaries and insight narratives from Excel data
- ▶ Fact-check claims against uploaded reference materials with inline citations

## 8.4 Business Strategy (7 Use Cases)

- ▶ Run structured strategic planning workshops via conversational AI facilitation
- ▶ Build SWOT, Porter's Five Forces, BCG Matrix, and Blue Ocean analyses on demand
- ▶ Draft investor pitch narratives and financial model commentary for Series A–C decks
- ▶ Evaluate business model trade-offs using first-principles reasoning
- ▶ Simulate customer objections and prepare sales counterarguments by segment

- ▶ Analyse pricing strategies across competitive scenarios and elasticity assumptions
- ▶ Build organisational design recommendations for growth-stage companies

### 8.5 Legal, HR & Compliance (8 Use Cases)

- ▶ Review contracts and flag high-risk clauses with severity ratings and remediation options
- ▶ Draft employment agreements, NDAs, partnership contracts, and board resolutions
- ▶ Summarise regulatory updates and map their impact on organisational workflows
- ▶ Build compliance training content for banking, healthcare, and fintech sectors
- ▶ Generate inclusive, clarity-optimised job descriptions from role outlines
- ▶ Analyse employee survey data and extract thematic insights with sentiment breakdown
- ▶ Create HR policy documents, employee handbooks, and code-of-conduct frameworks
- ▶ Prepare board reports, investor updates, and shareholder communications

### 8.6 Education & Learning (7 Use Cases)

- ▶ Generate personalised study plans and learning roadmaps for any subject or certification
- ▶ Explain complex concepts at adjustable difficulty levels from beginner to PhD
- ▶ Create quiz questions, assessment rubrics, and answer keys from uploaded course material
- ▶ Tutor students through mathematics, science, coding, and humanities in real time
- ▶ Provide detailed, rubric-based feedback on essays and written assignments
- ▶ Translate textbooks and research papers for non-native language learners
- ▶ Build interactive case studies, simulations, and Socratic dialogues for curricula

### 8.7 Usage Pattern Snapshot

<b>What Claude Users Actually Do</b> <small>Task distribution across claude.ai conversations (2025)</small>		
Metric	Value	Performance
Coding & Mathematical Tasks	34%	34%
Writing & Content Creation	24%	24%
Research & Analysis	20%	20%
Bug Fixing (Most Common Task)	10%	10%
Business & Strategy	7%	7%
Other / Miscellaneous	5%	5%

**CHAPTER 09**

# Safety, Ethics & Responsible AI Use

Understanding what Claude will and will not do — and why — is essential for every serious user. Claude's ethical framework is not a set of arbitrary restrictions; it is a carefully designed system for ensuring one of the world's most powerful AI models remains beneficial to humanity. Understanding this framework makes you a better, more effective user.

## 9.1 What Claude Always Does

- ▶ Acknowledges uncertainty and says clearly when it does not know something
- ▶ Provides balanced perspectives on contested political and social questions
- ▶ Protects user privacy — does not store or share personal information beyond the session
- ▶ Follows data deletion protocols — API data retained for maximum 30 days, not used for training
- ▶ Flags when a question requires professional expertise (legal, medical, financial)
- ▶ Allows users to correct its memory and remove any stored information on request

## 9.2 What Claude Will Not Do

- ▶ Generate content designed to harm, deceive, or manipulate specific individuals
- ▶ Provide detailed technical instructions for weapons of mass destruction
- ▶ Create sexual content involving minors under any circumstances whatsoever
- ▶ Write functional malware, ransomware, spyware, or network exploit code
- ▶ Assist with targeted harassment, coordinated hate speech, or incitement to violence
- ▶ Fabricate citations, academic credentials, or attribute statements to real people falsely

## 9.3 Safety Performance Metrics

Safety Metric	Performance
Toxicity filtering success rate	<b>97.2%</b>
AI transparency compliance	<b>99.1% (US Federal Review, Q1 2025)</b>
Red-team failure rate (high-risk categories)	<b>0.3% across 17 simulations</b>
Constitutional AI academic citations	<b>25 peer-reviewed papers</b>
Ethics Board independent members	<b>14 (researchers, ethicists, human rights experts)</b>
Engineering codebase for ethical alignment	<b>20% of total</b>

Countries with opt-out training data controls	<b>48 countries</b>
---	---------------------

## 9.4 Organisational Governance Principles

Densight Labs recommends the following six governance principles for any enterprise Claude deployment:

1. Establish an AI Acceptable Use Policy before deployment — define permitted and prohibited use cases explicitly, and require acknowledgement from all users.
2. Implement role-based access controls: not all employees should have equal access to Claude's full capability set.
3. Maintain human-in-the-loop for all high-stakes decisions. Claude's outputs inform decisions; humans make them.
4. Run regular AI audits to assess output quality, detect bias patterns, and ensure alignment with organisational values.
5. Invest in prompt engineering training for all users. Untrained users generate poor outputs and may misuse the system.
6. Never input genuinely confidential data into Claude.ai unless on an Enterprise plan with a signed Business Associate Agreement and data processing addendum.

### CHAPTER 10

# Power User Tips, Productivity & The Future

## 10.1 The Densight Labs Claude Productivity Stack

After deploying Claude across dozens of enterprise engagements and training thousands of professionals, Densight Labs has identified the highest-leverage habits of top Claude power users:

Habit	Why It Matters
<b>Maintain a Prompt Library</b>	Save your best prompts in a shared document. Reuse and iterate — never start from scratch.
<b>Use Projects for recurring work</b>	A Project with standing instructions eliminates 5–10 minutes of context-setting per session.
<b>Upload before you ask</b>	Always upload the relevant document, data, or code before asking. Context transforms outputs.
<b>Request structured outputs</b>	Ask for tables, JSON, or numbered lists — they are easier to copy, share, and process downstream.

<b>Use the Regenerate function</b>	Claude generates multiple variations on request. Ask for 3 options when exploring ideas.
<b>Separate tasks by conversation</b>	Unrelated tasks in separate sessions. Mixed long conversations degrade context quality.
<b>Rate responses</b>	The thumbs-down button directly improves Claude for everyone. Use it when outputs miss the mark.

## 10.2 Claude vs. The Competition

Dimension	Claude	ChatGPT	Gemini	Grok
Writing Quality	★★★★★	★★★★☆	★★★★☆	★★★★☆
Coding Accuracy	★★★★★	★★★★☆	★★★★☆	★★★★☆
Context Window	<b>200K tokens</b>	128K tokens	1M tokens	128K tokens
Safety & Ethics	★★★★★	★★★★☆	★★★★☆	★★★★☆
Enterprise Adoption	<b>29% share</b>	Dominant	Growing	Niche
API Data Privacy	<b>Not trained on</b>	Opt-out req.	Mixed	Limited info
User Satisfaction	<b>92%   4.6/5</b>	High	High	Moderate

## 10.3 The Future of Claude

Capability Area	Expected Direction
Agentic Autonomy	Multi-step, multi-day autonomous agents that plan, execute, and self-correct across complex workflows.
Extended Memory	Persistent long-term memory across all conversations and Projects — true ongoing AI collaboration.
Deeper Multimodality	Integrated image generation, video analysis, and audio understanding in the core Claude experience.
Real-Time Voice	Full conversational voice interface with natural turn-taking, emotion awareness, and live transcription.
Domain Specialisation	Purpose-built Claude variants fine-tuned for legal, medical, financial, and engineering domains.

Global Expansion	Deeper multilingual capability and localised compliance frameworks for emerging markets.
------------------	--

## 10.4 A Note from the Author

### From Numan Ahmad — Founder & CEO, Densight Labs

At Densight Labs, we believe we are living through the most important technological transition in human history. The professionals, businesses, and institutions that learn to collaborate with AI — not just use it — will define the next decade. Claude is not a tool you use once and forget. It is a collaborator you grow with. The more context you give it, the better it understands your work. The more you refine your prompts, the sharper its outputs become. I wrote this guide because I wanted every professional in Pakistan, across the GCC, and around the world to have access to world-class AI education. That is Densight Labs' mission: to turn AI consumers into AI builders. Master Claude, and you have mastered one of the most powerful thinking tools in human history.

## CHAPTER 11

# Claude Projects — Your Persistent AI Workspace

Most people use Claude the same way every time: open a new chat, re-explain all their context, get an answer, and lose everything when the conversation ends. Claude Projects solves this completely. A Project is a persistent, intelligent workspace where Claude permanently remembers your context, documents, custom instructions, and preferences — behaving like a specialist colleague who already knows your work every single session.

### The Core Insight

The difference between regular Claude chat and a Claude Project is the difference between hiring a freelancer for a single task and having a senior full-time employee who has read every document in your company and knows exactly how you like things done. Projects turn Claude from a stateless question-answering tool into a personalised AI expert on your specific domain.

## 11.1 What a Project Actually Contains

Component	What It Does
Custom Instructions	Your permanent system prompt for this Project. Defines Claude's role, tone, constraints, and approach. Applies to every conversation automatically.

<b>Knowledge Base</b>	Upload PDFs, DOCX, XLSX, CSV, TXT, HTML, and code files. Claude references them in every chat — no re-uploading needed. RAG-powered on paid plans for 10× capacity.
<b>Conversation History</b>	All chats within this Project are stored together and linked to the same context. Browse, search, and continue any previous session.
<b>Team Sharing</b>	Share Projects with specific colleagues or your entire organisation. Set Can View or Can Edit permissions per member.
<b>Context Window</b>	Each Project has a 200,000-token context window — equivalent to a 500-page book. Paid plans expand this 10× via RAG.

## 11.2 Step-by-Step: Creating Your First Project

### Step 1 — Create the Project

Open Claude.ai and click 'Projects' in the left sidebar. Click 'New Project' (or the orange + button). Give it a specific, descriptive name — not 'Project 1' but something like 'Q4 Sales Campaign', 'Backend API Codebase', or 'Client Acme: Legal Review'. A clear name makes the Project instantly recognisable when you have ten running in parallel.

### Step 2 — Write Your Custom Instructions

This is the most important step and the one most users skip. Your custom instructions are the standing system prompt — they define who Claude is inside this Project and how it should always behave. Click 'Set Project Instructions' and write them thoughtfully.

#### Custom Instructions Template (Copy & Adapt)

You are [Role] working with [Your Name/Team] at [Company/Context]. Your job is to [Core Task].  
 Tone & Style: [e.g. 'Formal and precise, like a McKinsey consultant. No hedging, no disclaimers. Direct answers only.']  
 Always: [e.g. 'Structure long outputs with clear headings. End every analysis with a numbered action plan.']  
 Never: [e.g. 'Do not recommend external consultants. Do not use jargon without explaining it. Never exceed 400 words for summaries.']  
 Background context: [Paste 2-3 sentences about your project, company, or domain that Claude should always know.]

### Step 3 — Build the Knowledge Base

In the Project Knowledge section, click the + button to upload files. Add any documents Claude should permanently reference: your company's style guide, product documentation, competitive landscape reports, your team's SOPs, relevant research papers, or your codebase. Claude processes each file and makes it searchable across every conversation in the Project.

File Type	Best Used For	Pro Tip
-----------	---------------	---------

PDF	Contracts, reports, research papers, manuals	Name files descriptively — Claude sees filenames
DOCX	Brand guides, SOPs, strategy documents	Plain text extracts better than image-heavy docs
XLSX / CSV	Data tables, KPI dashboards, lists	Keep spreadsheets clean — remove merged cells
TXT / MD	Code, meeting notes, raw text context	Markdown formats beautifully in responses
Code files	Codebases, configs, scripts	Upload multiple files — Claude reads them together

#### Step 4 — Start Chatting (Without Re-explaining Context)

Every conversation you open inside this Project automatically inherits your custom instructions and knowledge base. You can jump straight into the work without any setup. Ask Claude to 'analyse the Q3 report' and it already knows which Q3 report. Ask it to 'write in our brand voice' and it has already read your brand guide.

#### Step 5 — Share with Your Team

On Team or Enterprise plans, click 'Share Project' and add colleagues by email. Assign Can View (read-only) or Can Edit (full access) permissions. Team members see the same knowledge base and instruction set, ensuring every person on the team works with the exact same AI configuration — no drift, no inconsistency.

### 11.3 The Best Projects for Every Role

Role	Recommended Projects	What to Upload
CEO / Founder	Investor Relations · Strategy Ops · Board Prep	Cap table, pitch deck, financials, board minutes
Sales Manager	Prospect Research · Outreach Copy · Deal Reviews	ICP definition, product one-pager, competitor matrix
Software Engineer	Codebase Assistant · PR Reviews · API Docs	CLAUDE.md, key source files, architecture diagrams
Lawyer	Client Matter · Contract Templates · Regulatory Watch	Precedents, applicable statutes, client brief
Marketer	Brand Voice · Campaign Planning · SEO Research	Brand guidelines, content calendar, audience personas

Researcher	Literature Review · Data Analysis · Paper Drafting	Uploaded papers, datasets, methodology notes
HR Manager	Hiring · Onboarding · Policy Management	Job descriptions, onboarding checklist, HR policies

## 11.4 Advanced Projects: The Master Prompt Architecture

Power users structure their custom instructions as a 'Master Prompt' — a comprehensive document that turns Claude into a domain expert tuned specifically to their workflow. A well-designed Master Prompt has five sections:

7. Identity & Role — Who Claude is in this context and what its primary mission is
8. Background Context — 3–5 sentences about your company, project, or domain that Claude must always know
9. Behavioural Rules — Explicit do's and don'ts covering tone, length, structure, and reasoning style
10. Output Formats — Templates or examples showing exactly how Claude should format responses
11. Escalation Protocol — What Claude should do when it lacks information or faces an ambiguous question

### Real-World Example: Sales Project Master Prompt

You are a B2B Sales Intelligence Analyst at Densight Labs. Your job is to help the sales team with prospect research, outreach drafting, and deal analysis for mid-market companies in Pakistan, the GCC, and the US. Context: Densight Labs is an AI consulting and education company. Our ICP is Operations Directors and CTOs at 200–2,000 employee companies in banking, fintech, and technology. Always: Lead with the most actionable insight. Use bullet points for lists of more than 3 items. End every prospect analysis with a recommended first outreach angle. Never: Use generic AI phrases like 'certainly' or 'of course'. Never recommend reaching out on platforms other than LinkedIn and email. Never write outreach longer than 150 words. Format: For prospect profiles use: Company | Industry | Revenue | Key Pain Point | Recommended Angle.

## 11.5 Projects vs. Regular Chat — When to Use Which

Situation	Use Projects	Use Regular Chat
Recurring tasks in same domain	✓ Always	✗ Lose context each time
Working with specific documents	✓ Upload once, reference forever	✗ Must re-upload every session
Team-shared workflows	✓ Shared knowledge base	✗ Not shareable
One-off quick questions	✗ Overkill	✓ Faster for one-shot tasks

Brand/tone consistency needed	✓ Custom instructions lock it in	✗ Must re-specify each chat
Exploring a new topic you know nothing about	✗ No documents to upload	✓ Use general chat

CHAPTER 12

# Claude Code — The Developer's Autonomous Partner

Claude Code is Anthropic's command-line AI coding agent. Unlike chatbots that answer coding questions, Claude Code directly reads your files, runs commands, edits your codebase, manages git workflows, and delegates complex tasks to parallel sub-agents — all from natural language instructions in your terminal. It is the fastest-growing enterprise software product in history, reaching \$2.5 billion in annualized revenue within months of launch.

<b>\$2.5B</b> Annualized Revenue (Early 2026)	<b>30hrs+</b> Sustained Autonomous Coding	<b>80.9%</b> SWE-bench Score (Industry #1)
--	--	---

## 12.1 Installation & Setup

### Requirements

- ▶ macOS 10.15+, Ubuntu 20.04+/Debian 10+, or Windows 10+ via WSL or Git Bash
- ▶ Node.js 18+ (check with: node --version)
- ▶ Active Claude Pro, Max, Team, or Enterprise subscription
- ▶ 4 GB RAM minimum, internet connection

### Install via npm (Universal)

```
npm install -g @anthropic-ai/claude-code
claude # Launch interactive session
claude auth login # Authenticate your account
```

### Install via Homebrew (macOS)

```
brew install claude-code
brew upgrade claude-code # Update to latest version
```

### Desktop App (No Terminal Needed)

Claude Code is also available inside the Claude Desktop app under the 'Code' tab — same engine as the CLI, with a graphical interface. Preview running servers, review code diffs visually, and monitor pull request status without leaving the app. The Desktop and CLI share configuration files (CLAUDE.md, MCP servers, skills, hooks), so you can switch between them seamlessly.

## 12.2 Essential CLI Commands Reference

Command	What It Does
claude	Start interactive REPL session in current directory
claude "fix the auth bug"	Start session pre-seeded with a specific prompt
claude -c	Continue your most recent conversation
claude -p "summarise README"	Non-interactive print mode — outputs and exits
claude --model claude-opus-4-6	Choose a specific model for this session
claude update	Update Claude Code to the latest version
claude doctor	Diagnose setup issues, check version & environment
claude auth login	Authenticate your Anthropic account
claude auth status	Check current authentication status

## 12.3 In-Session Slash Commands

Slash Command	Purpose
/init	Scan your project and auto-generate a CLAUDE.md file with project context
/clear	Wipe conversation history — use when switching to a new task in the same project
/compact	Summarise current conversation to free context window — useful on long sessions
/config	Open settings menu: model selection, tool permissions, behaviour toggles

/help	Show all available commands including custom slash commands you have defined
/fast	Toggle fast mode — same Opus model at up to 2.5× faster output (higher token cost)
/install-github-app	Connect Claude to GitHub for automatic PR reviews on every pull request you open
@filename	Reference a specific file in your prompt — e.g. @src/auth.ts fix the token expiry bug
!command	Run a shell command from within the Claude session — e.g. !git status

## 12.4 The CLAUDE.md File — Your Project Memory

CLAUDE.md is the single most powerful configuration tool in Claude Code. It is a Markdown file you create at the root of your project that Claude automatically reads at the start of every session. It acts as permanent memory for your codebase — no more explaining your project structure, coding standards, or team conventions on every session.

### Run /init to auto-generate a starter CLAUDE.md

Claude Code will scan your project, read your package.json, directory structure, and existing config files, and produce a starter CLAUDE.md tailored to your project. Edit and expand it from there.

A production CLAUDE.md should include:

- ▶ Project overview — what the app does, tech stack, architecture summary
- ▶ File structure map — which directories contain what
- ▶ Coding standards — naming conventions, patterns to follow, patterns to avoid
- ▶ Frequently used commands — how to run tests, start dev server, build, deploy
- ▶ Known gotchas — legacy decisions, environment quirks, things not to touch
- ▶ MCP integrations — which external tools are configured and their purpose

## 12.5 The 4-Mode Workflow: How Power Users Use Claude Code

Mode	When to Use	How to Invoke
Explore	First pass on unfamiliar code	"Read src/auth/ and explain the authentication flow. List all files involved."
Plan	Before major features or refactors	"Think hard about how to add OAuth2. Create a detailed plan before writing any code."

<b>Code</b>	Implementing planned changes	"Implement the plan from above. Start with src/auth/oauth.ts then update the index."
<b>Review</b>	After changes, before commit	"Review the changes you just made against our coding standards in CLAUDE.md."

## 12.6 Skills, Sub-Agents & MCP — The Power Stack

### Skills in Claude Code

Skills are reusable instruction sets stored in `.claude/skills/` in your project. Each skill is a `SKILL.md` file that activates automatically when Claude detects a relevant task, or can be invoked explicitly. Skills encode your team's domain expertise into Claude — a Security Review skill, a TypeScript Standards skill, or a Database Migration skill. Unlike slash commands (which you invoke), skills are intelligent enough to self-activate.

#### Example: Creating a Custom Security Review Skill

Create the file: `.claude/skills/security-review/SKILL.md`  
 Content: `name: security-review | description: Automatically activates on any task touching authentication, authorisation, API endpoints, or data handling. Reviews for OWASP Top 10, SQL injection, XSS, insecure direct object references, and hardcoded secrets. Always outputs a severity-rated findings table.Result: Every time you ask Claude to review login code or modify an API endpoint, this skill activates automatically.`

### Sub-Agents

When Claude Code encounters a task that benefits from parallel execution or isolated context, it can spawn sub-agents. A sub-agent starts with a clean context window, works on a specific subtask, and reports back. This architecture prevents context bloat on long sessions and enables parallel workflows — three sub-agents researching three competitors simultaneously, or two agents fixing bugs in separate modules at the same time.

### MCP (Model Context Protocol) Integrations

MCP connects Claude Code to your external toolchain. With 300+ integrations available, Claude Code can query your production database, create GitHub issues, review PRs, check Sentry error logs, update Notion docs, and interact with virtually any API — all from natural language in your terminal.

Add MCP Server	Command
GitHub	<code>claude mcp add github -- gh-mcp</code>
Slack	<code>claude mcp add slack -- slack-mcp</code>
PostgreSQL	<code>claude mcp add postgres -- pg-mcp --connection-string \$DATABASE_URL</code>

Jira	<code>claude mcp add jira -- jira-mcp --host your.atlassian.net</code>
Custom API	<code>claude mcp add myapi -- node /path/to/mcp-server.js</code>

## 12.7 Claude Code: 10 Real-World Workflows

12. Fix a GitHub issue end-to-end: "Read issue #42 via MCP, understand the requirements, implement the fix, write tests, and open a PR linked to the issue."
13. Full feature build: "Build a user authentication system with JWT, refresh tokens, and rate limiting. Follow the patterns in src/auth/. Write tests for all edge cases."
14. Legacy codebase onboarding: "Explore the entire codebase, map the architecture, identify technical debt, and produce a prioritised refactoring roadmap."
15. Database migration: "Write a safe migration script to add the new user\_preferences table. Include rollback. Test against the staging schema in schema.sql."
16. API documentation: "Read every endpoint in src/api/ and generate a full OpenAPI 3.0 spec with request/response examples for each route."
17. Dependency audit: "Scan all dependencies for known CVEs using npm audit, propose updates, and implement the non-breaking ones automatically."
18. Performance profiling: "Analyse the database queries in the reports module, identify the N+1 queries, and rewrite them with proper JOINS."
19. Test suite generation: "Read every function in src/utills/ and write comprehensive pytest unit tests covering success cases, edge cases, and failure modes."
20. Code review automation: "Review every file touched in this PR against our CLAUDE.md standards. Output a table: File | Finding | Severity | Recommendation."
21. Multi-agent research: "Spawn three sub-agents to research Competitor A, B, and C in parallel. Compile findings into a competitive intelligence report."

### Model Selection Tip for Claude Code

Claude Code automatically uses Opus until you hit 50% of your usage allocation, then switches to Sonnet for cost efficiency. For straightforward tasks (formatting, simple bug fixes), Sonnet is plenty. For architectural decisions, complex refactors, and multi-file reasoning tasks, switch back to Opus with: `--model claude-opus-4-6`. Use `/fast` mode when iterating quickly — same quality, up to 2.5× faster output.

## CHAPTER 13

# Cowork & Skills — AI That Works While You Sleep

Claude Code gave developers an autonomous AI partner in the terminal. Cowork does the same thing for everyone else. It is the most significant product shift in Claude's history: a non-technical user can now point Claude at a folder, describe an outcome, step away, and come

back to finished work — organised files, synthesised reports, extracted data, drafted documents — without writing a single line of code or prompt.

**The One-Sentence Summary**  
 Regular Claude chat: you ask, Claude responds, you act. Cowork: you describe the outcome, Claude plans, executes, and delivers — you approve at key decision points and collect the finished result.

### 13.1 Cowork at a Glance

Dimension	Cowork	Regular Claude Chat
Mode	<b>Autonomous / Agentic</b>	Conversational / One-turn
File Access	<b>Direct read/write to your local folders</b>	Manual upload only, per session
Task Length	<b>Multi-step, 30+ min tasks</b>	Single responses
User Involvement	<b>Set task → approve plan → collect result</b>	Active back-and-forth required
Scheduling	<b>Daily/weekly automated tasks</b>	Manual trigger only
Parallel Execution	<b>Multiple sub-agents simultaneously</b>	One thread
Data Storage	<b>Local device only — never sent to Anthropic</b>	Server-side (standard data policy)

### 13.2 Requirements & Installation

Requirement	Detail
Subscription	<b>Pro, Max, Team, or Enterprise (research preview)</b>
Platform	<b>macOS or Windows (full feature parity as of February 2026)</b>
App Required	<b>Claude Desktop app — not available on web or mobile</b>
Internet	<b>Required throughout the session</b>
App Open	<b>Desktop app must remain open while Cowork is running</b>
Download	<b><a href="https://claude.com/download">claude.com/download</a> — select your OS, install, sign in</b>

**To access Cowork:** Open Claude Desktop → look for the three-tab selector at the top of the window → click 'Cowork'. You will see the task input interface with a 'Work in a Folder' checkbox. That checkbox is the key — it grants Claude access to your local files.

## 13.3 Your First Cowork Task — Step by Step

### Step 1 — Open Cowork Mode

Launch Claude Desktop. Click the 'Cowork' tab at the top of the window (alongside 'Chat' and 'Code'). You will see a task input panel rather than a chat interface.

### Step 2 — Grant Folder Access

Check the 'Work in a Folder' box. Select the folder Claude should have access to. Claude can only read, write, and create files within this folder — it cannot touch anything outside it. This is your safety boundary. Choose a dedicated work folder, not your entire home directory.

### Step 3 — Describe the Outcome, Not the Steps

This is the most important mindset shift. In regular chat, you ask step-by-step questions. In Cowork, you describe the finished result you want. Be specific about the output, not the process.

#### Weak vs. Strong Cowork Prompts

WEAK: 'Help me with my files'  
 STRONG: 'Scan my Downloads folder. Organise all files into subfolders by type (PDFs, Images, Spreadsheets, Installers, Other). Rename files using the format YYYY-MM-DD\_DescriptiveName. Flag any files over 500MB for my review. Show me the plan before making any changes.'  
 The strong prompt specifies: what to scan, what to produce, naming conventions, edge cases, and a review checkpoint.

### Step 4 — Review Claude's Plan

Before doing anything, Cowork shows you its execution plan. Read it. This is your opportunity to redirect before any files are touched. Claude will not proceed without your approval. You can modify the plan, add constraints, or cancel entirely.

### Step 5 — Monitor Progress (or Step Away)

Once you approve, Claude executes. A progress panel shows each step in real time — which files it has read, what it is creating, what sub-tasks it has spawned. You can watch, intervene at any point by pressing Escape, or simply close the panel and return to your work. Cowork runs in the background.

### Step 6 — Collect Your Output

When the task completes, Cowork delivers finished files directly into your folder. A summary panel shows what was created, modified, or flagged. Review the outputs. If anything needs adjustment, describe the change and Cowork continues from where it left off.

### 13.4 Global & Folder Instructions

Like Claude Code's CLAUDE.md, Cowork supports standing instructions that apply automatically to every session — no need to re-specify preferences every time.

- ▶ Global Instructions: Apply to every Cowork session regardless of folder. Set via Settings → Cowork → Edit. Include your name, role, preferred output formats, and general working style.
- ▶ Folder Instructions: Apply only when you work in a specific folder. Set by asking Claude to save folder instructions during a session, or via Settings. Include project-specific context, naming conventions, and constraints for that folder.

#### Global Instructions Template

My name is [Name]. I work as [Role] at [Company]. When creating documents, always use our brand style: [describe]. For spreadsheets, include an executive summary row at the top. For research tasks, always cite sources with URLs. Never permanently delete files — move them to a '\_Archive' subfolder instead. Flag anything requiring legal review with [LEGAL REVIEW NEEDED].

### 13.5 Scheduled Tasks

One of Cowork's most powerful features is the ability to schedule tasks that run automatically — daily briefings, weekly reports, monthly file organisation — without any manual trigger. This turns Claude into a genuine recurring workflow engine, not just an on-demand tool.

To schedule a task: Start any Cowork task and type /schedule. Claude will ask for the frequency (daily, weekly, monthly, or a specific time). Confirm the schedule. The task then runs automatically according to the cadence you set, delivering results to your folder each time.

Scheduled Task	Frequency	What Cowork Delivers
Daily Briefing	Every weekday morning	Summary of Slack messages, calendar, and priority tasks
Weekly Report Draft	Every Friday afternoon	Draft weekly update doc from your notes folder
Monthly File Cleanup	1st of each month	Organise Downloads, archive old files, flag duplicates
Expense Report	End of each month	Extract all receipt data from folder into XLSX

Competitive Monitor	Weekly	Pull competitor pricing/news via connectors, update tracker
---------------------	--------	---

### 13.6 Cowork + Connectors + Claude in Chrome

Cowork becomes dramatically more powerful when combined with Anthropic's connector ecosystem and Claude in Chrome:

#### Connectors

Connectors link Cowork to your cloud tools — Slack, Notion, Figma, Google Drive, and more. A task like 'Generate a weekly briefing from our Slack channels and project tracker' becomes possible because Cowork can pull live data from both tools, synthesise it, and write a formatted document to your folder — all without leaving your desktop.

#### Claude in Chrome

Pair Cowork with Claude in Chrome (a separate connector that gives Claude browser control) and tasks requiring web navigation become possible: pulling data from web dashboards, filling out forms, extracting competitor pricing from websites, submitting reports to online portals. Claude can now complete end-to-end tasks that span your local files, your cloud tools, and the open web simultaneously.

#### Real Enterprise Use Case

A CFO at a mid-market company uses Cowork with a scheduled weekly task: Pull Q-to-date revenue data from Salesforce (connector), grab the macro indicators from FRED (Claude in Chrome), compare company growth against macro trends, and produce a formatted 2-page board-ready PDF in my Reports folder by 8am every Monday. Zero manual work. The Monday report is waiting when they arrive.

### 13.7 Plugins for Cowork

Plugins are pre-packaged extensions that bundle skills, connectors, slash commands, and sub-agents into a single install — turning Claude into a specialist for a specific role, team, or workflow. Think of them as 'apps for Cowork'. Install once, and Claude has everything it needs to perform that specialised function.

- ▶ Install plugins via Settings → Cowork → Plugins, or within any Cowork session
- ▶ The official Anthropic plugin marketplace launched with 36 curated plugins; community marketplaces are growing rapidly
- ▶ Build custom plugins for your team to standardise workflows and share expertise across the organisation
- ▶ Plugins can bundle your internal brand standards, approved tools, naming conventions, and sub-agent configurations

### 13.8 Skills in Claude.ai (Non-Developer Skills Feature)

The Skills feature in Claude.ai (the web and mobile interface — not to be confused with Claude Code skills) was launched in October 2025 and allows any user — no coding required — to create reusable task templates that apply a fixed set of instructions to recurring workflows automatically.

What Skills Do	Example
Lock tone & format	A 'LinkedIn Post' skill always applies your brand voice, hashtag rules, and character limit
Pre-load context	A 'Client Brief Analyser' skill pre-loads your agency's analysis framework
Define output structure	A 'Sales Email' skill always outputs subject line, email body, and P.S. line separately
Set reasoning approach	A 'Strategy Analyser' skill always applies Porter's Five Forces before concluding

To create a Skill in Claude.ai: open Settings → Skills → Create New Skill. Name it, write the instruction set (what Claude should always do when this skill is active), and save. Activate a skill by selecting it from the Skills menu before starting a conversation. Skills can be combined with Projects for maximum consistency.

### 13.9 Choosing the Right Mode for Every Task

Task Type	Chat	Projects	Cowork	Claude Code
Quick one-off question	✓ Best	Fine	Overkill	Overkill
Recurring work with same context	Inefficient	✓ Best	Good	Good for code
Multi-step file/folder task	Not possible	Not possible	✓ Best	Code only
Writing / editing / research	Good	✓ Best with context	Good	Possible but slow
Coding & engineering tasks	Basic only	Good	Not ideal	✓ Best
Autonomous long-running task	Not possible	Not possible	✓ Best	✓ Best for code

Team-shared AI workflow	Manual	✓ Best	Good	CLAUDE.md sharing
Scheduled recurring automation	Not possible	Not possible	✓ Best	Via CI/CD hooks

APPENDIX

## Quick Reference Card

### A1 — Model API Strings

Model	API String	Best Use
Claude Opus 4.6	claude-opus-4-6	Complex reasoning, research
Claude Sonnet 4.6	claude-sonnet-4-6	Everyday tasks, coding
Claude Haiku 4.5	claude-haiku-4-5-20251001	High-volume automation

### A2 — Product Mode Decision Matrix

I need to...	Chat	Projects	Cowork	Claude Code
Ask a quick question	✓			
Work in a domain consistently		✓		
Organise / process local files			✓	
Write / debug code				✓
Automate a recurring workflow			✓	✓
Share AI setup with my team		✓		✓ CLAUDE.md

### A3 — Key Statistics

Metric	Value (2025–2026)
--------	-------------------

Monthly Active Users	<b>30 million+</b>
Annualized Revenue (Run-Rate)	<b>\$14 billion (February 2026)</b>
Anthropic Valuation	<b>\$380 billion (Series G)</b>
Countries Available	<b>159 countries</b>
Context Window	<b>200,000 tokens</b>
Response Accuracy (Claude 3.5)	<b>98.3%</b>
Fortune 100 Adoption	<b>70%</b>
SWE-bench Coding Score	<b>80.9% (industry-leading)</b>
Claude Code Revenue (Run-Rate)	<b>\$2.5 billion (early 2026)</b>
MCP Integrations Available	<b>6,000+ apps, 300+ developer integrations</b>
Cowork Launch Date	<b>January 2026 (macOS) · February 2026 (Windows)</b>

## A4 — Essential Resources

- claude.ai**      <https://claude.ai> — Claude.ai Web, Mobile & Desktop Interface
- code.claude.com**      <https://code.claude.com> — Claude Code Documentation
- docs.claude.com**      <https://docs.claude.com> — Official Developer Documentation & API Reference
- support.claude.ai**      <https://support.claude.ai> — User Support, Release Notes & Help Centre
- claude.com/download**      <https://claude.com/download> — Download Claude Desktop (macOS & Windows)
- anthropic.com**      <https://www.anthropic.com> — Anthropic Research, Safety Papers & Blog
- densightlabs.com**      <https://densightlabs.com> — Densight Labs: AI Education & Enterprise Consulting

**© 2026 Densight Labs (DENSIGHT SMC-PRIVATE LIMITED)**  
 Authored by Numan Ahmad, Founder & CEO · All rights reserved.  
*Statistics sourced from Anthropic, DemandSage, SQMagazine, AI Business Weekly, DataCamp, and Anthropic support documentation (2025–2026).*