

# Paralegals and Legal Assistants

AI now drafts first-pass documents, summarises discovery and pulls case law, while paralegals keep filings, clients and deadlines on track.

| EXPOSURE                                   | WINDOW                                           | ADOPTION TODAY  | READING                                                                                                                                         |
|--------------------------------------------|--------------------------------------------------|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <div>40</div> <div>Moderate exposure</div> | <div>4–9 yrs</div> <div>until change lands</div> | <div>High</div> | <div>Augmented more than re-<br/>placed.</div> <div>Exposure is the share of today's work AI<br/>can plausibly take on within the window.</div> |

§ 01WHAT IS HAPPENING

## What is happening to paralegals and legal assistants

### Impact

Generative tools built into practice-management and research platforms now produce first drafts of routine correspondence, pleadings and contract clauses, summarise depositions and medical records, and run case-law searches that used to take an afternoon. Document review platforms classify and prioritise discovery sets so that a paralegal checks the machine's work rather than reading every page. The day-to-day shifts from producing text to checking it, keeping the matter organised, managing court deadlines and dealing directly with clients, witnesses and court staff.

## Risk

Moderate exposure: routine research and drafting is augmented; value shifts to case management, judgement and client contact.

Tasks that automate most readily are first-draft document assembly, legal research summaries, document review for relevance and privilege, and the extraction of key terms from contracts and records. Tasks that stay human include managing the procedural calendar, filing with courts whose rules vary by jurisdiction, interviewing clients and witnesses, verifying that an AI-generated citation actually exists and says what it claims, and the judgement about what matters to the lawyer running the case. Occupation-level measures place the role in a high official exposure tier, but the measured applicability of current language models to its specific task mix is comparatively low, which is why the score sits in the moderate band. Over the 4-9 year window expect fewer junior drafting hours per matter and more paralegals supervising AI output, with headcount pressure concentrated in high-volume, document-heavy practices such as personal injury, insurance defence and conveyancing.

## Sector readiness

## Rapid Integration via Practice-Management Software

Legal AI has arrived mainly through tools firms already pay for: Clio, Thomson Reuters and LexisNexis have each put generative assistants into their standard products, and e-discovery vendors have added AI review. Large firms and in-house teams are furthest along; small firms and public-sector legal offices are adopting more slowly and often informally through general-purpose chatbots. Courts and professional bodies are now issuing rules on disclosure and verification of AI-assisted work, which is pushing firms towards formal policies.

## § 02 POSITION

## Where you stand

- i Position yourself as the verification layer: the person who checks AI-generated research and drafts against the primary sources before anything leaves the firm.
- ii Build depth in a practice area's procedure, court rules and filing systems, which AI tools handle poorly and lawyers rely on heavily.
- iii Take ownership of client and witness communication, where a careful human still does the work and the firm's reputation is on the line.

## § 03 ACTIONS

6 POINTS

## What this means for you

Concrete changes to how the work gets done, in the order you are likely to meet them.

- 01 **Learn the firm's tools properly.** Get formal training on Clio Duo, CoCounsel or whichever assistant your firm licenses, and become the person colleagues ask when it misbehaves. Fluency with the tools is now part of the job description.
- 02 **Verify every citation.** Treat any case, statute or quotation an AI tool produces as unverified until you have opened the source. Courts have sanctioned lawyers over fabricated citations, and the paralegal who prevents that is indispensable.
- 03 **Own the procedural calendar.** Deadlines, service rules and local court requirements are where matters are actually won or lost on the admin side, and current AI handles them badly. Make this your domain.
- 04 **Move up the discovery stack.** Instead of first-pass review, learn to set up review protocols, train the classifier, run quality-control samples and document the process defensibly.
- 05 **Keep the client relationship.** Clients remember who returned their call and explained what happens next. That contact is not being automated and it makes you harder to cut.
- 06 **Document what the AI did.** Keep notes on which tools touched which drafts and what you changed. Disclosure rules are tightening and a clear record protects you and the firm.

## § 04 CAUSES

6 DRIVERS

## What is pushing this change

- 01 **Generative drafting in practice software.** Clio Duo, CoCounsel and similar assistants produce first drafts of routine letters, pleadings and clauses from the matter file, removing much of the keyboard time that once defined junior legal work.
- 02 **AI-assisted legal research.** Lexis+ AI and Westlaw Precision return summarised answers with linked authorities, cutting research hours but raising the need for a human to check that the authority says what the summary claims.
- 03 **Technology-assisted document review.** Platforms such as Relativity aiR classify documents for relevance and privilege, so large review projects need fewer people reading and more people supervising.
- 04 **Record and deposition summarisation.** Long medical records, transcripts and contracts can now be summarised and key terms extracted in minutes, a task that was a staple of paralegal billable time.
- 05 **Client and court pressure on cost.** Clients increasingly refuse to pay for routine drafting hours, and courts are piloting their own AI tools, pushing firms to automate what they can.
- 06 **Verification and disclosure rules.** Professional-conduct rules on checking and disclosing AI-assisted work create a new human task, verification, that often lands on paralegals.

# Impact by sector

The headline figure is an average. Where you work changes the picture.

|                             |                                                                                                                                                                                |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Large law firms             | Highest adoption, with enterprise licences for Harvey, CoCounsel and e-discovery AI; paralegals here are already supervising output rather than producing it.                  |
| Small and solo practices    | Adoption is slower and more informal, often via general chatbots and built-in features of Clio or similar; the paralegal frequently becomes the firm's de facto AI specialist. |
| Corporate legal departments | Contract review and extraction tools are widely deployed, so in-house legal assistants spend more time on matter intake, vendor management and compliance tracking.            |
| Government and legal aid    | Budget and procurement constraints mean slower deployment, but high caseloads make these offices strong candidates for AI summarisation once governance questions are settled. |

# Skills to build

The skills that keep the human part of this work valuable as the routine part is automated.

- 01

**Source verification.** The ability to check an AI-generated authority or summary against the primary record quickly and reliably is now the core quality-control skill. Practise by auditing tool output against the source until it is habit.
- 02

**E-discovery and review protocol design.** Knowing how to set up a review, train a classifier and run defensible quality-control samples moves you from reviewer to project lead. Vendor certifications from Relativity and similar platforms are the usual route.
- 03

**Court procedure and jurisdictional rules.** Local filing rules, service requirements and deadline calculation remain human work and are where errors are most costly. Build a personal reference and keep it current.
- 04

**Prompting and tool configuration.** Getting a useful first draft out of a legal assistant depends on supplying the right matter context and constraints. Learn each tool's templates and limits rather than typing free-form requests.
- 05

**Client communication.** Explaining process, managing expectations and gathering facts from anxious clients is not automatable and is valued by lawyers who would rather not do it. Ask for more of this work.

06 **Data handling and confidentiality.** Understanding what can and cannot be put into which tool, and how client data is stored, is now a professional obligation. Read your firm's AI policy and the vendor's data terms.

§ 07 INSTRUMENTS

6 ENTRIES

## Tools in use

KINDS OF TOOL WORTH KNOWING

- 01 **Harvey.** Enterprise legal AI used mainly by large firms and in-house teams for research, drafting and contract analysis.
- 02 **General-purpose assistants (ChatGPT, Claude, Microsoft 365 Copilot).** Common in small firms for summarising and drafting non-confidential material; check your firm's policy before entering client data.

NAMED TOOLS ALREADY IN USE

|                                                                                                                                               |       |
|-----------------------------------------------------------------------------------------------------------------------------------------------|-------|
| <b>Clio Duo</b>                                                                                                                               | VISIT |
| Generative assistant inside the Clio practice-management platform that drafts from the matter file and answers questions about case data.     |       |
| <b>CoCounsel</b>                                                                                                                              | VISIT |
| Thomson Reuters' legal AI assistant for research, document review, summarisation and drafting, widely licensed by firms.                      |       |
| <b>Lexis+ AI</b>                                                                                                                              | VISIT |
| LexisNexis research platform with conversational search, drafting and summarisation linked to its case-law and statute database.              |       |
| <b>Relativity aiR</b>                                                                                                                         | VISIT |
| Generative AI review in the Relativity e-discovery platform that classifies documents for relevance and privilege and explains its reasoning. |       |

§ 08 EXAMPLES

3 EXAMPLES

## In practice

Ways people in this role are already using AI, and what they get from it.

|                                               |      |                                                                                                                                                                                                                 |
|-----------------------------------------------|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Medical record chronology</b><br>EXAMPLE 1 | HOW  | A personal-injury paralegal uploads several hundred pages of treatment records to a summarisation tool, which produces a draft chronology that the paralegal then checks against the source pages and corrects. |
|                                               | GAIN | A task that took days is finished in hours, with the paralegal's time spent on accuracy rather than transcription.                                                                                              |

## First-draft discovery responses

EXAMPLE 2

HOW

A litigation assistant uses CoCounsel to generate draft responses and objections to a set of interrogatories from the matter file, then tailors them with the supervising lawyer.

GAIN

Lawyer review starts from a structured draft instead of a blank page, shortening turnaround on routine discovery.

## Prioritised document review

EXAMPLE 3

HOW

On a large production, the paralegal sets up Relativity aiR to rank documents by likely relevance, reviews the top tier personally and samples the rest for quality control.

GAIN

The team reviews a fraction of the documents by hand while maintaining a defensible record of the process.

## § 09 CONTEXT

# How this role compares

Three neighbouring roles chosen to show the direction of travel.

## Clerical Assistants

MORE EXPOSED

AI IMPACT

General clerical work such as data entry, scheduling and document filing is being automated more directly than paralegal work, which carries more judgement and procedure.

WORK MOVES TO

Clerical roles are moving towards exception handling and system oversight rather than routine processing.

## Compliance Officers

DIFFERENT SKILLS, GROWING

AI IMPACT

AI monitors transactions and policies but regulatory interpretation and accountability remain human, and demand is rising as rules on AI itself multiply.

WORK MOVES TO

Paralegals with a regulatory practice background can move into compliance roles that pay more and are growing.

## Lawyers

COMPLEMENTARY, LESS EXPOSED

AI IMPACT

AI drafts and researches for lawyers as well, but advocacy, strategy and professional accountability for advice stay with the qualified practitioner.

WORK MOVES TO

The paralegal-lawyer relationship is shifting towards the paralegal managing tools and process so the lawyer can focus on judgement and client strategy.

## Closing judgement

If you are a paralegal, the drafting and research that filled your early career is being handed to software, but the organisation, verification and client work that makes a matter run is not. The paralegals who do well in this period will be the ones who learn the firm's AI tools properly, become the person who catches the hallucinated citation before it reaches a judge, and take on more of the case-management and client-facing work that lawyers are happy to delegate. Treat the technology as a junior you supervise, not a rival, and your position in the firm gets stronger rather than weaker.

## Evidence and revisions

What the exposure figure rests on, what changed when it was last revised, and the published work cited for this role.

SCORE

40

WINDOW

4-9 years (unchanged)

The 5 October 2026 review held the score.

Exposure Index v2. Inputs: task applicability 11/100 (Microsoft AI applicability score 0.06 for Paralegals and legal assistants); observed usage 39/100 (Anthropic observed exposure 0.29); official exposure tier 70/100 (BLS: high); labour-market trajectory not yet mapped for this occupation, so its weight was spread across the other inputs; published adoption rating 70/100 (high adoption). Weighted base 40.3. Final score 40. New report: the window of 4-9 years is set from the score band.

HOW THE FIGURE IS BUILT (EXPOSURE INDEX V2)

|                                                                       |              |      |
|-----------------------------------------------------------------------|--------------|------|
| Task applicability — Microsoft Research, AI applicability score       | 11 × 39%     | 4.4  |
| Observed usage — Anthropic Economic Index, observed exposure          | 39 × 22%     | 8.7  |
| Official exposure tier — US BLS AI-exposure category                  | 70 × 22%     | 15.6 |
| Labour-market trajectory — US BLS projected employment change 2025–35 | not measured | —    |
| Published adoption rating — This report's adoption level              | 70 × 17%     | 11.7 |
| Weighted base                                                         |              | 40.3 |
| Exposure score                                                        |              | 40   |

MEASURES BEHIND THE SCORE

**US Bureau of Labor Statistics** · *Employment Projections 2025–35 and AI Exposure Categories* (27 August 2026)

AI-exposure tier: High. Projected employment change not yet mapped for this occupation. Matched to Paralegals and legal assistants.

PUBLISHER PDF ARCHIVED COPY

**Microsoft Research** · *Working with AI: Measuring the Applicability of Generative AI to Occupations* (10 July 2025)

AI applicability score 0.06 for SOC 23-2011; scaled to 11/100 as the task-applicability input.

PUBLISHER PDF ARCHIVED COPY

**Anthropic** · *Anthropic Economic Index report: Cadences* (26 June 2026)

Observed exposure 0.29 for SOC 23-2011; scaled to 39/100 as the observed-usage input.

PUBLISHER PDF

**UK Department for Science, Innovation and Technology** · *Assessment of AI capabilities and the impact on the UK labour market* (28 January 2026)

UK context: around 70% of UK workers are in occupations with tasks AI could perform or enhance, above the US average; a one-standard-deviation rise in exposure was associated with a 3.9% fall in UK job postings.

PUBLISHER ARCHIVED COPY

Archived copies are served only where the licence permits; otherwise the link goes to the publisher. Full research library:

<https://www.careerguard.ai/sources>

## § 12 APPENDIX

## Method and sources

Each report was written from a large body of published research. The exposure score itself is computed, not written: it is the CareerGuard Exposure Index, a weighted average of occupation-level measures from the US Bureau of Labor Statistics (AI-exposure classification and 2025–35 projections), Microsoft Research (AI applicability scores) and Anthropic (observed exposure), together with the adoption rating published on the report. The organisations and publications below are the standing literature behind the narrative sections. Every source, with dates, licences and archived copies where we are permitted to hold them, is catalogued in the research library.

Exposure Index v2 (October 2026). Each input is scaled to 0–100 and weighted: task applicability 35% (Microsoft AI applicability score ÷ 0.5), observed usage 20% (Anthropic observed exposure ÷ 0.75), official exposure tier 20% (BLS very high = 100, high = 70, moderate = 40, low = 10), labour-market trajectory 10% ( $50 - 2.5 \times \text{projected \% employment change}$ ), published adoption rating 15% (very high = 85, high = 70, medium-high = 55, medium = 40, low-medium = 25, low = 10). Inputs not measured for an occupation are dropped and the remaining weights renormalised. An editorial adjustment of at most  $\pm 12$  points is allowed only for automation channels the measures cannot see (robotics, self-service, machine vision, medical imaging, RPA/OCR, generative video) and is always logged with its reason. Scores are whole numbers, not rounded to five. The change window shifts one notch (a year at each end) per ten points of movement.

---

**World Economic Forum**

*The Future of Jobs Report series* — Employer survey of expected job growth and decline, skill shifts and technology adoption (2020, 2023 and 2025 editions); *Artificial Intelligence and the Future of Entry-Level Work* (2026).

*AI governance and transformation reports* — Frameworks on ethical AI, talent strategy and industry transformation.

---

**McKinsey Global Institute**

*AI, Automation and the Future of Work series* — Research quantifying automation potential by task, sector and demographic, from "Jobs Lost, Jobs Gained" to "Agents, robots, and us" (2025).

*Industry-specific reports* — Financial services, healthcare, manufacturing and others.

---

**PwC**

*Global AI Jobs Barometer* — Annual analysis of job postings and productivity by AI exposure (2024–2026 editions).

*Upskilling Hopes and Fears survey* — Employee perceptions and readiness.

---

**Microsoft Research and Anthropic**

*Working with AI* (2025); *Anthropic Economic Index* (2025–2026) — Occupation-level usage data from Copilot and Claude conversations, the two observed-usage measures behind the 2026 revision.

---

**Stanford Digital Economy Lab and Stanford HAI**

*Canaries in the Coal Mine?* (2025–2026); *AI Index Report* (annual) — Payroll evidence on early-career employment in exposed occupations; annual measurement of AI capability, investment and adoption.

---

**Deloitte**

*Human Capital Trends series* — Workforce, talent and HR technology trends.

*Tech Trends series* — Emerging technologies and their business implications.

---

**Accenture**

*Technology Vision series* — Forward-looking analysis of emerging technology, with emphasis on AI.

*Fjord Trends* — Design, innovation and human experience in a digital world.

---

**Boston Consulting Group**

*AI/ML insights and industry solutions* — "The AI Revolution in the Workplace" and related research.

---

**EY**

*AI and workforce reports* — Adoption, talent strategy and ethics.

---

**IBM Institute for Business Value**

*AI and automation studies* — Business models, workforce evolution and leadership.

---

**OECD**

*AI Policy Observatory* — International data and policy on AI, labour markets and skills.

*Employment Outlook* — Labour-market trends including technological impact (2023–2026 editions).

---

**International Labour Organization**

*Generative AI and Jobs: A Refined Global Index of Occupational Exposure* (2025) — Task-level exposure gradients for every ISCO occupation; successor to the 2023 global index.

---

**International Monetary Fund**

*Staff Discussion Notes on AI and work* (2024, 2026) — Complementarity framing: where AI augments and where it substitutes.

---

**UK Department for Science, Innovation and Technology**

*Assessment of AI capabilities and the impact on the UK labour market* (2026) — UK occupational exposure and job-posting evidence.

---

**Brookings Institution**

*AI and automation research* — Economic and social implications, displacement and skills.

---

**Yale Budget Lab and Goldman Sachs Research**

*Tracking the Impact of AI on the Labor Market; AI and the US labour market* (2026) — Aggregate labour-market monitoring; macro displacement estimates.

---

**Oxford University (Oxford Martin School)**

*The Future of Employment* — Frey & Osborne and subsequent research on susceptibility to automation.

---

---

## MIT Technology Review

*AI & Work* — Reporting on AI research and its implications for industries and jobs.

---

## Gartner

*Hype Cycle for Artificial Intelligence* — Maturity and adoption of AI technologies.

*Future of Work reports* — Workplace models and talent strategy.

---

## U.S. Bureau of Labor Statistics

*Employment Projections 2025–35; AI Exposure Categories; Occupational Outlook Handbook* — Ten-year employment projections and, from the 2025 cycle, an AI-exposure tier for every detailed occupation.

---

## Indeed Hiring Lab

*AI at Work Report (2025) and posting-market updates* — Skill-level transformation estimates and job-posting trends by occupation.

---

---

# II CORE AI AND MACHINE-LEARNING RESEARCH

---

## OpenAI

*Research papers, blog and API documentation* — Large language models, generative AI, safety and societal impact.

---

## Google DeepMind

*Research papers and blog* — Reinforcement learning, AI for science, AGI and ethics.

---

## Meta AI

*Research papers and blog* — Large language models, computer vision, AI for social good.

---

## Hugging Face

*Transformers library and model hub* — Open-source state-of-the-art NLP models.

---

## TensorFlow and PyTorch

*Documentation and community forums* — Core frameworks illustrating practical capability.

---

## arXiv

*cs.AI, cs.LG, cs.CV, cs.CL* — Pre-print research.

---

## NeurIPS and ICML

*Conference proceedings* — Top-tier academic research.

---

## ACM and IEEE

*Journals and proceedings* — ACM Computing Surveys; IEEE Transactions on AI.

---

## Kaggle

*Datasets and competition solutions* — Applied machine learning on real-world problems.

---

## The Alan Turing Institute

*Research and reports* — Responsible and applied AI.

---

---

# III ETHICAL AND RESPONSIBLE AI DEPLOYMENT

---

## NIST

*AI Risk Management Framework* — Voluntary framework for managing AI risk.

---

## European Commission

*AI Act* — Risk-tiered legal framework for AI.

*Ethics Guidelines for Trustworthy AI* — Principles for responsible development.

---

## Partnership on AI

*Research and best practice* — Responsible AI development.

---

## AI Now Institute

*Annual reports* — Social implications: power, inequality, rights.

---

---

**ACM FAccT**

*Proceedings* — Fairness, accountability and transparency.

---

**Data & Society**

*Publications* — Social implications of data-centric technology.

---

**WIPO**

*Conversation on IP and AI* — Intellectual-property implications of AI.

---

**IEEE Global Initiative on Ethics of AI/IS**

*Ethically Aligned Design* — Recommendations for ethical AI design.

---

**Center for AI and Digital Policy**

*Policy briefs* — Accountable AI policy.

---