

# COST EFFECTIVE



## ANDREW HOGAN EXPLAINS HOW TO USE AI IN THE THIN MARGIN ENVIRONMENT OF FIXED COSTS WORK

**Low-value claimant personal injury work now operates under relentless economic pressure. Fixed recoverable costs under CPR Part 45 reward speed, consistency and disciplined process, but leave little margin for wasted labour, duplication or error. In portal, Official Injury Claim (OIC) and small-claim environments, the difference between profit and loss is often measured in minutes per file. It is against that background that artificial intelligence has attracted so much interest from PI firms.**

The appeal is obvious. AI promises to reduce staff time, standardise drafting, chase evidence, manage deadlines and keep files moving. But the real question for a regulated PI practice is not whether AI can generate text or summarise documents; it is whether a firm can redesign the entire lifecycle of a low-value PI file so that human lawyers intervene only where legal judgment, valuation, ethics, client protection or regulatory responsibility genuinely require it.

That question has become sharper since the SRA authorised Garfield.Law as the first AI-based law firm in England and Wales. But personal injury presents a harder and riskier mix than many other areas of practice: medical evidence, causation, valuation, credibility and vulnerable clients sit at the centre of the work. The challenge is therefore not to build a 'robot solicitor', but to construct a production system in which automation handles routine tasks, while qualified lawyers retain ownership of the decisions that matter.

While AI can already handle much of the administrative, drafting, extraction, summarising, prompting, and workflow-management work in low-value PI, it cannot safely replace qualified human oversight at the points where liability, causation, limitation, quantum judgment, settlement advice, vulnerable clients, unusual facts, privilege, or court-facing accuracy are in play.

The commercially realistic model, therefore, is not a robot solicitor with no adults in the room. It is a highly automated production line with carefully designed human checkpoints, audit trails, escalation rules and supervision. That conclusion is reinforced by current SRA guidance on technology and supervision, and by ICO guidance that AI use must still comply with UK GDPR and be explainable to affected individuals.

**The difference between profit and loss is often measured in minutes per file**

**The key facts must be captured accurately into the case management system**

## The margin problem

The economics of low-value personal injury work are now stark. CPR Part 45 provides the fixed-cost structure, while PD 45 contains the relevant tables for RTA, EL/PL and fast-track cases. But the commercial reality varies sharply from one category of case to another. An OIC whiplash claim in the small claims environment is not the same thing as a portal EL claim, and neither is it the same as a matter that drops out into ordinary fast-track litigation. FRCs reward process discipline, speed and consistency. Margin is steadily consumed by staff time, supervision, software, medical agency fees, compliance burdens and acquisition cost. The capped success fee helps, but only up to a point, and is no cure for inefficiency.

That is why AI has become so attractive. Its promise is not that it performs legal judgment, but that it compresses routine labour, standardises output, and makes thin-margin work economically survivable.

### The solicitor's role

When people talk about PI practice there is a tendency to describe the solicitor's job as 'doing the law'. In truth, a low-value PI file is mostly a chain of tasks. Some of these call for legal judgment, but many do not. That matters, because if one is asking whether AI can profitably change the economics of this work, one must first be clear what the work actually is.

It begins with intake. The firm has to capture the source of the enquiry, identify the accident type, date and place, the prospective defendant, and the broad nature of the injury. Somebody has to form an initial view on eligibility, likely value, limitation, and whether there are any obvious warning signs: fraud indicators, inconsistent account, contributory negligence, foreign elements, vulnerable claimant, or a claim that plainly belongs in the wrong process. That routing question is now fundamental. A low-value road traffic claim may point towards the OIC service, whereas other low-value RTA, EL and PL claims may still belong in the Claims Portal, and some matters will fall outside both and require ordinary pre-action handling.

Then comes opening the file properly. Identity must be verified, and AML or KYC checks may be needed in the background. The retainer must be put in place. The CFA, client care papers, authorities, privacy information and deductions explanation must all be prepared and explained. The key facts must then be captured accurately into the case management system, because from that point onwards the file will live or die by the quality of its data. A lazy opening usually produces a disorderly case later on.

After that, the solicitor moves into evidence and analysis. The client's account has to be reduced into usable form. Witness details, photographs, receipts, earnings evidence, repair material and other documents must be obtained and chased. Medical records may need to be requested. A medical expert must be chosen and instructed. The resulting report must then be read with

care, not merely received. Someone needs to decide if the prognosis makes sense, if further evidence is needed, and whether the claim remains routine or is developing awkward features.

The file will proceed through its external process: CNF or equivalent portal entry, deadlines, liability response, offers, counter-offers, valuation, and if necessary, preparation for Stage 3, disposal hearing, or trial. At the end come settlement advice, billing, any success fee deduction, client accounting, reporting and closure. CPR Part 45 and PD 45 supply the fixed-cost framework against which all this labour must be made to pay. That is the real point. The file is not a single legal act. It is a production process, punctuated by moments of judgment.



**From that point onwards the file will live or die by the quality of its data**

### What AI can do at each stage

The question is not whether a machine can 'run a claim'. That is too grand, and too vague. The right question is narrower. At each stage of the file, what can be automated? What can be assisted? What must still be checked? And what should never be left to the machine at all? That is the practical test.

At the intake stage, AI is already very useful. It can run structured questionnaires, turn a rambling account into a usable attendance note, identify missing dates or documents, classify the apparent claim type, and suggest the next process step.

It is also good at flagging what is absent: no accident location, no registration number, no defendant details, no injury duration, no earnings information. That is administrative automation - precisely the sort of work that could be removed from expensive human hands. But it should not be trusted, without supervision, to decide that a claim is sound, that limitation is safe, or that a file belongs in OIC rather than the Claims Portal or ordinary pre-action handling. That routing decision may look clerical. It is not.

Onboarding is similar. AI can populate the retainer pack, draft the client care letter, explain the process in plain English and route the matter into the correct workflow. It can extract data once and use it many times. That is valuable, reducing duplication. But a machine should not be left to decide if the funding arrangement has been properly explained, if the client has really understood deductions, or if vulnerability requires a different approach. The solicitor must still own that part of the relationship. The SRA's point, put shortly, is that technology may be used, but the firm remains responsible for competence, systems and control.

Evidence handling is where AI becomes particularly attractive. It can create OCR (optical character recognition) records, rename and tag documents, sort photographs,

build chronologies, extract treatment dates, and summarise large bundles more quickly than most junior staff. It can also compare documents and flag inconsistency: the wrong date in the draft CNF, a mismatch between the client's account and the medical note, or a missing page in a records bundle. Used properly, that is not legal judgment at all. It is disciplined file housekeeping of a very high order. But once the question becomes whether a prior symptom breaks the chain of causation, whether an injury pattern is plausible, or whether the expert's prognosis should be accepted at face value, the machine has reached the edge of its proper role.

The same pattern appears in drafting, negotiation and trial preparation. AI can produce first drafts of letters, CNFs, witness statements, schedules, offers, chronologies, issue lists and hearing notes. It can suggest valuation ranges and response options; and it can assemble bundles and chase deadlines relentlessly. It is also very good at exception spotting: missing documents, slippage in portal deadlines, inconsistent dates, mismatched facts.

What it should not do is advise finally on settlement, advise on liability, consider contested causation, or put court-facing assertions into circulation without review. Those tasks remain the lawyer's work. The real value of the solicitor is now found not in typing the first draft, but in knowing when the routine case has ceased to be routine.

### Using general-purpose AI tools

How should PI lawyers use general-purpose tools such as ChatGPT, Claude, Gemini, Copilot and related tools? For a claimant PI firm, the starting point should not be that general-purpose tools are usable unless shown to be forbidden. It should be the reverse. They are presumptively unsafe for live client material unless the firm has put in place a proper enterprise deployment, a processor contract, a defensible UK GDPR basis for the processing, a special-category condition for the medical data, and governance strong enough to satisfy both the regulator and the firm's own conscience. The SRA's own risk material points directly at the problem. It identifies, as a specific threat, a staff member using an online AI such as ChatGPT to answer a question on a client's case, and says firms must make sure their use of AI protects confidentiality and legal privilege. It also says some firms prohibit all staff use of online systems beyond their direct control, while others impose separate rules for confidential information depending on the type of AI used. That is a long way from saying that ordinary public tools are just another office convenience.

Continued on P34



**The question is not whether a machine can 'run a claim'. That is too grand, and too vague**

That matters even more in personal injury than in many other practice areas, because the ordinary PI file is full of sensitive data. Medical records are special category data. The ICO makes clear that if you process special category data you must identify both a lawful basis under Article 6 and a separate condition under Article 9; if you cannot identify a condition, you cannot process the data at all. The ICO also says the processing must be necessary in a targeted and proportionate sense, not merely useful, convenient, or part of your preferred business model. It adds that high-risk processing requires a Data Protection Impact Assessment (DPIA), and that special category data makes a DPIA more likely, especially where the data is used to determine access to a service or benefit. In other words, a solicitor cannot simply say, 'I uploaded the client's medical records into ChatGPT because it saved time'. Convenience is not the test. Necessity, proportionality, safeguards and accountability are the tests.

### That is not legal judgment at all; it is disciplined file housekeeping of a very high order

#### Bespoke legal AI and PI workflow software

The sensible future for claimant PI is not a solicitor pasting medical records into a public chatbot. It is a controlled software stack built around the life of the file. In practice, that means a case management system at the centre, with an intake portal, OCR and document parsing, a secure document store, a retrieval layer pointing the model to the firm's own precedents and workflows, a rules engine for limitation, portal routing and missing-document checks, and human approval gates before anything important leaves the building.

The software would not be asked to 'do personal injury law' in the abstract. It would be asked to perform a narrow sequence of tasks: classify the claim, build the chronology, draft the retainer pack, prepare the portal narrative, summarise the medical evidence, flag inconsistencies, and escalate anything unusual. That is much more realistic, and much safer.

As to the model itself, there are really two routes. The first is to use an enterprise model through a managed provider; the second is self-host an 'open-weight model' which sits inside the firm's own environment.

### Once the question becomes whether an injury pattern is plausible... the machine has reached the edge of its proper role

## Building an in-house AI stack

If a claimant firm wanted to build its own AI-driven PI workflow, the sensible design would be modular. The existing case management system could (potentially) remain the system of record. Around it, the firm would build an intake portal for clients, a document-ingestion layer, a retrieval layer, an LLM layer, and a workflow engine.

The workflow layer would then push the result back into the CMS as draft attendance notes, chronologies, letters, task lists and exception flags. On a managed-cloud route, the firm would not need special fee-earner machines at all. Ordinary locked-down business laptops would do, because the heavy lifting sits in the cloud.

If the firm wants tighter control it can self-host an open-weight model, although the technology needed for this is very substantial. The real work, however, is not model hosting, but control. The firm would need an evaluation set of closed PI files, prompt libraries, approval rules, logging, cyber attack testing, role-based permissions, and a sampling regime for supervision. The workflow should be designed so that the model can draft, classify, summarise and flag, but cannot silently send, settle, or submit.

A good build would therefore include explicit human approval stages for client-facing advice, medico-legal instructions, settlement recommendations, witness statements and anything court-facing. The software can reduce labour, but must not be allowed to reduce responsibility.

### The real value of the solicitor is now found in knowing when the routine case has ceased to be routine

## Agents: what they are, and what they are not

An 'agent' in this setting is not just a chatbot that answers a single prompt; it is a software component that uses an LLM, has access to tools and external data, and can take a sequence of steps in pursuit of a goal. Microsoft's current description of an agent is an AI application that can reason about a request and take autonomous actions, including calling tools and accessing external data.

In the architecture described above, the agent(s) therefore sits above the model. The model provides language and reasoning, while the agent provides orchestration. It is the thing that decides, 'I have received a new PI enquiry; I must collect the missing accident details, open the intake workflow, request the authority forms, prepare the chronology, and route the file for approval'.

Used in that disciplined way, agents could add real value to a claimant PI practice. They are well suited to the dull but essential work that bleeds margin out of fixed-cost files: chasing unsigned retainers, prompting clients for photographs and wage slips, booking medico-legal appointments, updating the case management system, drafting routine letters, identifying missing records, and escalating a file when limitation is near or the medical material does not match the pleaded account. In the self-hosted version, it would sit on top of the firm's own retrieval layer and model server and do the same jobs, but within the firm's own controlled estate.

But an agent is not a synthetic fee-earner. Its strength is process, not judgment. It can move a file through a known sequence. It can break down badly when facts are incomplete, documents conflict, the case falls between process lanes, or the model confidently misreads the evidence.

Agents also have a particular failure mode of their own: because they can call tools and take several steps, an early mistake can propagate. A wrong classification at intake can cause the wrong workflow, the wrong portal path, the wrong deadlines and the wrong correspondence. The more autonomy one gives the agent, the more one amplifies that risk. Microsoft's own recent agent materials emphasise workflows, checkpointing and human-in-the-loop support for multi-step tasks. That is an admission of reality, not a weakness.

That is why agents fit the vision only if they are made servants of the system, not masters of it. They should be allowed to collect, sort, prompt, draft and flag. They should not be allowed silently to settle a claim, choose between competing medical interpretations, or make final court-facing assertions. The SRA has already warned firms to distinguish between casual use of online AI and formally adopted systems, and to make sure supervision and guidance reflect that difference.

In a regulated PI practice, an agent is best understood as a workflow clerk with tireless energy and no instinct for danger. It can save time, but cannot safely replace the lawyer who knows when the routine case has ceased to be routine.

### One flawed prompt, one bad workflow or one unsound assumption can contaminate hundreds of files

## Risks, constraints and regulatory limits

The case for AI in low-value PI is easy to overstate. The harder and more important exercise is to identify where it goes wrong.

The SRA is not banning firms from using AI, but it is making two things plain. First, the ordinary principles and standards still apply. Secondly, responsibility stays with the firm. The code of conduct for firms requires competent service, competent staff, and an effective system for supervising clients' matters. The SRA's supervision guidance adds that supervisors should have oversight of live work at key stages, should know each matter or monitor a meaningful sample, and should be able to give guidance on non-standard issues. In other words, a firm cannot shelter behind the phrase 'the AI did it'.

#### Data protection

A PI file is not ordinary commercial paperwork. It commonly contains identity documents, financial material and medical evidence, which is special category data. The ICO's guidance says AI deployment is a processing activity which requires a lawful basis, and where special category data is involved, an additional condition as well. It also says the organisation must separate the distinct processing purposes, identify controller and processor relationships, document its decisions, and in high-risk cases use a DPIA.

### Many PI clients need explanation, reassurance and judgment, not merely polished text

#### Confidentiality and privilege

The true danger is not simply what a vendor promises in a privacy notice. It is what the firm is actually doing with the material. The SRA has specifically identified as a threat a staff member using an online AI such as ChatGPT to answer a question on a client's case, and it warns firms to distinguish between casual use of public online tools and systems they have formally adopted and controlled. In PI work, where medical records and instructions may be privileged and plainly confidential, careless uploading is a regulatory own goal.

#### Hallucination and weak reasoning

A fabricated date, a non-existent case, or a neat-but-wrong summary of a medical report is intolerable in claimant litigation. Models often smooth over ambiguity rather than confront it. That matters most where liability is messy, causation is mixed, prior symptoms complicate valuation, or limitation is tight. These are exactly the cases in which automation looks confident and is most dangerous. The SRA expressly warns firms not to trust AI to judge its own accuracy.

#### Bias, outliers and client care

Chronic pain, overlapping symptoms, exaggeration indicators, vulnerable clients and unusual accident mechanisms are where process systems misfire. Scale then magnifies error. One flawed prompt, one bad workflow or one unsound assumption can contaminate hundreds of files.

The court risk is obvious: inaccurate witness statements, schedules or submissions damage credibility at once. The client-care risk is just as real. Many PI clients need explanation, reassurance and judgment, not merely polished text. A machine can help move the file. It cannot be trusted to carry the professional burden of representing the person.

## Can this model really work?

AI has real potential to improve the viability of low-value claimant PI work, and in some practices it may be the only realistic way of preserving margin under fixed recoverable costs. Much of a PI file consists of structured, repetitive process: intake, data capture, document handling, drafting, chasing and workflow control. In straightforward cases, AI can already perform or assist with a large proportion of that work more quickly, more cheaply and often more consistently than human staff.

But it would be a serious mistake to treat that capability as a licence for autonomous practice. The weakest points of current AI systems are precisely the points at which PI work stops being routine: disputed liability on untidy facts, blurred causation, non-standard valuation, limitation risk, and final settlement advice. These are not production tasks. They are exercises of judgment, responsibility and professional skill. No current system can safely take ownership of them, and regulators are clear that responsibility remains with the firm.

The sustainable model for claimant PI is therefore not replacement, but redesign. Firms should use AI to remove lawyers from routine production work and redeploy them upwards, into supervision, exception handling, valuation, client advice and court-facing accuracy. That requires disciplined process mapping, clean data, standardised prompts and templates, clear escalation rules, and visible human checkpoints. It also requires business-grade systems with proper governance, audit trails and data-protection safeguards, not casual use of consumer tools.

The firms most likely to succeed will not be those with the most impressive chatbot. They will be the firms with the best process design: those that understand where automation adds value, where it introduces risk, and where only a human lawyer will do.

AI can make low-value PI economically survivable. It cannot relieve solicitors of responsibility. In a fixed-cost world, the future belongs to supervised automation, not to the illusion of autonomous law.

### The weakest points of current AI systems are precisely the points at which PI work stops being routine

Andrew Hogan is a barrister at Hailsham Chambers, London and is author of a beginners guide to AI, 'Andrew and the Marvellous Analytical Engine', available on Amazon