

Industrializing Without Losing the Soul: What AI Changes and What It Cannot Replace


BY SIRISHA GUMMAREGULA AND DANIELLE NOONAN

June 1, 2026

AI has the capability to produce a near-infinite sea of output. That abundance is both the source of its value and the source of its governance challenges. The question for organizations, therefore, has matured from whether AI can generate a particular work product to whether the correct questions are being asked, how to channel its abundance into valuable work, and whether the result can be trusted.

Access to near-infinite knowledge is useless without the right question. Like a vast body of water without a channel, AI is abundant, powerful, and capable of carrying immense value, but also formless. Without direction, it can flow, flood, stagnate, or overwhelm an organization in plausible output. And as with channeled waterways, abundance does not direct itself; direction comes from human judgment and expertise.

Much of the market conversation is currently focused on which AI tool produces the best output. While that conversation is important, it is incomplete. The more durable questions are: who is directing the work? Who is framing the problem? Who is deciding what matters? Who is testing whether the result is complete, consistent and reliable? Who is asking whether the output should exist at all?



Credit: Nadia/Adobe Stock

Bruges, the medieval Flemish trading city whose preserved historic center still carries the texture of its golden age, offers a useful way to think about this. Its early wealth depended on water: access to the North Sea through a managed network of waterways helped make Bruges a center of trade, finance and artistic production. When that access diminished and commerce shifted elsewhere, Bruges lost commercial centrality, reducing the pressure to rebuild the city during later waves of industrial modernization. Over time, the same forces that diminished Bruges commercially helped preserve it architecturally.

It would be easy to turn that history into a cautionary tale about missing the industrial revolution. But Bruges' value does not lie in being bypassed by modernization. It lies in what endured: not the

city's medieval buildings alone, but the artisan-ship, stewardship and human judgment they still reveal. Unlike Bruges, organizations cannot leave preservation to circumstance.

Organizations should embrace the AI revolution in a way that sharpens human judgment rather than displacing it. AI deployed without that discipline produces a flood, not value. The goal is to use these systems to their full potential without allowing the human capacities that give them direction to atrophy. AI can industrialize production, but only disciplined human judgment can determine what is worth producing and whether the result can be trusted.

AI increases the volume and speed of production. It makes work faster, more scalable and more accessible. But scale is not value. A flood of competent output is still a flood. The scarce resource in an AI-enabled organization will not be access to information. It will be the quality of the judgment that gives form to information.

Direction, however, is only the first act of judgment. Once judgment gives the flow direction, it must also determine when the surface answer is insufficient and where the system must be pushed to excavate what is hidden beneath it. That is one of AI's most important possibilities. Used well, AI can surface the thought beneath the thought: the risk that is almost visible, the pattern distributed across documents, the inconsistency that no single reviewer would easily detect, the argument that has not yet found its form. But that excavation is not automatic. The tool surfaces what it is pointed at. Human judgment decides where to dig, what the finding means, and whether it matters.

In a complex acquisition, for example, an AI system may be deployed to review hundreds of agreements for material contracts, change-of-control triggers, and indemnification obligations. One facilities management agreement may contain a revenue-sharing obligation buried in a

schedule and framed in commercial rather than legal terms. The system may generate a summary that correctly captures the parties, term, and termination provisions, while failing to surface the revenue-sharing obligation because it falls outside the effective scope of the instruction framing. The output looks complete. The agreement appears in the index. No error is flagged. No gap is signaled. But the deal team is still operating with an incomplete risk picture. The practitioner who asks the better question, and who tests not only what the system found but what it may have failed to engage, can surface an exposure that output review alone would miss.

That is the productive tension at the heart of AI-enabled work. The value does not sit in output alone. It sits in the intelligence that directed the inquiry and judged the result. The risk is that, as AI handles more execution, the human capacity for judgment erodes quietly. Human review can become passive and deferential, less capable of detecting what the system has missed. Judgment does not fail all at once. It rusts. Review becomes a formal step rather than a substantive act. Supervision becomes an approval workflow. The organization retains the language of judgment while losing the practice of it.

In litigation, the cost of passive review can be especially high. An AI system may be deployed across a large document universe with a clear case theory, scope definition, risk thresholds, and escalation logic established at the outset. In a multi-step workflow, the original instruction context can become diluted as the system processes more information. By the end, the outputs may appear coherent while no longer reflecting the analytical framework defined at the beginning. The lawyer who tests the workflow at defined checkpoints against the original task definition can catch that drift before it compounds. The lawyer who merely reviews the final outputs may not.

These examples show why human judgment cannot be treated as a final-stage review function. Preserving judgment in AI-enabled work requires operational discipline, not individual instinct after the fact. Organizations must build workflows that develop, require, and test human judgment alongside AI capability. That means defining the task before the tool is used, preserving the original instruction framework throughout the workflow, testing for omissions and boundary failures, and requiring escalation where an output appears complete but the risk profile has not been validated. The point is not ceremonial presence “in the loop.” It is preserving the capacity to know what the loop is for.

This is where the Bruges analogy returns. Bruges’ value is not that it stayed untouched or pure. It is that something human, specific and difficult to reproduce endured. Its preserved beauty is not simply oldness. It is the accumulated result of artisanship, constraint, and judgment. Those are not qualities industrialized systems naturally preserve. They require deliberate protection.

The AI age will intensify the pressure to prioritize speed, scale, and volume over artisanship, constraint, and judgment. As production becomes abundant, discernment becomes the scarce resource. The question for organizations is not simply how much more they can produce with AI. It is whether they have built the discipline to keep human judgment active inside the process, making the work better than output alone.

That discipline is artisanal intelligence: the disciplined application of human judgment to ensure AI-enabled legal work is complete, consistent, and reliable. It is not nostalgia, nor is it resistance to technology. It is what turns AI-enabled output into

work that can be trusted. Bruges reminds us that abundance can create a golden age, but abundance alone does not decide what will endure. Water brought the world to the gates of Bruges. Human judgment built a city of lasting value.

AI will bring organizations more capability than they have ever had. The tool produces. The artisan decides what is worth producing, and whether it can be trusted.

Sirisha Gummaregula is CEO and founder of QuisLex, a leading alternative legal services provider with 22 years of experience in complex legal and compliance work for Global 500 corporations and Am Law 100 and Magic Circle firms. QuisLex holds a U.S. patent for measuring quality in legal workflows, was the first legal service provider to obtain ISO 9001 quality management certification, and has applied Six Sigma methodology to legal services since its founding.

Danielle Noonan is an associate vice president, legal services at QuisLex. A former litigation associate at Kaye Scholer, she now advises in-house legal teams on defensible protocols across the EDRM cycle, including applying AI to legal review and investigation workflows. She writes on AI in litigation, e-discovery, and legal ethics, with an emphasis on human judgment, validation, and defensibility.

