



At the time Joseph Becker wrote his article, "The Computer," for *Studies in Intelligence*, the DEC PDP-1 computer was state of the art. This restored example is in the Computer History Museum, Mountain View, California. (Alexey Komarov/Wikimedia)

We Have Been Here Before

Joseph Becker's Computer and the AI Question

Kenneth Malphurs

The author is a member of CIA's Senior Analytic Service.

The really unique feature of a digital computer is what is called its "logic," its ability to choose for itself one of a number of alternative procedures according to the outcome of previous computations. This feature is the one mainly responsible for the notion that digital computers are endowed with near-human or even superhuman qualities. The fancy is encouraged in the vocabulary used by the computer people: you "instruct" and "query" the machine in its own "language"; it "accepts," "differentiates," "searches its memory," "analyzes," even "evaluates." A seasoned computer operator will argue on occasion that the machine has a personality of its own, and his emotional involvement with the machine is such that research is being done in

man-machine relationships to arrive at the right mix of human factors for happy and efficient work with a machine as colleague or subordinate. Nevertheless, although it is true that man can evolve ways to make a digital computer perform operations that closely resemble human thought, and although the machine can digest more information than a man and process it faster and more accurately, the parallel with human skills should not be carried past the point of fanciful analogy. The machine does not "think"; it is driven through a predetermined set of operations.

— Joseph Becker, "The Computer," *Studies in Intelligence* 4, No. 4 (Winter 1960)

All statements of fact, opinion, or analysis expressed are those of the authors and do not reflect the official positions or views of the US government. Nothing in the contents should be construed as asserting or implying US government authentication of information or endorsement of the authors' views.

We Have Been Here Before

In winter 1960, *Studies in Intelligence* published “The Computer—Capabilities, Prospects, and Implications,” by Joseph Becker, a pioneer in information technology. Much of the article now reads like a tour through a technical museum, with Becker writing about punch cards, magnetic tape, and machines that occupied dedicated space and required specialist operators, but the article’s real subject was not hardware so much as the effect that a new tool might have on intelligence work.

Read today, Becker’s article does something more than illustrate an earlier stage of automation; it also brings into view a line of argument that now feels familiar.^a Becker was interested in what computers might do for storage, retrieval, dissemination, translation, and analysis, but he was equally interested in the difficulties that would follow their introduction: converting records into usable form, training personnel, reorganizing work, and resisting the temptation to confuse machine assistance with machine judgment.

That resemblance to present discussions of artificial intelligence suggested a simple exercise. I set three texts beside one another: Becker’s original article; an AI-generated short commentary comparing Becker’s discussion of computers with current discussion

of AI; and an AI-generated rewrite of Becker’s article in which “AI” was substituted for “computer” while the surrounding language was changed as little as possible. The point was not to stage a literary trick; it was to see what the comparison would reveal.

The result was clear enough: the analogy does not hold everywhere, and it should not be forced, but it holds often enough to be useful. Intelligence has faced disruptive technical change before, and although the present case is not identical to Becker’s, it is similar in one important respect: *Many of the hardest questions have less to do with the machine than with the institution that means to use it.*

Becker’s Article Still Speaks Plainly

Becker began in the right place. Computers, he wrote, were not magical devices that would spare intelligence officers the trouble of thinking; they were tools, and their value depended on how they were used. That point has lost none of its force, though it is often among the first to disappear in periods of excitement. New technologies attract evangelists and detractors in equal numbers, but Becker did neither. He wrote as if the machine were something to be evaluated, adopted where useful, and bound where necessary.

That cast of mind makes the article more durable than its technical particulars. Becker’s examples now sound dated, but the underlying problem is still recognizable. In one of the article’s most memorable cautions, he writes computers are not “glamorous Aladdin’s lamps to do our bidding while we recline at ease,” but “extensions of human skills.” Elsewhere, in a passage that survives the rewrite with remarkably little strain, he describes intelligence dissemination by matching incoming reports or cables against analyst “profiles” made up of selected words and phrases. His discussion of translation and abstracting is more tentative, but even there the direction is recognizable: the machine is moving into parts of the intelligence process that had previously been left to human labor.

The article becomes more interesting, however, when Becker leaves the machine itself and turns to the organization around it, because there the prose has scarcely aged. He anticipated problems of conversion, staffing, training, parallel operation, and work rearrangement; he argued that the organization had to be treated as a “total, integrated, functional system”; and he put the change on a planning horizon of five to seven years, with costs that might run into the tens of millions. He also observed that people would resist the prospect of the machine “taking over,” and

a. Joseph Becker, “The Computer—Capabilities, Prospects, and Implications,” *Studies in Intelligence* 4, no. 4 (1960). Approved for release on September 18, 1995.

How I Used AI in Writing This Article

I began with a very specific experiment rather than a vague request: I gave a frontier AI model Joseph Becker's original article and asked the model to do three things: comment on the similarities and differences between Becker's discussion and today's discussion of AI; rewrite the article by substituting "AI" for "computer" while changing as little else as possible; and then compare the two versions. That initial exercise produced the core insights for this piece, and from there I used AI less as a ghostwriter than as a working partner in revision. I asked it to draft an article built around the experiment, then worked back and forth with it on structure, title, tone, evidence, sentence rhythm, and ending. That approach let me focus more of my energy on what I wanted to say and how I wanted it to land. In practice, I was managing a process: testing arguments, pushing for clearer distinctions, rejecting language that felt wrong, and trying many more ways of phrasing an idea than I probably would have if I had been drafting every line alone. I still gave the article a final close review and made manual edits myself, so AI did not remove the work of writing; it made the work more iterative, faster, and more focused on the overall project—its message, purpose, and impact.

insisted that intelligence should be planning an "enlightened training program chronologically coordinated with the rest of the change-over effort." Those are not timeless truths about every profession, but they are very close to the questions our profession is again asking.

For that reason, Becker is worth reading now, not because he foretold AI, but because this article captures a pattern that is visible again in the present case. The first arguments are often inflated, while the later problems are usually mundane. What can the machine actually do? What must be reformatted or rewritten so that it can do it? Who will be trained? Who will check the output? What work changes shape as a result? Those were Becker's questions, and they are close enough

to our own to make the comparison more than a curiosity.

Rewrite was Crude and Useful

The AI rewrite of Becker's article was intentionally crude. The AI was not asked to produce a modern essay in its own voice; it was asked to make the smallest possible set of changes required to turn a 1960 article about computers into a plausible article about AI. That simple substitution proved more revealing than one might have expected.

Some passages survived with very little damage, and the pattern of survival is itself instructive. Becker's warning that the machine is no "Aladdin's lamp," his

insistence that it extends rather than replaces human capability, and his discussion of apprehension, retraining, and work redistribution all read plausibly in the rewritten version. So does the dissemination section, where incoming reporting is compared against analyst profiles; in the rewrite, that sounds very much like a contemporary description of alerting or recommendation built around the routing of new reports to the people most likely to need them. The same is true when Becker writes about the institutional effects of the computer—organizational disruption, parallel processing, staffing, and training—in terms that now feel only lightly dated.

Other passages did not survive so cleanly, and those failures were just as useful. The hardware passages, once "computer" became "AI," turned awkward or absurd, which was expected. More revealing was Becker's statement that the machine does not think but carries out predetermined operations. In the original article, the sentence is descriptive. In the rewrite, it becomes argumentative, because the contemporary reader hears it against a live debate about how to characterize systems that produce language, summarize documents, write code, and conduct a plausible dialogue. The sentence remained the same; the argument around it changed. The translation section is another revealing case. Becker wrote that smooth-prose machine translation

We Have Been Here Before

was “not yet a reality” and “may not be for many years to come.”

In the rewritten version, that line no longer reads as caution alone; it also reads as a forecast whose timetable proved too slow. That, more than the successful substitutions, is what made the exercise useful. It showed that continuity of institutional reaction does not imply identity of technology. AI is not just the old machine under a new name. What the experiment does show is narrower and, in some ways, more interesting: Becker’s institutional language often remains legible even when the technology underneath it has changed substantially. Exaggeration, anxiety, managerial confusion, uneven adoption, and bad implementation are not proofs of sameness, but they are recurring features of how organizations meet tools they believe may alter the conduct of their work.

How Today is Different

The differences are substantial and should be stated plainly.

The first is accessibility. Becker wrote of specialists, operators, programmers, and scarce equipment, whereas AI has entered offices through ordinary language. An analyst can try a system directly, without a programmer, an operator, or a visit to a machine room, and that changes the pattern of adoption while also making control

harder. Institutions are not deciding only whether to introduce a tool; in many cases they are deciding how to govern one that has already arrived.

The second is speed, though even here the matter is less simple than current rhetoric suggests. New AI systems can spread very quickly across a workforce, and capabilities that once would have moved slowly from laboratory to office can now appear in common tools in short order. But broad availability should not be confused with broad integration. Computing is costly, security restrictions remain real, data are often poorly organized, and workflows are rarely designed from the start with such tools in view. The visible spread of AI can therefore be rapid even while institutional absorption remains slow.

The third is reach. Becker wrote about translation and abstracting as difficult prospects. Current systems can summarize, draft, classify, translate, review long texts, compare documents, and search across language in ways that put them in direct contact with forms of work once treated as beyond machine competence. That does not mean the machine has become an analyst, but it does mean that some tasks once regarded as securely human are now open to technical assistance in a way Becker could only anticipate in outline.

The fourth is that the problem is now defensive as well as productive. Becker wrote about what computers could do for intelligence work, but AI must also be considered as an instrument available to adversaries, fraudsters, propagandists, and ordinary opportunists. The same systems that support search, review, translation, and drafting can also support deception, impersonation, phishing, malware production, and manipulation of public information at scale. Any serious program of adoption must account for both sides of that ledger.

These differences matter, and they keep the comparison honest. They do not cancel the historical point.

What the Exercise Shows About AI as a Writing Aid

My experiment had a second purpose: it was meant to test what an AI system could actually contribute to a piece of writing of this kind. The answer, at least in this case, is that it can do quite a lot, and do it quickly. It could read Becker’s article, identify the themes most likely to survive into the present, compare them with current arguments about AI, produce the substitution exercise, and draft a serviceable essay in an appropriately formal tone. Those are substantial advantages at the stage of comparison, synthesis, and early drafting, even if they fall well short of independent authorship.

Its weaknesses were equally plain. Left alone, it tended to make everything sound slightly more complete than it really was. It generalized where specificity was needed, liked smooth transitions, and produced balanced sentences even when the thought called for a harder edge. It was more comfortable sounding authoritative than admitting uncertainty. In that respect, it behaved like a clever drafter who had read widely, written quickly, and not yet learned the value of restraint.

That combination makes it useful, but only under supervision. For writing, the best present use is not authorship in the full sense; it is assistance: summarizing a source, comparing texts, proposing language, testing a formulation, or producing a draft that a human writer will cut apart and rebuild. The technology is strongest where the work is preparatory or comparative, and weakest where it must decide what matters, what should be omitted, or where an argument needs to become more precise rather than merely more fluent. At least for now.

That distinction is worth keeping in mind because prose is one of the domains in which AI can most easily create a false sense of mastery. Orderly paragraphs, plausible transitions, and an air of confidence are easy to mistake for judgment. They are not the same thing.

The Institutional Lesson

Becker ended by calling for planning, experimentation, and system-level analysis, which still seems the right response. For a reader trying to think clearly about AI in intelligence, that is Becker's value now: he helps separate what is genuinely new in the technology from what is familiar in the institution's response. The sensible position on AI is neither missionary enthusiasm nor easy dismissal; it is to examine the work itself: which tasks benefit from these systems, which tasks become more error-prone when they are introduced, and what must be changed in data, holdings, access rules, workflow, supervision, and training if the tool is to be used responsibly.

Three practical conclusions follow from this. First, training cannot be treated as an afterthought. Access is not competence. Officers who will use these systems seriously need some understanding of model behavior, provenance, retrieval, verification, and common modes of failure. They also need a clearer sense of when the tool should not be used.

Second, infrastructure matters more than demonstration. A good model attached to disordered data, inaccessible holdings, or a poor workflow will disappoint. So will a successful pilot that collapses when it meets security rules, records requirements, and the routine architecture of office work. The

institution's problem is rarely the model by itself. It is the fit between the model and the process into which it is introduced.

Third, changes—including dramatic shifts—in work should be expected and described soberly. Some tasks will contract, others will grow, and some roles will move toward review, correction, or orchestration rather than first production. New specialties will appear, and older boundaries will weaken. None of that is unprecedented. It is what happens when a profession acquires a tool that changes the relation between labor, information, and time.

Becker's 1960 article remains useful for one simple reason: it helps clarify what is and is not new in the present discussion. The obsolete machinery is obvious enough. Less obvious, until one reads the piece closely or runs the substitution exercise, is how much of the argument turns on matters that are still with us—training, workflow, conversion of information into usable form, organizational disruption, and the temptation either to overstate the machine or to recoil from it. In that sense, the article rewards institutions that think clearly about process and punishes those that mistake novelty for strategy.

We have, in that sense, been here before. Not before the same machine, and not under the same conditions, but before the same

We Have Been Here Before

professional problem: how to absorb a powerful new instrument without allowing fascination with capability to outrun standards, training, and judgment. That was true of the computer in Becker's day. It is true of AI now. The decisive question is not whether the machine is impressive. It is whether the institution using it is prepared to govern the change it brings.

Use of Artificial Intelligence. The commentary and rewrite discussed in this article were generated with AI as part of the experiment described herein. AI was also used as a drafting aid during preparation of the manuscript. The argument, selection of evidence, and final text are the author's own. ■