

Can AI Make Your Job More Interesting?

The symbiosis of artificial and human intelligence can help create more dynamic and rewarding places for both people and robots to work.

Sixty years ago, in an era of computer mainframes and slide rules, J.C.R. Licklider of the Defense Advanced Research Projects Agency (DARPA) outlined a bold vision of a mutually beneficial human-computer symbiosis, a partnership in which both humans and machines positively benefit. The vision remains largely unfulfilled. Today, humans and intelligent machines work alongside each other—say, in a robotics-enabled Amazon fulfillment center—but it’s hard to claim that this shoulder-to-shoulder work represents a symbiosis. Instead, it is humans who do the routine work of filling boxes, while algorithms collect the insights on product popularity.

Now, the COVID-19 pandemic is accelerating a transformation in the way companies operate. Central elements of this transformation include flexible work arrangements, more automation, and a push toward a “contactless” economy built on a backbone of ubiquitous data collection, artificial intelligence (AI), and human-robot systems. Amid these changes, will the future workplace largely be one of deskilled drudgery like an Amazon warehouse? Or can machines and humans work together, as Licklider imagined, in ways that bring out the best in each?

We are two technologists—researchers, investors, and explorers—who have spent our careers trying to invent or foster the automation of the future. As part of our work at DARPA, charged with framing and realizing long-term disruption, we are wrestling with the future of human-AI teaming and collective activity. But the more we dive into the complexity of mixing humans and AI, the less the common paradigms—teammates, tools,

replacements—seem to hold up. Instead, these advances make us rethink the very way we conceptualize intelligence itself.

We assert that the correct way to think about AI and the workplace is not as a challenge of management versus labor, or machines versus humans, but as a problem of mediating the interactions among system components: humans, AI, firms. The activity of these dynamic groups will be mediated by intelligent coordination mechanisms that match information transformation and human judgment to produce remarkable outcomes. This evolution will change the character of work and the firm as profoundly as industrial automation did in the nineteenth and twentieth centuries, with correspondingly profound implications for individuals, businesses, and policy-makers.

But is it good for the machines too?

Human-machine symbiosis is not just a process of seeking out and dividing labor between those things that people are better at and those that machines are better at. That framing has been incredibly useful because it results in ever-improving superior human productivity that’s directly measurable, and it drives cycles of continually improving machines, algorithms, and human interfaces. Both of us have developed technologies that fit this template: lab automation for drug discovery, software for optimizing engineering design, and improved coordination and planning tools for military operators.

But just adding automation into a factory or workplace is not symbiosis, because the machines gain little to nothing from the humans involved. And humans gain only to the extent that increasing productivity frees them up to do equally or more interesting work.

The idea of artificial intelligence or AI-enabled robotics as a teammate is more symbiotic. As everyday experience with voice assistants makes clear, AI can learn from human interactions and improve its skills, while human productivity increases through interaction with an increasingly capable AI. This individualistic view is compelling and intuitive. Its archetype is the chess “centaur,” a hybrid mash-up of a human player and an AI chess guide, which could triumph not only over grandmasters but also chess AI alone. Although AI has now achieved a level of superhuman play that no longer benefits from human judgment, at the time of the freestyle tournaments in the early 2010s, humans provided “meta-strategy” and judgment: managing time effectively, deciding when to shift from formulaic moves to those calculated by the chess AI to exploit strong positions, and crafting tournament strategies to exploit the foibles of the opponent’s chess AI that only a human player could recognize.

One of our DARPA experiments, which we called Alias, attempted to apply this teaming model to aviation, placing a centaur of a human mission commander and automated assistant into a cockpit that previously required two people. The critical insight was that the automation couldn’t replace the cognitive breadth of a human copilot dealing with uncertainty and ambiguity, but could perform better at many tasks where humans struggled, such as landing with a failed engine where controlling the flight path to preserve every watt of energy counts, something once believed to be uniquely suited to humans.

Alias was a technical success. It allowed a novice pilot to pick up an iPad and fly a complex million-dollar helicopter as easily as if it were a toy drone. And though it eliminated the need for standard piloting skills, it also revealed the crucial importance of different, undervalued skills. For example, executing a real transport flight to and from an oil rig required a pilot to understand nuanced customer objectives, translate these into mission parameters, and perform long-range planning. These airline captain aspects of piloting were so varied and abstract in real life that programming them proved futile, but they were easy for a human. Alias showed that for a truly synergistic teaming construct to work, one must succeed at deciding which roles to assign to humans and which to machines, and explicitly coordinate the combined activity.

This experience led to the disappointing realization that the teaming approach is likely to work only in a field such as aviation, where there are small numbers of humans and machines with clearly defined tasks and objectives. Even though piloting is considered highly skilled because of the intensive human training required, the task allocation model underpinning the concept of human-machine teaming is unlikely to carry over to endeavors with more ambiguous notions of how tasks relate to “good”

outcomes, or to problems in which learning what the tasks or outcomes should even be is central to making progress. Teaming, with its reductive assignment of particular tasks to specific competency, fails to capture this dynamism and resulting symbiotic learning by both human and machine.

An example of an alternative to teaming is a “superintelligence” model in which AI lifts the ability of individuals to collectively address hard problems that elude experts or even teams of experts. An example is work one of us sponsored using FoldIt and Mozak, citizen science “serious game” platforms that have facilitated the discovery of new therapeutics and enabled breakthroughs in neuroscience. The overall players collective is capable of better performance than AI or the players alone. The process is not symbiotic—the humans are not becoming experts in protein folding or neuroscience. Rather, the humans are simply better than the machine at certain activities such as searching for optimal protein geometries or seeing neural connections in a fuzzy image. The centralized “algorithmic manager” can effectively orchestrate and coordinate effort, continually improving outcomes through active manipulation of player attention and effort. Those interventions are enabled by pervasive surveillance of how every player is behaving from the moment they log in—continually measuring how every click, move, and chat comment relates to the quality of the well-defined desired outcome.

Most modern knowledge work can’t be neatly broken down into roles and tasks with unambiguous measures of performance. Further, there is an incredible variety of skills among humans; even within a single job description one might value Bob for being “creative” and Sally for being “detail-oriented.” Deciding who should be assigned a task requires not only a rigid definition of the task but also a solid understanding of the particular human being. Understanding of human cognition is still too rudimentary to operationalize the “humans are better at” construct beyond abstractions such as creativity, empathy, or judgment. All of these concepts hint at the diversity of human intelligence: people’s ability to craft and apply abstract models of the world around them, in the appropriate context, in order to achieve remarkable collaborative outcomes. So we suspect that it may never be possible to design a single human-machine interaction framework that equally suits all humans or that replaces or exceeds all types of human intelligence.

The future of meaningful work, in which humans are not reduced to assets that complete prescribed workflows with well defined productivity measures, will hinge on how information, AI, and human judgment are combined. Doing so in a truly symbiotic way requires a systems-level perspective, focused neither on the individual nor the crowd, but on the firm, markets, and economies. True symbiosis will accommodate and depend on human interaction, values, and social choices, with technology playing a largely hidden but profound role.

The singularity is not coming

It makes sense to combine people's unique talents and aggregate them. Firms exist because collectives can produce more than the same number of people acting individually. This insight was framed in economic terms by the British economist Ronald Coase, who recognized that transaction costs and externalities drive a firm's decision of whether to buy a product or make it in-house. Firms enable both faster and better decisions under uncertainty—making predictions about future business—by bringing the right expertise in-house, such as planners, financial analysts, and project managers.

The internet, smartphones, and AI have become the backbone of the economy in large part because they have nearly eliminated transaction costs and transmission time for information exchange. This, in turn, makes prediction easy and cheap as we see in everyday use of Google, Facebook, or Netflix, which (for better or worse) aggregate vast amounts of cheaply acquired information and process it with AI to predict which advertisements will appeal to you, which conspiracy theory is most intriguing, and which TV show is most conducive to binge watching.

As a result, it's unsurprising that long-standing forms of interactions between individuals, firms, and markets have evolved radically over the past 15 years. These trends were predicted in the 1990s by the management researcher Thomas Malone, who foresaw an economy in which electronic marketplaces would replace and complement the firm, even if he didn't specifically foresee eBay and Shopify. The relationship between individuals and the firm is evolving in precisely this way: consider the employment relationship of Uber drivers to Uber itself, and their interactions with an app-based AI manager that uses behavioral economics principles to maximize productivity of thousands of drivers.

Though AI systems that underpin gig economy firms such as Uber are unequivocally not symbiotic—they are maximizing efficiency and manipulating extrinsic rewards for drivers, who in turn have tried to game the app for better pay—they show that the future isn't about better machines, or smarter humans, or even amazing centaur teams combining both. As reasoning shifts from people to artificially intelligent systems that can, in the words of the computer scientists David Parkes and Michael Wellman, “learn our preferences, overcome our decision biases, and make complex cost-benefit trade-offs,” our research shows that achieving true symbiosis requires change in the basic economic institutions (e.g., firms and governments) that mediate everyday transactions.

In other words, the commonly offered vision of AI moving inexorably toward a singularity where it will finally overtake human capacity is simplistic and limiting. On the contrary, technology is causing the diversity of

different forms of intelligence to explode. Now the real economic and technological opportunity lies in matching these diverse forms of intelligence into collaborative groups to tackle the problems and opportunities that society faces. This opportunity will require intelligent coordination systems—mediators, not managers—that comprise what we call the “intelligence economy,” knitting together markets, AI-enabled infrastructure, firms, and institutions to match information transformation and human judgment to advance human aspirations.

A glimpse of what this looks like at small scales can be found in a robotics company that one of us cofounded. In tackling the problem of warehouse robotics, instead of just focusing on a smarter robot, we built a tool, called Pivotal, to mediate the work between different robots and different humans. From one perspective, Pivotal took the idea of a distributed gig-work marketplace from Uber and Lyft and applied it in industrial settings. But rather than focusing on extracting maximum productivity via a twenty-first century version of the dehumanized assembly line advanced by Frederick Taylor in the late nineteenth century, the system was designed to give humans a choice in what they do, to help them see how their work is contributing to a larger overall goal, and to give them a voice in making it better. As with a classic Taylorist approach, all the work required to fulfill a set of customer orders from warehouse stock or sort packages from incoming trucks to outgoing trucks could be parceled out as a set of tasks or missions. But instead of being assigned work, different workers with different skills and various robots with different abilities would bid on what they felt that they could do or wanted to do.

The results were remarkable. An Amazon warehouse is a model of the principles of industrial revolution efficiency—every day, the human workers do the same basic shelf-picking task all day long. With Pivotal, though, the exact way to complete a task like filling an order could evolve from day to day. Workers often ended up finding variety in their jobs, handling unusual tasks such as catching a stray bird, clearing obstacles, and applying their abstract problem-solving skills to understanding why a box had two conflicting labels and deciding which one was right. Robots tended to settle into repeatable patterns with gradual performance improvement until their engineers analyzed the data and pushed out new features. This entire operation was implemented using a data-driven framework, so that machine learning could benefit from human insights and human training of AI algorithms, and humans could analyze the unexpected behaviors and emerging patterns and innovate—both on the shop floor and in the engineering office.

In many ways, Pivotal enabled an updated version of the famous Toyota assembly line, where any worker can pull the cord and stop the process, so they are engaged in the outcome, not just the isolated task. Instead of viewing the

worker as low-skilled, this process is open to the fact that the worker can have valuable insights and observations that were not apparent to the engineers who designed the original process. The symbiotic human-AI system was designed to emphasize autonomy, responsibility, competence, and diversity of intelligence.

It's the intelligence economy, stupid

At the heart of the intelligence economy concept is the idea of AI-augmented markets as a replacement for the firm of today, and in the Pivotal or Uber examples, the market takes the form of an auction. But as we consider a world of AIs and humans that continually evolve and learn from each other, it's not clear that an auction is the best form, or why we should expect an intelligent market to even maintain the same form over time. Recent academic work and one of our DARPA efforts (called Agile Teams) are beginning to explore using AI techniques to design systems to mediate between knowledge workers and AI systems so that there are beneficial incentives and outcomes for all of the players,

In the PPE vignette, that might entail AI adaptively providing preferential access to AI to help with design of new drones or space at certain warehouses in service of the more socially beneficial, but less profitable, cargo within the mediation platform. It also makes clear that expecting an intelligence economy to simply emerge in a way that ensures mutually beneficial outcomes and correctly accounts for shocks and externalities is naïve. Ensuring that it does will require the concerted and collaborative effort of policy-makers, executives, funding agencies, and researchers spanning economics, organizational theory, and computer science.

The pandemic has revealed how brittle existing institutions are to disruptions, and how AI built to maximize efficiency via nineteenth century approaches to human labor exacerbated that brittleness. Successful executives and management researchers will find new ways to rapidly identify opportunities, assemble and cultivate collectives of human and AI talent, and seek competitive advantage in a world that is fundamentally more dynamic through AI-enabled mediation. The future should not be one of digital

The evolution of AI will change the character of work and the firm as profoundly as industrial automation did in the nineteenth and twentieth centuries, with correspondingly profound implications for individuals, businesses, and policy-makers.

while also ensuring the resulting group is resilient even in the face of unexpected events such as the absence of a worker or a sudden shift in objectives.

In Agile Teams we are exploring what a future logistics “firm” that delivers directly to customers using drones might look like. In this experiment, human and AI business strategists, operations planners, and autonomous drone designers work together through an AI-augmented platform that mediates their interactions. In one vignette inspired by recent events, an urgent and unexpected need for personal protective equipment (PPE) delivery catalyzes the formation of a new ad hoc team that complements the existing one built to deliver regular cargo, but requiring a very different drone fleet and warehousing approach. The vignette exposes some important questions. How to best elicit ways to frame the business strategy or drone design problems? How does the mediation platform prioritize the PPE need, which is less profitable, but of greater social benefit, than the existing business model? How are other teams assisting with the PPE challenge, and in doing so are they actually shaping the competitive landscape to their advantage at the same time?

The research in Agile Teams exposes deep questions of how exactly to encode what we value in the systems we build: incentives and outcomes are modulated and amplified by AI.

Taylorism, but instead an AI-infused, “open” organization capable of both adaptation and scale.

A mix of data-driven AI and diverse human intelligence participating in an economy, combined with intelligent AI-mediated markets, will mark a profound shift in how collective goals are achieved. Existing research on the “future of work” is fixated on automation and team-based paradigms. Our vision of symbiosis highlights largely unexplored problems at the interfaces of economics, social science, and AI research. Traditional approaches to human-computer interaction, economic modeling, and understanding organizational performance are inadequate for the intelligence economy.

The future machinery of democracy?

Policy-makers should anticipate this future by providing incentives for the emerging intelligence economy to internalize desired social outcomes including equality of opportunity, fairness of outcomes, and assurance of competition. This focus on outcomes is often difficult in policy, where law must express the means. Many of the coming policy debates will be considerably more complex than today's regulatory debates. AI-enabled platform firms such as Google, where strong network effects lead to

dominant market share, and Airbnb, with its implications for the character of neighborhoods, foreshadow the challenges ahead, but our belief is that a focus on symbiosis rather than competition between humans and machines can be a powerful framing for meeting these challenges.

The right tool for policy-makers to address the challenges of the intelligence economy is the branch of economics called mechanism design. Mechanisms are what institutions use to allocate resources when the information needed to make those allocation decisions is dispersed and privately held. Kidney exchange is one of the best-known examples of the application of mechanism design; in this case, kidney donors and patients in need of a kidney transplant need to be matched. Clearinghouses exist to make these matches, but transplant centers have misaligned incentives. Though it is simpler and more profitable for each center to maximize the number of transplants internally and then exchange what can't be matched through the clearinghouse, that locally optimal outcome is actually worse overall, leading to fewer matches for the greater community. Through a properly designed mechanism—the rules and processes by which the transplant centers must operate with the clearinghouse—each transplant center gets more matches than it would have alone while also preventing gaming the system and providing higher benefit to the community by fairly allocating scarce resources.

The mechanism for optimizing the social benefit of kidney exchanges resulted from the work of skilled designers. We are interested in whether such a design process can be handled by AI, and one of us is now exploring how to use AI to find the right mediation mechanism to achieve a desired outcome in a specific context. Applied to policy, this leads to some of the most provocative implications of our vision of symbiosis.

We posit the best means to ensure a beneficent intelligence economy is itself a form of our vision of symbiosis, combining AI and humans to craft entirely new forms of policy and institutions. The legislative process is often reactive and prescriptive: a pandemic triggers an unemployment crisis, and Congress debates the specific processes and actions that should be written into law in response. An alternative would be to write desired outcomes into law (an acceptable unemployment threshold) accompanied by a supporting mechanism (such as flowing federal dollars to state unemployment agencies and tax-incentivization of business hiring) that could be automatically regulated according to an algorithm until an acceptable level of unemployment is again reached.

Though this legislative approach faces the challenge of needing to be implemented before a crisis, it has the advantage of moving debate to the more politically palatable terms of acceptable levels of unemployment, with AI dynamically mediating between competing forms of support. This outcome-based concept is less about ceding democracy

to machines, and more about recognizing the need to complement human-driven institutions with algorithmic help to navigate increasing complexity. The concept of AI-assisted policy has been explored for taxation schemes that seek to balance the seemingly competing interests of equality and productivity, and there is a growing movement exploring auctions and other algorithmic mechanisms for such fraught issues as income inequality, economic stagnation, and political strife.

By using AI to improve coordination among humans, teams, firms, and AIs, a mediation-driven approach can lead to more resilient and fairer outcomes for all parties. By catalyzing large-scale change in how people work, today's challenges in adapting to the pandemic will provide the opportunity to take the first steps toward a more symbiotic intelligence economy. We are bullish on a future that thrives on diversity and lets everyone find a productive and fulfilling application for their minds. In our future, AI does not overtake human intelligence one sad day, but instead is the only technology capable of helping unleash the true diversity of humanity's collective intellect, helping society cooperate more effectively on a better future.

John Paschkewitz is a program manager at the Defense Advanced Research Projects Agency, with a portfolio focused on technologies and systems to accelerate adaptability, increase resilience and improve competitive advantage. The views expressed are his own. Dan Patt is a fellow at the Hudson Institute and supports both strategy at the AI firm STR and automation investment at THL Partners. He is a former DARPA program manager and deputy director, as well as a former chief executive officer at a robotics company.

RECOMMENDED READING

- Chen Hajaj, John P. Dickerson, Avinatan Hassidim, Tuomas Sandholm, and David Sarne, "Strategy-Proof and Efficient Kidney Exchange Using a Credit Mechanism," Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (2015).
- Zoë Hitzig, Lily Hu, and Salomé Viljoen, "The Technological Politics of Mechanism Design," *The University of Chicago Law Review* 87, no. 1 (2019): 95–108.
- David C. Parkes and Michael P. Wellman, "Economic Reasoning and Artificial Intelligence," *Science* 349, no. 6245 (2015): 267–272.
- Roberto Mangabeira Unger, "The Knowledge Economy: A Critique of the Dominant View," *American Affairs* 4, no. 3 (2020): 48–76.
- Stephan Zheng, Alexander Trott, Sunil Srinivasa, Nikhil Naik, Melvin Gruesbeck, David C. Parkes, and Richard Socher, "The AI Economist: Improving Equality and Productivity with AI-Driven Tax Policies," arXiv:2004.13332 (April 28, 2020).