Research-Technology Management

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/urtm20

The Rise of Human-Machine Collaboration

An Interview with Tom Davenport *Jim Euchner talks with Tom Davenport about applications of machine learning and the role of people in the future of automation.*

Tom Davenport & Jim Euchner

To cite this article: Tom Davenport & Jim Euchner (2023) The Rise of Human-Machine Collaboration, Research-Technology Management, 66:1, 11-15, DOI: 10.1080/08956308.2023.2142435

To link to this article: https://doi.org/10.1080/08956308.2023.2142435

Published online: 20 Dec 2022.

Submit your article to this journal

Article views: 44

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=urtm20

The Rise of Human-Machine Collaboration

An Interview with Tom Davenport

Jim Euchner talks with Tom Davenport about applications of machine learning and the role of people in the future of automation.

Tom Davenport and Jim Euchner

Jim Euchner [JE]: Artificial intelligence (AI) has been around for a long time. I am old enough to remember the last wave of enthusiasm around expert systems. What's new?

Tom Davenport [TD]: The tremendous rise in the application of machine learning (ML), I think, is the biggest driver of change. While ML isn't really new overall, we can do a lot more of it because we have a lot more data that we can use to train the ML models. Deep learning algorithms are also relatively new, and they do some things really well that couldn't be done before, things like image recognition and voice and text analysis. In some ways, we have just achieved critical mass. For whatever reason, expert systems never really achieved critical mass, and the field went into another AI winter. While a recession, which we seem to be headed for, could diminish the amount of startup activity related to AI, I don't think it's going to lead to any sort of AI winter. I think we're at the point where there's no turning back.

JE: The perception is that expert systems didn't deliver enough value in their day, but my experience was that, if you got your application right and you focused on solving the whole problem—the social problem, the integration problem—you could deliver significant value. Do you think people blamed the technology for the lack of a holistic approach?

Tom Davenport is a distinguished professor of information technology and management at Babson College, visiting professor at Oxford Business School, Fellow of the MIT Digital Economy Initiative, and senior advisor to Deloitte's AI practice. He has published 23 books and over 300 articles for *Harvard Business Review* and many other publications. He has been named one of the world's Top 25 Consultants, one of the top 3 business/technology analysts, one of the 100 most influential people in the IT industry, and one of the top 50 business school professors in the world. tdavenport@babson.edu

Jim Euchner is editor-in-chief of *Research-Technology Management* and Honorary Professor at Aston University (UK). He previously held senior management positions in innovation leadership at Goodyear Tire and Rubber Company, Pitney Bowes, and Bell Atlantic. He holds BS and MS degrees in mechanical and aerospace engineering from Cornell and Princeton Universities, respectively, and an MBA from Southern Methodist University. He is the author of *Lean Startup in Large Organizations*. euchner@iriweb.org

DOI: 10.1080/08956308.2023.2142435

Copyright © 2022, Innovation Research Interchange.
Published by Taylor & Francis. All rights reserved.

TD: My impression is that they worked quite well up to a certain level of complexity of the problem, and then they became a little too brittle. Large expert systems could have thousands of rules, and they tended to interact with each other in strange ways. And if you needed to update them, it could be onerous. I worked with a company in Boston that had developed an expert system for financial planning, and it worked quite well, but if you changed the tax code, a huge amount of work was required to update the rules.

JE: A number of the examples in your newest book, *Working with AI*, incorporated image and language recognition, some of the most impressive capabilities of machine learning. Is this where the power is?

TD: When we were writing the book, we didn't really look for any particular type of AI technology, but those you mentioned have moved into the world of practice pretty quickly. It was surprising to me how many we found, and also, how many had mixed AI types. It became clear to us that the AI capabilities have to run on a platform, including a data platform, a workflow platform, and so on. By the way, rule-based (expert) systems are quite common in workflow systems now, so the different modes of AI are all coexisting, it seems.

JE: Would you say the biggest shift has been in the scope of problems that AI can address? What has stayed the same is that the challenges are still in integration with other systems, work design associated with the new systems, and change management.

TD: When we started doing research on this book, we didn't realize what a big deal it can be to put these systems into production. A lot of companies are just experimenting with AI, and impressive experiments are typically not that hard. But if you're integrating AI with an existing process and an existing set of systems and requiring your people to use it, it's a lot of work—and it's slow work. I mean, technology itself changes pretty fast, but the deployment of new technology is a slow process. In our book we say that it "takes a village" of people to make AI work in the real world—not only the developers and frontline workers who will be affected, but their supervisors and senior management and

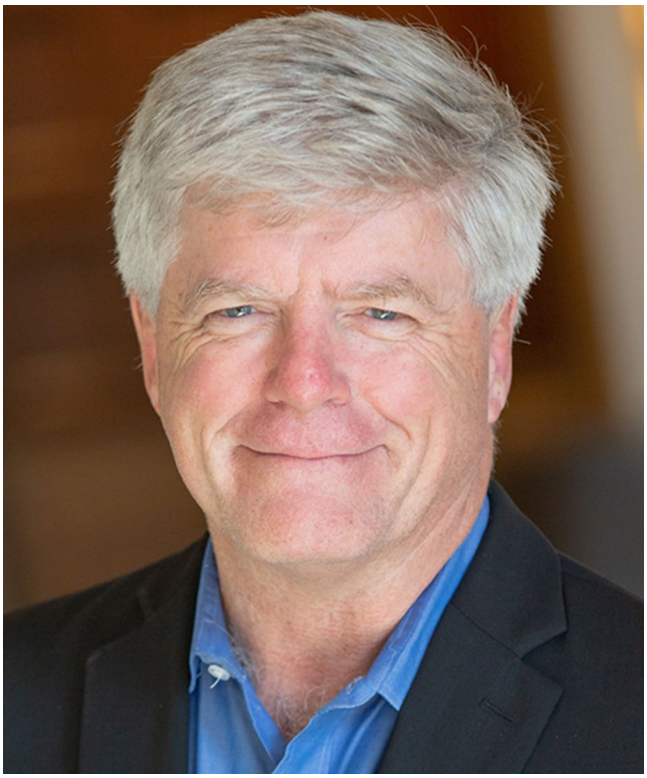

Tom Davenport is a distinguished professor at Babson College.

vendors and the IT people who do the integration, and so on. It's a challenge.

JE: One of the major observations in your book is AI is not really about automation, it's more about augmentation. You make the point that it will be humans plus machines that will be delivering productive work. Can you say more about how you think about that?

TD: The book is primarily about augmentation. I had written a previous book on augmentation called *Only Humans Need Apply*, which I coauthored with Julia Kirby. For this book, we wanted to find many more examples of augmentation and learn from those examples what works.

We didn't have any problems finding good examples. I think they are one-offs, for the most part, however. We didn't find, in most companies, an organizational structure oriented to creating them. Maybe it's because we are still in the early days of augmented intelligence, but these examples usually come from somebody in an organization figuring out that something was possible and pursuing it. My guess is that there will be a lot more of this sort of thing before long.

Ideally, we will find those circumstances where humans and machines can be collaborative partners, rather than one being in charge all the time.

JE: How do you perceive that work changing as a result of augmented intelligence? I have noticed many instances where individual tasks get automated, and the humans are left scrambling to pick up the pieces when the automation fails. That's a tough, unsatisfying job, and I'm not sure it creates a good customer experience either. If you automate tasks, it's different than if you redesign work to take advantage of technology. You know this well from your earlier work in re-engineering. What are your thoughts about work design and technology?

TD: There is a whole range of possibilities for the role that the AI can play in organizations. It can be the prime driver of the work, and the humans can pick up the ball when AI drops it; or the human can be the primary driver of the work, and the system can check it or prioritize it. Ideally, we will find those circumstances where humans and machines can be collaborative partners, rather than one being in charge all the time. But I think we're still finding our way there.

I'm looking now at what some people are calling generative AI systems, like OpenAI's GPT-3. It's pretty clear that in order to make generative AI work, you need a human at the front end to create good prompts and then a human at the back end to look at the results and decide whether they really got what was wanted. Some of the results are always junk, but some of them are actually pretty clever.

Good generative AI takes a real partnership. You may have heard about that artist who entered one of his AI-generated images into an art show. He ended up winning, and some of the other artists complained, but his argument was, "Hey, I did a lot of work with this. It took me two weeks of work to get the right prompts and to then to Photoshop some of the elements to make the final product." It was, I think, a real collaboration.

JE: Which is more effective from a business perspective? Is it the human in the loop or the workflow in control that works best? One of the reasons that the automation of tasks is so prevalent is that the ROI seems pretty obvious: we cut a minute out of this task, a half a minute out of that task, we aggregate it up and multiply by N. How do people talk about the business justification of the two types of automation?

TD: I'm not sure we've done systematic comparisons. In highly structured, repetitive work, it's probably better to do automation. In manufacturing, even robots that are not very smart can do quite well. For less structured work, more creative, and knowledge-oriented work, I think we'll be better off keeping a human in the loop. I think that the combination leads to capability, innovation, and intelligence levels higher than either the machine or the human alone could provide.

The world evolves as the technology gets better. I used to use freestyle chess as an example of the power of collaboration because the best chess players in the world were neither

humans alone or machines alone, but humans working with machines. Now it appears that machines have taken over that space again.

JE: I did not know that. Chess has been a cat and mouse game. For a long time, of course, humans were better than machines. Then computers got close. The first big, powerful chess programs were knowledge-based system, built on the expertise of grand masters. Then brute force learning systems won out, and Deep Blue beat Garry Kasparov and became world champion. Kasparov then showed that Centaurs, as they were called, could beat both the best humans and the best machines. Maybe now, deep learning and generative algorithms are tipping the scales again. Interesting.

TD: Seems like it. Maybe what will be the best chess players in the future will be humans with really smart chips implanted in their brains. It's part of a general pattern. As technology gets better, it takes on additional tasks and, in time, becomes better than what humans can do. But, for the most part, however, I don't think that we're there yet.

JE: What's the best example that you came across in your many studies of the human in the loop design?

TD: I would say one of the best is the Morgan Stanley "Next Best Action" system for financial advice. The human advisor is still in control. The system recommends investment ideas; the advisor chooses which ones to send to the client. At Morgan Stanley, advisors have the option to choose whether even to use the system, which is an interesting idea. It turns out that people who use it tend to be more successful with their clients, however. The typical advisor at Morgan Stanley has 200 clients, so they need help in giving them personalized investment ideas. The interesting lesson from that case study is that the AI generates good ideas, but success with the client comes from a sense that the advisor is looking out for them and thinking about the investment world on their behalf. That is what really seems to make the difference.

JE: That's an example of the human providing the relationship, the broader context, making trade-offs that are not easily quantified.

TD: Stitch Fix is another interesting example. Stitch Fix predicts the clothing that people will want and sends it to them on a subscription basis. The key to success is being good enough to minimize returns. The human stylist can override the AI in most situations, but they're discouraged from doing so because the AI, in general, has a better track record. But if the customer sends in a note with specific information, the AI system is not going to take that into account.

JE: Is the AI helping with design or helping primarily with the targeting of customers and product selection?

TD: It's primarily selection and matching with the customer's historical preferences. To gather data, they have a little application like Tinder where you see a lot of fashion items and you swipe left or swipe right. They get a lot of insight from that. They also use data on what you keep and what you send back.

JE: This is an instance for which the AI is reinventing the industry. It's not an instance of say, helping the clerk in Macy's to improve service to the customer; it's anticipating what people want before they know they want it.

TD: Yes. It's a totally new business model.

JE: If you are a developer thinking about developing systems with the perspective of augmentation vs. automation, what do you do? What do people do that works to make a good human + machine system?

TD: Ideally, you involve the workers from the beginning in the design of the system and decisions about what's going to be done by humans and what's going to be done by machines. We have some examples of that in the book. One is the digital life underwriter. The developers not only talked to the underwriters up front about what would make sense for the system to do, but they had meetings every week. It was a way not only to improve the system and the underwriting process, but also to engage the underwriters and provide socialization for people who work remotely all the time.

One of the other key issues is the need for high levels of experience on the part of the human in these human + machine systems. There are a few cases where the AI really helps entry level workers, but for the most part, augmented intelligence helps people with experience. It raises a question: if we automate, how are we going to get the experienced workers needed for the future?

JE: There's an historical thread about work that parallels technology development but is focused on the quality of work life movement and the sociotechnical systems movement, each of which has its own set of principles. Do you find those principles get applied in the development of augmented intelligence?

TD: In a word, no. I think that they should be, but they are not. We studied a few cases where companies hired what are effectively sociotechnical design companies. IDEO does this, for example, but people are just forgetting the fact that these are issues we've been wrestling with for a long time.

I was talking to some people at Shell, and the guy who's the head of implementing AI for their business was also a re-engineering guy, and he's thinking a lot about how to design a process that takes both human capabilities and AI into account. But this is still pretty rare.

JE: The very first expert system I worked on was in the oil industry, designing cements for oil wells. One of the things

There are a few cases where the AI really helps entry level workers, but for the most part, augmented intelligence helps people with experience.

I learned was that you can design a system that works very well in getting the right answer, but it still might be rejected by the people in the field. In our case, we created a social problem for field engineers who had previously just relied on cementing companies for the design. Our system put them in a difficult social situation. They weren't experts in cementing, and now they had to adjudicate between the system and a rep from the cementing company.

TD: Oh, interesting.

JE: They said, "Well, now I have to choose between his design and your design. How do I do that?" We had to redesign the system to make it critique the service company design and make recommendations about how to improve it. We learned that you have to take into account the social context of a system, not just its usability.

You mentioned the use of participatory design as a tool in your book. Did you come across instances where designers got deeply into participative design, not just for upfront input, but throughout the design process?

TD: Only the digital life underwriter application. They didn't really think of it as a participatory design, but they had ongoing involvement and there was a lot pivoting.

JE: In one of the closing chapters of the book, you focus on the unique qualities of humans. wherein what areas do human beings add the most value. What advice do you have for people who want to develop really good systems that make use of both computing and human capabilities?

TD: My first book on this topic, *Only Humans Need Apply*, focused on that question. We identified five ways for humans to work relative to machines. One is for people to really understand how the machine works and to pick up the ball when the system drops it. We called that "stepping in." Another we called, "stepping up," which put the human in a supervisory role to decide when more or less reliance on the system was appropriate. I think of that as a hedge-fund model; the human is monitoring how the portfolio is doing and deciding when changes are necessary in the AI systems that do the trading.

A third was "stepping aside," where people seek out work that machines can't do—although that's always changing. Fourth was "stepping narrowly," where the role of the human is very focused and not economical to automate, and the last one was "stepping forward," where people help to develop new AI systems and capabilities.

Most of the jobs are in the "step in" category, where you try to learn what the machines are capable of and what they're good and not good at. I think it's unfortunate that in most companies, this kind of automation is undertaken at an individual level—one-off activities.

JE: One of the challenges of this kind of work is that you have to have a good appreciation of what the technology can do, but you also have to have a mindset around the work that you're trying to understand.

You get lucky if someone's got both sides of the equation. In one IT organization I managed, we hired retirees who were shepherds about the operations for developers and helped them to understand the work from the worker's perspective. What have you seen developers do so that they are able to get themselves into the mindset of the work they're trying to support?

TD: Some have suggested that we need more AI-oriented ethnographers who can understand how the work is being done today. I think I agree with that. The problem with ethnographers, though, is they're mostly interested in describing the present details of work without really thinking much about how we can change it and improve it. I do think we likely need to come back to some of the re-engineering approaches to figure out how AI and other emerging technologies can enable new ways of working. But we shouldn't use the top down approach to re-engineering we used in the past; it should be more participatory.

JE: I'd like to ask you a few questions about the ethics of some of the AI systems, especially the ethics associated with black box learning algorithms. What do you see as the biggest ethical concern?

TD: We didn't encounter many ethical concerns in our research. One was in the Wilmington, North Carolina, police department where they were using a "precision policing" solution. The vendor, whom I generally have a lot of respect for, is Shot Spotter. They were developing systems to tell cops that if you want to catch bad guys, you should probably spend more time here than there. Policing is, obviously, politically potent, so you have to be quite careful about automating functions. There's a natural tendency of systems to recommend focusing on the areas that have the most problems, and in many cases, those are going to be heavily minority populations with low incomes. You have to be quite careful about what the system is recommending and ensure that it is not biased against disadvantaged populations.

Over the last year or two, I've definitely seen more attention being paid to AI ethics-related issues.

JE: I think that there is an awakening about the human element and AI. I think that we are only beginning to think about the ethical issues related to work and work design.

I do think we likely need to come back to some of the re-engineering approaches to figure out how AI and other emerging technologies can enable new ways of working.

TD: I think that's true. A few years ago, a few books raised the issue as it related to impacts on employment. Two prominent books were *No Humans Need Apply* and *The Rise of the Robot.* They predicted that the time would come when we weren't going to have any jobs anymore. There was also a big Oxford study that concluded that 47 percent of US jobs wouldn't exist in 10 years.

JE: But others will replace them. . .

TD: Exactly. Probably, anyway.

JE: What do you see in 5 or 10 years for augmented intelligence?

TD: A lot of new jobs are being formed that are in this space, but the role of the human with respect to the machine will change. I do think that AI will continue to eat away at the margins of what we do, so we'll have to work harder to add value as individuals to focus on those places where humans add value. Eventually, maybe, AI will do everything better than we do. Then all bets are off, but I think that you and I will be dead by then.

JE: Unless AI and medicine discover new ways of keeping us alive! Your next book is about the effect of AI and augmented intelligence on business models. Can you give us a little preview of your findings in that?

TD: It's about companies that are very aggressive in their use of AI. In many cases, what they do is to create new business models as well as new products and services. One example is the Morgan Stanley system I described earlier. Another one is a system for automated estimate creation for cars that have been in collisions from a company called CCC Intelligent Solutions. But to me, one of the most interesting things is how new ecosystems of companies that use AI are emerging, first to aggregate the data needed for an application, and then to provide value to customers on the basis of that data.

Airbus, for example, has created the Skywise ecosystem of all of the airlines that use Airbus equipment. It connects all of the Airbus customers. They use Palantir's AI to integrate the data and then add value to it with things like predictive maintenance, fuel consumption optimization, and route optimization and scheduling. Those are the kinds of things that the ecosystem does jointly.

We found a lot of cases of the large scale use of AI for operational transformation, too. Shell and Kroger are both big users of AI. Many companies are also using AI to shape customer behavior, which was pioneered—with some negative outcomes—by the social media companies. Now there are companies that are trying to change health behaviors using AI.

JE: When you look at the incumbent companies that are successful with AI, would you say that success is based mainly on how aggressive they are or how systemic they are or how willing they are to look at their underlying business model? What's driving the AI agenda of successful incumbents?

TD: I think it's all of the above. It's senior leaders deciding that AI is a transformational technology. It's the awareness that digital native companies are exploiting AI aggressively and getting great results.

One of the poster children in our book is the DBS Bank in Singapore. The CEO there, Piyush Gupta, is really impressive. He has declared that "We're not really competing with other banks. We're going to be competing with Google and Amazon. Not only do we have to use this stuff; we have to develop it ourselves. We can't just buy stuff from everybody else." So far, he has focused on operational transformation—no branches, fantastic chatbots. You never have to visit a branch! I think success requires that senior level commitment. It's got to be broad and deep and it requires that you make a big investment.

JE: And manage the organizational issues.

TD: Yes. You need to get a lot of people involved in helping to create these solutions, and you need to educate people about what the transformation means. Capital One is one of the companies focused on engaging a broad range of people They're trying to either hire or retrain 5,000 machine learning engineers in one year. It's just extraordinary.

JE: One final question. It seems to me that to be successful with AI, we have to figure out how to make visible the social side of technology—the things that people currently skip, like being out in the field, engaging the whole village internally. We have to be able to make the business case for spending time on work and the social systems of work, and we just don't have the language yet. We also need tools that embody these factors, tools that can scaffold for a broader set of design principles.

TD: I still think that we are lacking systems that are really good at supporting product development in that sense, whether you're talking about a physical product or a data product or a software product. We need tools that are collaborative and agile and get a broad cross-section of people involved. We need something a little more structured than Slack, but less structured than the product development systems that we've become used to. That would be really helpful. Some companies are making great progress, but we are still learning how to do this well, and we need help in spreading the wealth from AI to more companies and people.