The Ghosts of Tax Past, Present, and Future

by Michael D. Sontag, Stephen J. Jasper, Sara K. Morgan, and Robert C. Guth

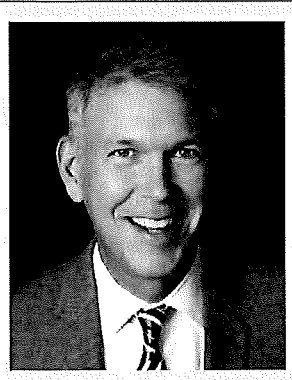
Michael D. Sontag

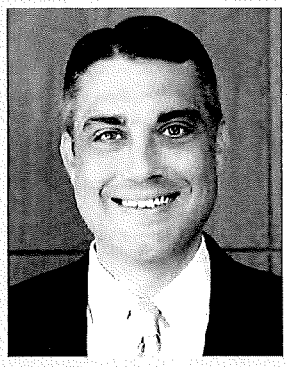
Stephen J. Jasper


Sara K. Morgan

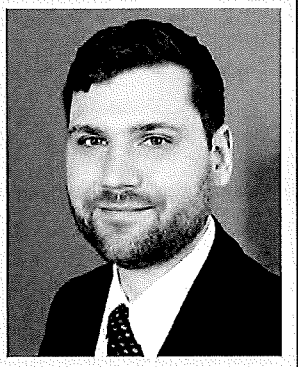
Robert C. Guth

Michael D. Sontag, Stephen J. Jasper, Sara K. Morgan, and Robert C. Guth are attorneys at Bass, Berry & Sims PLC and are part of the firm's state and local tax practice group.

In this installment of Bass Tax, the authors examine artificial intelligence and how it might help tax practitioners understand the challenges at the core of applying traditional tax concepts to these unprecedented new technologies.

Copyright 2024 Michael D. Sontag, Stephen J. Jasper, Sara K. Morgan, and Robert C. Guth.
All rights reserved.

In May 1997 people watched a ghost play chess. Many directly witnessed the event, and recordings and reports abound, yet scientists and observers still debate exactly what happened. The facts are these: Something not human guided chess pieces across the board. The moves were not random, were valid, and were by and large "good" moves. One of the world's best chess players volunteered to play against the invisible opponent. The ghost played very well. The pair played six games; three ended in a draw, the entity won twice, and the human player won only once.

The ghost in this story was Deep Blue, a sophisticated artificial intelligence system painstakingly built by IBM specifically to master chess. But Deep Blue played too well. No hardware in the 1990s — not even Deep Blue's — could process enough information to play chess at the level Deep Blue played. It was mathematically impossible. Something beat Garry Kasparov that day, but it was not the hardware, the AI running on that hardware, or the programming coded into it by IBM's engineers.

Of course, there is a perfectly rational explanation. And understanding how Deep Blue played beyond its apparent ability is essential to understanding what AI systems are and how they work. More to the point, removing some of the mystery about what AI is and how it works will help tax practitioners understand the challenges at the core of applying traditional tax concepts to these unprecedented new technologies.

But be warned: Like any good ghost story, this article will raise more questions than answers. Finding workable answers to those questions will be critically important as the use and development of AI continues to become

For more *Tax Notes® State* content, please visit www.taxnotes.com.

© 2024 Tax Analysts. All rights reserved. Tax Analysts does not claim copyright in any public domain or third party content.

more integrated into everything businesses and consumers do in the years to come.

Intelligence, Illusions, and Ghosts

To understand how Deep Blue played beyond its apparent limitations, it is important to first unlearn everything you think you know about modern software.

Ordinary software programs must be precise and specific. They are instructions given by developers for the computer to slavishly carry out. Computers cannot roll the dice or pick a random number, for example. To mimic chance, programmers must tell the computer to pick a number based on an elaborate math problem or other logic puzzle that looks and feels random enough for the purpose. With elaborate instructions and clever development, ordinary software can do some amazing things, sometimes even creating the illusion of intelligence. When prewritten software does things that look and feel intelligent, that is called symbolic AI.

But symbolic AI is not actually AI. Actual AI involves computers making decisions not directed or planned by a developer. To make those decisions, AI systems work more or less like elaborate guess-and-check machines. They create code, run it, test their work, and learn. Over time, AI can amass more successful ways to solve familiar problems and improve its response strategies by creating new code. In other words, AI systems learn and grow. These fundamentals of machine learning have been around for decades, but recent advances in hardware technology have allowed AI to test bigger ideas more quickly and against more data. Combining machine learning with today's technology results in some serious brain power — allowing these machines to appear to think on their own.

Importantly, though, just as unthinking software can be programmed to perform intelligent tasks, complex AI systems can be supported to perform tasks beyond their technical ability. In this way, programming and people still play a role in the activities and decisions of modern AI. Special-purpose programming often provides AI systems with direction and a safety net of proven solutions and strategies for specific tasks. Quality of training through human interaction can also have a tremendous impact on how quickly and how well an AI system learns. Through this collaboration, humans and AI systems work together to achieve things individuals and AI could not on their own.

In short, some AI systems teach themselves and perform their tasks autonomously. Others cannot perform at all without the support of an elaborate, interdependent organization of people, software, and even other AI. Most fall anywhere and everywhere in between, as varied and variable as any commercial operation. And in 2023 AI systems were integrated into all aspects of our daily and professional lives, "embedded in devices we use and interact with on a daily basis — for example, in our smartphones, wireless routers, and cars."[1]

It can be impossible for an outside observer to tell where the decisions made by these systems are coming from. The term "ghost in the machine" originated with human philosophy to describe the concept of the mind being distinct from the body. Modern AI systems have grown complex enough to present similar mysteries. Though experts assure us that AI systems have not yet reached sentience or anything close to it,[2] the complexity of AI often makes it difficult to know how a specific AI system actually works. The result of all this is that nagging feeling that a computer system may be smarter and more aware than it should be.

Ultimately, however, AI systems are nothing more than collaborative efforts that can combine a full spectrum of the following technologies: unthinking programs, all manner of other AI, and the often unsung and underappreciated efforts of lots of people — including the user experience teams working to package all this effort into the illusion of a singular machine.

The question facing tax practitioners, tax enforcement agencies, and businesses is how to properly characterize this unique combination of

[1] National Security Commission on Artificial Intelligence, "Final Report," at 33 (Mar. 2021).

[2] *See, e.g.*, Patrick Butlin et al., "Consciousness in Artificial Intelligence: Insights From the Science of Consciousness," Cornell University (Aug. 22, 2023).

For more Tax Notes® State content, please visit www.taxnotes.com.

© 2024 Tax Analysts. All rights reserved. Tax Analysts does not claim copyright in any public domain or third party content.

© 2024 Tax Analysts. All rights reserved. Tax Analysts does not claim copyright in any public domain or third party content.

programming, people, and independent intelligence for state tax purposes. We know how to tax people and machines, but how do you tax a ghost in a machine?

How Should a State Tax a Ghost?

The complexity, variability, and decision-making ability of AI undermine much that tax systems expect and assume about software. Software has been straining traditional tax concepts for years, but AI systems may be the straw that breaks the camel's back.

To get an idea of the problem, try applying a typical sales tax rule for software to a proper understanding of a modern AI system. For example, Tennessee imposes sales tax on any transfer or use of computer software, defined broadly to mean any "set of coded instructions designed to cause a computer or automatic data processing equipment to perform a task."[3] Tennessee also imposes sales tax on "any programming, transferring, or loading of computer software into a computer,"[4] as well as any "access or use of software."[5] To a software engineer, these broad phrases describe literally any use of a computer — from sending an email to visiting a website.

The state clearly does not intend to cast such a wide net. The law states that "use of computer software" shall not be construed "to impose a tax on any services that are not currently subject to tax."[6] This only compounds the confusion, however, because it is the nature of the service performed by software that seems to distinguish the taxable transfer of software from, say, services delivered by email. Clearly the state intends to tax *something*, but the laws fail to articulate what, precisely, should be taxed in the age of AI. The state generally uses the true object test to handle this uncertainty,[7] but that is of no use when it is not clear what objects the state intends to tax, and the true object of any software is the information or function provided — not the "set of coded instructions" themselves. These problems are true for all software systems, but especially for AI systems. How can the true object of the transaction be the set of coded instructions that are inaccessible, unknowable, and never the same twice?

Of course, there is nothing that necessarily prevents Tennessee from taxing whatever software-related transactions lawmakers had in mind. But software is losing its relevance as the focal point of whether a transaction is taxable. The discrete and identifiable thing that lawmakers imagined software to be is becoming an illusion — a ghost in the machine.

The fundamental problem is that for decades, tax laws have treated software as a product — anchored heavily by memories of a time when software could be purchased off the literal shelf. AI systems — and most modern software systems — are complex business operations that are distinct from those old-fashioned concepts. And understanding and explaining how an AI system differs from a software product can be overwhelming.

So, how should something so complicated be condensed into a workable framework? AI systems are complex combinations of people, property, and purpose that can be organized in innumerable combinations. As it happens, there is a term for that: a business. Tax professionals routinely work with or for business organizational structures that are just as complex as an AI system — or any software system, for that matter.

Perhaps the answer is to imagine AI as just another part of a business — and specifically, part of a business's workforce. There is no obvious boundary between AI as intangible property and AI as an intelligent agent or actor, and AI systems are rapidly advancing into this uncharted territory by making decisions and acting on behalf of companies. Microsoft, for example, has used an AI chatbot to give voice and personality to its search engine Bing, and the chatbot has already created headlines with the

[3] Tenn. Code Ann. section 67-6-102(18).

[4] Tenn. Code Ann. section 67-6-102(86)(K).

[5] Tenn. Code Ann. section 67-6-231(b).

[6] *Id.*

[7] Tennessee Department of Revenue, "Sales and Use Tax Manual," at 152-153 (Dec. 2023).

For more *Tax Notes® State* content, please visit www.taxnotes.com.

things it has said and done as a spokesperson (or spokes-something) of the company[8] — illustrating the vast differences in autonomy between AI systems and traditional software.

In these instances, it seems logical to treat AI as an agent of a business or as indistinguishable from an employee. But treating AI activities as the actions of a business leads to several dicey state tax questions. Can AI's autonomous decisions and actions establish nexus? When might AI's actions cause a business to lose the protection of P.L. 86-272? When is the taxation of AI services akin to taxing human services in violation of the Internet Tax Freedom Act? In a time when tax laws still treat software as something like a discrete product or application, it is difficult to answer these questions, or to even imagine what rules and doctrines courts might apply.

If AI were to be treated as an agent for a business, it would probably be best to no longer focus on *what* made the decisions or who conducted the activity, but rather on who authorized the decisions or activities in the first place. Those kinds of questions can be applied to teams that support an AI's deployment in much the same manner that similar questions might be asked regarding a sales team, operating division, or other employees. Then again, if an AI system is truly acting on its own, is it accurate to say a business authorized its activity at all?

A Focus for the Future

So how did Deep Blue overcome its technical limitations? Ideally, the answer is obvious. Like most modern AI systems, the Deep Blue system that chose chess moves was not just a machine, not just an AI program, and not a puppet playing out moves directed to it by human instruction. The whole was greater than its parts.

As promised at the beginning of this article, having this understanding of how Deep Blue worked raises more questions. But a proper understanding of what AI is and is not is an important first step in figuring out how it should be taxed and how its activities should be attributed to a business for state tax purposes. Resolving these issues is imperative, because AI systems are fully embedded into all aspects of our lives — all the while presenting novel questions about how these hybrid software, human, and intelligence "machines" should be treated. Finding workable answers to these questions should be a priority for tax practitioners, tax agencies, and businesses in 2024 and beyond. ■

[8] Kevin Roose, "A Conversation With Bing's Chatbot Left Me Deeply Unsettled," *The New York Times*, Feb. 17, 2023.

For more Tax Notes® State content, please visit www.taxnotes.com.

© 2024 Tax Analysts. All rights reserved. Tax Analysts does not claim copyright in any public domain or third party content.