



# **JTC Resource Bulletin**

---

## Introduction to AI for Courts

---

Version 1.0  
Adopted 27 March 2020

## Abstract

Many court technology systems today are leveraging some forms of AI. However, there are many more use cases where AI technologies might advantageously be brought to bear in the court setting.

## Document History and Version Control

Version	Date Approved	Approved by	Brief Description
1.0	3/27/2020	JTC	Initial release.

## Acknowledgments

This document is a product of the Joint Technology Committee (JTC) established by the Conference of State Court Administrators (COSCA), the National Association for Court Management (NACM) and the National Center for State Courts (NCSC).



### **JTC Mission:**

To improve the administration of justice through technology

### **Augmented Intelligence Focus Group:**

---

IV Ashton  
LegalServer

Karl Branting  
MITRE

David Byers  
Arizona Supreme Court

Roger Rand  
Multnomah Circuit Court (Oregon)

Justice Deno Himonas  
Utah Supreme Court

Anthony Capizzi  
Montgomery County Juvenile Court (Ohio)

Chris Draper  
Trokt

Brett Howard  
Orange County Superior Court (California)

Casey Kennedy  
Texas Office of Court Administration

Jack McCarthy  
New Jersey Courts

Nicolas Vermeys  
University of Montreal, Cyberjustice Laboratory

David Slayton  
Texas Office of Court Administration

Gary Marchant  
Arizona State University College of Law

Anil Sharma  
IBM Federal

Amy Schmitz  
University of Missouri-Columbia School of Law

### **National Center for State Courts**

---

Paul Embley  
Jim Harris  
Di Graski

Compiled and written by Lise Embley, technical writer, National Center for State Courts

### **Additional contributors**

---

Siva Appavoo, New Jersey Courts  
Cindy Guerra, Palm Beach County, Florida  
Dory Reiling, Ph.D. Mag.lur., Retired Senior Judge, Amsterdam District Court (Netherlands)

## **Joint Technology Committee:**

---

### **COSCA Appointments**

David Slayton (Co-Chair)  
Texas Office of Court Administration

David K. Byers  
Arizona Supreme Court

Jeff Shorba  
Minnesota Judicial Branch

Rodney Maile  
Administrative Office of the Courts, Hawaii

Laurie Dudgeon  
Kentucky Administrative Office of the Courts

### **NCSC Appointments**

Justice Deno Himonas  
Utah Supreme Court

The Honorable Samuel A. Thumma  
Arizona Court of Appeals

### **Ex-officio Appointments**

Joseph D.K. Wheeler  
IJS Courts Advisory Committee

### **NACM Appointments**

Kevin Bowling (Co-Chair)  
Michigan 20<sup>th</sup> Judicial Circuit Court

Paul DeLosh  
Supreme Court of Virginia

Roger Rand  
Multnomah Circuit Court, Oregon

Kelly C. Steele  
Florida Ninth Judicial Circuit Court

Jeffrey Tsunekawa  
Texas Office of Court Administration

### **CITOC Appointments**

Jorge Basto  
Judicial Council of Georgia

Casey Kennedy  
Texas Office of Court Administration

### **NCSC Staff**

Paul Embley  
Jim Harris

## Table of Contents

Abstract .....	ii
Document History and Version Control .....	ii
Acknowledgments .....	iii
Table of Contents .....	v
Executive Summary .....	vi
Introduction .....	1
AI Technologies .....	2
Visual Perception .....	2
Optical Character Recognition .....	3
Natural Language Processing .....	3
Symbolic AI .....	4
Machine Learning .....	4
Where are humans in “the loop”? .....	5
Vehicle Automation .....	5
Chatbots .....	6
Where are courts in “the loop”? .....	7
Automated Docketing .....	7
Legal Information/Assistance .....	7
Beyond Efficiency .....	8
Common Sense and Ethics .....	8
Conclusion .....	10

## Executive Summary

AI (artificial or augmented intelligence) is trendy. Every internet search, Siri response, Uber ride, etc. is made easier, cheaper, faster, more accurate, and more convenient through technologies encompassed in the term AI. AI is the new “normal.” Many court technology systems today are already leveraging some forms of AI. However, there are many more use cases where AI technologies might advantageously be brought to bear in the court setting. This paper explains AI basics for non-technical court personnel.

## AI Technologies

AI encompasses a broad range of technologies, some of which are already so widely accepted (e.g., spell check) that they are often no longer included in the growing list of AI-enabled applications. A single system may leverage one or more AI technologies:

### Visual Perception

Facial Recognition Technology (FRT), radar, Light Detection And Ranging (LIDAR) and ultrasonic sensors give machines visual perception - the ability to sense and interpret objects. Application: robotic vacuum cleaners, self-driving cars. Judges in Marion County, Oregon sign into court systems via facial recognition.

### Optical Character Recognition

Electronically capturing information from typed, handwritten or printed text is Optical Character Recognition (OCR). Application: USPS mail routing, online check deposit. In Palm Beach County, Florida, OCR is being used to scan incoming e-filed documents to docket them automatically.

### Natural Language Processing

Digital assistants like Alexa, Siri, Cortana, and Google Assistant use Natural Language Processing (NLP) in a conversational interface to extract meaning from human speech. Application: chatbots. JAI, the New Jersey courts' chat bot, leverages AI to respond appropriately to increasingly complex questions.

### Symbolic AI

Symbolic AI closely reflects classic legal reasoning: step 1, write down all the rules; step 2, apply relevant rules to individual fact patterns to reach a conclusion. Application: Legal navigator websites like [Florida Law Help](#) and [Colorado Resource Network](#) for senior citizens

## Machine Learning

Humans define a business goal and present a set of known cases or data to train the system, then computers determine the optimal algorithms to achieve the goal. Application: helps New Jersey’s chatbot respond to an increasingly broad range of inquiries. Several courts are exploring the potential to identify “red flags” in guardianship/conservatorship cases, triage cases for automated case management, and provide mediation suggestions

## Where are humans in “the loop”?

Since no automated process is perfect, the risks associated with AI systems are mitigated through a variety of automation strategies that define human oversight.

human-in-the-loop      no response/decision without human involvement

human-on-the-loop      humans can override an automated action before it occurs

human-out-of-the-loop      systems learn from users and deliver automated responses

## Where are courts in “the loop”?

Courts are using AI to handle repetitive, non-value-add processes like auto-docketing and to deliver higher quality, more efficient service to the public through chatbots.

## Beyond Efficiency

AI tools can eliminate or streamline many manual processes, allowing the same number of staff to better serve more members of the public. AI can also help identify previously unobserved correlations more efficiently.

## Common Sense and Ethics

Caution is important in the use of AI. Before deciding which human processes to offload onto AI-enabled systems, planners should consider both how the AI system could produce errors, and how consequential those errors could be. The US Department of Defense (DoD) Ethical Principles for AI could be applied to court uses, as well.

## Conclusion

Augmented Intelligence is helping courts do some things that humans do, only better and faster. Courts that are leveraging AI effectively are starting small, building both confidence and expertise.

## Introduction

AI is trendy. The acronym is used interchangeably to represent “Artificial” Intelligence or “Augmented” Intelligence, but the terms are not synonymous. The difference in those two interpretations of AI is both subtle and significant. While “augmented” signals the power of technology to enhance human capabilities, “artificial” carries the uncomfortable connotation of something fake or cheaply concocted to approximate a more costly but desirable option. (Think: artificial vanilla flavoring.) Because AI processes and algorithms can be shrouded in complexity or by intellectual property protections, “artificial” intelligence can feel like a façade used to conceal something that if seen clearly would be undesirable. “Augmented” Intelligence is both more accurate and more palatable. For courts, AI is Augmented Intelligence, the use of technology to do what humans do, only faster and better.

In spite of the current hype, the concept of AI is essentially as old as computers. When reel-to-reel tape and punch-card computers became available in the 1950s, scientific thought leaders and sci-fi fiction authors alike were contemplating the use of machines to simulate human thinking. Even so, the prevalence of AI in ordinary daily life today would likely impress the most visionary of those thought leaders.

Every internet search, Siri response, Amazon Prime purchase, FedEx delivery, airline flight, and Uber ride is made easier, cheaper, faster, more accurate, and more convenient through technologies encompassed in the term AI. There are many types and applications of AI. Most customer service call processing centers use some form of *Speech Recognition* and *Natural Language Processing* (NLP) to route callers to the right resources. Algorithms can better predict outcomes and trends using *Machine Learning*. Even paper processes including “snail mail,” paper bank checks, and hardcopy tax filings are processed by *Optical Character Recognition*. AI is the new “normal”: it is already routine and ubiquitous in the lives of most Americans. The applications and importance of AI in all aspects of our lives is expected to grow rapidly over the next few decades.

Many court technology systems today are already leveraging one or more “flavors” of AI. However, there are many more use cases where AI technologies might advantageously be brought to bear in the court setting. The purpose of this paper is to explain AI basics for non-technical court personnel to help facilitate conversations with technology providers, as well as to identify current and potential beneficial court uses.



## AI Technologies

AI encompasses a broad range of technologies. “AI is an ocean.”<sup>1</sup> At the same time, many AI-enabled applications like spell check and internet search technologies are so widely accepted that they are often no longer included in the growing list of AI-enabled applications.

A single system may leverage one or more of those technologies. The following is a brief overview of some of those use cases, arranged roughly by AI subspecialty.

### Visual Perception

Multiple technologies including Facial Recognition Technology (FRT), radar, Light Detection And Ranging (LIDAR) and ultrasonic sensors give machines visual perception - the ability to sense and interpret objects. Millions of households use robotic vacuum cleaners with sensors and mapping capabilities that help the vacuums navigate around obstacles like couches and puppies while effectively cleaning floors. Much less common but probably more often discussed are self-driving cars, which use a combination of visual perception technologies to avoid obstacles including pedestrians, road hazards, and other cars.

High end hotels and casinos use facial recognition to identify and greet wealthy patrons as well as problem gamblers. Most smartphone manufacturers now employ facial recognition and/or biometric authentication mechanisms to help users quickly confirm financial transactions and easily unlock their phones without entering a pin. The same technology powers popular social media and photo management software features that automatically tag photos. Law enforcement agencies are now using a controversial app<sup>2</sup> that combs through billions of images scraped from social media sites and gathered from DMV photo records to identify unknown individuals. The app has already solved crimes ranging from credit card theft to child sexual exploitation and murder.

London Metropolitan Police are deploying facial recognition cameras in key locations to help identify wanted suspects and missing children.<sup>3</sup> Judges in Marion County, Oregon sign into court systems via facial recognition, improving security as well as efficiency: password issues no longer cause delays.<sup>4</sup>

---

<sup>1</sup> Interview with Siva Appavoo, New Jersey Courts. 18 February 2020.

<sup>2</sup> Kashmir Hill, [The Secretive Company That Might End Privacy as We Know It](#), *New York Times*. 18 January 2020. Web.

<sup>3</sup> [“Met Police to deploy facial recognition cameras.”](#) *BBC News Service*. 30 January 2020. Web.

<sup>4</sup> Email correspondence with Multnomah County Circuit Court IT manager, 18 March 2020.

## Optical Character Recognition

One of the oldest AI technologies, Optical Character Recognition is the electronic process used to capture information from typed, printed, or handwritten text. Its origins date to the early 20<sup>th</sup> century in "reading machines" for the blind and a device used to encode telegraph messages.<sup>5</sup> The United States Postal Service has used OCR, not humans, to sort mail for at least ten years. Smartphone apps use the same underlying technology to capture information from paper checks deposited digitally. Increasingly sophisticated OCR now deciphers handwriting to glean the payee and amount, as well as bank routing code and account number information encoded along the bottom of paper checks – the few that are still being written. Indeed, OCR is one of those AI technologies that has become so ubiquitous many experts no longer consider it AI.

In Palm Beach County, Florida, OCR is being used to scan incoming e-filed documents to docket them automatically. The system verifies the case number and extracts and captures the document title with other required information that is then passed automatically to the court's case management system.

## Natural Language Processing

Digital assistants like Alexa, Siri, Cortana, and Google Assistant use Natural Language Processing (NLP) in a conversational interface to extract meaning from human speech. NLP also facilitates text-based language assessment, for example, email spam filters and grammar/spell-check features. Automated customer service, language translation apps, search engines, and virtual assistants all use NLP.

Speech recognition is only the first step in facilitating a conversation. Accurately assessing meaning in human speech (and human-created text) requires the ability to interpret slang and sarcasm, analyze emotion, and detect cultural nuances. Statements that appear to be superficially identical can mean very different things; to initiate the appropriate response (and/or avoid delivering a completely inappropriate response), chatbots and other NLP tools must be able to discern context. For example, "I don't know what you are talking about" could mean either "I don't understand" or "you are mistaken."

Sophisticated language processing capabilities including sentiment analysis and automated summarizing are being employed in some Online Dispute Resolution

---

<sup>5</sup> Schantz, Herbert F. (1982). *The history of OCR, optical character recognition*. [Manchester Center, Vt.]: Recognition Technologies Users Association. ISBN 9780943072012. Internet Archive.

(ODR) apps to help de-escalate interactions between parties. NLP capabilities are frequently used in phone-based customer service apps, and text-based chat bots are becoming increasingly common. Since text is the preferred communication mechanism for some demographics,<sup>6</sup> text chatbots could become particularly effective tools for courts.

These AI applications and more may be enabled with one of two ‘learning’ approaches: Symbolic AI, the dominant AI strategy for over 30 years, or Machine Learning, which is becoming increasingly common in today’s systems.

## Symbolic AI

Expert systems leverage ontologies that define the meaning and relationship of things (think: tagging) and algorithms, or step-by-step procedures for arriving at an answer, applying business rules deductively to new cases. Symbolic AI is the most familiar and therefore comfortable to the judicial branch because it closely reflects classic legal reasoning: step 1, write down all the rules; step 2, apply relevant rules to individual fact patterns to reach a conclusion.

Courts today use Symbolic AI in guided questionnaires for document generation (decision trees based upon business rules), workflow engines for automating case management, most risk/needs instruments, and new legal navigators like [Florida Law Help](#) and [Colorado Resource Network](#) for senior citizens.

## Machine Learning

In machine learning (also referred to as inferential AI), humans define a business goal and present a set of known cases or data to train the system, then computers determine the optimal algorithms to achieve the goal. Inferential AI is inductive: computers calculate probabilities, like IBM’s Watson playing *Jeopardy* or AI tools used in healthcare to detect harmful bacteria in blood or malignancies in mammograms. Inferential AI is succeeding because of the availability of Big Data and massive computing power enabled by artificial neural networks.

Several courts are exploring potential uses of Inferential AI to identify “red flags” in guardianship/conservatorship cases, triage cases for automated case management, and provide mediation suggestions for online dispute resolution. Others are gathering data to begin developing and training ODR tools that may not actually be implemented for several years.

---

<sup>6</sup> See Amanda Mull, “[Talk to People on the Telephone – It’s time to start calling your friends again.](#)” *The Atlantic*. 16 September 2019. Web.

## Where are humans in “the loop”?

Since no automated process is perfect, the risks associated with AI systems are mitigated through a variety of automation strategies that define human oversight. Pim Haselager, Associate Professor at Radboud University in the Netherlands describes that oversight in three basic levels:<sup>7</sup>

- human-in-the-loop
- human-on-the-loop
- human-out-of-the-loop

Broadly speaking, **human-in-the-loop** means no response or decision is delivered without some form of human involvement. That involvement may take the form of a human reviewing and approving a machine-generated response to a user question before the answer is delivered. The first step in “training” an AI system is often human “in-the-loop” until the system responds so accurately that the human no longer adds value.

On the other end of the spectrum of human involvement, systems learn from user questions, identifying correlations more efficiently than humans and delivering automated responses directly to the user, leaving humans **out-of-the-loop**. Between those extremes lie the bulk of current AI implementations: **human-on-the-loop** AI, where humans have the power to intervene and override an automated action before it occurs.

Vehicle automation and chat bots help illustrate the variations possible in the implementation of AI.

### Vehicle Automation

In the spectrum of vehicle automation, **human-in-the-loop** systems are “driver assist” features that warn of vehicle proximity to other vehicles/objects: following another vehicle too closely, sensing vehicles in the “blind spot” when making lane changes, warning of objects in the road, measuring space to the curb in parallel parking, etc. The driver receives information from the automation to take evasive action. In **human-out-of-the-loop**, the car is driverless and the vehicle fully autonomous.

---

<sup>7</sup> Pim Haselager presentation at the [17<sup>th</sup> International Conference on Artificial Intelligence and the Law \(ICAIL\)](#). Montreal (Quebec), Canada. June 2019.




Illustration courtesy Chris Draper

<b>Human “in the loop”</b> No action taken without human action/affirmation.	<b>Human “on the loop”</b> Action taken unless a human intervenes.	<b>Human “out of the loop”</b> Humans are not involved from input data to action.
“Driver assist” or partial automation strategies including lane change warnings, emergency braking, adaptive cruise control, and parallel parking assistance.	Autonomous vehicle design strategy SAE Level 3 or 4 <sup>8</sup> , which requires the driver to continually monitor the system and take action to prevent a dangerous condition.	Autonomous vehicle design strategy SAE Level 5, which prevents the human from taking any action associated with driving the vehicle (e.g., cars built without steering wheels or brake pedals).

## Chatbots

A chatbot (chat+robot) is software that simulates human conversation. Service organizations, including courts, often use chatbots to provide initial customer service or technical support. Interacting with humans via voice or text, chatbots today help people make appointments with medical providers, reset system passwords, evaluate cell phone data plans, and much more. Like the variations of technology used to improve driver safety, chatbots can be implemented in ways that involve humans to varying degrees - in, on, and out of “the loop.”

<b>Human “in the loop”</b>	<b>Human “on the loop”</b>	<b>Human “out of the loop”</b>
A messaging plugin mimics existing customer service question-response guidance.  Or messaging software tags and filters user messages and suggests options that a human customer service representative chooses from to send a response to the customer.	The system’s experience creates new connections between questions and answers that were not previously interpreted to be related. Appropriate answers are delivered automatically; inappropriate answers can be prevented by a human overseeing the system.	3rd Generation Chatbots employ unsupervised machine learning in NLP routines that can automatically update language tags which define the effectiveness of the NLP being used. This type of chatbot is not in any broad commercial use as of 2020.

<sup>8</sup> See *SAE J3016 Levels of Driving Automation* at <https://www.sae.org/news/2019/01/sae-updates-j3016-automated-driving-graphic>

## Where are courts in “the loop”?

The use of algorithms in the justice space to predict recidivism and the likelihood of future violent offences may be the most familiar and controversial uses of court AI. Much work remains to ensure those kinds of algorithms are unbiased and explainable. But AI isn't just predictive. Courts are using AI to handle repetitive, non-value-add processes and to deliver higher quality, more efficient service to the public.

### Automated Docketing

In Palm Beach County, Florida, AI-empowered software is classifying and docketing e-filed documents. The court started with three low risk/high volume case types, progressively expanding the variety and complexity of cases as they developed expertise with the robotic process automation (RPA) technology. The bots – each with its own name and user login -- classify incoming e-filings, extract info from tagged fields, and docket them in the court's case management system. Today, 68 case types representing nearly a third of all Palm Beach County's e-filed documents are being docketed automatically.

When the court first launched the system, humans double-checked 100% of the bots' work to verify accuracy (Human “in the loop”). That turned out to be more about reassuring the humans than about ensuring quality: the bots make fewer errors than human clerks and bot errors are ultimately an indication of a human programming/set up error. When errors are discovered and corrected (human “on-the-loop”), the robot never repeats the mistake (an accomplishment most humans could not claim). Today, humans review 15% of all filings, whether docketed by a human or a bot.

As the bots' human handlers have become more adept at using the software, the bots have been assigned increasingly complex work. Using “learning by example,” the bots have been taught to recognize and handle some kinds of filings that have additional circuit requirements: the bots look up the judge assigned to those cases and automatically email relevant documents (human “out of the loop”).

### Legal Information/Assistance

New Jersey state courts are handling an increasing number of public inquiries using a chatbot they've named JIA. To build the system, court staff assembled Q&A pairs using website FAQs, standard operating procedures, manuals, and other existing information resources. Through a carefully phased rollout, the court has developed the necessary internal staff expertise while ensuring the quality and accuracy of information provided to the public.

To train and test the system, JIA was initially released to AOC central office staff only. Staff manually entered all call center inquiries without necessarily utilizing the answers JIA supplied (Human “in the loop”). On a daily basis, court staff reviewed a report of all inquiries and answers, adjusting JIA responses and adding question variations to train the system. When JIA was responding at an 80% accuracy level, the system was then released to more than 10,000 state court staff. Response accuracy again dropped to about 30% as state court staff asked questions the central office staff had not anticipated. Additional Q&A pathways were added until JIA was responding with 80% accuracy. The system was then added to the court’s website making it live to the public (human “on-the-loop”) without any kind of formal announcement. As the volume of text inquiries has increased, the court has seen a roughly corresponding decrease in the number of call center calls.

## Beyond Efficiency

There are many AI tools which are both reliable and widely accepted, and the number and uses of such tools will continue to expand over time. While current AI tools still require significant, skilled human effort to set up and monitor, using these tools effectively and thoughtfully can significantly reduce overall workload and increase staff effectiveness. AI tools can eliminate or streamline many manual processes, allowing the same number of staff to better serve more members of the public. For example, Palm Beach County’s bots are now docketing about 12,000 filings/week - the work of 22 FTEs. Through attrition, the county has eliminated many entry-level positions, freeing up budget for better-paying, more skilled positions and giving greater emphasis to human interaction in more complex matters.

AI can also help identify previously unobserved correlations more efficiently. For example, one court used AI statistical modeling to analyze all traffic tickets for the last 15 years. The models revealed something unexpected: women were less likely to be cited than men for the same infraction, and when cited, were more likely to be offered a plea deal.<sup>9</sup> Humans must ultimately interpret results and determine the appropriate action to take on AI-identified patterns.

## Common Sense and Ethics

Caution is important in the use of AI. Before deciding which human processes to offload to AI-enabled systems, planners should consider both how the AI system could produce

---

<sup>9</sup> Personal communication, 18 February 2020.

errors, and how consequential those errors could be. The risks associated with AI decision making are not fundamentally different than those of human decision making:

- Both humans and AI make decisions based on the data they are provided.
- Both humans and AI ingest a wide range of data types and sources, continually refining their assumptions and conclusions.
- Based on their experience, both humans and AI develop “biases” (i.e., predictions that appear inequitable) and “intuition” (i.e. biases that cast doubt on a prediction).

Yet AI can use its more scalable computational capacity to arrive at these outcomes more quickly. Human biases and intuition develop far more slowly and irreversibly based on inputs that are far harder to assess.

Everything that happens in the criminal-justice system involves a human in some way, and every time a human is involved, there’s always this potential for bias... We already have black boxes making decisions for us all the time, but they just happen to be sitting in black robes.<sup>10</sup>

Ensuring appropriate, unbiased, ethical use of AI in the public sector is of grave concern. The US Department of Defense (DoD) recently developed and officially adopted Ethical Principles for AI:<sup>11</sup>

1. **Responsible.** The Department’s personnel will exercise appropriate levels of judgment and care, while remaining responsible for the development, deployment, and use of AI capabilities.
2. **Equitable.** The Department will take deliberate steps to minimize unintended bias in AI capabilities.
3. **Traceable.** The Department’s AI capabilities will be developed and deployed such that relevant personnel possess an appropriate understanding of the technology, development processes, and operational methods applicable to AI capabilities, including with transparent and auditable methodologies, data sources, and design procedure and documentation.
4. **Reliable.** The Department’s AI capabilities will have explicit, well-defined uses, and the safety, security, and effectiveness of such capabilities will be subject to testing and assurance within those defined uses across their entire lifecycles.

---

<sup>10</sup> Sharad Goel, as quoted by Derek Thompson in “[Should We Be Afraid of AI in the Criminal-Justice System? Many states and cities are putting Americans’ fates in the hands of algorithms.](#)” *The Atlantic*. 20 June 2019. Web.

<sup>11</sup> See [US Department of Defense Adopts Ethical Principles for Artificial Intelligence](#).



5. **Governable.** The Department will design and engineer AI capabilities to fulfill their intended functions while possessing the ability to detect and avoid unintended consequences, and the ability to disengage or deactivate deployed systems that demonstrate unintended behavior.

These principles are not unique to the DoD. Indeed, they could be applied to any public sector organization including courts.

## Conclusion

*Augmented* Intelligence is helping courts do some things that humans do, only better, faster, and at a lower cost. AI also has the power to improve decision making, reduce bias, and summarize information to make it more timely and accessible to courts and the public they serve.

While significant efficiencies can be gained by moving away from manual or human-in-the-loop processes, it is important to note that it is currently impossible (not to mention undesirable) to completely remove humans from systems. That said, each method of human oversight carries its own challenges. While human-on-the-loop AI systems may feel “safer” because a human could/should catch mistakes, how the human engages with the system can significantly impact the quality of this control. For example, if a human is tasked with passively monitoring a self-driving car as it navigates traffic unassisted, it would be almost impossible for that individual to remain focused enough to notice and respond to a pedestrian jay-walking.<sup>12</sup> Alternatively, if a human has the ability to override system decisions without any check on the human’s biases, repeated overrides could fundamentally alter the AI system’s quality. AI systems must be continually monitored, trained, and optimized – just like human processes – to ensure appropriate outputs.

Courts that are leveraging AI effectively are starting small, building both confidence and expertise. They are offloading repetitive, non-value-add tasks and business processes to AI, freeing people up to do what people do best. And some are using AI to discover and analyze patterns in existing processes, predict future patterns, and develop innovative tools to deliver on aspirational improvements in the justice process.

For more information, contact NCSC at [technology@ncsc.org](mailto:technology@ncsc.org).

---

<sup>12</sup> Lee, Dave. “[Uber Self-Driving Crash 'Mostly Caused by Human Error'](#).” *BBC News*, BBC, 20 Nov. 2019, [www.bbc.com/news/technology-50484172](http://www.bbc.com/news/technology-50484172).