Applied Human Factors and Ergonomics (AHFE2025), Vol. 199, 2025, 1125–1143
https://doi.org/10.54941/ahfe1006925

Making Legal Decisions and Managing Legal Risks in an AI World: A Better Approach

John Celona[1], Olof Heggemann[2], and Adam Seiver[3]

[1]Decision Analysis Associates, LLC, San Carlos, CA 94070, United States
[2]Eperoto AB, Regeringsgatan 29, 111 53 Stockholm, Sweden
[3]Philips Emergency Care, Bothell, WA 98021, United States

ABSTRACT

Advances in AI have propelled to the forefront of legal risk management and decision-making. To understand how to make better use AI, we present, first, a new conceptual model for understanding human cognition, decision-making, and action (including for risk management), and how it can be improved with decision analysis. Second, we review salient developments in AI to understand at a conceptual level how it works, what it is best for, and its limitations and shortcomings. We conducted detailed research on current AI and decision analysis offerings for legal risk management and decision-making and summarize our findings. Lastly, we present a model for better integrating human operators and AI for future decision-making and risk management.

Keywords: Artificial intelligence, AI, Decision analysis, Predictive analytics, Legal risk management, Cognition, Decision making

INTRODUCTION

Advances in AI have propelled it to the forefront of applications efforts in many different fields, including legal decision-making and managing legal risks. Assertions are made that AI tools can used to analyze past case outcomes, judicial rulings, and legal precedents to predict the likely trajectory and outcomes of litigation. In this paper, we will present the results of our search for evidence of actual use, along with an evaluation of the strengths and weaknesses of such use.

However, there are fundamental conceptual and practical limitations to this approach, including the fact that almost all cases are resolved through settlement with a resolution that remains confidential and is not publicly reported. Further, use of AI for legal decisions raises significant risks inherent to the way AI works (e.g., algorithmic bias, lack of transparency, etc.) and ethical, regulatory, and due process concerns. (Villasenor et al., 2019; Loyola, 2018).

A better approach is to use AI for the purposes it is eminently suitable and superior for and to use Decision Analysis for making decisions in risky, complex situations like managing legal risks and making legal decisions. (Howard et al., 2018; McNamee et al., 2007; Celona, 2016). Decision

Analysis was developed over 50 years ago for making better decisions in complex and risky situations and is standard and required in a number of high-stakes, high-risk industries like drug development and oil and gas development. However, its application for legal decision-making and risk management is very limited for reasons discussed in this paper.

We present an approach for understanding and applying AI and Decision Analysis appropriately in the legal and other areas, including a survey of current AI and Decision Analysis analytical tools for legal. This paper will build upon the conceptual foundation described earlier (Celona, 2021; Celona, 2020) and present a practical and theoretically sound approach for integrating AI and Decision Analysis in future technology for decision making and risk management in general and, in particular, in the legal arena.

UNDERSTANDING AND IMPROVING HUMAN DECISION-MAKING AND RISK MANAGEMENT

A lifetime and a library can be devoted to the study of human cognition and still leave unanswered questions. For our purpose here, we focus on cognition and decision-making to frame a discussion of making better legal decisions and better managing legal risks.

Kahnemann and Tversky (Tversky et al., 1974; Kahnemann, 2011) first described the modern concepts of Type 1 (fast intuitive) and Type 2 (slow deliberative) thinking, though its roots go back to Plato's separation of reason and appetite and Aristotle's division of the soul into rational and irrational parts.

Type 1 thinking, also referred to as System 1 or the affective decision system works quickly, easily and intuitively, using reasoning methods that are difficult to trace or parse. It governs emotions, trust, empathy and action, and is specialized for understanding others' behavior and motivations. It works well most of the time, but predictably goes astray with novelty, complexity or uncertainty.

In contrast, Type 2 thinking (likewise similarly referred to as System 2 or the deliberative decision system), is slower and more effort to invoke, but less prone to error. It uses various means of abstract reasoning, including logical reasoning as codified in the rules of logic, and visualization (including 3-dimensional visualization). It uses transparent, logical reasoning and requires concentration and focus—at the risk of frame blindness, or solving the wrong problem.

However, as any student of logical reasoning (including legal reasoning) knows, logic can easily be manipulated. Change or assume the premise or the conditional and any answer can be arrived at. Accordingly, logic is not persuasive nor does it lead to action unless the source is trusted and also intuitively persuasive.

The core of Kahnemann and Tversky's work is describing where the heuristics and biases employed by Type 1 thinking to make decisions diverge from the results of applying Type 2 thinking, producing an answer that logically is "wrong." Kahnemann addresses this issue in his newer work (see Chapter 25: Bernoulli's Errors and following) by concluding that his and

Tversky's prior work is wrong because it differs from the results produced by untrained intuitive decision makers.

Besides heuristics and biases, another major confounding our intuitive decision-making processes is emotions and stress and, perhaps most importantly, fear of bad things (ranging from getting eaten by a tiger to financial loss), and fear of risk of those bad things happening. As was explored in depth as it pertained to the Covid pandemic and responses to it (Celona, 2021), evolution programmed us with a rather limited range of responses to threats (Figure 1). When not faced by an imminent threat, other emotions like greed (or, its more civilized cousin, desire for gain) kick in. Thus the financial markets famously and unpredictably oscillate between fear and greed, despite enormous sums, brainpower, and computing resources being deployed to analyze and predict them. Further, people loathe uncertainty and seek to eliminate by any means possible, such as by constructing explanatory narratives or by controlling what they are able to control with little or no regard to whether their actions actually control the risk in question.

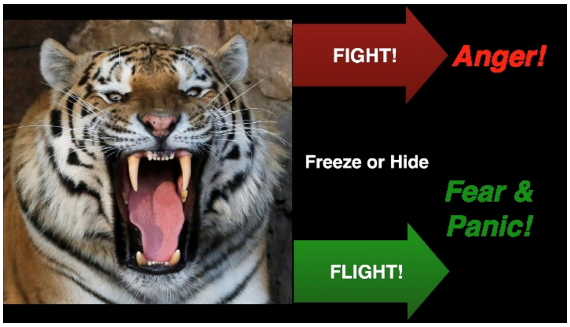

Figure 1: Intuitive behaviors and alternatives for threats.

The point that seems missing from these analyses is that Type 1 and Type 2 thinking are not separate "computing" systems sitting in separate boxes. They are part of one, evolved organism (us). Evolution perhaps provides a perspective on why—like all other parts of our body and thinking "machine," they must have evolved to work together.

A previous book chapter explored the possible impact of evolutionary processes on ethics. (Celona, 2020) Similarly, evolution may provide a solution to the apparent dissonance between intuition and logic. As evolved beings, our cognition has evolved along with the rest of us, and it makes no sense from an evolutionary perspective to evolve capabilities that, rather than working together, diverge from one another or—even worse—disagree. What if your senses of touch and your vision disagreed with each other?

Accordingly, Celona presented a programmable co-processor model of cognition in in which our "intuitive supercomputer" (Type 1 thinking) is in the driver's seat and governs action, but it and our logical co-processor (Type 2) can exchange information and "jobs" like the task of considering a decision as shown in Figure 2.

For most situations, our intuitive supercomputer perceives the situation, instantaneously matches a behavior pattern and acts, resulting in either a helpful action or mistake, as shown in Figure 3. Those mistakes include the heuristics and biases described by Kahnemann and Tversky which result in a decision differing from what logic might indicate.

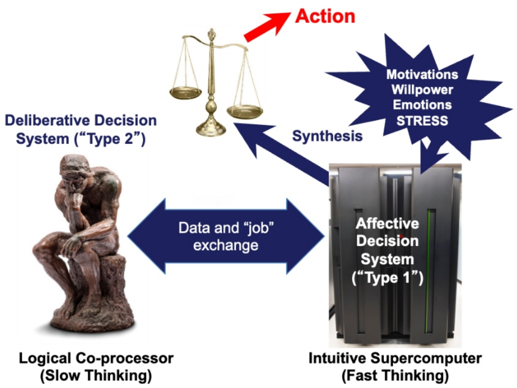

Figure 2: A programmable co-processor model of cognition.

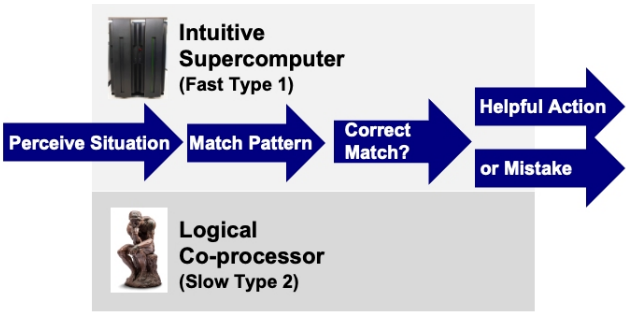

Figure 3: Normal automatic decision and response by our intuitive supercomputer.

However, the two modes of thinking can also communicate and reprogram each other. Intuition can recognize a situation where a new behavior pattern would be helpful, "download" the problem to deliberative thinking, create and practice a new pattern, and by practice "upload" the pattern to intuition so it can be executed quickly and automatically in the future, as illustrated in Figure 4.

Examples of this abound in everyday life. People learning to swim must learn a new breathing pattern (a quick "bite" of air, then continuous exhale through the nose) then, with practice, it becomes automatic and attention can be focused on stroke, body position in the water, etc. Pilots and scuba divers must learn to rely on instruments for depth and bubbles in the water or a horizon gauge for orientation because humans—unlike birds—have no innate sense of which way is up. Likewise, medical students laboriously peruse texts and practice with patients to start, then, over time, can instantly recognize and treat many conditions. It is the same process with progression in any discipline from novice to expert.

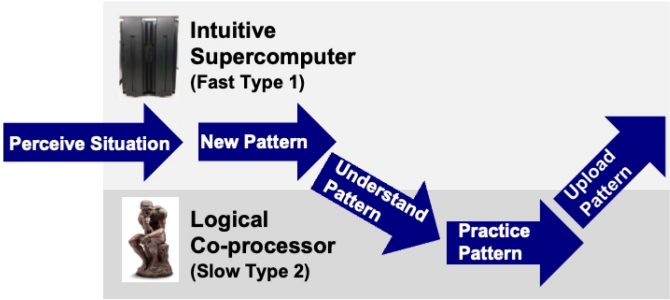

Figure 4: Intuition and logic working together to create a new situation response.

Similarly, with decision analysis, we can reprogram and upgrade our intuitive decision-making to arrive at decisions we both logically know are correct and intuitively can explain and justify. This process is illustrated in Figure 5.

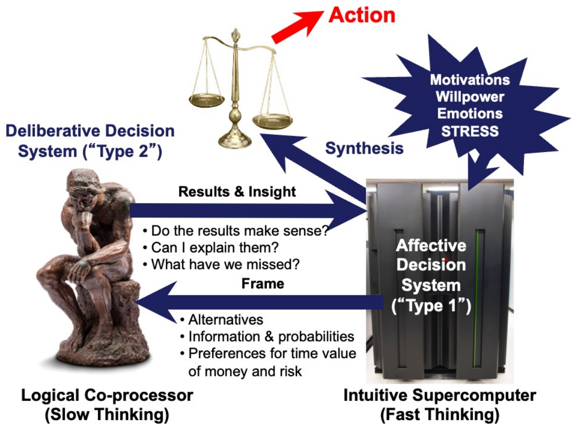

Figure 5: Decision analysis frames the interaction between intuition and logic to make better decisions.

Intuition provides the frame of the problem: which alternatives should be considered? What information do we have and what probabilities should be use? What are the decision-maker(s)' preferences for time value of money and risk?

These inputs go into a logical calculation model (often a decision tree, sometimes also linked to an Excel model for further calculations on scenario values) which produces a variation of outputs: mean values, probability distributions, the preferred alternative in particular scenarios, and so on.

Then, our intuitive supercomputer provides a review and check on the results: do the results makes sense? Can we explain them qualitatively? Is there anything we missed?

Thus, with practice, we learn to make better decisions and no longer employ our intuitive heuristics and biases where we know they would lead us to a course of action we would not logically agree with. Similarly, we learn to think about and deal differently with risk rather than instinctive loathing and irrationally telling ourselves we are controlling it with an explanatory narrative or by controlling something else. We employ logic, but also our intuition to provide a check and guidance on logic's inputs and outputs.

APPLYING AI TO DECISION-MAKING AND RISK MANAGEMENT

To tackle application of AI to decision-making and risk management, we first very briefly review some salient aspects of AI's development and mechanisms. Various sources cover the development of AI with differing focus and level of detail (Bringsjord, 2020). For our purposes, there are two key points.

First, early attempts to develop "artificial intelligence" comparable to human intelligence went through repeated cycles of optimism, generous funding, and bold predictions followed by disappointment and deep funding cuts, leading to the "AI winters" of 1974–1980 and the 1990's. One intractable problem was how to give a machine "commonsense knowledge"—the way our intuitive supercomputer provides a common sense check on deliberative reasoning. This type of intelligence with generalizability ability and common sense knowledge is now termed "Artificial General Intelligence" and remains theoretical only (Google Cloud, 2025).

Current AI efforts are now also termed "Artificial Narrow Intelligence" (ANI), focus on specific tasks, are based on mathematics and logic, and follow the fundamental structure of computer programs since their creation: algorithms + data structures = programs. Where probability is also introduced, it follows formal rules of probability (including Bayesian probability).

Second, the boom beginning about 2005 is based on huge growth in the amount of data available and comparable increases in computing power—not on any theoretical breakthroughs in how to program common sense knowledge or generalizable ability (or the ability to learn). Various sources chronicle the growth in the global data sphere and the progression in the number of transistors on a chip (which is a measure of computing power).

These capabilities are mind-boggling compared with what was possible only a few decades ago. The Apollo 11 moon rocket computer had 4 kb of RAM and could execute 40,000 instructions per second. The first author calculates that the iPhone in his pocket has 8 million times the Apollo computer's RAM and 875 million times its processing power.

For the non-technical, we summarize the salient AI techniques today briefly as follows (and apologies to technical folks who find these too simplistic). Artificial Intelligence is code which analyzes its environment and takes steps to maximize its chances of achieving "success" (as defined in the code). Machine Learning is an application of AI which uses algorithms and statistical models to perform a task by relying on patterns and inference to predict data not present. Deep Learning is a type of machine learning which uses multiple layers to progressively extract features from raw input. Neural networks are a type of machine learning system with nodes like a neural system which exchange data to arrive at conclusions.

These techniques are not new. Neural networks were first described by Walter Pitts and Warren McCulloch in 1943. They were simply not feasible with the amount of data and computing power then available. The current boom began about 2005 as large amounts of data and cheaper, faster computers made machine learning and deep learning techniques feasible for the first time.

The more recent boom began about 2017 with the introduction of Large Language Models (LLM's) such as ChatGPT which are based on "transformer architecture." Transformer architecture is another application of neural networks in which text is converted to a numerical representation which allows very complex calculations to be made. The results are quite astonishing, but they are still just a calculation. Pinar Seyhan Demirdag, co-founder of Cuebric (an AI company that creates AI-generated backgrounds for films), is an AI-optimist, but named her company Cuebric in tribute to Stanley Kubrick, in whose film "2001: A Space Odyssey" Hal the computer goes rogue. As she said in an episode of 60 Minutes: "Just because a machine can do faster calculations, comparisons, and analytical solution creations, that doesn't make it smarter than you. It simply computates faster. In my world, in my belief, smarts have to do with your capacity to love, create, expand, transcend… these are qualities that no machine can ever bear. They are reserved to only humans (60 Minutes, 2024).

Accordingly, if we look conceptually at what how an AI works as compared to human cognition (Figure 2), the result of the calculation is simply displayed or acted upon (depending on the program), without any intuitive check on the result, as shown in Figure 6.

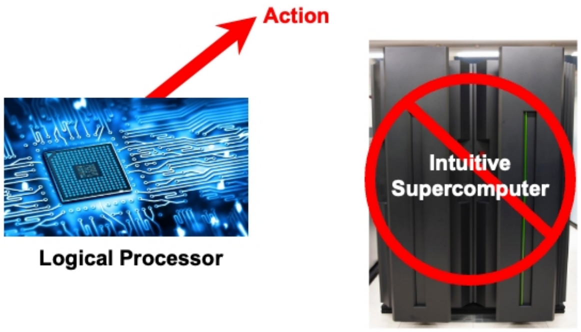

Figure 6: The man-made computer or AI.

The computer or AI will exactly follow the instructions written for it and cannot deviated from the instructions or written objectives, though the results may be quite surprising. As my co-author on one book Vint Cerf put it, "a bug is when the computer does what you told it to do instead of what you want it to do."

AI shines at applying exactly what it is: dizzying calculations on prodigious amounts of data to produce a result, and has been invaluable in preparation of this paper. So it is with risk-management. An AI can definitely review and summarize vast amounts of data and produce a variety of results, including the probabilities on average for a particular instance of data, but it can't tell you how your instance of data (such as your case) might fall relative to the sample or—more importantly—what you can do to change the odds. That requires insight beyond content present in the data, and, with present or currently conceived technology, that can only be supplied by an intuitive human supercomputer. Still, we will review current applications and their claims.

Decision-making is a far more difficult task because it requires conceiving and evaluating alternatives based on incomplete and uncertain present data and no data from the future (which is what we are really interested in!). More critically, what the data from the future will be is contingent on what we do now. An AI tackles decision-making—with often unanticipated results.

The problem is daunting, and may be unsolvable in some domains. The first author listened to one presentation by a researcher working on a U.S. Department of Defense application describing the efforts to successively identify and implement correct solutions for the many "corner cases" where the algorithm chose an incorrect course of action. With solutions implemented for each set, a new and just as large set was revealed. As he put it, "The corner cases are not decreasing."

And how would the AI come up with a new and better alternative, or know what the decision-maker(s) preferences are regarding the risk in various choices? That requires an integration of a decision analysis framework with AI to create a complimentary interface between the decision-maker(s) and the AI. That we'll tackle in the next section following a survey of current offerings.

CURRENT DECISION ANALYSIS AND AI TOOLS FOR LEGAL

For results of our research into current general-purpose and legal-specific decision analysis software and current analytics and generative AI/LLMs for legal, please see Appendices here. We searched for software and evidence with a variety of prompts to a number of search engines and ChatGTP. We will summarize and comment on our findings here.

General purpose decision tree software has been around for decades. It is aimed at users who desire basic functionality of being able to build and run a decision tree and see the answers, rather than more advanced tree structuring or analysis. Larger trees are sometimes handled by running a Monte Carlo simulation although, with current technology, it is often feasible to simply calculate all the scenario values for a decision tree with thousands of endpoints or more. Emphasis is on being able to structure and run larger trees, often with an interface to Excel for more complex models of scenario values.

These applications suffice for many users and are a great way to start applying and using decision analysis, but make it difficult to deal with what decision analysis co-founder Ron Howard described as the one concept he would like people to understand if there was only one: conditional probability. Conditional probabilities are almost always present, and usually critical when they are. A key part of the search for alternatives is finding ones that can shift the odds in your favor. These packages lack a simple way to enter conditional probability and analyze its implications.

Although decision analysis has been routinely covered in analytical methods classes in engineering and business schools for decades, it is not currently taught in law schools and, for this reason, few attorneys are familiar with it. Efforts are under way to introduce attorneys to decision analysis. For example, the Early Dispute Resolution Institute teaches decision analysis as part of its dispute resolution training, and the American Arbitration Institute will soon be offering it.

Similarly, descriptive data analytics has been around for decades, but has been "super-charged" by AI developments and applications. As specified in more detail in the appendices, a number of applications are still purely descriptive (what happened), others are trying to be predictive (what will happen in your case), but are very limited. For example, in its own video on how it works and why it is useful, Lex Machina shows that, out of 149 cases relevant to the example case, only 3 went to trial (LexisNexis Legal, 2025). One does not need to be a statistician to know that "sample" is not very helpful. As has happened repeatedly with AI, the hype for AI capability for legal (Hallene et al., 2024) has gotten far ahead of the actual capabilities. The results of our searches are detailed in the Appendices but, to summarize here, we found little to no evidence of use of AI as described.

Like using traditional statistics for prediction, AI works well when (1) you have "good" data; (2) a stable system producing consistent data; and (3) you are predicting large numbers of outcomes. It also helps when you have many opportunities to tune the algorithm (as in online marketing and social media targeting). Decision analysis still seems the best for prescribing what you should do with your case. AI prescriptive and unassisted decision-making, at the current time, remains quite fraught. Please see the detailed results here for more details.

TOWARDS AI HELPING PEOPLE MORE EFFECTIVELY

Given our understanding of the human and artificial intelligences (Figures 2 and 6), and the issues with AI making sole decisions in complex or uncertain situations, the question is how to integrate the two effectively to make better decisions than either the unaided person or unaided computer. Given AI's abilities, we believe their is great potential in a thoughtfully structured interaction.

The solutions is perhaps surprisingly simple: use people for what they are really good at and AI for what it is really good at. Figure 5 provides the road map.

People need to specify the framework for the decision: allowable alternatives, information (including Bayesian probabilities), and preferences for risk and time value of money. These constrain the possible actions of the AI. Then apply the AI to processing all the necessary data and choosing alternatives. The preferences describe how the designers wish the system to value the various alternatives, with the system choosing the highest or lowest value alternative as specified. This valuation can include probability weighted values. If outcome valuation is nonlinear, a utility function can capture the nonlinearity and maximizing mean utility can replace mean value as a decision criteria.

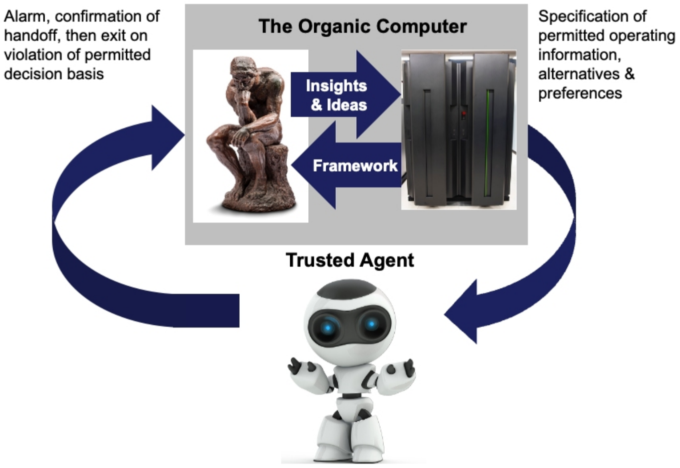

Figure 7: A possibly better way to integrate organic computers with "trusted agents".

The system can operate autonomously so long as it stays within the allowable decision basis. On any condition violating the permitting decision (action) basis, the systems needs send an alarm to the operator, confirm handoff of operational control, then exit operational control. This process is illustrated in Figure 7.

This would be a departure from current practices in several ways. First, the operating range of the AI would be clearly defined. Second, human operators are free to attend to other things so long as the AI is acting within its parameters. No need to simply watch the system all the time, which is a recipe for attention to wander and disaster. Third, this is far more specific way to target and constrain the AI's actions rather than current practice of specifying an objective function, which the optimization algorithm can find surprising and unintended ways of maximizing (like a self-driving car accelerating into an obstacle).

Fourth, upon detection of violation of permissible conditions of operation, the system would alarm and confirm human engagement before exiting control. The alert would include information on the alarming condition and possibly the next decision required. In other words, "when something goes wrong, as it will sooner or later, autonomous systems must allow other machine or human teammates to intervene, correct, or terminate actions in a timely and appropriate manner, ensuring directability" [emphasis in original] (Defense 2016 on p. 15). The exit conditions would need to be specifically defined for the system, and would depend on what the system is tasked at, what the components in the system are, and the magnitude of consequences from a potential mistake.

Exit conditions could include an AI predictive model calculating a confidence score (usually a percentage) on a prediction and triggering an action if the score is below a set threshold. Actions could be varied depending on how much below the threshold the score is. Another exit condition could be detecting changing data possibly resulting from sensor drift beyond a specified range (i.e., indicating the sensors need to be recalibrated), or that operating results are outside specified data quality metrics—or some combination of these. Or it could be that certain alternatives defined in the decision basis are set to require operator review before actually implementing them. Another possibility is the system identifying that it is progressing down a path in the tree rarely expected to be encountered, and therefore requiring operator review before proceeding down the path.

Likely a graded handoff protocol would be needed, with grading on both the alarm the system sends and on the required operator response. System alert levels could be (1) Concern: review for one or successive results below confidence thresholds, data drift, etc. (2) Alarm: system is operating repeatedly at the fringes of or beyond expected operating range. Serious malfunction possible. And (3) Emergency: system operating in prohibited operating range. Implementing emergency protocols. Immediate operative intervention required.

The operator responses could be at three levels: (1) instruct the system to continue operating as designed; create and save variance reports for analysis (for this response and the following responses). (2) Continue to operate, but transfer decision control on certain actions to the operator until further notice. Or, (3) implement minimal system maintenance procedures and go into standby mode; no further system initiated actions. Operator controls all actions during this procedure.

For each of these responses, the operator would need to acknowledge handoff of the specified level of control and continue provide needed inputs until the system is instructed otherwise.

Subsequently, the followed conditions would need to be established and verified to restore full system operation and designed level of autonomy: (1) all system predictions are within confidence thresholds. Do quality metrics need to be reset? (2) Analysis of variance reports and reexamination of decision basis indicates the basis is still appropriate for continued operation, or modify the decision basis as needed. (3) Decision actuation mechanisms still operating within tolerances.

Depending on the level of suspended system function, three actions might be needed to restore full system function: (A) restore data gathering, analysis, and predictions; (B) restore alternative evaluation and display recommended course of action, and (C) restore full designed level of autonomous control.

Humans are far better than robots dealing with unanticipated situations and figuring out what to do. The trick is to harness that ability in concert with the robot and let the robot handle the monotony (which they are impervious to!) of routine decisions and actions, and to alert as to the required level of human intervention, review, control, and remediation when the system is operating outside of its designed operating parameters.

CONCLUSION

Human risk management and decision-making has well-studied issues (heuristics and biases, emotional response to risks, etc.), but can be improved by making better use of our built in and programmable co-processor modes of cognition: the intuitive supercomputer (Type 1) and the deliberative, logical co-processor (Type 2). Decision analysis provides the framework for effective interaction between these two modes of thinking.

AI has made tremendous strides over the few decades based on enormous growth in available data and processing power, but remains based on mathematics and logic developed over the decades, some back in the 1940's and 1950's. A more recent development is Large Language Models (LLM's), which convert text into a mathematical representation (a "token"), upon which very sophisticated calculations can be made to generate a response. AI remains limited by its core functionality: mathematics, logic, and calculation only. Capabilities of commonsense knowledge and generalizable ability remain theoretical only, though AI researchers have sought ways to create these since the field's inception.

AI is incredibly valuable when applied for its core abilities to analyze and summarize data, and was invaluable in the preparation of this paper. It is superior to human oversight when applied to well-understood, routine decisions where human attention tires and wanders.

There are many applications offering decision analysis and AI functionality for legal risk management and decision-making. See the Appendices here for further detail. They vary in their technology offerings and whether they offer descriptive, predictive, or prescriptive results. AI offerings are highly suspect when they offer predictions of what will happen in your case for reasons discussed above and, as has been the case repeatedly in the development of AI, the hype far exceeds the actual capabilities. Decision analysis remains the only methodology to offer solid prescriptive results for risk management and decision-making.

AI and decision analysis can be effectively integrated for better decision-making in complex or uncertain situations, and a model is presented for doing so.

ACKNOWLEDGMENT

The authors would like to acknowledge and thank the respective "angel" or "angels" in our lives who continue to inspire, motivate, and energize us to tackle projects like writing this paper.

APPENDIX A: GENERAL PURPOSE DECISION ANALYSIS SOFTWARE

Product, Website & Company	Functionalities	Comments
DPL Professional www.syncopation.com	Influence diagrams, sensitivity analysis, decision trees and more, delivered as desktop software for Windows.	A series of desktop software applications around decision analysis. Originally developed by senior staff of Price Waterhouse Cooper and then expanded.

Product, Website & Company	Functionalities	Comments
PrecisionTree and @RISK www.lumivero.co m Lumivero (USA)	Influence diagrams, sensitivity analysis, decision trees, Monte Carlo simulations and more. Both software integrate with each other and run entirely and within Excel for Windows desktop computers.	Originally developed and sold by Palisade, now owned and distributed by Lumivero. Lumivero was formed in 2022 through the merger and tool consolidation of Palisade, Cambridge Econometrics, and others.
TreeAge Pro www.treeage.com TreeAge Software LLC (USA)	Influence diagrams, sensitivity analysis, Decision trees, Monte Carlo calculations and more provided as a desktop application for PC and Mac.	TreeAge Software LLC is a US-based company that develops TreeAge Pro. Used in healthcare for cost- effectiveness and clinical modeling and includes legal materials and models.
Silver Decisions silverdecisions.pl Originated from Warsaw School of Economics (Poland)	Free, web-based, open-source decision tree software supporting probability assignments, payoff calculations, sensitivity analysis, and user-defined expressions. Designed for educational applied decision analysis, it runs directly in browser without installation.	Developed as an educational and practical tool, especially for teaching and simple decision analysis.
TreePlan Decision Analysis Toolkit treeplan.com TreePlan Software (USA)	A collection of Excel add-ins for creating decision trees within Excel. Allows users to build, edit, and evaluate decision trees with probabilities, sensitivity analysis, payoffs, etc. in Excel for Windows.	Developed by Professor Michael R. Middleton, USF, and is used in business schools and industry training to teach decision analysis and spreadsheet modelling.
Analytica llumina.com/ Lumina Decision Systems (USA)	General-purpose decision analysis and modelling software. Supports influence diagrams, Monte Carlo simulation, optimization, sensitivity analysis, and more.	Used across industries including energy, environment, healthcare, and policy.

APPENDIX B: LEGAL SPECIFIC DECISION ANALYSIS SOFTWARE

Product, Website & Company	Functionalities	Comments
Eperoto www.eperoto.com Eperoto AB (Sweden)	Decision analysis software tailored for litigation and arbitration. Allows lawyers to build case scenarios with claims, probabilities, costs, and risks; includes decision trees, simplified sensitivity analysis, and settlement range assessments.	Founded in Sweden, developed specifically for the legal sector. Used by law firms, insurers, and in-house counsel to evaluate litigation risk and strategy. Combines classic decision analysis with features adapted to legal practice.
Settleindex www.settleindex. com Settleindex LTD	Online platform for case risk analysis and settlement strategy. Allows lawyers to model case outcomes, probabilities, costs and similar,	UK-based tool developed with a focus on litigation settlement strategy.
Mitigaze www.mitigaze.co m Mitigaze	Decision analysis tools with functionalities based around decision trees, sensitivity analysis. risk tolerance and visualizations.	Mitigaze (Previously Litigaze) is an internationally operating software company based in Melbourne, Australia.

APPENDIX C: LEGAL ANALYTICS FOR LITIGATION

Product, Website & Company	Functionalities	Comments
Pre/Dicta www.pre- dicta.com Pre/Dicta Inc (USA)	Predictive analytics platform covering millions of U.S. civil cases. According to the company's website, its product uses judge- specific data and hundreds of variables to forecast rulings (e.g. motions to dismiss), timelines, and comparative outcomes across judges and venues. The outcome analysis for the case assessed seems to be based on structured inputs of key parameters (court, judge, law firm acting as counsel etc.).	Pre/Dicta describes itself as a predictive analytics tool focused on judicial behavior. Its website highlights accuracy in forecasting rulings and timelines. In 2024, the company acquired Gavelytics, a state court analytics provider.
Lex Machina www.lexisnexis.c om/en-us/ products/lex-machina.page LexisNexis (USA)	Legal analytics platform that mines U.S. court dockets and filings. Provides data on case outcomes, damages, timing, remedies, judicial behavior, motion outcomes, and also analytics on parties, law firms, and attorneys. Does not appear to predict outcomes but provides a database that can be filtered to understand the outcomes of similar cases. See also Lexis Nexis Context.	Lex Machina began as a Stanford Law School project and was acquired by LexisNexis in 2015. It is possibly the leading legal analytics platform in the U.S., with coverage of federal courts, and is frequently referenced in legal media and scholarship as a primary example of applied legal analytics.
Bloomberg Law www.bloombergl aw.com Bloomberg Industry Group USA	According to its website, the solutions offers comprehensive legal research platform with integrated analytics. Provides statistics on judges, courts, law firms, and attorneys, along with docket insights and case timelines. Appears to offer structured data to support case assessment without generating explicit predictive outcome estimates.	Bloomberg Law is a full- service legal research and intelligence platform, with analytics included as part of a broader suite of tools. It combines research, news, and analytics in one system. It is not focused on outcome prediction
Trellis Law www.trellis.law Trellis Research, Inc (USA)	According to its website, Trellis provides analytics and insights from U.S. state trial court records. It offers searchable access to dockets and rulings, judicial biographies, motion outcomes, and judge-specific statistics to support case assessment and strategy. Does not appear to provide prediction of outcomes.	Trellis focuses on U.S. state trial courts, a level of the judiciary often less accessible than federal courts. Legal media describe it as useful for motion practice and judge research.

Product, Website & Company	Functionalities	Comments
Blue J www.bluej.com Blue J Legal Inc (Canada)	Earlier versions of Blue J:s website emphasized predictive analytics and outcome probabilities for tax cases. Currently the website states that Blue J provides AI-assisted legal research and drafting tools for tax and employment law by using generative AI trained on curated legal sources to deliver verifiable answers with inline citations and links to primary materials.	Blue J originated from research at the University of Toronto Faculty of Law and was developed into a commercial platform by Blue J Legal Inc. It is primarily used in Canada and the United States, with its focus on tax law.
Predictice www.predictice.c om Predictice SAS (France)	According to its website, Predictice provides litigation analytics for the French market. It offers statistics on case outcomes, average damages, time to judgment, and judge- specific trends, using a large database of French court decisions to support litigation risk assessment and predict outcomes.	French company focused on litigation analytics. According to its website and other legal news sites has been adopted by French law firms and insurers, particularly for estimating damages and case duration. Predictice has recently announced it will combine with the French company Doctrine.
Solomonic www.solomonic.c o.uk Solomonic Ltd (United Kingdom)	According to its website Solomonic provides litigation analytics for the English courts. It offers data on case outcomes, judge behavior, claim types, and durations, drawing on court filings and judgments to support litigation strategy, risk assessments and outcome prediction.	Solomonic is based in the UK and focuses on English court data. It reports that it has been adopted by law firms and litigation funders.
Case Law Analytics www.caselawanal ytics.com Case Law Analytics SAS (France)	According to its website, Case Law Analytics applies probabilistic modelling to assess litigation risk. It provides estimates of case outcomes, damages, and timelines, based on statistical analysis of court decisions. The platform appears to be used mainly in France.	Case Law Analytics was founded in 2017 in France, focusing on probabilistic modelling of litigation outcomes. In 2023, the company was acquired by LexisNexis, which announced plans to integrate its risk assessment and AI tools into the broader LexisNexis analytics offerings
Fastcase/Docket Alarm Acquired by vLex (below)	Legal research platform (Fastcase) combined with docket search and analytics (Docket Alarm). Provides access to federal and state court dockets, case law, and litigation analytics such as motion grant rates, timelines, and party/attorney profiles.Does not appear to predict outcomes but provides a database that can be filtered to understand the outcomes of similar cases.	Fastcase and Docket Alarm are established tools within U.S. legal research and analytics. They appear to be mainly used by practitioners for monitoring cases and analyzing litigation activity and are now part of vLex.

Product, Website & Company	Functionalities	Comments
Westlaw Edge âŁ“ Litigation Analytics legal.thomsonreuters.com/en/ products/westlaw-edge/litigation- analytics Thomson Reuters (USA)	Litigation analytics module within Westlaw Edge. Provides data and statistics on judges, courts, law firms, attorneys, motion outcomes, damages, and case timelines across U.S. federal and state courts. Integrates with the broader Westlaw research environment.	Developed by Thomson Reuters as part of its Westlaw suite. The tool focuses on descriptive analytics to support litigation strategy and research, rather than generating predictive outcome estimates.
vLex (Vincent AI & Litigation Analytics) vlex.com vLex (global, headquartered in Spain/USA) (Acquired by Clio 2025) See also Fastcase/Docket Alarm	According to its website, vLex provides global legal research and analytics through its Vincent AI platform, which combines case law, legislation, and secondary sources across multiple jurisdictions. It includes legal research assistance, citation linking, document summarization, and, in some regions, litigation analytics derived from court data.	vLex positions itself as an integrated research and analytics platform following its acquisition of Fastcase and Docket Alarm. While it incorporates AI- driven document analysis and search, it focuses on information retrieval and descriptive analytics rather than predictive outcome modelling.

APPENDIX D: GENERATIVE AI/LLM SOLUTIONS AND OTHER SOLUTIONS

Product, Website & Company	Functionalities	Comments
Lexis Nexis Context lexisnexis.com/ en-us/products/ context.page Lexis Nexis	Lexis Nexis Context appear to draw its main information from the Lexis Nexis /Lex Machina databases but appear to provide a different type of solution. According to website, the product uses language analysis of past opinions to provide insights into how judges rule and how experts have been relied upon by the courts.	Developed as a complementary product to Lex Machina. Provides descriptive, judge- and expert-focused analytics. Positioned as part of the Lexis+ analytics suite.
Harvey www.harvey.ai Counsel AI	According to its website, Harvey provides legal research with cited answers; contract analysis and drafting assistance; document review and due diligence; workflow automation through a custom workflow builder; and AI agents that can perform legal tasks across documents and queries.	There is no clear information on whether Harvey's system aims to predict case outcomes. Large, domain-specific language models like Harvey can, however, provide structured feedback on arguments and evidence and indicate whether particular claims appear consistent or logically supported within a given legal context.

Product, Website & Company	Functionalities	Comments
Legora legora.com Legora AB.	According to its website, Legora offers AI-based tools for legal research, document review, and drafting assistance. The platform appears to focus on improving efficiency in legal work through automation and text analysis rather than predictive analytics.	There is no clear information on whether Legora's system seeks to predict case outcomes. Tools of this type generally focus on analyzing legal texts and arguments, offering feedback on consistency, reasoning, and document quality, and will likely also assess whether particular claims appear consistent or logically supported within a given legal context.
ChatGPT chatgpt.com OpenAI	According to itself, when prompted, Chat GPT provides generative-AI capabilities for text generation, reasoning, and dialogue. It can analyze and summaries documents, answer complex questions, assist in drafting and research, and perform structured reasoning tasks such as scenario analysis or probabilistic modelling when prompted.	According to itself, Chat GPT is a general-purpose LLM by OpenAI. Provides qualitative assessments of argument strength and helps structure probabilistic scenarios based on the materials provided by the user. It does not access proprietary court datasets or produce statistically verified outcome predictions; outputs are non-authoritative and require legal review.

APPENDIX E: DECISION ANALYSIS AND LEGAL AI SOFTWARE

This appendix summarizes key Decision Analysis (DA) and Legal AI software systems across three analytical dimensions: core modeling capabilities, sectoral orientation, and overall design philosophy. It is intended as a reference appendix or comparative overview for research on AI-assisted decision analysis and legal risk management tools. Categorizations are difficult to make because creative users can employ the tools in various ways, so we've tried to focus on the main or intended use.

Table 1: Summary matrix by sector and predictive orientation.

Category	Primary Sector(s)	Main Focus	Analytics Type	Representative Tools
Classical Decision Analysis	Engineering, healthcare, policy	Quantitative decision modeling	Prescriptive	DPL, TreeAge, Analytica, PrecisionTree
Legal Decision Analysis	Law, insurance, arbitration	Case scenario analysis	Prescriptive	Eperoto, Mitigaze, SettleIndex
Legal Analytics (Data Mining) & Legal Research	Litigation, courts, law firms	Historical data, trends, judge behavior	Descriptive	Lex Machina, vLex, Westlaw Edge, Bloomberg Law, Trellis
Predictive Legal Analytics	Litigation risk, judicial outcomes	Statistical or ML-based forecasts	Predictive	Pre/Dicta, Predictice, Case Law Analytics
Hybrid Research & Analytics Platforms	Multi-jurisdictional legal research	Information retrieval + analytics	Descriptive / AI-assisted	Lexis Context

Continued

Table 1: Continued

Category	Primary Sector(s)	Main Focus	Analytics Type	Representative Tools
Generative Legal AI (LLMs)	Legal research, drafting, document analysis	Text generation, reasoning, workflow automation	AI-driven qualitative	Harvey, Blue J, ChatGPT, Legora

Table 2: Tiered grouping by capability and design philosophy.

Tier	Designation	Description	Examples
Tier 1 – Classical DA Engines	Quantitative, model-based systems	Built around explicit probability and utility models. Used in engineering, economics, and healthcare.	DPL, PrecisionTree, TreeAge, Analytica
Tier 2 – Legal DA Hybrids	Legal case adaptation of DA	Extend classical DA with legal-specific variables (claims, damages, settlement).	Eperoto, SettleIndex, Mitigaze
Tier 3 – Legal Analytics Systems	Empirical, data-mining platforms	Use large-scale case data to identify patterns and statistics, often descriptive.	Lex Machina, Bloomberg Law, Trellis
Tier 4 – Predictive Legal AI	Machine- learning litigation predictors	Apply ML/statistics to forecast outcomes and judge behavior.	Pre/Dicta, Case Law Analytics, Predictice
Tier 5 – Generative / Reasoning AI	LLM-driven systems for text and logic tasks	Perform synthesis, summarization, and scenario reasoning, not explicit modeling.	Blue J, Harvey, Legora, ChatGPT
Tier 6 – Legacy / Niche Systems	Low activity or uncertain functionality	Inactive or not predictive.	Premonition, VizLegal, Theo AI

REFERENCES

60 Minutes: Nvidia CEO Jensen Huang and the 2 trillion dollar company powering today's AI at 12:35. 60 Minutes, April 29, 2024.

Bringsjord, S., and Govindarajulu, N. S.: Artificial Intelligence. The Stanford Encyclopedia of Philosophy. (2020). https://plato.stanford.edu/archives/sum2020/entries/artificial-intelligence/.

Celona, J.: Artificial Intelligence, Data & Decisions in the Pandemic. Accepted for publication at The 2021 World Congress in Computer Science, Computer Engineering & Applied Computing (CSCE 2021). Paper ID# ICA6179. Published? See https://www.american-cse.org/csce2021/publisher.

Celona J. N. (2021) Decisions, Data, Panic and Probability in the Pandemic: Finding Humanity's Best Path Forward. In: Arai K. (eds) Advances in Information and Communication. FICC 2021. Advances in Intelligent Systems and Computing, vol. 1364. Springer, Cham. https://doi.org/10.1007/978–3-030–73103-8_6

Celona, J. N.: A Programmable Coprocessor Model of Human Cognition: Implications for Human Factors, Systems Interaction, and Healthcare. I. L. Nunes (Ed.): AHFE 2020, AISC 1207, pp. 100–107, 2020.

Celona, J. N.: Evolutionary Ethics: A Potentially Helpful Framework in Engineering a Better Society. Chapter 23 in: Abbas, A. E. (ed.): Next Generation Ethics: Engineering a Better Society. Cambridge University Press, Cambridge, United Kingdom (2020).

Celona, J. N.: Winning at Litigation through Decision Analysis. Springer Nature, Springer International Publishing AG, Switzerland (2016).

Defense Science Board, "Summer Study on Autonomy", p. 14. Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, US Department of Defense, Washington, D. C. (2016).

Google Cloud: What is artificial general intelligence? https://cloud.google.com/discover/what-is-artificial-general-intelligence Retrieved on October 9, 2025.

Hallene, A. and Allen, J: Use AI for Predictive Analytics in Litigation. The American Bar Association: Voice of Experience (October 16, 2024).

Howard, R. A. and Abbas, A. E.: Foundations of Decision Analysis. Pearson Education Ltd., London (2016).

Kahneman, D.: Thinking: Fast and Slow. Farrar, Straus and Giroux. New York, NY (2011).

LexisNexis Legal: What Lex Machina Can Do For You. Youtube: https://www.youtube.com/watch?v=uWsyB-jBB8s.

Loyola University of Chicago Center for Digital Ethics and Policy: Sentence by the Numbers: The Scary Truth Behind Risk Assessment Algorithms (May 7, 2018). https://www.luc.edu/digitalethics/researchinitiatives/essays/archive/2018/sentencebynumbersthescarytruthbehindriskassessmentalgorithms/

McNamee, P., and Celona, J. N.: Decision Analysis for the Professional. 4th ed. SmartOrg, Inc., Menlo Park, California (2007).

Tversky, A., and Kahneman, D.: Judgement Under Uncertainty: Heuristics and Biases. Science, September 27, 1974. Vol. 185, pp. 1124–1131. American Association for the Advancement of Science (1974).

Villasenor, J., and Foggo, V.: Algorithms and Sentencing: What does due process Require? The Brookings Institution (March 21, 2019. https://www.brookings.edu/articles/algorithms-and-sentencing-what-does-due-process-require/