Presentation Notes
AI & Labor-Management Arbitration

Daniel Rainey
daniel@danielrainey.us
http://danielrainey.us/
Holistic Solutions, Inc.
National Center for Technology and Dispute Resolution
International Council for Online Dispute Resolution

Monday, November 6, 2023
9:30 a.m.
NAA Labor Arbitration Conference

———

First of all, thanks for the invitation to come and speak with you today. I apologize for having to do this online - I have an accelerating issue with my knees (soon to be replaced) that gets in the way of minor life functions, like standing and walking, and I am just coming off of a bout with COVID..

At any rate, I'll just take a few minutes to go over some of the technology developments that I see on the horizon that may impact your work as labor advocates, management advocates, and arbitrators.

When we got together to organize this panel, we agreed that I would start things off by briefly focusing on Artificial Intelligence and giving an overview of what may be coming in the near future as regards AI on the labor side, the management side, and the practice of arbitration.
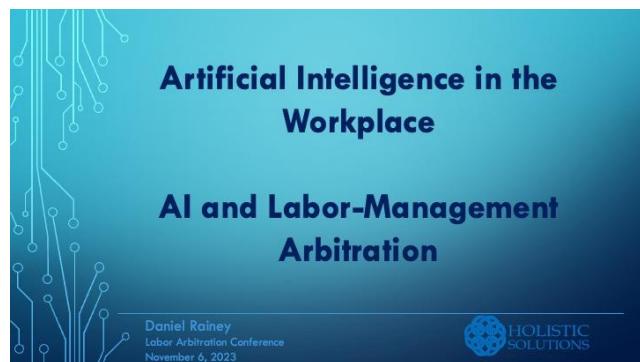
But before I get into that, I noticed a coincidental issue that I thought was interesting given the topic at hand and the timing today. In 1658, James Ussher, the Archbishop of Armagh, Primate of all Ireland, and Vice Chancellor of Trinity College, Dublin, published in "The Annals of the World" his calculation regarding the date on which Adam and Eve were expelled from the Garden of Eden. By his calculation, it happened on November 10, in 4004 BC. One of his colleagues, John Lightfoot, also a Vice Chancellor at Trinity, did his own calculation and determined that the event occurred on a Monday, at 9:00 a.m. So here we are, on a Monday, just before November 10, and just after 9:00 a.m. on a Monday, talking about what will be a momentous movement from "normal" (BAI - before AI) to the wider world that AI will offer. I'm not suggesting that pre-pandemic normal was a Garden of Eden, but I am suggesting that things are going to look very different "out here" as we leave the old normal.
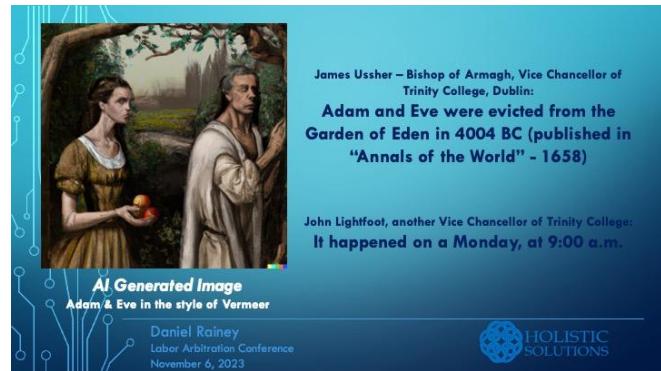
I'll start with a brief note about AI generally - What do we mean by "Artificial Intelligence?"

Development of and use of various algorithmic approaches to dispute resolution that have been wrongly called artificial intelligence is not at all new. Back in 2001, Ethan Katsh and Janet Ripkin coined the term "Fourth Party" to refer to the impact of online technology on dispute resolution generally, and Colin Rule, who has been called the "Godfather of ODR," created the algorithmic poster child for commercial ODR a couple of decades ago at eBay/PayPal. That model has been adapted widely in the access-to-justice space. That "AI" is basically an algorithm that takes one up a decision tree in a very bounded universe, starting with the few things that can go wrong (it didn't arrive, it wasn't what you said it was, etc.), adding the few things that can be done about the problem (give a refund, send another, etc.), thereby resolving disputes by the millions with no human intervention.
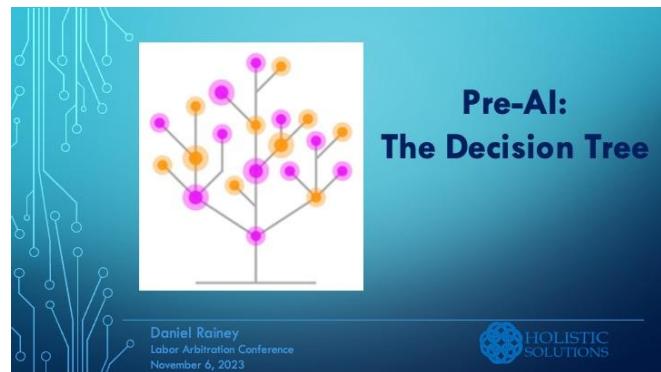
The ability to develop and use the massive Generative Pre-Trained Transformers and the Large Language Models that burst on the public scene with the roll out of ChatGPT *is* new. This AI goes beyond algorithms and trains the application on large data sets so that it can respond in a more nuanced way, using a much larger, if not unbounded, universe of data.
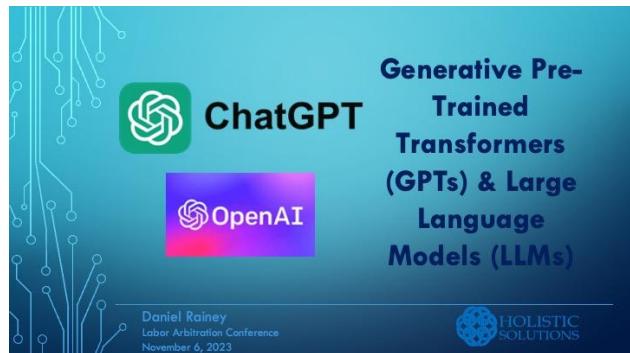
The reactions to this advanced and advancing form of AI have vacillated between outright ecstasy and doomsday anxiety, with a good dose of "gee whiz" and "hmm, how can I use this?" thrown in.

Generative Pre-Trained AI applications (like ChatGPT) and the Large Language Models that are out there now are an impressive leap forward from text searchable databases. There are ten major AI platforms that Stanford researchers included in a review of AI that I'll mention later - and there are more on the way. The bottom line for purposes of this presentation is that AI programs are dependent on the data set from which they learn. For labor, management, and arbitrators, what difference does this make?

There are a range of possible uses for AI on the labor side, but not surprisingly the one that jumps out immediately is, as a long time labor advocate told me, "AI as an extension of electronic research." If I can be allowed a personal example, when I was at the National Mediation Board, Ken Megill and I created what we called "The Knowledge Store." Basically, it is a word-searchable database of railroad arbitration decisions going back decades. The significance of the Knowledge Store is that before we created that database, management advocates had a smaller database of cases they could use for research, but which they did not share with the labor side. The Knowledge Store was open to everyone, making it possible to pull up awards by subject matter, arbitrator name, carrier, etc., so labor and management advocates (and arbitrators) could see what kind of arguments had been made in the past, which one's worked, and how various arbitrators seemed to rule.
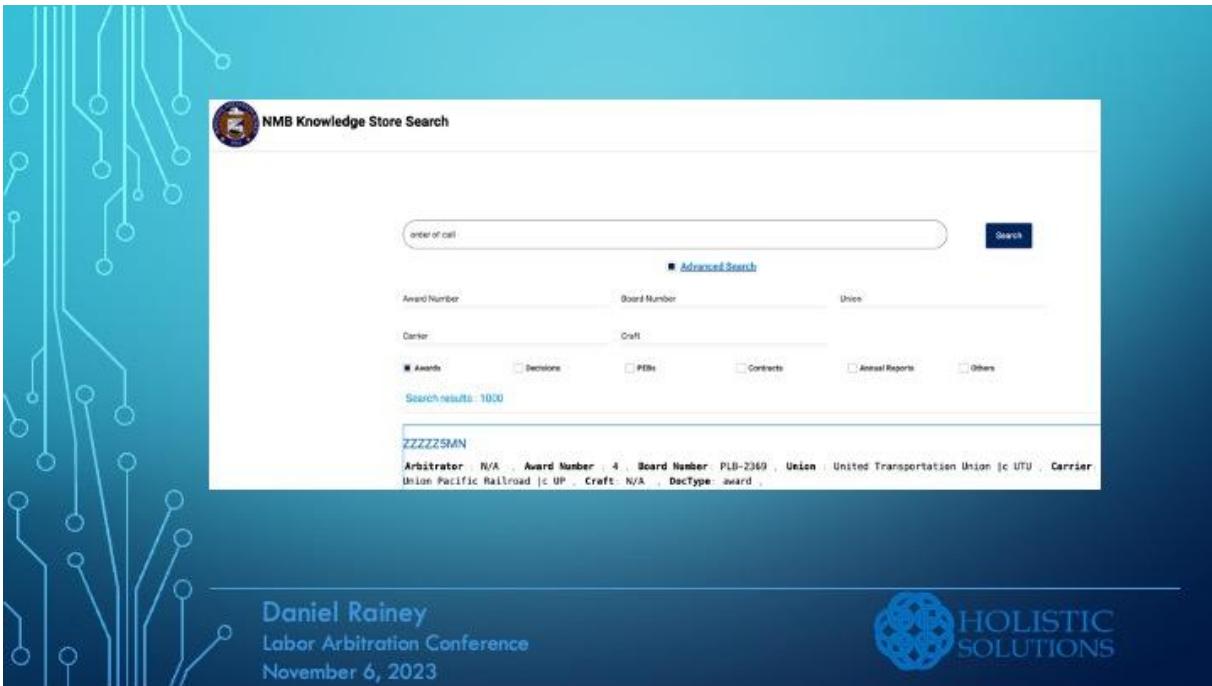
That's a fairly simple use of technology to allow research in a very specific venue. In this example, the word search on "order of call" pulled up about 1000 cases that had that focus, allowing the researcher to see who was involved, who arbitrated the case, and the details of the award. It's very helpful, but it dumps a lot of data on the user, who has to figure it out him or herself.

If I were at the NMB now, I'd be working on integrating AI into the Knowledge Store, training the AI on the decades of data from past railroad arbitration awards, so that it could respond to more abstract prompts, and be able to pull up arguments that might not fit in a tight word-search category but which might have relevance to a particular case. I could, for example, ask the AI to write a draft submission based on the arguments that have worked in the past on order of call cases, or any other case that might be relevant, with the particular arbitrator who was going to hear the case.

For the large but still bounded universe that is railroad arbitration, this approach should work well if the data set is drawn from that universe - data sets from other venues would be less helpful, or perhaps even harmful.

Leaving aside my hypothetical about the Knowledge Store and AI, we now have the availability of commercially offered AI platforms with general data sets used for training. For one example, have a look at AtLaw, an online arbitration preparation platform offered by a law firm and currently available internationally. The Guangzhou Arbitration Commission in China has created an "assistant" that is used directly in arbitration cases. As China Daily reports, "At the close of the hearing, the AI assistant wrote to the disputants: "Today's hearing has come to an end. I am currently analyzing the trial data and the ruling opinion will be sent to the arbitration tribunal via email in 5 minutes." In addition to drafting an award, the Commission argues that, "the assistant can improve the efficiency of resolving disputes by nearly four times by completing procedural tasks including intelligent acceptance of cases, multilingual realtime translation, blockchain recognition of evidence and inputting viewpoints and statements . . . ."
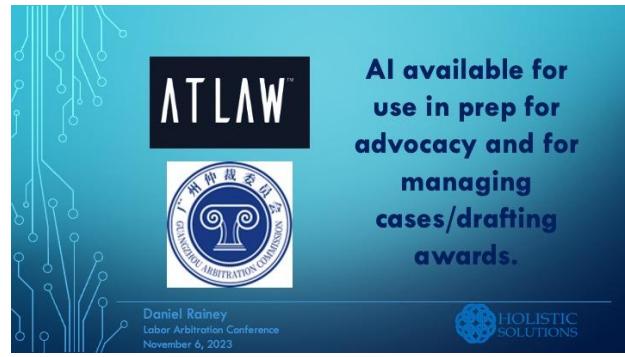
As I mentioned above, Stanford University just published a review of major AI platforms with a report card that reflects thirteen factors related to transparency of the platforms, including transparency of the data set on which they are trained. The *best* score was in the low 50's out of 100 - with most being very reluctant to disclose, or refusing to disclose, the training data set. The full review of AI platforms can be found at https://crfm.stanford.edu/fmti/
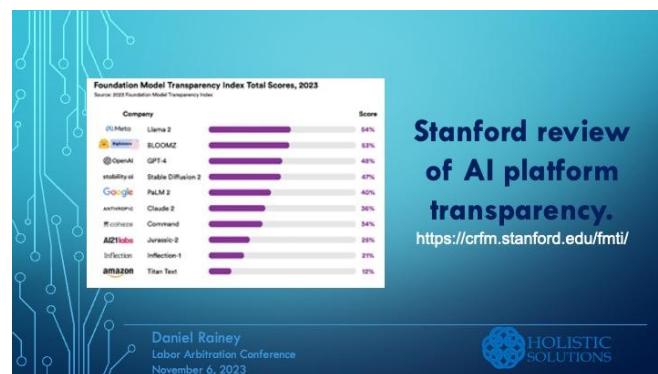
For the moment what this means is that currently available AI should be used carefully for research and for writing drafts of arguments in specialized venues.

That will change quickly as time goes by. Organizations like WestLaw and LexisNexis that have offered searchable databases for a fee have either added or will quickly add AI capabilities to their existing services, and organizations like AAA, with a century of awards and data, will get

into the game as quickly as possible. Ultimately, AI service providers will focus narrowly on practice areas and make their output very useful, and some providers outside of China will, undoubtedly, move in the direction of robo-arbitrators.

In addition to research and drafting arguments, other uses of AI for the labor side might include data analysis, member engagement and communication, and preparation for developing negotiation positions - and uses we can't predict that will arise from integration of AI into regular use.

On the management side, the same research assistance and drafting assistance is probably the first impact of AI that arbitrators will see. In addition, management will be able to use AI for things like predictive analysis, negotiations, workforce planning, employee engagement, performance monitoring, legal compliance, etc.

The worry that labor advocates have expressed to me is that use of AI to deepen the dependence on data analysis and statistics will erode the "just cause" standard for judging management behavior - basically encouraging mechanistically applied policies that streamline consideration and erode the practice of individual consideration. That's probably a reasonable fear.

For arbitrators, depending on the arena in which you normally play, there are now available AI platforms that will assist with research and case law review, and of course with case management and client communication. They exist now, and more are on the way.

On the horizon, there will be increasing availability of AI that will offer decision support, and award drafting. Already, general AI programs like ChatGPT are being used as drafting support - as Lord Chief Justice Birss (in the UK) said, ChatGPT is already "jolly useful in drafting rulings."
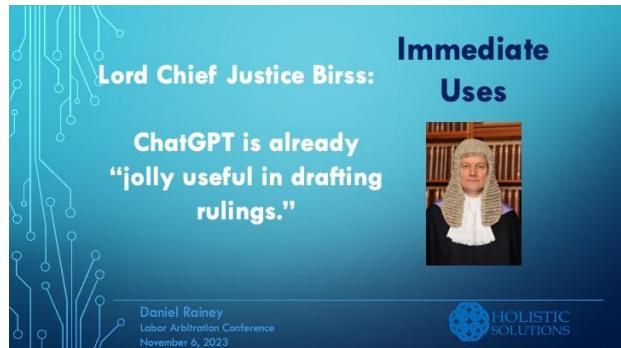
In one case, however, lawyers using ChatGPT to draft a court submission were sanctioned because of the AI's "hallucinations" - basically making up citations. The lawyers assured the judge they were unaware that AI could lie. The judge was not amused.

As AI programs come on the market, labor advocates, management advocates, and arbitrators should be keenly aware of the transparency of the systems offered for sale or service, and what data sets are being used to train the AI. The developers of generally available AI platforms, like ChatGPT, are reluctant to reveal the data set on which their systems are trained, for some understandable reasons. First, letting other developers know what data set I am using may give the other developers a clue about how to compete with me. Second, developers are reasonably afraid of being sued for intellectual property violations - this is already happening with book authors and with news organizations. Basically, if I don't reveal that I'm using your work to train my AI, you can't sue me. Third, there is an apocalyptic notion (not entirely without merit) that openly making AI platforms transparent will encourage bad actors to take advantage of the transparently communicated information to make better nefarious systems.

My basic message is that more targeted AI is on the way, that Fourth Party use of AI will eventually begin to merge with Third Party work (as co-mediators, second chairs, and as primary actors). The targeted AI, using very specific data sets as training material, means that users can be a bit more secure in the knowledge that the AI is learning on material germane to the venue in which it is being used. But, still, users should demand as much transparency as possible.

If you want more in-depth discussions of AI and arbitration, I suggest you check out
https://odr.info/ for the schedule of this year's CyberWeek, which runs online this week, from
November 6-10. There are three arbitration-themed sessions that will be offered live, and
recorded for later posting on odr.info.