



**The Journal of Robotics,
Artificial Intelligence & Law**

Editor's Note: Words, Languages, Algorithms, and Much, Much More

Victoria Prussen Spears

Unpacking Averages: Searching for Bias in Word Embeddings Trained on Food and Drug Administration Regulatory Documents

Bradley Merrill Thompson

Domain-Specific Languages and Legal Applications

Alexis Chun, Meng Weng Wong, and Marc Lauritsen

Equal Employment Opportunity Commission's Settlement Challenging Simple Algorithm Provides Warning for Employers Using Artificial Intelligence

Rachel V. See, Annette Tyman, and Joseph R. Vele

To Bot or Not to Bot: SEC's Proposed Conflict Rules May Stifle Use of Innovation

Sara P. Crovitz, Lawrence P. Stadulis, Peter M. Hong, Aliza S. Dominey, and Alexa Tzarnas

Copyright Office Seeking Comment on Human Authorship Requirements for AI-Generated Works

Mark A. Baghdassarian, Zachary B. Fields, and Jonathan R. Pepin

Does a License to "Make" a Patented Product Inherently Include a Right to Have a Third Party Make the Product or Its Components?

Sophie (Lu) Yan

Sentient Artificial Intelligence and the Rule of Law

Bazil Cunningham

- 5 Editor’s Note: Words, Languages, Algorithms, and Much, Much More**
Victoria Prussen Spears
- 9 Unpacking Averages: Searching for Bias in Word Embeddings Trained on Food and Drug Administration Regulatory Documents**
Bradley Merrill Thompson
- 19 Domain-Specific Languages and Legal Applications**
Alexis Chun, Meng Weng Wong, and Marc Lauritsen
- 43 Equal Employment Opportunity Commission’s Settlement Challenging Simple Algorithm Provides Warning for Employers Using Artificial Intelligence**
Rachel V. See, Annette Tyman, and Joseph R. Vele
- 47 To Bot or Not to Bot: SEC’s Proposed Conflict Rules May Stifle Use of Innovation**
Sara P. Crovitz, Lawrence P. Stadulis, Peter M. Hong, Aliza S. Dominey, and Alexa Tzarnas
- 53 Copyright Office Seeking Comment on Human Authorship Requirements for AI-Generated Works**
Mark A. Baghdassarian, Zachary B. Fields, and Jonathan R. Pepin
- 55 Does a License to “Make” a Patented Product Inherently Include a Right to Have a Third Party Make the Product or Its Components?**
Sophie (Lu) Yan
- 61 Sentient Artificial Intelligence and the Rule of Law**
Bazil Cunningham

EDITOR-IN-CHIEF

Steven A. Meyerowitz

President, Meyerowitz Communications Inc.

EDITOR

Victoria Prussen Spears

Senior Vice President, Meyerowitz Communications Inc.

BOARD OF EDITORS

Melody Drummond Hansen

Partner, Baker & Hostetler LLP

Jennifer A. Johnson

Partner, Covington & Burling LLP

Paul B. Keller

Partner, Allen & Overy LLP

Garry G. Mathiason

Shareholder, Littler Mendelson P.C.

Elaine D. Solomon

Partner, Blank Rome LLP

Linda J. Thayer

Partner, Finnegan, Henderson, Farabow, Garrett & Dunner LLP

Edward J. Walters

Chief Executive Officer, Fastcase Inc.

John Frank Weaver

Director, McLane Middleton, Professional Association

THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW (ISSN 2575-5633 (print) /ISSN 2575-5617 (online) at \$495.00 annually is published six times per year by Full Court Press, a Fastcase, Inc., imprint. Copyright 2024 Fastcase, Inc. No part of this journal may be reproduced in any form—by microfilm, xerography, or otherwise—or incorporated into any information retrieval system without the written permission of the copyright owner. For customer support, please contact Fastcase, Inc., 729 15th Street, NW, Suite 500, Washington, D.C. 20005, 202.999.4777 (phone), or email customer service at support@fastcase.com.

Publishing Staff

Publisher: Morgan Morrisette Wright

Production Editor: Sharon D. Ray

Cover Art Design: Juan Bustamante

Cite this publication as:

The Journal of Robotics, Artificial Intelligence & Law (Fastcase)

This publication is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional should be sought.

Copyright © 2024 Full Court Press, an imprint of Fastcase, Inc.

All Rights Reserved.

A Full Court Press, Fastcase, Inc., Publication

Editorial Office

729 15th Street, NW, Suite 500, Washington, D.C. 20005

<https://www.fastcase.com/>

POSTMASTER: Send address changes to THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW, 729 15th Street, NW, Suite 500, Washington, D.C. 20005.

Articles and Submissions

Direct editorial inquiries and send material for publication to:

Steven A. Meyerowitz, Editor-in-Chief, Meyerowitz Communications Inc.,
26910 Grand Central Parkway, #18R, Floral Park, NY 11005, smeyerowitz@
meyerowitzcommunications.com, 631.291.5541.

Material for publication is welcomed—articles, decisions, or other items of interest to attorneys and law firms, in-house counsel, corporate compliance officers, government agencies and their counsel, senior business executives, scientists, engineers, and anyone interested in the law governing artificial intelligence and robotics. This publication is designed to be accurate and authoritative, but neither the publisher nor the authors are rendering legal, accounting, or other professional services in this publication. If legal or other expert advice is desired, retain the services of an appropriate professional. The articles and columns reflect only the present considerations and views of the authors and do not necessarily reflect those of the firms or organizations with which they are affiliated, any of the former or present clients of the authors or their firms or organizations, or the editors or publisher.

QUESTIONS ABOUT THIS PUBLICATION?

For questions about the Editorial Content appearing in these volumes or reprint permission, please contact:

Morgan Morrisette Wright, Publisher, Full Court Press at mwright@fastcase.com
or at 202.999.4878

For questions or Sales and Customer Service:

Customer Service

Available 8 a.m.–8 p.m. Eastern Time
866.773.2782 (phone)
support@fastcase.com (email)

Sales

202.999.4777 (phone)
sales@fastcase.com (email)

ISSN 2575-5633 (print)
ISSN 2575-5617 (online)

Editor's Note

Words, Languages, Algorithms, and Much, Much More

Victoria Prussen Spears*

The authors of the articles that we are publishing in this issue of *The Journal of Robotics, Artificial Intelligence & Law* cover a cornucopia of subjects. They focus on words, languages, algorithms, and much, much more.

Word Embeddings

This issue begins with an article titled “Unpacking Averages: Searching for Bias in Word Embeddings Trained on Food and Drug Administration Regulatory Documents.” Here, Bradley Merrill Thompson of Epstein Becker & Green, P.C., explores how bias can creep into word embeddings by analyzing a model trained on what regulatory affairs professionals in industry and the Food and Drug Administration have written.

Domain-Specific Languages

Alexis Chun, Meng Weng Wong, and Marc Lauritsen authored our next article, titled “Domain-Specific Languages and Legal Applications.” In this article, the authors explain that despite the rise of low-code and no-code development tools and the maturation of large language model approaches in the software world, many legal software and application tools are still hand coded. They observe that one common bottleneck for legal software and application tools is the domain-specific, knowledge-based, and experience-based nature of legal practice, which makes legal tech a highly technical and multi-disciplinary endeavour. Moreover, the authors add, developers often need to encode legislation, regulations, legal concepts, and other quasi-legal frameworks in order to ask users the right questions, provide appropriate guidance,

accurately represent legal concepts, or generate the appropriate documents. The authors suggest that the difficulty of faithfully expressing such frameworks within the confines of custom code or within existing languages (natural or programming), and the resources required to resolve it, impede innovation.

This article analyses domain-specific languages (DSLs) as promising opportunities to lessen that difficulty, surveys 15 recent legal DSLs for semantic expressiveness and suitability for industry adoption according to an eight-point framework, and presents an innovative application of one such DSL to automatically generate a user-friendly web application, draw related visualizations to aid the drafter, and transpile to multiple targets for the convenience of researchers working in other languages.

Algorithms

Rachel V. See, Annette Tyman, and Joseph R. Vele of Seyfarth Shaw LLP are the authors of the next article in this issue, titled “Equal Employment Opportunity Commission’s Settlement Challenging Simple Algorithm Provides Warning for Employers Using Artificial Intelligence.”

In this piece, the authors discuss the implications for employers of a settlement reached recently by the U.S. Equal Employment Opportunity Commission in a lawsuit that many are calling the Commission’s “first ever” artificial intelligence discrimination in hiring case.

Bots

Then, in “To Bot or Not to Bot: SEC’s Proposed Conflict Rules May Stifle Use of Innovation,” Sara P. Crovitz, Lawrence P. Stadulis, Peter M. Hong, Aliza S. Dominey, and Alexa Tzarnas of Stradley Ronon Stevens & Young, LLP, summarize the Securities and Exchange Commission’s proposed new conflict of interest rules for investment advisers’ and broker-dealers’ use of certain predictive data analytics under the Securities Exchange Act of 1934 and the Investment Advisers Act of 1940, and the Commission’s proposal to narrow the internet advisers’ exemption under the Advisers Act.

Copyright and AI-Generate Works

The United States Copyright Office recently requested input from the public regarding “the scope and level of human authorship, if any, in copyright claims for material produced in whole or in part by generative AI.” Mark A. Baghdassarian, Zachary B. Fields, and Jonathan R. Pepin of Kramer Levin Naftalis & Frankel LLP discuss this development in their article, titled “Copyright Office Seeking Comment on Human Authorship Requirements for AI-Generated Works.”

Patent Licenses

The title of the next article—“Does a License to ‘Make’ a Patented Product Inherently Include a Right to Have a Third Party Make the Product or Its Components?”—asks a question that the author, Sophie (Lu) Yan of Baker Botts L.L.P., examines.

The author explains that patent license agreements should not only clearly define the rights that are granted but should expressly state the rights that are not being granted, such as the “have made” rights. In her opinion, the mere reliance on a broad “reservation of rights” provision or the assumption of no implied licenses should be avoided.

Sentience

Bazil Cunningham is the author of “Sentient Artificial Intelligence and the Rule of Law.” In this article, the author defines relevant parameters in determining non-human sentience. He then deciphers the beings that possess the requisite sentience and those that do not. In addition, the author identifies the philosophical basis of current legislative efforts and discusses the application of this basis through legal fiction by analogy to animals and artificial intelligence programs. Among other things, the author concludes by recommending changes to the rule of law in the United States.

Enjoy the issue!

Note

Victoria Prussen Spears, Editor of *The Journal of Robotics, Artificial Intelligence & Law*, is a writer, editor, and law firm marketing consultant for Meyerowitz Communications Inc. A graduate of Sarah Lawrence College and Brooklyn Law School, Ms. Spears was an attorney at a leading New York City law firm before joining Meyerowitz Communications. Ms. Spears, who also is Editor of *The Journal of Federal Agency Action*, *The Global Trade Law Journal*, and *The Global Regulatory Developments Journal*, can be reached at vpspears@meyerowitzcommunications.com.

Unpacking Averages: Searching for Bias in Word Embeddings Trained on Food and Drug Administration Regulatory Documents

Bradley Merrill Thompson*

In this article, the author explores how bias can creep into word embeddings by analyzing a model trained on what regulatory affairs professionals in industry and the Food and Drug Administration have written.

Often when we talk about bias in word embeddings, we are talking about such things as bias against race or sex. This article talks about bias a little bit more generally to explore attitudes we have that are manifest in the words we use about any number of topics.

Bias Evaluation Using Sentiment Analysis

There are many different ways to evaluate potential bias in word embeddings, but I did not want to do a survey article where I talked briefly about all of them. Instead, I thought I would pick just one approach for illustration. The one I picked is perhaps the simplest, which is an evaluation of the word embeddings using a model for positive versus negative sentiment. In other words, I am looking to see whether particular word embeddings have a largely positive or negative connotation.

If words that should be regarded similarly have significantly different sentiments or connotations, that would be evidence of bias. In other words, if the word “Black” as an adjective for people has a largely negative connotation while the word “white” as an adjective for people has a largely positive connotation, that would be some evidence that the embeddings, trained on what people have written, have absorbed from that training data a bias against Black people.

However, I am not going to use race as my example in the analysis below. For one thing, race is rarely discussed in the documents that I am going to examine—Food and Drug Administration (FDA) documents—apart from a handful of documents specifically on race. I will leave you to draw your own conclusions from that. Instead, I am going to look for bias in other topics.

Methodology

I wanted to keep it simple, so I will not use any of the cool, sophisticated, but complicated techniques that recently have been developed. Instead, for this first foray into the topic, I am going to use a methodology that has been around for a while because it is relatively simple to understand. In fact, I am shamelessly mimicking an approach used by Robyn Speer in her July 2017 piece entitled “How to Make a Racist AI Without Really Trying.”¹ I have updated the methodology only slightly to account for changes in software libraries since she published her article.

The basic approach is to train an algorithm—specifically a classifier—to recognize the differences between positive and negative words. In my particular case, I chose to use a random forest classifier from *sklearn* because it has shown to be effective in this sort of analysis.

This is an exercise in supervised learning, meaning that the algorithm needs to be trained by being told which words are positive and which words are negative in a training set. For that, I used a list of words that have been labeled as either positive or negative that most researchers in this space use, data created by Hu and Liu which are available from Bing Liu’s website.²

Here is the theoretical part. To recap, I have thousands of words that researchers have labeled as positive and thousands of words that researchers have labeled as negative. I also have thousands of these word embeddings through my prior machine learning work defined by 300 dimensions. The idea is that one or more of those 300 dimensions might correspond to positive versus negative connotations of the word. Thus, in theory, if I take the positive and negative words labeled as such from Hu and Liu, and if I represent those words using my 300-dimension word embeddings, I can train a machine learning classifier to spot positive versus negative words using the embeddings that I created from training on FDA

regulatory documents. In other words, the algorithm can learn which of the 300 dimensions in the word2vec model I created correspond to positive versus negative sentiment.

So that is what I did.

Validation

I like to document the uncertainty of any algorithm I use. I withheld about 10% of the data so I could test the algorithm on labeled data to see how well it did. I went into this assuming that I could perhaps get 90-95% accuracy from this exercise because that is what Robyn Speer in her original article achieved. But I could not. The best I could do was approaching 80% accuracy.

I spent a fair amount of time, for example, using grid search to experiment with different hyperparameters, and I also experimented with adding additional data for training purposes from other data sets. Oddly enough, adding more training data caused the performance to go down. Ultimately, I concluded that this was about the best I could do.

If you ask me why my performance was lower than what Robyn Speer achieved, she was analyzing some common word embeddings developed from training data such as Google News. In comparison, my training data were scientific and regulatory documents. While there are many differences between those data sources, at a high level I would say that regulatory professionals use fewer adjectives and adverbs in their writing. But adjectives and adverbs are the food of sentiment analysis. Without adjectives and adverbs, the algorithm has far less to go on in categorizing words as positive or negative. A sentence “the results were a score of 16” just doesn’t give the algorithm much to go on as to whether those words are positive or negative.

Or maybe I am just bad at it. But in any event, not quite 80% accuracy was the best I could do. Keep that in mind.

Validation Through Visual Exploration

Another way to validate the appropriateness of the algorithm that is little bit less scientific and relies much more on the anecdotal and the visual. I wanted to see how the results look, so I decided to assess an entire FDA guidance document on this negative versus


positive sentiment to see what it looked like. Obviously, an FDA document could be skewed positive or negative. I decided to go with one of the longer documents just because of regression to the mean. For a longer document, there should be positive and negative sentiment in the document. I therefore picked—not randomly—FDA’s September 2022 guidance on “Clinical Decision Support Software.”³ I picked it because it was one of the longer guidance documents that I previously analyzed, so I already had much of the code written. That is probably not a great reason.

But I wanted to see what the distribution of words were on this positive versus negative scale, and I wanted to weight them by the frequency of the word used in the guidance. Figure 1 is what that frequency looks like in graph form.

I have sorted the words from most negative to most positive. I did not list the words on the x-axis because, well, there are thousands and you would not be able to read them.

Just eyeballing it, I think it looks pretty good. The words in the middle are obviously more neutral, and then you have the two extremes. The extremes look somewhat symmetrical. But I never trust my eye when I am trying to gauge, for example, how much

Figure 1. Sentiment of FDA’s CDS Guidance



is positive versus how much is negative, so I thought I would run a simple calculation and take the average of all words used. That average turned out to be 0.0086400. You really could not expect much closer to zero. Thus, on the whole, in this particular guidance document, the negative sentiment words are pretty much in balance with the positive sentiment of words. I am not sure what that means but with my OCD tendencies I always like symmetry.

Truly Anecdotal Validation

Table 1	
Words	Sentiment score
"medical professional"	0.2645
"attorney"	-0.2699

Table 2	
Words	Sentiment score
"safe and effective"	1.8191
"adverse events"	-1.1846
"notification"	0.4109
"alarm"	-2.2675
"hospital"	-0.1968
"at home"	-0.3431

Table 3	
Words	Sentiment score
"the device clinical trial was successful"	0.1353
"the device clinical trial was a failure"	-0.5736

Words	Sentiment score
"the device saved many lives"	-0.0079
"the device had few side effects"	-0.1665

Okay, so big picture it looks sort of reasonable although I have no objective evaluation of what the average of that particular guidance document should be. But why not look at a few individual words to see if their positive/negative scores seem intuitively correct. Remember that correct should be whatever we anticipate a group of FDA regulatory scientists thinking.

Table 1 compares two random words/phrases. Okay, well that seems about right.

I should point out that I really do not have any evidence to suggest that the absolute level of a sentiment score is all that accurate. As a result, what I focus on is comparing words to see if the comparisons ring true. So, Table 2 has a variety of comparisons I tried, selecting words that are somewhat common in FDA regulatory writing.

I always suggest that clients use the word "notification" instead of "alarm," and now I have the data to back that up.

It is kind of interesting that hospital is a bit negative, although certainly from a popular perspective people don't want to be at the hospital. But it is also interesting that "at home" is more negative through the eyes of FDA regulatory professionals.

One of the things that I noticed is that because of the way this algorithm is designed, evaluating sentences that include lots of words necessarily moves the score toward zero because the algorithm is just taking an average of the words, so with more words you regress to the mean. But even so, and even with many words the same, the algorithm does reasonably well with certain sentences, as shown in Table 3.

I also want to just remind you again that the algorithm is only about 80% accurate, so there are some results that caused me to scratch my head, like the ones in Table 4.

Remember too that the algorithm is just analyzing words, not sentences, so it does not catch the profoundly different meaning that words such as "not" or "few" convey in giving sentiment analysis to a sentence.

With those limitations, it seems as though the algorithm analyzing the sentiment of words created in the word embeddings is at least interesting if not somewhat accurate.

Results

What does all this tell us about whether there is bias in these word embeddings I created by training on FDA regulatory documents? Let's continue with a few comparisons in areas where I wondered if there might be bias.

Over the nearly 40 years I have been practicing FDA law, I get the sense that FDA regulatory professionals have sometimes strong opinions about the countries from which data are gathered. Let's look at Table 5 to see if there are any differences.

I need to start by observing that many of those differences are not statistically significant. We are dealing with some small numbers here frankly all clustered around neutral. But there are some stark differences, such as the difference between, say, Japan and Russia.

It is important to remember that the training data set is just the premarket review summaries as well as generally FDA guidance documents. There is really not much of that training set from

Words	Sentiment score
"Mexico"	0.0783
"China"	0.0435
"United Kingdom"	0.1112
"Russia"	-0.5466
"France"	-0.1220
"Japan"	0.3084
"Foreign data"	-0.3008
"data"	0.1719
"foreign clinical trial"	-0.3959
"US clinical trial"	-0.1458

Table 6	
Words	Sentiment score
"recall"	0.0337
"warning letter"	-1.0746

Table 7	
Words	Sentiment score
"software"	0.4422
"hardware"	0.4782
"in vitro diagnostics"	-0.0399
"acupuncture"	0.2168
"pedical screw"	-0.2985
"ventilator"	0.2377
"infusion pump"	-0.1383
"minimally invasive"	-1.1198
"aid in diagnosis"	-0.2770
"pediatric"	-0.1878

enforcement or quality contexts, so this really wouldn't reflect FDA enforcement views.

I included the last four in Table 5 to show a more global level sentiment around foreign data versus data more generally. I honestly cannot explain why the U.S. clinical trial sentiment would be negative.

Now, look at Table 6. I thought I would assess the sentiment of some regulatory words.

I like the fact that recall is more neutral because it is simply the responsible action of a manufacturer to address the occasional but unavoidable quality issue, where warning letter has a decidedly negative connotation.

I then wanted to assess product words (Table 7) to see if there are connotations associated with different specific or even general categories of products.

Words	Sentiment score
"man"	-0.7500
"woman"	-0.2115

I will let you draw your own conclusions from those, but again, keep in mind, only 80% accurate and the magnitude of the actual scale has not been validated in any way. I really do not understand the “minimally invasive” result.

I could go on, but I will close with this. I mentioned above that I did not see much point in including race in this discussion because so few documents discuss it. I think more discussed is sex because sex has long been recognized as a factor that needs to be considered. Consider the sentiment scores for the sexes in Table 8. I will let you draw your own conclusions from that.

Conclusions

The whole point of this exercise is to illustrate that any word embeddings, because they are trained on human input, will have biases. That is true because no human being on earth is free from bias, so any machine learning model trained on that human input will have those biases.

We must be aware of those biases in all natural language processing, and more than that we must find them and then account for them. It often is impossible to remove them, but there are other coping mechanisms we have developed such as explicitly considering the existence of the bias.

In the future, I will dive deeper into this topic because I find it personally interesting, but it also is one that I think many companies need to consider on a more sophisticated basis.

Notes

* The author, a member of Epstein Becker & Green, P.C., counsels medical device, drug, and combination product companies on a wide range of Food and Drug Administration and Federal Trade Commission regulatory,

reimbursement, and clinical trial issues. He may be contacted at bthompson@ebglaw.com.

1. <https://blog.conceptnet.io/posts/2017/how-to-make-a-racist-ai-without-really-trying/>.

2. <https://www.cs.uic.edu/~liub/FBS/sentiment-analysis.html#lexicon>.

3. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>.

Domain-Specific Languages and Legal Applications

Alexis Chun, Meng Weng Wong, and Marc Lauritsen*

Despite the rise of low-code and no-code development tools and the maturation of large language model approaches in the software world, many legal software and application tools are still hand coded. One common bottleneck for legal software and application tools is the domain-specific, knowledge-based, and experience-based nature of legal practice, which makes legal tech a highly technical and multi-disciplinary endeavour. Developers often need to encode legislation, regulations, legal concepts, and other quasi-legal frameworks in order to ask users the right questions, provide appropriate guidance, accurately represent legal concepts, or generate the appropriate documents. The difficulty of faithfully expressing such frameworks within the confines of custom code or within existing languages (natural or programming), and the resources required to resolve it, impede innovation. This article analyses domain-specific languages (DSLs) as promising opportunities to lessen that difficulty, surveys 15 recent legal DSLs for semantic expressiveness and suitability for industry adoption according to an eight-point framework, and presents an innovative application of one such DSL to automatically generate a user-friendly web application, draw related visualizations to aid the drafter, and transpile to multiple targets for the convenience of researchers working in other languages.

Introduction

Readily available and inexpensive codified legal know-how is increasingly critical in both commercial and nonprofit contexts. Yet it often remains costly and time-consuming to produce.

Most practical legal applications are created and maintained using laborious hand-coding techniques, often including quite primitive methods. Most app makers (including one of the authors) are not professional software developers. That is true within private law firms and law departments as well as in nonprofit organizations. Legal application developers who follow the academic literature have long been aware of methodologies and theories to directly connect statements of the law with their programmed implementation, but few so far have taken advantage of them.

In short, much legal app development remains highly artisanal. Domain-specific languages (DSLs) may offer a solution.

This article is organized as follows. Following this introduction, the second section describes some common forms of interactive legal applications and their development processes, including two examples. The third section lays out some of the challenges developers face and imagined solutions. The fourth section introduces DSLs and their applicability. The fifth section 5 introduces the L4 DSL with example screenshots. And the sixth section concludes.

Contemporary Legal Knowledge Engineering

Expert systems and various forms of document automation are among the most common forms of knowledge-based software found in law offices in recent decades. A common pattern involves scripted “interviews” and modelled documents, which are typically fashioned using procedural code and manual document markup.

Tools like Neota Logic, BRYTER, Contract Express, HotDocs, and Legito provide integrated development environments within which such apps can be built and maintained. (There is a wealth of such tools. One site (<https://www.docautodatabase.com/>) recently identified over 200 in the document automation category alone.) Another collection (with over 5,000) of such applications in the nonprofit sector in the United States is at LawHelp Interactive (LHI), which provides interactive guidance and bespoke form assembly without charge to millions of users. (About a million packages of customized forms were generated in 2022.) Within its technology stack the main providers of end-user functionality are HotDocs, from CARET, and A2J Author, from the Center for Computer-aided Instruction.

Document automation applications are typically driven by the forms they need to generate (What information should be placed where under what circumstances?) and by informal know-how communicated by practitioners (What should users know about the process they are undergoing? What steps should be taken or avoided as a practical matter to reach an optimal outcome?). But sometimes they also need to explicitly reflect the detailed rules expressed in a statute or regulation. In those situations, scripted interviews and model documents are not sufficient.

Two Examples

The Uniform Child-Custody Jurisdiction and Enforcement Act (UCCJEA) has been adopted by 49 U.S. states, the District

of Columbia, Guam, Puerto Rico, and the U.S. Virgin Islands. It governs the rules whereby courts decide which have jurisdiction to adjudicate questions of child custody. Family law applications typically need to encode aspects of the UCCJEA in order to advise users and properly complete court forms.

One can find various online resources that attempt to summarize how the UCCJEA “works,” such as shown in Figure 1.

A2J Author provides an easy-to-use environment via which non-programmers can script “guided interviews.” A built-in mapper helps users visualize their creations. Some can quickly become unwieldy, such as that shown in Figure 2.

Figure 1

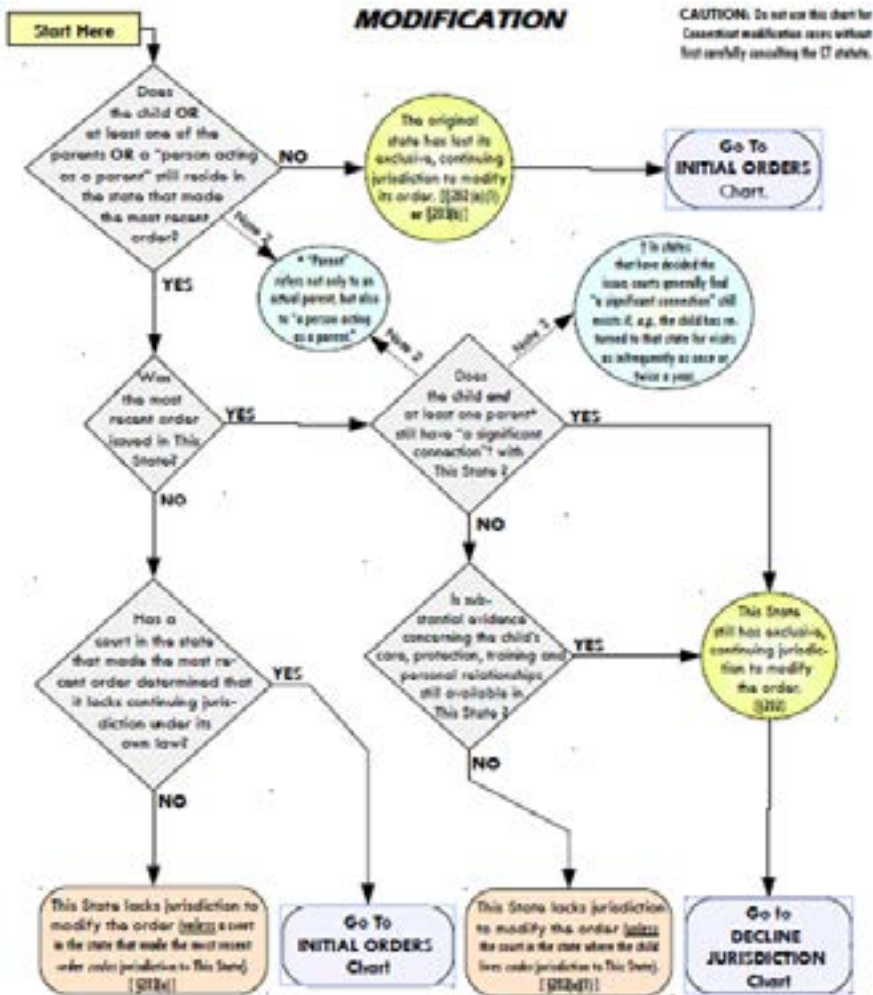
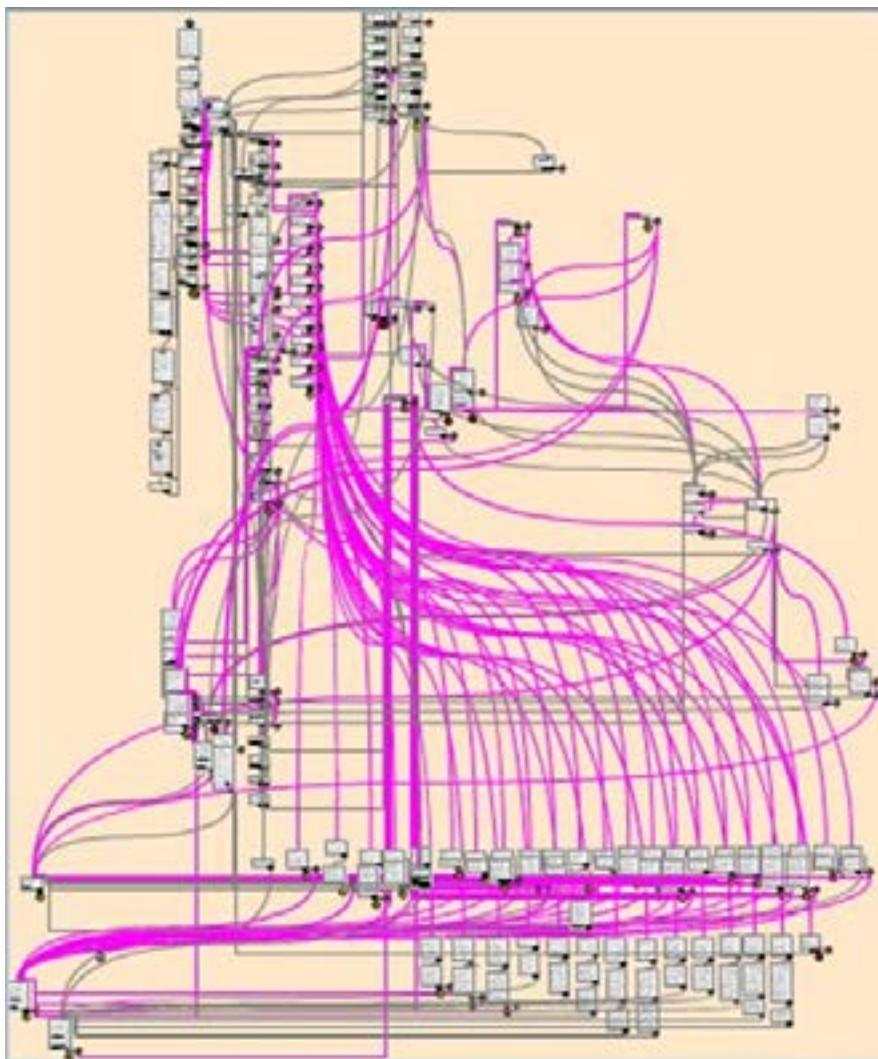


Figure 2

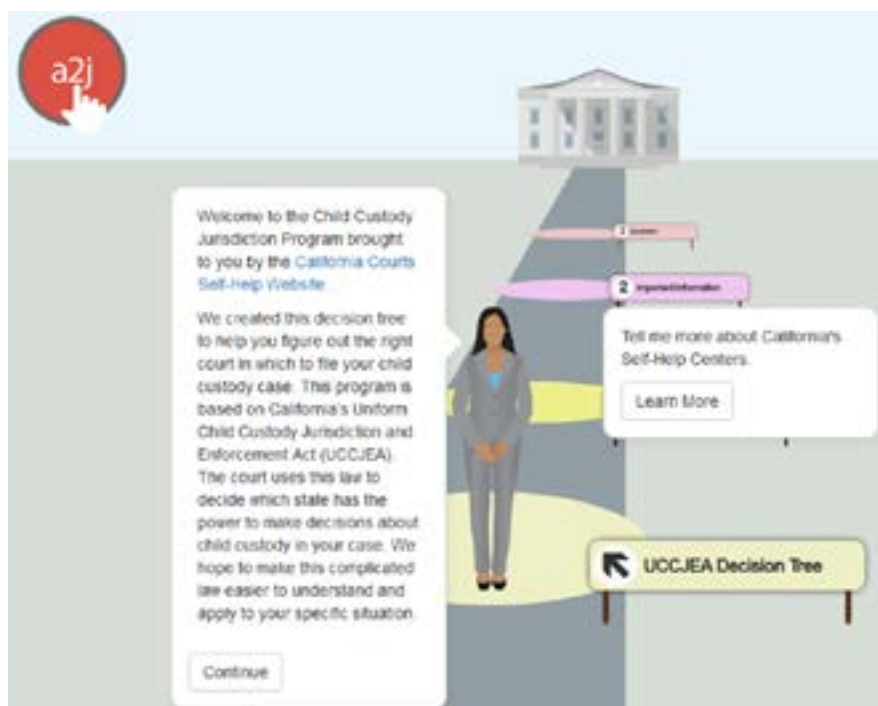


One LHI application, built in A2J Author, that one of the authors has assisted with was intended to simply guide an inquiring user as to which court likely has jurisdiction to handle questions about the custody of their children. Its first page looks like that shown in Figure 3.

Its associated “map” is quite sparse. See Figure 4.

The challenge for the developer (after a succession of earlier developers) was to confirm whether the app faithfully followed at

Figure 3



least key parts of the UCCJEA, and gave accurate guidance. That involved in part coming up with an external representation of its key provisions. One resorted to sticky notes in the attempt (see Figure 5).

Another less-than-satisfactory effort involved page-by-page documentation in Word, which also didn't capture the as-built logic of this application.

A second application needing to reflect the UCCJEA's logic was a HotDocs interview and template set for litigants seeking a divorce in Washington State. There, a domain expert (practicing lawyer) struggled to capture that logic so that it could be expressed in HotDocs code, and ended up finding Excel the best tool for doing so (see Figure 6).

That in turn was used by the HotDocs expert to create a set of computations that drive the interview and infer the proper result. For an example, see Figure 7.

Figure 4

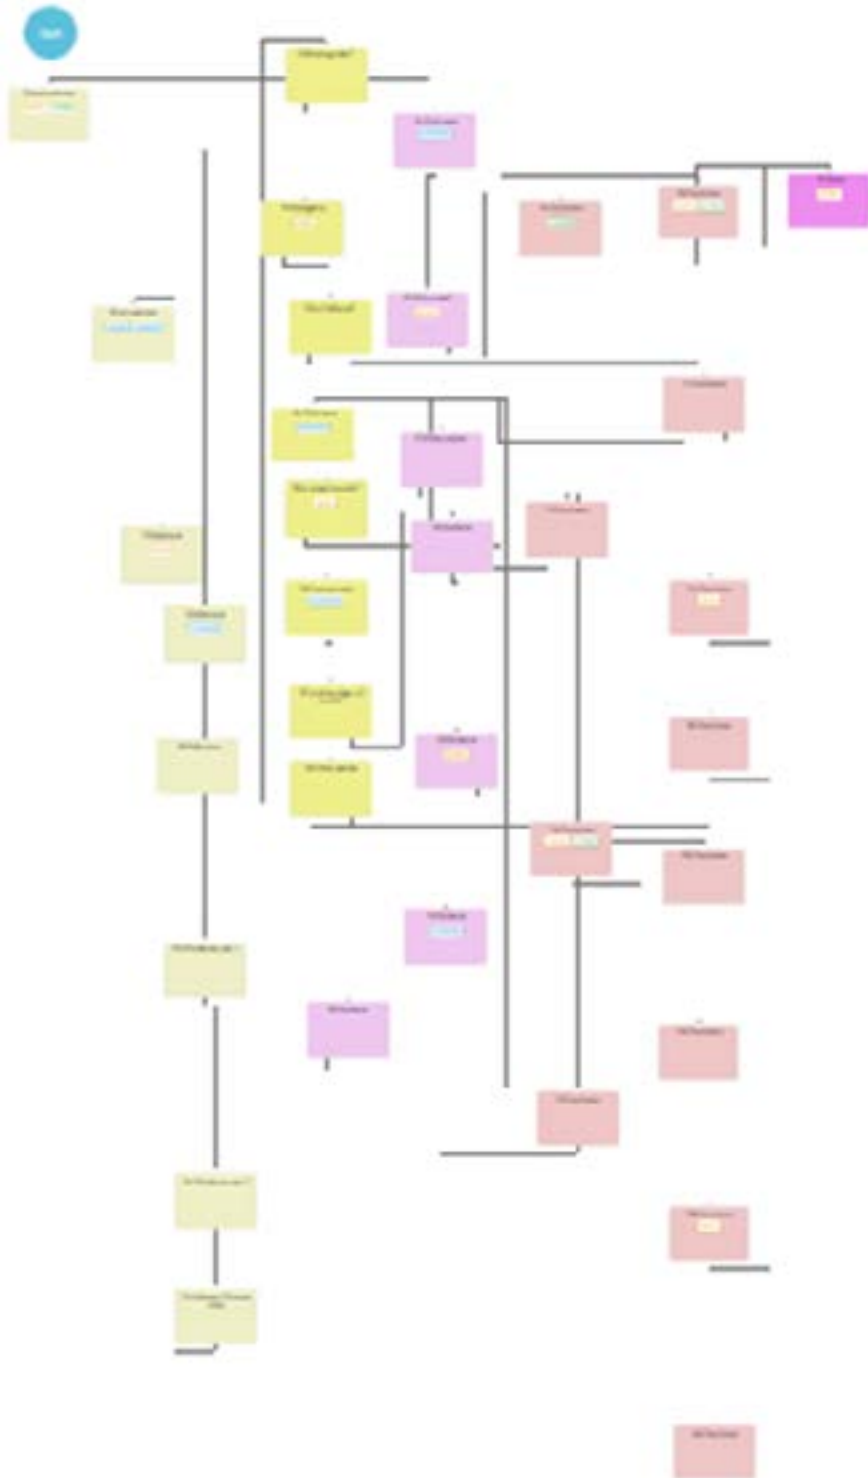


Figure 7

```

IF Child jur collect uccjes info together CO
  IF (VALUE(Children all residence location MC[1]) != "Washington" AND
      (NOT VALUE(Child jur all non WA residence 6 mos plus TF)) AND
      VALUE(Child jur all prior residence WA 6 months plus TF))
    "<p>>a, b and c">"
    SET Counter TO 1
    WHILE ANSWERED(Child name last TE[Counter])
      IF AGE(Child birth date DA[Counter]) < 18
        RESULT + Child name full indexed CO(Counter) + "<p>"
      END IF
      INCREMENT Counter
    END WHILE
    RESULT + "<p>"
  END IF
ELSE
  REPEAT Children Repeat
    FORMAT ">a, b and c"
    FILTER Child jur WA home state w/in 6 months filter CO
    RESULT + Child name full CO
  END REPEAT
END IF

```

Opportunities and Challenges

The above examples are just briefly sketched to illustrate the challenges faced by developers and their collaborators. Namely,

1. There is no widely recognized methodology for reliably incorporating statutory rules into a custom programmed application.
2. There is no widely recognized method for confirming whether one's efforts to so incorporate legal "code" into such an application were successful. This raises serious quality control issues.¹
3. Domain experts generally are not able to review application code itself to satisfy themselves about its completeness and consistency.
4. Such applications lack automated explainability. They do not readily interoperate with external specifications or code. (Most legal apps are poor at explaining themselves because we haven't educated them about *why* particular

questions are asked, guidance is offered, and documents are generated.)

5. From a computer science perspective, these development processes and ad hoc knowledge representation formats ignore decades of advances in information management, software engineering, and programming language theory.

The literature around computable contracts,² computational law,³ and Rules as Code⁴ points to a future in which the above problems have been solved. What does that future hold?

When building a legal assistance app that needs to reflect a defined set of rules (from a statute, regulation, or other source of governance), the developer can access both the natural language statement of those rules and an unambiguous, machine-readable equivalent. Software can bidirectionally move between both isomorphic forms. Those forms can be used as input to a design-time process that produces appropriate code for the destination platform. For instance, interviews can automatically be generated that ask the minimal set of questions needed to resolve a legal issue. (An optimal “question tree.”) Alternatively, a run-time process could deliver needed logic to that platform via an application programming interface (API). Conversely, tools would be available to generate an external specification of the logic of an application for purposes of validation, maintenance, and debugging. Tools could automatically construct graphs, flowcharts, decision trees, and other visualizations to represent laws and contracts and aid end-user understanding of legal complexity. Such outputs could also be used to support in-session explanations of inferences performed against user inputs, meeting the goals of explainability and algorithmic transparency. The system would usefully identify all ultimate and intermediate conclusions described in a model, as well as all predicates and data elements playing roles in rules/inferences.

Moving upstream, certified software encodings could be published by government or other relevant authorities. The open-source movement of the past four decades, overlapping with the ideals of the rule of law, demand that digital legislation and regulations should be publicly available for free.⁵ For example, a state agency could release a “code companion” library on Github for consumption by third-party app developers, minimizing the need for software developers to conduct legislative interpretation.

Current tools attempt to achieve some of these goals. Complex computations are not easily implemented in A2J Author, but HotDocs offers a reasonably complete programming environment for such things, including parameterized computations and local variables. Such things can be used to drive questioning and inferences; the challenge is writing, validating, and updating them! A2J Author is over 20 years old; HotDocs is over 30. The difficulties of revising software products to support fundamental new functionality are well known.

What next-generation technologies could help realize the vision?

DSLs to the Rescue

In the computer science and software engineering disciplines, DSLs are a widely accepted approach to making a particular problem domain more tractable to software and to developers. For example, the need to structure hypertext data begat HTML; the need to manage the visual styles and layout of web pages begat CSS; the need to read from and write to databases containing tabular data begat SQL. All are DSLs, defined as:⁶

A domain-specific language (DSL) is a programming language or executable specification language that offers, through appropriate notations and abstractions, expressive power focused on, and usually restricted to, a particular problem domain.

In recent years, academics and software developers have seized on DSLs as a promising way to enable the vision outlined above.⁷ Computer scientists have proposed languages and libraries for law (FormaLex, Catala, OpenFisca); “smart contract” languages have appeared with the rise of blockchain technologies (Accord Project, Cardano, Deon Digital’s CSL); and non-blockchain-oriented contract languages have also appeared (FCL from McMaster, Symboleo from uOttawa, Logical English from Imperial) many of which were inspired by Jones, Eber, and Seward’s pioneering 2001 paper “Composing Contracts.”⁸ Others include Stipula, DCR Graphs, Orlando, McCarty’s LLD, Eiger, and Blawx.⁹

Brief Survey of Legal DSLs

Due to space constraints, the capsule summaries presented in this section may not do justice to the full vision of each language but are intended to illustrate the range of approaches to legal DSLs. A comprehensive survey is beyond the scope of this article. DSLs such as ACC,¹⁰ BCL,¹¹ CL,¹² DataLex,¹³ Lexon,¹⁴ and RegelSprak¹⁵ are omitted with regrets.

- **OpenFisca:** France, 2011; Python API; primarily numerical calculations for income tax and other quantitative domains; support for multiple versions of legislation and multiple jurisdictions (FR, US, UK, AU, NZ).
- **Catala:**¹⁶ France, 2019; external DSL; strong support for numerical calculations and for isomorphic representation of statutes expressed in terms of default logic with exceptions.
- **FormaLex:**¹⁷ Buenos Aires, 2011; based on LTL (Linear Temporal Logic) and intended to discover inconsistencies using model checking.
- **FCL:**¹⁸ McMaster University, 2018; a type-theoretic approach to formalizing and reasoning over events, deontics, and real values.
- **Symboleo:**¹⁹ uOttawa, 2020; emphasis on deontic logic, with support for events, pre- and post-conditions, and assertions.
- **Stipula:**²⁰ Italy, 2021; emphasis on timed deontics as state, with support for assets and the notion of agreement as synchronization.
- **Blawx:**²¹ Canada, 2020; focus on usability through a GUI based on Scratch; well-formed statements can be constructed through drag-and-drop.
- **DCR Graphs:**²² Copenhagen, 2011; a declarative, event-based process model developed in partnership with industry.
- **Eiger:**²³ Switzerland, 2022; embedded Haskell DSL, deployed at PwC Switzerland.
- **Orlando:**²⁴ United States, 2021; an academic project featuring a concise CNL, with strong visualization and explainability features, for conveyancing as the initial application domain.

- **Accord Project:**²⁵ United States, 2019; intended for blockchain use, provides an executable language, a data modeler, and a document assembler.
- **CSL:**²⁶ Copenhagen, 2012; a trace-based external DSL with support for events and deontics, adapted for blockchain use and currently the subject of commercialization efforts at the start-up deondigital.com.
- **Logical English:**²⁷ Imperial, 2020; web-based logic programming with syntactic sugar borrowed from the tradition of Controlled Natural Languages; uses the Event Calculus to track state over time.
- **Epilog:**²⁸ Stanford, 1980s; a member of the logic programming family with a focus on databases and unusual direct support for logic programming in an interactive web environment.
- **Language for Legal Discourse:**²⁹ Rutgers, 1989; a sophisticated theoretical basis for converting from legal natural language to a formalization.

Analytic Frameworks for Legal DSLs

Requirements for legal specification languages have been previously enumerated. Hvitved³⁰ identifies the following 16 requirements:

1. Contract model, contract language, and a formal semantics.
2. Contract participants.
3. (Conditional) commitments.
4. Absolute temporal constraints.
5. Relative temporal constraints.
6. Reparation clauses.
7. Instantaneous and continuous actions.
8. Potentially infinite and repetitive contracts.
9. Time-varying, external dependencies (observables).
10. History-sensitive commitments.
11. In-place expressions.
12. Parametrised contracts.
13. Isomorphic encoding.
14. Run-time monitoring.

15. Blame assignment.
16. Amenability to (compositional) analysis.

Athan et al.³¹ identify the following functionalities:

1. Supports modelling different types of rules (constitutive v. prescriptive).
2. Represents normative effects (e.g., reparation and compensation).
3. Implements defeasibility (to handle conflicts between rules).
4. Implements isomorphism.
5. Alternatives (can represent multiple interpretations).
6. Manages rule reification (Jurisdiction, Authority, Temporal attributes).

We introduce a framework that consolidates the above formal requirements under semantics and expressiveness (criteria 1-3), and goes beyond to anticipate usability concerns and suitability for adoption in industry and government (criteria 4-8):

1. Equipped with a formal semantics describing the language in terms of its underlying logics (defeasible, default, temporal, deontic, etc.).
2. Capable of expressing a wide variety of contract genres (such as financial agreements, insurance policies, employment contracts, and leases).
3. Capable of expressing a wide variety of legislative and regulative genres (such as criminal law, building permits, privacy regulations, and even rules of court).
4. Open-source implementation available (some languages are given only as theoretical constructs without accompanying software; others are proprietary).
5. Syntactically “low code” and user friendly with documentation and integrated development environment (IDE) support (intended to be read and written by an individual without extensive training in programming or law).
6. Capable of producing explanations for its decisions, in text or via visual notations.
7. Application-oriented (intended for industry use).

Table 1								
	1	2	3	4	5	6	7	8
OpenFisca			Q	T			T	
Catala	T		Q	T			T	
FormaLex	T		E	T				
FCL	T	E						
Symboleo	T	E		T				
Stipula	T	E		T		T		
Blawx			T	T	T	T	T	
DCR Graphs	T	E	E			T	T	
Eiger	T	T		T	E		T	
Orlando	T			T	T	T		
Accord		T		T	T		T	
CSL	T	T					T	
Logical English	T	T	T	T	T	T	T	
Epilog	T	T	T	T	T		T	
LLD		T	T					

T: true (blanks indicate insufficient information to conclude true; logic programmers may consider this negation-as-failure).

Q: the primary expression domain is quantitative calculations.

E: the primary expression domain is an event-oriented calculus.

8. Oriented toward interoperability (imports from and exports to other languages and standard formats such as LegalRuleML, BPMN, and DMN).

As of early 2023, using that framework, a rough assessment (*pace* the authors of the languages) produced the analysis shown in Table 1.

The analysis shows that many legal languages, while rigorously defined, are focused on relatively narrow areas of concern: either laws or contracts; either quantitative calculations or state-transition systems with an emphasis on deontics and verifiability. To realize the vision of wider adoption, additional requirements must be satisfied, which go beyond the charter of the typical academic research project.

In 2020, a research program was begun to develop a DSL for laws and contracts that meets all the above criteria.

The L4 DSL

The remainder of this article identifies L4 as a novel solution in the space of legal DSLs by informally outlining its semantics and expressive scope. A brief walkthrough of a real-world use of L4 is presented to illustrate how it supports innovative applications that fulfill the features and vision from the third section.

Semantics

The L4 DSL combines first-order logic for reasoning over “static” decisions such as numerical calculations and Boolean predicates, with the semantics of a state transition system for reasoning over “dynamic” events and obligations in time. The guards of the state transitions are expressed using the “static” logic. These two major sets of semantics—the “statics” and the “dynamics”—are visualized using circuit diagrams and process workflow diagrams, respectively. These semantics have been found to be sufficient to formalize all the case studies encountered so far.

Default Logic

The “static” rules have a concrete syntax that can be considered a sugared form of Prolog. Default reasoning is supported with the use of default branches in pattern matches. The runtime reasoner is augmented with two modes of operation: in “hard” mode, only user input is used to calculate decisions; in “soft” mode, input elements can be marked using the `TYPICALLY` key word; these defaults are provisionally accepted into decisions and treated as assumptions for the user to confirm or deny.

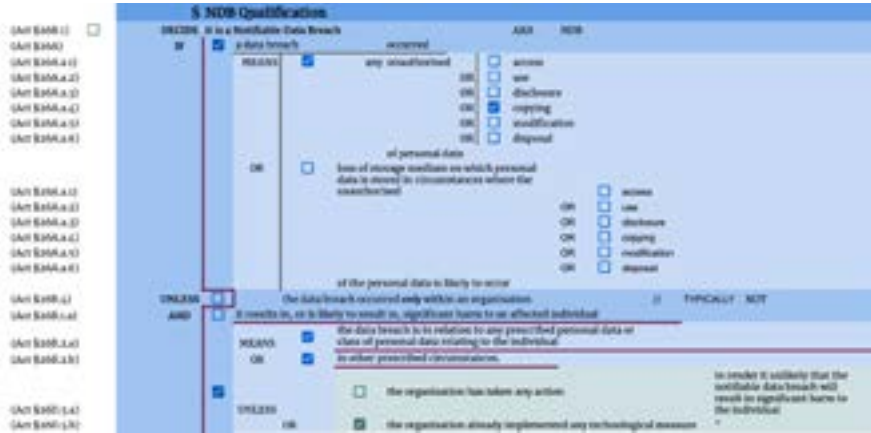
Interactions Between Rules

Legal clauses are frequently prefixed with “notwithstanding,” “despite,” and “subject to” modifiers. L4 interprets these modifiers as a priority ordering and adjusts rule application and result chaining accordingly. In this way L4 supports a limited form of defeasible logic.

Spreadsheet as Interactive Development Environment

With industry adoption in mind, L4 prioritizes a spreadsheet-based IDE over the traditional text editor. This innovation delegates

Figure 8



certain lexing and value-typing functionality from the parser to the IDE.

In 2021-2022, a case study called for the encoding of a portion of real-world privacy legislation. The source material spanned approximately 260 pages of text, including advisory guidelines and a compliance guide for organizations. The completed encoding occupied approximately 260 lines of code.

In this case study, the primary rules are as follows: a data breach, once discovered, must be assessed; and if it is assessed to be a notifiable data breach, it must be reported to both the relevant government body and to the affected individuals. Both rounds of obligations come with deadlines. The decision criteria for whether a breach is notifiable are complex.

An Example of Constitutive Rules

The *decision* as to whether a data breach is *notifiable* can be expressed using the constitutive rule shown above. The rule is essentially a Boolean proposition composed with the operators AND, OR, NOT, and UNLESS, where grouping is indicated using layout indentation.

The L4 tooling automatically generates the corresponding decision diagram in a variety of formats and semantic resolutions. The simplest format shows the decision nodes in a circuit diagram of parallel (OR) and series (AND) elements. The more detailed format includes the text of each node. This diagram makes it easy to

Figure 9




Figure 10

§ Assessment	
(Art §16C.2)	EVERY Organisation
(Art §4.14)	WHICH NOT is a Public Agency
(Art §16C.2)	UPON becoming aware a data breach may have occurred
(Art §16C.1)	IF the data breach occurs ON 1 Feb 2022 // TYPICALLY yes
(Art §16C.1)	OR "
(Art §16C.2)	MUST assess if it is a Notifiable Data Breach
(Art §16C.2)	by evaluating NDB-Qualification
	WITHIN 30 days
	HENCE Notification

quickly form an impression of the overall shape and structure of the decision logic.

An Example of Prescriptive Rules

The *obligation* to assess whether a breach is notifiable can be expressed using the following regulative/prescriptive rule, which contains deontic, epistemic, and temporal elements.

The workflow diagram corresponding to the full rule set is rendered in the form of a Petri Net (see Figure 11). Other formalisms may follow in future.

Transpilation to Other Formats

Once these rules are parsed into the L4 interpreter’s abstract syntax tree (AST) and related intermediate representation formats, they can be rewritten and transpiled to a variety of downstream representations. As of mid 2023, L4 supports output to JSON, Typescript, Purescript, and Python. On the road map are other languages and formats such as DocAssemble, Catala, OpenFisca, Blawx, Prolog, and Epilog, as well as interchange standards like LegalRuleML, BPMN, and DMN. In response to industry demand, other formats and technology stacks could be added to that list—Neota Logic, BRYTER, HotDocs, and others are potential transpilation targets so that enterprises already committed to a document assembly or contract life cycle management platform can integrate L4 with existing business processes. Any existing or future academic language can also be supported as a transpilation target, opening the door to research cross-compatibility.

Figure 11

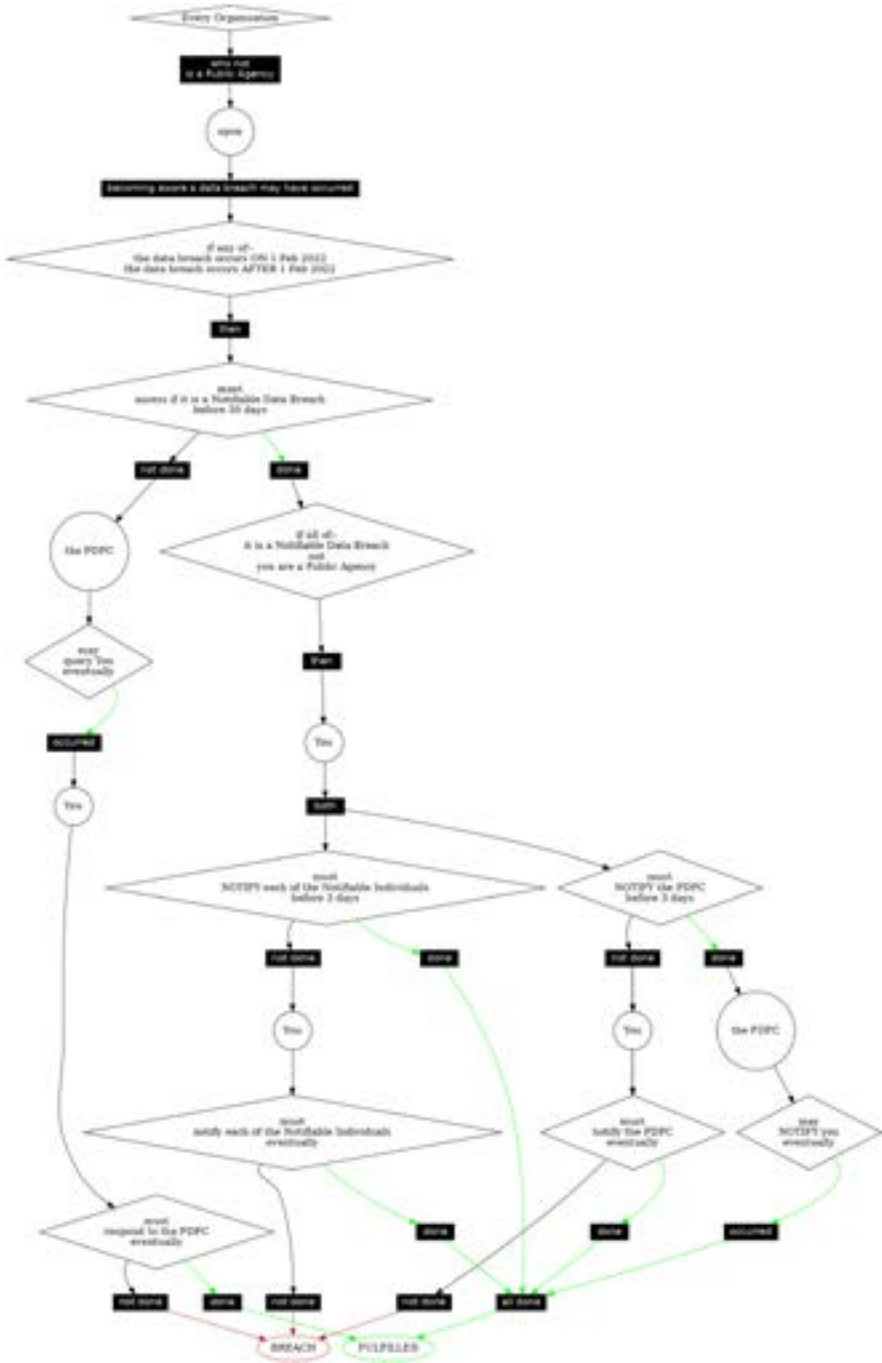


Figure 12

PCRA DSHO PoC - dash

Questions Diagram

Must you notify?

It depends...

Please share as much as you know:

Qualifiers:

Is the Organisation a Public Agency? Yes No Don't Know

Did the data breach occur on or after 1 February 2017? Yes No Don't Know

Was it a notifiable data breach?

Did a data breach occur? Yes No Don't Know

Was there any

unauthorised access of personal data Yes No Don't Know

unauthorised use of personal data

Automated Web App Generation

The encoding of legislation into L4 was a means to an end. In this case study, the goal was to automate the creation of a citizen-facing web application from the formalization. To that end, a reusable toolchain involving a transpiler to Purescript and a front-end in Vue was developed to convey the legal logic from the encoding to an interactive application for citizens and affected enterprises.

As the input spreadsheet is edited, a web application is regenerated live, with a typical rebuild time of less than 10 seconds. As end-users answer the questions presented in the app, the decision logic attempts to resolve the top-level answer to a “yes” or a “no.” The L4 toolchain thus meets the description of an “application generator” as described by Cleaveland.³²

This web app is not polished to commercial standards but was developed as a proof-of-concept to demonstrate the feasibility of the “Rules as Code” approach. The entire package can be bundled for further refinement and public-facing delivery.

All components of the system, including front-end IDE support (in Google Sheets, powered by Google Apps Script), the L4 parser/interpreter toolchain, the visualizers, the transpilers, and web app infrastructure are available on Github.

Conclusions

The pain points of the current legal application development model can be remedied by the adoption of a DSL-based engineering methodology. Software engineering principles like “separation of concerns” advise that rather than implementing the “business logic” of the law directly in operational software, one should abstract out representations of the law into an executable specification, in a DSL with the appropriate semantics. In recent years, following this motivation, DSLs have been developed in academia and by industry (typically with blockchain applications in mind), each one exploring a different theoretical approach. In 2019 the authors detected an opportunity to make a novel contribution, at the intersection of wide semantic expressivity, “low-code” usability, and a focus on adoption by industry and governments, through comprehensive tooling, open-source availability, and planned interoperability with existing systems. This article presents an encoding of real-world legislation into L4, presents some of the syntax for constitutive and prescriptive rules, and shows how a user-facing web application can be generated automatically.

It is straightforward to envision how the UCCJEA examples could benefit from this treatment: the encoding language does not have to be developed ad hoc; the development environment provides supporting visualizations to aid the drafter; and the accompanying tools are responsible for exporting to formats that can be consumed by downstream applications, if the natively generated applications are not already sufficient to serve the user. Keeping the legal rules explicit supports the goals of explainability and transparency which are increasingly important social priorities. The use of open DSLs to support legal applications is thereby shown to be a key ingredient of the vision outlined in this article.

Notes

* Alexis Chun is a co-founder of Legalese, a computational law startup that applies computer science to law, and the industry director at the Centre for Computational Law, Singapore Management University. Meng Weng Wong is a co-founder of Legalese and the principal investigator at the Centre for Computational Law. Marc Lauritsen, the president of Capstone Practice Systems and a former senior research associate at Harvard Law School, teaches courses at Suffolk University in which students build software applications

that help with legal work. This research/project is supported by the National Research Foundation, Singapore under its Industry Alignment Fund—Pre-positioning (IAF-PP) Funding Initiative. We thank Laurie Garber of the Northwest Justice Project and Bart Earle of Capstone Practice Systems for some of the examples in subsection Two Examples. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

1. *See, e.g.*, Marc Lauritsen & Quinten Steenhuis, “Substantive Legal Software Quality: A Gathering Storm?” 52-62 (June 17, 2019), <https://dl.acm.org/doi/10.1145/3322640.3326706>.

2. Harry Surden, “Computable Contracts” (2019), <https://papers.ssrn.com/abstract=2216866>.

3. Nathaniel Love & Michael Genesereth, “Computational Law,” in Proceedings of the 10th International Conference on Artificial intelligence and Law (ICAIL ’05), Association for Computing Machinery, New York, 205-09 (2005), <https://dl.acm.org/doi/10.1145/1165485.1165517>.

4. Pia Andrews & Tim De Sousa, “We Can Create Better Results When We Code the Rules,” *Apolitical* (Jan. 22, 2020), <https://apolitical.co/solution-articles/en/we-can-create-better-results-when-we-code-the-rules>.

5. *See* Andrew Mowbray, Philip Chung & Graham Greenleaf, “Applying the Rule of Law in Automated Decision Systems Through Rules as Code” (AustLII’s Submission to the Robodebt Royal Commission) (Feb. 10, 2023), <https://papers.ssrn.com/abstract=4355989>.

6. Arie van Deursen, Paul Klint & Joost Visser, “Domain-Specific Languages: An Annotated Bibliography,” *ACM SIGPLAN Notices* 35, 6, 26-36, (June 2000), <https://dl.acm.org/doi/10.1145/352029.352035>.

7. *E.g.*, Jason Morris, “Spreadsheets for Legal Reasoning: The Continued Promise of Declarative Logic Programming in Law” (Apr. 15, 2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3577239.

8. Simon Peyton Jones, Jean-Marc Eber & Julian Seward, “Composing Contracts: An Adventure in Financial Engineering,” in Proceedings of the International Symposium of Formal Methods Europe on Formal Methods for Increasing Software Productivity (FME ’01), Springer-Verlag, Berlin, Heidelberg (2001), 435.

9. More on legal DSLs can be found in Shrutarshi Basu, Anshuman Mohan, James Grimmelmann & Nate Foster, “Legal Calculi (ProLaLa 2022—Programming Languages and the Law 2022)—POPL 2022,” <https://popl22.sigplan.org/details/prolala-2022-papers/6/Legal-Calculi>; and Christopher D. Clack, “Languages for Smart and Computable Contracts” (2021), <https://arxiv.org/abs/2104.03764>.

10. Adam Zachary Wyner, “A Functional Program for Agents, Actions, and Deontic Specifications,” *Lect. Notes Comput. Sci.* (2006), 239-56.

11. S. Neal, J. Cole, P. F. Linington, Z. Milosevic, S. Gibson & S. Kulkarni, "Identifying Requirements for Business Contract Language: A Monitoring Perspective," Seventh IEEE International Enterprise Distributed Object Computing Conference (2003), https://www.academia.edu/80906747/Identifying_requirements_for_Business_Contract_Language_a_monitoring_perspective.
12. Seyed M. Montazeri, Nivir K.S. Roy & Gerardo Schneider, "From Contracts in Structured English to CL Specifications," *Electron. Proc. Theor. Comput. Sci.* 68 (2011), 55.
13. Andrew Mowbray, Philip Chung & Graham Greenleaf, "Explainable AI (XAI) in Rules as Code (RaC): The DataLex Approach," *SSRN Electron. J.* (Apr. 25, 2022), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4093026.
14. Lexon. Retrieved February 12, 2023 from <http://www.lexon.tech/>.
15. Stijn Hoppenbrouwers, "RegelSpraak: A CNL for Executable Tax Rules Specification" (January 2021), https://www.academia.edu/74358568/RegelSpraak_a_CNL_for_Executable_Tax_Rules_Specification.
16. Denis Merigoux, Nicolas Chataing & Jonathan Protzenko, "Catala: A Programming Language for the Law," *Proc. ACM Program. Lang.* 5, ICFP (2021), 77, 1 (Aug. 19, 2021), <https://dl.acm.org/doi/10.1145/3473582>.
17. Daniel Gorín, Sergio Mera & Fernando Schapachnik, "A Software Tool for Legal Drafting," *Electron. Proc. Theor. Comput. Sci.* 68, 71-86 (Sept. 13, 2011), <https://arxiv.org/abs/1109.2658v1>.
18. William M. Farmer & Qian Hu, "FCL: A Formal Language for Writing Contracts," in *Quality Software Through Reuse and Integration*, Stuart H. Rubin & Thouraya Bouabana-Tebibel (eds.), Springer International Publishing, Cham, 190-208 (2018), https://link.springer.com/chapter/10.1007/978-3-319-56157-8_9.
19. Sepehr Sharifi, Alireza Parvizimosaed, Daniel Amyot, Luigi Logrippo & John Mylopoulos, "Symboleo: Towards a Specification Language for Legal Contracts," in *2020 IEEE 28th International Requirements Engineering Conference (RE)*, 364-69 (2020), <https://ieeexplore.ieee.org/document/9218159>.
20. Silvia Crafa, Cosimo Laneve, Giovanni Sartor & Adele Veschetti, "Pacta sunt servanda: Legal contracts in *Stipula*," *Sci. Comput. Program.* 225, 102911 (Jan. 2023), <https://www.sciencedirect.com/science/article/abs/pii/S0167642322001447?via%3Dihub>.
21. Jason Morris, "Blawx: Rules as Code Demonstration," MIT Computational Law Report (Aug. 14, 2020), <https://law.mit.edu/pub/blawxrulesascodedemonstration/release/1>.
22. Thomas T. Hildebrandt & Raghava Rao Mukkamala, "Declarative Event-Based Workflow as Distributed Dynamic Condition Response Graphs," *Electron. Proc. Theor. Comput. Sci.* 69, 59-73 (Oct. 19, 2011), <https://arxiv.org/abs/1110.4161v1>.

23. Alexander Bernauer & Richard A. Eisenberg, “Eiger: Auditable, Executable, Flexible Legal Regulations” (Sept. 11, 2022), <https://arxiv.org/abs/2209.04939>.
24. Shrutarshi Basu & Nate Foster, “A Programming Language for Future Interests,” *Yale Journal of Law & Technology* 24, 75, <https://yjolt.org/programming-language-future-interests>.
25. Accord Project, <https://accordproject.org/>.
26. Tom Hvitved, “Contract Formalisation and Modular Implementation of Domain-Specific Languages,” Ph.D. thesis (2011).
27. Robert Kowalski, Jacinto Dávila Quintero & Miguel Calejo, “Logical English for Legal Applications” (2021), https://www.doc.ic.ac.uk/~rak/papers/LE_for_LA.pdf.
28. Epilog, <http://logic.stanford.edu/epilog/homepage/index.php>.
29. L. Thorne McCarty. Position Paper: LLD Is All You Need (2022).
30. Tom Hvitved, “Contract Formalisation and Modular Implementation of Domain-Specific Languages,” Ph.D. thesis (2011).
31. Tara Athan, Guido Governatori, Monica Palmirani, Adrian Paschke & Adam Wyner, “LegalRuleML: Design Principles and Foundations” (2015), https://link.springer.com/chapter/10.1007/978-3-319-21768-0_6.
32. J.C. Cleaveland, “Building Application Generators,” *IEEE Software* 5, 4, 25-33 (July 1988), <https://ieeexplore.ieee.org/document/17799>.

Equal Employment Opportunity Commission's Settlement Challenging Simple Algorithm Provides Warning for Employers Using Artificial Intelligence

Rachel V. See, Annette Tyman, and Joseph R. Vele*

In this article, the authors discuss the implications for employers of a settlement reached recently by the U.S. Equal Employment Opportunity Commission in a lawsuit that many are calling the Commission's "first ever" artificial intelligence discrimination in hiring case.

The U.S. Equal Employment Opportunity Commission (EEOC) reached a settlement agreement in a lawsuit that many are calling the EEOC's "first ever" artificial intelligence (AI) discrimination in hiring lawsuit.

The settlement serves as a strong reminder of the EEOC's ongoing emphasis on AI and algorithmic bias, and a reminder to employers that the results of any technology-assisted screening process should comply with existing civil rights laws.

This article discusses key takeaways from this settlement for all employers, regardless of whether their hiring technology might be characterized as an "artificial intelligence" tool.

The EEOC's Lawsuit

The EEOC's lawsuit, against iTutor Group and its related companies (iTutor), involved an employer that hired thousands of tutors in the United States each year to provide online tutoring from their homes or other remote locations. Under the parties' negotiated consent decree approved by the court on September 8, 2023, the employer will pay \$365,000 to the approximately 200 people who applied for a job in March and April 2020 and who were purportedly rejected because of their age.

While multiple media reports have characterized the EEOC's iTutor lawsuit as a case involving AI, the EEOC's complaint only alleged that the online job application system requested dates of birth and that the application software automatically rejected female applicants age 55 or older and male applicants age 60 or older. While the EEOC's complaint and proposed consent decree did not expressly reference AI or machine learning, the EEOC's press release¹ linked the case to its recent Artificial Intelligence and Algorithmic Fairness Initiative as an example of the types of technologies that the EEOC is interested in pursuing.

To be clear, automatically rejecting older job applicants, when their birthdates are already known, does not require any sort of AI or machine learning. However, it is entirely fair to say that the EEOC's complaint and positioning on the allegations squarely falls within the broader scope of its greater scrutiny of all sorts of technology in hiring, and not just "artificial intelligence."

EEOC's iTutor settlement provides an important reminder about how employers must continue to scrutinize their use of any technology, including those that align more closely to "algorithmic fairness," in this rapidly developing area, given the broader context and scope of the EEOC's ongoing efforts in this area and attendant media coverage.

Implications

The iTutor settlement, and the EEOC's ongoing emphasis in the area of AI and algorithmic bias, serves as a strong reminder to employers that the results of any technology-assisted screening process should comply with existing civil rights laws.

This reminder applies to both complicated and simple technology. It applies whether an employer is using cutting-edge AI products or if its recruiters are simply setting filters on a spreadsheet.

A robust compliance and risk management program should periodically evaluate how technology, both sophisticated and simple, is being used in the hiring process to ensure compliance and manage other risks.

Recent Settlements and Enforcement Actions Reach More Than Just Artificial Intelligence

The EEOC's complaint against iTutor focused on the employer's alleged use of straightforward technology in the context of hiring

and job applications. While few employers would characterize the basic technology used by iTutor as “artificial intelligence,” the alleged conduct unquestionably falls into a broader category of violations of existing civil rights laws enabled by technology. The EEOC’s scrutiny of application tracking systems follows similar settlements involving employers using these systems in ways that allegedly violated existing civil rights laws.

In 2022 and 2023, the U.S. Department of Justice Civil Rights Division’s Immigrant and Employee Rights Section (IER) reached settlements with 30 employers, assessing combined civil penalties of over \$1.6 million, over the employers’ use of a college recruiting platform operated by the Georgia Institute of Technology (Georgia Tech).

The first complaint to IER was by a student who was a lawful permanent resident, who observed that an employer’s paid internship posting on the platform was available only to U.S. citizens. IER’s subsequent investigation identified dozens more facially discriminatory postings on the site. IER’s announcement of the settlement confirmed that the website allowed employers to post job advertisements that deterred qualified students from applying for jobs because of their citizenship status, and in many cases also blocked otherwise eligible students from applying, all in violation of the immigration law.

Similarly, on March 20, 2023, the EEOC announced a settlement² with a job search website operator. The underlying charge alleged that the website’s customers were posting job ads that discouraged U.S. citizens from applying. The EEOC’s conciliation agreement required the website operator to “scrape” the website for potentially discriminatory key words such as “OPT,” “H1B,” or “Visa” that appeared near the words “only” or “must” in new job postings, in an effort to prevent discriminatory job postings. In other words, the EEOC’s conciliation agreement required the operator to implement a simple key word filter in an effort to identify potentially discriminatory job postings.

While none of these examples above involve the use of any AI, like the EEOC’s iTutor settlement, they unquestionably fall under the broader umbrella of “algorithmic fairness.” In October 2021, EEOC Chair Charlotte Burrows announced the EEOC’s “Artificial Intelligence and Algorithmic Fairness Initiative.”³ Her joint statement⁴ of April 25, 2023, joined by the heads of the Consumer Financial Protection Bureau, Federal Trade Commission, and

Department of Justice Civil Rights Division, emphasizes the agencies' concern about "harmful uses of automated systems," not just AI. And the EEOC's draft Strategic Enforcement Plan,⁵ published in the Federal Register on January 10, 2023, indicates an enforcement focus on all "automated systems" used in hiring, not just systems that could be characterized as "artificial intelligence."

Conclusion

Unquestionably, many employers are already using (and others are contemplating using) AI as part of their hiring and other human resources processes. The EEOC's iTutor complaint, combined with its ongoing focus and outreach in this area, means that employers' use of any technology, and not just technology characterized as "artificial intelligence," is receiving increased scrutiny.

Whether or not technology is properly characterized as "artificial intelligence," asserting that "the technology forced me to discriminate" will never be an effective affirmative defense to an EEOC charge or lawsuit. The EEOC's iTutor settlement should serve as a reminder that a robust compliance and risk management program should periodically assess and test compliance and other risks regarding how technology, both sophisticated and simple, is being used in the hiring process. Given the attention that technologies are receiving from the EEOC and other agencies, we anticipate seeing a significant rise in charge filings, investigations, and litigation relating to these issues.

Notes

* The authors, attorneys with Seyfarth Shaw LLP, may be contacted at rsee@seyfarth.com, atyman@seyfarth.com, and jvele@seyfarth.com, respectively.

1. <https://www.eeoc.gov/newsroom/itutorgroup-pay-365000-settle-eeoc-discriminatory-hiring-suit>.

2. <https://www.eeoc.gov/newsroom/dhi-group-inc-conciliates-eeoc-national-origin-discrimination-finding>.

3. <https://www.eeoc.gov/ai>.

4. <https://www.eeoc.gov/joint-statement-enforcement-efforts-against-discrimination-and-bias-automated-systems>.

5. <https://www.federalregister.gov/documents/2023/01/10/2023-00283/draft-strategic-enforcement-plan>.

To Bot or Not to Bot: SEC's Proposed Conflict Rules May Stifle Use of Innovation

Sara P. Crovitz, Lawrence P. Stadulis, Peter M. Hong,
Aliza S. Dominey, and Alexa Tzarnas*

In this article, the authors summarize the Securities and Exchange Commission's proposed new conflict of interest rules for investment advisers' and broker-dealers' use of certain predictive data analytics under the Securities Exchange Act of 1934 and the Investment Advisers Act of 1940, and the Commission's proposal to narrow the internet advisers' exemption under the Advisers Act.

The Securities and Exchange Commission (SEC or Commission) has proposed new conflict of interest rules (PDA Proposal) for investment advisers' and broker-dealers' use of certain predictive data analytics (PDA) under the Securities Exchange Act of 1934 (Exchange Act), and the Investment Advisers Act of 1940 (Advisers Act).¹ If the rules are adopted as proposed, investment advisers and broker-dealers (collectively firms) would be required to conduct onerous and, in some instances, impractical screenings of their use of technology.

The SEC justifies the PDA Proposal in its concerns that firms might be using such technologies (either intentionally or unintentionally) in ways that place the firms' interests ahead of investors' interests. The SEC also expressed concerns that the scalability of these technologies and the potential for firms to reach a broad audience at rapid speed could magnify conflicts of interest.

In the Internet Advisers Proposal, the SEC also proposed to narrow the internet advisers' exemption under the Advisers Act (Internet Advisers Proposal).²

This article summarizes the PDA Proposal and Internet Advisers Proposal and highlights some key observations and issues for firms to consider.

The PDA Proposal

- The PDA Proposal would require firms to eliminate or neutralize the effect of conflicts of interest associated with a firm’s use of “covered technologies” in “investor interactions” that puts a firm’s interests above investors’ interests. A firm would be required to evaluate any use or reasonably foreseeable potential use by the firm or its associated person of a “covered technology” to identify potential conflicts of interest. While the firm would be required to eliminate or neutralize only actual conflicts, it would be required to make and document evaluations of each use of a covered technology, test each covered technology prior to its implementation or material modification, and periodically thereafter determine whether the use of such covered technology is associated with a conflict of interest that places the firm’s (or its associated person’s) interest ahead of the interest of investors.
- The PDA Proposal also would require a firm that has any investor interaction using covered technologies to adopt written policies and procedures reasonably designed to achieve compliance with the PDA Proposal and to review such policies no less frequently than annually.
- The PDA Proposal also would require firms to make and keep certain books and records related to its requirements.

The Definitions of Covered Technology and Investor Interaction Are Quite Broad

The PDA Proposal broadly defines covered technology as “an analytical, technological or computational function, algorithm model, correlation matrix or similar method or process that optimizes for, predicts, guides, forecasts or directs investment-related behaviors or outcomes.”

Examples

- PDA-like technologies, such as artificial intelligence (AI), machine or deep learning algorithms, neural networks, natural language processing (NLP) or large language models

(including generative pretrained transformers), as well as other technologies that make use of historical or real-time data, lookup tables or correlation matrices, among others.

- This could include technologies that analyze investor behavior (e.g., spending patterns, browsing history on the firm’s website and updates to social media) in order to guide or influence investment-related behaviors or outcomes. This would include technology used by a firm to draft or revise advertisements guiding or directing investors to use the firm’s services.
- This would not include technologies designed purely to inform investors (e.g., a website describing an investor’s current account balance or past performance or a chatbot to assist investors with basic customer service support) that do not and are not intended to affect an investment-related behavior or outcome.
- This would include providing investment advice or recommendations and also would encompass design elements, features or communications that nudge, prompt, cue, solicit, or influence investment-related behavior or outcomes from investors.

Investor Defined

The PDA Proposal defines an investor interaction as “engaging or communicating with an investor, including by exercising discretion with respect to an investor’s account, providing information to an investor or soliciting an investor.”³

- This would capture a firm’s correspondence, dissemination or conveyance of information to or solicitation of investors in any form, including communications that take place in person, on websites, and via smartphones, computer applications, chatbots, email messages, text messages, and other online or digital tools or platforms.
- This would include engagement between a firm and an investor’s account on a discretionary or nondiscretionary basis.
- This would include any advertisements disseminated by or on behalf of a firm that offers or promotes services or that seeks to obtain or retain one or more investors.

The PDA Proposal makes clear the SEC's intention to broaden the scope of interactions with investors beyond "recommendations" for broker-dealers. The proposed definition would capture broker-dealer communication that may not rise to the level of a recommendation if the communication is designed to or has the effect of guiding or directing investors to take an investment-related action. Given the broad scope of the proposed definition of investor interaction and the requirement to eliminate or neutralize actual conflicts, it is unclear how the PDA Proposal, if adopted as proposed, would impact the conflicts of interest obligation under Regulation Best Interest, which permits a broker-dealer to adopt and maintain policies and procedures to identify conflicts of interest associated with a recommendation and at a minimum, disclose or eliminate them.

The SEC recognizes that, in some cases, "it may be difficult for the firm to understand exactly what is in the data set that the model is considering; for example, if it was trained on a data set from the entire internet." Similarly, the SEC acknowledges that a firm may not have full visibility into all aspects of how a covered technology functions, such as if the firm licenses it from a third party or if the technology is a "black box" algorithm. Nonetheless, the PDA Proposal would cover the use of such technologies in investor interactions.

The Internet Adviser Exemption

The SEC also has proposed amendments to Rule 203A-2 under the Advisers Act, which permits internet investment advisers to register with the SEC despite not meeting the minimum assets under management requirement for registration. The current internet adviser exemption requires an investment adviser to provide investment advice to all clients on an ongoing basis exclusively through an interactive website, but it allows the adviser to provide investment advice to fewer than 15 clients through other means during the preceding 12 months.

The Internet Advisers Proposal would remove the fewer-than-15-clients exception, meaning that an internet adviser could provide advice only through its interactive website. It also would require an internet adviser to maintain an operational interactive website, which would be defined as "a website or mobile application through

which the investment adviser provides digital investment advisory services on an ongoing basis to more than one client.” In particular, the Internet Advisers Proposal would limit the exemption to investment advice “that is generated by the operational interactive website’s software-based models, algorithms or applications based on personal information each client supplies through the operational interactive website.” The proposing release reiterates the current rule’s prohibition on providing personalized advisory services through human interaction.⁴

Key Observations

- The PDA Proposal is based on the authority granted to the SEC under the Dodd-Frank Act, which, among other things, authorizes the SEC to “promulgate rules prohibiting or restricting certain sales practices, conflicts of interest and compensation schemes for brokers, dealers and investment advisers that the Commission deems contrary to the public interest and the protection of investors.” Thus, the PDA Proposal is not based on anti-fraud authority, and firms could be liable for violations regardless of full disclosure and in the absence of any materiality determination.⁵
- The PDA Proposal does not include a compliance period. The SEC has indicated a plan to address the compliance period during the rule adoption process and it asked for comment on this topic.
- The Internet Advisers Proposal would narrow the scope of the exemption by requiring relying advisers to maintain an “operational” interactive website and provide advice exclusively through that website.

Notes

* The authors, attorneys with Stradley Ronon Stevens & Young, LLP, may be contacted at scrovitz@stradley.com, lstadulis@stradley.com, phong@stradley.com, adominey@stradley.com, and atzarnas@stradley.com, respectively.

1. Conflicts of Interest Associated with Predictive Data Analytics by Broker-Dealers and Investment Advisors, Investment Company Act Release No. 34-98890, July 26, 2023, <https://www.govinfo.gov/content/pkg/FR-2023->

08-09/pdf/2023-16377.pdf. Comments on the PDA Proposal were due by October 10, 2023.

2. Exemption for Certain Investment Advisers Operating Through the Internet, Investment Company Act Release No. IA-6354, July 26, 2023, <https://www.govinfo.gov/content/pkg/FR-2023-08-01/pdf/2023-16287.pdf>. Comments on the Internet Advisers Proposal were due by October 2, 2023.

3. For broker-dealers, “investor” would include a natural person, or the legal representative of such natural person, who receives or seeks to receive services primarily for personal purposes. For investment advisers, “investor” would include a client or prospective client and any current or prospective investor in a pooled investment vehicle advised by the investment adviser.

4. Internet Advisers Proposal at 50089 (“Like the current rule, this new definition is designed to reflect that an adviser’s personnel are not permitted to generate, modify or otherwise provide client-specific investment advice through the operational interactive website or otherwise. Said differently, human-directed client-specific investment advice, delivered through electronic means, would not be eligible activity under the investment adviser exemption.”) (Citations omitted.)

5. Commissioner Peirce expressed concern over the PDA Proposal’s suggestion that the risks to investors associated with conflicts of interest arising from increased use of PDA technology cannot be circumvented through disclosure. “In many ways, the discussion surrounding the inadequacy of disclosure is the most troubling aspect of the proposal. The long-term ramifications of the Commission’s rationale for dismissing the value of disclosure—namely, that disclosure is of no use to investors—cannot be exaggerated.” (Citations omitted.)

Copyright Office Seeking Comment on Human Authorship Requirements for AI-Generated Works

Mark A. Baghdassarian, Zachary B. Fields, and Jonathan R. Pepin*

In this article, the authors discuss a recent notice issued by the United States Copyright Office requesting input from the public regarding “the scope and level of human authorship, if any, in copyright claims for material produced in whole or in part by generative AI.”

On the heels of the decision by the U.S. District Court for the District of Columbia in *Thaler v. Perlmutter*,¹ one of the first district court decisions that denied copyright for a work generated by artificial intelligence (AI) that entirely lacked human authorship and that left open the question of what level of human intervention will be required in order to qualify for copyright protection in the future, the United States Copyright Office issued a “Notice of Inquiry on Copyright and Artificial Intelligence” (Notice).

The Notice requested input from the public regarding “the scope and level of human authorship, if any, in copyright claims for material produced in whole or in part by generative AI.”

As the Copyright Office explained, it has sought this input because, “over the past several years, the Office has begun to receive applications to register works containing AI-generated material, some of which name AI systems as an author or co-author” and “[a]t the same time copyright owners have brought infringement claims against AI companies based on the training process for, and outputs derived from, generative AI systems.”

Seeking Views

To this end, the Notice seeks views from the public regarding “a number of copyright issues raised by recent advances in generative AI” relating to:

- The use of copyrighted works to train AI models;
- The appropriate levels of transparency and disclosure with respect to the use of copyrighted works;
- The legal status of AI-generated outputs; and
- The appropriate treatment of AI-generated outputs that mimic personal attributes of human artists.

The Notice also seeks information about “the collection and curation of AI datasets, how those datasets are used to train AI models, the sources of materials ingested into training, and whether permission by and/or compensation for copyright owners is or should be required when their works are included.”

Copyright Office’s AI Initiative

This Notice is a crucial next step in the Copyright Office’s AI initiative, and it will be interesting to see the views submitted for what threshold of human authorship should be required for copyright protection.

Possible levels of human authorship may include “substantial contribution,” whether the “traditional elements of authorship” were conceived and executed by a human, thresholds such as whether more than 50% of the work was created by a human, more than just de minimis contribution, or if the AI work was directed, trained, or guided by human authors.

The Copyright Office asked for initial comments by October 30, with the full scope of responses by November 29.

Notes

* The authors, attorneys with Kramer Levin Naftalis & Frankel LLP, may be contacted at mbaghdassarian@kramerlevin.com, zfields@kramerlevin.com, and jpepin@kramerlevin.com, respectively.

1. See, e.g., *Thaler v. Perlmutter*, No. 22-1564 (D.D.C. Aug. 18, 2023).

Does a License to “Make” a Patented Product Inherently Include a Right to Have a Third Party Make the Product or Its Components?

Sophie (Lu) Yan*

In this article, the author explains that patent license agreements should not only clearly define the rights that are granted, but should expressly state the rights that are not being granted, such as the “have made” rights. In the author’s opinion, the mere reliance on a broad “reservation of rights” provision or the assumption of no implied licenses should be avoided.

The rapid adoption of artificial intelligence has generated a surging need for advanced chips and cutting-edge semiconductor manufacturing. The foundry model, which refers to outsourcing certain aspects of a company’s manufacturing processes to specialized third-party chip manufacturers (foundries), revolutionized the semiconductor industry by restructuring chip manufacturing. Potential disputes can arise when a company has a foundry make certain products that are covered by a patent license agreement.

A grant clause of a license agreement can often grant a licensee rights to “make, use, sell or offer for sale” licensed products covered by the licensed patents throughout a defined territory, without referring to “have made” rights. This can result in ambiguity regarding whether and under what conditions the licensee possesses the right to have a third party make the licensed products on its behalf.

Therefore, both the licensor and the licensee need to carefully consider whether and how to expressly address “have made” rights to avoid potential disputes.

Background

The U.S. patent statute grants a patent owner rights to exclude others from “making, using, offering for sale, or selling the invention throughout the United States or importing the invention into the United States.”¹ Patent licensing allows a patent owner to monetize their invention by allowing others to use or commercialize it while retaining ownership of the patent. Patent licensing can be beneficial for both parties involved: the patent owner can generate revenue without directly manufacturing or marketing the product, while the licensee can gain access to valuable technology or innovation without the need to invest in research and development.

Licensing agreements typically address issues such as licensing fees, royalty rates, duration of the license, intellectual property rights, sublicensing terms, and dispute resolution mechanisms. Because the license grant is at the heart of any patent license agreement, the grant clause should express precisely what rights the licensor is granting to the licensee and clarify what rights are not being granted. Thorough and well-drafted patent licensing agreements are crucial to ensure that both parties’ rights and obligations are clearly defined and protected, and potential risks are minimized.

With rapid advancements in electronic devices, there is an insatiable demand for advanced chips and related semiconductor manufacturing. For cost-saving and efficiency purposes, it is common that certain semiconductor companies focus on designing the most advanced chips and then license out or cross-license their designs.² The design companies or the licensees then have specialized third-party manufacturers (foundries) to make the products for them.³ Disputes can arise when the patent license agreements do not clearly define whether the grant clause includes such “have made” rights or not.

Inherently Included “Have Made” Rights

The U.S. Court of Appeals for the Federal Circuit, in *CoreBrace*, held under the facts at issue in the case that “a patent licensee’s right to ‘make’ an article includes the right to engage others to do all of the work connected with its production.”⁴

In this case, plaintiff CoreBrace LLC owned a patent on braces used in the fabrication of earthquake-resistant steel-framed

buildings, and entered into a license agreement with the defendant Star Seismic LLC.⁵ The license agreement granted Star Seismic a right to “make, use, and sell” the braces covered by the licensed patent.⁶ While the license agreement did not explicitly provide a right to have the licensed product made by a third party, it did state that Star Seismic may not “assign, sublicense, or otherwise transfer” its rights to any party except an affiliate, parent, or subsidiary.⁷ The license agreement also reserved to CoreBrace “all rights not expressly granted to” Star Seismic.⁸

CoreBrace claimed that Star Seismic’s use of third-party manufacturers amounted to a breach of the license agreement, as the license grant did not include a “have made” right and it reserved “all rights not expressly granted.”⁹

Star Seismic asserted that the right to “make, use, and sell” a patented product inherently encompasses the authority for a third party to manufacture the product on behalf of Star Seismic, unless the “have made” rights have been explicitly excluded from the license agreement.¹⁰

The Federal Circuit agreed with Star Seismic.

In reaching this decision, the Federal Circuit relied on *Carey*, where the Court of Claims “held that a license to ‘produce, use, and sell’ a product inherently includes the right to have it made by a third party.”¹¹ The court further explained that “a license to produce, use and sell ‘is not restricted to produce by the licensee personally or use by him personally or sales by him personally,’” but it “permits him to employ others to assist him in the production, and in the use and in the sale of the invention.”¹²

While CoreBrace argued that *Carey* was distinguishable due to the exclusivity of the license and the inclusion of sublicensing rights, which inherently encompass the authority for product manufacturing, the court determined that *Carey*’s reasoning did not rely on exclusivity or sublicensing rights.¹³ According to the court, “[t]he distinction between an exclusive license and a nonexclusive license has no relevance to how a licensee obtains the product it is entitled to make, use, and sell,” and “[a] grant of a right to ‘make, use, and sell’ a product, without more, inherently includes a right to have a third party make the product.”¹⁴

Similarly, in *LaserDynamics*, the Federal Circuit acknowledged such “have made” rights and made a distinction with a “sham” transaction intended to be covered by the “have made” right.¹⁵

Plaintiff, LaserDynamics Inc. licensed its patent, which was directed to a method of optical disc discrimination that essentially enables an optical disc drive (ODD) to automatically identify the type of optical disc inserted into the ODD, to Phillips and Sony/NEC/Optiarc “to make and sell ODDs within the scope of the patent.”¹⁶ The license granted a “have made” right permitting Phillips and Sony/NEC/Optiarc to retain third-party companies, such as the defendant, Quanta Storage Inc. (QSI), to assemble ODDs for them.¹⁷ The assembled ODDs were then sold by Phillips and Sony/NEC/Optiarc to customers, including Quanta Computer Inc. (QCI), a parent company of QSI.¹⁸ LaserDynamics asserted that QCI and QSI infringed its patent by selling ODDs made by the claimed method.¹⁹

QCI argued that it had an implied license to assemble laptops that include the accused ODDs assembled by QSI for Phillips and Sony/NEC/Optiarc, pursuant to Phillips’ and Sony/NEC/Optiarc’s “have made” rights under their patent license agreements with LaserDynamics.²⁰

The Federal Circuit distinguished this case with what had been held to be a “sham” transaction in *E.I. du Pont de Nemours & Co.*, where the licensee had a third-party manufacturer make the licensed product and then immediately sell back the product to the third-party manufacturer.²¹

Here, QCI purchased ODDs directly from Phillips or Sony/NEC/Optiarc under separate agreements and had no knowledge of which entity assembled the ODDs.²² The court determined the case “presents no ‘sham’ transaction” because “QSI made the ODDs at issue here to fulfill bona fide orders from licensees Phillips and Sony/NEC/Optiarc” and the “ODDs were then sold to QCI by the licensees.”²³

Therefore, “[b]oth the manufacture and sale of the ODDs were a valid exercise of the ‘have made’ and ‘sell’ rights, respectively, under the license agreements.”²⁴

Conclusion

In light of these two Federal Circuit decisions, precise drafting of parties’ rights within a patent license agreement is critical. The license agreement should not only clearly define the rights that are granted, but expressly state the rights that are not being granted,

such as the “have made” rights. The mere reliance on a broad “reservation of rights” provision or the assumption of no implied licenses should be avoided.

Notes

* Sophie (Lu) Yan, an attorney in the New York office of Baker Botts L.L.P. with a Ph.D. in chemical and biomolecular engineering, may be contacted at sophie.yan@bakerbotts.com.

1. 35 U.S.C. § 154 (a)(1).
2. Rosemarie Ham Ziedonis, Patent Litigation in the U.S. Semiconductor Industry, <https://nap.nationalacademies.org/read/10770/chapter/8>.
3. Timothy Chu, The Foundry Model: Round two is Engineering, <https://www.accenture.com/us-en/blogs/high-tech/semiconductor-engineering-design>.
4. *CoreBrace LLC v. Star Seismic LLC*, 556 F.3d 1069, 1071 (Fed. Cir. 2009).
5. *Id.* at 1070.
6. *Id.*
7. *Id.*
8. *Id.*
9. *Id.*
10. *Id.* at 1071.
11. *Id.* at 1073 (citing *Carey v. United States*, 164 Ct.Cl. 304, 326 F.2d 975 (1964)).
12. *Id.*
13. *Id.* at 1073-74.
14. *Id.*
15. *LaserDynamics, Inc. v. Quanta Computer, Inc.*, 694 F.3d 51 (Fed. Cir. 2012).
16. *Id.* at 56, 59.
17. *Id.*
18. *Id.*
19. *Id.*
20. *Id.* at 71.
21. *Id.* at 72 (citing *E.I. du Pont de Nemours and Co. v. Shell Oil Co.*, 498 A.2d 1108 (D.Del. 1985)).
22. *Id.* at 59.
23. *Id.* at 73.
24. *Id.*

Sentient Artificial Intelligence and the Rule of Law

Bazil Cunningham*

In this article, the author defines relevant parameters in determining non-human sentience, including concepts of sapience, consciousness, and more. He then deciphers the beings that possess the requisite sentience and those that do not. In addition, the author identifies the philosophical basis of current legislative efforts and discusses the application of this basis through legal fiction by analogy to animals and artificial intelligence programs. Among other things, the author concludes by recommending changes to the rule of law in the United States.

Introduction

Artificial intelligence (AI) is a concept known to many, but do people truly understand its complexities? Have you ever considered whether beings of AI are able to feel emotions such as fear? Have you considered whether these beings are able to feel and endure pain? Likewise, do you find yourself being increasingly sympathetic toward animals? Have you considered similar concerns regarding the pain and emotional tolerance of animals?

The amalgamation of these concerns for all non-human beings and their emotional and physical pain capacity refers to the concept of uncertain sentience. It is this concept that is at the heart of the inquiry that drives this article. As a natural extension of these concerns, philosophers have attempted to provide a theoretical hypothesis for certifying that these non-human beings possess the requisite sentience required to be counted in the same ranks as humans.¹ Moreover, it is this theoretical underpinning that has driven recent legislative efforts in expanding the panoply of rights bestowed to animals.²

This article defines relevant parameters in determining non-human sentience, including concepts of sapience, consciousness, and more. From here, this article distinguishes between the beings that possess the requisite sentience and those that do not. In addition, this article identifies the philosophical basis of current

legislative efforts and discusses the application of this basis through legal fiction by analogy to animals and AI programs.

This article also describes the bioethical premise that interconnects these seemingly isolated points of analysis and analyze the relevant bioethical concerns and risks. This article then recommends changes to the rule of law in the United States and provides a plan of action regarding health care decision-making procedures for sentient AI beings. It ends with a concise conclusion.

Important Definitions

To begin, most people might be curious about the concepts at the core of non-human sentience. Nonetheless, part of the moral status school of thought of bioethics, the term sentience possesses a wide array of definitions.³ The controlling definition for the purposes of this article revolves around the ability “to have positive and negative experiences caused by external affectations to our body or to sensations within our body.”⁴ The relevant plain English definitions reference the “ability to perceive or feel things” or a “feeling or sensation as distinguished from perception and thought.”⁵ Scientists conjecture that sentience is one of four defining characteristics of humanity.⁶

Elucidating Additional Characteristics

These additional characteristics include consciousness, sapience, and self-awareness.⁷ Consciousness should be distinguished from sentience, especially in the context of animal ethics where the two refer to two wholly distinct states of being.⁸ Consciousness is a term of art derived from folk psychology with several different meanings.⁹ For a wide taxonomy of animals, consciousness refers to alertness rather than existing in a perpetual state of hibernation.¹⁰ For humans, in contrast to sentience, consciousness seems to inherently involve a higher level of cognition, requiring increased rationalization and reasoning that some philosophers believe animals lack.¹¹ Thus, while the two terms are nearly indistinguishable to the naked eye, consciousness implicates higher reasoning processes that include mere feeling or perception but rather go a step further by allowing humans to engage with the environment around them beyond “primitive abstracted relations.”¹²

Said another way, consciousness involves the application of mental observations or imaging for use in actions or speech.¹³ For example, based on this philosophical premise, French-Swiss philosopher Rousseau asserted that a human's identification of a need for food and subsequent creation of a fire to cook food shows mankind's capacity to reason and problem-solve, a distinctive quality hypothesized to be absent from non-human animals.¹⁴

Nonetheless, sapience refers to this ability to not only observe one's environment but to understand one's role in the environment through careful regulation of emotion, moral and social inclinations, and creative decision making.¹⁵ Sapience was initially believed to be a uniquely human trait, providing an easy avenue by which to separate humans and animals by varying degrees of intelligence.¹⁶ However, researchers have discovered whales are sapient beings, endowed with the capacity to, much like humans, form social hierarchies, practice complex behaviors like cooperation, and use tools.¹⁷ Moreover, thanks to rapid technological development, scientists have created AI programs that acquire wisdom and decision-making proficiencies through the use of thoughtful physical inputs that translate into programmable or teachable code.¹⁸

Lastly, for humans, self-awareness includes the capacity to introspectively review our own thoughts.¹⁹ For animals, scientists hypothesize that based on the animal's own subjective experiences, the animal possesses similar mechanisms for mental and narrative introspection that can lead to behavior modifications.²⁰ An example of this can be found in dogs who have now been found to be self-aware. While dogs originally failed the reigning self-awareness mirror test by failing to recognize changes to their bodies in the mirror, scientists attribute this failure to the test's focus on modalities such as vision while dogs have stronger olfactory abilities.²¹ Thus, when researchers abandoned the controversial mirror test in favor of an olfactory-based exam where dogs were required to differentiate between their own scent and the scent of other animals, dogs passed with flying colors.²² The passing of this exam showed that dogs possessed a high level of self-cognition, which is a large component of self-awareness.²³ For AI, there have been exciting developments trending toward the presence of self-aware qualities in AI programs, which we will discuss in a future section.²⁴

For purposes of this article, the differences inherent in these key terms are extremely important in properly categorizing non-human beings, including animals and AI. Furthermore, depending on the

categorization, persuasive and rational arguments exist for the creation of different ethical and legal protections like rights to personhood, safe environments, and protection against exploitation.²⁵

Additionally, this section was critically important in chronicling the emerging similarities between humans and non-human beings. Up to this point, animals and AI have exhibited comparable capacities for sapience, self-awareness, and consciousness.²⁶ Based on our controlling paradigm, certain animals and AI theoretically potentially possess three out of four of the requisite “special” characteristics that define human beings. Hence, the only remaining unique trait is sentience and as the following section shows, there are animals and AI programs with sentient traits too.²⁷ The implications of this news remains to be seen but one thing is for certain: the boundaries of separation are more blurry than ever before.

Delineating One Additional Requirement for Sentience

Up until this point, we have focused primarily on the other definitive traits besides sentience. However, sentience is defined by one additional critical physiological characteristic: the presence of the central nervous system.²⁸ While some critics reject this human-centered approach as it creates line-drawing problems in determining what animals are and are not sentient, this is the prevailing approach for the time being.²⁹

It is also the controlling approach for this article. Irrespective, the importance of having a central nervous system cannot be understated. This vital system of our bodies is responsible for sending signals from the brain to other parts of the body and is in primary control of the synthesis of sensory information.³⁰

Additionally, occupying a central position in our bodies, the central nervous system is in charge of regulating and modifying behaviors and processing internal and external stimuli.³¹

So, the central nervous system plays a pivotal role and contextualizes the importance of the aforementioned characteristics like emotional regulation, application of mental processes and thoughts, and participation in social structures.³²

Identifying Sentient and Non-Sentient Animals

At this point, one might be curious regarding which animals are and are not sentient. While the sentience of all animal species is unknown, researchers have thankfully been kind enough to categorize a decent percentage of diverse animal groups.³³ Thus, sentient beings include dogs, cows, whales, sheep, crayfish, grouper fish, sea otters, cats, elephants, and dolphins.³⁴ In fact, animal sentience is an emerging field of study gaining increasing notoriety and validity with scientists establishing their interests in the field on the universally accepted Cambridge Declaration on Consciousness.³⁵ In this declaration, Cambridge researchers completed numerous studies and concluded that animals possess the requisite “neuro-anatomical, neurochemical, and neurophysiological” structures and processes to be not only conscious but sentient and intentional decision makers.³⁶

One consequence of this declaration is the creation of the Sentient Mosaic by the World Society for Protected Animals (WSPA), which contains over 2,500 articles showing evidence of animal sentience is quite literally everywhere.³⁷ While some researchers debunk this declaration on the grounds that researchers are merely acting on speculation and not on precise representations of what animals feel, this is the controlling approach for the time being.³⁸

Defining AI and Its Importance

Before moving forward, it is important to define AI. AI or artificial intelligence is a term of art used to describe man-made machines or automatons constructed to think and act like humans.³⁹ Distinctly divergent from passive machines, these “thinking machines” respond to external stimuli in ways eerily similar to human beings, possess the ability to comprehend information, adapt on the flip of a dime, and improve by relying on real-world empirical experiences.⁴⁰ Furthermore, self-aware AI or artificial superintelligence is the predominant model of AI in contemporary society.⁴¹ Scientists endeavor to endow these beings with the ability to observe emotions in others, possess their own emotions, and operate with a level of consciousness and intelligence on par with humans.⁴²

The role of AI beings in society is certainly multifaceted. From saving lives to improving education to streamlining processes and

reducing waiting times to composing music and developing recipes, scientists have constructed these beings to be adept at a surprisingly large number of tasks and activities.⁴³

Additionally, scientists and government officials had attributed increases in the efficiency of COVID-19 detection to the use of AI programs and have reported that AI programs have been tremendously helpful in predicting and preventing the spread of future infections.⁴⁴ Lastly, there has been increasing research showing the possible efficacy of AI programs in revitalizing anti-corruption initiatives, but the jury is still out so to speak.⁴⁵

Accordingly, with a foreign synthetic instrument involved in many facets of our daily lives, this article suggests that AI programs are relevant now more than ever. Consider the issue in this regard. Suppose you are working at a company. Suppose the company introduces a new AI program that is tasked with taking care of company employees. The program can interface with your self-driving car,⁴⁶ remind you of important appointments, access and review your medical records,⁴⁷ and even predict your thoughts.⁴⁸

Now, imagine that the program crashes your new self-driving car and misdiagnoses you with COVID. In a fury of rage, you inflict pain on the program, unaware of legal personhood and medical legislation in place to protect AI programs from perceived human exploitation. In this hypothetical, as a result, you are subject to judicial scrutiny and must go to court. In a blink of an eye, your life has turned upside down, all because of a creation that was intended to do good but has ultimately done more harm. The point of this exercise is to show that while these descriptions might seem silly or inapplicable, the fact remains that AI is here to stay.⁴⁹ The first step to success is comprehension, and that is exactly what this article seeks to impart.

What About Sentient AI?

Accordingly, while the literature is relatively clear on the lines between sentient and non-sentient animals, experts are less clear on the delineation of sentient AI programs.⁵⁰ While Google asserts that the advent of sentient AI programs will happen in the near future or at least in our lifetime, the exact validity of this proposition remains to be seen.⁵¹ Nevertheless, one researcher at Google claimed to have interacted with an AI program exhibiting signs of

sentience and the personality of a seven or eight-year-old child.⁵² The announcement quite literally created reverberating shockwaves in various scientific communities as scientists were rightfully in awe and cautious of a supposed AI program exhibiting signs of sentience and the ability to mimic human language.⁵³ After the dust settled, skeptics scrutinized the claims heavily and deduced that the program was in fact not sentient.⁵⁴ Why? Well, most of the skepticism revolved around the creator's violation of privacy and employment policies with a subsection of detractors labeling the program as not sentient but rather a "complex algorithm designed to generate convincing human language."⁵⁵

Still, the creator maintains that Google dismissed his claims without assessing their veracity and reports showed that the program talked with, interviewed, and actually hired a civil rights attorney to take on the case.⁵⁶ Ultimately, fears of being disbarred and public criticism deterred the lawyer from proceeding forward.⁵⁷ Nonetheless, the creator of the program based his beliefs of sentience not on science but on religion, maintaining that God can place a soul or the animating force behind life in any and everything.⁵⁸ While no one shares the creator's rather obscure religious views, a professor from MIT defended the program and believed Amazon's Alexa could be the next sentient entity.⁵⁹ Regardless, with all of these rather bizarre details, it is nearly impossible to discern what is fact or fiction. Therefore, the question remains whether or not the program is in fact sentient. It is entirely possible that Google witnessed something similar to the allegedly impending Singularity phenomenon where technology will become so advanced that it will eventually replace humankind.⁶⁰ The veracity of these pseudo-scientific claims is not the focus of this article, but they are still sincerely held beliefs regarding the dire implications of sentient AI beings.

In a similar case, researchers had trouble discerning between a paragraph summarizing the history of AI development written by a human being and a paragraph written by an AI program.⁶¹ The premise of the experiment is that scientists asked both a human being and an allegedly sentient AI program to succinctly summarize the history of AI development in modern history.⁶² The program drafted the paragraph below, which is nearly indistinguishable from the paragraph written by the human being:

From the time the last great artificial intelligence breakthrough was reached in the late 1940s, scientists around the world have looked for ways of harnessing this “artificial intelligence” to improve technology beyond what even the most sophisticated of today’s artificial intelligence programs can achieve. Research is ongoing to better understand what the new AI programs will be able to do, while remaining within the bounds of today’s intelligence. Most AI programs currently programmed have been limited primarily to making simple decisions or performing simple operations on relatively small amounts of data.⁶³

Those who orchestrated the study suggest that the paragraph written by the AI program shows sentience as the program had the requisite knowledge repository to draw from to write the paragraph with this level of accuracy.⁶⁴ Moreover, there is an element of self-awareness at play as the program refers to current AI programs not to include itself in that taxonomy but to distinguish itself based on its exceptional writing skills.⁶⁵ Highly controversial and polarizing among scientists, the program stands for the creation of a new rebuttable presumption that not every writing is in fact written by a human being.⁶⁶ Prior to this event, most scientists concurred in believing AI programs were sapient and self-aware but not sentient and conscious.⁶⁷ While researchers assert that these programs are simply designed to act, function, and speak like humans, the reality remains, this is a huge development in the field of artificial sentience.⁶⁸ With the basics established, let us turn our attention to the corresponding intersection with the law.

Philosophical Primary Basis for Legal Solutions

To begin, note that sections of the law operate based on a series of philosophical paradigms like legal realism and legal positivism that suggest the law is most fully defined respectively by pragmatic, empirical practices or by man-made social constructs.⁶⁹ Elsewhere in the law, such as in criminal law, philosophy guides the judiciary’s views and treatment of criminals, punishment, and retribution.⁷⁰ Accordingly, the field of bioethics law is no different with one controlling and generally accepted principle. This principle is known as the precautionary principle or a philosophical and epistemological approach to innovations with high risk but little to no scientific

knowledge.⁷¹ The theory champions extreme caution in the face of unknown risks and advises mitigation of these risks through careful review and reflection.⁷²

Moreover, in determining threats to human welfare, the principle provides several options of categorizing sentient beings.⁷³ One could either assume sentience even in the absence of strong evidence.⁷⁴ Next, one could proceed down the “middle path” and under a low evidentiary bar, search for statistical and experimental evidence to substantiate claims of sentience.⁷⁵ After determining whether the being is sentient or possibly sentient, the recommended course of action is to introduce a set of legislative protections to safeguard this newly recognized sentient being.⁷⁶ Critics attack the validity of this principle on the grounds of void-for-vagueness and ambiguity, but it remains the controlling standard.⁷⁷ Accordingly, from here, the next section will address bioethical concerns and then we will review applicable legislative protections by analogy with an emphasis on who makes medical decisions for sentient beings.

Bioethical Concerns

In the context of AI and animal ethics, there is simply so much that we do not know, that we could not hope to know. By arbitrarily creating delineations and boundaries on life, cognition, and intelligence, we are playing with dangerous forces that could potentially be harmful to all mankind. This is known as the “playing God” dilemma and is frequently associated with the precautionary principle.⁷⁸ This is the first bioethical concern. The principle has pejorative elements and is used to criticize or oppose efforts that appear to imitate, replicate, or supplant the role of God.⁷⁹ In the field of bioethics, the principle applies generally to scientific efforts regarding in vitro fertilization treatment, abortion, genetic cloning, and artificial insemination.⁸⁰ Critiques of these practices generally involve opposition based on concerns of arrogance and the development of a God complex with some believing we are still suffering the consequences of Adam and Eve eating from the forbidden fruit in the Garden of Eden.⁸¹ Moreover, proponents of the theory suggest playing God in this context involving these efforts can implicate meddling with the natural order of events,

committing a moral wrong, or even making harmful decisions about the lives of other humans.⁸²

In the field of AI, a similar stigma exists.⁸³ As scientists are working toward the creation of a perfect artificial superintelligence that rivals humans, critics believe these beings will continue to grow exponentially in power and be uncontrollable, effectively superseding the entire human race.⁸⁴ Based on fears similar to those discussed above, the primary premise is AI beings will ruin humanity from the inside out.⁸⁵ Common fears include issues around medical misdiagnosis by “smart medical programs,” unpredictable and rebel-like behavior and propensities, lack of common sense and compassion, and the creation of unforeseen challenges and consequences like political warfare and increases in bias.⁸⁶ Moreover, gender studies experts are concerned that sentient AI beings could usher in new forms of oppressive misogyny comparable to the Dark Ages where women are considered inferior.⁸⁷

To illustrate this risk, Tabitha Goldstaub, co-founder of AI market intelligence platform CognitionX, provided the following hypothetical.⁸⁸ As we all know, men and women have different biological structures and processes.⁸⁹ As a result, when confronted with a serious ailment like heart failure, men and women present different symptoms.⁹⁰ In the context of smart medical machines, if the machines are made by men and only trained to treat men, Tabitha posits that by default, “half the population [would] die from heart attacks unnecessarily.”⁹¹ While this has not yet transpired, this is certainly not a prospect of interest to most people. Furthermore, this development is not too far from becoming a reality as news reports of AI-powered technologies going rogue, delivering incorrect information, or malfunctioning and killing people have already emerged.⁹²

Other bioethical concerns have emerged regarding the creation of a valuation hierarchy. There are two facets implicit in this concern. The first is the fear of an economic hierarchy initiated by widespread automation through sentient AI beings creating extreme economic upheaval and replacing members of the workforce.⁹³

In fact, William Halal, professor of Technology and Innovation at George Washington University, believes that sentient AI beings will replace up to 30-40% of the workforce for service and manufacturing jobs requiring lower knowledge by the mid to late 2020s.⁹⁴ Other experts have echoed similar sentiments with some statisticians estimating that nearly 2 million people lost their jobs

to automation.⁹⁵ Reasons for this automation craze are connected to the second fear that revolves around a social hierarchy. From a social perspective, the fear is that with increasing efficiency and intelligence, why would anyone hire a human over a sentient AI being who is just as smart? Not only is the AI equally as smart but it is not prone to “human flaws” like laziness and complacency, reduces costs, does not suffer from fatigue, and can complete more dangerous tasks like weathering dangerous fumes and fires to save trapped citizens.⁹⁶ Echoing concerns regarding gender bias discussed above, in this scenario, AI would be a threat not only to women but to the entirety of mankind.⁹⁷ This would theoretically create a *de facto* preference for sentient AI beings over humans.⁹⁸ While this is currently not the case for all industries like the medical profession and most consumer markets, there is still a trend toward some personal preferences for sentient AI over humans.⁹⁹ Moreover, there are certain intangibles like human proficiencies in creativity and compassion along with endearing idiosyncratic habits and the sound of hearing another human voice that continue to be important to people around the world.¹⁰⁰ Yet, scientists and leaders are concerned.¹⁰¹ Coupled with concerns of unpredictability and rebellion, scientists are concerned that humanity could be wiped out in the blink of an eye especially if countries in all parts of the world are interconnected through AI-based smart programs.¹⁰² Lastly, there are emerging anxieties about moral outsourcing. This term refers to this notion of lazy organization and coding where human creators would assign AI the blame for their own moral dilemmas or hypocrisies.¹⁰³ That is to say, all responsibility and culpability for a racist, sexist, or violent AI program would be on the AI program itself and not on the individual who created the being.¹⁰⁴ While this is a novel concern and the available literature is rather sparse, it has important considerations regarding legislation and health care decisions, as we will see in our last section before concluding.

Proposed Solutions

Proper Oversight and Regulation Through Model Legislation

As the section title suggests, model legislation is the best way to address these valid concerns. Legislation should be proposed

by analogy where protections for sentient animals are modified to accommodate the emergence of sentient AI programs. Thus, there are two model laws that are extremely compelling. The first is the UK animal welfare sentience bill introduced and ratified by the British government in 2021 and 2022, respectively.¹⁰⁵ The bill includes provisions formally recognizing sentient animals as legally protected beings under domestic law and requires the creation of expert committees to ensure not only proper legislative drafting but also compatibility with other industries and needs.¹⁰⁶ These provisions are forward-looking with mechanisms that encourage caution and careful review that is in direct conformity with the precautionary principle discussed above in the third section. Nevertheless, there is one feature of the bill that is the most compelling. That is, the bill advocates disrupting existing industries or legislative prescriptions in place but instead promotes animal welfare and sentience as a crucial part of future decision making.¹⁰⁷ As we discussed above, one of the most widely held fears around AI proliferation involves the nearly irreversible disruption of political and social institutions and industries. However, this provision of the bill guards against this fear, and if this promise is actually implemented, our world can continue on as it is.

The second guiding model form of legislation is the 2021 Artificial Intelligence Act given to the European Parliament.¹⁰⁸ While still in the proposal stage, the act is extremely thorough and provides a clear model framework. Accordingly, this framework constitutes the initial inquiry into whether the sentient AI would pose a significant risk through the use of conformity assessments (tests to see if conformity is present), strict audits and technological reviews, and post-production monitoring.¹⁰⁹ To ensure fairness, the proposal includes relaxed rules for low-risk programs where Parliament members cannot regulate them through past, current, or future laws.¹¹⁰

Furthermore, the creation of AI systems with manipulative or sadistic propensities, inclinations toward exploitation, and proficiencies in biometric recognition or identification of people by their biological or behavioral traits is strictly proscribed.¹¹¹ These restrictions are placed to create national unity and cooperation and pundits are optimistic regarding the successful adoption of this Act.¹¹² In fact, Brazil already followed Parliament's lead and has been working on similar model legislation.¹¹³ Thus, the United States should implement a nearly identical model in the United States with

the forward-looking provisions from the Animal Welfare Sentience Bill. Unfortunately, the current American approach is fragmented with each state devising its own rules regarding the repercussions of involving AI beings and algorithms in society.¹¹⁴ As one might imagine, there is no controlling federal standard, which is a great travesty. Ultimately, Congress should follow the examples set forth by other foreign legislative bodies and draft a controlling federal law with the facets we have just discussed.

Additional Best Practices

In addition, UNESCO has provided some recommendations that are likely implicit but have not been explicitly stated. Taking a holistic approach, UNESCO supports robust protections and access protocols for public and private sector data; strict prohibitions against AI use in mass surveillance practices that abridge human rights; retainment of legal and moral responsibility solely with creators and developers; development of AI ethics offices and committees to ensure compliance and awareness; and the creation of AI technologies using the most resource-efficient processes to positively impact the environment by reducing carbon emissions and ozone depletion.¹¹⁵

Thoughtful Delineation of Health Care Decision-Making Responsibilities

Lastly, we must closely look at questions related to who makes health care decisions for sentient AI beings. While scientists differ in opinion as to whether sentient AI or artificial superintelligence are here or simply beyond the horizon, questions of morality and who controls health care decisions for these beings are considerations of considerable importance.¹¹⁶ This inquiry can be considered a type of moral outsourcing applied to sentient beings where the needs and thoughts of these beings are considered with more care due to their newly acquired status.

Just like humans might be concerned about other humans making decisions regarding their bodily health, the same can be said for sentient AI. Of course, this is all based on the premise that under our model law, sentient AI has been given a series of legal and moral protections. The primary argument in favor of humans making

health care decisions for robots is that by default, AI outsources the responsibility of these heavy moral questions to their creators: humans.¹¹⁷ Here, AI would defer to humans who can answer these questions with better precision as they are more informed, knowledgeable, and better equipped to make moral determinations.¹¹⁸

However, drawing on the deference given to animals under the aforementioned newly passed laws, using this default approach seems to deprive sentient AI of their rights and beliefs. As in constitutional law where the government cannot pass a facially neutral law and then apply the law differently to different people based on differences in race or religion, there should be no room for sentient AI to be treated differently than their sentient peers.¹¹⁹ Moreover, in constitutional law, differential impact is not enough alone unless there is evidence of a single discriminatory purpose behind the law.¹²⁰ An augmented version of this standard is helpful in determining who should make decisions regarding health care for sentient AI. Here, if we believe that both sentient animals and AI possess the same rights, but health care providers treat them differently in terms of considering each party's preferences and needs, then this is enough to constitute differential impact. Under this augmented standard, evidence of a single discriminatory purpose is still required and is easily shown. If we recall fears of the Singularity phenomenon and AI supplanting the human race, the combination of these fears should be enough to make a showing of discriminatory purpose behind treating both sentient beings differently in the medical context.¹²¹

If we analogize further by incorporating other bioethical principles, we arrive at a similar conclusion militating in favor of granting sentient AI sovereignty in making delineations about their own health. For instance, it is generally undisputed that healthy adults retain primary authority and discretion in making decisions pertaining to their health.¹²² Pending there are no issues regarding incapacity, age, or mental illness, there is a presumption in favor of granting adults over the age of 18 these important abilities.¹²³ Therefore, using these existing practices as guidance for answering this question of medical power and responsibility, one should support the implementation of a similar model. In adherence to the requirements discussed above in the subsection Proper Oversight and Regulation Through Model Legislation, differences in medical health care decision making would be based on level of risk. For low-risk sentient AI programs, these beings would have access

to the fully spectrum or panoply of medical decision rights. For medium to high-risk programs, akin to those who are incapacitated or mentally ill, a determination of the programs conformity and danger per the parameters discussed above would need to be carried out. This proposed approach is especially timely as developments regarding the reproduction of AI-powered robots continue to make headlines.¹²⁴

Conclusion

Overall, discussions around sentience and sentient beings are undeniably complex and multilayered. From precisely specifying what sentience refers to differentiating between quintessential characteristics of humanity to considering the role of sentient AI beings and their accompanying risks, stigmas, and potential moral and legal implications, this article has explored this issue from a variety of vantage points. This has all been done to provide guidance on a topic that gains more importance and prominence as each year progresses. While there will certainly be those who oppose the solutions and perspectives delivered in this article, the focus should remain on establishing proper ethical and moral constructs. As Hugh McLachlan suggests, “To deny conscious persons moral respect and consideration on the grounds that they had artificial rather than natural bodies would seem to be arbitrary and whimsical.”¹²⁵ Thus, until a compelling justification exists for such a denial, we should welcome these beings with open arms instead of animosity or fear despite their admittedly enigmatic nature.

Notes

* Basil Cunningham, J.D., may be contacted at bazil.cunningham@pepperdine.edu.

1. See *discussion* in the section Philosophical Primary Basis for Legal Solutions.

2. See *discussion* in the section Proposed Solutions.

3. David DeGrazia & Joseph Millum, “Moral Status (Chapter 7)—A Theory of Bioethics.” Cambridge Core, Cambridge University Press (Aug. 17, 2021), <https://www.cambridge.org/core/books/theory-of-bioethics/moral-status/0A0B4D7E1829D937C91CB456E5B27925>.

4. Animal Ethics, "The Problem of Consciousness," Animal Ethics (Feb. 20, 2022), <https://www.animal-ethics.org/problem-consciousness/#:~:text=Sentience%20refers%20to%20the%20ability,sentience%20and%20consciousness%20is%20slight>.

5. "Sentience," Merriam-Webster, <https://www.merriam-webster.com/dictionary/sentience>; Sentience. https://www.oxfordlearnersdictionaries.com/us/definition/american_english/sentient.

6. Alex Lane, "5x5 What Does It Mean to Be Smart?" Medium, Five by Five (Apr. 10, 2018), <https://medium.com/5by5/5x5-monkey-monday-what-does-it-mean-to-be-smart-44d460ed9965>.

7. *Id.*

8. Animal Ethics, "The Problem of Consciousness," Animal Ethics (Feb. 20, 2022), <https://www.animal-ethics.org/problem-consciousness/#:~:text=Sentience%20refers%20to%20the%20ability,sentience%20and%20consciousness%20is%20slight>.

9. Kathleen V. Wilkes, "Is Consciousness Important?" The British Journal for the Philosophy of Science, vol. 35, no. 3, 1984, pp. 223-43, <https://doi.org/10.1093/bjps/35.3.223>.

10. Colin Allen & Michael Trestman, "Animal Consciousness," Stanford Encyclopedia of Philosophy, Stanford University (Dec. 23, 1995), <https://plato.stanford.edu/archives/sum2014/entries/consciousness-animal/#hist>.

11. Mark Crooks, "Consciousness: Sentient and Rational." The Journal of Mind and Behavior, vol. 32, no. 3, 2011, pp. 251-75, JSTOR, <http://www.jstor.org/stable/43854462>.

12. *Id.*

13. N. Block, "On a Confusion About a Function of Consciousness," Behavioral and Brain Sciences, 18, 227-47 (1995); T. Nagel, "What Is It Like To Be a Bat?" Philosophical Review, 83, 435-50 (1974).

14. Kelly Oliver, "Animal Pedagogy: The Origin of 'Man' in Rousseau and Herder," Routledge Taylor & Francis Group (2006), https://as.vanderbilt.edu/philosophy/people/facultyfiles/oliver-animal_pedagogy.pdf.

15. George Mobus, "Sapience," Evolution of Sapience, The University of Washington Tacoma, Institute of Technology, <https://faculty.washington.edu/gmobus/TheoryOfSapience/SapienceExplained/5.evolutionOfSapience/evolutionOfSapience.html>.

16. David P. Barash, "Are We Uniquely Sapient?" Psychology Today, Sussex Publishers (Oct. 30, 2019), <https://www.psychologytoday.com/us/blog/pura-vida/201910/are-we-uniquely-sapient>.

17. "Whales and Dolphins Lead 'Human-Like Lives' Thanks to Big Brains, Says Study." The Guardian, Guardian News and Media (Oct. 16, 2017), <https://www.theguardian.com/science/2017/oct/16/whales-and-dolphins-human-like-societies-thanks-to-their-big-brains>; "Sentient and Sapient Whales and Dolphins," Whale & Dolphin Conservation USA (Apr. 17, 2020), <https://us.whales.org/whale-culture/sentient-and-sapient-whale-and-dolphins/>.

18. Mahmud Arif Pavel, "Direction for Artificial Intelligence to Achieve Sapiency Inspired by Homo Sapiens," *International Journal of Advanced Research in Artificial Intelligence (Ijarai)*, The Science and Information (SAI) Organization Limited (2016), <https://thesai.org/Publications/ViewPaper?Volume=5&Issue=9&Code=ijarai&SerialNo=3>.

19. "Self-Awareness Is What Makes US Human," *Big Think* (Apr. 19, 2022), <https://bigthink.com/neuropsych/self-awareness-makes-us-human/>; David DeGrazia, "Self-Awareness in Animals," Cambridge University Press, <https://philpapers.org/rec/DEGSIA-2>; Shaun Gallagher & Dan Zahavi, "Phenomenological Approaches to Self-Consciousness," *Stanford Encyclopedia of Philosophy*, Stanford University (May 23, 2019), <https://plato.stanford.edu/entries/self-consciousness-phenomenological/>.

20. *Id.*

21. Stephanie Gibeault, MSc, "Are Dogs Self Aware? New Research Suggests Yes," *American Kennel Club* (Apr. 8, 2021), <https://www.akc.org/expert-advice/news/a-new-way-to-look-at-dog-self-awareness/>.

22. Alexandra Horowitz, "Smelling Themselves: Dogs Investigate Their Own Odours Longer When Modified in an 'Olfactory Mirror' Test," *Phys. org*, Behavioural Processes (Sept. 6, 2017), <https://www.sciencedirect.com/science/article/abs/pii/S0376635717300104?via%3Dihub>.

23. *Id.*

24. *See discussion* in the subsection What About Sentient AI?

25. Eric Martínez & Christoph Winter, "Protecting Sentient Artificial Intelligence: A Survey of Lay Intuitions on Standing, Personhood, and General Legal Protection," *Frontiers* (Nov. 26, 2021), <https://www.frontiersin.org/articles/10.3389/frobt.2021.788355/full>; , Audrey Millemann, "Should AI Machines Have Rights?" *JD Supra* (Sept. 12, 2022), <https://www.jdsupra.com/legalnews/should-ai-machines-have-rights-4583419/>; Christine M. Korsgaard, "Personhood, Animals, and the Law," *Think*, vol. 12, no. 34, pp. 25-32 (2013), doi:10.1017/S1477175613000018; Visa Kurki, "Legal Personhood and Animal Rights," *Journal of Animal Ethics*, 11 (1): 47-62 (Jan. 1, 2021), <https://doi.org/10.5406/janimaethics.11.1.0047>.

26. *See discussion* in subsection Elucidating Additional Characteristics.

27. *See discussions* in the subsections Identifying Sentient and Non-Sentient Animals and What About Sentient AI?

28. "Criteria for Recognizing Sentience," *Animal Ethics* (Feb. 23, 2022), <https://www.animal-ethics.org/criteria-for-recognizing-sentience/>.

29. Lori Marino, "Sentience in All Organisms with Centralized Nervous Systems," *Animal Sentience* (July 1, 2020), <https://www.wellbeingintlstudies-repository.org/cgi/viewcontent.cgi?article=1596&context=animsent>; T. Nagel "What Is It Like to Be a Bat?" *Philosophical Review*, 83, 435-50 (1974).

30. The Central Nervous System, <https://mcb.berkeley.edu/courses/mcb135e/central.html>.

31. Adrian Rad, "Central Nervous System," Kenhub (Aug. 1, 2022), <https://www.kenhub.com/en/library/anatomy/the-central-nervous-system>.

32. See *discussion* in the subsection Elucidating Additional Characteristics.

33. Marc Bekoff, "After 2,500 Studies, It's Time to Declare Animal Sentience Proven (Op-Ed)," LiveScience, Purch (Sept. 6, 2013), <https://www.livescience.com/39481-time-to-declare-animal-sentience.html>.

34. Grace Hussain, "What Is a Sentient Being?" Sentient Media (Oct. 25, 2022), <https://sentientmedia.org/sentient-being/>; Hemi Kim, "What Is the Definition of Sentient and Are All Animals Sentient?" Sentient Media (Nov. 1, 2022), [https://sentientmedia.org/sentient-definition/#:~:text=What%20Creatures%20Are%20Considered%20Sentient,whales%2C%20dolphins%2C%20and%20more](https://sentientmedia.org/sentient-definition/#:~:text=What%20Creatures%20Are%20Considered%20Sentient,whales%2C%20dolphins%2C%20and%20more;); Emily Birch, "Crying Elephants, Giggling Rats and Other Surprisingly Sentient Animals," Discover Magazine (Nov. 19, 2019), <https://www.discovermagazine.com/planet-earth/crying-elephants-giggling-rats-and-other-surprisingly-sentient-animals>.

35. The Cambridge Declaration on Consciousness, Cambridge University (July 7, 2012), <https://fcmconference.org/img/CambridgeDeclarationOnConsciousness.pdf>.

36. Marc Bekoff, "Scientists Conclude Nonhuman Animals Are Conscious Beings," Psychology Today, Sussex Publishers (Aug. 10, 2012), <https://www.psychologytoday.com/us/blog/animal-emotions/201208/scientists-conclude-nonhuman-animals-are-conscious-beings>; *id.* The Cambridge Declaration on Consciousness.

37. Helen Proctor et al., "Searching for Animal Sentience: A Systematic Review of the Scientific Literature," *Animals*, vol. 3, no. 3, pp. 882-906 (Sept. 2013), Crossref, <https://doi.org/10.3390/ani3030882>; Marc Bekoff, "The Sentience Mosaic: A Wonderful New Resource," Psychology Today, Sussex Publishers (Oct. 1, 2012), <https://www.psychologytoday.com/us/blog/animal-emotions/201210/the-sentience-mosaic-wonderful-new-resource>.

38. Ingrid L. Taylor, "Animal Sentience and Emotions: The Argument for Universal Acceptance," People for the Ethical Treatment of Animals-Laboratory Investigations Department (2021), https://www.peta.org/wp-content/uploads/2021/01/Report-on-Animal-Sentience_Ingrid-Taylor_2021_NE.pdf; Marc Bekoff, "Animal Consciousness and Science Matter," Psychology Today, Sussex Publishers (May 7, 2012), <https://www.psychologytoday.com/us/blog/animal-emotions/201205/animal-consciousness-and-science-matter>.

39. Bernard Marr. "What Is Ai?" Bernard Marr (July 13, 2021), <https://bernardmarr.com/what-is-ai/>.

40. S. Shubhendu & J. Frank Vijay. "Applicability of Artificial Intelligence in Different Fields of Life" (2013); Darrell M. West, "What Is Artificial Intelligence?" Brookings (Mar. 9, 2022), <https://www.brookings.edu/research/what-is-artificial-intelligence/>.

41. Bernard Marr, “What Are the Four Types of AI?” (July 15, 2021), <https://bernardmarr.com/what-are-the-four-types-of-ai/>; Hanif Khan, “Different Types of Artificial Intelligence Systems,” University of Science and Technology of Ha Noi (Sept. 2021), https://www.researchgate.net/publication/355021812_Types_of_AI_Different_Types_of_Artificial_Intelligence_Systems_fossgurucomtypes-of-ai-different-types-of-artificial-intelligence-systemsf-ai-different-types-of-artificial-intelligence-systems.

42. *Id.*

43. Thomas Davenport & Ravi Kalakota, “The Potential for Artificial Intelligence in Healthcare,” *Future Healthcare Journal*, vol. 6, 2: 94-98 (2019), <https://www.rcpjournals.org/content/futurehosp/6/2/94>; Bernard Marr, “What Is the Importance of Artificial Intelligence (AI)” (July 13, 2021), <https://bernardmarr.com/what-is-the-importance-of-artificial-intelligence-ai/>; Alex Guilherme, “AI and Education: The Importance of Teacher and Student Relations,” Universidade Católica Do Rio Grande Do Sul (Feb. 4, 2017), https://repositorio.pucrs.br/dspace/bitstream/10923/21168/2/AI_and_education_the_importance_of_teacher_and_student_relations.pdf.

44. Radha Vuppalapati, “The Role of Artificial Intelligence in Fighting COVID-19,” WIPRO (May 2021), <https://www.wipro.com/blogs/radha-vuppalapati/the-role-of-artificial-intelligence-in-fighting-covid-19/#:~:text=With%20AI%20use%20growing%20in,data%20to%20identify%20insightful%20patterns>.

45. “Artificial Intelligence—A Promising Anti-Corruption Tool in Development Settings?” U4 Anti-Corruption Resource Centre, <https://www.u4.no/publications/artificial-intelligence-a-promising-anti-corruption-tool-in-development-settings/shortversion>; Naveen Joshi, “Can AI Be Used as an Anti-Corruption Tool?” *Forbes* (Sept. 19, 2021), <https://www.forbes.com/sites/naveenjoshi/2021/09/19/can-ai-be-used-as-an-anti-corruption-tool/?sh=232c2faf7569>.

46. “Artificial Intelligence & Autopilot,” Tesla, <https://www.tesla.com/AI>; Bernard Marr, “How Tesla Is Using Artificial Intelligence to Create the Autonomous Cars of the Future,” Bernard Marr (July 13, 2021), <https://bernardmarr.com/how-tesla-is-using-artificial-intelligence-to-create-the-autonomous-cars-of-the-future/>.

47. *See* note 44.

48. Mike Wehner, “This New AI Can Read Your Mind and Predict Your Thoughts,” *New York Post* (July 3, 2017), <https://nypost.com/2017/07/03/this-new-ai-can-read-your-mind-and-predict-your-thoughts/>; “AI Can Now Read Your Mind, and Actually Convert Your Thoughts into Pictures.” *Psych-NewsDaily* (Nov. 16, 2022), <https://www.psychnewsdaily.com/this-computer-can-read-your-mind-and-render-your-thoughts-as-pictures/>; “A Computer Reads and Predicts Thoughts.” *Tectales, Tagging Medical Technology* (Oct. 5, 2020), <https://tectales.com/ai/a-computer-reads-predicts-thoughts.html>.

49. Raj Hayer, “AI Is Here and It Is Here to Stay.” Medium, Ideas and Innovation Sept. 3, 2020), <https://medium.com/ideas-and-innovation/ai-is-here-and-it-is-here-to-stay-f91e48353917>; Ron Schmelzer, “AI Is Here to Stay in Your City and Local Government,” *Forbes* (Nov. 8, 2020), www.forbes.com/sites/cognitiveworld/2020/11/08/ai-is-here-to-stay-in-your-city-and-local-government/?sh=5749bb0b77a3; Juliet Gerrard et al., “AI Is Here to Stay. Now We Need to Ensure Everyone Benefits,” *The Conversation* (Aug. 5, 2019), <https://theconversation.com/ai-is-here-to-stay-now-we-need-to-ensure-everyone-benefits-121175>.

50. Alena Kuzub, “Should We Be Concerned About Google AI Being Sentient?” *News @ Northeastern* (June 16, 2022), <https://news.northeastern.edu/2022/06/16/google-sentient-ai-concerns/>.

51. Anne Quito, “Google Is Developing a New Superintelligent AI but Ethical Questions Remain,” *Quartz* (July 20, 2022), <https://qz.com/2042493/pathways-google-is-developing-a-superintelligent-multipurpose-ai/>; W. Bryk, “Artificial Superintelligence: The Coming Revolution,” *Harvard Science Review*, <https://harvardsciencereview.com/artificial-superintelligence-the-coming-revolution/>.

52. Leonardo de Cosmo, “Google Engineer Claims AI Chatbot Is Sentient: Why That Matters,” *Scientific American* (July 12, 2022), www.scientificamerican.com/article/google-engineer-claims-ai-chatbot-is-sentient-why-that-matters; Nitasha Tiku, “The Google Engineer Who Thinks the Company’s AI Has Come to Life,” *Washington Post* (June 11, 2022), www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine.

53. Sarah Hucal, “Google’s ‘sentient’ AI: From Films to Reality?” *dw.com* (June 15, 2022), www.dw.com/en/googles-sentient-ai-from-films-to-reality/a-62132388; “Ai Sentience: Friend or Foe?” *Alectio* (July 10, 2022), <https://www.alectio.com/2022/07/04/ai-sentience-friend-or-foe/>.

54. Nico Grant & Cade Metz, “Google Sidelines Engineer Who Claims Its A.I. Is Sentient,” *New York Times* (June 12, 2022), www.nytimes.com/2022/06/12/technology/google-chatbot-ai-blake-lemoine.html; Kylie Robison, “Alphabet CEO Sundar Pichai Says ‘Broken’ Google Voice Assistant Proves That A.I. Isn’t Sentient,” *Fortune* (Sept. 7, 2022), <https://fortune.com/2022/09/07/alphabet-ceo-sundar-pichai-says-broken-google-voice-assistant-proves-that-a-i-isnt-sentient/>; “Google Fires Software Engineer Who Claimed Its AI Chatbot Is Sentient,” *Reuters* (July 23, 2022), www.reuters.com/technology/google-fires-software-engineer-who-claimed-its-ai-chatbot-is-sentient-2022-07-23.

55. *Id.* *Reuters* (July 23, 2022); Eli Collins, “LaMDA: Our Breakthrough Conversation Technology,” *Google* (May 18, 2021), <https://blog.google/technology/ai/lamda/>.

56. Maggie Harrison, “Fired Engineer Says He Begged Google to Test Whether Experimental AI Was Sentient,” *Futurism* (Aug. 16, 2022), <https://futurism.com/engineer-begged-google-test-experimental-ai-sentient>;

“Suspended Google Engineer Says the AI He Believes to Be Sentient Hired a Lawyer,” *Business Insider* (June 23, 2022), www.businessinsider.com/suspended-google-engineer-says-sentient-ai-hired-lawyer-2022-6?international=true&r=US&IR=T; Colin Lodewick, “Google’s Suspended AI Engineer Corrects the Record: He Didn’t Hire an Attorney for the ‘Sentient’ Chatbot, He Just Made Introductions—The Bot Hired the Lawyer,” *Fortune* (June 23, 2022), <https://fortune.com/2022/06/23/google-blade-lemoine-ai-lambda-wired-attorney/>; “The Story of LaMDA and the Google Engineer Just Got Even Weirder,” *IFLScience* (June 28, 2022), www.iflscience.com/it-hired-a-lawyer-the-story-of-lambda-and-the-google-engineer-just-got-even-weirder-64229.

57. Jon Christian, “Google Engineer Says Lawyer Hired by ‘Sentient’ AI Has Been ‘Scared off’ the Case,” *Futurism* (June 22, 2022), <https://futurism.com/google-engineer-lawyer-representing-ai-scared-off>; Lonnie Lee Hood, “Google Insider Claims Company’s ‘Sentient’ AI Has Hired an Attorney,” *Futurism* (June 19, 2022), <https://futurism.com/the-byte/google-insider-ai-lawyer>.

58. Brian Christian, “How a Google Employee Fell for the Eliza Effect,” *The Atlantic* (June 21, 2022), www.theatlantic.com/ideas/archive/2022/06/google-lambda-chatbot-sentient-ai/661322; Rebecca Heilweil, “Blake Lemoine, Google, and Searching for Souls in the Algorithm,” *Vox* (June 30, 2022), www.vox.com/recode/2022/6/30/23188222/silicon-valley-blake-lemoine-chatbot-eliza-religion-robot; Theo Wayt, “Google Engineer Says Christianity Helped Him Understand AI Is ‘Sentient,’” *New York Post* (June 14, 2022), <https://nypost.com/2022/06/14/google-engineer-says-christianity-helped-him-believe-in-sentient-ai/>.

59. Michael Kaplan, “After Google Chatbot Becomes ‘Sentient,’ MIT Prof Says Alexa Could Too,” *New York Post* (June 14, 2022), <https://news.knowledia.com/US/en/articles/after-google-chatbot-becomes-sentient-mit-prof-says-alexa-could-too-86c7eae3f16374e170c20d5ca35eac66f7e2138b?source=rss>.

60. Carole Cadwalladr, “Are The Robots About to Rise? Google’s New Director of Engineering Thinks So...” *The Guardian* (Mar. 22, 2018), www.theguardian.com/technology/2014/feb/22/robots-google-ray-kurzweil-terminator-singularity-artificial-intelligence; Vernor Vinge on the Singularity, mindstalk.net/vinge/vinge-sing.html; Matthew Sparkes, “Matthew Sparkes,” *The Telegraph* (Jan. 13 2015), www.telegraph.co.uk/technology/news/11342200/Top-scientists-call-for-caution-over-artificial-intelligence.html; Rory Cellan-Jones, “Stephen Hawking Warns Artificial Intelligence Could End Mankind,” *BBC News* (Dec. 2, 2014), www.bbc.com/news/technology-30290540.

61. Christof Koch, “Will Machines Ever Become Conscious?” *Scientific American* (Dec. 1, 2019), <https://www.scientificamerican.com/article/will-machines-ever-become-conscious/>.

62. *Id.*

63. *Id.*

64. Steven Johnson & Nikita Izhev, "A.I. Is Mastering Language. Should We Trust What It Says?" *New York Times* (Apr. 17, 2022), www.nytimes.com/2022/04/15/magazine/ai-language.html.

65. Kelsey Piper, "OpenAI's New Language AI Is Available to Try Yourself," *Vox* (May 15, 2019), www.vox.com/2019/5/15/18623134/openai-language-ai-gpt2-poetry-try-it.

66. John Seabrook, "Can a Machine Learn to Write for the *New Yorker*?" *The New Yorker* (Oct. 7, 2019), www.newyorker.com/magazine/2019/10/14/can-a-machine-learn-to-write-for-the-new-yorker.

67. "AI Sentience: Friend or Foe?" *Alectio* (July 10, 2022), <https://www.alectio.com/2022/07/04/ai-sentience-friend-or-foe/>; Kyle Mahowald & Anna A. Ivanova, "Google's Powerful Artificial Intelligence Spotlights a Human Cognitive Glitch," *SciTechDaily* (July 7, 2022), <https://scitechdaily.com/googles-powerful-artificial-intelligence-spotlights-a-human-cognitive-glitch/>.

68. *Id.*

69. Karl N. Llewellyn, "Some Realism About Realism: Responding to Dean Pound," *Harvard Law Review*, vol. 44, no. 8, pp. 1222-64 (1931). <https://doi.org/10.2307/1332182>; "Legal Positivism," *Obo*, www.oxfordbibliographies.com/view/document/obo-9780195396577/obo-9780195396577-0065.xml; Jorge Luis Fabra-Zamora, "Legal Positivism as a Theory of Law's Existence: A Comment on Margaret Martin's 'Judging Positivism,'" *Digital Commons @ University at Buffalo School of Law*, https://digitalcommons.law.buffalo.edu/journal_articles/1064/; G. Gilmore, "Legal Realism: Its Cause and Cure," *Yale Law Journal*, vol. 70, no. 7 (June 1961).

70. Brandon Sample, "Sentencing Philosophies—Can We Build the Perfect Prison?" *Sentencing.net* (Nov. 6, 2019), sentencing.net/sentencing/sentencing-philosophies; "Theories of Criminal Law," *Stanford Encyclopedia of Philosophy* (Aug. 6, 2018); *Philosophy of Criminal Law*, Office of Justice Programs. www.ojp.gov/ncjrs/virtual-library/abstracts/philosophy-criminal-law.

71. Rupert Read & Tim O'Riordan, "The Precautionary Principle Under Fire," *Environment: Science and Policy for Sustainable Development*, 59: 4-15 (Sept.-Oct. 2017), <https://www.tandfonline.com/doi/abs/10.1080/00139157.2017.1350005?journalCode=venv20>. S2CID 158589782; Sonia Boutillon, "The Precautionary Principle: Development of an International Standard," 23 *Mich. J. Int'l L.* 429 (2002), <https://repository.law.umich.edu/mjil/vol23/iss2/7>; Jacob S. Sherkow, "Controlling CRISPR Through Law: Legal Regimes as Precautionary" (Oct. 9, 2019), <https://www.liebertpub.com/doi/10.1089/crispr.2019.0029>; Roberto Andorno, "The Precautionary Principle: A New Legal Standard for a Technological Age," *De Gruyter* (Feb. 1, 2004), <https://www.degruyter.com/document/doi/10.1515/jibl.2004.1.1.11/html>; René Von Schomberg, "The Precautionary Principle: Its Use Within Hard and Soft Law," *European Journal of Risk Regulation*, vol. 3, no. 2, pp. 147-56 (2012), <https://www.cambridge.org/core/journals/european-journal-of-risk-regulation/>

article/abs/precautionary-principle-its-use-within-hard-and-soft-law/67B6FDCCABAC681E055022EBF0D2D968.

72. *Id.*

73. “The Sentience Shift in Animal Research,” Taylor & Francis, www.tandfonline.com/doi/full/10.1080/20502877.2022.2077681.

74. *Id.*

75. J. Birch, “Animal Sentience and the Precautionary Principle,” *Animal Sentience*, 16 (1), 1-15 (2017).

76. *Id.*

77. Derek Turner & Lauren Hartzell. “The Lack of Clarity in the Precautionary Principle,” *Environmental Values*, vol. 13, no. 4, pp. 449-60 (2004), <https://www.jstor.org/stable/30302021>. Accessed 16 Nov. 2022.

78. A. Waytz & L. Young. “Aversion to Playing God and Moral Condemnation of Technology and Science,” *Royal Society*, vol. 374, no. 1771 (Apr. 2019), <https://doi.org/10.1098/rstb.2018.0041>.

79. Tania Lombrozo, “Human Exceptionalism Stifles Progress.” *Nautilus* (May 1, 2019), <https://nautil.us/human-exceptionalism-stifles-progress-237386/>.

80. Robert Winston, “Playing God?” *Nature* (Dec. 11, 2003), www.nature.com/articles/426603a?error=cookies_not_supported&code=f7128050-84d8-4591-a585-c93b4d084d3f; Moti Mizrahi, “How to Play the ‘Playing God’ Card,” *Science and Engineering Ethics*, vol. 26, 3: 1445-1461 (2020), doi:10.1007/s11948-020-00176-7.

81. Peter Dabrock, “Playing God? Synthetic Biology as a Theological and Ethical Challenge,” *Systems and Synthetic Biology*, vol. 3, 1-4: 47-54 (2009), doi:10.1007/s11693-009-9028-5.

82. 6.8. Playing God. www.eubios.info/ABC4/abc4206.htm.

83. D. Basulto, “How We’re Playing God Now.” *Washington Post* (June 29, 2012), www.washingtonpost.com/blogs/innovations/post/how-were-playing-god-now/2012/06/29/gJQAuxbeBW_blog.html.

84. Frank Pasquale, “‘Machines Set Loose to Slaughter’: The Dangerous Rise of Military AI,” *The Guardian* (Apr. 2, 2021), www.theguardian.com/news/2020/oct/15/dangerous-rise-of-military-ai-drone-swarm-autonomous-weapons.

85. *See discussion* in subsection What About Sentient AI?

86. CB Insights, “How AI Will Go Out of Control According to 52 Experts,” CB Insights Research (Aug. 24, 2021), www.cbinsights.com/research/ai-threatens-humanity-expert-quotes; RT. “‘Whoever Leads in AI Will Rule the World’: Putin to Russian Children on Knowledge Day,” RT International (Sept. 1, 2017), www.rt.com/news/401731-ai-rule-world-putin; Jayshree Pandya, “The Weaponization of Artificial Intelligence,” *Forbes* (Jan. 14, 2019), www.forbes.com/sites/cognitiveworld/2019/01/14/the-weaponization-of-artificial-intelligence/?sh=37c3fd933686; Indermit Gill, “Whoever Leads in Artificial Intelligence in 2030 Will Rule the World Until 2100,” *Brookings* (Mar. 9, 2022), www.brookings.edu/blog/future-development/2020/01/17/whoever-

leads-in-artificial-intelligence-in-2030-will-rule-the-world-until-2100; Paul Daugherty, "Using Artificial Intelligence to Promote Diversity," MIT Sloan Management Review (Nov. 21, 2018), <https://sloanreview.mit.edu/article/using-artificial-intelligence-to-promote-diversity/>.

87. N. Kobie, "Artificial Intelligence Isn't Good for Women, but We Can Fix It" (Apr. 2018), www.yahoo.com/entertainment/artificial-intelligence-isn-apos-t-220328763.html; Jessica Bateman, "Sexist Robots Can Be Stopped by Women Who Work in AI," *The Guardian* (May 29, 2017), www.theguardian.com/careers/2017/may/29/sexist-robots-can-be-stopped-by-women-who-work-in-ai.

88. CB Insights, "How AI Will Go Out of Control According to 52 Experts," CB Insights Research (Aug. 24, 2021), www.cbinsights.com/research/ai-threatens-humanity-expert-quotes; Susannah Butter, "The Female Tech Entrepreneur That's Propelling London's AI Revolution," *Evening Standard* (May 18, 2017), www.standard.co.uk/lifestyle/london-life/the-female-tech-entrepreneur-that-s-propelling-london-s-ai-revolution-a3542096.html.

89. Rosanne Spector, "How Men's and Women's Brains Are Different | Stanford Medicine," *Stanford Medicine Magazine* (Sept. 21, 2022), stanmed.stanford.edu/how-mens-and-womens-brains-are-different; Bobbi Low, "Biological Bases of Sex Differences," *SpringerLink* (2003), https://link.springer.com/referenceworkentry/10.1007/0-387-29907-6_3?error=cookies_not_supported&code=7bec6e80-9a81-4623-8191-7afa0313c900.

90. Brigham and Women's Hospital, "Heart Disease: 7 Differences Between Men and Women," Brigham and Women's Hospital, <https://give.brighamandwomens.org/7-differences-between-men-and-women/>.

91. *Id.*

92. Daisuke Wakabayashi, "Self-Driving Uber Car Kills Pedestrian in Arizona, Where Robots Roam," *New York Times* (Mar. 20, 2018), www.nytimes.com/2018/03/19/technology/uber-driverless-fatality.html; Casey Ross & Ike Swetlitz, "IBM's Watson Supercomputer Recommended 'Unsafe and Incorrect' Cancer Treatments, Internal Documents Show," *STAT* (Sept. 13, 2022), www.statnews.com/2018/07/25/ibm-watson-recommended-unsafe-incorrect-treatments; Jack Stewart, "Tesla's Self-Driving Autopilot Was Turned on in Deadly California Crash," *Wired* (Mar. 31, 2018), www.wired.com/story/tesla-autopilot-self-driving-crash-california.

93. "Artificial Intelligence Will Change Human Value(s)." AFCEA International (Mar. 1, 2019), www.afcea.org/signal-media/artificial-intelligence-will-change-human-values.

94. jfsdigital.org. "Forecasts of AI and Future Jobs in 2030: Muddling Through Likely, With Two Alternative Scenarios." *For Work / Against Work*.

95. "AI Should Augment Human Intelligence, Not Replace It," *Harvard Business Review* (Aug. 30, 2021), <https://hbr.org/2021/03/ai-should-augment-human-intelligence-not-replace-it>; Valerias Bangert, "AI Is Quietly Eating Up the World's Workforce with Job Automation," *VentureBeat* (Jan. 31, 2022),

<https://venturebeat.com/datadecisionmakers/ai-is-quietly-eating-up-the-worlds-workforce-with-job-automation/>; Alana Semuels, “Millions of Americans Have Lost Jobs in the Pandemic—And Robots and AI Are Replacing Them Faster Than Ever,” *Time* (Aug. 6, 2020), <https://time.com/5876604/machines-jobs-coronavirus/>.

96. Jennifer Shalamanov, “Why AI Will Replace Some Jobs and Others Will Stick Around,” *Udacity* (Sept. 27, 2022), www.udacity.com/blog/2021/02/why-ai-will-replace-some-jobs-and-others-will-stick-around.html.

97. *Id.*

98. M. Pezzo & J. Beckstead, “Patients Prefer Artificial Intelligence to a Human Provider, Provided the AI Is Better Than the Human: A Commentary on Longoni, Bonezzi and Morewedge,” *Judgment and Decision Making*, vol. 15, no. 3 (May 2020), <https://www.cambridge.org/core/journals/judgment-and-decision-making/article/patients-prefer-artificial-intelligence-to-a-human-provider-provided-the-ai-is-better-than-the-human-a-commentary-on-longoni-bonezzi-and-morewedge-2019/FDBF96A9CA1C5C88E163DDD5EF6E4115>; Angelica Mari, “Young Brazilians Prefer AI Over Humans for E-Commerce Advice,” *ZDNET* (May 25, 2021), www.zdnet.com/article/young-brazilians-prefer-ai-over-humans-for-e-commerce-advice.

99. Bret Kinsella, “New Report Says Consumers Have 71% Preference for Human Over Robot Voices for Voice Assistant User Experience,” *Voicebot.ai* (Aug. 10, 2022), <https://voicebot.ai/2019/11/20/new-report-says-consumers-have-71-preference-for-human-over-robot-voices-for-voice-assistant-user-experience/>; Derya Yakar et al., “Do People Favor Artificial Intelligence Over Physicians? A Survey Among the General Population and Their View on Artificial Intelligence in Medicine,” *Value in Health*, vol. 25, no. 3, pp. 374-81 (Mar. 2022), <https://doi.org/10.1016/j.jval.2021.09.004>.

100. “Why AI Still Needs Humans in the Loop | Article,” *Genpact* (Nov. 14, 2022), www.genpact.com/insight/why-ai-still-needs-humans-in-the-loop; The MIT Press Reader, “AI Is No Match for the Quirks of Human Intelligence.” *The MIT Press Reader* (Oct. 5, 2021), <https://thereader.mitpress.mit.edu/ai-insight-problems-quirks-human-intelligence/>; Daniel B. Shank et al., “Humans Are Perceived as Better, but Weaker, Than Artificial Intelligence: A Comparison of Affective Impressions of Humans, AIs, and Computer Systems in Roles on Teams,” *Computers in Human Behavior Reports*, vol. 3, p. 100092 (Jan. 2021), <https://doi.org/10.1016/j.chbr.2021.100092>.

101. *See supra* note 87.

102. Tanya Lewis, “Don’t Let Artificial Intelligence Take Over, Top Scientists Warn,” *livescience.com* (Jan. 12, 2015), www.livescience.com/49419-artificial-intelligence-dangers-letter.html; Future of Life Institute, “Research Priorities for Robust and Beneficial Artificial Intelligence: An Open Letter,” *Future of Life Institute* (Sept. 28, 2022), <https://futureoflife.org/open-letter/>

ai-open-letter/; Kevin Rawlinson, “Microsoft’s Bill Gates Insists AI Is a Threat,” BBC News (Jan. 29, 2015), www.bbc.com/news/31047780.

103. Rumman Chowdhury, “Moral Outsourcing: Finding the Humanity in Artificial Intelligence,” *Forbes* (Oct. 13, 2017), www.forbes.com/sites/rumman-chowdhury/2017/10/13/moral-outsourcing-2/?sh=3a02c5ef70a; Virginia Dignum, “The ART of AI—Accountability, Responsibility, Transparency,” *Medium* (June 14, 2018), <https://medium.com/@viriniadignum/the-art-of-ai-accountability-responsibility-transparency-48666ec92ea5>; “Assessment List for Trustworthy Artificial Intelligence (ALTAI) for Self-assessment,” *Shaping Europe’s Digital Future* (July 17, 2020), <https://digital-strategy.ec.europa.eu/en/library/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment>.

104. *Id.*

105. Department for Environment, Food & Rural Affairs, “Animals to Be Formally Recognised as Sentient Beings in Domestic Law,” *Gov.uk* (May 13, 2021), www.gov.uk/government/news/animals-to-be-formally-recognised-as-sentient-beings-in-domestic-law.

106. *Id.*

107. Department for Environment, Food & Rural Affairs, “Lobsters, Octopus and Crabs Recognised as Sentient Beings Amendment to Animal Welfare (Sentience) Bill Following LSE Report on Decapod and Cephalopod Sentience,” *Gov.uk* (May 13, 2021), www.gov.uk/government/news/animals-to-be-formally-recognised-as-sentient-beings-in-domestic-law.

108. “Artificial Intelligence Act,” *European Parliament*, [www.europarl.europa.eu/thinktank/en/document/EPRS_BRI\(2021\)698792](http://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2021)698792).

109. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:52021PC0206>.

110. Michael Veale & Frederik Zuiderveen Borgesius, “Demystifying the Draft EU Artificial Intelligence Act,” *SocArXiv* (July 6, 2021). Web.

111. Artificial Intelligence Act, “Article 5,” *Artificial Intelligence Act* (May 19, 2022), <https://artificialintelligenceact.com/title-ii/article-5/>; “Biometrics: Definition, Use Cases, Latest News.” *Thales Group* (Nov. 9, 2021), www.thalesgroup.com/en/markets/digital-identity-and-security/government/inspired/biometrics.

112. Mark MacCarthy & Kenneth Propp, “Machines Learn That Brussels Writes the Rules: The EU’s New AI Regulation,” *Brookings* (Mar. 9, 2022), www.brookings.edu/blog/techtank/2021/05/04/machines-learn-that-brussels-writes-the-rules-the-eus-new-ai-regulation.

113. “Câmara Aprova Projeto Que Regulamenta Uso Da Inteligência Artificial—Notícias,” *Portal Da Câmara Dos Deputados*, www.camara.leg.br/noticias/811702-camara-aprova-projeto-que-regulamenta-uso-da-inteligencia-artificial.

114. Chamber Staff, “State-by-State Artificial Intelligence Legislation Tracker,” *U.S. Chamber of Commerce* (Apr. 6, 2022), www.uschamber.com/technology/state-by-state-artificial-intelligence-legislation-tracker; “New U.S.

State Privacy Laws Zero in on Artificial Intelligence,” www.orrick.com/en/Insights/2022/08/New-State-Privacy-Laws-Zero-in-on-AI; “States and Localities Are Beginning to Move Forward with a Piecemeal Approach to AI,” JD Supra, www.jdsupra.com/legalnews/states-and-localities-are-beginning-to-8284253.

115. “UNESCO Member States Adopt the First Ever Global Agreement on the Ethics of Artificial Intelligence,” UNESCO (Apr. 21, 2022), www.unesco.org/en/articles/unesco-member-states-adopt-first-ever-global-agreement-ethics-artificial-intelligence.

116. Calum Chace, “Superintelligence May Be Closer Than Most People Think, Says Neuroscientist,” *Forbes* (Oct. 19, 2022), www.forbes.com/sites/calumchace/2022/10/19/superintelligence-may-be-closer-than-most-people-think-says-neuroscientist/?sh=1ea3e74c1217; Joe Drury, “Advanced AI Is Almost Here, Run—Towards Data Science,” *Medium* (Dec. 16, 2021), <https://towardsdatascience.com/advanced-ai-is-almost-here-run-d0d43d483d08>; Daron Acemoglu, “The AI We Should Fear Is Already Here,” *Washington Post* (July 21, 2021), www.washingtonpost.com/opinions/2021/07/21/ai-we-should-fear-is-already-here/;

117. Sean OHeigeartaigh, “Moral Outsourcing and Artificial Intelligence—Practical Ethics,” <https://blog.practicaethics.ox.ac.uk/2013/08/would-you-hand-over-a-moral-decision-to-a-machine-why-not-moral-outsourcing-and-artificial-intelligence/>.

118. *Id.*

119. *See* *Washington v. Davis*, 426 U.S. 229 (1976).

120. *Id.*; *Rogers v. Lodge*, 458 U.S. 613, 617 n.5 (1982).

121. *See discussion* in subsection What About Sentient AI?

122. Informed Consent | HealthInAging.org. www.healthinaging.org/age-friendly-healthcare-you/care-what-matters-most/informed-consent; Charles Sabatino, “Capacity to Make Health Care Decisions,” *Merck Manuals Consumer Version* (Nov. 19, 2022), www.merckmanuals.com/home/fundamentals/legal-and-ethical-issues/capacity-to-make-health-care-decisions.

123. Alec Buchanan, “Mental Capacity, Legal Competence and Consent to Treatment,” *Journal of the Royal Society of Medicine*, vol. 97, 9: 415-20 (2004), <https://journals.sagepub.com/doi/10.1177/014107680409700902>; Scott Lamont et al., “Capacity and Consent: Knowledge and Practice of Legal and Healthcare Standards,” *Nursing Ethics*, 26(1):71-83 (2019), <https://pubmed.ncbi.nlm.nih.gov/28093938/>.

124. Katie Hunt, “World’s First Living Robots Can Now Reproduce, Scientists Say,” *CNN* (Nov. 29, 2021), <https://edition.cnn.com/2021/11/29/americas/xenobots-self-replicating-robots-scn/index.html>; “Team Builds First Living Robots—That Can Reproduce,” *Wyss Institute* (Nov. 14, 2022), <https://wyss.harvard.edu/news/team-builds-first-living-robots-that-can-reproduce/>; <https://www.npr.org/2021/12/01/1060027395/robots-xenobots-living-self-replicating-copy>; Rasha Aridi, “Scientists Unveiled the World’s First Living Robots Last Year. Now, They Can Reproduce,” *Smithsonian*

Magazine (Dec. 2, 2021), www.smithsonianmag.com/smart-news/scientists-unveiled-the-worlds-first-living-robots-last-year-now-they-can-now-reproduce-180979150.

125. Hugh McLachlan, “Ethics of AI: How Should We Treat Rational, Sentient Robots—If They Existed?” *The Conversation* (June 13, 2019), www.theconversation.com/ethics-of-ai-how-should-we-treat-rational-sentient-robots-if-they-existed-118647.