A Primer on Artificial Intelligence and AI-Powered Financial Analysis

Zain Kalson

Ohio Northern University
Law Review

Lead Articles

A Primer on Artificial Intelligence and AI-Powered Financial Analysis

ZAIN KALSON*

I. INTRODUCTION

Undisclosed accounts, millions in dissipated assets, and thousands of pages of bank statements.[1] High-networth divorces, corporate investigations, and complex insolvency cases often contain vast amounts of financial data.[2] In a sea of transactions, all it takes is finding one lousy bank transfer, an accidental credit card purchase, or a suspicious check for the house of cards to come crashing down and the case to settle favorably.[3]

Traditionally, a team would spend hundreds of hours crowded around a banker's box full of discovery, manually typing bank statements into Excel.[4] Each time a number did not add up, the team would scramble to determine if it was a data entry error or an omission by opposing counsel.

Today, machine learning models are getting powerful enough to streamline both data entry and the first leg of analysis. Using these models to handle the grunt work, attorneys can filter out the noise, curate the best evidence for their arguments, and focus on high-value work.

---

* Zain is the CEO of Etheia, a legal tech company that helps attorneys and forensic accountants uncover hidden assets and analyze bank statement expenses.

1. Chris Fitzgerald & Sean Renshaw, *Fraud investigations: Uncovering hidden financial accounts and assets*, RSM (Nov. 30, 2017), https://rsmus.com/insights/services/risk-fraud-cybersecurity/fraud-investigations-uncovering-hidden-bank-and-other-financial-accounts.html.
2. *Id*.
3. *Id*.
4. *Id*.

II. HOW MACHINES LEARN

Over eighty years ago, Konrad Zuse invented the first program-controlled computer.[5] It was programmed by punch tape, and instructions were fed in raw binary (1s and 0s).[6] As early programs became more complex and binary programming became unwieldy, the need for abstraction grew. In the 1950's, high-level programming languages like Common Business-Oriented Language ("COBOL") took hold of the industry.[7] Instead of writing in binary ("000101010011"), these languages allowed engineers to create human-readable representations of data and information, like "5 + 3" or "x + y".[8]

With time, people began creating more complex programs. Using traditional software approaches, every edge case would have to be accounted for.[9] If there were a million unique scenarios that each had a different impact on the output, each scenario would necessitate custom code.[10]

Today, computing power has advanced dramatically and people are trying to represent, model, and predict complex relationships with messy, real-world data. In the context of autonomous vehicles, for example, there might be millions or billions of unique situations one runs into on the road, and it's practically impossible to manually account for every situation, geography, or time of day. In these types of systems, where the potential inputs and outputs are limitless, artificial intelligence can perform much better than traditional methods.

At its core, artificial intelligence ("AI") "is technology that enables computers and machines to simulate human intelligence and problem-solving capabilities."[11] Instead of explicitly programming every possible input and output, AI models can learn underlying patterns, including edge cases, by observing vast amounts of data in a process known as training.[12] In the prior example of self-driving cars, the model may learn general rules, like stopping at a red light, but also exceptions like cautiously proceeding if a law enforcement official waves you through a red light.

---

5. *Konrad Zuse*, IEEE COMPUTER SOCIETY, https://www.computer.org/profiles/konrad-zuse (last visited July 29, 2024).

6. Raul Rojas, *The Z1: Architecture and Algorithms of Konrad Zuse's First Computer*, ARXIV (June 2014), https://arxiv.org/abs/1406.1886.

7. *The 1950s and 1960s*, IEEE COMPUTER SOCIETY, https://www.computer.org/about/cs-history/1950-1969 (last visited July 29, 2024).

8. *Id*.

9. *Id*.

10. Andrew Bonham, *A Modern Dilemma: When to Use Rules vs. Machine Learning*, CAPITAL ONE (Aug. 5, 2020), https://www.capitalone.com/tech/machine-learning/rules-vs-machine-learning/.

11. *What is artificial intelligence (AI)?*, IBM, https://www.ibm.com/topics/artificial-intelligence (last vistied July 29, 2024).

12. Michael Chen, *What Is AI Model Training & Why Is It Important?*, OCI (Dec. 6, 2023), https://www.oracle.com/ artificial-intelligence/ai-model-training/.

The training process always starts with a dataset, which serves as the ground truth.[13] Depending on the type of model, this dataset may be annotated, or labeled, by a human.[14] Take an example of a dataset with a million pictures of cats and a million pictures of dogs. If the dataset was labeled, each image would be associated with a tag of "cat" or "dog" by a human. In supervised models, these labels are used as correct answers.[15] If the output model is given a picture of a cat that was not in the training data, it should nonetheless predict the label of "cat".[16] In unsupervised models, no annotations are provided, and the model attempts to learn underlying structures and relationships in the data.[17] An unsupervised model might be used to generate novel pictures of cats and dogs.[18]

The quality and scope of the dataset directly correlates with the performance of the resulting model.[19] If the goal of a model is to generate all forms of natural language, one would want to include diverse data, ideally at the scale of the internet. GPT-2, one of the precursors to the models used in ChatGPT, was trained on hundreds of millions of words.[20] A specialized model identifying cats and dogs in images would need training samples with different breeds, positions, activities, and backgrounds to be accurate.

At first glance, it may seem puzzling how patterns get learned and embedded into the model. Today, most models use deep learning, which is a process that extracts features or characteristics of the data through increasing layers of abstraction.[21]

To conceptualize how models understand complex data, one can imagine a rectangular grid of millions or billions of parameters, or dials. Each dial can be turned up or down to increase or decrease its impact on the output. Columns in this rectangular grid represent layers in the network. In an image model, the earliest layers may learn to detect things like lines, shapes, primary colors, and stripes.[22] As data flows through the model and to later layers, more nuanced features, like complex shapes, scenes, and objects, become

---

13. *Id*.
14. *Id*.
15. *Supervised vs. unsupervised learning: What's the difference?*, GOOGLE CLOUD, https://cloud.google.com/discover/supervised-vs-unsupervised-learning (last visited July 29, 2024).
16. *Id*.
17. *Id*.
18. *Id*.
19. *What is AI?*, ISO, https://www.iso.org/artificial-intelligence/what-is-ai#toc6 (last visited July 29, 2024).
20. Alec Radford, et al., *gpt-2-output-dataset*, GITHUB, https://github.com/openai/gpt-2-output-dataset (last visited July 29, 2024).
21. *What is AI?*, *supra* note 19.
22. Tim Dettmers, *Deep Learning in a Nutshell: Core Concepts*, NVIDIA (Nov. 3, 2015), https://developer.nvidia.com/blog/deep-learning-nutshell-core-concepts/#deep-learning.

represented in the parameters.[23] In a model discriminating between cats and dogs, later layers may focus on the shape of the face, the structure of the ears, and the qualities of the fur.[24]

Nuanced understanding comes from having well-adjusted parameters that appropriately model the relationships in the data.[25] To adjust these parameters, the training process goes through every sample in the dataset, identifies how the parameters can be adjusted for maximum total accuracy, and makes those adjustments.[26] After all the samples have been processed, the dials are tuned so that the model produces reasonable and generally correct outputs.[27]

Critically, this training process can embed bias that exists in the dataset into the final model.[28] In the previous example, if all the images of dogs were outdoors and all the cats were indoors, then the model might learn to associate walls, windows, and carpets with cats. When an indoor picture of a dog is provided, the model might confidently predict that it is a cat because it learned the wrong features.

III. NATURAL LANGUAGE PROCESSING

Natural Language Processing ("NLP") is a field in machine learning that enables machines to have a complete, nuanced, and rich understanding of language.[29] Human language is filled with homonyms, homographs, and intent.[30] In NLP, models can decipher rich intent by representing words with embeddings, or vectors in a high-dimensional space.[31] For example, in the sentence "I ate an orange orange," models can discriminate between subjects and descriptors, the model can understand that the first instance of orange is describing the color of the second instance. Imagine a graph with two axes. In this example, one axis is cost and the other is speed. The phrase "car" is neutral in intent, appearing at the origin (point 0,0). BMW is a German vehicle brand associated with luxury and speed, and BMW cars tend to be

---

23. *Id*.
24. *Id*.
25. Andrew Ng and Kian Katanforoosh, *Deep Learning*, STANFORD, https://cs229.stanford.edu/summer2020/cs229-notes-deep_learning.pdf (last Updated July 22, 2019).
26. *Id*.
27. *Id*.
28. Adam Zewe, *Can machine-learning models overcome biased datasets*?, MASSACHUSETTS INSTITUTE OF TECHNOLOGY (Feb. 21, 2022), https://news.mit.edu/2022/machine-learning-biased-data-0221.
29. *What is Natural Language Processing* (*NLP*)?, AMAZON WEB SERVICES, https://aws.amazon.com/what-is/nlp/ (last visited July 29, 2024).
30. *Id*.
31. Joel Barnard, *What are word embeddings?*, IBM (January 23, 2024), https://www.ibm.com/topics/word-embeddings.

expensive.[32] The phrase, "BMW car," will likely appear up and to the right of the center, around point 4,5, because the BMW vector "adds" both cost and speed when used as a descriptor. If the graph had another axis for association with Germany, the phrase "BMW car" would be more positive on that axis than the "car", too. If there were yet another axis representing cheap construction, "BMW car" would appear to the left, more negative than that of "car". Already at four dimensions, this graph is nearly impossible to visualize. Today, many NLP models have thousands of dimensions.[33]

Because of the vast complexity of language, most models today start training from a base model checkpoint while fine-tuning it for specialized use cases.[34] One of the most common base models for NLP is Bidirectional Encoder Representations from Transformers ("BERT"), which was trained on thousands of books and Wikipedia.[35] Models based off BERT have been adapted for a variety of use cases, which are described in the following subheadings.[36]

A. Language Translation

Language translation has historically been challenging because words in one language may not directly map to words in another language. For example, the English idiom "break a leg" would not make sense in many other languages if the words were literally translated. AI-powered machine translation allows one to translate speaker intent and deeper meanings accurately.[37]

B. Text Classification

Text classification models allow data to be tagged with a set of pre-trained categories.[38] These categories can be almost anything, so long as there is enough training data, text classification models are highly effective.[39] One benefit to using classification models over Large Language Models

---

32. *Luxury SUVs, Sedans, Coupes, Convertibles, & Crossovers*, BMW USA, https://www.bmwusa.com/ (last visited Aug. 5, 2024).

33. Jacob Devlin, et al., *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*, ARXIV, https://arxiv.org/abs/1810.04805 (last modified May 24, 2019).

34. *Id*.

35. Devlin, *supra* note 33; *BERT base model (uncased)*, HUGGING FACE, https://huggingface.co/google-bert/bert-base-uncased (last visited Aug. 5, 2024).

36. Devlin, *supra* note 33.

37. Jinhua Zhu, et al., *Incorporating BERT into Neural Machine Translation*, OPEN REVIEW (2020) https://openreview.net/pdf?id=Hyl7ygStwB; *Machine Translation*, NVIDIA, https://catalog.ngc.nvidia.com/orgs/nvidia/collections/machinetranslation (last modified May 17, 2024).

38. Chen, *supra* note 12.

39. *See generally* Zhu, *supra* note 37.

("L.L.M.'s") like GPT-4 is that these models have deterministic outputs.[40] In other words, it is impossible for a classification model to produce categories that were not in the training dataset, whereas that robustness is difficult to produce in L.L.M.'s.[41]

C. Named Entity Recognition ("N.E.R.")

N.E.R. models allow users to quickly identify people, places, dates, and more from text.[42] For example, in the phrase "John Doe works at ACME Inc," an N.E.R. model would identify John Doe as a person and ACME Inc. as a company. When dealing with highly unstructured data, such as discovery documents, business memos, emails, text messages, or transcripts, identifying named entities can streamline the review processes.

D. Sentiment Analysis

Sentiment analysis allows for a nuanced understanding of the emotion and tone behind the text.[43] At a foundational level, models can understand if the text is positive, negative, or neutral, and more sophisticated models can be trained to detect complex emotions such as sarcasm.[44]

E. Generative AI

Text generation models are excellent generalists at a variety of language tasks. With L.L.M.'s, generative AI models can also be powerful few-shot or zero-shot learners, where when given a few examples, the model can perform language and reasoning tasks.[45] Generally, few-shot and zero-shot learning is less accurate than models trained with a large dataset.[46]

F. Summarization and Question Answering

Text summarization and question-answering models come in a variety of forms. Some use L.L.M.'s while others rely on more traditional machine

---

40. *See generally* Takeshi Kojima, et al., *Large Language Models are Zero-Shot Reasoners*, ARXIV, https://arxiv.org/pdf/2205.11916 (last modified Jan. 29, 2023).

41. Shyamai Anadkat, *How to make your completions outputs consistent with the new seed parameter*, OPENAI COOKBOOK (Nov. 5, 2023), https://cookbook.openai.com /examples/reproducible_outputs_with_the_seed_parameter.

42. Jing Li et al., *A Survey on Deep Learning for Named Entity Recognition*, ARXIV, https://arxiv.org/abs/1812.09449 (last modified Mar. 18, 2020).

43. *See What is Sentiment Analysis?*, AWS AMAZON, https://aws.amazon.com/what-is/sentiment-analysis/ (last visited Aug. 5, 2024).

44. *Id*.

45. *See* Kojima, *supra* note 40.

46. *Id*.

learning approaches.[47] By understanding what information in a text is most important, these models can distill long-form content.[48]

IV. COMPUTER VISION

To machines, images and videos are just pixels on a screen. Understanding people, objects, text, and places from those images is a big part of Computer Vision ("C.V."). Image models today are very robust, but video models still have a long way to go to be production ready. Similar to how text can be embedded into a high-dimensional space, visual data can be too.[49] In a series of four images, including two images of cats, one image of a dog, and one image of a plane, the images of the cats would be close together and the dog would likely be closer to the cats than the plane.

Vision Transformer ("ViT") is a popular base model for image tasks.[50] Some use cases of ViT and other models can be fine-tuned on, to include those listed below.

A. Image Classification

Similar to text classification, image classification models can assign tags, or labels based on predefined categories. With enough training data, the model will learn to distinguish features between the different categories to perform classifications accurately.[51]

B. Object Detection

Object detection refers to defining classes of objects and recognizing the objects in images. Common examples might be recognizing stop signs, pets, or phones in images. These models can become extremely sophisticated and precise. For example, segmentation models can individually pick out the exact pixels occupied by an object, whereas standard approaches simply draw a box around the detected object.[52] This can also be used for extremely

47. *Id*.
48. Randy DeFauw, *Use a generative AI foundation model for summarization and question answering using your own data*, AWS AMAZON (July 29, 2023), https://aws.amazon.com/blogs/machine-learning/use-a-generative-ai-foundation-model-for-summarization-and-question-answering-using-your-own-data/.
49. Omar Espejel, *Getting Started With Embeddings*, HUGGING FACE (June 23, 2022), https://huggingface.co/blog/getting-started-with-embeddings.
50. *See generally Vision Transformer and MLP-Mixer Architectures*, GITHUB, https://github.com/google-research/vision_transformer (last visited Aug. 5, 2024).
51. *ML Practicim: Image Classification*, GOOGLE, https://developers.google.com/machine-learning/practica/image-classification/ (last modified July 18, 2022).
52. Segment Anything Model (SAM): a new AI model from Meta AI that can "cut out" any object, in any image, with a single click, SEGMENT ANYTHING, https://segment-anything.com/ (last visited Aug. 5, 2024).

precise image editing like changing the background or removing a person orobject from the image.[53]

C. Object Tracking

With moving entities, such as people, animals, and cars, it may become useful to identify sequences of detections. In other words, if someone walked from one side of a security camera's coverage to the other side, object tracking models can lock in on individual entities and track their change in position across space.[54]

D. Image Search

Models can perform image similarity searches by finding other images that are close in the embeddings space.[55] Having properly trained embeddings is critical to ensure the search works well for the specified use case.[56]

E. Image Captioning

One powerful element of AI models today is their ability to identify what is important in an image, this is also called attention.[57] Using a multimodal embeddings space, which will be discussed later in the article, models can describe key areas of an image in natural language.[58]

F. Text Detection

Oftentimes, documents are not stored in the native text form. Scanned PDFs may visually present text, but that text is not selectable. Converting a screenshot, scan, or picture of text into a machine-readable form is a complex process that requires building specialized object detection and classification models. This process is often called Optical Character Recognition

53. *See generally* Lingyi Hong, et al., *OneTracker: Unifying Visual Object Tracking with Foundation Models and Efficient Tuning*, ARXIV (Mar. 14, 2024), https://arxiv.org/pdf/2403.09634.
54. *Id*.
55. Nilesh Verma & Sai Kumar, *Deep Image Search – AI-Based Image Search Engine*, GITHUB, https://github.com/TechyNilesh/DeepImageSearch?tab=readme-ov-file (last visited Aug. 5, 2024).
56. Verma & Kumar, *supra* note 55; Hong, *supra* note 53.
57. *Image captioning with visual attention*, TENSORFLOW, https://www.tensorflow.org/text/tutorials/image_captioning?hl=en (last visited Aug. 5, 2024).
58. Hong, *supra* note 53.

("O.C.R.").[59] As one of the first uses of machine learning, O.C.R. today is very robust, with many models exceeding 99.5% accuracy on typed data.[60]

G. Table Detection

Sometimes, structured data is embedded in documents. Bank statements and medical records often have table data. While O.C.R. can recognize text within these tables, O.C.R. cannot meaningfully extract the table itself.[61] Specific table detection models can look for visual cues as to the structure of the table, like borders or cell spacing, and combined with O.C.R. results, table data can be directly extracted into a structured form, such as Excel.[62]

H. Generative AI

As with text, images and videos can be generated from unsupervised models. Much work has been done to allow these generations to be steerable from a natural language prompt.[63]

V. MULTIMODAL MODELS

Multimodal models allow you to train on a dataset with multiple types of data. For example, one could train a document processing model that can use both an image of the document and the text from the document to then categorize it. By training with multiple modalities, a complex embedding space can be achieved. The text "cat" is likely to be similarly embedded to an image of a cat, and this depth of understanding can allow models to ingest text, images, or audio and output in either or all of the modalities![64]

VI. AUTOMATED SORTING

While it is not always the case, it is possible to receive discovery responses in disorganized, paper form from opposing counsel. Additionally, some banks will only respond to subpoenas in paper form. Given these issues, the financial component of a case file can quickly become messy and unmanageable. With complex cases often having tens of thousands of

---

59. Jill Daley & Esther Adediran, *What is Optical Character Recognition? OCR Explained by Google*, GOOGLE CLOUD (Sep. 13, 2023), https://cloud.google.com/blog/products/ai-machine-learning/what-is-ocr/.

60. Cem Dilmegani, *OCR in 2024: Benchmarking Text Extraction/Capture* Accuracy, AI MULTIPLE RESEARCH, https://research.aimultiple.com/ocr-accuracy/ (last updated Feb. 14, 2024).

61. *See generally* Brandon Smock & Rohith Pesala, *Table Transformer (TATR)*, GITHUB, https://github.com/microsoft/table-transformer (last visited Aug. 5, 2024).

62. *Id*.

63. *DALL-E 3*, OPEN AI, https://openai.com/index/dall-e-3/ (last visited Aug. 5, 2024).

64. Rohit Girdhar, et al., *IMAGEBIND: One Embedding Space To Bind Them All*, ARXIV, https://arxiv.org/pdf/2305.05665 (last modified May 31, 2023).

documents, it can take weeks of manual work to sort by hand. Once digitized through a process such as O.C.R, AI models can automatically sort documents by either classifying each page, for example, page 1 is Bank of America, page 2 is Chase, and so on, or clustering similar pages together, for example, pages 1 and 7 are of the same bank, pages 2-6 are for another bank, and so on.[65] Models can categorize documents based on the document's text, the visual appearance, and structure of the document, or a hybrid approach.[66]

It may also be useful to adopt fine-tuned, specific classification models for different types of cases. In divorce proceedings, it might be relevant to have categories for text messages and household bills. However, in a bankruptcy case, it may be useful to have a designated category for loans and creditor paperwork.

Once the documents are sorted, they can be further structured using N.E.R. and sentiment analysis models.[67] For example, it may accelerate the document review process to have every page tagged with the companies, the people mentioned, and the tone. Combining different types of models can allow for a highly robust Technology Assisted Review ("TAR"), where individuals can filter documents based on a variety of criteria.[68]

VII. DATA EXTRACTION/RECONCILIATION

Once the financial documents have been sorted, the documents can be analyzed in-depth using AI. Financial statements contain tables with transactions, beginning and ending statement dates, beginning and ending balances, and account numbers, all of which can be extracted. State-of-the-art results usually require millions of data points for training and multiple specialized models.[69] There may be one model for detecting tables, one for detecting rows and columns inside of tables, and one for extracting key values like statement period and account number.[70]

There is a host of pre-built, open-source datasets and models for table and key-value extraction, though these offerings use publicly-available data, like research papers, SEC filings, and other data, or synthetic data, like artificially

---

65. *See generally* Rajiv Shah, et al., *Accelerating Document AI*, HUGGING FACE (Nov. 21, 2022), https://huggingface.co/blog/document-ai.

66. *Id*.

67. *See supra* notes 42 & 43 and accompanying text.

68. *See generally Myths and facts about technology-assisted review*, THOMSON REUTERS, https://legal.thomsonreuters.com/en/insights/articles/myths-and-facts-about-technology-assisted-review (last visited Aug. 5, 2024).

69. *See generally* Chen, *supra* note 12.

70. Prithiv S., *OCR for data extraction from bank statements*, NANONETS, https://nanonets.com/blog/ocr-for-data-extraction-from-bank-statements/ (last updated Mar. 27, 2024); Sanjana Ramachandran & Neetha K., *What is financial data extraction?*, NANONETS, https://nanonets.com/blog/financial-data-extraction/ (last updated Mar. 25, 2024).

generated bank statements.[71] To create models that are performant on financial statements, one can fine-tune these open-source offerings with their own custom data.

After all the information has been extracted, traditional software such as Microsoft Excel can be used to reconcile the data by summing the beginning balance with all the transactional activity in each statement, this creates a calculated balance. If the calculated balance matches the ending balance and the statement is reconciled, then one can have high confidence that all transactions were properly extracted.[72] Using conditional formatting, it is possible to visually indicate when the balances match, quickly providing input to the user that the reconciliation was performed correctly.[73] Doing this at scale with high-quality O.C.R. and data extraction, one can reconcile hundreds of thousands of transactions in an hour as opposed to weeks.

Another component of the data processing pipeline may include check O.C.R. and reconciliation.[74] Oftentimes, financial institutions will include images of cashed checks in the statement. By reading the "Pay To The Order Of" line, one can glean valuable information about the recipient of a given transaction. Using a dedicated check recognition object detection model and an O.C.R. model trained specifically on handwriting, one can add the check payee into the description field.[75] In commercial matters, a similar process may take place with invoices or other business records that can provide additional context as to the purpose of certain transactions.

VIII. TRANSACTION ENRICHMENT

With a reconciled Excel file, one can begin to analyze the transactions. This primarily consists of identifying merchants and categorizing transactions.

For this transaction, "Purchase Authorized 10/21 Wal-Mart SC AH2148 Card 1234," the merchant would be Wal-Mart Supercenter, and the category would be Grocery. Adding these as columns in Excel allows one to quickly filter transactions and produce summary-level information. In production systems, one might expect hundreds of categories to allow for fully customizable searches.

---

71. *See* Minghao Li, *TableBank: A Benchmark Dataset for Table Detection and Recognition*, GITHUB, https://github.com/doc-analysis/TableBank (last visited June 1, 2024).

72. Tim Maxwell, *How to Reconcile Your Bank Statements*, EXPERIAN (Feb. 28, 2024), https://www.experian.com/blogs/ask-experian/how-to-reconcile-your-bank-statements/.

73. *Id*.

74. *Id*.

75. *See generally* A K Nain & Sayak Paul, *Handwriting recognition*, KERAS, https://keras.io/examples/vision/handwriting_recognition/ (last modified July 7, 2023).

Depending on the type of litigation, different categorization models may be needed. In a marital waste claim, a forensic accountant may create a report with only hotels, family business distributions, luxury goods, and remittances.[76] In an insolvency case, it might be relevant to see all bank transfers to identify accounts that still may contain assets. In commercial matters, it's possible to build a categorization model that can partially reconstruct a cash-basis business accounting using only the bank/credit card statements.[77]

Traditional approaches to entity recognition and transaction categorization use lookups to match specific keywords, for example categorize as grocery if Wal-Mart is in the description, though these fail to work well with broad or non-specific transactions and require having data on any business that could show up.[78] Machine learning models trained specifically on financial data can perform much better.[79]

One key element of a transaction categorization system is injecting context.[80] With millions of potential merchants, there needs to be a way to associate types of merchants with categories. For example, if Wal-Mart is categorized as a grocery store, one would want Kroger, Albertsons, Whole Foods, and other supermarkets to be similarly classified, even if they did not appear in the dataset used for training. To do this, the transaction categorization model can be injected with additional context such as North American Industry Classification System ("NAICS") codes or business database information.[81] By providing this supplementary information, the model learns to use additional features to make better predictions, rather than just the transaction description.[82]

IX. CLOUD VS ON-PREMISES

Certain cases may have unique confidentiality and data sensitivity requirements.[83] Sometimes, all parties agree to never process data on the

76. Barbara A. Ruth & Michael A. Gillen, *The Role of Forensic Accountants in Divorce Engagements*, DUANEMORRIS (Feb. 2, 2010), https://www.duanemorris.com/articles/forensic_accountants_divorce_3548.html.

77. *Id*.

78. *See generally* Jing Li et al., *supra* note 42.

79. *See generally Forensic analytics in fraud investigations*, DELOITTE, https://www2.deloitte.com/us/en/pages/advisory/articles/forensic-analytics-in-fraud-investigations.html (last visited Aug. 5, 2024).

80. *See generally* Jing Li et al., *supra* note 42.

81. *See generally North American Industry Classification System*, UNITED STATES CENSUS BUREAU, https://www.census.gov/naics/ (last visited June 1, 2024).

82. *See generally* Jing Li et al., *supra* note 42.

83. *See generally* Michael Richey, *Secure Couriers Vital Role in Securing Confidential Information*, TEXAS DEFENSE FORCE, https://www.txdf.org/resources-and-articles/securing-confidential-information-the-vital-role-of-secure-co uriers (last visited June 1, 2024).

cloud.[84] This means that discovery documents can never be sent to a third party via the Internet. To meet this requirement, some vendors offer solutions that can run on one's own computer.[85] By performing local inference, the risk of data leakage is minimized, and the end user has much greater governance, or control over the data.[86] However, some larger models are unable to run on end users' systems, so vendors have to distill models, which reduces their complexity.[87] While this may allow models to run on local systems, users will generally have to deal with reduced accuracy.

For highly sensitive cases, an airgap, an alternative solution, is a computer system disconnected from the internet.[88] With the internet being the primary mode of cybersecurity breaches, it becomes exceedingly difficult for an unauthorized individual to gain access to sensitive data on the system.[89] Generally, users will send data over encrypted drives via a secure courier transfer, where the physical data storage medium is always being monitored by security personnel.[90] While these are extreme security measures and only ever used on a minority of cases, they can allow for the most confidential and highly restricted cases to benefit from advances in AI.

X. CONCERNS AROUND TRAINING DATA AND MODEL OWNERSHIP

As previously established, AI models are extremely data-hungry.[91] Without a large and diverse dataset for training, results will always be suboptimal.[92] For many vendors, in order to build a large dataset, the vendors outsource much of their data labeling to third-party, offshore workforces.[93] Depending on the type of data being labeled, the consequences can range from negligible to extremely severe.

Often, there is very little oversight over these third-party workforces, and in some instances, labeling companies are given the right to use all data provided to them for their own models.[94] In the legal world, this can be a massive breach of confidentiality. Vendors must provide robust Data

84. *Id*.
85. *Id*.
86. *Id*.
87. *Id*.
88. *Air Gap*, NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY, https://csrc.nist.gov/glossary/term/air_gap (last visited June 2, 2024).
89. *See generally* Richey, *supra* note 83.
90. *Id*.
91. *See supra* note 12 and accompanying text.
92. Chen, *supra* note 12.
93. *Id*.
94. *See generally California Consumer Privacy Act (CCPA)*, CALIFORNIA DEPARTMENT OF JUSTICE OFFICE OF THE ATTORNEY GENERAL, https://oag.ca.gov/system/files/attachments/press_releases/CCPA%20Fact%20Sheet%20%2800000002%29.pdf (last visited June 1, 2024).

Protection Agreements ("DPA") to their customers. These agreements may: (1) contractually guarantee data privacy rights, (2) contractually guarantee security policies and best practices, (3) explicitly incorporate privacy laws such as CCPA or GDPR,[95] (4) require subcontractors to have an equally protective DPA, (5) restrict overseas data transfers,[96] or (6) require that data labeling be done in-house.

In some instances, vendors may "WRAP" third-party services.[97] Instead of building functionality in-house, the vendors will utilize third-party models to infer customer data.[98] While there is nothing inherently wrong with this practice, it is important to understand the data privacy and governance policies of those third parties, if applicable. Unchecked, this can easily lead to the data risks previously mentioned.

XI. CONCLUSIONS

With machine learning, practitioners can substantially increase the leverage they have over their practice. What once would take weeks, can now be done in hours. Using AI to uncover the truth, attorneys can spend less time bogged down by administrative tasks and more time achieving positive outcomes for their clients.

95. *California Consumer Privacy Act (CCPA)*, *supra* note 94; *See General Data Protection Regulation*, INTERSOFT CONSULTING, https://gdpr-info.eu/ (last visited Aug. 5, 2024).

96. *See generally Cross-Border Data Policy Principles*, GLOBAL DATA ALLIANCE, https://globaldataalliance.org/wp-content/uploads/2021/07/03022021gdacrossborderdatapolicyprinciples.pdf (last visited Aug. 5, 2024).

97. George Fitzmaurice, *What is WRAP and how can it help train AI more efficiently?*, ITPRO (Aug. 26, 2024), https://www.itpro.com/technology/artificial-intelligence/what-is-wrap-ai.

98. Derek Wood, *Challenges When Building and Monetizing an Inference Model*, Duality (Apr. 4, 2024), https://dualitytech.com/blog/challenges-when-building-and-monetizing-inference-models/.