AI Patents: Maximizing Protection and Licensing Value

Introductions

Robert Plotkin

John Brenner

Robert Plotkin, Patent Attorney

- Patent attorney specializing in software patents for 25+ years
- Recognized leader in AI patents
- Co-founding partner of Blueshift IP, a boutique patent firm
- Represents universities nationwide
- Clients have licensed and sold patent portfolios worldwide
- Alumnus of MIT (Computer Science) and Boston University School of Law
- Author of *The Genie in the Machine* (book on AI patents, Stanford University Press, 2009)

John Brenner

- 25+ years consulting services with various aspects of early-stage technology companies.
  - In conjunction with:
    - 7 years as a technology manager at Cornell University.
    - 2 years as VP microfluidics at Kionix (a Cornell startup).
    - 13 years and counting as VP of IP management & licensing at Rheonix (a Kionix spinout).
- Alumnus of the University of New Hampshire (BS – Plant Science), Syracuse University (MS) & Cornell University (MBA).

Let's Get Started!

The conversation begins.

Summary of AI

AI today often involves using an automatically learned model, rather than human-designed rules or algorithms, to perform processes.

- Refers to seemingly intelligent behavior by machines
- No clear definition, but signs to look for in disclosures:
  - "Machine learning"
  - "Deep learning"
  - "Neural networks," e.g., "convolutional neural networks (CNNs)"
  - "Training a model" (e.g., a language model)
  - "Evolutionary algorithm"
  - "Genetic algorithm"
  - "Generative adversarial network (GAN)"

- Common functions performed by AI
  - Learning and matching patterns, e.g.:
    - image recognition, speech recognition
    - anti-patterns (e.g., software viruses)
  - Predicting: locations, disease, evolution of a phenomenon (e.g., weather, financial markets, traffic)
  - Optimizing: scheduling, logistics, travel, space
  - Controlling: robots, vehicles, software

What could/should a disclosure look like?

The conversation continues.

A thousand words.

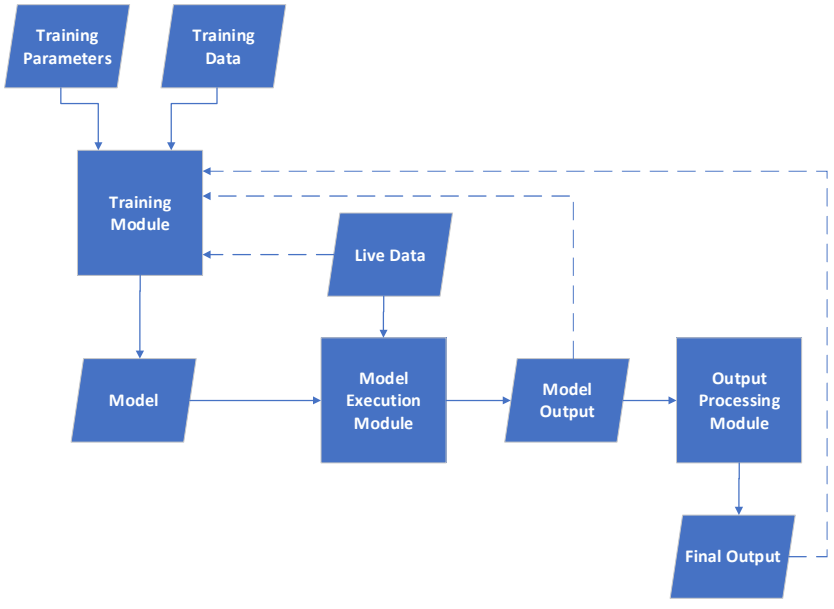

Summary of the disclosure

- Go beneath the surface of the disclosure (disclosure diligence)
  - Was any AI used to develop what has been disclosed?
  - Is any AI used when practicing what has been disclosed?
- How to best prepare a disclosure for an AI innovation?
  - Consider each of the basic ways of claiming an AI innovation and include at least a high-level description of each relevant aspect of the innovation
  - If a model was trained, best to include written description and enabling details of how that model was trained, including description of training data, model, model parameters, etc.
  - If the selection/organization of training data is critical to the invention, describe that selection/organization.
  - Many AI innovations are now cross-disciplinary, e.g., computer science and biology. A cross-disciplinary team may be useful for strategy and/or drafting.

A hypothetical

- Problem is how to more accurately find locations where geological deposits are located.
- We trained a neural network using geological mining training data that included proprietary data generated in a study and [Data Source D1].
- We applied the neural network to new geological data obtained from [Data Source D2] to produce output.
- We applied method [Y] to the output to [filter/sort/summarize] the output, and thereby learned [Z].
- We applied algorithm [A] to output [Z] to produce final results indicating high-probability sources of mineral [V].
- We could update the model above over time based on new geological data.

Any questions yet?

Summary of claiming

- Some examples of subject matter to claim are:
  - a method of training a model
  - a method of using the trained model
  - a method of updating the model based on uses of the model
  - the output of the model (possibly using a product-by-process claim)
  - the model itself (possibly using a product-by-process claim)
- Important to explore each of the above and ask:
  - Is it patentable?
  - What is its licensing value?
- How is the evolution of the AI captured in claims?
  - Will the claims only cover the initial version of any of the above, or will they also cover updates (e.g., to the model) over time?
- Do you focus on the application of the AI? Or the AI itself?
  - See above. Could be either or both. It depends on what is innovative in each case, and the value of the claims in each case.
  - Always worth considering whether any of the above aspects should be retained as trade secrets. Could patent some, keep others as trade secrets.
- Should you divide the claims early?
  - In general, software examiners (who are most likely to be assigned to these patent applications, even if they relate to life sciences, e.g., AI for drug discovery) tend to divide whenever possible these days. Therefore, it is best to divide early.
  - Must be careful to provide sufficient support in the specification for future claims, even if they are not filed at the outset.
  - If multiple issued patents are desired soon, then consider filing multiple applications simultaneously or in rapid succession.

Summary of claiming - continued

- Short, clear claims in which elements are easy to map to products and in which infringement is easy to detect. Easier said than done! Easier to say what **not** to do, e.g.:
  - Avoid overly long claims with unnecessary details
  - Avoid field-of-use restrictions to maintain breadth
  - Avoid preambles with unnecessary restrictions
  - Avoid limitations directed to specific types of AI (e.g., neural networks) except when necessary
- As with any claim, go through every word in each independent claim and ask:
  - Is this necessary? Could someone implement the invention without this?
  - Could this term be rephrased more broadly?
  - "Automatically does X based on large data sets"
- Make sure the specification has support for claims that aren't included in the initial application, in order to keep future options open
  - Satisfying the written description and enablement requirements for AI-related inventions can be tricky.
  - Need to evaluate on a case-by-case basis in close collaboration with the inventors.
  - For example, may need to disclose data, experiments, and training in detail – more like a chemical or biological patent application than a typical software patent application.

Any questions?

Summary of Licensing

- The training model?
  - Consider the similarity to life science with data in place of a plasmid or cell line.
    - Who owns the data and what obligations are owed the data owner?
    - Is the data a part of what is licensed if it is only used to train the model?
  - How to handle new data to update the training model?
- The model?
  - What is the role of the licensor in respect to updating the training model upstream of the license?
  - Can you license the model exclusively to multiple parties who then diverge into their own fields-of-use as the model evolves?
- Or applications?
  - How to write the license to account for the divergence of the license over time?

- Is there a claiming format that helps lead to licensing partners?
  - Consider adding further independent or dependent claims which include
    - Industry specific terminology
    - Field-of-use restrictions
    - Output directed at commercial applications
  - Ask the question – How will licensees likely implement this?
    - Then cover those specific cases in dependent claims and/or separate independent claims.
- Improvements?
  - How are improvements defined? Any difference at all?
  - How are improvements related to the evolution of the model?
- How can licensees monitor for infringement?
  - How do you see into the black box?
  - Clues to watch for – "proprietary machine learned model", "automatically does (X) in conjunction with (Y) datasets" or solves the same problem as the claimed invention automatically.

Any final questions?

Thank you.