

Virtual Coffee with Consulting: Boise Valley ARMA Edition

Fundamentals of Artificial Intelligence for
Non-Computer Scientists



March 10th Q & A

Questions and Answers from Fundamentals of Artificial Intelligence for Non-Computer Scientists – Answered by Maura Grossman

1. **Question:** Could you speak to challenges of categorization in multi-use / multi-user and ambiguous contexts?

Answer: Different users will tend to code documents differently. That's inherent in all document review and we've known that for many years. Most state-of-the-art algorithms like logistic regression and support vector machines can handle some noise and inconsistency as long as it is not extreme. But if the categories are so ambiguous that a human would have trouble distinguishing them, so will an AI tool.

2. **Question:** What needs to be in place to ensure adequate levels of transparency and explainability are in place when using AI?

Answer: While there is a movement in AI called Explainable AI or XAI, I tend to be a little less concerned with transparency and explainability than with demonstrated validity and reliability. We get on airplanes all the time without understanding how they work; we take medicines without understanding their biological and pharmacological mechanisms. We do that because there has been someone responsible for testing and certifying them as safe for their intended purposes (i.e., the FAA or the FDA). We do not have enough of that type of validation in AI and it is critical, especially where explanation is impossible. Some algorithms—like the deep learning that I talked about—are opaque

even to their developers and simply cannot be explained (yet). Then the question becomes one of tradeoffs: Would you prefer a more accurate algorithm to diagnose your brain tumor or are you willing to give up 15% accuracy to use a more transparent algorithm? The answer to this may depend on context. There may be some uses (e.g., sentencing) where explanation is the more important value.

3. **Question:** I'd love to get recommendations on where to start some basic self-directed education on understanding AI tools and how to use them in a general context. Could you provide some guidance here?

Answer: I'd probably begin with some online webinars on introduction to AI or to machine learning. Basic computer science textbooks are probably too hard for a non-computer scientist. This blog is a good one to follow for all things AI: <https://swisscognitive.ch/>. You can find a copy of the paper I wrote where we used machine learning and beat the Virginia State Archivist at his own game [here](#).