White Paper

Artificial Intelligence
Opportunities for innovation and advancement

By Darby Kidwell

MACEDON
TECHNOLOGIES

Table of Contents

Defining Artificial Intelligence

Artificial intelligence (AI) is an idea with many connotations behind it. When we think of what AI is, we understand that it is a powerful concept that has helped solve some of the most complex problems computer scientists face. The complexity of these solutions and the broadness of their applications tends to obfuscate our ability to see how artificial intelligence blends with, and complements, the business applications that exist today. Throughout this white paper, we will discuss what artificial intelligence is; its effects on businesses and the benefits AI has when incorporated into enterprise systems; and how to begin with the incorporation and adoption of AI into current business applications.

Understanding Artifical Intelligence

Artificial intelligence is more than just self-driving cars or image detection. It is a way to gain insights into data. As such, any mathematical model that helps to gain insights or provide classifications, can be considered a machine learning algorithm. A few commonly studied statistical models are: linear regression, logistic regression, decision trees, k-nearest neighbors, random forests, or neural networks.

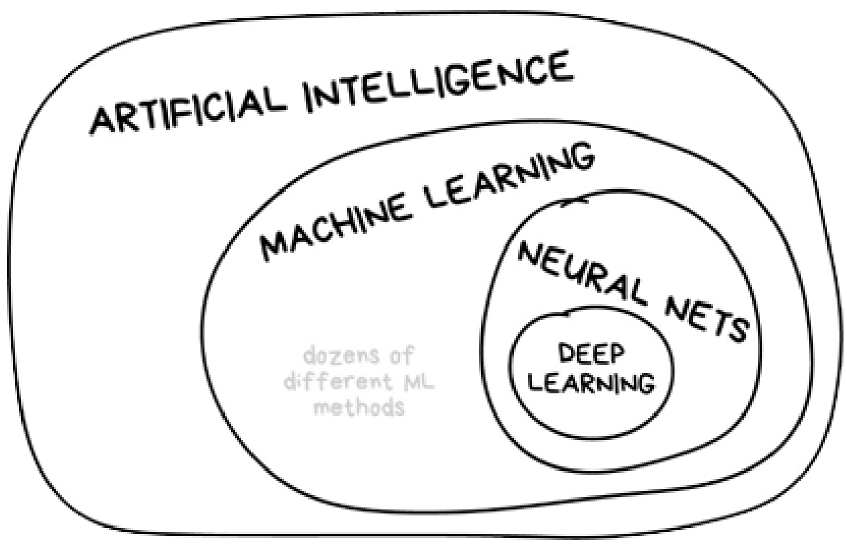

These models, like all machine learning models, take in a set of training data, and attempt to find a heuristic to predict an answer. The heuristic built by the model and how each algorithm optimizes for it varies. For example, in a linear regression model, the heuristic created is a linear equation, where the distance from each data point to this linear line is minimized. However, for a k-Nearest Neighbor algorithm, the heuristic built is a set of nodes where related nodes are all connected, so that when novel test data is provided, the algorithm can find the k closest features to this novel data and classify this element as the feature that was most in common with it.

Implementation Space

There are two fundamental things needed in order to start any machine learning (ML) endeavor. The first is the requisite amount of data to train the models and the second is an understanding of which algorithms to use for your model, as each approach tends to have its pros and cons. Fortunately for application developers, many of these problems have been wrapped into libraries or are their own products. This opens up the implementation space to two areas: incorporating pre-existing ML solutions into an application or implementing and training a custom ML model that can be incorporated into an application.

Pre-existing ML Solutions

While ML solutions tend to be highly customized, it is worth noting that many pre-existing ML solutions solve their respective problems extremely well and have broad applications. Google itself offers a suite of AI algorithms, which range from vision and object detection, to speech-to-text and language translation, and even dialogue and conversation creation. Leveraging these existing ML solutions into current business applications can in-and-of themselves have transformative results.

Custom ML Solutions

However, while these ML solutions offer much, they don't solve all problems and so, in these circumstances, custom ML solutions may be

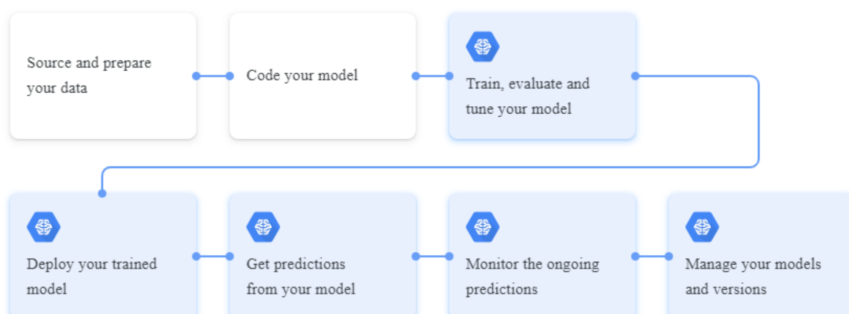

ML workflow

the answer. As previously noted, the first thing needed when creating custom ML algorithms is data. Lots of data. In addition, the second thing needed is an understanding of which model would fit the described problem. There is no golden-rule for choosing models, and many problems require testing implementations of various models to find the one with the highest efficacy. Despite no golden rule, there are still heuristics that can be followed, such as neural networks which tends to do best with ample data and complex feature detection (think looking for features in an image that identify what is pictured); whereas, logistic regressions tend to work better on more bound datasets that have simpler classification goals (think finding the probability of infants requiring ventilation based on birth weight, gestational age, and mother's age.[1]

Incorporation

Machine learning algorithms are hard, and understanding where and how to incorporate these models into products, let alone business applications, requires a certain level of understanding of these solutions and creativity in their incorporations. Knowing that artificial intelligence and machine learning are more than neural networks and are rather statistical models used to predict an outcome helps to remove the grandiose ideas and ground thoughts in the actual business applications companies work with every day. When thinking of where to start and how to incorporate these models into business applications, there are two important things to remember: **1) start small and then 2) work your way up.**

Starting Small

When starting small, building off the work of others is an important concept. This is no less true in the realm of machine learning. Many broad application machine learning products already exist and are robust. Google's suite of AI products tends to be a good jumping in point. These products offer API's that let an application pass in its own data and the API provides an answer on pre-trained models. These models include chat-bots, frequently-asked-question answering mechanisms, image classification, and text translation.

Using Appian itself as a case study, one can see how they found ways to incorporate these ideas into everyday business applications. A self-reported use case was incorporating AI into their ticketing triage system. Prior to the implementation of AI, an individual would make their best guess on which engineering group a ticket belonged with and if that was wrong, then that engineering group would make their best guess on what group it belonged to. This would repeat until finally the right engineering team was identified. Each individual group lacked insight into the entire platform and thus had a hard time identifying which group a ticket should be assigned to. However, training an ML model with the entire dataset helped to improve the rate at which tickets were successfully assigned. This solution likely used Google's text classification AI in order to group the text in tickets into specific categories (in this case teams) in which the ticket could be assigned.

Expanding AI Practice

Once small adoptions and incorporations of AI have taken place, one is more prepared to take larger leaps into the AI realm. When reviewing systems for areas that could benefit from ML, larger wins come from targeted implementations of ML that tie into the business mission. Data here tends to be the key. In order to create these types of business-specific ML solutions, training an ML model is the next step and in order to train these models, there needs to be sufficient data for the model to find patterns.

Let’s take an example of fraudulent claim detection. In the past, workers would review claims in order to find these fraudulent ones manually. From this, these companies would have millions of claims where each claim could be categorized as fraudulent or not. This data could be massaged into feature sets (for example, frequency of visits, location of visit, amount billed, etc) that could be used in a prediction model that would give the probability that a claim is fraudulent or not. While it might be hard for people to quantify the relationships between these chosen feature sets, ML models tend to be good at finding relationships between these variables or feature sets. So now, this workflow that used to take swaths of people reviewing claims to find fraudulent ones, can now have workers only review claims that have a predefined probability of being fraudulent.

Benefits

These examples also touch on the power and value of ML incorporated into business applications - that is time saved for workers and improved quality of the results. Both examples given demonstrate how ML models can help save time for the workforce. In Appian’s use case, engineers no longer had to go through trial and error of assigning bugs and were saved the time so they could work on the solutions to these bugs instead. For the fraudulent claim detection use case, we can see that the sheer number of claims needed to be reviewed by workers was drastically reduced to a more manageable size. The value of ML comes from the insights gained. Our ability to automate processes based on these insights is a major boon to business applications. Whether the insight is what category a certain piece of data belongs to, or the ability to flag a piece of data is irrelevant so that workers do not need to review it, finding a highly repetitive task which requires human insight, and further automating this task, will always lead to value in the long run.

Conclusion

Artificial intelligence at its core is nothing more than a predictive model. It answers questions like “what is the likely temperature tomorrow” or “what is the probability that the picture I’m looking at is a picture of a cat?” These simpler questions can be built upon each other to answer tougher and deeper questions, such as “what’s the sentiment of this message” or “what’s the best response to provide for a given question.”

The answers or insights into these questions allow the applications of today to evolve into the applications of tomorrow. Operations that once required human intervention can now have predictive models trained to provide answers. These machine learning models provide value through the insights derived and the speed at which computers can reach these insights. Operations viewed as human-centric tasks or requiring fleets of humans to solve, such as triaging tickets to the correct development group or asking a customer for identification and searching for their records, are now easily done with a single AI model. This allows workers to focus on more important tasks, whether that is providing better customer experiences, or giving better predictions or risk assessments to leaders so that they have more time to take action.

Rome wasn't built in a day and neither should an organization's AI suite. When looking for opportunities, think small and look for ways to leverage pre-built models. Once users and application designers have become familiar and comfortable with these small incorporations, then start to work towards loftier goals as the infrastructure and organizational knowledge of AI increases.

About the Author

Darby Kidwell is an Appian Enterprise Architect who has been working at Macedon for the past 8 years. He graduated from the University of Virginia with an engineering degree. He has a passion for machine learning algorithms and studies and develops them on the side.

[1] Logistic Regression Analysis; Julien H. Hoffman:
https://www.sciencedirect.com/topics/medicine-and-dentistry/logistic-regression-analysis

[2] Logistic Regression -- Detailed Overview; Saishruthi Swaminathan:
https://towardsdatascience.com/logistic-regression-detailed-overview-46c4da4303bc

[3] https://en.wikipedia.org/wiki/Linear_regression

[4] https://blog.bigml.com/2016/09/28/logistic-regression-versus-decision-trees/

[5] https://en.wikipedia.org/wiki/Decision_tree_learning

Contact: (571) 526-4281
info@macedontechnologies.com