deepinstinct
BEFORE YOU KNOW IT

DEEP LEARNING CHECKLIST FOR CYBER SECURITY

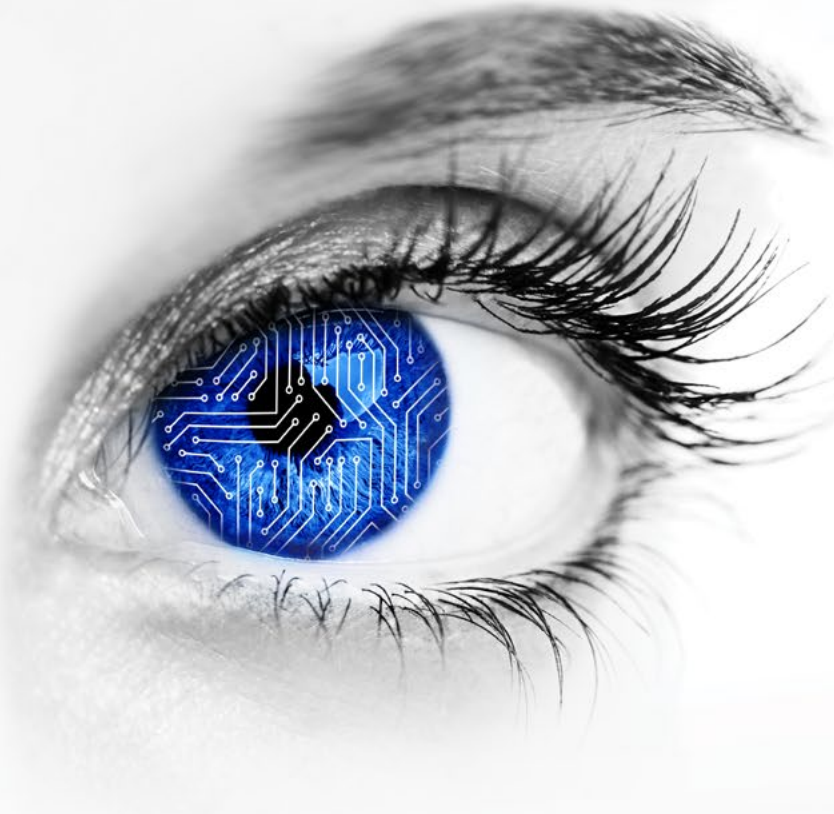

Artificial Intelligence, machine learning, and deep learning are all buzzwords dominating any conversation about computer processing capabilities and cause large scale confusion. Therefore, this checklist was prepared in the growing need to understand the key differentiators between deep learning and machine learning.

Deep learning, an advanced subset of machine learning, is an artificial intelligence method that imitates the way the human brain works in the sense of processing data and creating patterns for use in decision making.

Deep learning utilizes a hierarchical level of artificial neural networks to carry out the learning process involved in machine learning. The artificial neural networks are built like the human brain, with neuron nodes connected like an interconnected web.

The advantage of deep learning over other forms of machine learning is the end-to-end processing of data. By end-to-end we refer to three aspects;

1. The elimination of the feature engineering phase, an inherent part of machine learning, which involves a human expert identifying and selecting the features for analysis

2. The analysis of all the available data in the training sample

3. Encompassing of representation learning the ability of a model to get input low level features (such as characters in a text) and to transform these raw features to high level features (such as words and sentences) and predict based on these higher level features.

The combination of these three factors contributes to deep learning's greater level of determining accuracy.

When considering a cybersecurity product, it isn't always easy to clearly understand whether what you are looking at is based on deep learning technology or is based on artificial intelligence and machine learning.

Therefore, here at Deep Instinct™, the first company to apply deep learning to cybersecurity, we put together a "deep learning checklist for cybersecurity". It includes the key questions you should ask to better understand the AI capabilities of any given product and easily differentiate between machine learning and deep learning features.

Who are the deep learning researchers behind the implementation?

Developing state-of-the-art deep learning solutions require deep learning experts, which are in short supply. By enquiring about the researchers involved, and learning about their experience in deep learning (mostly academic experience, since deep learning spawned directly from academic research groups), it is possible to get a better understanding of the expertise behind the solution.

What deep learning framework is being used?

Developing a deep learning framework is an extremely complex task, which only a few companies have

successfully accomplished. Publicly available deep learning frameworks are sufficient to meet the needs for most computer vision applications, but are inefficient for other applications, such as cyber security. Very few companies are currently practicing deep learning in this domain, due to the number of significant challenges (e.g. scanning thousands of files per second).

Is a specific domain expert required for the training process?

Unlike traditional machine learning methods where applied features must be identified by an expert and then hand coded per domain and data type, deep learning is applied directly to raw data without any required domain knowledge.

On the contrary, in deep learning, the features are identified by the algorithm itself. The process in which the algorithms are learning higher level representation of features in deeper layers is called representation learning.

How much time does it take to perform the training?

Deep learning algorithms take significantly more time to train than those of traditional machine learning. Traditional machine learning algorithms typically take from a few seconds to a few hours to train, while deep learning algorithms take just several hours to train.

Can the AI algorithm used handle any file format?

Traditional machine learning algorithms require different human engineered features for every file type (PDF, DOC, EXE, etc.). In contrast end-to-end deep learning models, using raw features such as raw byte content are agnostic to the file type.

To what extent is the algorithm susceptible to adversarial attacks?

Traditional machine learning uses engineered features. Those can be easily modified by attackers to bypass the AI model, as has already been documented with commercial Next-Gen AVs. End-to-end deep learning models, using raw features such as raw byte content that are more robust and resilient to such modifications.

What input data is fed into the models?

Deep learning is typically applied directly to raw data. Any answer that involves “feature extraction” or “manual preprocessing” suggests that machine learning is used.