Simon Koeten
Justin Atwell
WRTG 3035
3 March 2020

Neural Networks: Current Applications Glancing Towards the Future

Abstract:

Neural networks are a blooming field within the wide range of artificial intelligence. However, the theory and math has been around for over twenty years, while mankind's computing power has been growing dramatically since the inception of neural networks. Now in an age when using neural networks are feasible, this review takes a look at current applications and barriers to these networks. From this, the reader can gain an understanding of what we can do currently with neural networks, but more importantly where they can take us moving into the future.

Keywords:
neural networks, artificial intelligence(AI), processing power, data

History of Neural Networks:

Before even jumping into neural networks, it is important to understand their scope in relation to the broader umbrella of AI. Within AI, there are three levels to a machine's "intelligence". Artificial intelligence: a machine's ability, often performing a single task, to learn and improve its ability in that one task [7]. Artificial general intelligence: at this level, a machine would be considered of similar intelligence to a human (no AI is at this level) [7]. Artificial super intelligence: this level is what is commonly depicted in movies (Ex Machina), where an AI has surpassed human cognitive and functional capabilities (again no AI has reached this level) [7]. Neural networks lie somewhere between artificial intelligence and general intelligence, but lie much closer to the first. Further, amongst many AI techniques, neural networks are one of the most prominent and budding methods.

Additionally, what is interesting about neural networks, is that the theory of how they work has been around since before the 2000's [2]. The reason they are just now coming into fruition is because neural networks are both computationally very dense and require massive amounts of data. Solutions to these barriers have not been available until recently, as by Moore's Law, in which Gordon E. Moore observed that processing power doubles every two years, where in the current (2020), we have the means to actually implement this form of AI. It is important to note that Moore's law was no longer applicable starting around 2010, due to a limit on how small we can make transistors. However, processing power is still increasing (albeit at a slower rate), and additionally with quantum computing on the horizon, processing power can only grow. In order to understand why neural networks are so computationally expensive, we have to look into how these networks actually work.

How Neural Networks work:

In many ways, neural networks mimic how our brains communicate information, via neurons and connections between said neurons, human-kind's brainchild. A major difference between the two, our brains in a single cortex can have “140 million neurons, with tens of billions of connections between them” [1], compared to a basic neural network, which probably contains 100 nodes or “neurons” with thousands of connections. The difference is immense, which is part of a barrier to neural networks, processing power, computers simply cannot compute that much information in a reasonable amount of time. Continuing, the nodes within a neural network are divided into levels, with each node in contact to each of the nodes of the previous level and the next level [2]. The final level of nodes will eventually produce a probability of how certain the network is of the answer [1,2].

While the above describes the structure of neural networks, there is a lot more to them than just feeding in an input and expecting it to produce a correct output. Let us use an example of a neural network being able to recognize hand-written letters. Say a letter was inputted into the network, it would traverse through the different levels, once it reaches the last level (26 nodes for each letter), the AI would return the letter that has the highest probability (of being correct) within the last level of nodes. The way neural networks “learn” is by feeding thousands to millions of different images with an answer attached to deem if the network got it correct [3]. In sending in these data points, the different levels will change their node's “weight” based on if it got the answer correct or not [1,2]. The network will fine tune the weights for each node during each iteration of a data point or image in the letter example. This is called “training”, which is extremely reliant on good, clean, and accurate data [1,2,3]. Words like “good” and “clean” are very subjective, but in relation to neural networks, if data is not clean the network will not be able to accurately read in the data so its already processing something objectively wrong. This can often be another huge barrier to neural networks, as having dirty data can hinder progress or ability of the network [1,3]. Neural networks use a structure similar to how our own brains work, and similar to how humans learn, the networks need a lot of data to learn, what is right and wrong.

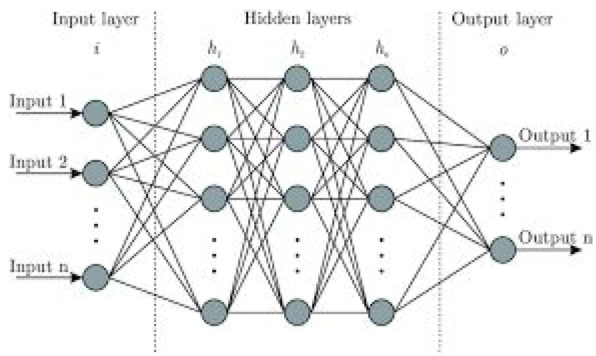

Current Applications:

Much like all methods of AI, there are two main types of neural networks: supervised versus unsupervised. With supervised neural networks, like in the letter example previously mentioned, there is a correct answer associated with the input, a “ground truth” [8]. Whereas unsupervised neural networks are focused on organizing similar inputs, in some sense creating its own trends or categorization of inputs, similar to clustering [8].

Supervised:

Many of the neural network examples people think about tend to be supervised networks. Many of the examples revolve around image processing [6], especially with many different levels of detail, these can be difficult problems to solve. The level of detail can go from recognizing digits/numbers, to locating faces in an image, to even determining gender using a neural network [6]. The examples are reliant on pre processed data where the network can determine its correctness, and then based on that improve the network. An example of associating images with a correct answer is Google’s CAPTCHA (distinguishes computers from humans), the pop up where it gives you an image divided into a 3x3 or 4x4 quadrant sets, asking you to select all photos with a traffic light or with a taxi. These methods are easy ways to associate an image with a correct answer. While these examples seem to be a bit rudimentary, there is an incredible amount of math, implementation, and training going along with these projects.

Another popular example of neural networks is Google’s DeepMind AlphaZero program introduced in 2017. The AlphaZero neural network system was applied to three board games: shogi, go, and chess. The system learned how to play by simply playing against itself [4]. Well, one might think that tends to be more unsupervised than supervised. The system actually uses winning or win percentage as a mark of “correctness”. In each different game, the AlphaZero system would play an unconventional and revolutionary style of play, learned by playing itself. Further, in each different game, Google’s program “outperformed” the respective hallmark programs after four hours of training in chess, and two hours of training in shogi [4]. Additionally, in chess, AlphaZero played against world champion Stockfish (world famous chess program), “winning 155 games and losing just six games out of 1,000” [4]. In shogi, AlphaZero beat Elmo (shogi’s comparative Stockfish) “91.2%” [4] of the time. Finally, in go, AlphaZero beat AlphaGo Zero, its own predecessor, “61%” [4] of the time. In each game, AlphaZero was worlds ahead of the previously known computer program champion, a testament to the power of neural networks.

Unsupervised:

Would it be a proper literature review about a subject in the tech world without mentioning Elon Musk? Probably, however the work being done by Musk and one of his newer companies Neuralink, could really change society as we know it. Neuralink is a company devoted to the integration of AI and, literally, the brain. The device will be integrated into the brain via threads to not only read the activity of neurons, but ultimately have the capabilities to write and send messages to neurons [5]. While the initial goal is to help in recovery from brain injuries such as epilepsy and paralysis, Musk has his eyes on “a sort of symbiosis with artificial intelligence” [5]. Musk eventually desires an intertwined future where humans can access an integrated AI already inside their head. The company is trying to get approval from the US FDA, but to integrate a device with a brain is no easy task. Once integrated, Neuralink will use unsupervised neural networks to recognize and categorize patterns in neural activity to improve functionality [5]. A neural network working in tandem with a vastly more complex neural network. This technology could be another societal breakthrough. Musk argues that we already sort of have access to this technology in the forms of our phones and laptops. The barrier there is the interface ability, which Musk believes Neuralink could provide. This could change civilization as we know it, but until then humans will have to rely on their own neural network: their brain.

Future of Neural Networks:

Looking towards the future, there is a lot to be excited about the impacts neural networks can have. Although, of course with any great technology, there are drawbacks, and for neural networks it boils down to processing power and obtaining clean data sets to train the networks[1,2,3]. As time moves forward, processing power naturally increases (Moore’s Law), however eventually there will be a limit to how small we can make transistors which leads to a limit on the number of transistors to fit on a microchip. It is also important to note that once we reach a limit with transistors, there are other means of computations, an example already mentioned, being quantum computing.

With respect to producing or receiving data for training, the data is out there (reported that roughly 90% of data has been created in the last two years [9]), companies and projects need to continue gathering relevant data. Because data is so crucial to the growth and learning of these networks, it is vital to have complete data, so incomplete data points have to be thrown out. Continuing, companies have tremendous amounts of data on almost all individuals, including you the reader. This is something to be wary about moving forwards, as companies like Cambridge Analicitca have sold user data (Cambridge sold data to Facebook) in order for other groups to manipulate or persuade users. Pair that data with a powerful tool such as neural networks, the effects can be amplified. So with personal data often being used in any and every industry, it will be important to use neural networks to serve others, rather than influence and exploit humans across the world.

With both of these barriers, their walls will not tople all at once, but with time, neural networks will continue to grow in terms of use and importance. Knowing how powerful this tool is, and how it will only grow in efficiency and effectiveness, it is vital to remember it is simply a tool. Companies should have an eye out to protect and serve users and their data. So, just like how the math for these networks was known pre-2000 and we are just scraping the surface of neural network uses in 2020, the future of this technology has the sky as the ceiling.

Conclusion:

Projects like Elon Musk' Neuralink to devices like an Amazon Alexa are using neural networks to innovate the bounds of human capabilities. The networks are being trained using the massive amounts of data, something only huge companies like Google, Amazon, Apple, and Facebook have, to make their technologies smarter. As we continue to the future, neural networks will become increasingly impactful being used as a vessel for anonymus machine learning, hopefully to learn the best way to assist human kind.

References:

[1] Nielsen, Michael A. “Neural Networks and Deep Learning.” *Neural Networks and Deep Learning*, Determination Press, Dec. 2019, neuralnetworksanddeeplearning.com/chap1.html.

[2] Duda, Richard O., et al. *Pattern Classification*. 2nd ed., Wiley-Interscience, 2001.

[3] Dreiseitl, Stephan, and Lucila Ohno-Machado. “Logistic Regression and Artificial Neural Network Classification Models: a Methodology Review.” *Journal of Biomedical Informatics*, Academic Press, 7 June 2003, www.sciencedirect.com/science/article/pii/S1532046403000340.

[4] Silver, David, et al. “AlphaZero: Shedding New Light on the Grand Games of Chess, Shogi and Go.” *Deepmind*, DeepMind, 6 Dec. 2018, deepmind.com/blog/article/alphazero-shedding-new-light-grand-games-chess-shogi-and-go.

[5] “Neuralink.” *Neuralink*, 2019, www.neuralink.com/.

[6] Murray, Alan F. *Applications of Neural Networks*. Springer, 2011.

[7] Ulbert, Sebastian. “The Difference Between Artificial Intelligence, General Intelligence, And Super Intelligence.” *Field Service Management Software*, Apr. 2017, www.coresystems.net/blog/the-difference-between-artifical-intelligence-artifical-general-intelligence-and-artifical-super-intelligence.

[8] Soni, Devin. “Supervised vs. Unsupervised Learning.” *Medium*, Towards Data Science, 16 July 2019, towardsdatascience.com/supervised-vs-unsupervised-learning-14f68e32ea8d.

[9] Marr, Bernard. “How Much Data Do We Create Every Day? The Mind-Blowing Stats Everyone Should Read.” *Forbes*, Forbes Magazine, 5 Sept. 2019, www.forbes.com/sites/bernardmarr/2018/05/21/how-much-data-do-we-create-every-day-the-mind-blowing-stats-everyone-should-read/#3af1aa7360ba.