

How You Should Read Research Papers According To Andrew Ng (Stanford Deep Learning Lectures)

medium.com/m/global-identity

Richmond Alake

August 30, 2021

towards
data science

Informational

Instructions on how to approach knowledge acquisition through published research papers by a recognized figure within the world of machine learning and education

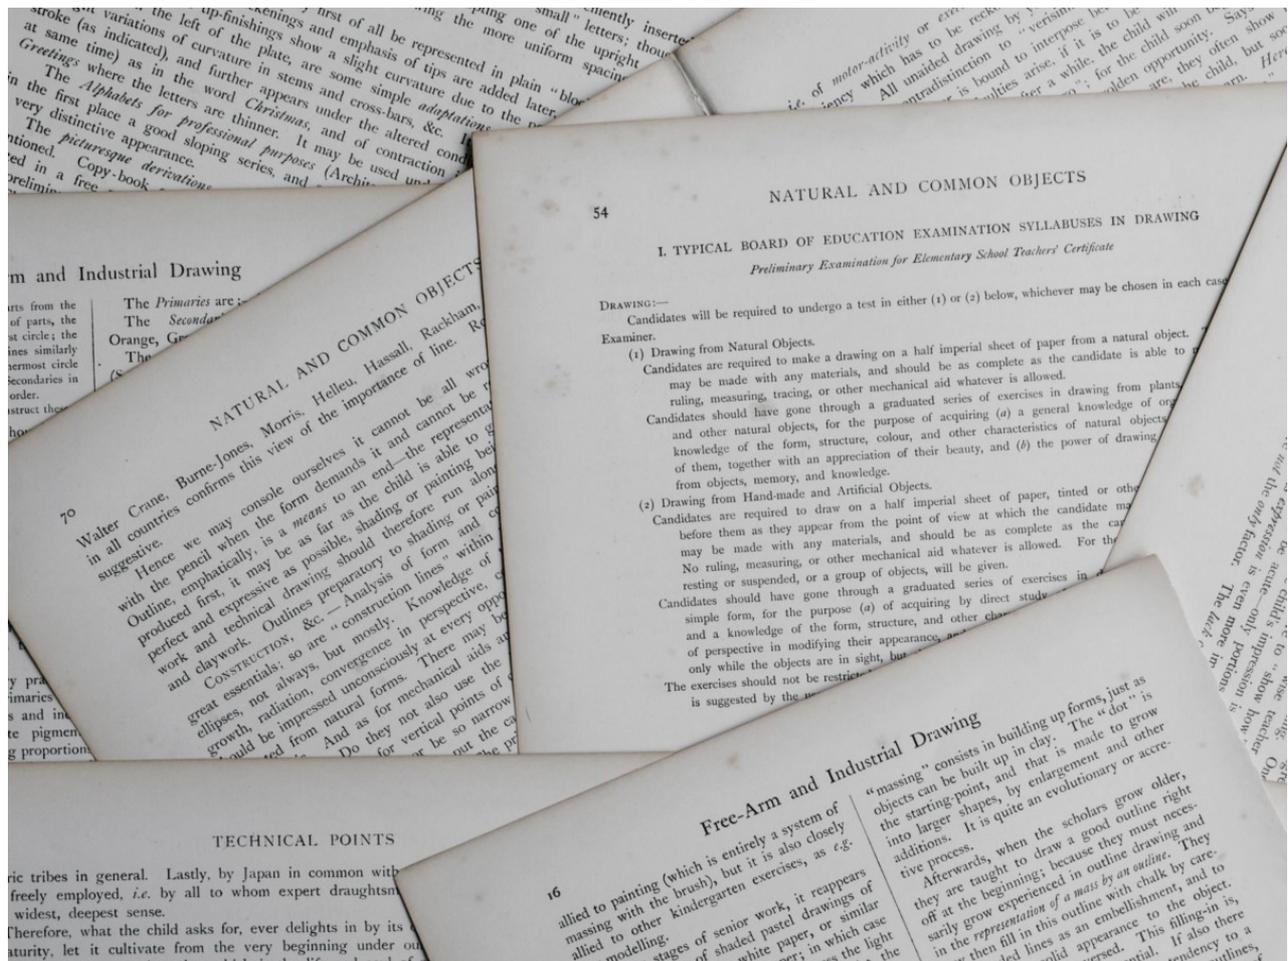


Photo by on

“Wisdom is not a product of schooling but of the lifelong attempt to acquire it.”

— Albert Einstein

Introduction

The ability to understand information produced by the individuals at the cutting edge of research within Artificial Intelligence and the Machine learning domain is a skill that every serious machine learning practitioner should acquire.

To stay relevant and increase your knowledge, machine learning practitioners need to have an academic mindset and habit. AI, ML and DL are evolving at a fast pace, and we have to equip ourselves with the knowledge to keep up with the field, knowledge that is only attainable through research papers.

This article will provide you with instructions on how to go through a research paper effectively, and also provide the following:

- A systematic approach to reading a collection of papers to gain knowledge within a domain
- How to properly read a research paper
- Useful online resources that can aid you in searching for papers and key information

For those who would like to get to the key content within this article, scroll down to the section titled “**Reading Research Papers**”.

First, Who is Andrew Ng?

The information I provide in this article was derived from a [Stanford lecture taught by Andrew Ng](#). I've also supplemented the information contained in this article with personal tips and information from resources on the internet.

But first, a brief introduction on [Andrew Ng](#).

[Andrew Ng](#) is probably the most known (*and watched*) machine learning teacher on the internet. He is also the co-founder of both [DeepLearning.ai](#) and [Coursera](#).

Apart from his ongoing work within online education, he's also a professor at [Stanford University](#).

More information on Andrew Ng is just a [google search](#) away.

It's natural for a person to pick up skills and habits demonstrated by individuals around their environment; this is why most PhD students will acquire the skill of effectively digesting the content of a research paper appropriately. This is somewhat a fact, and Andrew mentions it very early in the video referenced earlier.

But we are not PhD students, well some might be, but how do we normal individuals gain the required skills to read a research paper and understand its content wholeheartedly.

Reading Research Papers



Photo by [Christa Dodoo](#) on [Unsplash](#)

Specialization within the machine learning domain is favourable if you are talented. For example, having a generalist knowledge on the field of Computer vision is commendable, but having specialized knowledge and expertise within key techniques such as Pose estimation will be more appealing to companies and organization looking for practitioners within that domain.

So let's use Pose Estimation as a guide to how we would approach reading research papers related to the subject matter: pose estimation.

1. Assemble collections of resources that focus on the subject matter. Resources can come in the form of research papers, Medium articles, blog posts, videos, GitHub repository etc.

A quick [google search](#) on the phrase “pose estimation” will provide you with top resources that contain information in regards to the subject matter. At this initial step, the aim is to collate all resources that are relevant, such as YouTube videos, implementation documentations and of course research papers. Ideally, at this stage, there is no limit to the number of resources you consider important, but be sure to create a shortlist of papers, videos and articles that are useful.

2. In this next step, you will conduct a deep dive of any resource you deem relevant to the subject matter. It is crucial that there's a method to track the understanding of each shortlisted resources. Andrew Ng, suggests a table of resource plotted against your understanding level that looks similar to the table below.

Pose Estimation Research

Understanding level

Resource	<input checked="" type="checkbox"/> 10% - 20%	<input checked="" type="checkbox"/> 20% - 40%	<input checked="" type="checkbox"/> 40% - 60%	<input checked="" type="checkbox"/> 60% - 80%	<input checked="" type="checkbox"/> 80% - 100%
https://www.tensorflow.org/lite/models/pose_estimation/overview	<input type="checkbox"/>				
https://medium.com/tensorflow/track-human-poses-in-real-time-on-android-with-tensorflow-lite-e66d0f3e6f9e	<input type="checkbox"/>				
https://arxiv.org/pdf/1602.00134.pdf	<input type="checkbox"/>				
https://arxiv.org/pdf/1603.06937.pdf	<input type="checkbox"/>				
https://arxiv.org/pdf/1505.07427.pdf	<input type="checkbox"/>				
https://arxiv.org/pdf/1703.06870v3.pdf	<input type="checkbox"/>				
https://arxiv.org/pdf/1812.03595v3.pdf	<input type="checkbox"/>				

Understanding level table of resources By Richmond Alake

It is advisable to ensure you go through at least 10–20% of the content of each paper you have added to the list; this will ensure that you have been exposed to enough of the introductory content within an identified resource and are able to gauge its relevancy accurately.

For the more relevant papers/resources identified, it is expected that you progress to a higher level of understanding. Eventually, you will have identified some appropriate resources with content that you understand fully.

You are probably asking yourself, “what number of papers/resource is sufficient”.

Well, I don't have the answer, but Andrew does.

According to Andrew, an understanding of 5–20 papers will showcase a basic understanding within the subject matter, perhaps enough understanding to progress to implementation of techniques.

50–100 papers will primarily provide you with a very good understanding of the domain.

After going through the resources and extraction of vital information, your table might look something similar to what's shown below.

Pose Estimation Research

Understanding level

Resource	<input checked="" type="checkbox"/> 10% - 20%	<input checked="" type="checkbox"/> 20% - 40%	<input checked="" type="checkbox"/> 40% - 60%	<input checked="" type="checkbox"/> 60% - 80%	<input checked="" type="checkbox"/> 80% - 100%
https://www.tensorflow.org/lite/models/pose_estimation/overview	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
https://medium.com/tensorflow/track-human-poses-in-real-time-on-android-with-tensorflow-lite-e66d0f3e6f9e	<input checked="" type="checkbox"/>				
https://arxiv.org/pdf/1602.00134.pdf	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
https://arxiv.org/pdf/1603.06937.pdf	<input checked="" type="checkbox"/>				
https://arxiv.org/pdf/1505.07427.pdf	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
https://arxiv.org/pdf/1703.06870v3.pdf	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
https://arxiv.org/pdf/1812.03595v3.pdf	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Updated Understanding level table of resources By Richmond Alake

3. The third step is a quick tip from what I've observed works for me when trying to understand research papers. The third step is to take structured notes that summarises the key discoveries, findings and techniques within a paper, in your own words.

| The following steps will now be focused on how to read a single research paper.

Reading A Single Research Paper

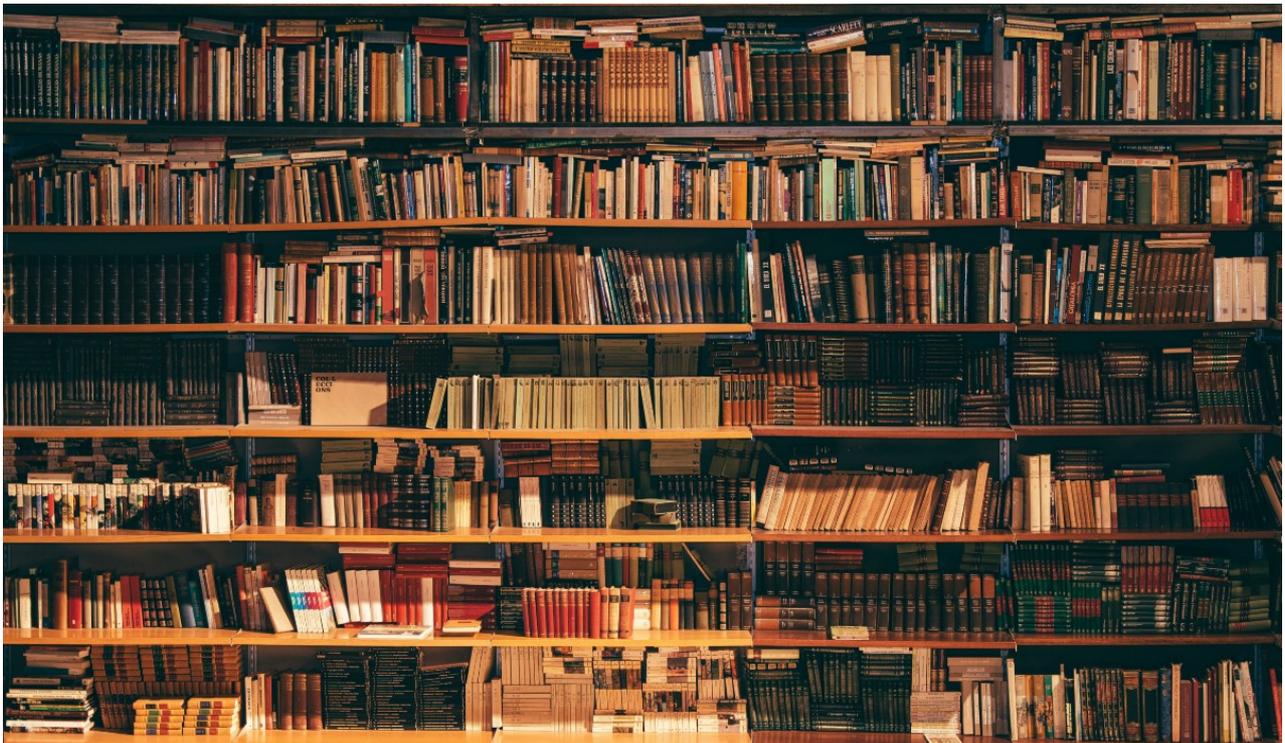


Photo by [Alfons Morales](#) on [Unsplash](#)

Reading for the purpose of understanding is not done through one pass of the contents within the paper. According to Andrew, reading a paper from the first word to the last word in one sitting might not be the best way to form an understanding.

Be prepared to go through a paper at least three times to have a good understanding of its content

4. In your first pass, start with reading the following sections within the paper: title, abstract and figures.

5. The second pass entails you reading the following sections: introduction, conclusion, another pass through figures and scan through the rest of the content.

The introduction and conclusion section of a paper contains clear and concise information on the content of the paper and a summary of any findings. The information presented in this section usually dismisses any supplementary information and only key information are included. This is beneficial to you as a reader as you get the vital information required to proceed to the other sections within the paper.

6. The third pass of the paper involves reading the whole sections within the paper but skipping any complicated maths or technique formulations that might be alien to you. During this pass, you can also skip any terms and terminologies that you do not understand or aren't familiar.

7. Those conducting in-depth research into a domain can take a few more passes. These additional passes will mainly be focused on an understanding of the maths, techniques and unknown terminologies presented within the paper.

For those who are generally reading research papers for informational and engineering purposes, then in-depth research might be very time consuming, especially if you have 20 more papers to get through.

I went through the process presented in this article with the [original paper](#) introducing the LeNet convolutional neural network, and I summarised the key content in notes which I later then converted to a series of Medium articles.

[Understanding and Implementing LeNet-5 CNN Architecture \(Deep Learning\)](#)

[In this article, we perform image classification on the MNIST dataset with custom implemented LeNet-5 neural network...](#)

towardsdatascience.com

(You Should) Understand Sub-Sampling Layers Within Deep Learning

Average pooling, max-pooling, sub-sampling, downsampling, are all phrases that you'll come across within Deep Learning...

towardsdatascience.com

Understanding Parameter Sharing (or weights replication) Within Convolutional Neural Networks

Parameter sharing or weights replication is a topic area that can be easily overlooked within deep learning studies...

towardsdatascience.com

Understand Local Receptive Fields In Convolutional Neural Networks

Ever wondered why all the neurons in a convolutional neural network aren't connected?

towardsdatascience.com

Questions To Ask Yourself

Andrew provides a set of questions that you should ask yourself as you read a paper. These questions generally will show you understand the critical information presented in a paper. I use the questions below as beacons to ensure I don't stray from the aim of understanding vital information.

They are as follow:

1. Describe what the authors of the paper aim to accomplish, or perhaps did achieve.
2. If a new approach/technique/method was introduced in a paper, what are the key elements of the newly proposed approach?
3. What content within the paper is useful to you?
4. What other references do you want to follow?

Additional Resources To Assist Research

Several online resources have made the discovery and retrieval of relevant information relatively easy. Below are examples of resources that will assist you in your search for pertinent information.

- [The Machine Learning Subreddit](#)
- [The Deep Learning Subreddit](#)
- [PapersWithCode](#)
- Top conferences such as [NIPS](#), [ICML](#), [ICLR](#)
- [Research Gate](#)

Conclusion

“Learn steadily rather than short burst for longevity.”

— Andrew Ng

I’m still relatively new to the field of Machine Learning and Computer Vision, there is a lot that I do not know (*that’s an understatement*). Still, I believe that if an individual is consistent in their search for knowledge, regardless of the domain, they’ll be rewarded with an understanding and expertise that surpasses the norm.

From the techniques introduced by Andrew Ng, I’ll be reading at least four research papers a month, reading to the point of understanding. I’ll be honest and say that the LeNet paper took me about a week and a half to complete wholeheartedly. But you get better and faster at reading and understanding research papers the more times you do it.

Andrew states in his video that he carries a batch of research papers around with him, intending to read them. Andrew is a prominent figure within the field of machine learning, and I believe emulating his habits and learning techniques can be advantageous to your learning journey.

I hope you found the article useful.

To connect with me or find more content similar to this article, do the following:

1. **[Support my writing](#)** by becoming a referred Medium member
2. Subscribe to my **[Email List](#)** for weekly newsletters
3. Connect and reach me on **[LinkedIn](#)**
4. Learn with me at **[O’Reilly](#)**.