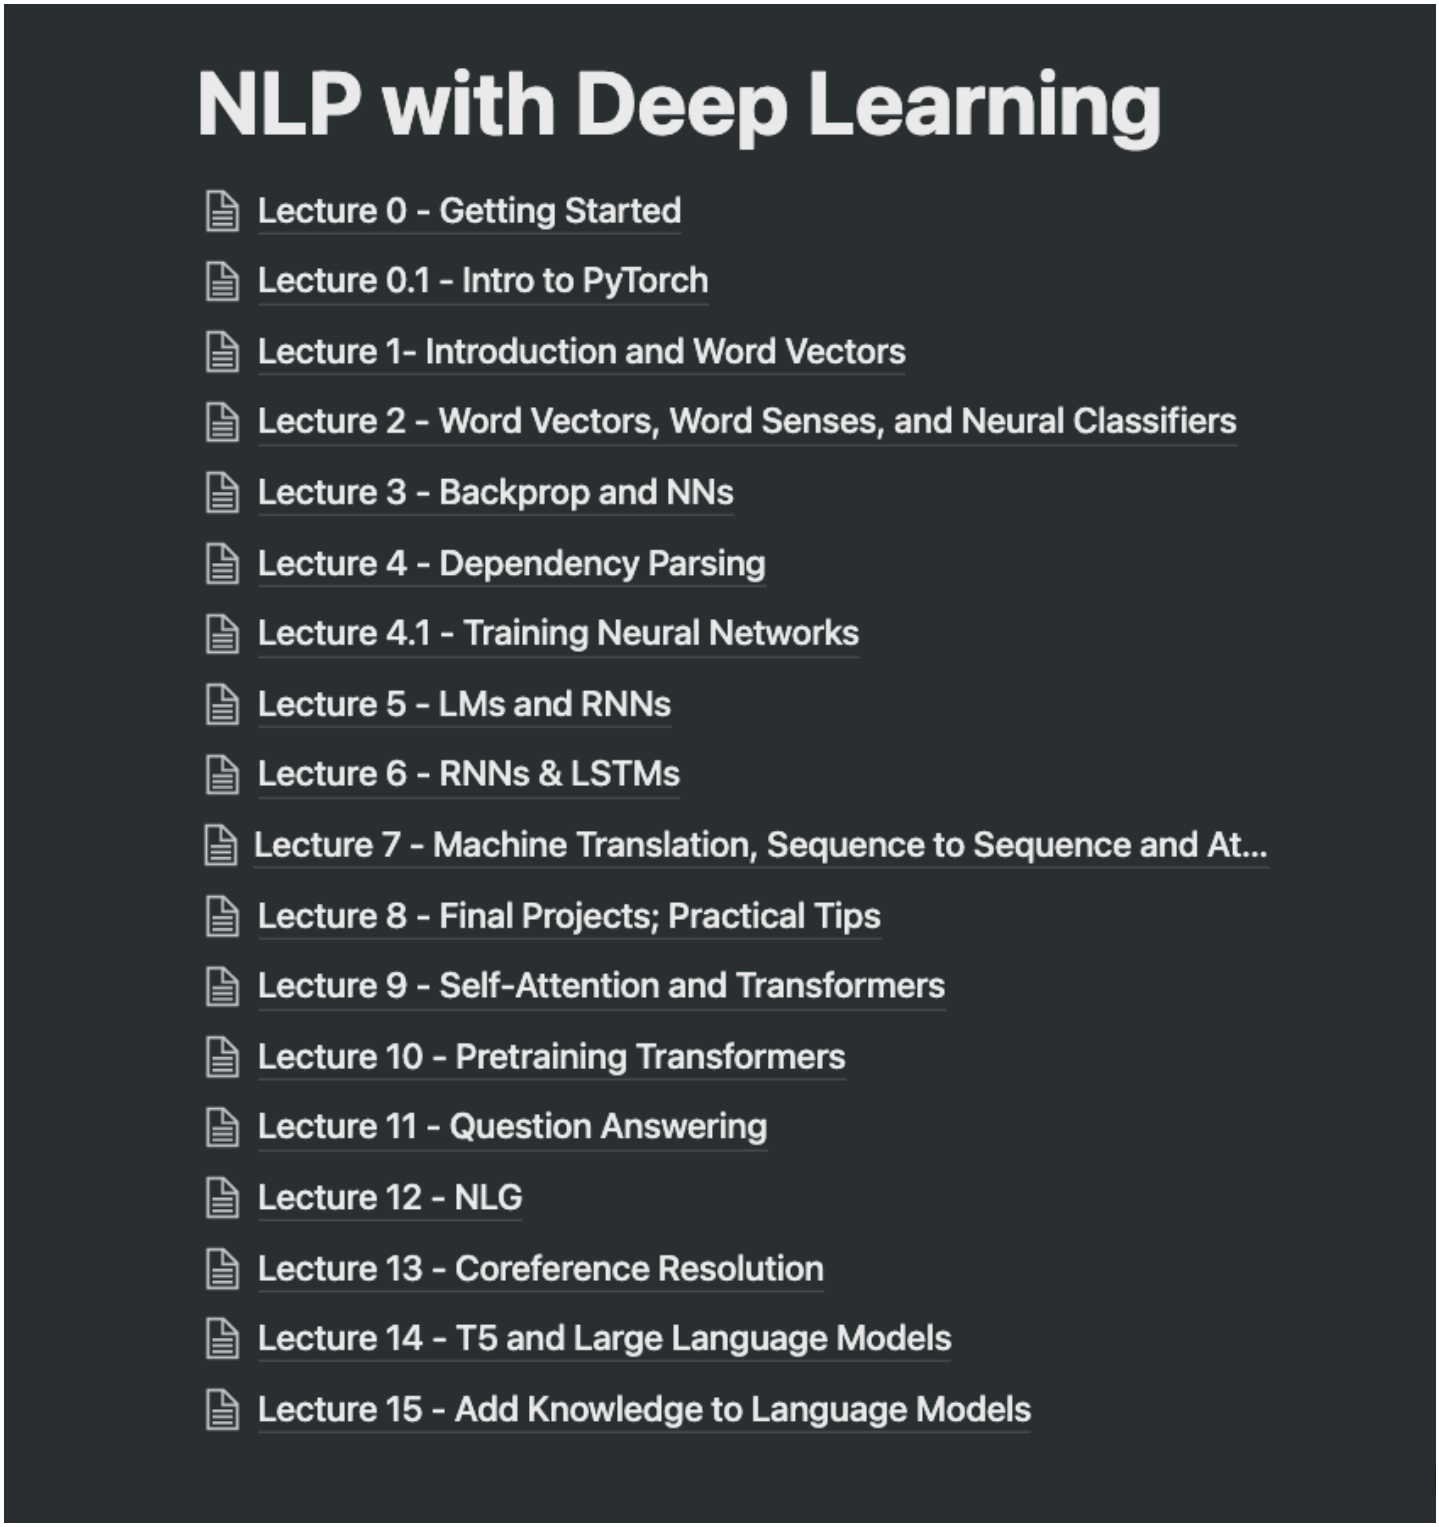
Twitter Thread by elvis

elvis
@omarsar0

The past month I've been writing detailed notes for the first 15 lectures of Stanford's NLP with Deep Learning. Notes contain code, equations, practical tips, references, etc.

As I tidy the notes, I need to figure out how to best publish them. Here are the topics covered so far:



I know there are a lot of you interested in these from what I gathered 1 month ago. I want to make sure they are high quality before publishing, so I will spend some time working on that. Stay tuned!

https://t.co/8LnkAZThlf

> I've been writing notes for the latest Deep Learning for NLP course by Stanford.
>
> For fun, I also started to add my own code snippets into the notes. I think this is a more efficient way to study: theory + code.
>
> Plan to share these notes soon. Stay tuned! pic.twitter.com/hWzZDORbl6
>
> — elvis (@omarsar0) January 14, 2022

Below is the course I've been auditing. My advice is you take it slow, there are some advanced concepts in the lectures. It took me 1 month (~3 hrs a day) to take rough notes for the first 15 lectures. Note that this is one semester of content.

https://t.co/d9yoNz7Bbu

I'm super excited about this project because my plan is to make the content more accessible so that a beginner can consume it more easily. It's tiring but I will keep at it because I know many of you will enjoy and find them useful. More announcements coming soon!

NLP is evolving so fast, so one idea with these notes is to create a live document that could be easily maintained by the community. Something like what we did before with NLP Overview: https://t.co/Y8Z1Svjn24

Let me know if you have any thoughts on this?