Machine Learning,ML

Prof. Carl Gustaf Jansson
Prof. Henrik Boström
Prof. Fredrik Kilander
Department of Computer Science and Engineering
KTH Royal Institute of Technology, Sweden

Lecture 4
Applications of Machine Learning

So, this is the fourth lecture of the first week of the Machine Learning course. The topic is Applications of Machine Learning. This topic is, say, more or less independent of many other topics of the course. It's placed early in the course in this position because, for many techniques we are going to talk about, we will refer to applications, and therefore you I think it to be a good idea to introduce an overview of applications for you to have something to refer to.

So, when you open a newspaper today or local television, it's not unlikely that you can see ads or other programs or articles which announce success stories of Artificial Intelligence applications. So the question is how can you make sense of this flood of new inventions that claim that all these products are becoming better because they applied Artificial Intelligence? I mean I saw recently announcements of new kinds of AI focused TV brands. I mean all the major products within the smartphone area, by companies like Huawei, Samsung, Qualcomm, they launched AI-powered smart phones. Even Burger King boasts that their ads are AI-written. Some of these announcements are pretty serious actually, but some of them are just crap. But one could say at least, with some certainty, is that a majority of the real stories - the real artificial intelligence success stories that you see in the media - relate to the application of Machine Learning, more or less only. So even if the ads claim that it’s an application of Artificial Intelligence, it's not actually any application of Artificial Intelligence in the broad sense. Plus we have described earlier this week that with the pretty narrow application of Machine Learning techniques it still might be very useful. But in order to understand what we actually are doing I think this is a very important fact. The second pretty certain fact is that out of the many Machine Learning success stories, a majority of these success stories that are announced today relate to Image Recognition or Speech Recognition. Maybe there are a lot of different kinds of applications which have been improved, but the source of improvement is that the Image Recognition in these systems have become better, or the Speech Recognition. So this is just to give you some sanity check of what is probably going on. But when you see these ads, some systems announced may be very broad - may be very broadly using Machine Learning, very broadly Artificial Intelligence, but it's a minority of the systems.

So in order to give you some aid in analyzing what is really going on with respect to applications, what I've tried to do on the current slide is to separate what I call general application sectors from what I call specific categories of Data Analysis. The general application sectors are, I would say, trivial or straightforward. They are well recognized - medical diagnosis, personalized treatments, drug design, driverless vehicles and household robots, personal assistants and recommender systems and navigators, adaptation of communication and social media services, marketing and sales, optimization of technical processes, monitoring and surveillance, financial services, cyber security, and machine translation. All these are rather common-sense application sectors. And when you look at the announcements of progress, of course always these relate to some of these. What is a little more innovative on this slide is the introduction of what is called specific categories of Data Analysis, because these categories represent some kind of middle level of analysis - which means that you not only say that you make this really great medical diagnosis system using Machine Learning, but you really instead you say "Ok, there is this medical diagnosis system using Image Recognition, or using data mining for large data set, or using text mining", and then indirectly you say that "Oh in the emignator image recognition system you use machine learning", because that's what's actually the truth. So what I claim here is that it has more explanatory power to introduce these two levels and explain the relation between Machine Learning and the real application sectors in two steps. And the five categories I here introduce in the Data Analysis - Image Recognition, which I see as the foremost sector, the category for the moment given everything that's happened at that is the point, Speech Recognition is also very important, Data Mining for large datasets is the third, Text Mining of large document collections is the fourth, and Dynamic Adaption on technical systems is the fifth. In the coming slides I will try to say a few words about each of these categories.

So starting with Image Recognition or Computer Vision, which is another word for it. Image Recognition or Computer Vision deals with how computers can be made to gain high-level understandings from digital images or videos. So the area seeks to automate tasks that the human visual system is assumed to solve. So, we could say we want to mimic here, somehow, the human visual system, not entirely but partly. So Image Recognition in my view is currently the most successful application domain for Machine Learning techniques. Many of the real success stories we hear about is coming from this area. Recent advances in medical diagnosis, for example for different kinds of cancers, and the key components in self-driving vehicles, are primarily due to the progress in Image Recognition. So if we look at some sub-domains, we can see what is important is to detect objects, it is important to detect events in images. We also have the issue not only about static objects, but also about moving objects - so actually tracking videos and doing motion analysis. Another part is not just looking at objects - static or moving - but also looking at whole scenes, including many objects, moving or not. A fourth issue here which is a growing area is also how you can work with images - how you can restore images if you have a partially damaged image or a bad image - how can you restore it to a state where it's useful for different purposes. And of course there are many phases to this kind of system, so of course you

have to have a good system for acquiring images in the first place for pre-processing them. A very important part is of course Feature Extraction, because an object on an image may have thousands and millions of features. And of course that'll make sense to some analysis or learning based on all these possible features that you get from extraction or selection. So then comes the phase of detection. Maybe you'll have to divide up the image into different parts, into different segments. In order to simplify the work comes high-level processing. And finally the stage of decision making based on what image you have analyzed and found.

So, Speech Recognition enables the recognition and translation of spoken language into computer readable text. In speech recognition, both Acoustic Modeling and Language Modeling or Language Engineering are important parts in the modern statistically-based Speech Recognition systems. Actually, Speech Recognition in my opinion is currently one of the most, or I will say the second most (after Image Recognition) successful application domain for Machine Learning techniques today. There are a lot of advances reported on Personal Assistants today. Also Machine Translation systems get much more usable. And also for various communication systems there are speech interfaces that become more and more used and accepted today. And I would say that in most of these cases it's not so much the underlying reasoning about the domains, it's not so much about the true language understanding, but it's rather the Speech Recognition that has been improved that makes these systems more accepted and useful today. There are two kinds of recognition systems - on one hand you have Speaker Independent system that could be used by any person that happened to use them, or you have Speaker Dependent systems where you actually train the systems based on a specific person's voice – so when your tailor the system and fine-tune the recognition for that specific person. And also related to that, of course, is the technology that can be used for identifying speakers.

So the third category is Data Mining, or synonymously Data Analytics for Large Data Sets. Data Mining is a process that is used to extract actionable data patterns from large data sets. "Actionable" is a pretty important word here because you don't do these analyses for fun, you do it with the purpose and many times, or most of the times, the motivation for doing data mining is a business oriented motivation and to support a wide range of business decisions that could be ranging from operational to strategic. Data Mining is primarily applied in weak theory domains to discover new patterns, and it is also related to other areas like Data Discovery, Knowledge Discovery, Business Intelligence and Data Warehouses. These terms or concepts are more or less synonyms. So the data set analyzed in Data Mining may be user input to e-commerce system - that's a very important source today. Or such data could be also user input to personal assistants, to recommender systems, to main systems, or social media. It could be data collected from industrial processes, from production processes. And it could be data collected from online learning systems, where there is a specific term for this - Learning Analytics. But it could also be large factual, financial and statistical databases. For speech and for image recognition that we discussed earlier, at the moment the artificial neural network techniques are preferred as a technique and have had many successes. But for Data Mining, still a variety of different Machine

Learning techniques can apply and there is no such clear bias on which technique is the most successful of them.

So the fourth category is Text Mining or Text Analytics for Large Document Collections. We still talk about large sets of data, but in the last category when we talked about Data Mining, we mostly referred to structured data. Here we have just normal texts. So the goal for Text Mining or Text Analytics is the process of exploring and analyzing large amounts of unstructured text data aided by software that can identify concepts, patterns, topics, keywords, and other attributes in this text data. So, Text Mining needs a combination of Natural Language Engineering techniques and Data Science techniques or Machine Learning. Typical document collections targeted by Text Mining are - business related documents, scientific publications, governmental reports, and News streams. Multiple language issues is a key thing here because, of course when you analyze the topic you cannot just say we will look at articles in English, or in Chinese, or Japanese. You have to handle multiple languages, so therefore multiple translation issues are crucial in text mining. Many applications need incremental Text Mining for new documents to augment knowledge-based systems or expert systems. So, for example, you have an expert system in medicine, or in law. For example, those systems gets old after a while, there comes new information all the time in document form and in order to keep the systems up-to-date, you actually need to extract the relevant new facts from the new document and augment the systems based on those.

So the last category of data analysis here I call Adaptive Control of Technical Systems. I mean it is actually a different species of technology we are talking about here, because Adaptive Control is the control method used by a controller which must adapt to a controlled system. The parameter of the systems are typically uncertain in the beginning of the process. Of course the purpose is, during the process, to set values on the parameters to acquire an optimal solution. Analytic Adaptive Control methods are of course focused on truly optimal solutions and algorithms for the exact computation of those solutions. What machine learning can contribute with here are methods that can function and create good solutions even if optimality cannot be guaranteed, and the most relevant Machine Learning technique in this context is Reinforcement Learning, which then actually can function in an absence of a mathematical model in this case. One example of this is robot learning which studies how, for example, a robot can acquire new skills or adapt to an environment during learning. So Reinforcement Learning really looks into how software and hardware agents take actions in an environment so as to maximize its capability. Of course it's needed that the results of earlier actions are systematically logged or documented so that these outcomes can then be fed back into the system to allocate some kind of credit or blame to earlier actions, and modifying the rules of the system and thereby optimizing the future performance.

So, having presented these five categories of Data Analysis that are all dependent on Machine Learning or Data Science statistics, I now want to make one example to advocate that this way of

analyzing things have an explanatory value. So, just look at medicine and applications in medicine, and what I do in this slide is I try to exemplify that in medicine you can have applications of all these five categories that make sense. So you can have one application in medicine where you actually use image analysis for diagnosis of breast cancer. That's what you do and you can have a great success in doing this. The second example in medicine is a big issue - how to efficiently file medical records in clinical work, and to have a really good speech recognition system that helps doctors and nurses to file medical information in the records in an efficient way. It is a good help, and that is another kind of application. The third application is to use Data Mining to explore all the large clinical databases that exist to find patterns. But still, clinical databases are structured data, it's not images. So it's another form of analysis needed. And the fourth part here is mining text because in the medical area there are a huge number of new articles published all the time, and it's a big pressure on professionals in medicine to follow the advances through the publications. So a system that helps the professionals to explore all new information in these publications by doing text mining of these is also a very useful system. And of course these kinds of findings could also be used to update medical expert systems as an indirect effect. The fifth area, as you know today there are many experiments with surgical robots and of course surgical robots also have to be trained to do the optimal movements or perfect movements. So training of robot movements for surgical robots is clearly a case of the fifth category. So by this slide I want to illustrate that it has some explanatory value not to just say “Oh, we have this fantastic application in medicine and it uses Machine Learning”. It makes more sense to first analyze what we are doing in one of these five categories and then in the second stage say “Oh, the image analysis is done by using Machine Learning”.

So this was the end of this lecture. I thank you for your attention, and the next lecture will be on the Tutorial regarding the assignments for this week.