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 1
Introduction to the Machine Learning Course

Dear Students,
Welcome to this NPTEL course on machine learning, I'm Carl Gustaf Jansson, professor at KTH. I'm responsible for developing and running this course. This is the first lecture of the first week of this course and the purpose of this lecture is to relate machine learning to some other important relevant areas, but also to give you an overview of the content of the course. The purpose of this course is to give a broad introduction to machine learning, as being a subfield of artificial intelligence and computer science. The course is not aimed to be a professional course in the sense of giving hands-on knowledge on specific algorithms, coding in specific languages and apply to specific problems. You can learn many such alternative courses. Machine learning has imported much relevant knowledge from statistics and probability theory. This is not a general course on these fields either. This course will hopefully make you better computer scientists, who are able to handle data analysis problems but it will not make you a general statistician or data scientist. Even if the focus of this course is machine learning in the context of artificial intelligence I would like to say a few words about some neighboring areas like statistics, data science and big data. And first of all on the slide you see now you can get a little overview of the situation of the coming slides, I will say a few words about the different areas. So essentially today case against something happening and actually some involves in machine learning get a better spreading and also a new area emerged which is often called big data. Actually based on the growth, rapid growth of data in the world with of course of a great need to get this and a lot data analysed and I would say that these two things together triggered some urge to widen the concept of statistics somewhat within you were a traveller concept and term data science. So let's first turn to the area of statistics the oldest area in is stronger which has a very long history. As statistical concessional doubt with a range of activities concerning data from data collection, data harvesting, data modelling up to data presentation for decision-makers and in the middle there are data analysis and of course the big overlap between machine learning, statistics is in this middle, data analysis part. And of course in particular what is very useful and needed in machine learning is input from the traditional area of statistics with respect to mathematical statistics which is the core of data analysis which in turn is based on probability theory. So a lot of input has come this way into machine learning. Also one can say that there are different kinds of statistics or the descriptive statistics who only concern with really a good way

describing examples-statistical examples taken, without any ambition to draw any wider conclusions from it from inferential statistics where the ambition is to really draw wider conclusions on whole populations, and of course machine learning is much more dependent on call inferential statistics. If we turn now to data science this area doesn't present much new, it was launched as an umbrella concept for statistics actually augmented with some other data analysts, analysis techniques developing within the big data field and with advances from machine learning and as for statistics data science cover also the range of activities from harvesting up to the visualization appreciation and with the same old overlap in the data analysis stage. Actually the concept was coined around 1996-98 which is like 20 years ago but then it was really boosted in this last decade and it got a lot of publicity in various places, an example of that is this very special article from the Harvard business review called the "Data scientist: the sexiest job of the 21st century" which contributed a lot to this new areas fame and of course many young people looked for jobs and also of course the relevant education in this new area. In contrast to data science, the term big data primary refers to storage, maintenance and access to data even if the borderline the design is some rather blurred. The big data area is based on more traditional areas, one is very large databases others are data warehousing and distributed databases. There was born in late 90'ies right after the fast data growth due to explosion development or harvested sensor data as well as data from the web the parallel growth of traditional databases in a variety of sectors. We speak about big data, we talk about terabytes at least 10 upto 12 and already we talk about zettabytes 10 upto 21, but the size is not the only relevant parameter at the various variety of data types, the quality, velocity of generation are also important factors to increase the complexity. This was an excursion in some neighboring areas. So now we turn back to machine learning and an overview of this course and as you have understood we have eight weeks and the first three weeks have an introductory character, so the second week will be a more abstract description of a characterization of learning problems and the third week will focus on different kinds of representations. The motivation for having this kind of part of the course is that all kind of learning have to take place and within some kind of representation and therefore it's crucial to have some knowledge about those kinds of representations that are most common in artificial intelligence. If you already have a good view good knowledge on that, of course you can go very fast on that part of the course. The three middle weeks of the course four to six are kind of the core parts of the course, so two weeks focus on learning in symbolic representations the first only inductive learning in situations where you have no theory or rather weak theory, that backs up the learning process while in the second week we will look at situation where we really have a pretty strong theory in order to maintain theme of the domain you're working in sense learned on the border what is already now. Week six is done as a contrast learning in a sip sub sub symbolic representation particular to issue the URLA code and of course also here there is a distinction between situation where you have some rest prior theory or not that it doesn't seem relevant to divide that up in the same fashion as for wait for a week for it. So finally week seven and eight focus on goes in two directions, actually week seven looks into the context of cognitive science even if in theory machine learning algorithm could just be pure computer science algorithm, normally the way one design this kind of techniques is based upon some intuitions and some knowledge and some inspiration from the way natural systems works to contrast to artificial. So this week we'll look into some models of learning from various cognitive sciences that has very much inspired the way we do much machine learning. Finally week seven and it gives an overview of the current very vast set of tools and resources that are available at this point for doing working machine learning, it's not an easy area of it because of the current very strong industrial interest which promotes the development of tools and resources very much. So therefore the hope is to bring some order and some overview or what happens in this area. This slide is aimed at giving you a clearer figure out the order of how one should follow this course, I mean of course the bottom line is always possible and meaningful and may be advisable to go week by week, from week one to week eight and but as I understand there is not always as strong preconditions anywhere. So i just want to say a few things here was a that week one week two then week three are strongly advisable to follow in the beginning, because they give introduction to the area and the abstractions and the general knowledge that is supposed to be a background for that. And then a week four, five and six day they could be taken in different orders but still as this course is set up it's more favorable to even following week four, five and six and in the given order because many things that turn up in week six in the some symbolic representations in a way mirror somehow what really happens in this symbolic case in week four and five and as it's now set up is treated maybe more in detail and more thoroughly during this week's. What is more independent is of course week seven and eight because they are extensions of the course in a more theoretical and more background direction or in a more practical direction, so of course you could probably feel free to look into those whenever you like during your work with the course. So now I'm going to turn to the individual weeks of the course so my intention is to shortly comment and the content each of the week to give you a flavor and a feeling for what you can anticipate to come and during the course. I will not go into depth, I will make just a few observations and comments for each part. So this week we will focus mostly in the coming two lectures on artificial intelligence, we'll talk both about the roots of the area sixty years ago and what's happened then and how the profile of the area looking at that point and then I will make a more general overview of the areas as such and place machine learning within the whole area finally there will be a lecture about applications of machine learning and theoretically that lecture could be placed anywhere within the course but I've chosen to put it here very early because whatever technique we were going to talk about it will always refer to some applications so therefore it's good to have some reference point of this kind early in the course, and finally I want to comment on the general structure of each week so every week will you have an introduction area lecture and which sometimes could be a little short or sometimes a little longer and then there will always be a last lecture which we call tutorial and tutorial here means preparations and introduction of the assignments for that week. So in the second week of the course the goal is to characterize the various situations which machine learning tries to tackles and I'm going to be very general that a majority of machine learning takes want to tackle situations where you have a lot of set of examples and you study those examples and from those examples you want to be able to classify new examples or build some abstractions. So while on the lecture is trying to sort out what is an example, what is an abstraction, how can you characterize an example and so on. Then of course there are different modes of working here so one way of working is what we call supervised where all examples are pre classified so for example we say and let's feed the system with a number of images of a cat and all the examples are cats and overseeing that a system do is to build up an abstraction of a cat. The contrast to that is unsupervised which means that we feed the system with any image let's say any animal or any category, and then it's up to the system to sort out what is a cat, was its dog what is something else which is a more difficult task. Also one can play with the way you use the examples so you if you want to supervise better or more you can say these are positive examples, these are negative examples, these are extreme examples, these are typical examples which also facilitate or simplify the learning process. And our distinction that is important is whether you have all examples they are existing in the beginning or whether all the examples come one at a time in a continuous fashion, the later is also more difficult to handle and another aspect is whether you learn and let's say in the blue so only thing that they say says your examples and you build a small theory of abstractions from that the other extreme means you already have a strong theory a really strong knowledge base which you can already use for problem solving, and then what you want to do is to improve that knowledge base which means that you learn on the border of prior knowledge. Yeah what an example of that is what call reinforcement learning which is a special case for examples it used a lot for example for training robots it's an example, where you want to learn a robot to move the arm in an optimum way by of course have to start with some basic movement schemes so there is a series for moving but you can improve the movement by adopting a number of parameters but of course this means that you have to give the system feedback all the time on good movement bad movement and so on and this is a specific kind I would say a strong theory learning. Finally the most ambitious thing here is to learn representation and by that it means that normally machine learning you have a fixed representation you have a fixed model of the domain you're working with then you always learn within that model, however more and more people understand that it's not very realistic because systems have to be robust and you cannot think about everything from the start so therefore, there is also need for the system to learn the representation themselves learn which kind of objects which you can see there which kind of features we could see there and so on, which is then another step ahead in the advancement on the field. So the third week will be about various representations used for machine learning and one can say there are two parts here first we can look at the typical symbolic representation that additionally being used as the basis for building learning algorithm and to such a decision trees and bayesian networks also one can add and that will also be referred to there is this we can work simply with arrays of different dimensionality which is the third symbolic way of handling this. And the other group are sub symbolic and that's the typical type of representation in artificial neural network and that we will spend some time on, we also mention another kind of simple symbolic representation genetic algorithms we will not spend so much time on that on the course here at all but we will mention it and now and then because it complements the picture yeah. Finally I would comment on a slight complication we have here because of course in artificial intelligence we used we want to use machine learning to improve our programs or our problem-solving programs and actually most systems built historically in artificial intelligence are based on multi-paradigm programming a kind of combination of logic programming, functional programming, object-oriented programming and of course logically it would be so that when we learn we want to modify those programs. But slightly contradictory here most work in machine learning have been focused on the earlier forms of representation not them so much about learning functional logic and object-oriented program as such, but of course it's important you know about this contradiction and it's also important to know the properties of these representations because in the end results from learning processes had to be fed into thee the existing systems. Now we turn to the core part of the course three weeks where we will look at different kinds of algorithms, that has been important in the development of machine learning and in the first week, week four we will look at inductive learning algorithms that work on symbolic representations but also primarily with in cases where we are weak theories in the sense that we are almost none or very little prior knowledge that guides the learning process so the first case here is supervised learning this is the classical case when we learn from pre classified examples and build up a well-defined set of defined abstractions. The opposite case is unsupervised learning where we have non classified examples and this algorithms themselves have to sort and categorize these example and build the appropriate abstractions. The intermediate case here we will look at is called instance based learning, in the two first cases the algorithms build concrete explicit abstractions out from the examples and store these abstractions while in the last case we have a system philosophy where we won't store all examples and do not build any abstractions that means of course we have to have a structure of examples and when new examples come in we have to relate them to the old and place them in the structure. So in week 5 we will still look at learning a symbolic representation but look at a few cases where we combine the inductive learning with the existence of strong theories and the first case I want to mention is inductive logic programming. I already mentioned when we talked about representation that it was a kind of contradiction that many machine learning algorithm worked in representations that were not the most important and common ones for really building systems in artificial intelligence. Logic programming is one of the very strong paradigms that has been used to build many systems in artificial intelligence and 20 years ago the idea come up that why not we engineer our inductive algorithms so that they can directly learn logic programming elements, and the first step is of course only to mimic what one did earlier in the other form of representation but of course the next step of is that because we now know now do everything in the same formalism we can more easily combine prior knowledge or prior theories expressed in logic programming with the new kind of abstractions that we built up from the inductions of examples and so this is essentially the key idea logic programming. Another aspect is of course that for learning in cases with no prior theory we need often many examples because many examples may be bad, may have noise may be typical or not typical and so on, so this means that in order to play at there's noise and different of my examples we need many examples, however we could also manage to learn from fewer examples but then in a way we have to make a sanity check of the examples and screen the examples in the base in the light of the prior knowledge, so this means that if you know proprietary existing strong theory we can validate the soundness of our examples in the light of that theory and thereby we could still utilize induction from these new examples even if there are few. There are few other things that will be talked about in this week we will talk little about learning by analogy which essentially we have two domains and we have a knowledge base or a theory in one domain and we have a much weaker theory in another domain but we can assume that there are similarities between two domain so therefore we can infer abstractions in the second domain from the already existing one in the first. Finally there is also a need in many kinds of system to explicitly and conveniently directly add new abstractions not having to reprogram the whole system, but in this case no induction takes place its just that new knowledge pieces are easily added to the system that was referred to by the term being told. And finally we have reinforcement learning where we essentially also have a prior theory for example a prime theory for guiding the moment of a robot but then we want to optimize the functionality of that example that example of robot by incremental learning. So week six and this course covers machine learning algorithms in artificial neural networks and there is a history of this that no one can say that that the structure of lectures is based upon this history so we start by talking about something called perceptron which was a very early attempt of an artificial neural network people were very positive in all already in that 1960s about the successes one could have using this technique however there was a backlash people showed that the original designs wasn't really possible to go forward with and then one can say it was a big long period where not much happened, so but in 1986 the like 20 years later this area revived again and there were some key results in the early 1980s which showed that networks with many layers actually the perceptron has just typically a single layer initially but if we build networks with many layers, then and special kinds of propagating values forward and backwards in these networks, one could show successful results. And so this happens in a 1980s and we will talk about what happened in that period. The next step in this process is that in order to really come forward and solve real problem with this technique you cannot have general networks you have to anyway adopt the networks so they fit certain kinds of issues and certain kinds of problems, and image recognition has been of the one of the key application areas and there are special techniques that we will cover, an example is convolution networks which has specific properties that makes them suitable for handling each problems in image recognition settings. Also the initial networks was not ideal for modeling temporal phenomena like sequences of sequences of events and in order to handle that another kind of network type appeared called recurrent networks. And these two kinds of networks had a large importance for the development and we've also worked further looking at more of these kinds of recurrent networks examples of that and field networks and boltzmann machines. So since the 90's there have been a continuous development of this but I would say the next very strong revival of this area was once in the in the early 2000's where this area blossom and essentially the term used today for the state of the art of this technology is essentially deep learning, I mean deep learning was coined a long time ago and the reason for the for the name is that the successful neural networks need to have many layers that have to be deep in a sense, but today of course this term covers a lot more because it encompass many of the recent advances or the algorithmic side, and I would say all the one and very important characteristics of what you can do today in deep learning is to take the kind of last step in the learning problem which is to learn representations that is not only learn within a fixed model, but really incrementally modify and develop the model itself.

Week seven of the course and one can say it is a separate part of this course the reason why is there is to really show to you that machine learning advances have not happened in isolation as just parting to computer science many times the designs have been inspired by insights fetched from other areas and which one can characterize as cognitive science and part of cognitive science are of course cognitive psychology, neuroscience, tribology, linguistics and so on and then during this week we will primarily look at three things, we will look at some models in cognitive psychology, we have to do how we learn categories as humans we conceptualize the world but also modern cecum named cognitive psychology that describes how we as humans solve problems and how we develop expertise so that's one thing another thing we will locate is some model from your science which is actually models primarily we will look into how our human perception works. And the third thing we will look at it during this week is looking at theory of evolution because in a large scale evolution is a learning process and one can say on an abstract level so they think things we will look into in two week 7. Finally for week number eight we will talk about tools and resources useful for doing machine learning work when an area is new, there are no most master special tools support systems available but when area comes into the limelight, as machine learning of time today suddenly there is a great interest for all kinds of tool and language provider to see to that that area is supported by their tool or language. So today there is explosion of initiatives of various kind to support work in machine learning and of course this also has influenced the success of the area lately. So the goal of this week is to really try to sort out for you what kind of tools are they or what kind and so on because, it is a kind of jungle today and so not so easy to understand what is what and so in this week we will look at four categories of things. So the first category is named general technical computing tools and these are general toolboxes existed for a long time. A typical example of such a toolbox is Mathematica and which traditionally supported a lot of things but it didn't support machine learning but today does so this is an example of a kind of tool but of course others in the same category and we will go into that. And the next category is term here general support software support which means all kinds of languages, support systems, libraries and so on that can help you to work in this area. And just take a very simple example so for example language like python is very popular today, first of all because it's a very commonly used language very popular language to use for programming but also that when I made some effort to easily integrate function in that language that facilitate working machine learning. The third category is more dedicated platform that's been developed specifically for this purpose, for the purpose of supporting machine learning. And there are a lot of companies who develop those, one well-known example there is for example google that has developed a platform called tensorflow that is used in many cases but there is there is really a large set of similar platforms. Finally something which is very important is that there are easily accessible and openly available datasets which you can test your systems and algorithms on and this also now starting to emerge a large set of such data sets well knows that data sets it's image net which supplies a large number of images for all kinds of common nouns. So before we end this introductory video want to say a few things about literature, so first of all I want to say that this course has no dedicated course literature, not one single book that that we follow. There are so many books in machine learning and as you have understood hopefully now the area is pretty old, so there are very old books there are very new books and there are very the books that are like in-between. I would not recommend you at this point to read the very old books those from 1980's and so. I would recommend the new ones because some of them are very good, some of them are always bestsellers at least within the context of an academic era like this, but of course the new ones you probably have to buy because you may maybe significant evil burrow of them is absolutely not possible to access them on the way but for the books in between and then I given you an example here the book machine learning by Tom Mitchell it was originally published in1997. But there are also later editions of course but for this kind of book older editions may be accessible of what because they are not bestsellers anymore so there are no interest to the to the publishers. So for me maybe like a default book that could be useful you want to read more about the area is this book machine learning by Tom Mitchell. Newer books I also given you two examples here there is a book from 2014 which is a general introduction by Ethem Alpaydin but there also a book where much talked about to her this deep learning book by Yoshua Bengio and his associates. For this course on every week you will get a number of literal references but this literal reference will typically not be books and they will typically not be survey articles either. So maybe a special feature of this course is when I give you the literal references in the various weeks I will mostly give you reference to original articles, and or articles that represent the kind of breakthrough in this area at various periods and various point in time. So this is the end of the introduction thanks for your attention, the topic of the next lecture will be the foundation of artificial intelligence and machine learning.