Fundamentals of Artificial Intelligence
Prof. Shyamanta M. Hazarika
Department of Mechanical Engineering
Indian Institute of Technology – Guwahati

Lecture – 27
Introduction to Machine Learning

Welcome to fundamentals of artificial intelligence today we start our final module of the course this module is on machine learning. Machine learning addresses the fundamental question of how to build computer programs that could learn automatically from experience. In fact machine learning is the most rapidly expanding technical field in the intersection of statistics and computer science and lies at the core of artificial intelligence.

The rapid progress in machine learning is because of A. development of new algorithms and theory of learning and B. the ongoing explosion in availability of data and low cost computation. The adoption of machine learning methods in science, engineering and technology and commerce have actually led to data intensive decision making in domains as diverse as healthcare, manufacturing, education financial projections, agriculture policing and many others.

Machine learning is concerned with two fundamental problems one that I have already mentioned is the problem of creating computer programs that could learn automatically from experience. The second is about identifying statistical computational information theoretic laws that govern learning in such computational systems in particular and in any organization in general.

The idea of machine learning is to explore the questions underlying the theory of learning as also working on means and ends to take forward the highly practical computational systems that have been fielded in many applications. Machine learning within AI is the fundamental choice for problems such as in computer vision speech recognition and robot control. So today we would briefly introduce machine learning and try to understand machine learning vis-a-vis artificial intelligence.

(Refer Slide Time: 04:00)

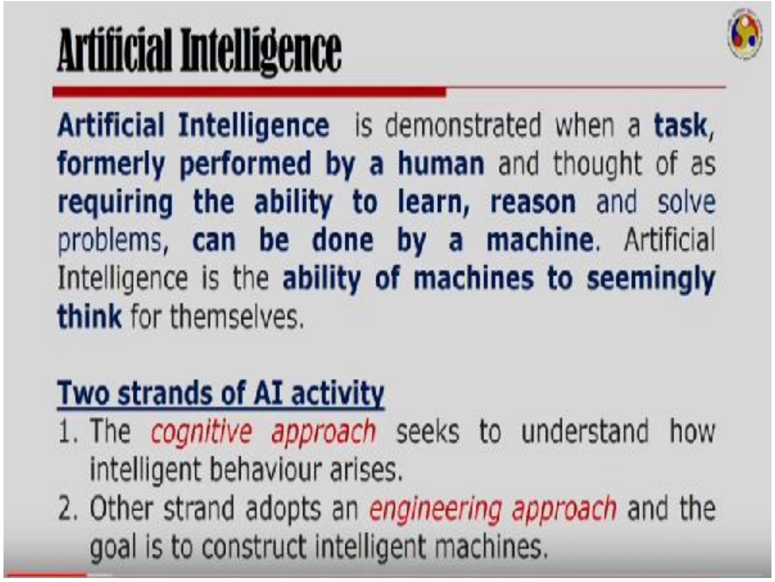

Kindly recall that in our first lecture we had given a definition of AI and defined AI as the ability to do a task which is more easily done by a human. So AI is demonstrated when a task which formerly performed by a human and thought of as requiring the ability to learn, reason and solve can be done by a machine. Now artificial intelligence is the ability of machines to seemingly think for themselves.

There are two strands of artificial intelligence activity one which is referred to as the cognitive approach actually seeks to understand how the intelligent behavior arises and the other strand actually adopts a more engineering approach and the goal is to construct intelligent machines. Now whenever we are talking of intelligence recall that we say intelligence is coming from knowledge and intelligence is the ability to solve problems.

(Refer Slide Time: 05:29)

Now that intelligence which gives you the ability to solve problems could be looked at being accrued from the knowledge that we gain through experience and therefore we can think of 2 broad approaches to solving problems. One is to treat every problem to be solved using first principles. So here what I mean by first principles is by the agent to solve a problem by reasoning about its actions exploring combinations choosing one that lead to the solution.

Whereas the second approach is to Harnish the knowledge that's gleaned from experience or from other agents. Now this is not that watertight in themselves it is not that we will look at problem solving only by the first approach or only by the second approach. We have seen that many of the algorithms that we have discussed so far are driven by some knowledge of the domain. Even if we are searching through our problem space.

Therefore problem-solving demands an eclectic mix of search methods that is operating upon different knowledge representation and that is what makes machine learning so very important because intelligence requires knowledge and the question is where does a system get that knowledge from.

(Refer Slide Time: 07:12)

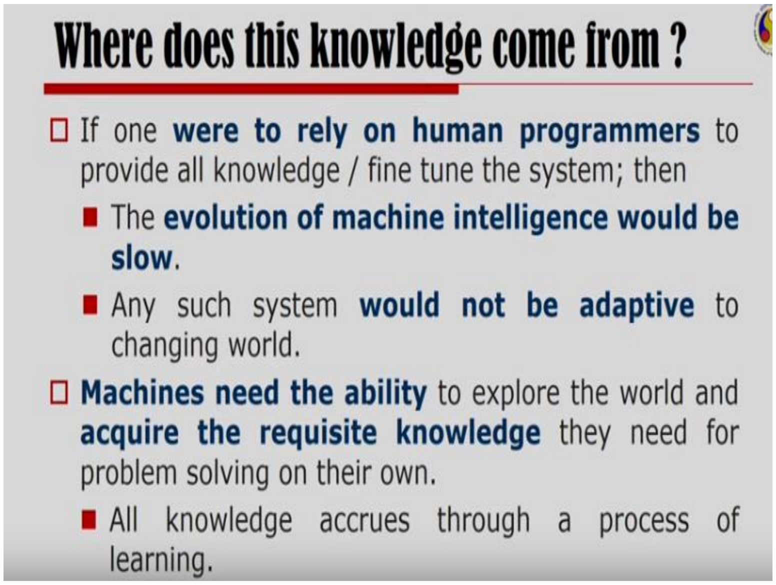

So if one way to just rely on the human programmers to provide all the knowledge or fine-tune the system. Then the very evolution of machine intelligence would have been very slow and any such system further would not be adaptive to the changing world for if we have given the system knowledge apriori and when the system is reasoning the world changes then a system could not adapt if it was pre-programmed.

Therein lies the need for the machine to be able to explore the world and acquire the requisite knowledge they need for the problem-solving. So such knowledge accrues through a process of learning and that is what makes machine learning so very important.

(Refer Slide Time: 08:10)

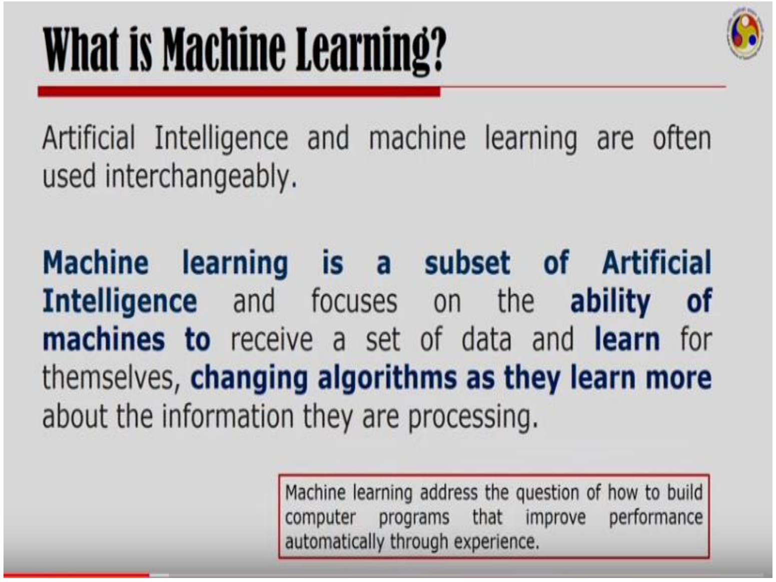

Now let us try and understand what is machine learning as I was telling you the very crust of anything being artificially intelligent is based on the fact that I should be able to learn about my environment and therefore these two terms machine learning and artificial intelligence are interwined and in fact are often used interchangeably.

Machine learning is a subset of artificial intelligence where the focus is on the ability of machines to actually receive a set of data and learn for themselves now when I say learn I mean change the algorithm as they get more information or as we will see later, we call it as they gain more experience. Now the very idea of machine learning is to get to the fundamental question of building computer programs that could improve performance.

(Refer Slide Time: 09:26)

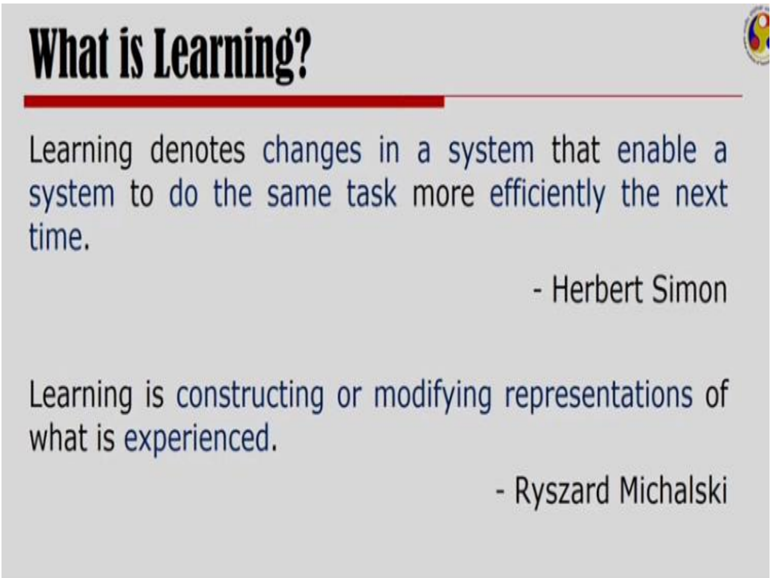

Let us try to understand where learning comes from here or what is learning for that matter. Now according to Herbert Simon learning is about changes in the system that enable a particular system to do the same task more efficiently the next time. So this is a definition you can correlate very well like let us say you have got a bicycle and you are trying to learn how to ride a bicycle. So in your first attempt, you are unable to control the steering, the handle of the bicycle and do the pedaling simultaneously and as you do it over time you gain more experience and you are able to efficiently do it. So when you are able to do the riding efficiently, we say you have learned how to ride a bicycle. So that is what Herbert Simon emphasizes on that learning is about some changes in a system that will enable a system to do the same task more efficiently the next time and according to Michalski learning is about constructing or modifying the representation of what is experienced under such a interpretation of learning we would try to understand

(Refer Slide Time: 10:58)

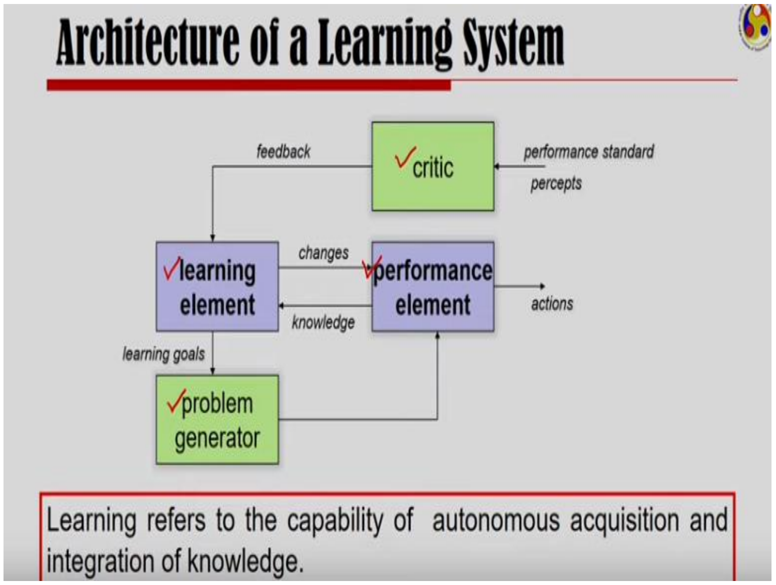

machine learning but then first let us look at the architecture of a learning system. So here learning refers to the capability of any autonomous acquisition and integration of knowledge into the system. A typical learning system has 4 elements a learning element, a problem generator, a performance element and a critic. We will look at each of them one by one and then a learning system cannot stay all alone it needs to interact with the environment to learn.

So there is the environment, the environment gives me some performance standard percepts and the agent that learns is able to act on the environment through a group of actions.

(Refer Slide Time: 12:01)

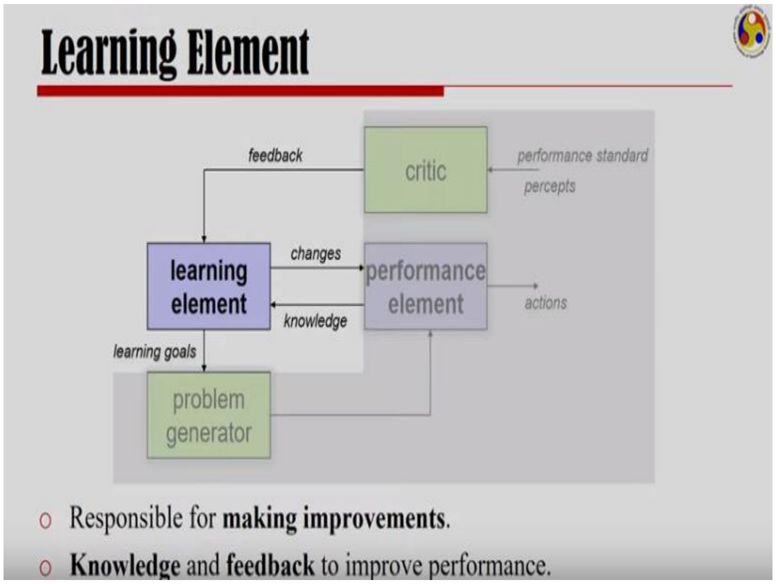

Now let us look at the learning element. The learning element is actually responsible for making improvements in the system as it learns. So the learning element gets the knowledge and it gets the feedback on how it is doing from the critic and thereby makes improvements in its own system, these improvements that it makes is nothing but what we refer to as learning.

(Refer Slide Time: 12:37)

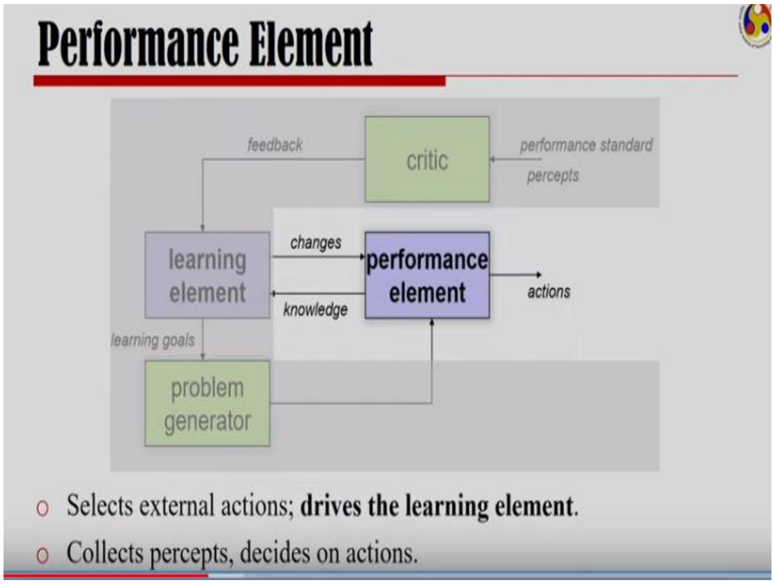

The performance element on the other hand actually drives the learning element for it selects what external actions are to be done on the environment and decides on the actions as it operates.

(Refer Slide Time: 12:58)

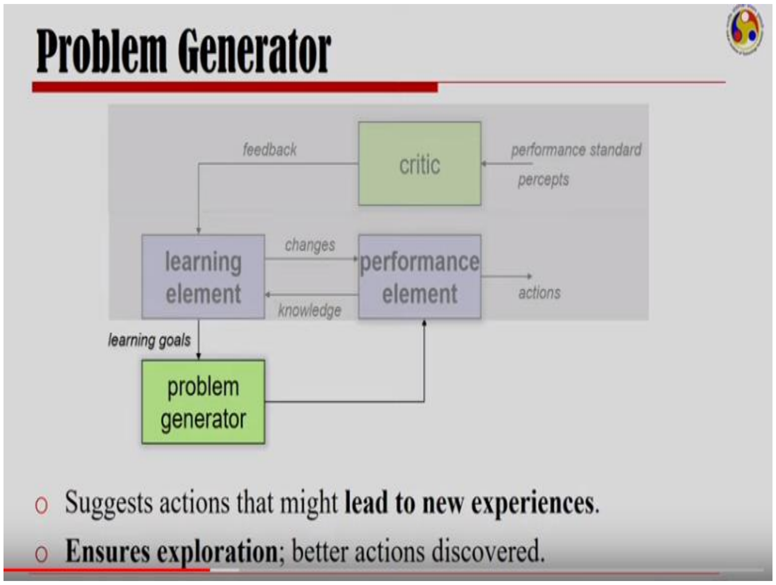

Now the problem generator possibly is the most important of the elements in a learning system for the problem generator is one which suggests actions that might lead to new experiences and these new experiences then might lead to learning. So the problem generator in a sense ensures exploration and allows us to have better actions discovered in the system, see how newer things are done that is get new experiences and then and only then we could learn.

(Refer Slide Time: 13:41)

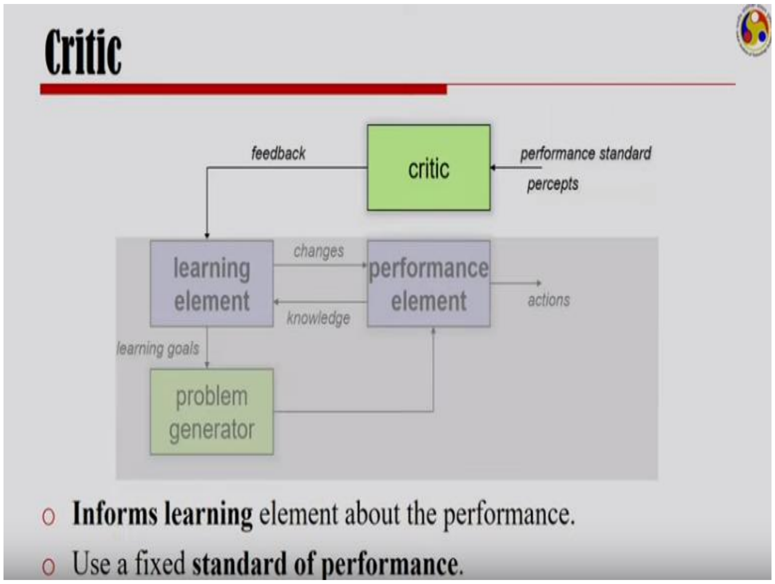

The other important element in learning is the critic now the critic is the one which actually informs the learning element on how it is doing. So the critic tells the learning element about its performance and in a way the critic is the one which actually measures what is happening to the learning system in terms of standards of performance. So given a system like this we can see that the problem generator gives us new experience. The performance element decides on the actions and generates new knowledge. The learning element then knows what are its learning goals and also gets a feedback from the critic on how it is doing, and this new experience is added on to make any changes if required and we say learning happens.

(Refer Slide Time: 14:56)

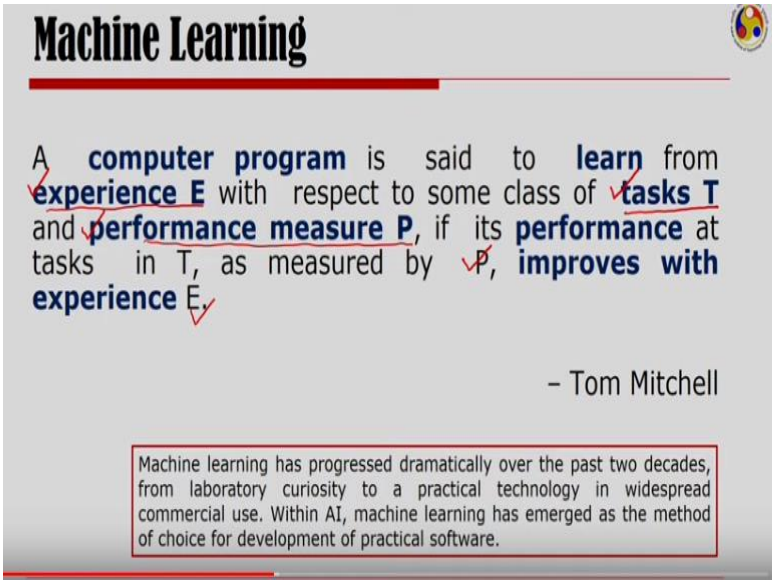

Now to relate this to the definition of machine learning let us look at the definition given by Tom Mitchell in the 90s. So he says a computer program is said to learn from experience E with respect to a class of tasks T. So we are trying to learn some tasks T by gathering experience E and our performance is measured in terms of P. Now if we can see that the P improves as we get more experience that is P improves with E, we say the computer program is learning. So here it is about the task T the program is trying to accomplish and we are talking of experience E and the performance of the task is measured in terms of P. So the program is said to learn if P improves with experience E. Now with this very simple definition of about generating computer programs that could improve its performance automatically from experience, machine learning has dramatically over the past decade moved from laboratory curiosity to a practical technology in widespread commercial use. Now as I already mentioned within artificial intelligence machine learning has emerged as the method of choice for development of a number of practical software in areas as diverse as computer vision, speech recognition, natural language processing ,robot control and many others.

(Refer Slide Time: 17:13)

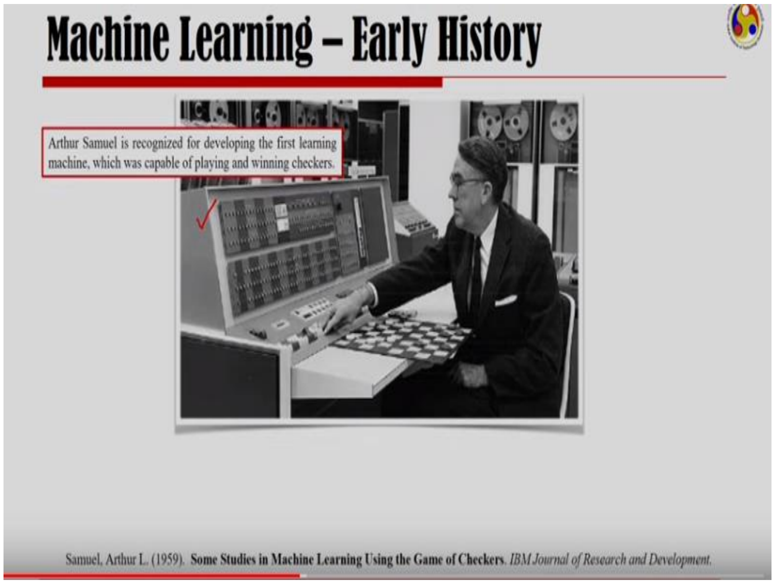

Let us now quickly look at the early history of machine learning. So here on your screen you can see Arthur Samuel playing with an IBM 700 computer here and he is playing a game of checkers against the computer. Now Arthur is recognized for developing the first machine which was capable of playing and winning checkers, till then IBM 700 was not that very popular. Arthur playing it live on television against the IBM 700 a game of checkers was exciting for the period and Arthur actually was the one who coined the term machine learning.

Possibly the Samuels checker playing program was among the worlds first successful self learning program and as such a very early demonstration of the fundamental concept of artificial intelligence. According to Arthur Samuel machine learning is the field of study that gives computer the ability to learn without being explicitly programmed.

So you do not need to explicitly program them but then if the computers can learn how to play the game then we say that it has learnt. Now this paper by Arthur some studies in machine learning using the game of checkers is possibly one of the first papers on machine learning which was published in 1959.

(Refer Slide Time: 19:13)

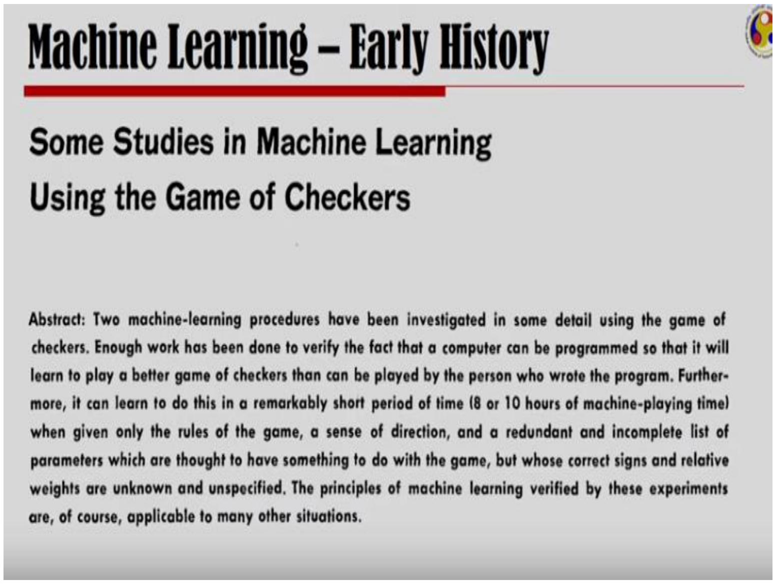

And I think it is important for us to look at that paper. So here is the paper from the IBM Journal.

(Refer Slide Time: 19:21)

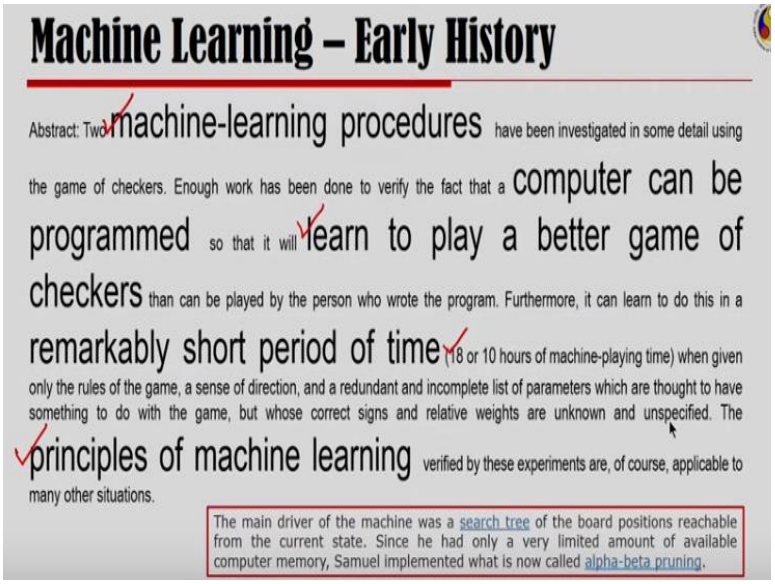

Let us read the abstract for it is interesting. So in the abstract he says he have given 2 machine learning procedures and he have shown that the computer can be programmed to learn to play and this is exciting for he says that it can be programmed to learn to play a better game of checkers then can be played by the person who wrote the program and this was done in a remarkably short period of time some 18 or 10 hours of machine playing time and this paper sets the principles of machine learning.

Now just for our own curiosity it is important for us to know that the algorithm of alpha-beta pruning that we have learned while we were discussing game playing in this course was actually implemented by Samuel in his Checkers playing program first. So the main driver of the machine was a search tree of the board positions reachable from the current state and since Sam had a very limited amount of available computer memory. Samuel implemented what is now called alpha-beta pruning in his checker playing program.

(Refer Slide Time: 20:52)

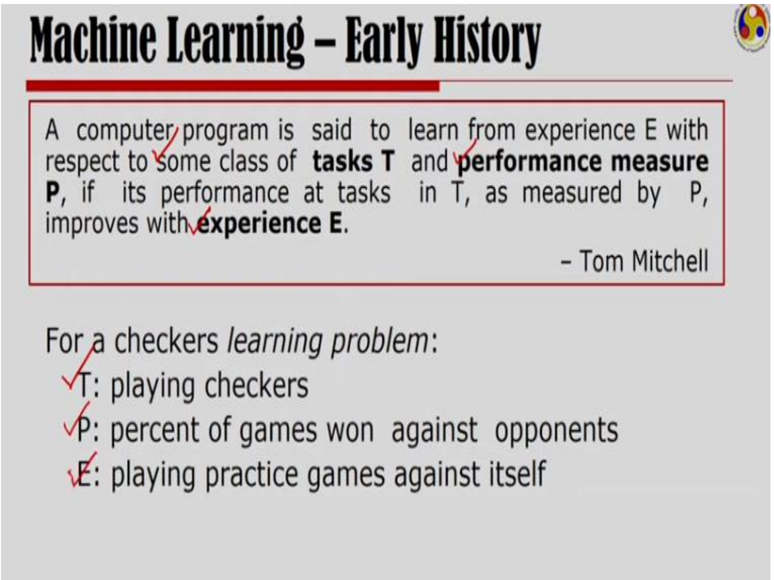

So that is early history of machine learning and if you look at Tom Mitchells definition vis-a-vis the checker learning program then Tom Mitchells definition talks of some class of tasks T. So here it is playing checkers is our task and he talks of a performance measure P. So here it is about the percentage of games won against the opponent and experience E is about playing practice games against itself. So this is how we can look at the checker playing program as our machine learning program.

(Refer Slide Time: 21:39)

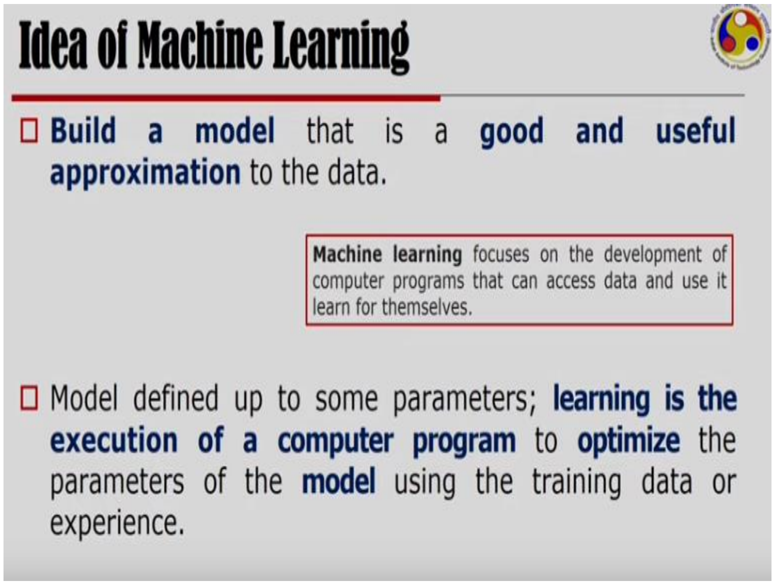

Now the idea of machine learning which must be clearly understood is about building a model that is a good and useful approximation to the data that is the idea and it is about development of programs that can access data and use it to learn for themselves. So the model is defined up to some parameters is what we are looking for and learning is the execution of a computer program to optimize the parameters of the model using the training data or experience.

(Refer Slide Time: 22:23)

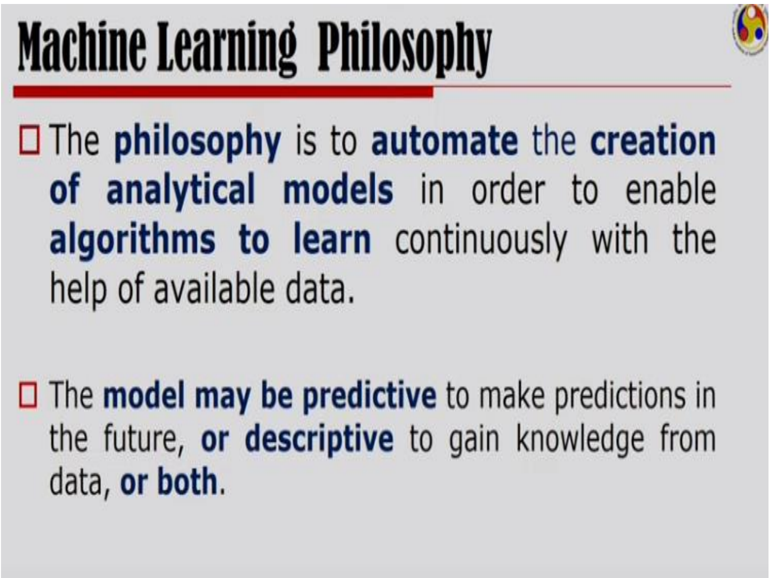

The philosophy is to automate the creation of the analytical model in order to enable algorithms to learn continuously and one needs to understand that in machine learning the model may be predictive to make predictions in the future or it could be descriptive to gain knowledge from data or both. One needs to really realize that machine learning is about identifying patterns and making decisions with minimal human intervention that is the philosophy of machine learning.

(Refer Slide Time: 23:05)

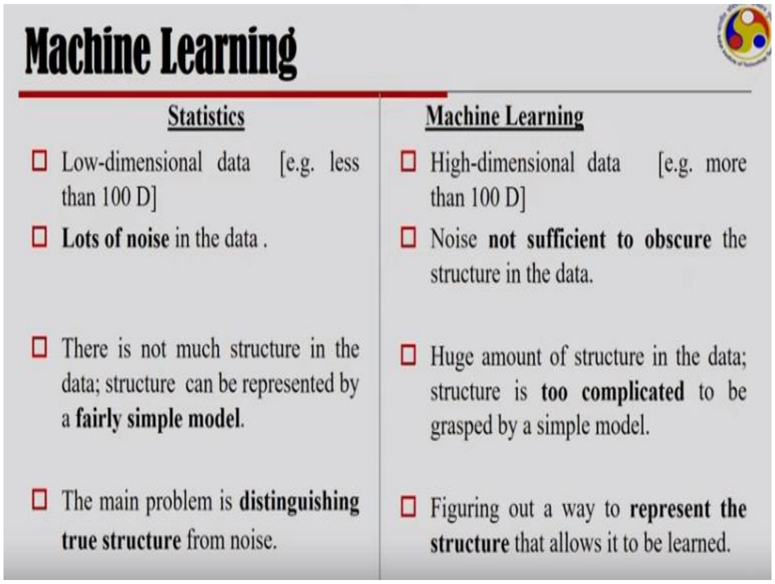

So I started my talk saying machine learning is possibly the most rapidly growing technical area in the intersection of statistics and computer science and at the core of artificial intelligence. So it is important that we look at how is statistics and machine learning related to one another or different from one another for that matter. In statistics we are usually involved with low dimensional data whereas in machine learning we are concerned with high dimensional data.

Noise in the data is something that we did not welcome in statistics here in machine learning we thrive, and we consider noise not sufficient to obscure the structure in the data. In statistics there is not much structure in the data structure can be represented by a fairly simple model, whereas in machine learning we are concerned with huge amount of structure in the data and the structure is too complicated to be grasped by a simple model.

The main problem is distinguishing true structure from noise in statistics, whereas in machine learning the main problem is about figuring out a way to represent the structure so that the structure will somehow allow it to be learned.

(Refer Slide Time: 24:40)

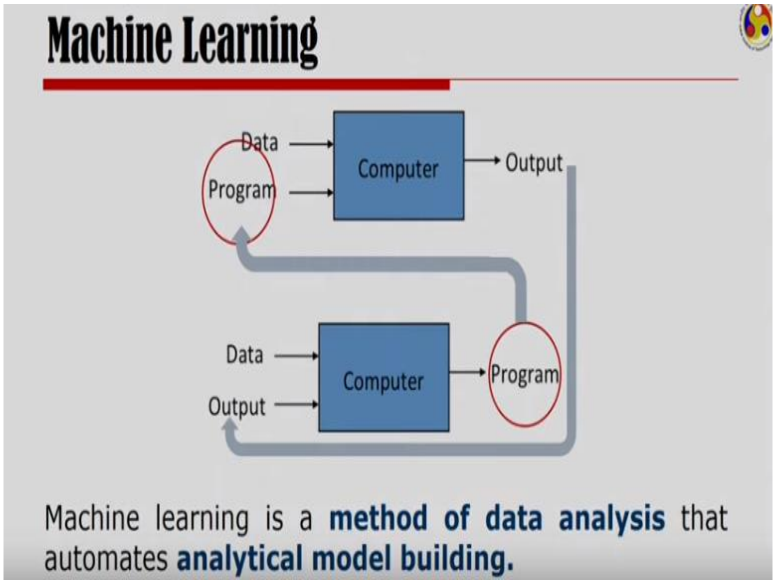

So given that let us now have a look at the relationship or the difference between traditional programming and programming as we understand in machine learning. Recall that in traditional programming you have data you have a computer and you give the program and we get our expected output. So traditional program is about having a data in a program to generate some output whereas in machine learning we have data and output and what we are interested is in getting to a program or getting to an analytical model.

And this is what we are interested in machine learning now we can take this program that we get in machine learning plug it in to the traditional program and generate more output. So that then we can drive the output of the machine learning program and this is how we can take forward machine learning. So machine learning is a method of data analysis that somehow automates the analytical model building process.

(Refer Slide Time: 26:11)

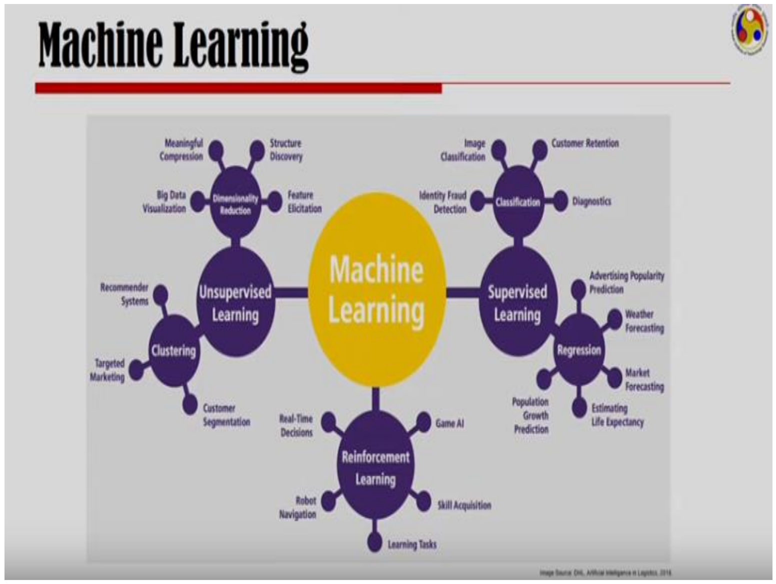

Now having looked at the very idea of machine learning and the philosophy behind it let us quickly look at the different types of machine learning that is available. Machine learning can be seen under 3 categories one called the supervised learning, another the unsupervised learning and third the reinforcement learning. We will look at each of them one by one.

(Refer Slide Time: 26:44)

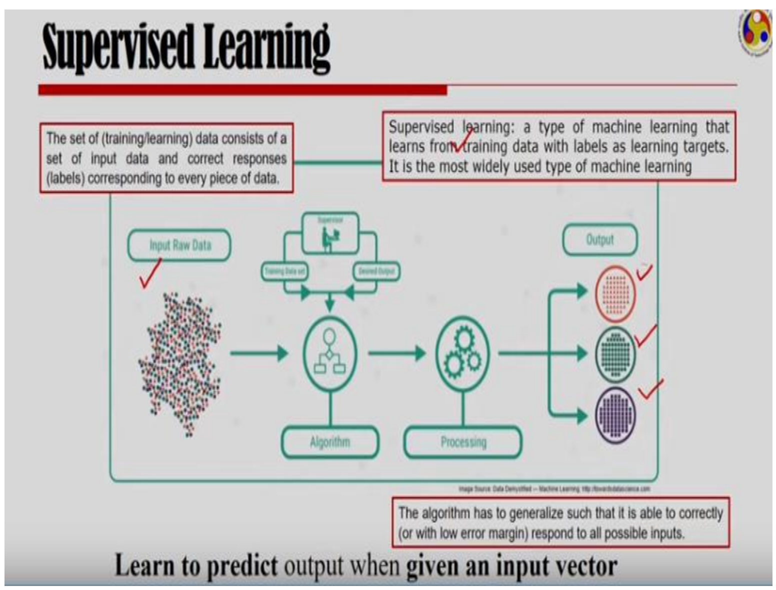

So supervised learning is about learning to predict output when we are given an input vector, what it means is that in this type of machine learning we learn from training data that is data where there are labels and so we have raw data coming to us but the raw data is with labels and therefore when we are trying to put it into different categories somehow for the beginning portion we know what each of them is and using that training data we try to learn.

So the training data with labels as training targets is used to learn and thereafter, we can use this analytical model for categorizing data that it has not seen so far. So the set of training or learning data consists of set of input data and correct responses to every piece of data and supervised learning the algorithm has to generalize such that it is able to correctly or with very low error margin respond to all possible inputs.

So in this example here I have certain circles, triangles and squares and at the first some form of supervisory control tells me what is each of them because each of them comes with the correct label and then puts it into each category and later when I see something which I have not seen so far, I am able to tell what it is, so this is supervised learning.

(Refer Slide Time: 28:36)

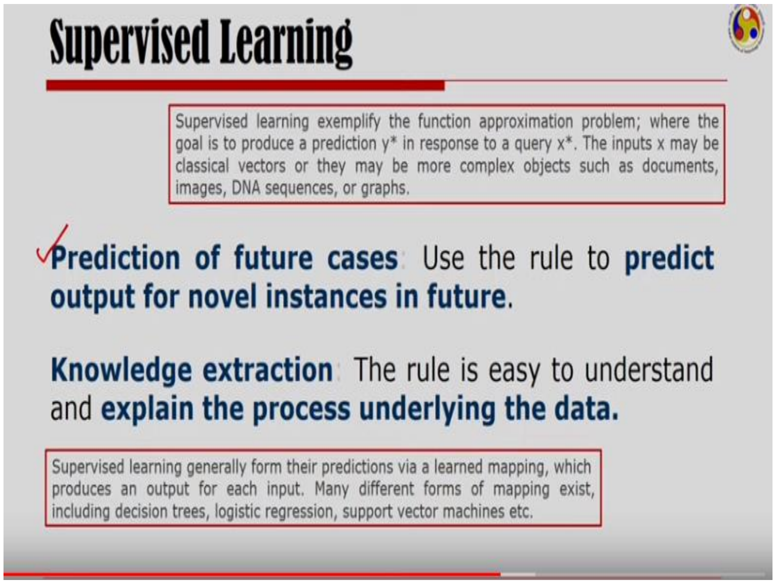

Supervised learning actually exemplify the function approximation problem. So here the goal is to produce a prediction in response to a query, the inputs maybe classical vectors or they may be more complex objects such as documents, images, DNA sequences or graphs. So we should be able to tell what they are or produce a prediction for them. So there is the rule to predict output for novel instances in the future or we should have a rule that is easy to understand and explain the process underlying the data.

So supervised learning generally forms their prediction via a learned mapping that is a mapping that produces an output for each input. Many different forms of mapping exist including things like decision trees, logistic regression, support vector machines etc., which we will see during the course of our discussion in machine learning. Supervised learning can be divided into two types.

(Refer Slide Time: 30:01)

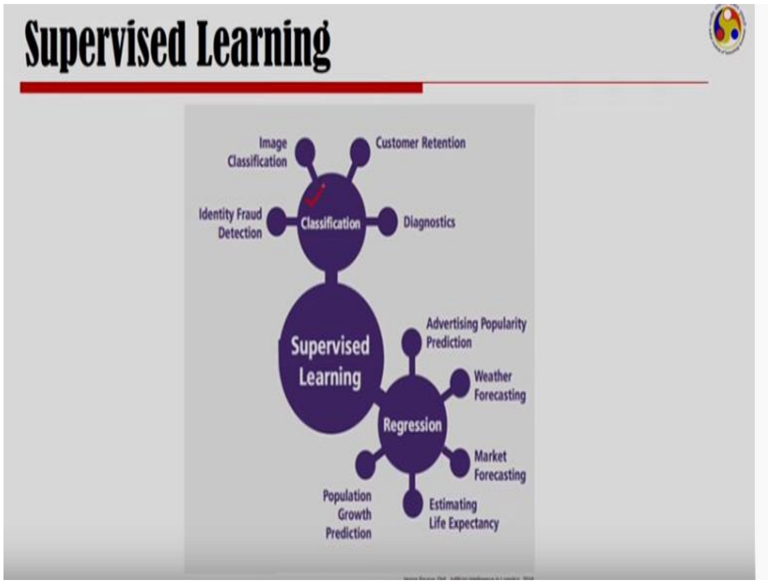

One called classification and the other called regression.

(Refer Slide Time: 30:08)

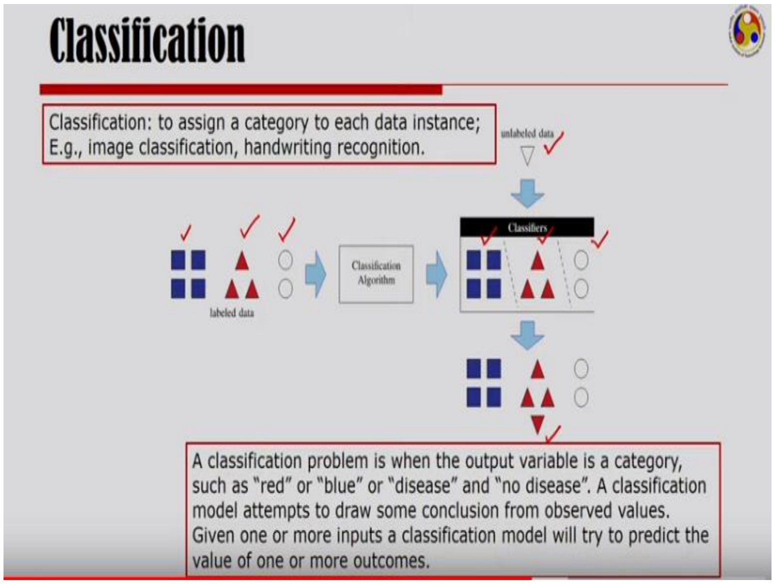

Now classification is about assigning a category to each data instance like I can classify handwriting, or I can classify images and here is a diagrammatic representation of the classification algorithm I have labeled data which are squares, triangles and circles and I have an classification algorithm. So the classification algorithm actually assigns a category to squares another category to triangles and another category to circles and whenever I have an unlabeled data coming to me the algorithm is strong enough to put it in the right category.

So a classification problem is when the output variable is a category like it could be red, blue or it could be disease or no disease and a classification model as we have seen here attempts to draw some conclusion from the observed values. Given one or more inputs a classification model will try to predict the value of one or more outcomes.

(Refer Slide Time: 31:24)

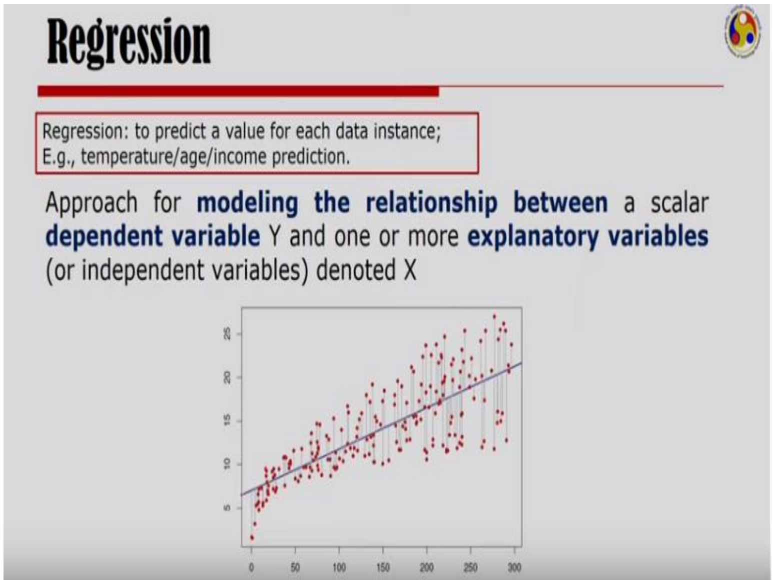

So the other form of supervised learning is referred to as regression. So regression is about predicting of value for each data instance. So it is an approach for modeling the relationship between a scalar dependent variable and one or more explanatory variables. So if I have a variable here on the x-axis and a variable here on the y-axis, the relationship between them given by these points here I can think of explaining this via this simple line here so that would be a regression.

(Refer Slide Time: 32:08)

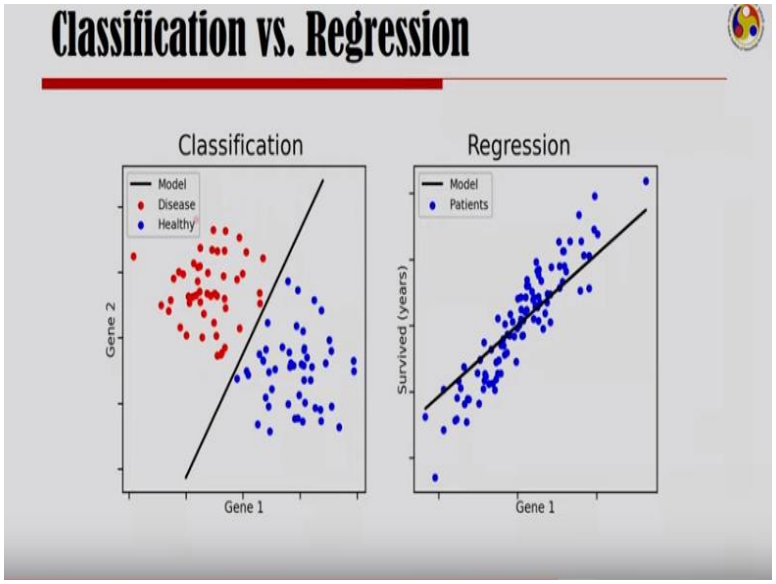

Now if you look at classification. Classification is about drawing boundaries between two things. So here I have on your screen plus and circle and classification is about drawing a boundary between them on the other hand on the screen on your right now I have a number of points given via circles and I am trying to fit in a line which is shown in dotted red as the line that explains the XY variables.

So taking a more interesting example if I have genes which are either diseased or healthy and they are plotted here I can draw a classification that separates them out, whereas if I have a relationship between some gene and its survival in the patient then I can draw a regression between them.

(Refer Slide Time: 33:07)

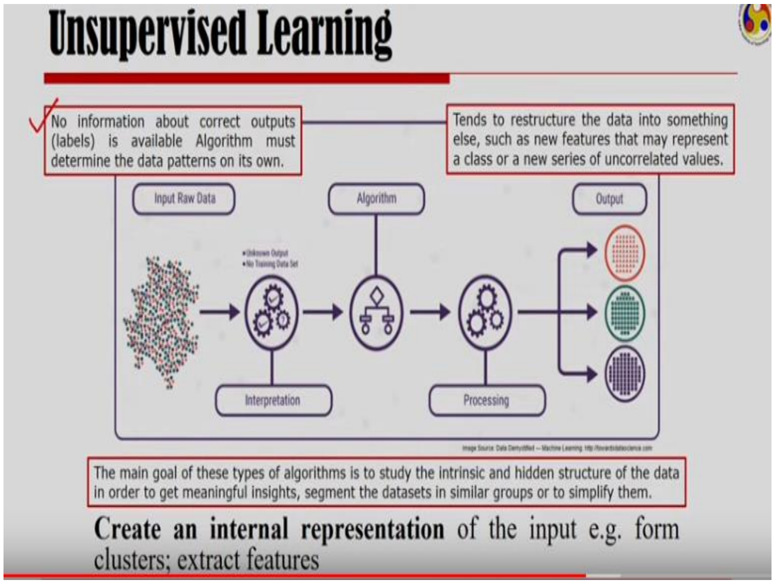

The other type of machine learning is called unsupervised learning. So unsupervised learning is about creating an internal representation of the input for example getting to the clusters or extracting features. Here no information about the correct output labels is available and the algorithm must determine the data patterns on its own. So the input is without any labeling and the algorithm needs to figure out what this input is.

So it tends to restructure the data into something else such as new features that may represent a class or a new series of uncorrelated values and the main goal of unsupervised machine learning is to study the intrinsic and hidden structure of the data in order to get meaningful insights and segment the data sets into similar groups or to simplify them.

(Refer Slide Time: 34:14)

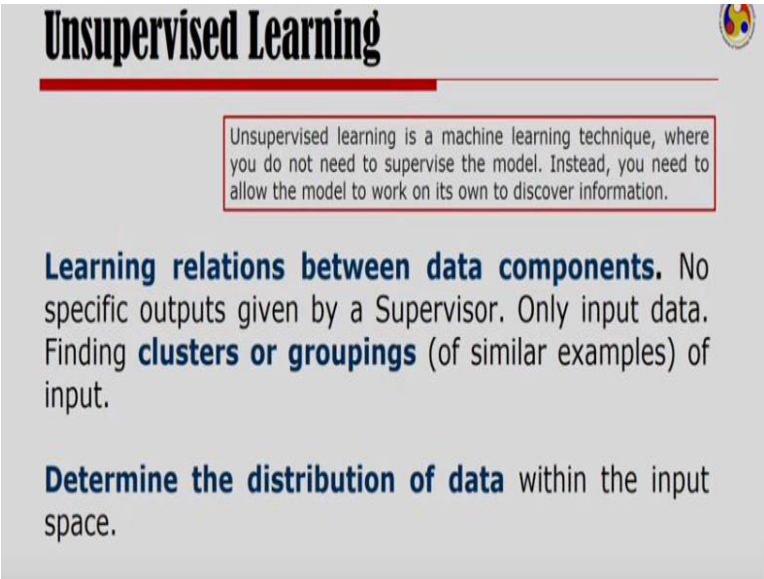

So unsupervised learning is where we do not have a supervision during the creation of the model instead one needs to allow the model to work on its own to discover information. So it is about learning relations between data components no specific outputs are given only the input data is there and its about finding clusters or groupings of input and determine the distribution of data within the input space.

(Refer Slide Time: 34:50)

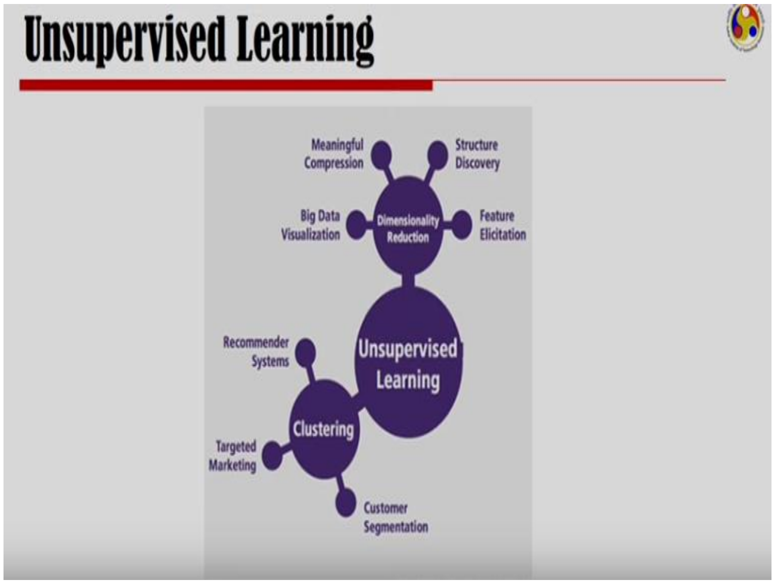

Unsupervised learning is either seen as clustering or dimensionality reduction.

(Refer Slide Time: 34:58)

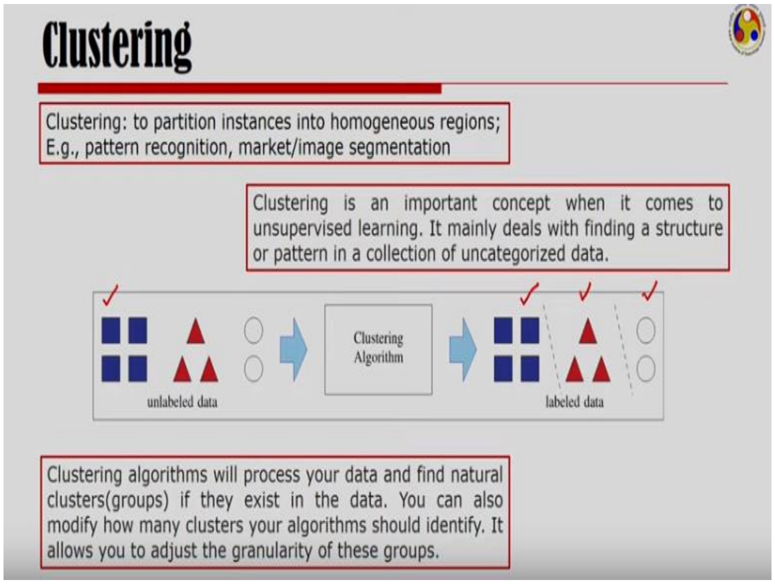

Let us look at each of them one-by-one. So clustering is about getting to partition instances of homogeneous reasons. So we would be thinking of getting to a real segment of an image or getting to a pattern. So clustering is an important concept when it comes to unsupervised learning, its about finding a structure or pattern in a collection of uncategorized data. Here is the graphical representation of the clustering algorithm.

So unlabeled data comes to it and the algorithm somehow is able to draw boundaries and find out that I have three clusters one of rectangles, one of triangles and one of circles. So clustering algorithms will process the data and find the natural clusters if they exist in the data and can modify the number of clusters you need to identify, and clustering algorithms also allow one to adjust the granularity of these groups. We will see clustering in more detail in our discussion in the lecture on unsupervised learning.

(Refer Slide Time: 36:20)

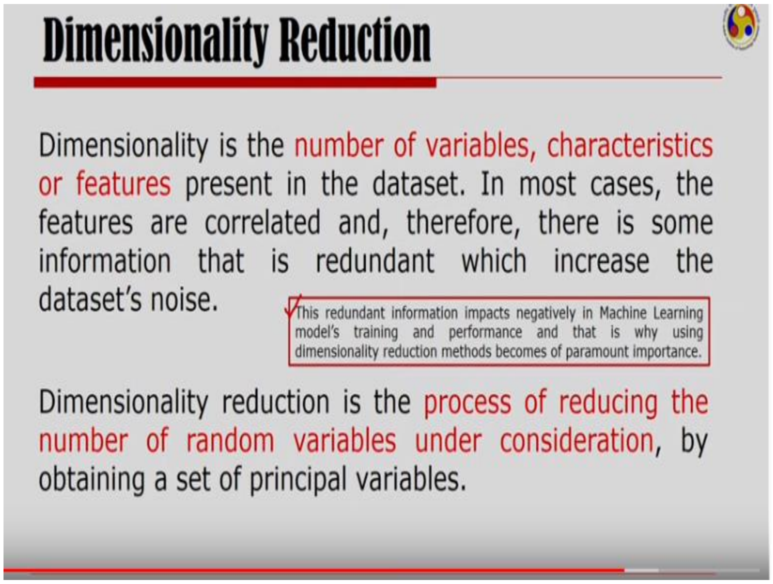

Dimensionality on the other hand is the number of variables characteristics or features present in the data set. In most cases the features are correlated and therefore there is some information that is redundant which increases the datasets noise. Now one needs to realize that this redundant information actually impacts negatively in the machine learning model, training and performance and that is why dimensionality reduction methods are of paramount importance.

What we mean by dimensionality reduction is the process of reducing the number of random variables under consideration by obtaining a set of principal variables in the problem under consideration.

(Refer Slide Time: 37:10)

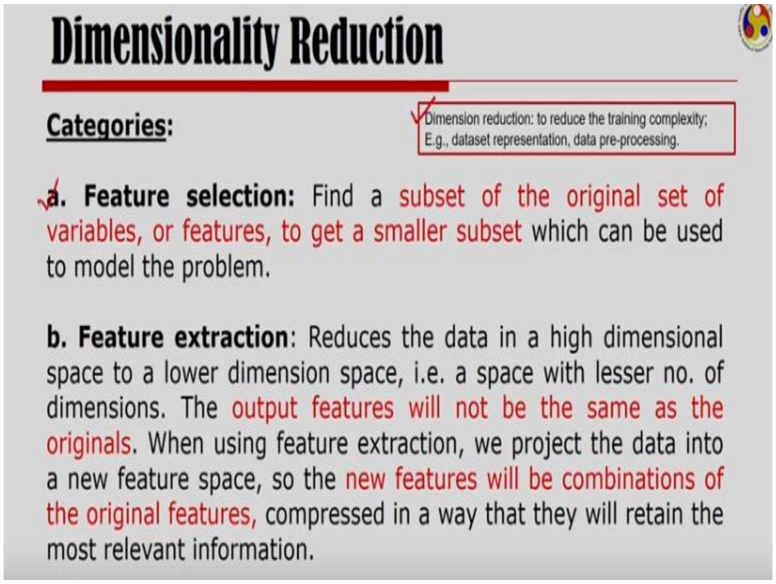

So dimension reduction is about reducing the training complexity for example the data representation and dimensionality reduction is either in terms of feature selection where we find a subset of the original set of variables or features to get a smaller subset which can then be used to model the problem or feature extraction which is about reducing the data in a high dimensional space to a lower dimensional space and interestingly the output features will not be the same as the originals.

What we do in feature extraction is we project the data into a new feature space. So that the new feature space will be combination of the original features but somehow compress in a way that they will retain the most relevant information and would not have redundant information to the extent possible.

(Refer Slide Time: 38:17)

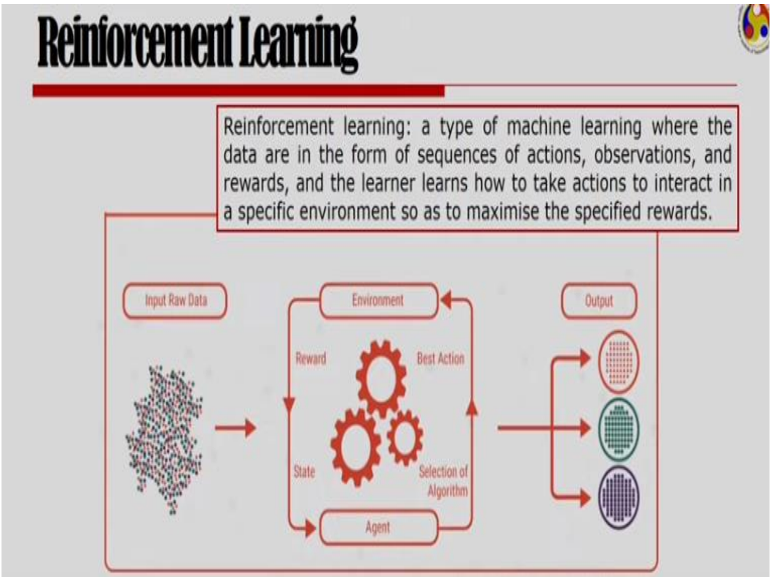

Now we come to the third form of machine learning which is reinforcement learning. Reinforcement learning is the type of machine learning where the data are in the form of sequence of actions, observations and rewards and the learner learns how to take actions to interact in a specific environment, so as to maximize the specified rewards. So you learn action to maximize payoff. So in this case you have as I was trying to explain an agent and given the best action on the environment by the agent you reward the agent and then slowly the actions are reinforced.

(Refer Slide Time: 39:09)

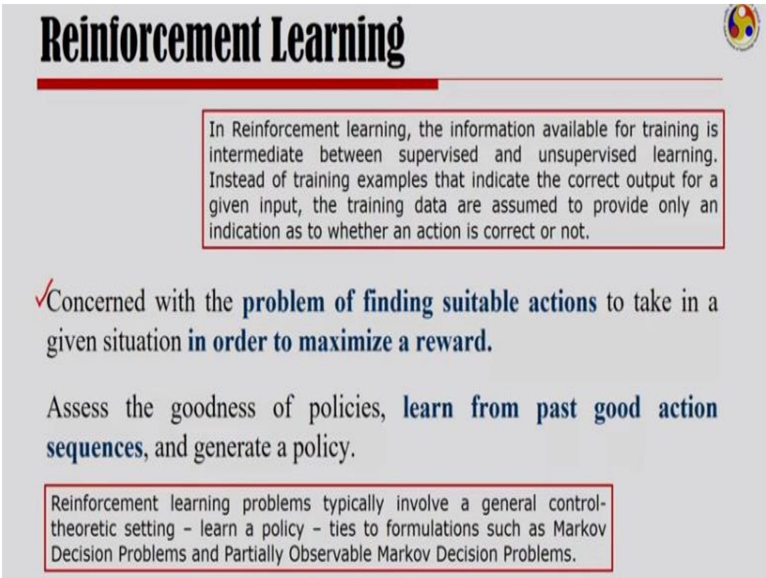

So in reinforcement learning the information available for training is actually intermediate between supervised and unsupervised learning. So instead of training examples that indicate the correct output for a given input. The training data are here assumed to provide only an indication as to whether an action is correct or not. So we are concerned with the problem of finding suitable actions to take in a given situation.

So as to maximize the reward and we access the goodness of policies that is we learn from past good action sequences and generate a policy. So now you can see that this idea of machine learning is actually about a general control theoretic setting of learning a policy and ties to formulations that we have discussed while we were discussing planning and decision-making. So reinforcement learning has ties to formulations such as the Markov decision problem and the partially observable Markov decision problems and it is about finding optimal policy and therefore has ties to value iteration and policy iteration algorithms that we have discussed before.

(Refer Slide Time: 40:47)

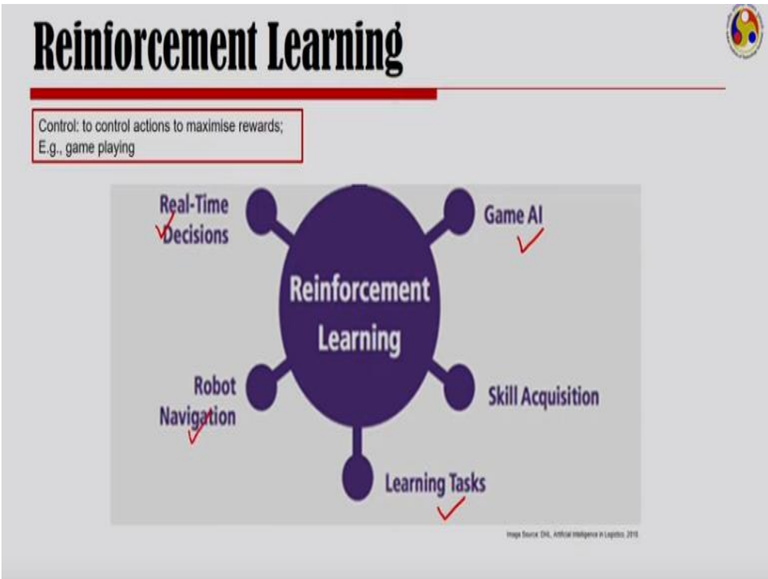

Now reinforcement learning is used in a number of areas like for real-time decisions, for gaming in AI, for learning tasks, for robot navigation, for human machine interaction and in all of this the idea is to maximize rewards for a given action and try to learn the policy.

(Refer Slide Time: 41:21)

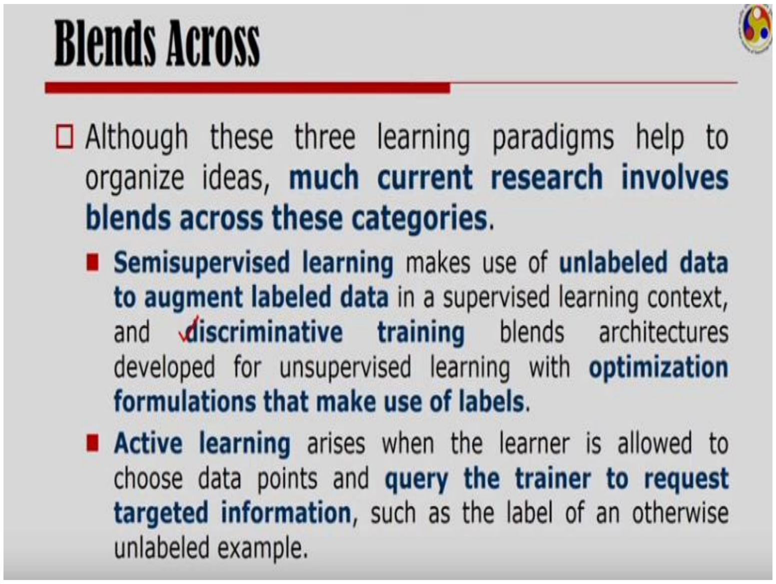

Now although these 3 learning paradigms help to organize ideas however much current research involves blends across these categories. Therefore, we have semi-supervised learning that makes use of unlabeled data to augment label data in a supervised learning context and then we have what is called discriminative training. That is, it blends architectures developed for unsupervised learning with optimization formulations that make use of labels.

There are other ideas such as active learning which is about the learner being allowed to choose data points and query the trainer to request targeted information such as the level of an otherwise unlabeled example.

(Refer Slide Time: 42:18)

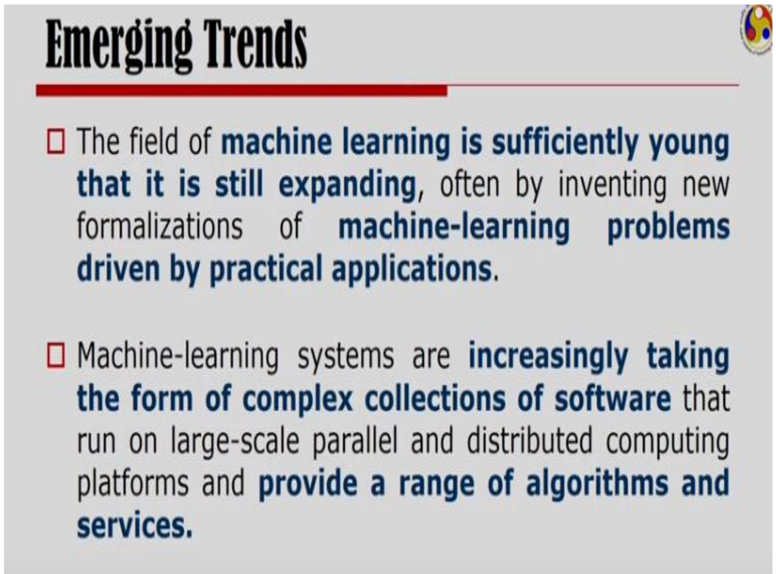

The field of machine learning is sufficiently young, and it is still expanding often by inventing new formalizations of machine learning problems driven by practical applications. Machine learning systems are increasingly taking the form of complex collections of software that run on large-scale parallel and distributed computing platforms and thus machine learning is driving the Internet of Things and provide a range of algorithms and services.

In the course of this module we will look at supervised learning unsupervised learning and reinforcement learning and also look at a very interesting area of machine learning called deep learning. Now in this portion of the module I shall follow three different books from what has been notified earlier.

(Refer Slide Time: 43:24)

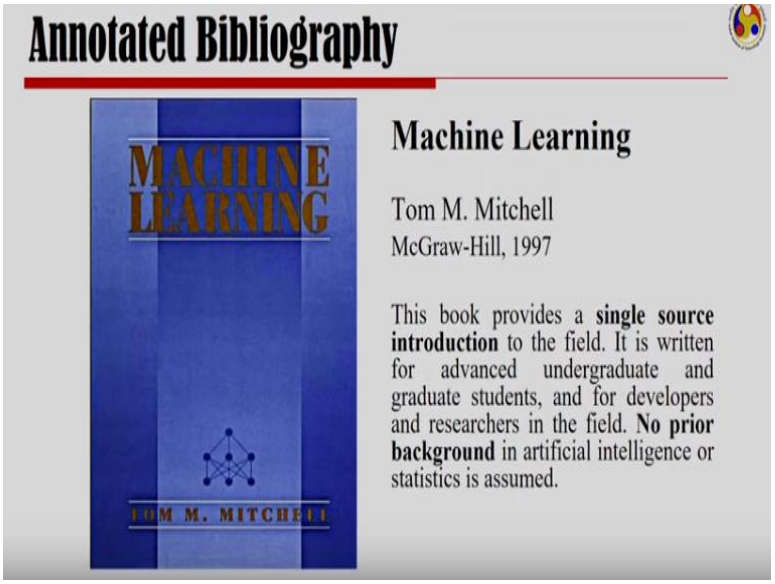

So the first book would be by Tom Mitchell and this book provides a single source introduction to the field it is written for advanced undergraduates and graduate students and the best part is this book does not assume any prior background in AI or statistics.

(Refer Slide Time: 43:43)

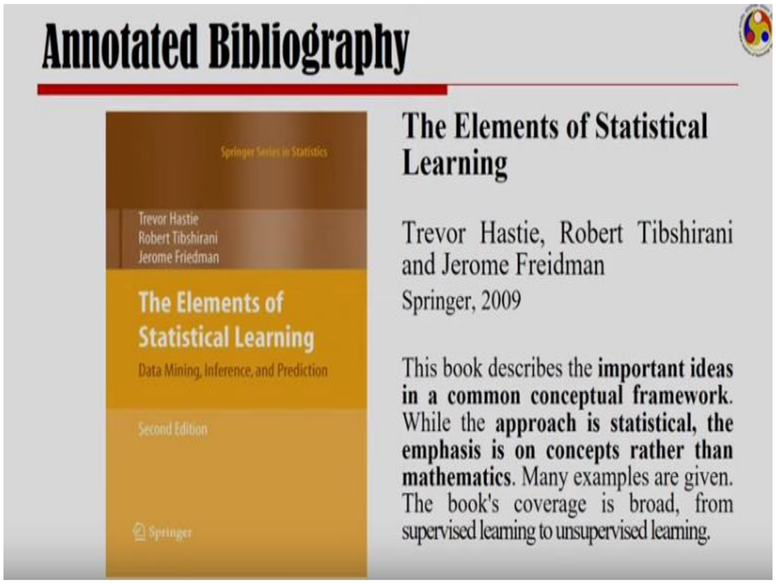

The next book would be a book that describes important ideas in a common conceptual framework and the book is called the elements of statistical learning by Hastie and others. So this book the approach is statistical the emphasis is on concepts rather than mathematics that is why I have chosen this book.

(Refer Slide Time: 44:11)

And the third book that we will refer to is the book called machine learning, the new AI by Alpaydin this is a concise overview of machine learning. Computer programs that learn from data which underlies the application that include recommendation systems, face recognition and driverless cars. Thank you very much.