International Journal on Recent Researches in Science, Engineering & Technology (IJRRSET)
A Journal Established in early 2000 as National journal and upgraded to International journal in 2013 and is in existence for the last 10 years. It is run by Retired Professors from NIT, Trichy. Journal Indexed in JIR, DIIF and SJIF.
Available online at: www.jrrset.com

ISSN (Print) : 2347-6729
ISSN (Online) : 2348-3105

JIR IF : 2.54
SJIF IF : 4.334
Cosmos: 5.395

Volume 8, Issue 5 - May 2020 - Pages 19 - 28

Machine Learning and its Types of Learning

[1]Shashidhar T Halakatti [2]Sangamesh S K [3]Pavitra M Gadhar
Assistant Professor
RTE Society's Rural Engineering College Hulkoti, Karnataka-582101, India

Abstract : The main idea of this paper is to create attentiveness among upcoming research scholars about the latest technology, specifically in the domain of machine learning which finds vide area of applications in data analytics and artificial intelligence.

Key Words: Machine Learning, Supervised Learning, Unsupervised Learning.

I. INTRODUCTION

The field of machine learning is concerned with the question of how to construct computer programs that automatically improve with experience. In recent years many successful machine learning applications have been developed, ranging from data-mining programs that learn to detect fraudulent credit card transactions, to information-filtering systems that learn users' reading preferences, to autonomous vehicles that learn to drive on public highways. At the same time, there have been important advances in the theory and algorithms that form the foundations of this field.

A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.

For example, a computer program that learns to play checkers might improve its performance as measured by its ability to win at the class of tasks involving playing checkers games, through experience obtained by playing games against itself. In general, to have a well-defined learning problem, we must identity these three features: the class of tasks, the measure of performance to be improved, and the source of experience.

A checkers learning problem:

Task T: playing checkers

0 Performance measure P: percent of games won against opponents

Training experience E: playing practice games against itself.

Machine learning is a large field of study that overlaps with and inherits ideas from many related fields such as artificial intelligence. The focus of the field is learning, that is, acquiring skills or knowledge from experience. Most commonly, this means synthesizing useful concepts from historical data. As such, there are many different types of learning that

International Journal on Recent Researches in Science, Engineering & Technology (IJRRSET)
A Journal Established in early 2000 as National journal and upgraded to International journal in 2013 and is in existence for the last 10 years. It is run by Retired Professors from NIT, Trichy. Journal Indexed in JIR, DIIF and SJIF.
Available online at: www.jrrset.com
ISSN (Print) : 2347-6729
ISSN (Online) : 2348-3105
JIR IF : 2.54
SJIF IF : 4.334
Cosmos: 5.395
Volume 8, Issue 5 - May 2020 - Pages 19 - 28

you may encounter as a practitioner in the field of machine learning: from whole fields of study to specific techniques.

II. TYPES OF LEARNING

Given that the focus of the field of machine learning is "learning," there are many types that you may encounter as a practitioner. Some types of learning describe whole subfields of study comprised of many different types of algorithms such as "supervised learning." Others describe powerful techniques that you can use on your projects, such as "transfer learning."

There are perhaps 14 types of learning that you must be familiar with as a machine learning practitioner; they are:

Learning Problems
1. Supervised Learning
2. Unsupervised Learning
3. Reinforcement Learning

Hybrid Learning Problems
1. Semi-Supervised Learning
2. Self-Supervised Learning
3. Multi-Instance Learning

Statistical Inference
1. Inductive Learning
2. Deductive Inference
3. Transductive Learning

Learning Techniques
1. Multi-Task Learning
2. Active Learning
3. Online Learning
4. Transfer Learning
5. Ensemble Learning

III. LEARNING PROBLEMS

1. Supervised Learning

Supervised learning describes a class of problem that involves using a model to learn a mapping between input examples and the target variable. Models are fit on training data comprised of inputs and outputs and used to make predictions on test sets where only the inputs are provided and the outputs from the model are compared to the withheld target variables and used to estimate the skill of the model.There are two main types of supervised learning problems: they are classification that involves predicting a class label and regression that involves predicting a numerical value.

Classification: Supervised learning problem that involves predicting a class label.

Regression: Supervised learning problem that involves predicting a numerical label.

International Journal on Recent Researches in Science, Engineering & Technology (IJRRSET)
A Journal Established in early 2000 as National journal and upgraded to International journal in 2013 and is in existence for the last 10 years. It is run by Retired Professors from NIT, Trichy. Journal Indexed in JIR, DIIF and SJIF.
Available online at: www.jrrset.com
ISSN (Print) : 2347-6729
ISSN (Online) : 2348-3105
JIR IF : 2.54
SJIF IF : 4.334
Cosmos: 5.395
Volume 8, Issue 5 - May 2020 - Pages 19 - 28

Both classification and regression problems may have one or more input variables and input variables may be any data type, such as numerical or categorical.

An example of a classification problem would be the MNIST handwritten digits dataset where the inputs are images of handwritten digits (pixel data) and the output is a class label for what digit the image represents (numbers 0 to 9). An example of a regression problem would be the Boston house prices dataset where the inputs are variables that describe a neighborhood and the output is a house price in dollars. Some machine learning algorithms are described as "supervised" machine learning algorithms as they are designed for supervised machine learning problems. Popular examples include: decision trees, support vector machines, and many more.

Algorithms are referred to as "supervised" because they learn by making predictions given examples of input data, and the models are supervised and corrected via an algorithm to better predict the expected target outputs in the training dataset. Some algorithms may be specifically designed for classification (such as logistic regression) or regression (such as linear regression) and some may be used for both types of problems with minor modifications (such as artificial neural networks).

2. Unsupervised Learning

Unsupervised learning describes a class of problems that involves using a model to describe or extract relationships in data.

Compared to supervised learning, unsupervised learning operates upon only the input data without outputs or target variables. As such, unsupervised learning does not have a teacher correcting the model, as in the case of supervised learning.

There are many types of unsupervised learning, although there are two main problems that are often encountered by a practitioner: they are clustering that involves finding groups in the data and density estimation that involves summarizing the distribution of data.

Clustering: Unsupervised learning problem that involves finding groups in data.

Density Estimation: Unsupervised learning problem that involves summarizing the distribution of data.

An example of a clustering algorithm is k-Means where k refers to the number of clusters to discover in the data. An example of a density estimation algorithm is Kernel Density Estimation that involves using small groups of closely related data samples to estimate the distribution for new points in the problem space.

Clustering and density estimation may be performed to learn about the patterns in the data.

Additional unsupervised methods may also be used, such as visualization that involves graphing or plotting data in different ways and projection methods that involves reducing the dimensionality of the data.

International Journal on Recent Researches in Science, Engineering & Technology (IJRRSET)
A Journal Established in early 2000 as National journal and upgraded to International journal in 2013 and is in existence for the last 10 years. It is run by Retired Professors from NIT, Trichy. Journal Indexed in JIR, DIIF and SJIF.
Available online at: www.jrrset.com
ISSN (Print) : 2347-6729
ISSN (Online) : 2348-3105
JIR IF : 2.54
SJIF IF : 4.334
Cosmos: 5.395
Volume 8, Issue 5 - May 2020 - Pages 19 - 28

Visualization: Unsupervised learning problem that involves creating plots of data.

Projection: Unsupervised learning problem that involves creating lower-dimensional representations of data.

An example of a visualization technique would be a scatter plot matrix that creates one scatter plot of each pair of variables in the dataset. An example of a projection method would be Principal Component Analysis that involves summarizing a dataset in terms of eigenvalues and eigenvectors, with linear dependencies removed.

3. Reinforcement Learning

Reinforcement learning describes a class of problems where an agent operates in an environment and must learn to operate using feedback.

The use of an environment means that there is no fixed training dataset, rather a goal or set of goals that an agent is required to achieve, actions they may perform, and feedback about performance toward the goal.

It is similar to supervised learning in that the model has some response from which to learn, although the feedback may be delayed and statistically noisy, making it challenging for the agent or model to connect cause and effect.

An example of a reinforcement problem is playing a game where the agent has the goal of getting a high score and can make moves in the game and received feedback in terms of punishments or rewards.

Impressive recent results include the use of reinforcement in Google's AlphaGo in out-performing the world's top Go player.

Some popular examples of reinforcement learning algorithms include Q-learning, temporal-difference learning, and deep reinforcement learning.

IV. HYBRID LEARNING PROBLEMS

The lines between unsupervised and supervised learning is blurry, and there are many hybrid approaches that draw from each field of study.

In this section, we will take a closer look at some of the more common hybrid fields of study: semi-supervised, self-supervised, and multi-instance learning.

1. Semi-Supervised Learning

Semi-supervised learning is supervised learning where the training data contains very few labeled examples and a large number of unlabeled examples.

The goal of a semi-supervised learning model is to make effective use of all of the available data, not just the labelled data like in supervised learning.

International Journal on Recent Researches in Science, Engineering & Technology (IJRRSET)
A Journal Established in early 2000 as National journal and upgraded to International journal in 2013 and is in existence for the last 10 years. It is run by Retired Professors from NIT, Trichy. Journal Indexed in JIR, DIIF and SJIF.
Available online at: www.jrrset.com
ISSN (Print) : 2347-6729
ISSN (Online) : 2348-3105
JIR IF : 2.54
SJIF IF : 4.334
Cosmos: 5.395
Volume 8, Issue 5 - May 2020 - Pages 19 - 28

Making effective use of unlabelled data may require the use of or inspiration from unsupervised methods such as clustering and density estimation. Once groups or patterns are discovered, supervised methods or ideas from supervised learning may be used to label the unlabeled examples or apply labels to unlabeled representations later used for prediction.

It is common for many real-world supervised learning problems to be examples of semi-supervised learning problems given the expense or computational cost for labeling examples. For example, classifying photographs requires a dataset of photographs that have already been labeled by human operators.

Many problems from the fields of computer vision (image data), natural language processing (text data), and automatic speech recognition (audio data) fall into this category and cannot be easily addressed using standard supervised learning methods.

2. Self-Supervised Learning

Self-supervised learning refers to an unsupervised learning problem that is framed as a supervised learning problem in order to apply supervised learning algorithms to solve it.

Supervised learning algorithms are used to solve an alternate or pretext task, the result of which is a model or representation that can be used in the solution of the original (actual) modeling problem.

A common example of self-supervised learning is computer vision where a corpus of unlabeled images is available and can be used to train a supervised model, such as making images grayscale and having a model predict a color representation (colorization) or removing blocks of the image and have a model predict the missing parts (inpainting).

A general example of self-supervised learning algorithms are autoencoders. These are a type of neural network that is used to create a compact or compressed representation of an input sample. They achieve this via a model that has an encoder and a decoder element separated by a bottleneck that represents the internal compact representation of the input.

These autoencoder models are trained by providing the input to the model as both input and the target output, requiring that the model reproduce the input by first encoding it to a compressed representation then decoding it back to the original. Once trained, the decoder is discarded and the encoder is used as needed to create compact representations of input.

Although autoencoders are trained using a supervised learning method, they solve an unsupervised learning problem, namely, they are a type of projection method for reducing the dimensionality of input data.

Another example of self-supervised learning is generative adversarial networks, or GANs. These are generative models that are most commonly used for creating synthetic photographs using only a collection of unlabeled examples from the target domain.

International Journal on Recent Researches in Science, Engineering & Technology (IJRRSET)
A Journal Established in early 2000 as National journal and upgraded to International journal in 2013 and is in existence for the last 10 years. It is run by Retired Professors from NIT, Trichy. Journal Indexed in JIR, DIIF and SJIF.
Available online at: www.jrrset.com
ISSN (Print) : 2347-6729
ISSN (Online) : 2348-3105
JIR IF : 2.54
SJIF IF : 4.334
Cosmos: 5.395
Volume 8, Issue 5 - May 2020 - Pages 19 - 28

GAN models are trained indirectly via a separate discriminator model that classifies examples of photos from the domain as real or fake (generated), the result of which is fed back to update the GAN model and encourage it to generate more realistic photos on the next iteration.

3. Multi-Instance Learning

Multi-instance learning is a supervised learning problem where individual examples are unlabeled; instead, bags or groups of samples are labeled.

Instances are in "bags" rather than sets because a given instance may be present one or more times, e.g. duplicates.

Modeling involves using knowledge that one or some of the instances in a bag are associated with a target label, and to predict the label for new bags in the future given their composition of multiple unlabeled examples.

imple methods, such as assigning class labels to individual instances and using standard supervised learning algorithms, often work as a good first step.

V. Statistical Inference

Inference refers to reaching an outcome or decision.

In machine learning, fitting a model and making a prediction are both types of inference.

There are different paradigms for inference that may be used as a framework for understanding how some machine learning algorithms work or how some learning problems may be approached.

Some examples of approaches to learning are inductive, deductive, and transductive learning and inference.

1. Inductive Learning

Inductive learning involves using evidence to determine the outcome.

Inductive reasoning refers to using specific cases to determine general outcomes, e.g. specific to general.

Most machine learning models learn using a type of inductive inference or inductive reasoning where general rules (the model) are learned from specific historical examples (the data).

Fitting a machine learning model is a process of induction. The model is a generalization of the specific examples in the training dataset.

A model or hypothesis is made about the problem using the training data, and it is believed to hold over new unseen data later when the model is used.

International Journal on Recent Researches in Science, Engineering & Technology (IJRRSET)
A Journal Established in early 2000 as National journal and upgraded to International journal in 2013 and is in existence for the last 10 years. It is run by Retired Professors from NIT, Trichy. Journal Indexed in JIR, DIIF and SJIF.
Available online at: www.jrrset.com
ISSN (Print) : 2347-6729
ISSN (Online) : 2348-3105
JIR IF : 2.54
SJIF IF : 4.334
Cosmos: 5.395
Volume 8, Issue 5 - May 2020 - Pages 19 - 28

2. Deductive Inference

Deduction or deductive inference refers to using general rules to determine specific outcomes.

We can better understand induction by contrasting it with deduction.

Deduction is the reverse of induction. If induction is going from the specific to the general, deduction is going from the general to the specific.

Deduction is a top-down type of reasoning that seeks for all premises to be met before determining the conclusion, whereas induction is a bottom-up type of reasoning that uses available data as evidence for an outcome.

In the context of machine learning, once we use induction to fit a model on a training dataset, the model can be used to make predictions. The use of the model is a type of deduction or deductive inference.

3. Transductive Learning

Transduction or transductive learning is used in the field of statistical learning theory to refer to predicting specific examples given specific examples from a domain.

It is different from induction that involves learning general rules from specific examples, e.g. specific to specific.

Unlike induction, no generalization is required; instead, specific examples are used directly. This may, in fact, be a simpler problem than induction to solve.

A classical example of a transductive algorithm is the k-Nearest Neighbors algorithm that does not model the training data, but instead uses it directly each time a prediction is required.

Contrasting Induction, Deduction, and Transduction:

We can contrast these three types of inference in the context of machine learning.

For example:

Induction: Learning a general model from specific examples.

Deduction: Using a model to make predictions.

Transduction: Using specific examples to make predictions.

VI. Learning Techniques

1. Multi-Task Learning

Multi-task learning is a type of supervised learning that involves fitting a model on one dataset that addresses multiple related problems.

International Journal on Recent Researches in Science, Engineering & Technology (IJRRSET)
A Journal Established in early 2000 as National journal and upgraded to International journal in 2013 and is in existence for the last 10 years. It is run by Retired Professors from NIT, Trichy. Journal Indexed in JIR, DIIF and SJIF.
Available online at: www.jrrset.com
ISSN (Print) : 2347-6729
ISSN (Online) : 2348-3105
JIR IF : 2.54
SJIF IF : 4.334
Cosmos: 5.395
Volume 8, Issue 5 - May 2020 - Pages 19 - 28

It involves devising a model that can be trained on multiple related tasks in such a way that the performance of the model is improved by training across the tasks as compared to being trained on any single task.

Multi-task learning can be a useful approach to problem-solving when there is an abundance of input data labeled for one task that can be shared with another task with much less labeled data.

For example, it is common for a multi-task learning problem to involve the same input patterns that may be used for multiple different outputs or supervised learning problems. In this setup, each output may be predicted by a different part of the model, allowing the core of the model to generalize across each task for the same inputs.

A popular example of multi-task learning is where the same word embedding is used to learn a distributed representation of words in text that is then shared across multiple different natural language processing supervised learning tasks.

2. Active Learning

Active learning is a technique where the model is able to query a human user operator during the learning process in order to resolve ambiguity during the learning process.

Active learning is a type of supervised learning and seeks to achieve the same or better performance of so-called "passive" supervised learning, although by being more efficient about what data is collected or used by the model.

It is not unreasonable to view active learning as an approach to solving semi-supervised learning problems, or an alternative paradigm for the same types of problems.

Active learning is a useful approach when there is not much data available and new data is expensive to collect or label.

The active learning process allows the sampling of the domain to be directed in a way that minimizes the number of samples and maximizes the effectiveness of the model.

3. Online Learning

Online learning involves using the data available and updating the model directly before a prediction is required or after the last observation was made.

Online learning is appropriate for those problems where observations are provided over time and where the probability distribution of observations is expected to also change over time. Therefore, the model is expected to change just as frequently in order to capture and harness those changes.

This approach is also used by algorithms where there may be more observations than can reasonably fit into memory, therefore, learning is performed incrementally over observations, such as a stream of data.

International Journal on Recent Researches in Science, Engineering & Technology (IJRRSET)
A Journal Established in early 2000 as National journal and upgraded to International journal in 2013 and is in existence for the last 10 years. It is run by Retired Professors from NIT, Trichy. Journal Indexed in JIR, DIIF and SJIF.
Available online at: www.jrrset.com
ISSN (Print) : 2347-6729
ISSN (Online) : 2348-3105
JIR IF : 2.54
SJIF IF : 4.334
Cosmos: 5.395
Volume 8, Issue 5 - May 2020 - Pages 19 - 28

Generally, online learning seeks to minimize “regret,” which is how well the model performed compared to how well it might have performed if all the available information was available as a batch.

One example of online learning is so-called stochastic or online gradient descent used to fit an artificial neural network.

4. Transfer Learning

Transfer learning is a type of learning where a model is first trained on one task, then some or all of the model is used as the starting point for a related task.

It is a useful approach on problems where there is a task related to the main task of interest and the related task has a large amount of data.

It is different from multi-task learning as the tasks are learned sequentially in transfer learning, whereas multi-task learning seeks good performance on all considered tasks by a single model at the same time in parallel.

An example is image classification, where a predictive model, such as an artificial neural network, can be trained on a large corpus of general images, and the weights of the model can be used as a starting point when training on a smaller more specific dataset, such as dogs and cats. The features already learned by the model on the broader task, such as extracting lines and patterns, will be helpful on the new related task.

As noted, transfer learning is particularly useful with models that are incrementally trained and an existing model can be used as a starting point for continued training, such as deep learning networks.

5. Ensemble Learning

Ensemble learning is an approach where two or more modes are fit on the same data and the predictions from each model are combined.

The objective of ensemble learning is to achieve better performance with the ensemble of models as compared to any individual model. This involves both deciding how to create models used in the ensemble and how to best combine the predictions from the ensemble members.

Ensemble learning is a useful approach for improving the predictive skill on a problem domain and to reduce the variance of stochastic learning algorithms, such as artificial neural networks.

Some examples of popular ensemble learning algorithms include: weighted average, stacked generalization (stacking), and bootstrap aggregation (bagging).

International Journal on Recent Researches in Science, Engineering & Technology (IJRRSET)
A Journal Established in early 2000 as National journal and upgraded to International journal in 2013 and is in existence for the last 10 years. It is run by Retired Professors from NIT, Trichy. Journal Indexed in JIR, DIIF and SJIF.
Available online at: www.jrrset.com
ISSN (Print) : 2347-6729
ISSN (Online) : 2348-3105
JIR IF : 2.54
SJIF IF : 4.334
Cosmos: 5.395
Volume 8, Issue 5 - May 2020 - Pages 19 - 28

References:

1. Pattern Recognition and Machine Learning, 2006.
2. Reinforcement Learning: An Introduction, 2nd edition, 2018.
3. Machine Learning: A Probabilistic Perspective, 2012.
4. Machine Learning, 1997.
5. The Nature of Statistical Learning Theory, 1995.
6. Foundations of Machine Learning, 2nd edition, 2018.
7.

Papers

8. Revisiting Self-Supervised Visual Representation Learning, 2019.
9. Active Learning Literature Survey, 2009