Online Machine Learning in Big Data Streams

András A. Benczúr and Levente Kocsis and Róbert Pálovics

Abstract The area of online machine learning in big data streams covers algorithms that are (1) distributed and (2) work from data streams with only a limited possibility to store past data. The first requirement mostly concerns software architectures and efficient algorithms. The second one also imposes nontrivial theoretical restrictions on the modeling methods: In the data stream model, older data is no longer available to revise earlier suboptimal modeling decisions as the fresh data arrives. In this article, we provide an overview of distributed software architectures and libraries as well as machine learning models for online learning. We highlight the most important ideas for classification, regression, recommendation, and unsupervised modeling from streaming data, and we show how they are implemented in various distributed data stream processing systems. This article is a reference material and not a survey. We do not attempt to be comprehensive in describing all existing methods and solutions; rather, we give pointers to the most important resources in the field. All related sub-fields, online algorithms, online learning, and distributed data processing are hugely dominant in current research and development with conceptually new research results and software components emerging at the time of writing. In this article, we refer to several survey results, both for distributed data processing and for online machine learning. Compared to past surveys, our article is different because we discuss recommender systems in extended detail.

1 Introduction

Big data analytics promise to deliver valuable business insights. However, this is difficult to realize using today's state-of-the-art technologies, given the flood of data generated from various sources. A few years ago, the term **fast data** [146] arose to

All authors

Institute for Computer Science and Control, Hungarian Academy of Sciences (MTA SZTAKI), e-mail: {benczur, kocsis, rpalovics}@sztaki.mta.hu

capture the idea that **streams** of data are generated at very high rates and that these need to be analyzed quickly in order to arrive at actionable intelligence.

Fast data can flood from network measurements, call records, web page visits, sensor readings, and so on [29]. The fact that such data arrives continuously in multiple, rapid, time-varying, possibly unpredictable, and unbounded streams appears to yield some fundamentally new research problems. Examples of such applications include financial applications [241], network monitoring [18, 1], security, sensor networks [66], Twitter analysis [25], and more [83].

Traditional data processing assumes that data is available for multiple access, even if in some cases it resides on disk and can only be processed in larger chunks. In this case, we say that the data is **at rest**, and we can perform **batch processing**. Database systems, for example, store large collections of data and allow users to initiate queries and transactions.

Fast data, or **data in motion** is closely connected to and in certain cases used as a synonym of the **data stream** computational model [172]. In this model, data arrives continuously in a potentially infinite stream that has to be processed by a resource-constrained system. The main restriction is that the main memory is small and can contain only a small portion of the stream, hence most of the data has to be immediately discarded after processing.

In one of the earliest papers that describe a system for data stream processing [1], the needs of monitoring applications are described. The tasks relevant for monitoring applications differ from conventional processing of data at rest in that the software system must process and react to continuous input from multiple sources. The authors introduce the data active, human passive model, in which the system permanently processes data to provide alerts for humans.

Needs and opportunities for **machine learning over fast data streams** are stimulated by a rapidly growing number of industrial, transactional, sensor and other applications [243]. The concept of online machine learning is summarized in one of the earliest overviews of the field [228]. The principal task is to learn a concept incrementally by processing labeled training examples one at a time. After each data instance, we can update the model, after which the instance is discarded.

In the data stream computational model, only a small portion of the data can be kept available for immediate analysis [110]. This has both algorithmic and statistical consequences for machine learning: Suboptimal decisions on earlier parts of the data may be difficult to unwind, and if needed, require low memory sampling and summarization procedures. For data streaming applications, incremental or online learning fits best.

> **Requirement 1**: Online learning updates its model after each data instance without access to all past data, hence the constraints of the data streaming computational model apply.

Data streaming is not just a technical restriction on machine learning: Fast data is not just about processing power but also about fast semantics. Large databases available for mining today have been gathered over months or years, and the underlying processes generating them have changed during this time, sometimes radically [119]. In data analysis tasks, fundamental properties of the data may change

quickly, which makes gradual manual model adjustment procedures inefficient and even infeasible [243]. Traditional, batch learners build static models from finite, static, identically distributed data sets. By contrast, stream learners need to build models that evolve over time. Processing will strongly depend on the order of examples generated from a continuous, non-stationary flow of data. Modeling is hence affected by potential concept drifts or changes in distribution [98].

Requirement 2: Adaptive machine learning models are needed to handle concept drift.

As an additional consequence of Requirement 2, adaptive learning also affects the way evaluation is performed. Potentially in every unit of time, the system may return predictions from different models, and we may receive too few predictions from a particular model to evaluate by traditional metrics. Instead, we have to define error measures that we can minimize in a feedback system: Predictions are made for a stream of objects one by one, and the correct answer is received immediately afterwards. A discrepancy between the prediction and the observed value serves as feedback, which may immediately trigger modifications to the model [228].

Finally, the third important aspect for online learning from big data is algorithmic. In order to cope with the volume of the data, processing has to be distributed. Clusters of machines are hard to manage, and hardware failure must be mitigated in the case of an application running on thousands of servers. Map-Reduce [68] was the first programming abstraction designed to manage the cluster, provide fault tolerance, and ease software development and debugging. While Map-Reduce is designed for batch processing, distributed data stream processing needs other solutions, such as communication between processing elements [174] via an interconnection topology [219]. For an outlook, mostly batch distributed data mining solutions are surveyed in [78].

Requirement 3: Online learning from big data has to be implemented in a distributed stream processing architecture.

Several surveys for online machine learning in general [91, 29, 206, 83, 92] and subfields [228, 220, 57, 165, 193, 131, 243, 7, 207, 99] have appeared recently. Our survey is different, first of all, in that we focus on three aspects: the data stream computational model, the adaptive methods for handling concept drift, and the distributed software architecture solutions for streaming. We elaborate on systems for machine learning by distributed data stream processing. In particular, we explore the idea of using Parameter Servers. We focus on nonstatic environments and give no convergence theorems for performance on identically distributed streams.

To the best of our knowledge, this is the first survey on online machine learning with an extensive discussion of recommender systems. Recommenders are important as they give a clear, industry-relevant example of Requirement 2. Note that in [243] it is observed that adaptive learning models are still rarely deployed in industry.

This paper is organized as follows: In Section 2 we give an overview of distributed software architectures for online machine learning. Then in Sections 3–6

we list some of the most important online learning models for supervised classification, reinforcement learning, recommendation, unsupervised analysis, and concept drift mitigation.

2 Processing Data Streams for Machine Learning

Data-intensive applications often work in the data stream computational model [17], in which the data is transient: Some or all of the input data is not available for random access from disk or memory. The data elements in the stream arrive online and can be read at most once; in case of a failure, it is possible that the data elements cannot be read at all. The system has no control over the order in which data elements arrive to be processed either within a data stream or across data streams.

The strongest constraint for processing data streams is the fact that once an element from a data stream has been processed, it has to be discarded or archived. Only selected past data elements can be accessed by storing them in memory, which is typically small relative to the size of the data streams. Many of the usual data processing operations would need random access to the data [17]: For example, only a subset of SQL queries can be served from the data stream. As surveyed in [172], data stream algorithms can tackle this constraint by a variety of strategies, including adaptive sampling in sliding windows, selecting representative distinct elements, and summarizing data in low-memory data structures, also known as sketches or synopses.

When designing online machine learning algorithms, we have to take several algorithmic and statistical considerations into account. The first problem we face is the restrictions of the computational model. As we cannot store all the input, we cannot unwind a decision made on past data. For example, we can use statistical tests to choose from competing hypotheses [73], which give theoretical guarantees in the case of identically distributed data.

A second problem with real stream learning tasks is that data typically changes over time, for example, due to concept drifts [228]. For changing distributions, we can even use the streaming computational model to our advantage: By permanent training, we can adapt to concept drift by overwriting model parameters based on insights from fresh data and thus forgetting the old distribution [89]. Possible means of concept drift adaptation, however, will depend on the task, the data, and the choice of the algorithm [228, 137].

In this section we give a brief overview of how different, mostly open source, software projects manage machine learning tasks over streaming data. We note that very active, ongoing research in the field may make parts of our description obsolete very quickly. First, in Section 2.1 we provide a summary of the most important distributed data stream processing engines with active development at the time of writing. Next, in Section 2.2 we give a taxonomy of the possible machine learning solutions with respect to support from streaming data and distributed processing. Since operation on shared-nothing distributed architectures is a key requirement for

big data processing solutions, we discuss the main machine learning parallelization strategies in Section 2.3. One particularly popular solution, the parameter server is also described in more detail. Finally, in Section 2.4 we list functionalities of stream learning libraries as of the time of writing.

2.1 Data Stream Processing Engines

Distributed stream learning libraries are usually either part of **data stream processing engines (DSPEs)** or built on top of them through interfaces. For a better understanding of the software architecture, we give an overview of DSPEs next.

Since the emergence of early systems such as Aurora [1, 15], several DSPEs have been developed. For a survey of DSPEs, see [194]. The main focus of the most recent, evolving DSPEs is to provide simplicity, scalability, stateful processing, and fault tolerance with fast recovery. In terms of simplicity, the most important goal is to overcome the need to integrate the DSPE with a batch processing engine as in the so-called lambda architecture described, among other places, in [167, 133].

In this section we focus on Apache Spark [236] and Apache Flink [43], since at the time of writing these DSPEs have the most active development for learning from streams. Other systems usually provide machine learning functionalities by interfacing with SparkML, a Spark-based machine learning library, or SAMOA [170], a stream learning library designed to work as a layer on top of general DSPEs. We list the functionality of these systems in Section 2.4.

Without attempting to be exhaustive, we list some other main DSPEs. Note that this field is very active and several systems with large impact on both DSPEs and stream learning have already stopped development. For example, several active projects have borrowed concepts from the **S4 project** [174], which retired in 2014. Proprietary systems are summarized in [103]; it appears to be the case that commercial developers as of yet have no special focus on data stream processing, hence our main goal is to give an overview of the open source solutions.

Storm [219] relies on the concept of the connection topology of the processing elements. The topology is allowed to contain cycles; however in case of a failure with cycles, neither the exactly-once nor the at-least-once processing condition can be enforced. **Samza** [175] provides a design for fast recovery after failures independent of the state size. As a new advantage, stateful streaming systems have the possibility to emulate batch operation, hence the need for the lambda architecture can be eliminated. **Beam** [12], based on the Google Cloud dataflow model [46, 11, 12], focuses on event time windows to process out-of-order data. Beam can be connected to the deep learning framework Tensorflow [67] for batch machine learning. Finally, while most of the above systems are distributed, we mention Esper [22] as a prominent single machine, in-memory DSPE.

One DSPE in the focus of this article, **Spark** [236] treats stream processing as a sequence of small batch computations. Records in the stream are collected into micro-batches (by time), and a short-lived batch job is scheduled to process each

micro-batch. The advantage of this approach is its intuitive transition from batch to streaming with straightforward fault tolerance and interaction with batch programs. The main disadvantage is higher latency due to the inherent scheduling overhead for the micro-batch.

Another DSPE of our choice, **Flink** [43] provides consistent managed state with exactly-once guarantees, while achieving high throughput and low latency, serving both batch and streaming tasks. Flink handles the streams event by event in a true streaming fashion through the underlying streaming (dataflow) runtime model. While this provides more fine-grained access to the stream, it does not come with a throughput overhead due to various runtime optimizations such as buffering of output records. The advantages are very low processing latency and a natural stateful computational model. The disadvantages are that fault tolerance and load balancing are more challenging to implement. Flink is primarily for stream processing. Flink batch tasks can be expressed by using loops in stream processing, as we will see in the next section.

2.2 Taxonomy of Machine Learning Tools
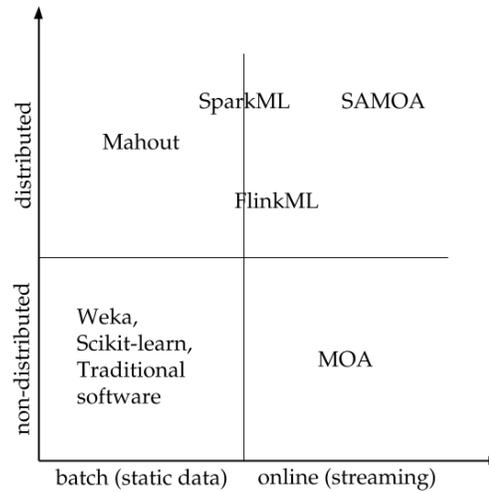
Fig. 1 Taxonomy of machine learning tools.

In this article, we survey online machine learning tools for big data. Our main focus are models, architectures, and software libraries that process streaming data over distributed, shared-nothing architectures. As shown in Fig. 1, the two key distinguishing features of machine learning tools are whether they are distributed and whether they are based on static or streaming data.

As the least restrictive of the four quadrants in Fig. 1, **batch non-distributed** tools can implement any method from the other quadrants. If scale permits, data streams can first be stored and then analyzed at rest, and distributed processing steps can be unfolded to run sequentially on a single-processor system. Traditional machine learning tools, for example, R, Weka, and scikit-learn, fall into this quadrant.

Distributed batch machine learning systems [78] typically implement algorithms by using the Map-Reduce principle [68]. Perhaps the best-known system is Mahout [180] built on top of Hadoop [227], but a very rich variety of solutions exists in this field. GraphLab [164], withdrawn from the market in 2016, was another mostly batch tool that has also influenced online systems.

Online learning solutions cover algorithms that immediately build models after seeing a relatively small portion of the data. By this requirement, we face the difficulty of not necessarily being able to undo a suboptimal decision made in an earlier stage, based on data that is no longer available for the algorithm.

The first library for online machine learning, MOA [27] collects a variety of models suitable for training online, most of which we describe in Sections 3 and 6.1. Based on MOA concepts, SAMOA [170] is a distributed framework—a special purpose DSPE—and library that provides distributed implementation for most MOA algorithms. Online learning recommender systems, both distributed and non-distributed, are also described in [185]. Another recent non-distributed online learning tool is described in [31].

For a **distributed online learning** software architecture, the underlying system needs to be a DSPE. In addition to SAMOA, which can be considered a DSPE itself, most DSPEs of the previous section can be used for distributed online learning. For example, Flink, Samza, and Storm implement interfaces to use SAMOA libraries. On the other hand, a DSPE can also implement batch machine learning algorithms: For example, the machine learning library SparkML is mostly batch and FlinkML is partly batch.

Combined batch and online machine learning solutions are of high practical relevance. We can train our models batch based on a precompiled sample, and then apply prediction for a live data stream. DSPE solutions are progressing in this area. One recent result, Clipper is a low-latency online prediction system [63] with a model abstraction layer that makes it easy to serve pre-trained models based on a large variety of machine learning frameworks.

Another batch approach to learning from time-changing data is to repeatedly apply a traditional learner to a sliding window of examples: As new examples arrive, they are inserted into the beginning of the window. Next, a corresponding number of examples are removed from the end of the window, and the learner is reapplied, as in early research on concept drift in learning from continuous data [228, 119]. Finally, as a combination with online learning methods, the batch trained model can be incrementally updated by a streaming algorithm [89].

2.3 Parallel Learning and the Parameter Server

In order to design distributed modeling algorithms, we have to elaborate on parallelization strategies. **Horizontal** or **data parallel** systems partition the training data and build separate models on subsets that may eventually get merged. Such a solution is applied, for example, for training XGBoost trees [53]. This approach, however, will typically depend on the partitioning and lead to heuristic or approximate solutions.

Vertical parallel or **model parallel training** solutions partition across attributes in the same data point, rather than partitioning the training data coming from the stream. Each training point is split into its constituting attributes, and each attribute is sent to a different processing element. For example, for linear models trained by gradient descent, coefficients can be stored and updated by accessing a distributed store [157]. As another example, the fitness of attributes for a split in a decision tree construction can be computed in parallel [170]. Further examples such as Tensorflow [67], Petuum [230], and MXNet [54] are also capable of model parallel training. The drawback is that in order to achieve good performance, there must be sufficient inherent parallelism in the modeling approach.
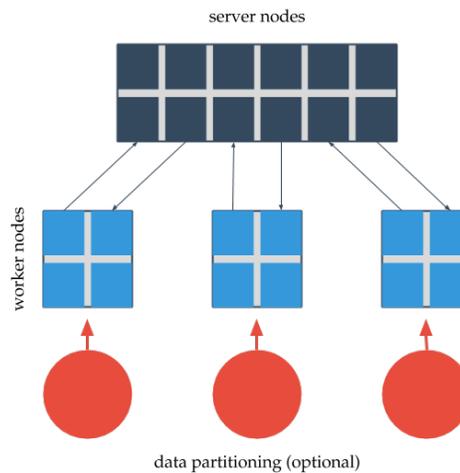
Fig. 2 The parameter server architecture for distributed machine learning.

The **parameter server** introduced in [209] is a popular way to implement model parallel training, with several variants developed [115, 159, 157, 158], including an application for deep learning [67]. The main idea is to simplify the development of distributed machine learning applications by allowing access to shared parameters as key–value pairs. As shown in Fig. 2, the compute elements are split into two subsets: Parameters are distributed across a group of **server nodes**, while data processing and computation are performed at **worker nodes**. Any node can both

push out its local parameters and pull parameters from remote nodes. Parallel tasks are typically asynchronous, but the algorithm designer can be flexible in choosing a consistency model.

Various parameter server systems are summarized best in [157]. The first generation [209] uses distributed key-value stores such as memcached [81] to store the parameters. More recent solutions [67, 159] implement specific parameter access interfaces. For example, as the first step towards a general platform, [115] design a table-based interface for stale synchronous operation, which they use to implement a wide variety of applications.

Most results [209, 115, 157] describe parameter servers for variants of regression as well as unsupervised modeling by Latent Dirichlet Allocation (see Section 6.1). Another popular use is classification by deep neural networks [67, 159]. An implementation for Multiple Additive Regression Trees is given in [240].

Recommendation algorithms can also be implemented by the parameter server principle. Batch implementations of asynchronous distributed stochastic gradient descent recommender algorithms are given in [115, 203]. A recent comparison of distributed recommenders, including an online one based on a Flink parameter server, is given in [185].

The parameter server idea fits the data stream as well as the batch computational models. For example, if it suffices to read the data only once and in input order for gradient descent, a batch parameter server such as in [157] immediately yields a streaming algorithm. Yet, most applications are for batch modeling only [209, 115, 157, 203, 240]. As examples of parameter servers for online learning, applications for reinforcement learning can be found in [173] and for recommenders in [185].

2.4 Stream Learning Libraries

We provide an overview of the present machine learning functionalities of the main open source engines. Note that all these engines are under active development with a very large number of components already available as research prototypes or pull requests. For this reason, we only want to give a representative overview of the main components.

2.4.1 MOA and SAMOA

SAMOA [170], a distributed online learning library, is based on the concepts of MOA [27], a single machine library. SAMOA provides model parallel implementations of specific algorithms by using the concepts of Storm [219]: Processing elements are connected in a loop-free topology, which connects the various pieces of user code. SAMOA can run on top of a DSPE with a flexible interface connection. SAMOA connectors are implemented for all active DSPEs in Section 2.1.

For classification, SAMOA [65, 170] provides the Vertical Hoeffding Tree (VHT), a distributed version of a streaming decision tree [73] (Section 3.3). For clustering, it includes an algorithm based on CluStream [9] (Section 6.1). The library also includes meta-algorithms such as bagging and boosting (Section 3.4).

2.4.2 Apache Spark

Spark has a very rich set of batch machine learning functionalities, including linear, tree, Support Vector Machine, and neural network models for classification and regression, ensemble methods as well as explicit and implicit alternating least squares for recommender systems, and many more listed at https://spark.apache.org/mllib/. However, the only streaming algorithm in Spark MLLib is a linear model with ongoing work regarding online recommender and Latent Dirichlet Allocation prototypes. Spark has no SAMOA connector yet.

Spark has several parameter server implementations, most of which are batch only. We mention two projects with recent activity. Glint is a parameter server based Latent Dirichlet Allocation (Section 6.1) implementation, which is described in [123]. Angel is a general parameter server [126] that has implementations for logistic regression, SVM (Section 3.2), matrix factorization (Section 5.3), Latent Dirichlet Allocation, and more. Angel supports synchronous, asynchronous, and stale synchronous processing [125].

2.4.3 Apache Flink

Flink has a loosely integrated set of machine learning components, most of which are collected at https://github.com/FlinkML. In addition to the SAMOA connector at https://github.com/apache/incubator-samoa/tree/master/samoa-flink, Flink has a true streaming parameter server implementation at https://github.com/FlinkML/flink-parameter-server, which includes a Passive Aggressive classifier (Section 3.2) and gradient descent matrix factorization (Section 5.3). Finally, Flink can serve online predictions by models trained on any system supporting the PMML standard [102], using the JPMML library at https://github.com/FlinkML/flink-jpmml.

In Flink, the parameter server is implemented as part of the data stream API at https://github.com/FlinkML/flink-parameter-server. Since the communication between workers and servers is two-way, the implementation involves loops in stream processing. As mentioned in Section 2.1, exactly-once processing and fault tolerance is conceptually difficult, and implementation is not yet complete as of the time of writing [42].

3 Classification and Regression

The goal of classification and regression is to predict the unknown label of a data instance based on its attributes or variables. The label is discrete for classification, and continuous for regression. Classification and regression in batch settings are well established, with several textbooks available in data mining [187] and machine learning [87].

For classification and regression over streaming data, the first two of the requirements in the Introduction play a key role. By the first requirement, data is transient, since it does not fit in main memory, hence algorithms that iterate over the entire data multiple times are ruled out. By the second requirement, we cannot assume that the examples are independent, identically distributed, and generated from a stationary distribution, hence different predictions by different models have to be applied and evaluated at different times. Special modeling tools are required to meet the two challenges, and known evaluation and comparison methods are not convenient yet [97].

An important, and perhaps the oldest, application of online learning is single trial EEG classification, in which the system learns from the online feedback of the experimental subjects [176]. These early experiments differ from the approach in this section in that the performance was only measured at the end of the experiments, and no systematic analysis of the classifier performance in time was conducted.

This section is organized as follows: First, we describe the difficulties and possibilities of evaluating online learning methods in Section 3.1. We cover the most important online classification and regression methods in the next subsections. We discuss the main linear models in Section 3.2, tree-based methods in Section 3.3, classifier ensembles in Section 3.4, Bayes models in Section 3.5, and finally, neural networks in Section 3.6. Other methods such as **nearest neighbor** [149] are known as well. Our list of online classification methods is not comprehensive; for an extended list, see the recent survey [91].

3.1 Evaluation

In an infinite data stream, data available for training and testing is potentially infinite. Hence holdout methods for selecting an independent test set seem viable at first glance, and are used in a large number of online machine learning research results. However, for online evaluation, we have no control over the order in which the data is processed, and the order is not independent of the data distribution. Since the distribution generating examples and the decision models evolve over time, cross-validation and other sampling strategies are not applicable [97].

Most studies of online learning determine overall loss as a sum of losses experienced by individual training examples. Based on this so-called predictive sequential, abbreviated as **prequential** method [64], we define **online training and evaluation** in the following steps:

1. Based on the next unlabeled instance in the stream, we cast a prediction.
2. As soon as the true label of this instance becomes available, we assess the prediction error.
3. We update the model with the most recently observed error calculated by comparing the predicted and the true labels before proceeding with the next item of the data stream.

Prequential error is known to be a pessimistic estimator, since it may be strongly influenced by the initial part of the sequence, when only a few examples have been processed, and the model quality is low. The effect of the beginning of the stream can be mitigated, for example, by forgetting mechanisms [97, 98].

A further issue in online evaluation is that in reality, labels may arrive with delay [243]. The majority of adaptive learning algorithms require true labels to be available immediately after casting the prediction. If true labels are delayed, we have to join two streams with time delay, one for the variables and one for the labels, which can make the implementation of the prequential evaluation scheme computationally challenging.

Error metrics that can be defined for individual data points can be applied for prequential evaluation. For example, the definition of mean squared error, a popular metric for regression is

$$MSE = \frac{1}{N} \sum_{i=1}^{N} (\hat{y}_i - y_i)^2, \tag{1}$$

where y_i is the actual and $\hat{y}_i$ is the predicted class label for data point i, and N is the current size of the data stream. Accuracy and error rate can be computed by a similar averaging formula.

For certain common metrics such as precision, recall, and true and false positive rates, the definition will involve the changing size of the set of positive, negative, or all instances, which makes the metrics difficult to interpret in a very long stream of inhomogeneous data. Although one definition of ROC AUC [82] is based on true and false positive rates, its online interpretation is described in [238].

3.2 Linear Models

Linear models in online machine learning date back to the perceptron algorithm [200]. The **perceptron learning** algorithm learns label y of d-dimensional input vector $\mathbf{x}$ as a linear combination of the coordinates,

$$\hat{y} = \mathbf{w} \cdot \mathbf{x}. \tag{2}$$

The prediction mechanism is based on an n-dimensional hyperplane of direction $\mathbf{w}$, which divides the instance space into two half-spaces. The margin of an example, $y \cdot \mathbf{w} \cdot \mathbf{x}$, is a signed value proportional to the distance between the instance and the

hyperplane. If the margin is positive, the instance is correctly classified; otherwise, it is incorrectly classified.

The perceptron can be trained by gradient descent for hinge loss as target function. If labels y take the values ± 1 and prediction $\hat{y}$ is defined by equation (2), hinge loss is equal to

$$\ell(\mathbf{w};(\mathbf{x},y)) = \begin{cases} 0 & \text{if } y \cdot \hat{y} \geq 1; \\ 1 - y \cdot \hat{y} & \text{otherwise,} \end{cases} \tag{3}$$

from which the gradient can be computed as follows: If prediction $\hat{y}$ has the correct sign, and its absolute value is at least 1, that is, hinge loss is 0, then there is no change; the algorithm is passive. Otherwise, the gradient for $\mathbf{w}$ is $-\mathbf{x} \cdot \hat{y}$, and the update rule for learning rate η is

$$\mathbf{w} \leftarrow \mathbf{w} + \eta \cdot \mathbf{x} \cdot \hat{y} \quad \text{if } y \cdot \hat{y} < 1. \tag{4}$$

The online gradient descent algorithm simply applies the above step to training examples in order [147]. By contrast, the batch gradient descent algorithm reads the input multiple times and usually repeatedly optimizes coefficients for the same instance. Batch gradient descent can be emulated by an online algorithm: We go through training examples one by one in an online fashion and repeat multiple times over the training data [45].

Based on the general idea of perceptron learning, several online linear models have been proposed; for a detailed overview, see [83]. The **Passive Aggressive (PA) classifier** [60] is a popular online linear model that works well in practice, for example, applied as the Gmail spam filter [2]. The main goal of the PA algorithm is to improve the convergence properties of the perceptron algorithm, the simplest numerical optimization procedure. Several improved online optimization procedures were proposed prior to PA: Kivinen and Warmuth [135] give an overview of numerous earlier additive and multiplicative online algorithms, all of which are solvable by gradient descent [160, 100, 62, 134]. Based on the experiments in [62, 60], the Passive Aggressive algorithm outperforms most of the earlier online linear classification methods. Perceptron learning and most of its successors apply both to classification and regression; in other words, the range of the actual label y can be both binary and continuous.

The PA algorithm solves a constrained optimization problem: We would like the new classifier to remain as close as possible to the current one while achieving at least a unit margin on the most recent example. We use the Euclidean distance of the classifiers, $||\mathbf{w} - \mathbf{w}'||^2$. PA optimizes for hinge loss, with the goal to keep the distance minimal and the value of the margin at least 1, that is, to solve the optimization problem

$$\mathbf{w} \leftarrow \text{argmin}_{\mathbf{w}'} ||\mathbf{w} - \mathbf{w}'||^2 \quad \text{s.t.} \quad \ell(\mathbf{w};(\mathbf{x},y)) = 0. \tag{5}$$

The algorithm is passive if loss is 0, and otherwise aggressive, since it enforces zero loss. As shown in [60], the solution of the optimization problem yields the update rule

$$\mathbf{w} \leftarrow \mathbf{w} + \ell(\mathbf{w}; (\mathbf{x}, y)) \cdot y \cdot \mathbf{x} / ||\mathbf{x}||^2. \tag{6}$$

In [60], variants of PA are described that introduce a slack variable $\xi \geq 0$ in equation (5) such that the margin must stay below ξ. Adding constant times ξ or ξ^2 to the minimization target in equation (5) leads to similar optimization problems. In the same paper, multi-class, cost-sensitive, and regression variants are described as well.

Although the class of linear predictors may seem restrictive, the pioneering work of Vapnik [222] and colleagues demonstrates that by using kernels one can employ highly nonlinear predictors as well. The **Support Vector Machine (SVM)** learns an optimal separating hyperplane $\mathbf{w}^*$ in a certain high-dimensional mapped space defined by the mapping $\varphi(\mathbf{x})$ over training vectors $\mathbf{x}$. Like with perceptron learning, the prediction has the form $\hat{y} = \mathbf{w}^* \cdot \varphi(\mathbf{x})$; however, $\mathbf{w}^*$ may be of very high dimensionality.

Potential problems of the very high dimensional mapped space are eliminated by the **kernel trick**. In the optimization procedure, $\mathbf{w}^*$ turns out to be the combination of the so-called support vectors, the subset SV of training instances $\mathbf{x}_i$ for $i \in SV$ that maximize the margin. We can obtain the parameters

$$\mathbf{w}^* = \sum_{i \in SV} \alpha_i y_i \varphi(\mathbf{x}_i), \tag{7}$$

where y_i are the labels and α_i can be found by maximizing the margin. By equation (7), we can reduce the prediction to computing inner products in the mapped space:

$$\mathbf{w}^* \cdot \varphi(\mathbf{x}) = \sum_{i \in SV} \alpha_i y_i \varphi(\mathbf{x}_i) \cdot \varphi(\mathbf{x}), \tag{8}$$

The main goal of online SVM is to maintain a set of support vectors and corresponding multipliers within the limits of available memory. Whenever an online SVM learner decides to add new support vectors, others need to be discarded first, and α_i need to be updated.

To decide which support vector to discard, the definition of the span and the S-span of the support vectors defined by Vapnik [221] can be used. The span of a support vector is the minimum distance of the vector from a certain set defined by all others. The S-span of a set of support vectors is the maximum span among them.

An online SVM method to update the set of support vectors from the data stream is proposed in [6], namely, maintaining a set of support vectors and multipliers that fit into the memory limit. For a misclassified new instance, the algorithm measures the S-span of all possible ways of replacing one old support vector with the new instance, and selects the best one. Multipliers are updated by the incremental learning procedure of [44].

Several other, similar online SVM optimization methods are known, for example [62, 61, 33, 34]. Online learning algorithms were also suggested as fast alternatives to SVM in [85], based on the online algorithm for calculating the maximum margin linear model of [86]. Online gradient descent optimizers for linear models often work with kernels as well [134].

3.3 Decision and Regression Trees

In a decision or regression tree, each internal node corresponds to a test on an attribute, while leaves contain predictors for classification or regression. To build a tree, decision tree induction algorithms iterate through each attribute and compute an information theoretic function such as entropy or Gini index for classification and variance for regression [187].

Online tree induction algorithms face the difficulty that the recursive tree construction steps cannot read past data. After a split decision is made, batch algorithms partition all data into child nodes and compute the required information theoretic function of the attributes separately in each child node. Online algorithms cannot partition past data; instead, they take advantage of the potentially infinite data and use fresh instances only in the newly created nodes.

Another problem with online tree construction is that a split decision has to be made at a certain point in time, without seeing future data. A popular solution is to use the Hoeffding criterion for a statistical guarantee that the selected split is optimal for future data as well. In the so-called Hoeffding Tree or Very Fast Decision Tree (VFDT) [73], attribute information theoretic functions are maintained over the stream. If the Hoeffding criterion is met, a split decision is made, and the attribute statistics over the new child nodes are computed based on the new data from the stream. Similar methods for regression trees are described in [13, 120]. For online vertical parallel distributed Hoeffding Trees, see [141].

One problem in decision tree construction on streaming data is the cost of maintaining attribute statistics for many-valued attributes. For such attributes, a low granularity histogram has to be maintained. In [127], a method is described that partitions the range of a numerical attribute into intervals and uses statistical tests to prune these intervals.

Another difficulty is caused by nonstationary data, since each new child node is processed based on a new portion of data from the stream. Decision trees over evolving streams are considered in [26, 24, 31]. We give an overview of general methods for nonstationary data in Section 6.2.

3.4 Ensemble Methods

Ensemble methods build multiple, potentially different models on potentially different subsets of instances and attributes [187]. The simplest example is **bagging** where several base models are trained based on samples with replacement. Sampling with replacement can be simulated online [181, 28]. A special ensemble technique, the online random forest algorithm is described in [70].

A highly successful ensemble technique is **boosting**, in which we generate a sequence of base models by incorporating the prediction error of the previous models when constructing the next. For example, in AdaBoost, an algorithm designed for online learning in its first description [84], the next classifier is trained by weight-

ing instances based on the function of the error of previous classifiers. In gradient boosting [53], the next model is trained on the residual, that is, the difference of the training label and the continuous predicted label of the previous classifiers.

For online boosting, the difference compared to batch boosting is that the performance of a base model cannot be observed on the entire training set, only on the examples seen earlier. Like with online decision tree induction, decisions need to be taken based only on part of the training data. Various online boosting algorithms are described in [181, 52, 23], based on the ideas of parallel boosting via approximation [79, 182, 150]. For gradient boosted trees [53], a recent, most successful classification method, the online version is described in [223].

3.5 Bayes Models

Bayesian networks were one of the earliest applications for online learning [88], with the first methods mostly using mini-batch updates [38]. Bayesian learning methods maintain conditional probability tables $P(\mathbf{x}|y)$ by counting, where $\mathbf{x}$ is a feature vector and y is its label. Considering the conditional probability tables and the class distribution as priors, the predicted class of an instance will be the one that maximizes the posterior probability computed by the Bayes rule [187].

The simplest, Naive Bayes model makes the "naive" assumption that each input variable is conditionally independent given the class label [75]. While this model works surprisingly well [74], a weaker assumption is needed in Bayesian networks [189], in which we represent each variable with a node in a directed acyclic graph. Rather than assuming independence naively, we assume that each variable is conditionally independent given its parents in the graph.

For online learning, it is easy to update the conditional probabilities both for Naive Bayes and for Bayesian networks [88], provided that they fit into internal memory. Methods for updating the network structure online are described, for example, in [88, 51].

3.6 Neural Networks

Neural networks and deep learning have shown great promise in many practical applications ranging from speech recognition [114] and visual object recognition [144] to text processing [59]. One of the earliest applications of online trained neural networks is EEG classification [108].

Gradient descent is perhaps the most commonly used optimization procedure for training neural networks [151], which naturally leads to online learning algorithms [129] as well. The traditional formulation of gradient descent is impractical for very large neural networks. A scalable online distributed version, Downpour

SGD [67], uses asynchronous parameter updates in conjunction with a parameter server (Section 2.3).

4 Reinforcement Learning

Reinforcement learning is an area of machine learning concerned with agents taking actions in an environment with the aim of maximizing some cumulative reward. It is different from supervised learning in that the environment does not provide a target behavior, only rewards depending on the actions taken.

The environment is typically assumed to be a Markov Decision Process (MDP). Formally, we assume a set of states, S, a set of actions, A, a transition probability function $P(s,a,s')$ denoting the probability of reaching state s' after taking action a in state s, and a reward function $R(s,a)$ denoting the immediate reward after taking action a in state s.

While there is a wide range of reinforcement learning algorithms (see, e.g., [214]), we focus here on algorithms that fit the streaming model and (possibly) deal with nonstationary environments. The streaming model of reinforcement learning is constrained not only by a continuous flow of input data, but also by a continuous requisite to take actions.

4.1 Algorithms for Stationary Environments

Most reinforcement learning algorithms estimate the value of feasible actions and build a policy based on that value (e.g., by choosing the actions with the highest estimates with some additional exploration). An alternative to value prediction methods are policy gradient methods that update a parameterized policy depending on the performance.

4.1.1 Value Prediction

The value of a state is the expected cumulative reward starting from a given state and following a particular policy. In a similar way, the action value is the expected reward starting from a given state with a particular action.

Value prediction methods estimate the value of the state or the value of the actions in particular states. In the former case, to build a policy from the estimated values, an additional transition model is needed as well. Such a model is provided for some domains (e.g., by the rules of a game), but in many cases the transition model needs to be learned as well. Action-values can be used directly for constructing a policy without the need for a model.

Temporal-difference (TD) learning learns the state value estimate $V(s)$ by the following update rule after each state transition (S_t, S_{t+1}):

$$V(S_t) \leftarrow (1-\alpha)V(S_t) + \alpha(R_t + \gamma V(S_{t+1})),$$

where α is a step-size, and γ is the discount factor. TD learning was used in one of the first breakthroughs for reinforcement learning, that is Tesauro's backgammon program [218].

The best-known action-value prediction algorithm is Q-learning [226]. For each occurrence of a transition (S_t, A_t, S_{t+1}), the algorithm updates the action-value $Q(S_t, A_t)$ by

$$Q(S_t, A_t) \leftarrow (1-\alpha)Q(S_t, A_t) + \alpha(R_t + \gamma \max_a Q(S_{t+1}, a)).$$

Q-learning using deep neural network to approximate the action-values has been successfully applied to playing some Atari games at human expert level [169].

Another algorithm that learns action-values is Sarsa [213]. For each sequence $S_t, A_t, S_{t+1}, A_{t+1}$, the algorithm updates its estimates by

$$Q(S_t, A_t) \leftarrow (1-\alpha)Q(S_t, A_t) + \alpha(R_t + \gamma Q(S_{t+1}, A_{t+1})).$$

Sarsa was successfully used by [121] for optimizing a DRAM memory controller.

The value prediction algorithms above were described with update rules for a tabular representation. In most cases, function approximation is used, and the update rules rely on a gradient step. Online enhancements of gradient descent as well as eligibility traces [214] can be applied to all variants.

4.1.2 Policy Gradient

While using value functions is more widespread, it is also possible to use a parameterized policy without relying on such functions. Parameterized policies are typically optimized by gradient ascent with respect to the performance of the policy.

A policy gradient algorithm, the REINFORCE algorithm [229] was used to optimize policy in a Go playing program that outperforms the best human players [208].

4.2 Algorithms for Nonstationary Environments

Most reinforcement learning algorithms, including those discussed in the previous section, assume that the environment does not change over time. While incremental algorithms such as Q-learning can adapt well to nonstationary environments, it may be necessary to devise more explicit exploration strategies that can cope with changes, for example, in reward distribution.

A special case of reinforcement learning is the multi-armed bandit problem. In this case, the agent repeatedly selects an action from K possible choices, obtaining a reward after each choice. This problem retains the notion of reward; however, there are no states and consequently no state-transitions. In the non-stochastic variant [16], the distribution of the rewards may change over time arbitrarily. Standard algorithms for this problem are Exp3 and its variants [16], which rely on an exponential selection algorithm, including some exploration terms as well. Contextual bandits extend the bandit setting with the notion of state (or context); however, state transitions are still missing. This framework was used for instance in [156] to select personalized new stories. We note that the distinguishing feature of recommendation in a bandit setting is that the user can provide feedback only on the recommended items.

5 Recommender Systems

Recommender systems [198] serve to predict user preferences regarding items such as music tracks (Spotify), movies (Netflix), products, books (Amazon), blogs, or microblogs (Twitter), as well as content on friends' and personal news feeds (Facebook).

Recommender systems can be categorized by the type of information they infer about users and items. Collaborative filtering [161, 202] builds models of past user-item interactions such as clicks, views, purchases, or ratings, while content-based filtering [163] recommends items that are similar in content, for example, share phrases in their text description. Context-aware recommenders [5] use additional information on the user and the interaction, for example user location and weather conditions. Recent events in a user session [138] serve as a special context.

A milestone in the research of recommendation algorithms, the Netflix Prize Competition [21] had high impact on research directions. The target of the contest was based on the one to five star ratings given by users, with one part of the data used for model training and the other for evaluation. As an impact of the competition, tasks now termed batch rating prediction were dominating research results.

Recommendation models rely on the feedback provided by the user, which can be **explicit**, such as one to five star movie ratings on Netflix [4]. However, most recommendation tasks are **implicit**, as the user provides no like or dislike information. Implicit feedback can be available in the form of time elapsed viewing an item or listening to a song, or in many cases, solely as a click or some other form of user interaction. In [190], the authors claim that 99% of recommendation industry tasks are implicit.

As a main difference between recommendation and classification, classifiers usually work independently of the event whose outcome they predict. Recommender systems, on the other hand, may directly influence observations: They present a ranked top list of items [71], and the user can only provide feedback for the items on the list. Moreover, real systems process data streams where users request one or a

few items at a time and get exposed to new information that may change their needs and taste when they return to the service next time. Furthermore, an online trained model may change and return completely different lists for the same user even for interactions very close in time.

By the above considerations, real recommender applications fall in the category of top item recommendation by online learning for implicit user feedback, a task that has received less attention in research so far. In this section, we show the main differences in evaluating such systems compared to both classifiers and batch systems, as well as describe the main data stream recommender algorithms.

Online recommenders seem more restricted than those that can iterate over the data set several times, and one could expect inferior quality from the online methods. By contrast, in [184, 89], surprisingly strong performance of online methods is measured.

As an early time-aware recommender system example, the item-based nearest neighbor [202] can be extended with time-decay [72]. Most of the early models, however, are time-consuming to compute, difficult to update from a data stream, and hence need periodical batch training. Probably the first result in this area, the idea of processing transactions in chronological order to incrementally train a recommendation model first appeared in [216, Section 3.5]. Streaming gradient descent matrix factorization methods were also proposed in [122, 14], who use Netflix and Movielens data and evaluate by RMSE.

The difficulty of evaluating streaming recommenders was first mentioned in [148], although the authors evaluated models by offline training and testing split. Ideas for online evaluation metrics appeared first in [183, 224, 184]. In [224], incremental algorithms are evaluated using recall. In [184], recall is shown to have undesirable properties, and other metrics for evaluating online learning recommenders are proposed.

Finally, we note that batch distributed recommender systems were surveyed in [130].

5.1 Prequential (Online) Evaluation for Recommenders

To train and evaluate a time-sensitive or online learning recommender, we can use the prequential or online evaluation framework that we defined for classifier evaluation. As seen in Fig. 3, online evaluation for a recommender system includes the following steps:

1. We query the recommender for a top-k recommendation for the active user.
2. We evaluate the list in question against the single relevant item that the user interacted with.
3. We allow the recommender to train on the revealed user-item interaction.

Since we can potentially re-train the model after every new event, the recommendation for the same user may be very different even at close points in time, as seen

user - item interactions in chronological order

top lists for user u at different times
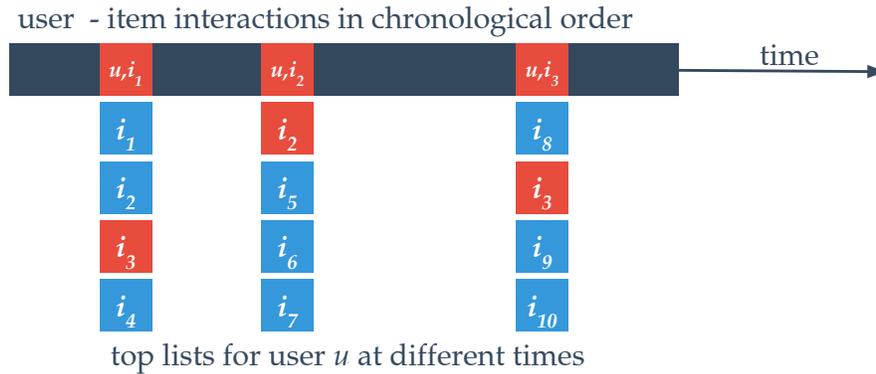
Fig. 3 Prequential evaluation of the online ranking prediction problem.

in Fig. 3. The standard recommender evaluation settings used in research cannot be applied, since there is always only a single relevant item in the ground truth.

In one of the possible recommender evaluation settings, the **rating prediction** problem, which is popular in research, we consider a user u and an item i. The actual user preference in connection with the item is expressed as a value r_{ui}, for which the system returns a prediction $\hat{r}_{ui}$. This explicit rating can be a scale such as one to five stars for a Netflix movie, while implicit rating can be the duration of viewing a Web page in seconds. Implicit rating is binary when the only information is whether the user interacted with the item (clicked, viewed, purchased) or not. Depending on whether r_{ui} is binary or scale, the same prequential metrics, such as error rate or MSE, can be applied as for classification or regression. For example, in the Netflix Prize competition, the target was the square root of MSE between the predicted and actual ratings, see Equation (1).

Another possible way to evaluate recommenders is **ranking prediction**, where performance metrics depend on the list of displayed items. We note that given rating prediction values $\hat{r}_{ui}$ for all i, in theory, ranking prediction can be solved by sorting the relevance score of all items. For certain models, heuristics to speed up the selection of the highest values of $\hat{r}_{ui}$ by candidate preselection exist [217].

To evaluate ranking prediction, we have to take into consideration two issues that do not exist for classifier evaluation. In the case of prequential evaluation, as shown in Fig. 3, the list for user u may change potentially after every interaction with u. As soon as u provides feedback for certain item i, we can change model parameters and the set of displayed items may change completely. Most of the batch ranking quality measures focus on the set of items consumed by the same user, under the assumption that the user is exposed to the same list of items throughout the evaluation. As this assumption does not hold, we need measures for individual user-item interactions.

Another issue regarding ranking prediction evaluation lies in a potential user-system interaction that affects quality scores. Typically, the set of items is very large, and users are only exposed to a relatively small subset, which is usually provided

by the system. The form of user feedback is usually a click on one or more of these items, which can be evaluated by computing the clickthrough rate. Since users cannot give feedback on items outside the list, the fair comparison of two algorithms that present different sets for the user can only be possible by relying on live user interaction. This fact is known by practitioners, who use **A/B testing** to compare the performance of different systems. In A/B testing, the live set of users is divided into groups that are exposed to the results of the different systems.

Most traditional ranking prediction metrics, to a certain level, rely on the assumption that the same user is exposed to the same list of items, and hence the interactions of the same user can be considered to be the unit for evaluation. For online evaluation, as noted in [184], the unit of evaluation will be a single interaction, which usually contains a single relevant item. Based on this modification, most batch metrics apply in online learning evaluation as well. Note that the metrics below apply not just in A/B testing, but also in experiments with frozen data, where user feedback is not necessarily available for the items returned by a given algorithm. For example, if the item consumed by the user in the frozen data is not returned by the algorithm, the observed relevance will be 0, which may not be the case if the same algorithm is applied in an A/B test. Note that attempts to evaluate research results by A/B testing have been made in the information retrieval community [20]; however, designing and implementing such experiments is cumbersome.

Next, we list several metrics for the quality of the ordered top-K list of items $L = \{i_1, i_2, ..., i_K\}$ against the items E consumed by the user. We will also explain how online evaluation metrics differ from their batch counterparts. For the discussion, we mostly follow [184].

Clickthrough rate is commonly used in the practice of recommender evaluation. It is defined as the ratio of clicks received for L:

$$\text{Clickthrough@K} = \begin{cases} 1 & \text{if } E \cap L \neq \emptyset; \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

For precision and recall, similar to Clickthrough, the actual order within L is unimportant:

$$\text{Precision@K} = \frac{|E \cap L|}{K}, \quad \text{Recall@K} = \frac{|E \cap L|}{|E|}. \tag{10}$$

For batch evaluation, E is the entire set of items with positive feedback from a given user who is exposed to the same L for each interaction. The overall batch system performance can be evaluated by averaging precision and recall over the set of users. For online evaluation, typically $|E| = 1$, where Precision@K is 0 or $1/K$ and Recall@K is 0 or 1 depending on whether the actual item in E is listed in L or not. Precision and recall are hence identical to clickthrough, up to a constant. As a consequence, the properties of online precision and recall are very different from their batch counterparts. The main reason for the difference lies in the averaging procedure of prequential evaluation: We cannot merge the events of the same user, instead, we average over the set of individual interactions.

Measures that consider the position of the relevant item i in L can give more refined performance indication. The first example is reciprocal rank:

$$RR@K = \begin{cases} 0 & \text{if } \mathrm{rank}(i) > K; \\ \dfrac{1}{\mathrm{rank}(i)} & \text{otherwise.} \end{cases} \tag{11}$$

Discounted cumulative gain (DCG) is defined similarly, as

$$\mathrm{DCG@K} = \sum_{k=1}^{K} \frac{\mathrm{rel}(i_k)}{\log_2(1+k)} \tag{12}$$

where $\mathrm{rel}(i_k)$ indicates the relevance of the i-th item in the list. For the implicit task, relevance is 1 if the user interacted with the item in the evaluation set, 0 otherwise. For batch evaluation, we can consider all interactions of the same user as one unit. If we define iDCG@K, the ideal maximum possible value of DCG@K for the given user, we can obtain nDCG@K, the normalized version of DCG@K, as

$$\mathrm{nDCG@K} = \frac{\mathrm{DCG@K}}{\mathrm{iDCG@K}}. \tag{13}$$

Note that for online learning, there is only one relevant item, hence $\mathrm{iDCG} = 1$. For emphasis, we usually use the name nDCG for batch and DCG for online evaluation.

5.2 Session-Based Recommendation

Previous items in user sessions constitute a very important context [113]. In e-commerce, the same user may return next time with a completely different intent and may want to see a product category completely different from the previous session. Algorithms that rely on recent interactions of the same user are called **session-based item-to-item** recommenders. The user session is special context, and it is the only information available for an item-to-item recommender. In fact, several practitioners [138, 190] argue that most of the recommendation tasks they face are without sufficient past user history. For example, users are often reluctant to create logins and prefer to browse anonymously. Moreover, they purchase certain types of goods (for example, expensive electronics) so rarely that their previous purchases will be insufficient to create a meaningful user profile. Whenever a long history of previous activities or purchases by the user is not available, recommenders may propose items that are similar to the most recent ones viewed in the actual user session.

Session-based recommendation can be served by very simple algorithms, most of which are inherently online. A comparison of the most important such online algorithms in terms of performance is available in [89]. Data stream processing algorithms can retain items from the most recently started sessions as long as they fit in their memory. Recommendation is based on the recent items viewed by the user in the actual shopping session. For example, we can record how often users visited item i after visiting another item j. Since fast update to transition frequencies is usually possible, the method is online.

In an even simpler algorithm that is not strictly session-based, we recommend the most popular recent items. This method can be considered batch or online depending on the granularity of the item frequency measurement update. Both algorithms can be personalized if we consider the frequency of past events involving the user. If items are arranged hierarchically (for example, music tracks by artist and genre), personal popularity and personal session data can involve the frequency of the artists or genres for recommending tracks. More session-based algorithms are described in [138].

5.3 Online Matrix Factorization
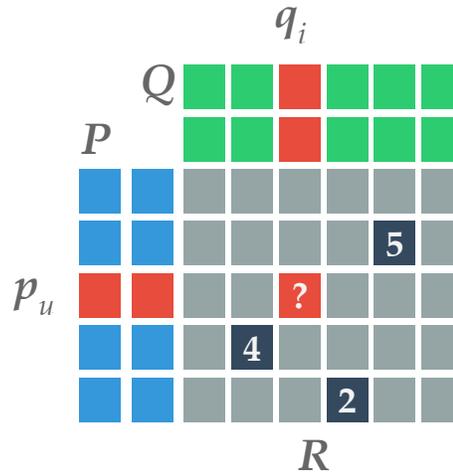
Fig. 4 Utility matrix R and the matrix factorization model built from matrices P and Q.

Most nontrivial online recommender algorithms are based on matrix factorization [140], a popular class of collaborative filtering methods. Given the user-item utility matrix R shown in Fig. 4, we model R by decomposing it into the two dense matrices P and Q. For a given user u, the corresponding row in P is user vector p_u. Similarly, for item i, the corresponding column of Q is item vector q_i. The predicted relevance of item i for user u is then

$$\hat{r}_{ui} = p_u q_i^T.$$ (14)

Note that we can extend the above model by scalar terms that describe the biased behavior of the users and the items [140].

One possibility to train model parameter matrices P and Q is by gradient descent [140, 90], which can be applied to online learning as well [184]. For a set of

interactions E, we optimize equation (1) for MSE as target function. In one step of gradient descent, we fit P and Q in equation (14) to one of the ratings in E. Unlike in batch training, where we can use the ratings several times in any order, in online learning, we have the most recent single item in E. In other words, in online gradient descent, we fit the model to the events one by one as they arrive in the data stream.

For a given (explicit or implicit) rating r_{ui}, the steps of gradient descent are as follows. First, we compute the gradient of objective function F with respect to the model parameters:

$$\frac{\partial F}{\partial p_u} = -2(r_{ui} - \hat{r}_{ui})q_i, \quad \frac{\partial F}{\partial q_i} = -2(r_{ui} - \hat{r}_{ui})p_u. \tag{15}$$

Next, we update the model parameters in opposite direction of the gradient, proportionally to learning rate η, as

$$p_u \leftarrow \eta(r_{ui} - \hat{r}_{ui})q_i,$$
$$q_i \leftarrow \eta(r_{ui} - \hat{r}_{ui})p_u.$$

Overfitting is usually avoided by adding a regularization term in the objective function [140].

In the case of implicit feedback, the known part of the utility matrix only contains elements with positive feedback. To fit a model, one requires negative feedback for training as well. Usually, such elements are selected by sampling from those that the user has not interacted with before [195]. We can also introduce confidence values for ratings and consider lower confidence for the artificial negative events [118].

Gradient descent can also be used in a mix of batch and online learning, for example, training batch models from scratch periodically, and continuing the training with online learning. We can also treat users and items differently, for example, updating user vectors more dynamically than item vectors, as first suggested by [216].

Another use of online gradient descent is to combine different recommendation models [184]. We can express the final prediction as the linear combination of the models in the ensemble whose parameters are the linear coefficients and the individual model parameters. In [184], two online gradient descent methods are described with regards to whether the derivative of the individual models is available, where all parameters can be trained through the derivative of the final model or otherwise by learning the coefficients and the individual models separately.

5.4 Variants of Matrix Factorization

Several variants of matrix factorization that can be trained by gradient descent both for batch and online learning tasks have been proposed. Bayesian Personalized Ranking [196] has top list quality as target instead of MSE. In asymmetric matrix factorization [188], we model the user by the sum of the item vectors the user rated in the past.

Recently, various factorization models have been developed that incorporate context information [113]. Context data can be modeled by introducing data tensor D instead of the rating matrix R. In a simplest case, the data includes a single piece of additional context information [197]: for example, music tracks can have artist as context.

Alternating Least Squares [140, 191] (ALS) is another optimization method for matrix factorization models, in which for a fixed Q, we compute the optimal P, then for a fixed P, the optimal Q, repeatedly until certain stopping criteria are met. Hidasi et al. [112, 111, 113] introduced several variants of ALS-based optimization schemes to incorporate context information. By incremental updating, ALS can also be used for online learning [109].

5.5 Summary of Recommendation by Online Learning

Recommendation differs from classification in that in recommendation, there are two types of objects, users, and items, and a prediction has to be made for their interaction. A practical recommender system displays a ranked list of a few items for which the user can give feedback. In an online learning system, the list shown to the same user at different times may change completely for two reasons. First, as in the prequential classifier training and evaluation setting, the list of recommendations may change because the model changes. Second, the user feedback we use for evaluation depends on the actual state of the model, since the user may have no means to express interest in an item not displayed. Hence for online learning evaluation, metrics that involve the notion of a volatile list have to be used.

Online learning is very powerful for recommender systems due to their advantage of having much more emphasis on recent events. For example, if we update models immediately for newly emerged users and items, trends are immediately detected. The power of online learning for recommendation may also be the result of updating user models with emphasis on recent events, which may be part of the current user session. User session is a highly relevant context for recommendation and most session-based methods are inherently online.

6 Additional Topics

In this final section, we give a brief overview of two additional topics, both of which are extensively covered in recent surveys. In Section 6.1, we describe unsupervised data mining methods, including clustering, frequent itemset mining, dimensionality reduction, and topic modeling. In Section 6.2, we describe the notion of the dataset drift, or in other terms, concept drift, and list the most important drift adapting methods. We only discuss representative results in these areas.

6.1 Unsupervised Data Mining

The most prominent class of unsupervised learning methods is **clustering** where instances have to be distributed into a finite set of clusters such that instances within the cluster are more similar to each other than to others in different clusters [187]. Batch clustering algorithms have been both studied and employed as data analysis tools for decades [124, 231]. One frequently applied clustering method is **k-means** [107] where cluster center selection and assignment to nearest centers are iteratively performed until convergence. Another is **DBSCAN** [77], a density-based method that groups points that are closely packed together.

Online clustering algorithms are surveyed among other places in [165, 131, 7, 207]. The majority of the most relevant methods are data stream versions of k-means or its variants such as k-medians [237, 35, 80, 177, 104, 9, 239, 96, 143, 3]. Another set of results describes the data stream implementation of DBSCAN [41, 55, 143]. Finally, an online hierarchical clustering algorithm that maintains similarity measures and hierarchically merges closest clusters is described in [199].

Finding **frequent itemsets** [10] is another central unsupervised data mining task, both static and streaming. In brief, for a table of transactions and items, the task is to find all subsets of items that occur together in transactions with at least a prescribed frequency. Several variants of the task are described in [8]. Online frequent itemset mining algorithms are surveyed in [57]. Algorithms based on counts of all past data in the stream [48, 101, 155, 234, 152] are also called landmark window based approaches. In some of these algorithms, time adaptivity is achieved by placing more importance on recent items [48, 101, 152]. Sliding window based approaches [47, 58, 49, 210, 56, 154, 232, 39] are particularly suitable for processing data with concept drift. For a comparative overview, see, for example, how MOA's algorithm was selected [192]. Note that a special subtask, finding frequent items in data streams, is already challenging and requires approximate data structures [50].

Principal component analysis (PCA) is a powerful tool for dimensionality reduction [128] based on matrix factorization. Online variants are based on ideas to incrementally update the matrix decomposition [37, 106, 36]. The first PCA algorithms suitable for online learning are based on neural networks [178, 201, 179]. Similar to the linear models in Section 3.2, PCA can also apply the kernel trick to involve nonlinear modeling [205]. Iterative **kernel PCA** is described in [132, 105] and online kernel PCA in [117]. We note that for nearest neighbor search in the low-dimensional space provided by PCA, the heuristics for selecting large inner products is applicable [217].

Probabilistic topic modeling fits complex hierarchical Bayesian models to large document collections. A topic model reveals latent semantic structure that can be used for many applications. While PCA-like models can also be used for latent semantic analysis [69], recently the so-called **Latent Dirichlet Allocation (LDA)** [32] has gained popularity. Most topic model parameters can only be inferred based on Markov Chain Monte Carlo sampling, a method difficult to implement for online learning. LDA inference is possible based on either online Gibbs sampling [211, 40]

or online stochastic optimization with a natural gradient step [116]. Several online LDA variants are described in [209, 115, 157, 235, 233, 123].

6.2 Concept Drift and Adaptive Learning

In dynamically changing and nonstationary environments, we often observe concept drift as the result of data distribution change over time. The phenomenon and mitigation of concept (or dataset) drift for online learning are surveyed in several articles [228, 220, 193, 243, 99]. The area of transfer learning where the (batch) training and the test sets are different [186], is closely related to concept drift [212], but more difficult in the sense that adaptation by learning part of the new data is not possible.

Adaptive learning refers to the technique of updating predictive models online to react to concept drifts. One of the earliest active learning systems is STAG-GER [204]. In [242], the main steps of online adaptive learning are summarized as (1) making assumptions about future distribution, (2) identifying change patterns, (3) designing mechanisms to make the learner adaptive, and (4) parameterizing the model at every time step.

A comprehensive categorization of concept drift adaptation techniques is found in [99]. Online learning algorithms can naturally adapt to evolving distributions. However, adaptation happens only as the old concepts are diluted due to the new incoming data, which is more suitable for gradual changes [162, 73]. For sudden changes, algorithms that maintain a sliding window of the last seen instances perform better [228, 93, 145]. Another option is to include explicit forgetting mechanisms [142, 136, 76]. The most important distinction is whether changes are explicitly or implicitly detected: **Trigger-based** methods aim at detecting when concept drift occurs to build a new model from scratch [93]. **Evolving learners**, by contrast, do not aim to detect changes but rather maintain the most accurate models at each time step. Evolving learners are method-specific, most of them based on ensemble methods [225, 139].

A few papers [168, 171] give overviews of different types of environmental changes and concept drifts based on speed, recurrence, and severity. Drift can happen gradually or suddenly, in isolation, in tendencies or seasonally, predictably or unpredictably, and its effect on classifier performance may or may not be severe. In [204, 93], several artificial data sets with different drift concepts, sudden or abrupt, and gradual changes are described.

A large variety of single classifier and ensemble models capable of handling concept drift are described in [220]. Perhaps the majority of the results consider tree-based methods [13]. For example, concept drift adaptive online decision trees based on a statistical change detector that works on sliding windows are described in [26, 24]. More examples include Bayesian models [94, 19], neural networks [95, 153], and SVM [215, 137]. Concept drift adaptation methods exist for clustering [199, 207]. Sliding window based data stream frequent itemset mining

is also adaptive [192]. Some of the results do not follow the data stream computational model but rather use computational resources with little restriction. One class of such methods are incremental algorithms with partial memory [166]. We also note that there is a MOA-based software system for concept drift detection [30].

References

1. Abadi, D.J., Carney, D., Çetintemel, U., Cherniack, M., Convey, C., Lee, S., Stonebraker, M., Tatbul, N., Zdonik, S.: Aurora: a new model and architecture for data stream management. The VLDB JournalThe International Journal on Very Large Data Bases **12**(2), 120–139 (2003)
2. Aberdeen, D., Pacovsky, O., Slater, A.: The learning behind gmail priority inbox. In: LCCC: NIPS 2010 Workshop on Learning on Cores, Clusters and Clouds (2010)
3. Ackermann, M.R., Märtens, M., Raupach, C., Swierkot, K., Lammersen, C., Sohler, C.: Streamkm++: A clustering algorithm for data streams. Journal of Experimental Algorithmics (JEA) **17**, 2–4 (2012)
4. Adhikari, V.K., Guo, Y., Hao, F., Varvello, M., Hilt, V., Steiner, M., Zhang, Z.L.: Unreeling netflix: Understanding and improving multi-cdn movie delivery. In: INFOCOM, 2012 Proceedings IEEE, pp. 1620–1628. IEEE (2012)
5. Adomavicius, G., Tuzhilin, A.: Context-aware recommender systems. In: Recommender systems handbook, pp. 217–253. Springer (2011)
6. Agarwal, S., Saradhi, V.V., Karnick, H.: Kernel-based online machine learning and support vector reduction. Neurocomputing **71**(7), 1230–1237 (2008)
7. Aggarwal, C.C.: A survey of stream clustering algorithms. Data Clustering: Algorithms and Applications p. 231 (2013)
8. Aggarwal, C.C., Han, J.: Frequent pattern mining. Springer (2014)
9. Aggarwal, C.C., Han, J., Wang, J., Yu, P.S.: A framework for clustering evolving data streams. In: Proceedings of the 29th international conference on Very large data bases-Volume 29, pp. 81–92. VLDB Endowment (2003)
10. Agrawal, R., Imielienski, T., Swami, A.: Mining association rules between sets of items in large databases. In: P. Bunemann, S. Jajodia (eds.) Proceedings of the 1993 ACM SIGMOD Conference on Managment of Data, pp. 207–216. ACM Press, New York (1993)
11. Akidau, T., Balikov, A., Bekiroğlu, K., Chernyak, S., Haberman, J., Lax, R., McVeety, S., Mills, D., Nordstrom, P., Whittle, S.: Millwheel: fault-tolerant stream processing at internet scale. Proceedings of the VLDB Endowment **6**(11), 1033–1044 (2013)
12. Akidau, T., Bradshaw, R., Chambers, C., Chernyak, S., Fernández-Moctezuma, R.J., Lax, R., McVeety, S., Mills, D., Perry, F., Schmidt, E., et al.: The dataflow model: a practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing. Proceedings of the VLDB Endowment **8**(12), 1792–1803 (2015)
13. Alberg, D., Last, M., Kandel, A.: Knowledge discovery in data streams with regression tree methods. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery **2**(1), 69–78 (2012)
14. Ali, M., Johnson, C.C., Tang, A.K.: Parallel collaborative filtering for streaming data. University of Texas Austin, Tech. Rep (2011)
15. Arasu, A., Babcock, B., Babu, S., Datar, M., Ito, K., Nishizawa, I., Rosenstein, J., Widom, J.: Stream: the stanford stream data manager (demonstration description). In: Proceedings of the 2003 ACM SIGMOD international conference on Management of data, pp. 665–665. ACM (2003)
16. Auer, P., Cesa-Bianchi, N., Freund, Y., Schapire, R.E.: The nonstochastic multiarmed bandit problem. SIAM journal on computing **32**(1), 48–77 (2002)

17. Babcock, B., Babu, S., Datar, M., Motwani, R., Widom, J.: Models and issues in data stream systems. In: Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, pp. 1–16. ACM (2002)
18. Babu, S., Widom, J.: Continuous queries over data streams. ACM Sigmod Record **30**(3), 109–120 (2001)
19. Bach, S., Maloof, M.: A bayesian approach to concept drift. In: Advances in neural information processing systems, pp. 127–135 (2010)
20. Balog, K., Kelly, L., Schuth, A.: Head first: Living labs for ad-hoc search evaluation. In: Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, pp. 1815–1818. ACM (2014)
21. Bennett, J., Lanning, S.: The netflix prize. In: KDD Cup and Workshop in conjunction with KDD 2007 (2007)
22. Bernhardt, T., Vasseur, A.: Esper: Event stream processing and correlation. ONJava, in http://www.onjava.com/lpt/a/6955, O'Reilly (2007)
23. Beygelzimer, A., Kale, S., Luo, H.: Optimal and adaptive algorithms for online boosting. In: Proceedings of the 32nd International Conference on Machine Learning (ICML-15), pp. 2323–2331 (2015)
24. Bifet, A.: Adaptive stream mining: Pattern learning and mining from evolving data streams. In: Proceedings of the 2010 conference on adaptive stream mining: Pattern learning and mining from evolving data streams, pp. 1–212. Ios Press (2010)
25. Bifet, A., Frank, E.: Sentiment knowledge discovery in twitter streaming data. In: International conference on discovery science, pp. 1–15. Springer (2010)
26. Bifet, A., Gavaldà, R.: Adaptive learning from evolving data streams. In: International Symposium on Intelligent Data Analysis, pp. 249–260. Springer (2009)
27. Bifet, A., Holmes, G., Kirkby, R., Pfahringer, B.: Moa: Massive online analysis. Journal of Machine Learning Research **11**(May), 1601–1604 (2010)
28. Bifet, A., Holmes, G., Pfahringer, B.: Leveraging bagging for evolving data streams. Machine Learning and Knowledge Discovery in Databases pp. 135–150 (2010)
29. Bifet, A., Kirkby, R., Pfahringer, B.: Data stream mining: a practical approach. Tech. rep., University of Waikato (2011)
30. Bifet, A., Read, J., Pfahringer, B., Holmes, G., Žliobaitė, I.: Cd-moa: change detection framework for massive online analysis. In: International Symposium on Intelligent Data Analysis, pp. 92–103. Springer (2013)
31. Bifet, A., Zhang, J., Fan, W., He, C., Zhang, J., Qian, J., Holmes, G., Pfahringer, B.: Extremely fast decision tree mining for evolving data streams. In: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1733–1742. ACM (2017)
32. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. Journal of machine Learning research **3**(Jan), 993–1022 (2003)
33. Bordes, A., Bottou, L.: The huller: a simple and efficient online svm. In: ECML, pp. 505–512. Springer (2005)
34. Bordes, A., Ertekin, S., Weston, J., Bottou, L.: Fast kernel classifiers with online and active learning. Journal of Machine Learning Research **6**(Sep), 1579–1619 (2005)
35. Bradley, P.S., Fayyad, U.M., Reina, C., et al.: Scaling clustering algorithms to large databases. In: KDD, pp. 9–15 (1998)
36. Brand, M.: Incremental singular value decomposition of uncertain data with missing values. Computer VisionECCV 2002 pp. 707–720 (2002)
37. Bunch, J.R., Nielsen, C.P.: Updating the singular value decomposition. Numerische Mathematik **31**(2), 111–129 (1978)
38. Buntine, W.: Theory refinement on bayesian networks. In: Proceedings of the Seventh conference on Uncertainty in Artificial Intelligence, pp. 52–60. Morgan Kaufmann Publishers Inc. (1991)
39. Calders, T., Dexters, N., Gillis, J.J., Goethals, B.: Mining frequent itemsets in a stream. Information Systems **39**, 233–255 (2014)

40. Canini, K., Shi, L., Griffiths, T.: Online inference of topics with latent dirichlet allocation. In: Artificial Intelligence and Statistics, pp. 65–72 (2009)
41. Cao, F., Estert, M., Qian, W., Zhou, A.: Density-based clustering over an evolving data stream with noise. In: Proceedings of the 2006 SIAM international conference on data mining, pp. 328–339. SIAM (2006)
42. Carbone, P., Ewen, S., Fóra, G., Haridi, S., Richter, S., Tzoumas, K.: State management in apache flink®: consistent stateful distributed stream processing. Proceedings of the VLDB Endowment **10**(12), 1718–1729 (2017)
43. Carbone, P., Katsifodimos, A., Ewen, S., Markl, V., Haridi, S., Tzoumas, K.: Apache flink: Stream and batch processing in a single engine. Bulletin of the IEEE Computer Society Technical Committee on Data Engineering **36**(4) (2015)
44. Cauwenberghs, G., Poggio, T.: Incremental and decremental support vector machine learning. In: Advances in neural information processing systems, pp. 409–415 (2001)
45. Cesa-Bianchi, N., Gentile, C.: Improved risk tail bounds for on-line algorithms. IEEE Transactions on Information Theory **54**(1), 386–390 (2008)
46. Chambers, C., Raniwala, A., Perry, F., Adams, S., Henry, R.R., Bradshaw, R., Weizenbaum, N.: Flumejava: easy, efficient data-parallel pipelines. ACM Sigplan Notices **45**(6), 363–375 (2010)
47. Chang, J.H., Lee, W.S.: estwin: adaptively monitoring the recent change of frequent itemsets over online data streams. In: Proceedings of the twelfth international conference on Information and knowledge management, pp. 536–539. ACM (2003)
48. Chang, J.H., Lee, W.S.: Finding recent frequent itemsets adaptively over online data streams. In: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 487–492. ACM (2003)
49. Chang, J.H., Lee, W.S.: Finding frequent itemsets over online data streams. Information and software technology **48**(7), 606–618 (2006)
50. Charikar, M., Chen, K., Farach-Colton, M.: Finding frequent items in data streams. Theoretical Computer Science **312**(1), 3–15 (2004)
51. Chen, R., Sivakumar, K., Kargupta, H.: An approach to online bayesian learning from multiple data streams. In: Workshop on Ubiquitous Data Mining for Mobile and Distributed Environments, Held in Conjunction with Joint 12th European Conference on Machine Learning (ECML'01) and 5th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD'01), Freiburg, Germany (2001)
52. Chen, S.T., Lin, H.T., Lu, C.J.: An online boosting algorithm with theoretical justifications. In: Proceedings of the 29th International Coference on International Conference on Machine Learning, pp. 1873–1880. Omnipress (2012)
53. Chen, T., Guestrin, C.: Xgboost: A scalable tree boosting system. In: Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining, pp. 785–794. ACM (2016)
54. Chen, T., Li, M., Li, Y., Lin, M., Wang, N., Wang, M., Xiao, T., Xu, B., Zhang, C., Zhang, Z.: MXNet: A flexible and efficient machine learning library for heterogeneous distributed systems. arXiv preprint arXiv:1512.01274 (2015)
55. Chen, Y., Tu, L.: Density-based clustering for real-time stream data. In: Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 133–142. ACM (2007)
56. Cheng, J., Ke, Y., Ng, W.: Maintaining frequent closed itemsets over a sliding window. Journal of Intelligent Information Systems **31**(3), 191–215 (2008)
57. Cheng, J., Ke, Y., Ng, W.: A survey on algorithms for mining frequent itemsets over data streams. Knowledge and Information Systems **16**(1), 1–27 (2008)
58. Chi, Y., Wang, H., Philip, S.Y., Muntz, R.R.: Catch the moment: maintaining closed frequent itemsets over a data stream sliding window. Knowledge and Information Systems **10**(3), 265–294 (2006)
59. Collobert, R., Weston, J.: A unified architecture for natural language processing: Deep neural networks with multitask learning. In: Proceedings of the 25th international conference on Machine learning, pp. 160–167. ACM (2008)

60. Crammer, K., Dekel, O., Keshet, J., Shalev-Shwartz, S., Singer, Y.: Online passive-aggressive algorithms. The Journal of Machine Learning Research **7**, 551–585 (2006)
61. Crammer, K., Kandola, J., Singer, Y.: Online classification on a budget. In: Advances in neural information processing systems, pp. 225–232 (2004)
62. Crammer, K., Singer, Y.: Ultraconservative online algorithms for multiclass problems. Journal of Machine Learning Research **3**(Jan), 951–991 (2003)
63. Crankshaw, D., Wang, X., Zhou, G., Franklin, M.J., Gonzalez, J.E., Stoica, I.: Clipper: A low-latency online prediction serving system. In: NSDI, pp. 613–627 (2017)
64. Dawid, A.P.: Present position and potential developments: Some personal views: Statistical theory: The prequential approach. Journal of the Royal Statistical Society. Series A (General) pp. 278–292 (1984)
65. De Francisci Morales, G.: Samoa: A platform for mining big data streams. In: Proceedings of the 22nd International Conference on World Wide Web, pp. 777–778. ACM (2013)
66. De Francisci Morales, G., Bifet, A., Khan, L., Gama, J., Fan, W.: Iot big data stream mining. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 2119–2120. ACM (2016)
67. Dean, J., Corrado, G., Monga, R., Chen, K., Devin, M., Mao, M., Senior, A., Tucker, P., Yang, K., Le, Q.V., et al.: Large scale distributed deep networks. In: Advances in neural information processing systems, pp. 1223–1231 (2012)
68. Dean, J., Ghemawat, S.: Mapreduce: Simplified data processing on large clusters. Communications of the ACM **51**(1), 107–113 (2008)
69. Deerwester, S.C., Dumais, S.T., Landauer, T.K., Furnas, G.W., Harshman, R.A.: Indexing by latent semantic analysis. Journal of the American Society of Information Science **41**(6), 391–407 (1990). URL citeseer.nj.nec.com/deerwester90indexing.html
70. Denil, M., Matheson, D., Nando, D.: Consistency of online random forests. In: Proceedings of the 30th International Conference on Machine Learning (ICML-13), pp. 1256–1264 (2013)
71. Deshpande, M., Karypis, G.: Item-based top-n recommendation algorithms. ACM Transactions on Information Systems (TOIS) **22**(1), 143–177 (2004)
72. Ding, Y., Li, X.: Time weight collaborative filtering. In: Proceedings of the 14th ACM international conference on Information and knowledge management, pp. 485–492. ACM (2005)
73. Domingos, P., Hulten, G.: Mining high-speed data streams. In: Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 71–80. ACM (2000)
74. Domingos, P., Pazzani, M.: On the optimality of the simple bayesian classifier under zero-one loss. Machine learning **29**(2), 103–130 (1997)
75. Duda, R.O., Hart, P.E., Stork, D.G.: Pattern classification. John Wiley & Sons (2012)
76. Elwell, R., Polikar, R.: Incremental learning in nonstationary environments with controlled forgetting. In: Neural Networks, 2009. IJCNN 2009. International Joint Conference on, pp. 771–778. IEEE (2009)
77. Ester, M., Kriegel, H.P., Sander, J., Xu, X., et al.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: Proc. ACM SigKDD, pp. 226–231 (1996)
78. Fan, W., Bifet, A.: Mining big data: current status, and forecast to the future. ACM SIGKDD Explorations Newsletter **14**(2), 1–5 (2013)
79. Fan, W., Stolfo, S.J., Zhang, J.: The application of adaboost for distributed, scalable and on-line learning. In: Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 362–366. ACM (1999)
80. Farnstrom, F., Lewis, J., Elkan, C.: Scalability for clustering algorithms revisited. ACM SIGKDD Explorations Newsletter **2**(1), 51–57 (2000)
81. Fitzpatrick, B.: Distributed caching with memcached. Linux journal **2004**(124), 5 (2004)
82. Fogarty, J., Baker, R.S., Hudson, S.E.: Case studies in the use of roc curve analysis for sensor-based estimates in human computer interaction. In: Proceedings of Graphics Interface 2005, GI '05, pp. 129–136. Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada (2005). URL http://portal.acm.org/citation.cfm?id=1089508.1089530

83. Fontenla-Romero, Ó., Guijarro-Berdiñas, B., Martinez-Rego, D., Pérez-Sánchez, B., Peteiro-Barral, D.: Online machine learning. Efficiency and Scalability Methods for Computational Intellect **27** (2013)

84. Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. In: European conference on computational learning theory, pp. 23–37. Springer (1995)

85. Freund, Y., Schapire, R.E.: Large margin classification using the perceptron algorithm. Machine learning **37**(3), 277–296 (1999)

86. Frie, T.T., Cristianini, N., Campbell, C.: The kernel-adatron algorithm: a fast and simple learning procedure for support vector machines. In: Machine Learning: Proceedings of the Fifteenth International Conference (ICML'98), pp. 188–196 (1998)

87. Friedman, J., Hastie, T., Tibshirani, R.: The elements of statistical learning, vol. 1. Springer series in statistics New York (2001)

88. Friedman, N., Goldszmidt, M.: Sequential update of bayesian network structure. In: Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence, pp. 165–174. Morgan Kaufmann Publishers Inc. (1997)

89. Frigó, E., Pálovics, R., Kelen, D., Benczúr, A.A., Kocsis, L.: Online ranking prediction in non-stationary environments. In: Proceedings of the 1st Workshop on Temporal Reasoning in Recommender Systems, co-located with 11th International Conference on Recommender Systems (2017)

90. Funk, S.: Netflix update: Try this at home http://sifter.org/simon/journal/20061211.html (2006). URL http://sifter.org/simon/journal/20061211.html

91. Gaber, M., Zaslavsky, A., Krishnaswamy, S.: A survey of classification methods in data streams. Data streams pp. 39–59 (2007)

92. Gaber, M.M., Gama, J., Krishnaswamy, S., Gomes, J.B., Stahl, F.: Data stream mining in ubiquitous environments: state-of-the-art and current directions. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery **4**(2), 116–138 (2014)

93. Gama, J., Medas, P., Castillo, G., Rodrigues, P.: Learning with drift detection. In: Brazilian Symposium on Artificial Intelligence, pp. 286–295. Springer (2004)

94. Gama, J., Rocha, R., Medas, P.: Accurate decision trees for mining high-speed data streams. In: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 523–528. ACM (2003)

95. Gama, J., Rodrigues, P.P.: Stream-based electricity load forecast. In: European Conference on Principles of Data Mining and Knowledge Discovery, pp. 446–453. Springer (2007)

96. Gama, J., Rodrigues, P.P., Lopes, L.: Clustering distributed sensor data streams using local processing and reduced communication. Intelligent Data Analysis **15**(1), 3–28 (2011)

97. Gama, J., Sebastião, R., Rodrigues, P.P.: Issues in evaluation of stream learning algorithms. In: Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 329–338. ACM (2009)

98. Gama, J., Sebastião, R., Rodrigues, P.P.: On evaluating stream learning algorithms. Machine learning **90**(3), 317–346 (2013)

99. Gama, J., Žliobaite, I., Bifet, A., Pechenizkiy, M., Bouchachia, A.: A survey on concept drift adaptation. ACM Computing Surveys (CSUR) **46**(4), 44 (2014)

100. Gentile, C.: A new approximate maximal margin classification algorithm. Journal of Machine Learning Research **2**(Dec), 213–242 (2001)

101. Giannella, C., Han, J., Pei, J., Yan, X., Yu, P.S.: Mining frequent patterns in data streams at multiple time granularities. Next generation data mining **212**, 191–212 (2003)

102. Grossman, R.L., Hornick, M.F., Meyer, G.: Data mining standards initiatives. Communications of the ACM **45**(8), 59–61 (2002)

103. Gualtieri, M., Rowan Curran, A., TaKeaways, K., To, M.T.B.P.P.: The forrester wave: Big data predictive analytics solutions, q1 2013. Forrester research (2013)

104. Guha, S., Meyerson, A., Mishra, N., Motwani, R., O'Callaghan, L.: Clustering data streams: Theory and practice. IEEE transactions on knowledge and data engineering **15**(3), 515–528 (2003)

105. Günter, S., Schraudolph, N.N., Vishwanathan, S.: Fast iterative kernel principal component analysis. Journal of Machine Learning Research **8**(Aug), 1893–1918 (2007)
106. Hall, P., Marshall, D., Martin, R.: Merging and splitting eigenspace models. IEEE Transactions on pattern analysis and machine intelligence **22**(9), 1042–1049 (2000)
107. Hartigan, J.A., Hartigan, J.: Clustering algorithms, vol. 209. Wiley New York (1975)
108. Haselsteiner, E., Pfurtscheller, G.: Using time-dependent neural networks for eeg classification. IEEE transactions on rehabilitation engineering **8**(4), 457–463 (2000)
109. He, X., Zhang, H., Kan, M.Y., Chua, T.S.: Fast matrix factorization for online recommendation with implicit feedback. In: Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, pp. 549–558. ACM (2016)
110. Henzinger, M.R., Raghavan, P., Rajagopalan, S.: Computing on data streams. External memory algorithms **50**, 107–118 (1998)
111. Hidasi, B.: Factorization models for context-aware recommendations. Infocommun J VI (4) pp. 27–34 (2014)
112. Hidasi, B., Tikk, D.: Fast als-based tensor factorization for context-aware recommendation from implicit feedback. In: Machine Learning and Knowledge Discovery in Databases, pp. 67–82. Springer (2012)
113. Hidasi, B., Tikk, D.: General factorization framework for context-aware recommendations. Data Mining and Knowledge Discovery **30**(2), 342–371 (2016)
114. Hinton, G., Deng, L., Yu, D., Dahl, G.E., Mohamed, A.r., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T.N., et al.: Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. IEEE Signal Processing Magazine **29**(6), 82–97 (2012)
115. Ho, Q., Cipar, J., Cui, H., Lee, S., Kim, J.K., Gibbons, P.B., Gibson, G.A., Ganger, G., Xing, E.P.: More effective distributed ml via a stale synchronous parallel parameter server. In: Advances in neural information processing systems, pp. 1223–1231 (2013)
116. Hoffman, M., Bach, F.R., Blei, D.M.: Online learning for latent dirichlet allocation. In: advances in neural information processing systems, pp. 856–864 (2010)
117. Honeine, P.: Online kernel principal component analysis: A reduced-order model. IEEE transactions on pattern analysis and machine intelligence **34**(9), 1814–1826 (2012)
118. Hu, Y., Koren, Y., Volinsky, C.: Collaborative filtering for implicit feedback datasets. In: Data Mining, 2008. ICDM'08. Eighth IEEE International Conference on, pp. 263–272. IEEE (2008)
119. Hulten, G., Spencer, L., Domingos, P.: Mining time-changing data streams. In: Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 97–106. ACM (2001)
120. Ikonomovska, E., Gama, J., Džeroski, S.: Online tree-based ensembles and option trees for regression on evolving data streams. Neurocomputing **150**, 458–470 (2015)
121. Ipek, E., Mutlu, O., Martínez, J.F., Caruana, R.: Self-optimizing memory controllers: A reinforcement learning approach. In: Computer Architecture, 2008. ISCA'08. 35th International Symposium on, pp. 39–50. IEEE (2008)
122. Isaacman, S., Ioannidis, S., Chaintreau, A., Martonosi, M.: Distributed rating prediction in user generated content streams. In: Proceedings of the fifth ACM conference on Recommender systems, pp. 69–76. ACM (2011)
123. Jagerman, R., Eickhoff, C., de Rijke, M.: Computing web-scale topic models using an asynchronous parameter server. In: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM (2017)
124. Jain, A.K., Murty, M.N., Flynn, P.J.: Data clustering: a review. ACM computing surveys (CSUR) **31**(3), 264–323 (1999)
125. Jiang, J., Cui, B., Zhang, C., Yu, L.: Heterogeneity-aware distributed parameter servers. In: Proceedings of the 2017 ACM International Conference on Management of Data, pp. 463–478. ACM (2017)
126. Jiang, J., Yu, L., Jiang, J., Liu, Y., Cui, B.: Angel: a new large-scale machine learning system. National Science Review nwx018 (2017)

127. Jin, R., Agrawal, G.: Efficient decision tree construction on streaming data. In: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 571–576. ACM (2003)
128. Jolliffe, I.T.: Principal component analysis and factor analysis. In: Principal component analysis, pp. 115–128. Springer (1986)
129. Juang, C.F., Lin, C.T.: An online self-constructing neural fuzzy inference network and its applications. IEEE transactions on Fuzzy Systems **6**(1), 12–32 (1998)
130. Karydi, E., Margaritis, K.: Parallel and distributed collaborative filtering: A survey. ACM Computing Surveys (CSUR) **49**(2), 37 (2016)
131. Kavitha, V., Punithavalli, M.: Clustering time series data stream-a literature survey. arXiv preprint arXiv:1005.4270 (2010)
132. Kim, K.I., Franz, M.O., Scholkopf, B.: Iterative kernel principal component analysis for image modeling. IEEE transactions on pattern analysis and machine intelligence **27**(9), 1351–1366 (2005)
133. Kiran, M., Murphy, P., Monga, I., Dugan, J., Baveja, S.S.: Lambda architecture for cost-effective batch and speed big data processing. In: Big Data (Big Data), 2015 IEEE International Conference on, pp. 2785–2792. IEEE (2015)
134. Kivinen, J., Smola, A.J., Williamson, R.C.: Online learning with kernels. IEEE transactions on signal processing **52**(8), 2165–2176 (2004)
135. Kivinen, J., Warmuth, M.K.: Exponentiated gradient versus gradient descent for linear predictors. Information and Computation **132**(1), 1–63 (1997)
136. Klinkenberg, R.: Learning drifting concepts: Example selection vs. example weighting. Intelligent Data Analysis **8**(3), 281–300 (2004)
137. Klinkenberg, R., Joachims, T.: Detecting concept drift with support vector machines. In: ICML, pp. 487–494 (2000)
138. Koenigstein, N., Koren, Y.: Towards scalable and accurate item-oriented recommendations. In: Proceedings of the 7th ACM conference on Recommender systems, pp. 419–422. ACM (2013)
139. Kolter, J.Z., Maloof, M.A.: Dynamic weighted majority: A new ensemble method for tracking concept drift. In: Data Mining, 2003. ICDM 2003. Third IEEE International Conference on, pp. 123–130. IEEE (2003)
140. Koren, Y., Bell, R., Volinsky, C.: Matrix factorization techniques for recommender systems. Computer **42**(8), 30–37 (2009)
141. Kourtellis, N., Morales, G.D.F., Bifet, A., Murdopo, A.: Vht: Vertical hoeffding tree. In: Big Data (Big Data), 2016 IEEE International Conference on, pp. 915–922. IEEE (2016)
142. Koychev, I.: Gradual forgetting for adaptation to concept drift. In: Proceedings of ECAI 2000 Workshop on Current Issues in Spatio-Temporal Reasoning (2000)
143. Kranen, P., Assent, I., Baldauf, C., Seidl, T.: The clustree: indexing micro-clusters for anytime stream mining. Knowledge and information systems **29**(2), 249–272 (2011)
144. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems, pp. 1097–1105 (2012)
145. Kuncheva, L.I., Žliobaitė, I.: On the window size for classification in changing environments. Intelligent Data Analysis **13**(6), 861–872 (2009)
146. Lam, W., Liu, L., Prasad, S., Rajaraman, A., Vacheri, Z., Doan, A.: Muppet: Mapreduce-style processing of fast data. Proceedings of the VLDB Endowment **5**(12), 1814–1825 (2012)
147. Langford, J., Li, L., Zhang, T.: Sparse online learning via truncated gradient. Journal of Machine Learning Research **10**(Mar), 777–801 (2009)
148. Lathia, N., Hailes, S., Capra, L.: Temporal collaborative filtering with adaptive neighbourhoods. In: Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval, pp. 796–797. ACM (2009)
149. Law, Y.N., Zaniolo, C.: An adaptive nearest neighbor classification algorithm for data streams. In: European Conference on Principles of Data Mining and Knowledge Discovery, pp. 108–120. Springer (2005)

150. Lazarevic, A., Obradovic, Z.: Boosting algorithms for parallel and distributed learning. Distributed and parallel databases **11**(2), 203–229 (2002)
151. LeCun, Y., Bottou, L., Orr, G.B., Müller, K.R.: Efficient backprop. In: Neural networks: Tricks of the trade, pp. 9–50. Springer (1998)
152. Lee, D., Lee, W.: Finding maximal frequent itemsets over online data streams adaptively. In: Data Mining, Fifth IEEE International Conference on, pp. 8–pp. IEEE (2005)
153. Leite, D., Costa, P., Gomide, F.: Evolving granular neural networks from fuzzy data streams. Neural Networks **38**, 1–16 (2013)
154. Li, H.F., Ho, C.C., Lee, S.Y.: Incremental updates of closed frequent itemsets over continuous data streams. Expert Systems with Applications **36**(2), 2451–2458 (2009)
155. Li, H.F., Lee, S.Y., Shan, M.K.: An efficient algorithm for mining frequent itemsets over the entire history of data streams. In: Proc. of First International Workshop on Knowledge Discovery in Data Streams, vol. 39 (2004)
156. Li, L., Chu, W., Langford, J., Schapire, R.E.: A contextual-bandit approach to personalized news article recommendation. In: Proceedings of the 19th international conference on World wide web, pp. 661–670. ACM (2010)
157. Li, M., Andersen, D.G., Park, J.W., Smola, A.J., Ahmed, A., Josifovski, V., Long, J., Shekita, E.J., Su, B.Y.: Scaling distributed machine learning with the parameter server. In: 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14), pp. 583–598. USENIX Association (2014)
158. Li, M., Andersen, D.G., Smola, A.J., Yu, K.: Communication efficient distributed machine learning with the parameter server. In: Advances in Neural Information Processing Systems, pp. 19–27 (2014)
159. Li, M., Zhou, L., Yang, Z., Li, A., Xia, F., Andersen, D.G., Smola, A.: Parameter server for distributed machine learning. In: Big Learning NIPS Workshop, vol. 6, p. 2 (2013)
160. Li, Y., Long, P.M.: The relaxed online maximum margin algorithm. Machine Learning **1**(46), 361–387 (2002)
161. Linden, G., Smith, B., York, J.: Amazon.com recommendations: item-to-item collaborative filtering. Internet Computing, IEEE **7**(1), 76–80 (2003)
162. Littlestone, N.: Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. Machine learning **2**(4), 285–318 (1988)
163. Lops, P., De Gemmis, M., Semeraro, G.: Content-based recommender systems: State of the art and trends. In: Recommender systems handbook, pp. 73–105. Springer (2011)
164. Low, Y., Bickson, D., Gonzalez, J., Guestrin, C., Kyrola, A., Hellerstein, J.M.: Distributed graphlab: a framework for machine learning and data mining in the cloud. Proceedings of the VLDB Endowment **5**(8), 716–727 (2012)
165. Mahdiraji, A.R.: Clustering data stream: A survey of algorithms. International Journal of Knowledge-based and Intelligent Engineering Systems **13**(2), 39–44 (2009)
166. Maloof, M.A., Michalski, R.S.: Incremental learning with partial instance memory. Artificial intelligence **154**(1-2), 95–126 (2004)
167. Marz, N., Warren, J.: Big Data: Principles and best practices of scalable realtime data systems. Manning Publications Co. (2015)
168. Minku, L.L., White, A.P., Yao, X.: The impact of diversity on online ensemble learning in the presence of concept drift. IEEE Transactions on Knowledge and Data Engineering **22**(5), 730–742 (2010)
169. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al.: Human-level control through deep reinforcement learning. Nature **518**(7540), 529–533 (2015)
170. Morales, G.D.F., Bifet, A.: Samoa: scalable advanced massive online analysis. Journal of Machine Learning Research **16**(1), 149–153 (2015)
171. Moreno-Torres, J.G., Raeder, T., Alaiz-RodríGuez, R., Chawla, N.V., Herrera, F.: A unifying view on dataset shift in classification. Pattern Recognition **45**(1), 521–530 (2012)
172. Muthukrishnan, S., et al.: Data streams: Algorithms and applications. Foundations and Trends® in Theoretical Computer Science **1**(2), 117–236 (2005)

173. Nair, A., Srinivasan, P., Blackwell, S., Alcicek, C., Fearon, R., De Maria, A., Panneershel-vam, V., Suleyman, M., Beattie, C., Petersen, S., et al.: Massively parallel methods for deep reinforcement learning. arXiv preprint arXiv:1507.04296 (2015)
174. Neumeyer, L., Robbins, B., Nair, A., Kesari, A.: S4: Distributed stream computing platform. In: Data Mining Workshops (ICDMW), 2010 IEEE International Conference on, pp. 170–177. IEEE (2010)
175. Noghabi, S.A., Paramasivam, K., Pan, Y., Ramesh, N., Bringhurst, J., Gupta, I., Campbell, R.H.: Samza: stateful scalable stream processing at linkedin. Proceedings of the VLDB Endowment **10**(12), 1634–1645 (2017)
176. Obermaier, B., Guger, C., Neuper, C., Pfurtscheller, G.: Hidden markov models for online classification of single trial eeg data. Pattern recognition letters **22**(12), 1299–1309 (2001)
177. O'callaghan, L., Mishra, N., Meyerson, A., Guha, S., Motwani, R.: Streaming-data algo-rithms for high-quality clustering. In: Data Engineering, 2002. Proceedings. 18th Interna-tional Conference on, pp. 685–694. IEEE (2002)
178. Oja, E.: Simplified neuron model as a principal component analyzer. Journal of mathematical biology **15**(3), 267–273 (1982)
179. Oja, E.: Principal components, minor components, and linear neural networks. Neural net-works **5**(6), 927–935 (1992)
180. Owen, S., Anil, R., Dunning, T., Friedman, E.: Mahout in Action. Manning Publications (2011)
181. Oza, N.C.: Online bagging and boosting. In: Systems, man and cybernetics, 2005 IEEE international conference on, vol. 3, pp. 2340–2345. IEEE (2005)
182. Palit, I., Reddy, C.K.: Scalable and parallel boosting with mapreduce. IEEE Transactions on Knowledge and Data Engineering **24**(10), 1904–1916 (2012)
183. Pálovics, R., Benczúr, A.A.: Temporal influence over the last.fm social network. In: Pro-ceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, pp. 486–493. ACM (2013)
184. Pálovics, R., Benczúr, A.A., Kocsis, L., Kiss, T., Frigó, E.: Exploiting temporal influence in online recommendation. In: Proceedings of the 8th ACM Conference on Recommender systems, pp. 273–280. ACM (2014)
185. Pálovics, R., Kelen, D., Benczúr, A.A.: Tutorial on open source online learning recom-menders. In: Proceedings of the Eleventh ACM Conference on Recommender Systems, pp. 400–401. ACM (2017)
186. Pan, S.J., Yang, Q.: A survey on transfer learning. IEEE Transactions on knowledge and data engineering **22**(10), 1345–1359 (2010)
187. Pang-Ning, T., Steinbach, M., Kumar, V., et al.: Introduction to data mining. WP Co (2006)
188. Paterek, A.: Improving regularized singular value decomposition for collaborative filtering. In: Proc. KDD Cup Workshop at SIGKDD'07, 13th ACM Int. Conf. on Knowledge Discovery and Data Mining, pp. 39–42 (2007)
189. Pearl, J.: Probabilistic reasoning in intelligent systems: networks of plausible inference. Mor-gan Kaufmann (2014)
190. Pilászy, I., Serény, A., Dózsa, G., Hidasi, B., Sári, A., Gub, J.: Neighbor methods vs matrix factorization - case studies of real-life recommendations. In: LSRS2015 at RECSYS (2015)
191. Pilászy, I., Zibriczky, D., Tikk, D.: Fast als-based matrix factorization for explicit and implicit feedback datasets. In: Proceedings of the fourth ACM conference on Recommender systems, pp. 71–78. ACM (2010)
192. Quadrana, M., Bifet, A., Gavalda, R.: An efficient closed frequent itemset miner for the moa stream mining system. AI Communications **28**(1), 143–158 (2015)
193. Quionero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N.D.: Dataset shift in machine learning. The MIT Press (2009)
194. Ranjan, R.: Streaming big data processing in datacenter clouds. IEEE Cloud Computing **1**(1), 78–83 (2014)
195. Rendle, S., Freudenthaler, C.: Improving pairwise learning for item recommendation from implicit feedback. In: Proceedings of the 7th ACM international conference on Web search and data mining, pp. 273–282. ACM (2014)

196. Rendle, S., Freudenthaler, C., Gantner, Z., Schmidt-Thieme, L.: Bpr: Bayesian personalized ranking from implicit feedback. In: Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, pp. 452–461. AUAI Press (2009)
197. Rendle, S., Schmidt-Thieme, L.: Pairwise interaction tensor factorization for personalized tag recommendation. In: Proceedings of the third ACM international conference on Web search and data mining, pp. 81–90. ACM (2010)
198. Ricci, F., Rokach, L., Shapira, B.: Introduction to recommender systems handbook. Springer (2011)
199. Rodrigues, P.P., Gama, J., Pedroso, J.P.: Odac: Hierarchical clustering of time series data streams. In: Proceedings of the 2006 SIAM International Conference on Data Mining, pp. 499–503. SIAM (2006)
200. Rosenblatt, F.: The perceptron: A probabilistic model for information storage and organization in the brain. Psychological review **65**(6), 386 (1958)
201. Sanger, T.D.: Optimal unsupervised learning in a single-layer linear feedforward neural network. Neural networks **2**(6), 459–473 (1989)
202. Sarwar, B., Karypis, G., Konstan, J., Reidl, J.: Item-based collaborative filtering recommendation algorithms. In: WWW '01: Proceedings of the 10th international conference on World Wide Web, pp. 285–295. ACM Press, New York, NY, USA (2001). DOI 10.1145/371920.372071. URL http://portal.acm.org/citation.cfm?id=372071
203. Schelter, S., Satuluri, V., Zadeh, R.B.: Factorbirda parameter server approach to distributed matrix factorization. In: NIPS 2014 Workshop on Distributed Machine Learning and Matrix Computations (2014)
204. Schlimmer, J.C., Granger, R.H.: Incremental learning from noisy data. Machine learning **1**(3), 317–354 (1986)
205. Schölkopf, B., Smola, A., Müller, K.R.: Nonlinear component analysis as a kernel eigenvalue problem. Neural computation **10**(5), 1299–1319 (1998)
206. Shalev-Shwartz, S., et al.: Online learning and online convex optimization. Foundations and Trends® in Machine Learning **4**(2), 107–194 (2012)
207. Silva, J.A., Faria, E.R., Barros, R.C., Hruschka, E.R., de Carvalho, A.C., Gama, J.: Data stream clustering: A survey. ACM Computing Surveys (CSUR) **46**(1), 13 (2013)
208. Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al.: Mastering the game of go with deep neural networks and tree search. Nature **529**(7587), 484–489 (2016)
209. Smola, A., Narayanamurthy, S.: An architecture for parallel topic models. Proceedings of the VLDB Endowment **3**(1-2), 703–710 (2010)
210. Song, G., Yang, D., Cui, B., Zheng, B., Liu, Y., Xie, K.: Claim: An efficient method for relaxed frequent closed itemsets mining over stream data. In: International Conference on Database Systems for Advanced Applications, pp. 664–675. Springer (2007)
211. Song, X., Lin, C.Y., Tseng, B.L., Sun, M.T.: Modeling and predicting personal information dissemination behavior. In: Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining, pp. 479–488. ACM (2005)
212. Storkey, A.: When training and test sets are different: characterizing learning transfer. Dataset shift in machine learning pp. 3–28 (2009)
213. Sutton, R.S.: Generalization in reinforcement learning: Successful examples using sparse coarse coding. In: Advances in neural information processing systems, pp. 1038–1044 (1996)
214. Sutton, R.S., Barto, A.G.: Reinforcement learning: An introduction. MIT press Cambridge (1998)
215. Syed, N.A., Liu, H., Sung, K.K.: Handling concept drifts in incremental learning with support vector machines. In: Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 317–321. ACM (1999)
216. Takács, G., Pilászy, I., Németh, B., Tikk, D.: Scalable collaborative filtering approaches for large recommender syst ems. Journal of machine learning research **10**(Mar), 623–656 (2009)
217. Teflioudi, C., Gemulla, R., Mykytiuk, O.: Lemp: Fast retrieval of large entries in a matrix product. In: Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, pp. 107–122. ACM (2015)

218. Tesauro, G.: Td-gammon: A self-teaching backgammon program. In: Applications of Neural Networks, pp. 267–285. Springer (1995)

219. Toshniwal, A., Taneja, S., Shukla, A., Ramasamy, K., Patel, J.M., Kulkarni, S., Jackson, J., Gade, K., Fu, M., Donham, J., et al.: Storm @ Twitter. In: Proceedings of the 2014 ACM SIGMOD international conference on Management of data, pp. 147–156. ACM (2014)

220. Tsymbal, A.: The problem of concept drift: definitions and related work. Tech. Rep. 2, Computer Science Department, Trinity College Dublin (2004)

221. Vapnik, V., Chapelle, O.: Bounds on error expectation for support vector machines. Neural computation **12**(9), 2013–2036 (2000)

222. Vapnik, V.N.: Statistical learning theory, vol. 1. Wiley New York (1998)

223. Vasiloudis, T., Beligianni, F., Morales, G.D.F.: Boostvht: Boosting distributed streaming decision trees. In: CIKM (2017)

224. Vinagre, J., Jorge, A.M., Gama, J.: Evaluation of recommender systems in streaming environments. In: Workshop on Recommender Systems Evaluation: Dimensions and Design (REDD 2014), held in conjunction with RecSys 2014. October 10, 2014, Silicon Valley, United States (2014)

225. Wang, H., Fan, W., Yu, P.S., Han, J.: Mining concept-drifting data streams using ensemble classifiers. In: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 226–235. AcM (2003)

226. Watkins, C.J., Dayan, P.: Q-learning. Machine learning **8**(3-4), 279–292 (1992)

227. White, T.: Hadoop: The Definitive Guide. Yahoo Press (2010)

228. Widmer, G., Kubat, M.: Learning in the presence of concept drift and hidden contexts. Machine learning **23**(1), 69–101 (1996)

229. Williams, R.J.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. Machine learning **8**(3-4), 229–256 (1992)

230. Xing, E.P., Ho, Q., Dai, W., Kim, J.K., Wei, J., Lee, S., Zheng, X., Xie, P., Kumar, A., Yu, Y.: Petuum: A new platform for distributed machine learning on big data. IEEE Transactions on Big Data **1**(2), 49–67 (2015)

231. Xu, R., Wunsch, D.: Clustering, vol. 10. John Wiley & Sons (2008)

232. Yen, S.J., Wu, C.W., Lee, Y.S., Tseng, V.S., Hsieh, C.H.: A fast algorithm for mining frequent closed itemsets over stream sliding window. In: Fuzzy Systems (FUZZ), 2011 IEEE International Conference on, pp. 996–1002. IEEE (2011)

233. Yu, H.F., Hsieh, C.J., Yun, H., Vishwanathan, S., Dhillon, I.S.: A scalable asynchronous distributed algorithm for topic modeling. In: Proceedings of the 24th International Conference on World Wide Web, pp. 1340–1350. International World Wide Web Conferences Steering Committee (2015)

234. Yu, J.X., Chong, Z., Lu, H., Zhou, A.: False positive or false negative: mining frequent itemsets from high speed transactional data streams. In: Proceedings of the Thirtieth international conference on Very large data bases-Volume 30, pp. 204–215. VLDB Endowment (2004)

235. Yuan, J., Gao, F., Ho, Q., Dai, W., Wei, J., Zheng, X., Xing, E.P., Liu, T.Y., Ma, W.Y.: Lightlda: Big topic models on modest computer clusters. In: Proceedings of the 24th International Conference on World Wide Web, pp. 1351–1361. International World Wide Web Conferences Steering Committee (2015)

236. Zaharia, M., Chowdhury, M., Franklin, M.J., Shenker, S., Stoica, I.: Spark: Cluster computing with working sets. HotCloud **10**(10-10), 95 (2010)

237. Zhang, T., Ramakrishnan, R., Livny, M.: Birch: an efficient data clustering method for very large databases. ACM Sigmod Record **25**(2), 103–114 (1996)

238. Zhao, P., Jin, R., Yang, T., Hoi, S.C.: Online auc maximization. In: Proceedings of the 28th international conference on machine learning (ICML-11), pp. 233–240 (2011)

239. Zhou, A., Cao, F., Qian, W., Jin, C.: Tracking clusters in evolving data streams over sliding windows. Knowledge and Information Systems **15**(2), 181–214 (2008)

240. Zhou, J., Cui, Q., Li, X., Zhao, P., Qu, S., Huang, J.: Psmart: Parameter server based multiple additive regression trees system. In: Proceedings of the 26th International Conference on World Wide Web Companion, pp. 879–880. International World Wide Web Conferences Steering Committee (2017)

241. Zhu, Y., Shasha, D.: Statstream: Statistical monitoring of thousands of data streams in real time. In: Proceedings of the 28th international conference on Very Large Data Bases, pp. 358–369. VLDB Endowment (2002)
242. Žliobaitė, I.: Learning under concept drift: an overview. Tech. rep., Vilnius University (2009)
243. Žliobaite, I., Bifet, A., Gaber, M., Gabrys, B., Gama, J., Minku, L., Musial, K.: Next challenges for adaptive learning systems. ACM SIGKDD Explorations Newsletter **14**(1), 48–55 (2012)