Flow-Based Programming For Machine Learning

Tanmaya Mahapatra (✉ tanmaya.mahapatra@tum.de)

Technische Universität München <https://orcid.org/0000-0002-7946-5497>

Syedda Nilofer Banoo

Technische Universität München: Technische Universität München

Research

Keywords: End-User Programming, Graphical Flows, Graphical Programming Tools, Machine learning as a service (MLaaS), Machine-Learning-Platform-as-a-Service (ML PaaS), Machine Learning Pipelines

Posted Date: July 27th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-707294/v1>

License: This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Flow-based Programming for Machine Learning

Tanmaya Mahapatra · Syeeda Nilofer Banoo

Received: July 11, 2021/ Revised: NA / Accepted: NA

Abstract Machine Learning (ML) has gained prominence and has tremendous applications in fields like medicine, biology, geography and astrophysics, to name a few. Arguably, in such areas it is used by domain experts, who are not necessarily skilled-programmers. Thus, it presents a steep learning curve for such domain experts in programming ML applications. To overcome this and foster widespread adoption of ML techniques, we propose to equip them with domain-specific graphical tools. Such tools, based on the principles of flow-based programming paradigm, would support the graphical composition of ML applications at a higher level of abstraction and auto-generation of target code. Accordingly, (i) we have modelled ML algorithms as composable components; (ii) described an approach to parse a flow created by connecting several such composable components and use an API-based code generation technique to generate the ML application. To demonstrate the feasibility of our conceptual approach, we have modelled the APIs of Apache Spark ML as composable components and validated it in three use-cases. The use-cases are designed to capture the ease of program specification at a higher abstraction level, easy parametrisation of ML APIs, auto-generation of the ML application and auto-validation of the generated model for better prediction accuracy.

Keywords End-User Programming · Graphical Flows · Graphical Programming Tools · Machine learning as a service (MLaaS) · Machine-Learning-Platform-as-a-Service (ML PaaS) · Machine Learning Pipelines

Tanmaya Mahapatra · Syeeda Nilofer Banoo
Lehrstuhl für Software und Systems Engineering, Fakultät für Informatik,
Technische Universität München, Boltzmannstraße 03, 85748 Garching b. München
E-mail: tanmaya.mahapatra@tum.de
Tanmaya Mahapatra
Birla Institute of Technology and Science, Department of Computer Science and Information
Systems, Vidya Vihar, Pilani - 333031, India.
E-mail: tanmaya.mahapatra@pilani.bits-pilani.ac.in

1 Introduction

Machine Learning (ML) is a scientific discipline that develops and makes use of a particular class of algorithms which are designed to solve problems without explicit programming [42]. The algorithms infer about patterns present in a dataset and learn how to solve a specific problem. This self-learning technique of computer systems has gained prominence and has vast application in the current era. The massive influx of data from the Internet and other sources creates a large bed of structured as well as unstructured datasets, where ML techniques can be leveraged to make meaningful correlations and automated decision-making. Nevertheless, ML techniques used by domain experts like traffic engineers or molecular biologists who are less-skilled programmers have to counter a steep learning curve, i.e. learn how to program and write a ML application from scratch using general-purpose, high-level languages like Java, Scala or Python. The learning curve hinders the widespread adoption of ML by researchers unless they are well-trained in programming. In response to this, we propose to equip less-skilled programmers who are domain-experts with graphical tools. In particular, we intend to support graphical specification of ML programs via flow-based programming paradigm and support auto-generation of target code, thereby shielding the user of such tools from nuances and complexities of programming. Such, graphical flow-based programming tools called mashup tools have been extensively used to simplify application development [10].

1.1 Contributions

Succinctly, the paper contributes to these aspects via:

1. We take the Java APIs of Spark ML operating on DataFrame [24], a popular ML library of Apache Spark [40,41,39], model them as composable components. Every component abstracts one or more underlying APIs such that they represent one unit of processing step in a ML application. The different parameters accepted by the underlying APIs are made available on the front-end while using a specific component to support easy parametrisation.
2. Development of a conceptual approach to parse a ML flow created by connecting several such components from step 1. The parsing ensures that the components are connected in an acceptable positional hierarchy such that it would generate target code which is compilable. The parsed user-flow is used to generate target ML code using principles of Model-Driven Software Development (MDSD). Model to text transformation is used, especially API based code generation techniques [32], to transform the graphical model to target code.
3. The conceptual approach is validated by designing three ML use-cases involving prediction using decision trees, anomaly detection with K-means clustering, and collaborative filtering techniques to develop a music recommender application. The use-cases demonstrate how such flows can be created by connecting different components from step 1 at a higher level of abstraction, parameters to various components can be configured with ease, automatic parsing of the user flow to give feedback to the user if a component has been used in a wrong position in a flow and finally automatic generation of ML application without

the end-user having to write any code. The user can split the initial dataset into training and testing datasets, specify a range for different model parameters for the system to iteratively generate models and test them till a model is produced with higher prediction accuracy.

1.2 Outline

The rest of the paper is structured in the following way: Section 2 summarizes the background while Section 3 discusses the related work. We give an overview of Spark and its machine learning library, i.e. Spark ML, different kinds of data transformation APIs available in Spark and our design choices to support only specific kind of APIs in Section 4. Section 5 describes our conceptual approach to support graphical flow-based ML programming at a higher level of abstraction involving modelling of APIs as components, flow-parsing and target code generation while Section 6 describes its realization. Section 7 validates the conceptual approach in three concrete use-cases. We compare our conceptual approach with existing works in Section 8 which is shortly followed by concluding remarks in Section 9.

2 Background

2.1 Machine Learning

In classical instruction-based programming, a machine processes datasets based on predefined rules. However, in ML programming machines are trained on large datasets to discover the inherent pattern and thereby auto-create the data processing rules. ML programming involves the usage of an input *dataset*, *features* or pieces of information in the dataset useful for problem-solving and a *ML model*. The features when passed to a ML algorithm for learning purposes outputs a ML model. ML algorithms have been broadly classified into two categories of *supervised* and *unsupervised* learning algorithms. In supervised learning, the system learns from training dataset with labelled data to predict future events. Examples include *classification* and *regression*. Regression provides continuous prediction while classification does non-continuous prediction. Unsupervised machine learning algorithms work on unlabelled data to find unknown patterns in the data. Unsupervised learning outputs data as clusters. Clusters are a grouping of similar data from the unlabelled input. Creation of a ML model involves specific steps like data collection, data processing, choosing a relevant ML algorithm, training the model on the training dataset, evaluating the model on testing dataset for prediction accuracy, tuning the parameters to enhance the prediction accuracy of the model and finally using the model to make predictions on new input data.

ML is often confused with Deep Learning (DL) [18]. The necessary steps involved to create a model are the same for both ML and DL. Nevertheless, there are many subtle differences between the two. First, we rely on a single algorithm in ML to predict while DL uses multiple ML algorithms in a layered network to make predictions. Second, ML typically works with structured data, while DL can work with unstructured data too. Third, in ML, manual intervention is required in

the form of model parameter tuning to improve prediction accuracy while DL self-improves to minimise error and increase prediction accuracy. Fourth, DL is suited well with the availability of massive amounts of data and is a more sophisticated technique in comparison to ML.

2.2 Machine Learning libraries

There are a plethora of ML libraries available, including TensorFlow [3,2], PyTorch [29], FlinkML [7], SparkML and scikit-learn [30], among others. TensorFlow has become one of the most prominent libraries for both ML as well as DL. It provides flexible APIs for different programming languages with support for easy creation of models by abstracting low-level details. PyTorch is another open-source ML library developed by Facebook. It is based on Torch [8], an open-source ML library used for scientific computation. This library provides several algorithms for DL applications like natural language processing and computer vision, among others via Python APIs. Similarly, Scikit-learn is an open-source ML framework based on SciPy [36], which includes lots of packages for scientific computing in Python. The framework has excellent support for traditional ML applications like classification, clustering, dimensionality reduction. Open-source computer vision (OpenCV) is a library providing ML algorithms and mainly used in the field of computer vision [9]. OpenCV is implemented in C++. However, it provides APIs in other languages like Java, Python, Haskell and many more. Apache Flink, the popular distributed stream processing platform provides ML APIs in the form of a library called FlinkML. The application designed using these APIs will run inside the Flink execution environment. Another prominent open-source library is Weka. These are some of the most widely used libraries and listing all the available all the ML libraries is beyond the scope of this paper. We, therefore, invite interested readers to refer to [27] for more comprehensive information about ML and DL libraries.

2.3 Flow-Based Programming

Flow-Based Programming (FBP) is a programming paradigm invented by J. Paul Rodker Morrison in the late 1960s [26]. It is an approach to develop applications where program steps communicate with each other by transmitting streams of data. The data flow is unidirectional. At one point in time, only one process can work on data. Each process/component is independent of others, and hence many flows of operation can be generated with different combinations of the components. The components are responsible only for the input data that they consume. Therefore, the input/output data formats are part of the specifications of the components in FBP. The components act as software black boxes and are as loosely coupled as possible in the FBP application which provides the flexibility to add new features without affecting the existing network of FBP components.

2.4 Model-Driven Software Development

MDSD abstracts away the domain-specific implementation from the design of the software systems [32]. The levels of abstraction in a model-driven approach help to communicate the design, scope, and intent of the software system to a broader audience, which increases the quality of the system overall. The models in MDSD are the abstract representation of real-world things that need to be understood before building a system. These models are transformed into platform-specific implementation through domain modelling languages. The MDSD can be compared to the transformation of a high-level programming language to machine code. MDSD often involves transforming the model into text which is popularly known as *code generation*. There are different kinds of code-generation techniques widely like templates and filtering, templates and meta-model, code weaving and API-based code generation among others [32]. API-based code generators are the most simple and the most popular. These simply provide an API with which the elements of the target platform or language can be generated. They are dependent on the abstract syntax of the target language and are always tied to that language. To generate target code in a new language we need new APIs working on the abstract syntax of the new target language.

3 Related work

Literature hardly indicated any significant research work done to support graphical ML programming at a higher level of abstraction and simultaneously explaining programming concepts necessary for such. Nevertheless, there are a number of relevant works in the literature as well as products in the market which support high-level ML programming like WEKA [16], Azure Machine Learning Studio [38], KNIME [5], Orange [11], BigML [6], mljar [25], RapidMiner [17], Streamanalytix [33], Lemonade [31] and Streamsets [34] among others. Out of these only Streamanalytix, Lemonade and Streamsets specifically deal with Spark ML. We compare our conceptual approach with these solutions in Section 9.

4 Apache Spark

4.1 APIs

A Resilient Distributed Dataset (RDD), an immutable distributed dataset [39], is the primary data abstraction in Spark. To manipulate data stored within the Spark runtime environment, Spark provides two kinds of APIs. First, is a direct coarse-grained transformation applied on RDDs directly using function handlers like a map, filter or groupby, among others. This involves writing custom low-level data transformation functions and invoking them via the handler functions. To simplify this kind of data processing, Spark introduced a layer of abstraction on top of RDD called a DataFrame [4]. A DataFrame is essentially a table or two-dimensional array-like structure with named columns. DataFrames can be built from existing RDDs, external databases or tables. RDDs can work on both structured and unstructured data while DataFrames strictly require either structured

or semi-structured data. It becomes easy to process data using named columns than directly working with data. The second set of APIs is the DataFrame based APIs which work on DataFrames and perform data transformation. These APIs take one or more parameters for fine-tuning their operation. A further improvement over the DataFrame based APIs is the Dataset APIs of Spark which are strictly typed in comparison to the untyped APIs of DataFrames.

The RDD-based APIs operating at low-level provide fine-grained control over data transformation. However, DataFrame and DataSet APIs offer a higher level of abstraction in comparison to RDD-based APIs. The APIs are domain-specific and developers do not have to write custom data transformation functions to use them.

4.2 ML with Spark

Spark supports distributed machine learning via:

Spark MLlib Spark MLlib, has been built on top of Spark Core using the RDD abstractions, offers a wide variety of machine learning and statistical algorithms. It supports various supervised, unsupervised and recommendation algorithms. Supervised learning algorithms include decision trees, random forest etc., while some of the unsupervised learning algorithms supported are k-means clustering, support vector machine etc.

Spark ML Spark ML is the successor of Spark MLlib and has been built on top of Spark SQL using the DataFrame abstraction. It offers Pipeline APIs for easy development, persistence and deployment of models. Practical machine learning scenarios involve different stages with each stage consuming data from preceding stage and producing data for the succeeding stage. Operational stages include transforming data into appropriate format required by the algorithm, converting categorical features into continuous features etc. Each operation involves invoking declarative APIs which transform DataFrame based on user inputs [24] and produce a new DataFrame for use in the next operation. Hence, a Spark ML pipeline is a sequence of stages consisting of either a *transformer* or an *estimator*. A transformer transforms one DataFrame to another often via operations like adding or modifying columns in the input DataFrame while an estimator is used to train on the input data and output a transformer which can be a ML model.

4.3 Design choice: API selection

We have chosen DataFrame based APIs for supporting flow-based ML programming because the RDD based APIs operate at low-level and necessitate the usage of custom written functions for bringing data transformations. This makes it challenging to come up with a generic method to parse a flow and ascertain if the components are connected in a way which leads to the generation of always-compilable target code. Additionally, it would demand strict type checking to be introduced at the tool-level to facilitate the creation of such flows. The other kind of API is the Dataset based APIs. These detect syntax as well as analytical errors

during compile time. Nonetheless, these are too restrictive and would render the conceptual approach described in this manuscript non-generic in nature and tied to specific use-cases only. In comparison to these, the DataFrame based APIs are untyped APIs which provide columnar access to the underlying datasets. These are easy to use domain-specific APIs and detect syntax errors at compile time. The analytical errors which might creep in during usage of such APIs can be avoided by designing checks to ensure that the named columns which a specific API is trying to access are indeed received in its input. Additionally, the Spark ML library (version 2.4.4) has been updated to use DataFrame APIs. Hence, it is an ideal choice to use DataFrame APIs over RDD APIs.

5 Conceptual approach

The conceptual approach to enable non-programmers easily specify a ML program graphically at a higher-level of abstraction makes use of principles of FBP and MDSD. It consists of several distinct steps as discussed below:

5.1 Design of modular components

A ML flow consists of a set of connected components. Hence, we need to design the foundational constituent or components first to realise flow-based ML programming. A component mostly does one data processing step in ML model creation and internally abstracts a specific API for delivering that functionality. However, modelling every single Spark ML API as a different component would defeat the very purpose of abstraction. With such a design, the components would have a one-to-one correspondence with the underlying APIs and programming syntax of the ML framework making it harder for less-skilled programmers to comprehend. Hence, we introduce our first design choice by grouping several APIs as one component such that the component represents a data processing operation at high-level and is understandable to end-users. Moreover, the abstracted APIs are invoked in an ordered fashion within the component to deliver the functionality. The different parameters accepted by the APIs used inside a component are made available on the front-end as component property which the user can configure to fine-tune the operation of the component. The essential parameters of the APIs are initialised with acceptable default settings unless overridden by the user of the component.

The second design choice is making the components as loosely coupled to each other possible and the achieving tight functional cohesion between the APIs used within a single component. This leaves space for a future extension where we can introduce new Spark ML APIs as components without interfering with the existing pool of components. To achieve this, we clearly define the input and output interface of every component and define positional hierarchy rules for each one of them. These rules help to decide whether a component can be used in a specific position in a flow or not. The flow-validation step (discussed in Section 5.2) checks this to ensure that the order in which the components are connected or rather the sequence in which the APIs are invoked would not lead to compile-time errors.

The third design choice introduces additional abstraction by hiding away parts of the ML application which are necessary for it to compile and run, written by

developers when coding from scratch, nevertheless, which do not directly correspond to the data processing logic of the application. An example would be the code responsible for initialising the Spark session or closing it within which the remaining data processing APIs get invoked. Another example can be configuring the Spark session like setting the application name, configuring the running environment mode (local or cluster), specifying the driver memory size and providing the binding address among many others. We handle these aspects at the back-end to enable the end-user of such a tool to focus solely on the business logic or data processing logic of the application. The code-generator, running at the back-end (discussed in Section 5.3), is responsible for adding such crucial parts and initialising required settings with sensible defaults to the final code to make it compilable. Nevertheless, the default settings can be overridden by the user from the front-end. For example, the “Start” component (discussed in Section 5.2) is a special component used by the user to mark the start of the flow which can be configured to fine-tune and explicitly override different property values of the Spark session as discussed above.

5.2 Flow specification and Flow-checking

As a second step, we take the components (discussed in Section 5.1) and make them available to the end-user via a programming tool. The programming tool must have an interactive graphical user interface consisting of a palette, a drawing area called canvas, a property pane and a message pane. The palette contains all available modular components which can be composed in a flow conforming to some standard flow composition rules. The actual flow composition takes place on the canvas when the user drags a component from the palette and places it on the canvas. The component, when present on the canvas and when selected, should display all its configurable properties in the property pane. The user can override default settings and provide custom settings for the operation of a component via this pane. The flow, while being composed on the canvas, is captured, converted into a directed acyclic graph (DAG) and checked for the correct order of the connected components. The flow-checking should be done whenever there is a state change. A state change occurs whenever something changes on the canvas. For example, a new component is dragged onto the canvas from the palette or connection between two already present components change, among other such change possibilities. Flow-checking checks the user flow for any potential irreconcilability with the compositional rules. Such a flow when passed to the back-end for code generation will generate target code which does not produce any compile-time errors. The components are *composable* in the form of a flow if they adhere to the following compositional rules:

1. A flow is a DAG consisting of a series of connected components.
2. It starts with a particular component called the “Start” component and ends with a specific component called the “Save Model” component.
3. Every component has its input and output interface adequately defined, and a component can be allowed to be used in a specific position in a flow if it is compatible with the output of its immediate predecessor and if it is permitted at that stage of processing. For example, the application of the ML algorithm

is only possible after feature extraction. Hence, the component corresponding to the ML algorithm and evaluation must be connected after the feature extraction component.

When the flow is marked complete and flow-checking has completed successfully, it is passed to the back-end to start the code generation process. Fig. 1 and Fig. 2 depict the flow-checking sequences for a typical ML flow leading to passing and failure of the process respectively. In the figure, it is assumed that the flow-checking begins after the flow is complete to make it easier for illustration purposes. Nonetheless, it is done whenever there is a change detected on the canvas.

Fig. 1: Flow-checking sequences for a valid ML flow

Fig. 2: Flow-checking sequences for an invalid ML flow

5.3 Model generation

The third step deals with the parsing of the user flow and generating the target code. The back-end consists of a parser, API library and a code generator. The parser is responsible for parsing the flow received and represent it in an intermediate representation, typically in the form of a DAG. Next, it traverses the DAG to ascertain the type of component used and checks in the library for their corresponding method implementations. In MDSD, there are many code generation techniques, but we have used the API-based code generation technique because it is simple to use and serves our purpose well. The only downside is that in API-based code generation technique, the API can generate code only for a specific platform or language as it inherently depends on the abstract syntax of the target language to function. In our case, the API-base code generation technique is restricted to generate Java code. The method implementations for every component in the library contain statements or specific APIs to generate the particular Spark API, which is represented by the component on the front-end. The DAG or intermediate flow representation along with the information of which method to invoke from the library for each vertex is passed to the code generator. The code generator has extra APIs to generate the necessary but abstracted portion of the ML code like the start of a spark session and inclusion of required libraries. Then, it calls the invokes the specific method implementation from the library for each vertex, i.e. invokes the APIs contained within them to generate the target Java Spark API to be used in the final target code. When the code generator does this for all the vertices of the DAG, the target Spark ML code in Java is generated. The final code is compiled and packaged to create the runnable Spark ML driver program. The driver program is the artefact which sent to a Spark cluster for execution.

5.4 Model evaluation and hyperparameter tuning

The final step is used to test the generated model on the testing dataset to check for prediction accuracy or performance of the model. The user should have specified different model parameters to tune during flow specification. If yes, then the code generator generates different ML models, each pertaining to a unique model parameter value. All the models are tested, and the best performant ML model is saved for final usage. For example, if a user is designing a ML model using k-means algorithm and has supplied a range of values for k as a parameter while specifying the flow, where k is the number of clusters user wants to create from the input dataset. In this case, the code generator generates different k-means models for the range of the k values and evaluates the compute cost for each k-means model. Compute cost is one of the ways to evaluate k-means clustering using Spark ML. The system displays the k value and the corresponding model evaluation score in ascending order for the user to select or go with the best performant model. The goal of the hyper-parameter tuning in case of k-means algorithm is to find the optimal k value for the given dataset. If the user provides only one k value, then the code generator generates only one model.

Fig. 3 illustrates the conceptual approach to generate ML code after the steps of user flow specification and flow-checking have been completed. Step 1 in the

figure corresponds to ideas discussed in Section 5.2 while steps 2-5 correspond to Section 5.3. Finally, step 6 corresponds to ideas discussed in Section 5.4.

Fig. 3: Generation of a ML model from a graphical user flow

6 Realisation

In this section, we describe the realisation of the conceptual approach described in Section 5 in the form of a graphical programming tool. In particular, we describe the functioning of the tool and its architecture.

Fig. 4: Interactive GUI of the prototype with palette and canvas

6.1 Prototype Overview

We have implemented a graphical programming tool with Spring Boot [37]. The application comes with a minimalistic interactive graphical user interface (GUI) built using Angular [14,15,12], a popular open-source web application framework. Fig. 4 shows the GUI of the prototype with its palette containing available ML components. It also has a canvas where the components are dragged and connected in the form of a flow and finally, a button to start the code generation process. Whenever a ML component is dragged onto the canvas, a property pane opens up allowing the user to configure its various parameters to fine-tune its operation by overriding the default values. Additionally, when something changes on the canvas, the entire flow is checked to ensure its correctness, and in case of any error, the user is provided feedback on the GUI. Fig. 5 depicts these aspects of the prototype. On press of the button, the flow configuration is captured in a JavaScript Object Notation (JSON) and send to the Spring Boot back-end via REST APIs. It is converted into a Data Transfer Object (DTO) before processing. Fowler describes such a DTO as “an object that carries data between processes to reduce the number of method calls.” [13]. At the back-end, the user flow is parsed, converted into an intermediate representation and passed to the code generator. The code generator auto-generates the target Spark ML application with the help of JavaPoet, an API-based Java code generator [1]. We have provided three video files demonstrating use case 7.1, use case 7.2 and use case 7.3 respectively as supplementary material with this work.

Fig. 5: Component property configuration and flow-checking with user feedback

6.2 Components

Components are the essential constituents of a flow. We have implemented around seven components relevant to our use cases out of which some of them are generic

and can be used in all the three use cases. A typical ML flow consists of four components. The first component marks the start of the flow and is responsible for the generating code related to Spark session initialisation. The second component is used to specify and pre-process the input data. It deals with dimensionality reduction, i.e. extracting a reduced set of non-redundant information from the dataset which can be fed to a ML algorithm and is sufficient to solve a given problem. The third component is a very specific training module which houses a specific ML algorithm like decision tree, k-means, among others. The fourth component trains the model, supports hypertuning of model parameters and saves the final selected model. Table 1 lists the prototyped components and summarises their functionalities.

Table 1: Supported graphical ML components with their functionality

SN.	Component Name	Functionality
1.	Start	Marks the start of the flow and help create a Spark session.
2.	Feature Extraction	Creates DataFrame from input data depending on the features selected.
3.	Feature Extraction From Text File	Transforms the input text data to DataFrame(s).
4.	Decision Tree	Processes the DataFrames to train a decision tree model and evaluates it.
5.	KMeans Clustering	Creates a k-means model and evaluates it. It works on given hyperparameters to find the most efficient model with minimum error in prediction.
6.	Collaborative Filtering	Creates a collaborative filtering model and evaluates it.
7.	Save Model	Marks the end of the flow and saves the ML model in a specified file path.

6.3 Working

In the back-end, the flow is parsed and represented in the form of a DAG, an intermediate representation format. It traverses all the vertices of the DAG to find out which components map to a specific module in an API library. A module in API library contains API statements written using JavaPoet which on execution generates target language statements. Collection of such statements from all such modules in the library initialised with user parameters, which the code-generator assembles, results in the final target code. This code is then compiled and packaged to produce a runnable application. In this entire process, the prototype has many components which interact and accomplish certain functionalities. The system architecture of the prototype with its major components and their interactions

has been summarised in the form of an Identify, Down, Aid, and Role (IDAR) graph in Fig. 6.

Fig. 6: IDAR

IDAR graphs are more simple and intuitive to comprehend about the system structure, hierarchy and communication between components than a traditional Unified Modelling Language (UML) diagram [28]. In an IDAR graph, the objects at higher-level control the objects situated below. Communication is either via a downstream command message (control messages) or an upstream non-command message, which is called notice. Separate subsystems are denoted by hexagonal boxes. Similarly, a dotted line ending with an arrow depicts the data flow in the system, and an arrow with a bubble on its tail shows an indirect method call. For a comprehensive understanding of IDAR, we suggest interested readers to refer [28] and [19].

The front-end has been depicted as a separate subsystem where the user creates the flow and it is validated. When the flow is marked as complete, the front-end sends it via a REST API to the back-end controller. The controller is the main component which houses all the REST APIs for interaction with the front-end, invoking other components and coordinating the whole process of code generation. On receiving a flow from the front-end, it invokes the parser which creates a DAG out of it, saves it in the local database, traverses all the vertices to find out which specific modules in the API library must be invoked for code generation. It passes this information back to the controller which invokes the code-generator. The code-generator invokes the necessary modules in the API library in the order of their connection in the user flow. These modules on execution produce target language statements which are then initialised with user supplied parameters. This is done for all the components the user has connected in the flow. All the generated target statements are assembled in the same order in which the modules were invoked which is the final target code. This code is passed back to the controller. Then, the controller invokes the model generator which takes the final code compiles and

packages it into a runnable Spark ML application. If the user had supplied a range of parameters for the ML model, then the code generator invokes the modules in the APIs and initialise the target statement with different set of parameters which leads to the production of a series of final target codes. Accordingly, the model generator compiles the entire set of generated codes to produce a set of runnable Spark applications. All of the generated applications are evaluated on test data for prediction accuracy. The best performing application/ML model is selected from the whole set.

7 Running Examples

In this section, we discuss three running examples/use cases to capture the replicability of the automatic code-generation from graphical Spark flows, which is the quintessence of our conceptual approach. The running examples are based on three ML algorithms, namely- Decision Tree, K-Means Clustering and Collaborative Filtering. The goal is to demonstrate that the end-user is able to create a runnable Spark program based on the ML algorithms mentioned without having to understand the detail of the Spark ML APIs or requiring any programming skills. Nevertheless, the user is expected to know the datasets on which a ML algorithm is to be applied. Additionally, the user is expected to know the label column and the features columns of the dataset.

7.1 Use Case 1: Predicting forest cover with decision trees

The first use case involved creating a ML Spark application based on the Decision tree algorithm. This supervised ML technique splits the input data depending on the model parameters to make a decision. We have used the covtype dataset in this example [35].

The dataset is available online in compressed CSV format. The dataset reports types of forest covering parcels of land in Colorado, USA. The dataset features describe the parcel of land in terms of its elevation, slope, soil type and the known forest type covers. The dataset has 54 features to describe the pieces of land and 581,012 examples. The forest covers have been categorised into seven different cover types. The categorised values range from 1 to 7. The trained model using the 54 features of the dataset and labelled data should learn forest cover type. The data is already structured, and therefore have used it directly as input in our example.

The end-user drags various graphical components from the palette to the canvas and connects them in the form of a flow, as shown in Figure 7. The specifics of the application like the input dataset, label column, ML model parameters are taken as inputs from the user. Part 3A in Figure 7 depicts the parameters required to create a decision tree model. Internally the parameters are specified to make calls to the Spark ML Decision Tree APIs. The flow creation specifications are sent as a JSON object to the back-end system. The decision tree model parameters comprise *impurity*, *depth of the tree* and *max bins*. The model parameters are crucial to the performance of the model. The impurity parameter minimises the probability of misclassification of the decision tree classifier. The Spark ML library

Fig. 7: Graphical flow to create a decision tree model for forest cover prediction

provides two impurity measures for classification, namely, *entropy* and *Gini*. The impurity is set to the estimator by using the *setImpurity* function provided by *DecisionTreeClassifier* Spark API. Similarly, max bins and max depth also play a vital role in finding an optimal decision tree classifier. The end-user can tweak with these settings until the desired outcome is achieved without having to understand the internals of Spark ML APIs or updating the code manually. This process of trying out different model parameters is called hyperparameter tuning. Listing 1 lists the feature extraction function auto-generated for our use case on covtype data. Listing 2 lists the automatic generated code of the decision tree function that includes an estimator and transformer steps of a ML Spark flow.

```

1 public static Dataset<Row> featureExtraction(SparkSession spark,
2 String filePath, String labelColName) {
3     Dataset<Row> df = spark.read()
4     .option("header", false)
5     .option("inferSchema", true).csv(filePath);
6
7     for (String c : df.columns()) {
8         df = df.withColumn(c, df.col(c).cast("double"));
9     }
10    df = df.withColumnRenamed(labelColName, "labelCol");
11    df = df.withColumn("labelCol", df.col("labelCol").minus(1));
12    return df;
13 }

```

Listing 1: Generated feature extraction code for use case 1

```

2 public static DecisionTreeClassificationModel decisionTree(String
3 impurity, int depth,
4 int maxBins, Dataset<Row> elbyuvroqk) {
5     DecisionTreeClassifier pscsjndfqw = new DecisionTreeClassifier()
6     .setLabelCol("labelCol")
7     .setFeaturesCol("features")
8     .setMaxDepth(depth)
9     .setImpurity(impurity)
10    .setMaxBins(maxBins);

```

```

10     DecisionTreeClassificationModel rwljasoxf = pscsjndfqw.fit(
        elbyuvroqk);
12     return rwljasoxf;
    }

```

Listing 2: Automatic generation of decision tree function

7.2 Use Case 2: Anomaly detection with K-means clustering

The second use case deals with the creation of a ML application involving K-Means. K-Means algorithm partitions the data into several clusters. Anomaly detection is often used to detect fraud, unusual behaviour or attack in the network. The unsupervised learning techniques are suitable for this kind of problems as they can learn the pattern. We have used the data set from KDD cup 1999. The data constitutes network packet data. The data contains 38 features, including a label column. We did not need the label data for applying K-Means, and therefore we removed the same in the feature extraction module. The graphical flow to create the Spark K-means application is depicted in Figure 8. The example model parameters necessary to design the model is depicted in part 3A of Figure 8. The data for the use case contained String columns. We have removed those during the feature extraction phase as the KMeans algorithm cannot process them. Additionally, the String columns would cause runtime error as VectorAssembler only supports numeric, boolean and vector types. The code for this operation is auto-generated when the user inputs the String column names in the feature extraction stage of the flow, as depicted in part 2A of Figure 8.

Fig. 8: Graphical flow to create a k-means clustering model for anomaly detection

The TrainModel component takes a range of K values. K is the number of clusters that the user wants to create from the input dataset. The system generates code to create KMeans models for the range of the K values and stores the compute the cost for each K-means model evaluation in an array in ascending order. Compute cost is one of the ways to evaluate K-Means clustering using

Spark ML. If the end-user provides only one K value, the system generates only one code to generate one model. The output of the K-Means Spark application for automatic hyperparameter tuning is shown in part 5 of Figure 8. The figure indicates different K values and their corresponding model evaluation score. The goal of the hyperparameter tuning in the case of the K-Means algorithm is to find the optimal K value for the given dataset. Listing 3 shows code generated by the back-end system for the graphical flow composed on the front-end.

```

public static KMeansModel kMeansClustering(String initMode, int
    lowK, int highK, int maxIter,
2  double distanceThreshold, int step, Dataset<Row> oiosnsutcx) {
    Map<Integer, Double> degklrfhyb = new LinkedHashMap<Integer,
        Double>();
4  Map<Integer, KMeansModel> mnqoxuuyhr = new LinkedHashMap<Integer
        , KMeansModel>();
    for(int iter = lowK; iter <= highK; iter +=step) {
6      KMeans pjtnfmyssi = new KMeans().setFeaturesCol("features")
        .setK(iter)
8      .setInitMode(initMode)
        .setMaxIter(maxIter)
10     .setTol(distanceThreshold)
        .setSeed(new Random().nextLong());
12     KMeansModel okvhgrwqrk = pjtnfmyssi.fit(oiosnsutcx);
        // Evaluate clustering.
14     Double ujyyccvqv = okvhgrwqrk.computeCost(oiosnsutcx);
        degklrfhyb.put(iter, ujyyccvqv);
16     mnqoxuuyhr.put(iter, okvhgrwqrk);
        System.out.println("*****Sum of Squared Errors = "+
            ujyyccvqv);
18     }
    Map<Integer, Double> vqtkeniici = degklrfhyb.entrySet()
20     .stream()
        .sorted(comparingByValue())
22     .collect(toMap(Map.Entry::getKey, Map.Entry::getValue, (e1, e2)
        -> e2,
        LinkedHashMap::new));
24     Integer hwnnwdbhq = vqtkeniici.entrySet().stream().findFirst().
        get().getKey();
    KMeansModel okvhgrwqrk = mnqoxuuyhr.get(hwnnwdbhq);
26     System.out.println("*****Optimum K = "+ hwnnwdbhq);
    System.out.println("*****Error with Optimum K = "+ degklrfhyb
        .get(hwnnwdbhq));
28     return okvhgrwqrk;
}

```

Listing 3: Function generated for applying KMeans algorithm

7.3 Use Case 3: Music recommender application

The third use case demonstrates the application of Spark Collaborative filtering APIs. The use case checks the code generation of loading data from text files and creating a custom “Rating” class from the input text file. The aim is to ensure that the generated application should load the input text file using the “Rating” class, train and evaluate the recommender model according to the end-user preferences. For this use case, we have used data published by Audioscrobbler, the

first music recommendation system for last.fm. The input data includes three text files, `user_artist_data.txt`, `artist_data.txt` and `artist_alias.txt`. The primary data file is `user_artist_data.txt` which contains user id, artist id and play count. The `artist_data.txt` file includes the names of each artist mapped to artist ID. The `artist_alias.txt` contains the map artist ID that is unknown misspellings. The data contains implicit feedback data; it does not contain any direct rating or feedback data from users.

Collaborative filtering (CF) is a model-based recommendation algorithm. The CF algorithm finds the hidden factors or latent factors about the user's preferences from the user's history data. The current data set fits for collaborative filtering application as we don't have any other information about users or artists. These type of data are sparse. The missing entries in the user-artist association matrix are learnt using the alternating least square (ALS) algorithm. At the moment, Spark ML library supports only model-based collaborative filtering. We have used Spark ML ALS estimator to train the recommendation system. The Spark ML collaborative filtering requires the user to develop a Rating Java class for parsing the main input data file while loading the raw data into a DataFrame. The Rating class is responsible for casting the raw data to respective types, and in our case, we also need the implementation to map the misspelt artist IDs to correct IDs. spark.ml package also provides API to set the cold start strategy for NaN (Not a Number) entries in the user data to mitigate the cold start problem. It is quite normal to have missing entries in such data; for example, maybe a user never rated a song, and a model can't learn about the user in the training phase. The dataset used for the evaluation may have entries for users that are missing in the training dataset. These problems are defined as cold start problem in recommender system design.

Fig. 9: Graphical flow to create a collaborative filtering model for music recommendation

The CM ML flow features extraction technique is a bit different from the previous two use cases. The Spark ML implementation requires the system to generate a Rating class file in Spark version 2.4.4. The Rating file adds better

control to the model designing. Listing 4 shows the Rating class auto-generated by the back-end system. Figure 9 depicts the assembling of graphical components required for CF application generation. The CF Spark program generation required FeatureExtractionFromTextFile component.

The user can create an application with auto-hyperparameter tuning through the graphical interface shown in part 3A of Figure 9. All the available Spark ALS API parameters settings have been considered to design the collaborative filtering module in our system. In the given dataset, the preference of the user is inferred from the given data. Hence the implicit preference box should be checked, or in case of sending direct JSON data, the field should contain the true value. Part 3A of Figure 9 depicts how to set the model parameters for CF application. The model parameters are crucial in making a sound recommender system. The user interface allows setting different parameters for training the model.

```

2  @Getter
  @Setter
  public static class Rating implements Serializable {
4      private Integer userId;

6      private Integer playcount;

8      private Integer artistId;

10     public Rating(Integer artistId, Integer playcount, Integer
        userId) {
12         this.artistId = artistId;
        this.playcount = playcount;
        this.userId = userId;
14     }

16     public static Rating parseRating(String str) {
        String[] fields = str.split(" ");
18         if (fields.length != 3) {
            throw new IllegalArgumentException("Each line must contain 3
                fields");
20         };
        Integer artistId = Integer.parseInt(fields[0]);
22         Integer playcount = Integer.parseInt(fields[1]);
        Integer userId = Integer.parseInt(fields[2]);
24         Integer finalArtistData = CollaborativeFiltering.
            artistAliasMap.getOrDefault(artistId, artistId);
        return new Rating(artistId, playcount, userId);
26     }
28 }

```

Listing 4: Rating class for parsing CF input data

The code generation for training the model using ALS is done depending on the input parameters from the enduser. Table 2 shows the available option for the end-user, and Listing 5 shows how the parameters are set for the ALS algorithm in Spark internally.

```

2  //rest of the code
  for( Integer rank: ranks){
    for( Double regParam: regParams) {
4      for( Double alpha: alphas) {

```

Table 2: Supported graphical ML components with their functionality

SN.	Parameter Name	Purpose
1.	numBlocks	parallelizing the computations by partitioning the data. default value is 10.
2.	rank	number of latent factors to be used in the model. default value is 10
3.	maxIter	maximum number of iteration to run the train model. default value is 10
4.	regParam	regularization parameter. default value is 1.0
5.	implicitPrefs	to specify if the data contains implicit or explicit feedback. default value is false, which means explicit feedback
6.	alpha	Only applicable when implicitPrefs is set to true. default value is 1.0
7.	userCol	setting of input data user column name in ALS algorithm
8.	itemCol	setting of input data item column name in ALS algorithm. in our use case, its artistId
9.	ratingCol	setting of input data rating column name in ALS algorithm

```

6      ALS als = new ALS()
      .setMaxIter(10)
      .setAlpha(alpha)
8      .setRegParam(regParam)
      .setImplicitPrefs(true)
10     .setRank(rank)
      .setUserCol("userId")
12     .setItemCol("artistId")
      .setRatingCol("count");
14     ALSModel alsModel = als.fit(training);
      alsModel.setColdStartStrategy("drop");
16     Dataset<Row> predictions = alsModel.transform(test);
      RegressionEvaluator evaluator = new RegressionEvaluator()
18     .setMetricName("rmse") //rmse, mse, mae, r2
      .setLabelCol("count")
20     .setPredictionCol("prediction");
      Double rmse = evaluator.evaluate(predictions);
22     System.out.println("Hyper Params = (" + rank + " " + alpha +
      " " + regParam + ")");
      System.out.println("Root-mean-square error = " + rmse);
24   }
26 }

```

Listing 5: Example of ALS model creation and evaluation

Spark 2.4.4 has four variants of metric for ALS model evaluation, namely rmse (Root Mean Square Error), mse (Mean Square Error), mae (Mean Absolute Error),

r2 (RSquared). The RegressionEvaluator uses one of these metrics to evaluate the model. The setting of the metric to this API is shown in Listing 5, and the output of the model evaluation is shown in part 5 of Figure 9. The Spark application outputs the least error and the corresponding model parameters used for ALS model creation. The output also shows the error calculated for other combination of parameters sent as input by the end-user. The end-user was able to create a runnable Spark application using a collaborative filtering algorithm with the given dataset. The user could customise the recommender system by providing different parameters to the flow components, which did not change or add any new source code to the existing system. The following observations summarise the essence of the running examples:

1. The modular approach of the system development helps add new ML functionality without affecting the existing behaviour of the system. The end-user can design applications from the list of the components provided through REST API interfaces. The components are independent of each other, and interaction between the components happen only through the input/output data to and from the components that help the end-user design Spark applications depending on the problem statement.
2. The flow-based programming approach hides the underlying Spark implementation from the end-users. The end-users do not have to learn Apache Spark ML library or functionality of Spark Data abstractions to implement an ML application using flow-based ML programming. The end-users can customise the Spark application by providing the specifications through the flow components. The end-users should only have a better understanding of the input dataset used for training a model. The flow-based programming makes it easier for users to customise their ML applications by providing specification through the graphical components.

8 Discussion

In this Section, we compare the conceptual approach of our flow-based programming for ML with the existing solutions as introduced in Section 3. The comparison criteria include:

1. *Graphical interface:* The graphical interfaces should generally be intuitive and easy to use for the end-users to navigate through a software application. The graphical programming interface with flow-based programming paradigm and options to customise the automatic code generation of ML application makes an ideal choice for users with less knowledge of data science or ML algorithms. As discussed in Section III, we have seen that almost all existing tools have a flow-based graphical interface implementation for creating ML model, except Rapidminer. However, Rapidminer provides a graphical wizard to initialise the input parameters. In our approach, we have a list of graphical components representing the steps of ML model creation. The drag and drop feature with feedback on the incorrect assembly of the components guides the user for submitting a logical flow to the back-end system. The graphical interface implemented supports customisation of the components and abstracts the underlying technologies used for automatic Spark application generation for the user.

2. *Target Frameworks*: The second criterion is the target frameworks used by the graphical tools that provide a high-level abstraction. The Deep learning studio offers an interface to implement only deep learning models. Both Microsoft Azure and Rapidminer visual platform supports an end to end automated ML application generation, hiding the underlying technology used from the user. However, the Microsoft Azure HDinsight service allows PySpark code snippet to be used through the Jupyter notebook interface of the tool to run the code on a Spark cluster. Lemonade and StreamSets support high-level abstraction of Spark ML to build ML models. Lemonade also uses the Keras platform to generate models for deep learning applications. On the contrary, the Streamanaltix tool uses multiple frameworks to create a ML model, such as Spark MLlib, Spark ML, PMML, TensorFlow, and H2O. In our solution, we support Spark ML as our target framework.
3. *Code generation*: The code generation capability is another exciting feature to compare whether the tool can create a native program in the target framework from the graphical flow created by the end-user or not. StreamSets and Microsoft Azure require customised code snippet from end-user to train the ML model. They don't generate any code, and the models are loaded directly in their environment. Lemonade and Streamanaltix generate native Apache Spark program, while Deep learning studio generates code in Keras from the graphical flow generated by the user. The user can edit the generated code and re-run the application. Our conceptual approach also generates Java source code for the Apache Spark program from the graphical flow created by the end-user.
4. *Code snippet input from user*: It's desirable to have a graphical tool that can support both expert and non-programmer to create ML applications without having to understand the underlying technology. The fourth criterion is to compare if the tool requires code snippets from the user for ML application creation. Mainly, the code snippet is required for some part of the application generation or the customisation of the program. For example, the StreamSets tool provides an extension to add ML feature by writing customise code in Scala or Python to the pipeline for generating the program. Tools like Rapidminer, Lemonade, Deep learning studio, and Streamanaltix do not require any input code snippet from the user to create the ML application program. While Microsoft Azure auto ML feature does not require any code-snippet from the user, it explicitly asks for a code snippet to create models to run in the Spark environment. The conceptual approach described in this manuscript does not require the user to write any code for Spark ML application generation.
5. *Data pre-processing*: As we already know, the performance of the ML model hugely depends on the quality of the input data. The collected data is usually not structured, requiring a bit of processing before applying the ML algorithms to the data. The manual preprocessing of this data is time-consuming and prone to errors. Having a visual pre-processing feature to the ML tool saves much time for the users. All the tools except Deep learning studio and Lemonade have a data pre-processing step that supports data cleansing through the graphical interface. Our conceptual approach does not include data pre-processing, although it does help essential cleansing and feature extraction from the input data.

6. *Ensemble Learning*: The sixth criterion is to compare whether the tools provide the ensemble learning method. Ensemble methods is a machine learning technique that combines several base models to produce one optimal predictive model. The ultimate goal of the machine learning technique is to find the optimum model that best predicts the desired outcome — tools like Streamanlytix, Rapidminer, Microsoft Azure auto ML support the ensemble learning method. This is a limitation in our current approach as it cannot automatically combine several base models to solve a specific use case.

Table 3 summarises the comparison of our conceptual approach with the existing solutions for the criteria discussed above.

Tools	Graphical Interface	Target Framework	Code-snippet required as input	Code generation for ML program	Include Data pre-processing	Ensemble method supported
Deep Learning Studio	Flow-based GUI	Keras	✗	✓	✗	✗
Microsoft Azure ML	Flow-based GUI	Spark ML(Python, R)	✓(✗for auto ML)	✗	✓	✓(auto ML), ✗(for Saprk application)
Streamanlytix	Graphical wizard-based	Spark ML, H2O, PMML	✗	✓	✓	✓
StreamSets	Flow-based GUI	Spark ML(Python & Scala)	✓	✗	✓	✗
Rapidminer	Graphical wizard-based	Unknown	✗	✗	✓	✓
Lemonade	Flow-based GUI	Spark ML (PySpark)	✗	✓	✗	✗
Our Solution	Flow-based GUI	Spark ML(Java)	✗	✓	✗	✗

Table 3: Comparison of existing solutions with our approach to automate ML application creation

Comments about previous attempts

Previously we had attempted to support programming of Spark applications via graphical flow-based programming paradigm [20, 21, 22, 23]. The previous work culminated in the doctoral dissertation of the last author. This work is an extension of the previous work. The main difference lies in the code-generation technique. Previously, we had used the API-based code generation technique to generate only

the basic skeleton of the Spark application. A library called ‘SparFlo’ [23] was developed, which contained generic method implementation of various Spark APIs. Codeweaving was used to invoke these generic method implementations inside the basic skeleton of the target Spark program. This ensured that the SparFlo library when supported by any graphical programming tool, would easily support Spark programming. Nevertheless, any changes or updates in Spark libraries would cause the release of a new version of the SparFlo library containing the latest generic method implementations of the Spark APIs. Hence, in this attempt, we have relied only on the API-based code generation technique, which eliminates our conceptual approach to develop a pre-packaged implementation of all Spark APIs. It also decouples from a specific Spark version as now we can independently parse a Spark version to generate relevant target source code.

9 Conclusion

The field of data science is challenging in many ways. First, the datasets are usually messy, and most of the time, the data scientists go into pre-processing the data and selecting features from the data. Second, ML algorithms apply multiple iterations to the dataset to train the model. The process of preparing a model is computationally expensive and time-consuming. Third, the real-world application of the generated model to the new data, such as fraud detection, recommends that models become part of the production service in real-time. The data scientists engage most of their time in understanding and analysing the data. They would want to try different ML applications and tweak the models to achieve the sought accuracy in data analytics. The modelling of such ML applications adds another level of difficulty. There should be a way to reuse the models while experimenting with the existing designed systems. Apache Spark framework combines both distributed computing with clusters and library to write ML applications on top of it. Nevertheless, writing good Spark programs requires the user to understand the Spark session, data abstractions, and transformations. Moreover, writing independent code for each ML application adds code redundancy. This paper enables end-users to create Spark ML applications by circumventing the tedious task of learning the Spark programming framework through a flow-based programming paradigm. Our main contributions include taking Java APIs of Spark ML operating on DataFrame, a popular ML library of Apache Spark, model them as composable components, and development of a conceptual approach to parse a ML flow created by connecting several such components. The conceptual approach has been validated by designing three ML use-cases involving prediction using decision trees, anomaly detection with K-means clustering, and collaborative filtering techniques to develop a music recommender application. The use-cases demonstrate how easily ML flows can be created graphically by connecting different components at a higher level of abstraction, parameters to various components can be configured with ease, automatic parsing of the user flow to give feedback to the user if a component has been used in a wrong position in a flow and finally automatic generation of ML application without the end-user having to write any code. In addition to this, our work lays the foundation for several future works. This includes data visualisation techniques to make it more promising for the end-users to work on ML problems. Automatic deployment of the ML model with one

click after the training phase would be another extension of our work. Design and implementation of a flow validation mechanism based on input/output of the flow component would make the system more flexible for future changes and generic for all kinds of flow design.

List of abbreviations

API: Application Programming Interface; DAG: Directed Acyclic Graph; DL: Deep Learning; DTO: Data Transfer Object; FBP: Flow-based Programming; GUI: Graphical User Interface; IDAR: Identify, Down, Aid, and Role; JSON: JavaScript Object Notation; MDSD: Model-Driven Software Development; ML: Machine Learning; OpenCV: Open Computer Vision; REST: REpresentational State Transfer; UML: Unified Modelling Language.

Declarations

Ethics approval and consent to participate

Not applicable

Consent for publication

Not applicable

Availability of data and material

Data sharing is not applicable to this article as no datasets were generated or analysed during the current study.

Competing interests

The authors declare that they have no competing interests.

Funding

This work was supported by the German Research Foundation (DFG) and the Technical University of Munich within the Open Access Publishing Funding Programme.

Authors' contributions

TM designed the conceptual approach, selected the use cases and wrote the entire manuscript. SB implemented the conceptual approach and evaluated the work against the use cases.

Acknowledgements

The authors are grateful to all their colleagues who actively provided insights, reviews and comments leading to the materialisation of this work.

Authors' information

SB holds a bachelor degree in computer science from KIIT, India. She recently completed her masters in computer science from the Technical University of Munich (TUM). In her master thesis, she has worked extensively to build tools to prototype Machine Learning applications. TM is an Assistant Professor at the Department of Computer Science and Information Systems at the Birla Institute of Technology and Science, Pilani. He holds a bachelors degree in Computer Engineering from NMIMS Mumbai, an M.Sc. degree in Software Systems Engineering from RWTH Aachen and a PhD degree in Computer Science from the Technical University of Munich. His research interest lies in supporting Big Data Analytics and ML via graphical flow-based programming paradigm to foster widespread adoption of data science.

References

1. JavaPoet. <https://github.com/square/javapoet>. [Online; accessed 18-May-2020]
2. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015). URL <https://www.tensorflow.org/>. Software available from tensorflow.org
3. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., Zheng, X.: Tensorflow: A system for large-scale machine learning. In: 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), pp. 265–283 (2016). URL <https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf>
4. Armbrust, M., Xin, R.S., Lian, C., Huai, Y., Liu, D., Bradley, J.K., Meng, X., Kaftan, T., Franklin, M.J., Ghodsi, A., Zaharia, M.: Spark sql: Relational data processing in spark. In: Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, SIGMOD '15, p. 1383–1394. Association for Computing Machinery, New York, NY, USA (2015). DOI 10.1145/2723372.2742797. URL <https://doi.org/10.1145/2723372.2742797>
5. Berthold, M.R., Cebon, N., Dill, F., Gabriel, T.R., Kötter, T., Meinl, T., Ohl, P., Thiel, K., Wiswedel, B.: Knime - the konstanz information miner: Version 2.0 and beyond. SIGKDD Explor. Newsl. **11**(1), 26–31 (2009). DOI 10.1145/1656274.1656280. URL <https://doi.org/10.1145/1656274.1656280>
6. BigML: Machine Learning that works. <https://static.bigml.com/pdf/BigML-Machine-Learning-Platform.pdf?ver=5b569df> (2020). [Online; accessed 06-June-2020]
7. Carbone, P., Katsifodimos, A., Ewen, S., Markl, V., Haridi, S., Tzoumas, K.: Apache flink™: Stream and batch processing in a single engine. IEEE Data Eng. Bull. **38**, 28–38 (2015)
8. Collobert, R., Bengio, S., Mariéthoz, J.: Torch: a modular machine learning software library. Idiap-RR Idiap-RR-46-2002, IDIAP (2002)

9. Culjak, I., Abram, D., Pribanic, T., Dzapo, H., Cifrek, M.: A brief introduction to opencv. In: 2012 Proceedings of the 35th International Convention MIPRO, pp. 1725–1730 (2012)
10. Daniel, F., Matera, M.: *Mashups: Concepts, Models and Architectures*. Springer Berlin Heidelberg, Berlin, Heidelberg (2014)
11. Demšar, J., Curk, T., Erjavec, A., Črt Gorup, Hočevar, T., Milutinović, M., Možina, M., Polajnar, M., Toplak, M., Starič, A., Štajdohar, M., Umek, L., Žagar, L., Žbontar, J., Žitnik, M., Zupan, B.: Orange: Data mining toolbox in python. *Journal of Machine Learning Research* **14**, 2349–2353 (2013). URL <http://jmlr.org/papers/v14/demsar13a.html>
12. Escott, K.R., Noble, J.: Design patterns for angular hotdraw. In: Proceedings of the 24th European Conference on Pattern Languages of Programs, EuroPLop '19. Association for Computing Machinery, New York, NY, USA (2019). DOI 10.1145/3361149.3361185. URL <https://doi.org/10.1145/3361149.3361185>
13. Fowler, M.: *Patterns of Enterprise Application Architecture*. Addison-Wesley Longman Publishing Co., Inc., USA (2002)
14. Freeman, A.: *Pro Angular 6*, 3rd edn. Apress, USA (2018)
15. Hajian, M.: *Progressive Web Apps with Angular: Create Responsive, Fast and Reliable PWAs Using Angular*, 1st edn. APress (2019)
16. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The weka data mining software: An update. *SIGKDD Explor. Newsl.* **11**(1), 10–18 (2009). DOI 10.1145/1656274.1656278. URL <https://doi.org/10.1145/1656274.1656278>
17. Jannach, D., Jugovac, M., Lerche, L.: Supporting the design of machine learning workflows with a recommendation system. *ACM Trans. Interact. Intell. Syst.* **6**(1) (2016). DOI 10.1145/2852082. URL <https://doi.org/10.1145/2852082>
18. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**, 436–44 (2015). DOI 10.1038/nature14539
19. M. A. Overton: The idar graph. *Commun. ACM* **60**(7), 40–45 (2017). DOI 10.1145/3079970. URL <https://doi.org/10.1145/3079970>
20. Mahapatra, T.: *High-level Graphical Programming for Big Data Applications*. Dissertation, Technische Universität München, München (2019). URL <http://mediatum.ub.tum.de/?id=1524977>
21. Mahapatra, T., Gerostathopoulos, I., Prehofer, C., Gore, S.G.: Graphical Spark Programming in IoT Mashup Tools. In: 2018 Fifth International Conference on Internet of Things: Systems, Management and Security, pp. 163–170 (2018). DOI 10.1109/IoTSMS.2018.8554665. URL <https://doi.org/10.1109/IoTSMS.2018.8554665>
22. Mahapatra, T., Prehofer, C.: aFlux: Graphical flow-based data analytics. *Software Impacts* **2**, 100007 (2019). DOI 10.1016/j.simpa.2019.100007. URL <https://doi.org/10.1016/j.simpa.2019.100007>
23. Mahapatra, T., Prehofer, C.: Graphical Flow-based Spark Programming. *Journal of Big Data* **7**(1), 4 (2020). DOI 10.1186/s40537-019-0273-5. URL <https://doi.org/10.1186/s40537-019-0273-5>
24. Meng, X., Bradley, J., Yavuz, B., Sparks, E., Venkataraman, S., Liu, D., Freeman, J., Tsai, D., Amde, M., Owen, S., Xin, D., Xin, R., Franklin, M.J., Zadeh, R., Zaharia, M., Talwalkar, A.: Mlib: Machine learning in apache spark. *J. Mach. Learn. Res.* **17**(1), 1235–1241 (2016)
25. mljar: Machine Learning for Humans! Automated machine learning platform. <https://mljar.com> (2018). [Online; accessed 18-May-2020]
26. Morrison, J.P.: *Flow-Based Programming*, 2nd Edition: A New Approach to Application Development. CreateSpace, Paramount, CA (2010)
27. Nguyen, G., Dlugolinsky, S., Bobák, M., Tran, V., López García, Á., Heredia, I., Malík, P., Hluchý, L.: Machine learning and deep learning frameworks and libraries for large-scale data mining: a survey. *Artificial Intelligence Review* **52**(1), 77–124 (2019). DOI 10.1007/s10462-018-09679-z. URL <https://doi.org/10.1007/s10462-018-09679-z>
28. Overton, M.A.: The idar graph. *Queue* **15**(2), 29–48 (2017). DOI 10.1145/3084693.3089807. URL <https://doi.org/10.1145/3084693.3089807>
29. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (eds.) *Advances in Neural Information Processing Systems* 32, pp.

- 8026–8037. Curran Associates, Inc. (2019). URL <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
30. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.* **12**(null), 2825–2830 (2011)
 31. Santos, W.d., Avelar, G.P., Ribeiro, M.H., Guedes, D., Meira Jr., W.: Scalable and efficient data analytics and mining with lemonade. *Proc. VLDB Endow.* **11**(12), 2070–2073 (2018). DOI 10.14778/3229863.3236262. URL <https://doi.org/10.14778/3229863.3236262>
 32. Stahl, T., Völter, M., Czarnecki, K.: *Model-Driven Software Development: Technology, Engineering, Management*. John Wiley & Sons, Inc., Hoboken, NJ, USA (2006)
 33. StreamAnalytix: Self-Service Data Flow and Analytics For Apache Spark. <https://www.streamanalytix.com> (2018). [Online; accessed 18-May-2020]
 34. StreamSets: DataOps for Modern Data Integration. <https://streamsets.com> (2018). [Online; accessed 18-May-2020]
 35. University of Irvine: UC Irvine Machine Learning Repository. <https://archive.ics.uci.edu/ml/machine-learning-databases/covtype/>
 36. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., Vijaykumar, A., Bardelli, A.P., Rothberg, A., Hilboll, A., Kloeckner, A., Scopatz, A., Lee, A., Rokem, A., Woods, C.N., Fulton, C., Masson, C., Häggström, C., Fitzgerald, C., Nicholson, D.A., Hagen, D.R., Pasechnik, D.V., Olivetti, E., Martin, E., Wieser, E., Silva, F., Lenders, F., Wilhelm, F., Young, G., Price, G.A., Ingold, G.L., Allen, G.E., Lee, G.R., Audren, H., Probst, I., Dietrich, J.P., Silterra, J., Webber, J.T., Slavič, J., Nothman, J., Buchner, J., Kulick, J., Schönberger, J.L., de Miranda Cardoso, J.V., Reimer, J., Harrington, J., Rodríguez, J.L.C., Nunez-Iglesias, J., Kuczynski, J., Tritz, K., Thoma, M., Newville, M., Kümmerer, M., Bolingbroke, M., Tartre, M., Pak, M., Smith, N.J., Nowaczyk, N., Shebanov, N., Pavlyk, O., Brodtkorb, P.A., Lee, P., McGibbon, R.T., Feldbauer, R., Lewis, S., Tygier, S., Sievert, S., Vigna, S., Peterson, S., More, S., Pudlik, T., Oshima, T., Pingel, T.J., Robitaille, T.P., Spura, T., Jones, T.R., Cera, T., Leslie, T., Zito, T., Krauss, T., Upadhyay, U., Halchenko, Y.O., Vázquez-Baeza, Y., 1.0 Contributors, S.: Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature Methods* **17**(3), 261–272 (2020). DOI 10.1038/s41592-019-0686-2. URL <https://doi.org/10.1038/s41592-019-0686-2>
 37. Walls, C.: *Spring Boot in Action*, 1st edn. Manning Publications Co., USA (2016)
 38. Washington, M.: *Azure Machine Learning Studio for The Non-Data Scientist: Learn How to Create Experiments, Operationalize Them Using Excel and Angular .Net Core ... Programs to Improve Predictive Results.*, 1st edn. CreateSpace Independent Publishing Platform, North Charleston, SC, USA (2017)
 39. Zaharia, M., Chowdhury, M., Das, T., Dave, A., Ma, J., McCauley, M., Franklin, M.J., Shenker, S., Stoica, I.: Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In: *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation, NSDI’12*, p. 2. USENIX Association, USA (2012)
 40. Zaharia, M., Chowdhury, M., Franklin, M.J., Shenker, S., Stoica, I.: Spark: Cluster computing with working sets. In: *Proceedings of the 2nd USENIX Conference on Hot Topics in Cloud Computing, HotCloud’10*, p. 10. USENIX Association, USA (2010)
 41. Zaharia, M., Xin, R.S., Wendell, P., Das, T., Armbrust, M., Dave, A., Meng, X., Rosen, J., Venkataraman, S., Franklin, M.J., Ghodsi, A., Gonzalez, J., Shenker, S., Stoica, I.: Apache spark: A unified engine for big data processing. *Commun. ACM* **59**(11), 56–65 (2016). DOI 10.1145/2934664. URL <https://doi.org/10.1145/2934664>
 42. Zecevic, P., Bonaci, M.: *Spark in Action*, 1st edn. Manning Publications Co., USA (2016)