**QUEEN'S
UNIVERSITY
BELFAST**

AI Back-End as a Service for Learning Switching of Mobile Apps between the Fog and the Cloud

Athanasopoulos, D., & Liu, D. (2021). AI Back-End as a Service for Learning Switching of Mobile Apps between the Fog and the Cloud. *IEEE Transactions on Services Computing*. <https://doi.org/10.1109/TSC.2021.3117927>

Published in:

IEEE Transactions on Services Computing

Document Version:

Peer reviewed version

Queen's University Belfast - Research Portal:

[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

Copyright 2021, IEEE.

This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

AI Back-End as a Service for Learning Switching of Mobile Apps between the Fog and the Cloud

Dionysis Athanasopoulos and Dewei Liu

Abstract—Given that cloud servers are usually remotely located from the devices of mobile apps, the end-users of the apps can face delays. The Fog has been introduced to augment the apps with machines located at the network edge close to the end-users. However, edge machines are usually resource constrained. Thus, the execution of online data-analytics on edge machines may not be feasible if the time complexity of the data-analytics algorithm is high. To overcome this, multiple instances of the back-end should be deployed on edge and remote machines. In this case, the research question is how the switching of the app among the instances of the back-end can be dynamically decided based on the response time of the service instances. To answer this, we contribute an AI approach that trains machine-learning models of the response time of service instances. Our approach extends a back-end as a service into an AI self-back-end as a service that self-decides at runtime the right edge/remote instance that achieves the lowest response-time. We evaluate the accuracy and the efficiency of our approach by using real-word machine-learning datasets on an existing auction app.

Index Terms—Mobile back-end as a service, Fog infrastructure, machine learning.

1 INTRODUCTION

Mobile apps that perform online data-analytics are rapidly increasing. Ofcom's Online UK Nation 2020 report states the percentage of users of mobile apps was significantly increased the last year¹. While we have high expectations from apps, mobile devices cannot meet these expectations due to their limited hardware-resources. In this case, the back-end of apps usually run computationally demanding online data-analytics on powerful remote servers by adopting the Web-service technology (a.k.a., mobile back-end as a service²). Data analytics refer to algorithms that mine knowledge from data (e.g., price predictions) [1].

However, the use of remote servers for performing data analytics may lead mobile apps to experience delays in receiving the analysis output. For instance, the network latency is not always predictable [2] and the end-users of mobile apps can face high end-to-end response times. The response time equals to the sum of the execution time of the analysis on a machine plus the latency of the underlying communication links.

The Fog has been introduced to reduce response times of mobile apps [3]. The Fog augments mobile apps with devices located near/at the network edge. In this way, the Fog has arrived to mitigate the latency issues that emerge from the usage of remote machines. Given edge devices are resource constrained compared to powerful remote servers, the execution of data analytics on edge devices is not always feasible for large datasets. To overcome this, the execution of data analytics should combine edge and remote machines³ [1]. In particular, multiple instances of the back-end should be deployed on both edge machines and remote machines.

The data analytics whose time complexity is low enough can be executed on edge machines. On the contrary, data analytics whose time complexity is high cannot be executed on edge machines. Thus, the *dynamic switching* of the front-end (e.g., GUI) among the instances of the back-end should be decided. However, the dynamic decision of the right instance is not straightforward because it depends on the dataset provided as input to the data-analysis algorithm. Thus, the following research question is raised:

“How the switching among the instances of a back-end can be decided based on the response time of the back-end?”

To answer this, the response time of a service back-end should get first calculated. However, data analytics implement complex algorithms that contain many interconnected data-structures whose response time depends on the particulars of the algorithms' utility (a.k.a., the size of input datasets). Moreover, the response time of a back-end further depends on the machines where the service instances have been deployed. Overall, we focus in this work on the following empirical factors that affect the response time of service instances: input datasets and underlying machines⁴.

The problem is that the response time of service instances cannot be necessarily estimated at the development time of mobile apps because the input datasets of the back-end and the used machines may have not been decided at the development time of the apps. Thus, the response time of service instances should be *predicted at runtime* based on the input datasets and the available edge/remote machines. If the predicted response-time is low, then the app should get bound to an edge instance of the back-end. On the contrary, if the predicted response-time is high, then the app should get bound to a remote instance of the back-end.

We meet approaches in the literature that estimate the response time of services, without analysing the source code of the service

• Dionysis Athanasopoulos and Dewei Liu are with the School of Electronics, Electrical Engineering & Computer Science, Queen's University of Belfast, Northern Ireland, UK. E-mails: {D.Athanasopoulos,dliu08}@qub.ac.uk

Manuscript received MONTH XX, 20XX; revised MONTH XX, 20XX.

1. https://www.ofcom.org.uk/_/data/assets/pdf/_file/0027/196407/online-nation-2020-report.pdf

2. <https://azure.microsoft.com/en-us/solutions/mobile>

3. We do not consider in this work the financial debt of hiring machines.

4. Please note we focus on the response time of service instances in this work (i.e., the round-trip time from the service client to the service instance) because we assume we do not have access to the service implementation and to the resources (e.g., CPU, RAM) consumed on the used machines. In this way, we cannot calculate performance metrics (like CPU, memory consumption) that are related to the performance of the service instances on the machines.

implementation. These approaches estimate the response time via analysing the Web API of a service (e.g., [4]). The restriction of the existing approaches is that they build mathematical expressions of the response time in a manner *tied to the independent variables* of the expressions. The independent variables usually include input/output data-types of Web APIs and/or the type of edge/remote machines. Given that the construction manner of a mathematical expression is tied to its independent variables, a new mechanism should be developed for different independent variables. Thus, we face the *challenge* in proposing a generic approach that builds expressions without being tied to independent variables.

To address this challenge, we contribute an artificial intelligence (AI) approach that takes any independent variable as input and constructs machine-learning models that capture the response time of service instances. We specify the architectural design and the algorithmic mechanisms of the AI approach. The approach is applied to mobile apps that follow the architectural pattern of the mobile back-end as a service. These apps usually include at least three interconnected components: the front-end, the controller, and the back-end. Our approach extends a back-end as a service into an AI back-end as a service. The AI back-end encapsulates the available edge/remote instances of the service back-end. The AI back-end further constructs machine-learning regression expressions of the response time of the edge instances and the remote instances of the service back-end.

The machine-learning technique implemented by the AI back-end has been designed in a generic way to use multiple machine-learning models (e.g., k-nearest neighbours, support vector machine). The AI back-end uses the trained machine-learning expressions to self-decide at runtime the edge/remote instance that achieves the lowest response-time. In other words, our approach further extends the AI back-end to an AI self-back-end.

We conduct a set of experiments where multiple machine-learning models were implemented. We compare and contrast the usage of four machine-learning models (neural networks, k-nearest neighbours, support vector machines, decision trees). We evaluate the accuracy of our approach via measuring the percentage of the correct switching-decisions made by the AI self-back-end. We further evaluate the efficiency of the approach via measuring the improvement in the response time of the app on all the switching decisions. We use in our experiments real-world datasets collected from UC Irvine machine-learning repository⁵. We also use a real-world auction mobile app as a case-study.

The rest of the paper is structured as follows. Section 2 categorises and compares the related approaches. Section 3 defines the conceptual model of mobile apps with back-end as a service. Section 4 describes the case study that we use as a running example in the paper. Section 5 defines the Fog infrastructure that we adopt in this work. Section 6 specifies the architecture of the AI self-back-end as a service. Section 7 details the mechanisms of the AI self-back-end. Section 8 presents the experimental evaluation of our approach. Finally, Section 9 summarises our contribution and discusses the future research-directions.

2 RELATED WORK

The existing approaches for executing mobile apps via the Fog generate deployment/execution plans to map instances of back-ends to edge/remote machines. These plans aim at meeting the

non-functional requirements for apps and/or at considering the runtime performance of apps and of the underlying Fog infrastructure. For instance, back-ends are mapped to machines only if the requirements for the back-ends (e.g., CPU requirements) can be met by the machine characteristics (e.g., CPU cores).

We categorise the existing approaches in two dimensions: (i) non-functional requirements vs. runtime performance; (ii) deployment time vs. runtime of apps. We form a categorisation matrix of four quadrants, Q_1 - Q_4 (Table 1), that are detailed in Sections 2.1.1-2.2.2. Our approach falls into Q_3 . We finally provide in Section 2.3 an overall comparison of the existing approaches against our approach to highlight the research gap in the literature.

2.1 Deployment Time of Mobile Apps

2.1.1 Non-Functional Requirements for Mobile Apps

[18] proposes an automated mechanism for generating deployment plans that meet the requirements for the latency and the bandwidth of communication links. [18] proposes a heuristic to get approximate solutions because the mapping of many back-ends to many machines is NP-hard (subgraph isomorphism problem). [19] proposes an automated mechanism for generating deployment plans that minimise the delay of a Fog infrastructure to serve requests to apps. To approximate optimal deployment plans, [19] adopts a mixed integer non-linear programming technique.

2.1.2 Runtime Performance of Fog Infrastructure

[9], [20] generates deployment plans that reduce the power consumption of apps on edge machines. [19] produces deployment plans that consider both the service delay of a Fog infrastructure and the power consumption of apps on edge/remote machines. [19] approximates optimal deployment plans by balancing the service time and the power consumption. [21] allocates service instances to edge machines to optimize the consumption of computation and bandwidth resources on the machines.

2.2 Runtime of Mobile Apps

2.2.1 Non-Functional Requirements for Mobile Apps

[5] proposes an approach for the migration a back-end from an edge machine to another if the response time of the back-end exceeds a time threshold. [11] presents an approach for the machine-learning modelling of the runtime performance of apps in function of (software or hardware) control knobs and the number of the service requests. [6] models and verifies (using discrete-time Markov chains) rule-based auto-scaling policies for cloud-based apps in function of the discrete elapsed time. [7], [8] propose approaches for modelling (using Kalman filters) the runtime performance of apps in function of the input data, the number of the service requests, and the discrete elapsed time. [4] contributes the conceptual model and the algorithmic mechanisms of an autonomous approach [22]. [4] predicts the response time of a back-end by building predictive mathematical models.

2.2.2 Runtime Performance of Fog Infrastructure

[12] proposes an analytical model for measuring the service delay of a Fog infrastructure in function of the elapsed discrete-time. Based on this model, [12] offloads computation from an edge machine to another in an online manner. [10] proposes an indexing of edge machines with respect to the trade-off that exist between the service delay and the power consumption. [14]

5. <http://archive.ics.uci.edu/ml/index.php>

TABLE 1
Categorization and comparison of the existing approaches that deploy/execute mobile apps via the Fog.

Runtime of apps	<i>Non-functional attribute:</i> latency [5], CPU [6], [7], [8] response time [4], [7], [8], [11] throughput [11] availability [11], reliability [11] <i>Model/Execution plan generation:</i> tree traversal [5] machine learning [4], [11] Markov chains [6] layered queuing models [7], [8]	<i>Performance indicator:</i> power consumption [9], [10] service delay [12], [13], resource consumption [13] computation cost [14], energy consumption [15], [16] fog-to-cloud bandwidth [13], latency [17] <i>Model/Execution plan generation:</i> sub-graph isomorphism [9] graph traversal [12], deep learning [15] indexing data-structure [10] bipartite matching [14], maximum independent set [16] integer non-linear programming [13] discrete time steps [10], [12] incoming traffic rate [13] energy consumption [15], service requests [17] multiple tasks of multiple users [16]
	<i>Model/plan in function of:</i> control knobs [11], input data [4] service requests [7], [8], [11] discrete time steps [6], [7], [8] input data size [7], [8]	<i>Model/plan in function of:</i>
Deployment	<i>Non-functional attribute:</i> latency [18] bandwidth [18], response time [19] subgraph isomorphism [18]	<i>Performance indicator:</i> power consumption [9], [19], [20] computation & bandwidth resources [21] subgraph isomorphism [9], integer programming [19] graph traversal [20], Game theory [21]
	<i>Deployment plan generation:</i>	<i>Deployment plan generation:</i>
Non-functional requirements for apps		Runtime performance of Fog infrastructure
Q ₃		Q ₄
Q ₁		Q ₂

proposes an offloading mechanism for migrating back-ends based on the end-users' budgets. [13] proposes a framework to decrease the service delay of a Fog infrastructure, the Fog-to-Cloud bandwidth, along with the resource usage on edge machines. [15] proposes a scheduling model to reduce the energy consumed by dynamic switching activities between edge/cloud machines. [16] proposes a multi-user multi-task offloading approach for the mobile edge cloud towards minimising the energy consumption on the edge/cloud machines. [17] proposes mechanisms for provisioning and auto-scaling resources on edge machines.

2.3 Comparing the Related Approaches

Only two approaches monitor the app execution to regenerate deployment plans [5], [12]. However, these approaches are reactive (i.e., they suspend the app execution) and lay between apps and operating systems (e.g., redeployment engines). Contrarily, our approach runs at the application layer and pro-actively (without suspending the app execution) self-decides what instance of the back-end will be invoked. Our past work proposed in [4] is a proactive approach too. However, [4] builds mathematical expressions of response time in a manner tied to the independent variables of the expressions. On the contrary, our current work is a generic approach that builds machine-learning models in a manner that is not tied to independent variables.

3 MOBILE BACK-END AS A SERVICE

Mobile apps that follow the architectural pattern of the mobile back-end as a service usually include at least three interconnected components: the front-end, the controller, and the back-end. The front-end of an app constitutes the part of the app that interacts with the end-users of the app (e.g., interaction via using a GUI). The back-end of an app is provided as a service and is mainly responsible for data analysis⁶. Each service exposes a public Web API that is formally defined in Section 3.1. A controller

Fig. 1. Mobile apps with (data-analytics) back-end as a service.

facilitates the interaction between the front-end and the back-end. The controller receives data from the front-end and forwards them for analysis to the back-end via using programming clients of the back-end to invoke service operations of the back-end.

The architecture of mobile apps that we just described is depicted in Fig. 1. In particular, Fig. 1 visualises the architecture of mobile apps by adopting the notation of the UML class-diagram⁷. The UML class-diagram is a graphical notation used to construct and visualise the classes of a system and the relationships among objects of the classes⁸.

3.1 Web API of Service Back-End

The Web APIs of service back-ends can be specified via using the RESTful [23] or the SOAP [24] Web-service technologies. To model the Web API of service back-ends independently of the

⁷ <https://www.omg.org/spec/UML>

⁶ The storage of the back-end, along with the synchronisation of the data stored across multiple back-end instances, are beyond the scope of the paper.

⁸ The arrows/links that we use in the UML diagrams of this paper represent the UML relationships of the association, the aggregation, and the composition. Association exists when a class uses references to object(s) of another class and is visualised by an arrow. Aggregation is a "has a" relationship between a composite class and its owned objects. Aggregation is visualised by using a hollow diamond shape on the composite class. Composition is a "has a" relationship indicating that when the composite object is finalised, its owned objects are also finalised. Composition is visualised by using a filled diamond shape on the composite class. The arrow/link of a UML relationship can be annotated with a number or an expression that corresponds to the multiplicity/cardinality of the relationship. In particular, the expression of "1" denotes that a class uses exactly one object of another class. The expression of "1..*" denotes that a class uses one or many objects of another class.

underlying technology and language, we define the Web API in the following generic way.

Definition 1 (Web API of service back-end). We define the Web API of a service back-end by the following tuple:

$$api = (n, ops, s, c) \quad (1)$$

The tuple comprises:

- the name, n , of the service
- the set, $ops = \{op_i\}$, of the operations of the service that correspond to the public operations of the service
- the technology style, s (RESTful or SOAP), of the Web API
- the credentials, c (e.g., JSON document), usually required by remote servers for authenticating the client identity. $\square$

Definition 2 (Web API operation). We define an operation op of the Web API of a service by the following tuple:

$$op = (n, in, out) \quad (2)$$

The tuple consists of the name, n , of the operation, the set, $in = \{p_i\}$, of the input parameters, and the set, $out = \{p_i\}$, of the output parameters. $\square$

An input/output parameter of an operation of a Web API can be a primitive parameter (Def. 3) or a complex parameter (Def. 4).

Definition 3 (Primitive parameter of an operation). We define a primitive parameter, p , of an operation by the following tuple:

$$p = (n, t) \quad (3)$$

The tuple comprises the name, n , of the parameter and the built-in data-type, t , of the parameter (e.g., `int`). $\square$

Definition 4 (Complex parameter of an operation). We define a complex parameter, p , of an operation by the following tuple:

$$p = (n, g, t) \quad (4)$$

The tuple comprises the name, n , of the parameter, the grouping data-structure, g , of t (e.g., a list), and the data-type of the (possibly absent) nested (primitive/complex) parameter, t . $\square$

4 CASE STUDY

Emma is an avid eBay user, always ready to snap up an auction bargain when she sees one. But Emma faces a difficulty on deciding the best end-price for the auction to maximise her profits. To this end, Emma envisions to use an auction app that can predict end-prices based on past bidding data in a timely manner [25]. Thus, Emma asked by Tom, who is a developer of mobile apps, to build an auction app for her.

Mobile back-end as a service. Tom built an app in which the back-end is deployed as a service on a cloud machine (2.2GHz vCPU, 3.75GB RAM). The data analysis performed by the back-end uses the k-means clustering algorithm to forecast the end-price of an online auction [26]. In particular, the algorithm of the data analysis organises the past auction prices into groups of similar prices and further employs regression trees to predict end-prices.

Architecture of the auction app. The architecture of the auction app is depicted in Fig. 2. According to Fig. 2, the front-end of the app includes two classes: the Bidding class and the Price Estimation class. The Bidding class uses the Price Estimation class to predict the end-price of the current auction.

Fig. 2. The architecture of the auction app with back-end as a service.

Following, the Price Estimation class interacts with the controller of the app and especially, the Price Estimation class interacts with the client of the RESTful Web API of the service back-end of the app. The service client invokes the Web API of the k-means service of the back-end via establishing an HTTP connection to the endpoint of the instance of the k-means service that has been deployed on a cloud machine.

Web API of the auction back-end. Regarding the RESTful Web API of the k-means service, the Web API provides the single operation, `cluster`. The parameters of this operation are the following: the clusters number, k , the clustering iterations, it , and a bidding dataset, d . The dataset is specified as a set of multi-dimensional data-points. In other words, the parameter of a bidding dataset is a complex parameter (Def. 4) that consists of the Java grouping data-structure, `Vector`, and further includes the nested complex parameter, `DataPoint`. The `DataPoint` parameter consists of the array data-structure that contains double numbers (a.k.a., `double[]`). Overall, the Web API of the k-means service (Web API, operation, and parameters) is depicted in Fig. 3 by using the definition of the Web API of a service back-end that was provided in Section 3.1 (Def. 1 – 4). The visualisation of the Web API in Fig. 3 adopts the notation of the UML class-diagram.

Problem description. Running the front-end of the app on a smartphone LG Nexus 5 (2.3 GHz CPU, 2GB RAM⁹) for small datasets (≤ 1000 bidding data), Emma was happy to see that the response times were ranging from 0.15 seconds to 34 seconds. However, running the app for larger datasets (≈ 5000 bidding data), she experienced the response time of ≈ 3.5 minutes. Emma could not accept this response time because Emma wants to determine her bid for a product quickly.

Sketching the solution to problem. Tom observed that the response time of the app was increased due to the network latency and especially the latency of transmitting many bidding data. To avoid this, Tom wants to augment the cloud instance of the back-end with a service instance that is deployed on the Fog (i.e., close to Emma's device). However, it is not clear to Tom how he should implement the *dynamic switching* of the front-end between the instances of the back-end. To help Tom in implementing the dynamic switching, we propose an automated AI-based approach that learns switching decisions between the Fog and the Cloud and dynamically makes switching decisions that achieve the lowest response-time. Before detailing our approach, we first specify below the concept of the Fog infrastructure that we adopt.

5 FOG INFRASTRUCTURE

The Fog infrastructure usually includes three groups of devices [27]. The first group encompasses mobile devices where the front-end of mobile apps has been installed. The second group consists

9. https://www.gsmarena.com/lg_nexus_5-5705.php

Fig. 3. The Web API of the k-means service of the auction back-end.

of edge devices used for deploying instances of the back-end of apps. These devices provide physical machines or virtual machines for the deployment of the service back-end of the apps and are usually more powerful than those of the first group. The devices of the first group and the second group usually communicate to each other through one-hop (or few-hops) wireless communication links. The communication links between the front-end and the edge instances of the back-end of an app have negligible latency. The third group comprises the virtual machines provided by powerful remote (e.g., cloud) devices. Overall, we define the Fog infrastructure that we adopt in this work as follows.

Definition 5 (Fog infrastructure). We define the Fog infrastructure by the following tuple:

$$I = (M_e, M_r) \quad (5)$$

The tuple includes the set, $M_e = \{m_{e_i}\}$, of edge machines, and the set, $M_r = \{m_{r_i}\}$, of remote machines. $\square$

Multiple instances of the same service back-end may have been deployed in (edge and/or remote) machines of the Fog infrastructure. These instances are accessible through different endpoint addresses. We define the notion of the service instance in the Fog infrastructure as follows.

Definition 6 (Service instance through the Fog). We define an instance, si , of a service back-end by the following tuple:

$$si = (api, uri, m, p) \quad (6)$$

The tuple includes:

- the api of the service (Def. 1)
- the endpoint address, uri , of si
- the machine, m (m_e or m_r), where si has been deployed
- machine-learning model, p , of response time of si on m .

Illustrative example. The front-end of the auction app is installed on a resource-constrained device (esp., a smartphone). An instance of the `k-means` service of the back-end of the app is deployed on an edge machine (i.e., laptop computer). Finally, a virtual machine on the cloud is used for deploying a

remote instance of the `k-means` service. The details about the Fog infrastructure that we used for the app are provided in the experimental evaluation of our approach (Section 8).

6 AI SELF-BACK-END AS A SERVICE

Our approach extends the back-end of mobile apps so that the back-end is converted into an AI self-back-end exposed as a service. The back-end is converted into an AI self-back-end because the back-end includes an AI-based autonomic (i.e., self-adaptive) component [22]. The autonomic component dynamically predicts the response time of the edge instances and the remote instances of the service back-end and decides to invoke the instances that achieves the lowest predicted response-time. To do so, the autonomic component could implement the classical version of the autonomic control-loop of self-adaptive software [22], as described in Section 6.1.

6.1 Classical Version of Autonomic Control-Loop

The classical version of the autonomic control-loop has already been implemented in our preliminary approach specified by our past work in [4]. This classical version of the control loop implements the Monitor-Analyze-Plan-Execute loop of self-adaptive software [22]. The entry point of the control loop is the Planning mechanism that predicts the response time of service instances and selects the instance with the lowest response-time. In particular, the Planning mechanism uses the constructed mathematical-expressions to predict the response time of service instances. The Execution mechanism invokes the selected service-instance via using a service proxy that is included between the autonomic component and the target service-instance. The Execution mechanism implements the service proxy. The Monitoring mechanism records the response time of the invocations to the edge/remote service-instances. The Analysis mechanism (re-)builds the mathematical expression of the response time of service instances.

However, as described in Section 2, the version of the autonomic control-loop that has been implemented in [4] has an important restriction. The restriction is that it captures the response time via building mathematical expressions in a manner that is tied to the independent variables of the expressions. Given that the construction manner of a mathematical expression is tied to its independent variables, a new mechanism should be developed for different independent variables. To propose a generic approach that builds mathematical expressions in a manner that is not tied to specific independent variables, we propose an AI-based version of the autonomic control-loop (Section 6.2). In a nutshell, the AI control-loop takes any independent variable as input and fits the independent variables to monitoring data for constructing machine-learning models of the response time of service instances.

6.2 AI Version of Autonomic Control-Loop

The AI version of the autonomic control-loop implements a machine-learning technique (e.g., k-nearest neighbours, support vector machine, decision tree, neural network) [26] that is used to pre-train machine-learning models of the response time of service instances. In other words, the switching decision is made based on pre-trained machine-learning expressions. To decide a service instance, the autonomic component implements an AI variation of the classical version of the autonomic control-loop that differs from the classical version in the following points (Section 6.2.1).

Fig. 4. Mobile apps extended with AI self-back-end as a service.

6.2.1 AI Control-Loop vs. Classical Control-Loop

This variation has some differences with respect to the classical version of the control loop that is implemented in [4]. The first difference is that the Planning mechanism of the AI version of the control loop now works in two separate modes: an offline training mode and an online prediction mode:

- 1) *Training mode*: the mechanism offline receives training data and trains a machine-learning model of the response time of edge/remote instances.
- 2) *Prediction mode*: the mechanism receives an input dataset in an online way and predicts the response time of service instances via using the trained machine-learning models.

In the classical version of the autonomic control-loop, the Planning mechanism works in a joint online training-prediction mode [4].

The second difference between the versions of the control loop is that the Monitoring mechanism is not needed because the response time of the available service-instances have been already been provided to the Planning mechanism in the training mode. The third difference is that the Analysis mechanism is not also needed because the machine-learning models of response time are not reconstructed. Finally, the fourth difference is that the Planning mechanism does not interact with the Execution mechanism when the Planning mechanism works in the training mode because no instance is invoked in the training mode. Having detailed the differences between the two versions of the control loop, we now move on to the description of the architecture of the AI version of the autonomic control-loop (Section 6.2.2).

6.2.2 Architecture of AI Autonomic Control-Loop

The architecture of the AI self-back-end as a service is depicted in Fig. 4. According to this architecture, the controller of an app interacts with the AI self-back-end. The AI self-back-end encapsulates and invokes the available edge instances and remote instances of the service back-end of the app (see the depicted instances at the bottom of Fig. 4). In other words, multiple instances of the back-end are deployed on the Fog infrastructure that is specified in Section 5.

In more detail, the Web API of the AI self-back-end implements the proxy design-pattern [28] because both the edge

Fig. 5. The auction app extended with an AI self-back-end as a service.

instances and the remote instances of the service back-end expose the original Web API of the back-end (see the *Service Proxy* abstract class in Fig. 4). The service proxy is implemented by a RESTful or a SOAP technological variations (see the *Technological Service Proxy* class in Fig. 4).

The AI self-back-end further implements the AI version of the autonomic control-loop that trains machine-learning models of response time and self-decides at runtime the edge/remote instance that will be invoked (see the *Autonomic Control-Loop* class in Fig. 4). The control loop is activated when the Web API of the AI self-back-end is invoked. The entry point of the control loop is the Planning mechanism that is invoked when the controller wishes to train machine-learning expressions or to predict the response time of service instances. The mechanisms of the AI version of the control loop are detailed in Section 7. Before describing the mechanisms, we illustrate the architecture of the AI self-back-end by revisiting the case study of the auction app.

Illustrative example. Fig. 5 presents the architecture of the auction app extended with an AI self-back-end as a service. In this figure, the Web API of the service back-end corresponds to the *kMeans* class visualised in Fig. 3. The *kMeansSelfBackEnd* class implements the *kMeans* Web API and interacts with the autonomic control-loop. The *kMeansSelfBackEnd* class further uses a RESTful service proxy to forward invocations to the edge/remote instances of the back-end. In other words, all the instances of the auction back-end have been deployed as RESTful Web services (they could have been deployed as SOAP-based Web services). Finally, the *kMeansSelfBackEnd* class has been exported as a RESTful service and the controller uses a RESTful client for invoking the *kMeansSelfBackEnd* class.

7 AI AUTONOMIC CONTROL-LOOP MECHANISMS

As described in Section 6.2, the AI autonomic control-loop includes two mechanisms out of the four mechanisms of the classical version of the autonomic control-loop: the Planning mechanism (Section 7.1) and the Execution mechanism (Section 7.2).

Both the Planning and the Execution mechanisms are deployed on an edge machine. In this way, the network/communication

latency between the controller of an app and the Planning mechanism is avoided, when the Planning mechanism is executed in prediction mode. In a similar way, the latency between the Planning mechanism and the Execution mechanism is avoided, when the Planning mechanism interacts with the Execution mechanism.

The AI autonomic control-loop does not include further mechanisms, apart from the Planning mechanism and the Execution mechanism. Especially, the current version of the AI autonomic control-loop *does not manage* the edge/remote instances of service back-end. In other words, the AI autonomic control-loop does not add/remove edge/remote instances or does not manage the availability/stability of the instances (leaving the dynamic management of the instances as future work). Thus, it makes the assumption that a set of edge/remote instances is available when an app is first installed on the device of an end-user. It further assumes that this set of edge/remote instances is not updated as the app is executed.

7.1 Planning Mechanism

The Planning mechanism works in a training mode (Section 7.1.1) or in a prediction mode (Section 7.1.2).

7.1.1 Training Mode

The Planning mechanism receives training/monitoring data and uses them to train one of the requested machine-learning models. In particular, the mechanism accepts one of the following types of machine-learning models as input: k-nearest neighbours, support-vector machines, decision tree, and neural network [26]. Then, the mechanism trains the requested model as follows.

K-nearest neighbours. The mechanism sets the number of the neighbors to be equal to 10 because, according to the experimental evaluation of the model of k-nearest neighbors (Section 8), the model achieves on average the highest accuracy for this number.

Support-vector machines. According to the experimental evaluation of the model of support-vector machines (Section 8), the model achieves on average the highest accuracy when the mechanism sets the gamma variable to the following value:

$$\gamma = \frac{1}{\text{indep. variables} + \text{variance of indep. variables}} \quad (7)$$

Decision trees. The mechanism sets the three parameters (criterion, max features, min impurity split) of the model of decision-tree models as follows. The criterion is the function to measure the quality of a split. The mechanism calculates the value of this criterion by using the Gini impurity [26]. The max number of features corresponds to the number of features to be considered in each split. The mechanism sets the max number of features to four. The min impurity split corresponds to the threshold for early stopping in the tree growth. The mechanism sets the min value of the impurity split to zero.

Neural networks. The mechanism implements a three-layer LeakyReLU neural network [26]. The number of the neurons at the first layer equals to the number of the independent variables (see below). For instance, the neural network for the Web API (Fig. 3) of the auction back-end includes four neurons. The number of the neurons at the second layer is calculated by the following expression: $\lceil \sqrt{\text{first layer neurons}} \rceil$. The single neuron of the third layer corresponds to the predicted response-time.

(In)dependent variables of the machine-learning models. The independent variables of all the models are related to the characteristics of the input datasets of the data-analytics algorithms. The datasets include the values for the primitive parameters

of the operations (Def. 3). We focus on primitive parameters exclusively because primitive parameters carry on real data. In the case of the k-means algorithm of the auction back-end, the independent variables include the following variables: the number of clusters, the number of clustering iterations, the number of data-points, and the dimension of data-points. The dependent variable corresponds to the response time of the operations of the Web API of edge/remote instances of the back-end.

Illustrative example. Revisiting Fig. 3, the Planning mechanism accepts as input: the clusters number, k , the clustering iterations, it , the dimension, d , of the parameter, `Vector`, the length of the parameter, `DataPoint` (along with the values of each one of the data points). As an example, the mechanism accepts the following dataset as input: 3 clusters, 100 iterations, and 1000 data points of 3 dimensions.

Training machine-learning models for service instances. The mechanism trains a separate machine-learning model for each edge/remote instance of a back-end.

Training a large number of service instances. The training mode of the mechanism is executed offline and does not affect the performance of an app. However, the mechanism avoids a large training-time in the case of large number of available edge/remote instances. In particular, the mechanism avoids a large training-time by training the expression of each instance in parallel. The parallel training applies a multi-thread technique by executing one thread for the training of each edge/remote instance. The parallel training of multiple machine-learning models is feasible because they are independent to each other.

7.1.2 Prediction Mode

The Planning mechanism receives the current datasets provided from the front-end to the AI self-back-end. Moreover, the mechanism accepts the desired kind of machine-learning model (KNN, SVM, decision tree, or neural network) as input. To make a prediction based on a trained model, the mechanism first restores the trained model. Then, the mechanism predicts the response times of all the service instances via using the requested machine-learning model. The mechanism selects the instance with the lowest response-time. Finally, the Planning mechanism interacts with the Execution mechanism to invoke the selected instance.

Making a large number of predictions. The prediction mode of the Planning mechanism is executed online and affects the performance of the app. However, the time that the mechanism needs to make the predictions for all the edge/remote instances scales in a linear way with the number of the instances. This time is practically small because the mechanism needs a time in the order of milliseconds to make predictions for a hundred of instances.

7.2 Execution Mechanism

The Execution mechanism implements a service proxy that forwards the invocations made by the front-end to the edge/remote instance selected by the Planning mechanism. In this way, the Execution mechanism binds the front-end of the app to the back-end in a seamless way, i.e. without suspending the app execution. The Execution mechanism provides two implementations of an abstract proxy: one implementation for RESTful Web services and one implementation for SOAP-based services.

8 EXPERIMENTAL EVALUATION

We evaluate the accuracy of our approach via measuring the percentages of the correct switching-decisions made by the AI

TABLE 2
Overall statistics of the datasets we used in the experiments.

Data-Points	Dimension of Data-Points	Number of Datasets
[1, 999]	[3, 14]	3131
[1000, 1999]	[3, 14]	1908
[2000, 2999]	[3, 14]	1916
[3000, 3999]	[3, 14]	1837
[4000, 4999]	[3, 12]	1891
[5000, 5999]	[3, 12]	157
[6000, 6999]	[3, 14]	164
[7000, 7999]	[3, 11]	169
[8000, 8999]	[3, 11]	162
[9000, 9999]	[3, 12]	196
[10000, 10999]	[7, 8]	46
Total number of the datasets:		11577

self-back-end (Section 8.2). We further evaluate the efficiency of our approach via measuring the improvement in the response time of an app when the AI self-back-end is used (Section 8.2). Prior to presenting the results, we set up our experiments (Section 8.1).

8.1 Experimental Setup

Mobile app. We extended an open-source auction app¹⁰ with a research-prototype (implemented in Java) of the AI self-back-end that is exposed as a RESTful Web service.

Datasets. We use datasets collected from publicly available UC Irvine machine-learning repository¹¹. Each dataset consists of multi-dimensional data-points. We selected datasets that can be used for online data-analytics. In other words, we selected 11577 datasets for which the execution time of the k-means algorithm on the used edge/remote machines is low¹². In particular, we selected datasets whose data-points' number is lower than 12000 and the data-points' dimension is at most 14. The overall statistics of the used datasets are presented in Table 2. The repository provides a small number of datasets that contain a very high number of data-points. The statistics of the datasets used in each machine-learning model are presented in Sections 8.2.1-8.2.4.

Fog infrastructure. The front-end of the auction app runs on a mobile phone. We further used two laptop devices. Each laptop provides an edge machine. The first laptop used for deploying an edge instance of the original service back-end of the app. The second laptop used for deploying an edge instance of the component of the AI autonomic control-loop. The mobile phone and the laptops are connected at the same wireless local-network. A remote instance of the original back-end is deployed on Microsoft Azure Cloud Service in the UK-south region. All the instances of the service back-end are exposed as RESTful Web services.

8.1.1 Accuracy Metrics

To measure the accuracy of our approach, we first calculate the precision and the recall of each switching decision [29]. In particular, we construct for each switching decision the confusion matrix demonstrated in Table 3. We then calculate the precision of all the remote (resp., edge) switching-decisions that equals to

TABLE 3
The confusion matrix for calculating precision and recall.

		Prediction	
		Remote	Edge
Correct Decision	Remote	True remote	False edge
	Edge	False remote	True edge

the percentage of the correct predictions over all the remote (resp., edge) switching-decisions, as defined below.

$$\text{remote precision} := \frac{\text{true remote}}{\text{true remote} + \text{false remote}} \quad (8)$$

$$\text{edge precision} := \frac{\text{true edge}}{\text{true edge} + \text{false edge}} \quad (9)$$

The recall of all the remote (resp., edge) switching-decisions equals to the percentage of the correct predictions over all the correct remote (resp., edge) switching-decisions and the wrong edge (resp., remote) switching-decisions, as defined below.

$$\text{remote recall} := \frac{\text{true remote}}{\text{true remote} + \text{false edge}} \quad (10)$$

$$\text{edge recall} := \frac{\text{true edge}}{\text{true edge} + \text{false remote}} \quad (11)$$

Finally, we measure the overall accuracy of our approach over all the switching decisions, as follows:

$$\text{Overall accuracy} := \frac{\text{correct switching decisions}}{\text{all switching decisions}} \quad (12)$$

8.1.2 Efficiency Metrics

To evaluate the efficiency of our approach, we take into account the fact that an app would use only remote instances of a back-end if we do not apply our approach. By applying our approach, there is an improvement in the response time of the app if a correct switching-decision is the invocation of an edge instance and the AI self-back-end decides the edge instance. In this case, the response time is reduced by the difference of the response time of the app using a remote instance from the response time of the app using the edge instance. Thus, the improvement in the response time (RTI) of the app for a single switching decision is defined below.

$$\text{edge RTI} := \frac{\text{remote RT} - \text{SelfBackEnd edge RT}}{\text{remote RT}} \quad (13)$$

If the correct switching-decision is the invocation of a remote instance and the AI self-back-end decides the remote instance, then there is no response-time improvement because the remote instance should had by default been invoked. We finally calculate the average improvement in the response time of an app for all the switching decisions, as follows.

$$\overline{\text{edge RTI}} := \frac{\sum \text{edge RTI}}{\text{edge switching decisions}} \quad (14)$$

10. <https://github.com/jagmohansingh/auction-system>

11. <http://archive.ics.uci.edu/ml/index.php>

12. The datasets we use in the experiments are available here: https://drive.google.com/file/d/1h9I9v_O9r03j3z-axXEmxI50RZs9F9An/view?usp=sharing

TABLE 4
The configuration of the AI self-back-end in the experiments to train/use various machine-learning models.

	Remote CPUs	Remote Memory	Edge CPUs	Edge Memory	Training Datasets	Testing Datasets
<i>Neural networks</i>	4	8GB	0.5	1GB	578	925
<i>K-nearest neighbours</i>	4	8GB	0.5	1GB	578	810
<i>Support-vector machines</i>	4	8GB	0.5	1GB	578	1540
<i>Decision trees</i>	4	8GB	0.5	1GB	578	1016

TABLE 5
The statistics of the testing datasets we used in neural networks.

Data-Points	Dimension of Data-Points	Number of Datasets
[1, 999]	[3, 14]	240
[1000, 1999]	[3, 9]	153
[2000, 2999]	[3, 8]	175
[3000, 3999]	[3, 11]	141
[4000, 4999]	[3, 11]	140
[5000, 5999]	[3, 12]	11
[6000, 6999]	[3, 8]	16
[7000, 7999]	[3, 8]	11
[8000, 8999]	[3, 8]	17
[9000, 9999]	[3, 8]	18
[10000, 10999]	8	3
Total number of the testing datasets:		925

Fig. 6. The overall accuracy of the AI self-back-end in neural networks.

TABLE 6
The confusion matrix and the precision/recall of neural networks.

	Remote	Edge		Precision	Recall
Correct Remote	715	52	Remote	98%	93%
Correct Edge	18	140	Edge	73%	89%

8.2 Experimental Results

We repeat below the same set of experiments four times, one time for each one of the following machine-learning models: neural networks, k-nearest neighbours, support-vector machines, and decision trees (Sections 8.2.1-8.2.4). To this end, we configure the AI self-back-end of the auction app to use one of the models at each time. Table 4 presents the hardware characteristics of the machines used for this experiment, along with the numbers of the training/testing datasets. We use the same numbers of CPUs and the same amount of RAM in all the machine-learning models for uniformity reasons. We also use the same number of training datasets for all the models. Concerning the testing datasets, we present a different number of training datasets for each model. The reason is because we want to present for each model, *representative datasets* that show the variation of the accuracy of the model (from the smallest accuracy to the highest accuracy). The number of the testing datasets that we present is different in each model because the accuracy of each model vary in a different manner from the other models.

8.2.1 AI Self-Back-End Uses Neural Networks

Experimental configuration. Table 5 presents the statistics of the 925 representative testing-datasets of this experiment. The datasets span all the available ranges of data-points and as expected, a small number of datasets contain a very high number of data-points.

Accuracy. Fig. 6 presents the overall accuracy of the neural-network model for each testing dataset (Eq. 12). 925 switching decisions made by the AI self-back-end in the prediction mode. Please note that we present in Fig. 6 the accuracy of the model for the datasets with IDs that are greater than 578 because we use the first 578 datasets for training the neural-network model. 825 out of 925 switching decisions are correct decisions. We depict at the left-hand side of the x-axis of the chart of Fig. 6, representative datasets for which the accuracy of the model varies/fluctuates. We depict at the right-hand side of the x-axis, approximately 250

representative datasets for which the accuracy of the model is stabilised. The overall accuracy of the model is 92%.

To get an insight in the switching decisions, we present in Table 6 the confusion matrices of the switching decisions. We further present in Table 6 the precision/recall of the model (Eq. 8-11). We observe the AI self-back-end that is configured to use neural networks achieves high precision ($\geq 73\%$) and very high recall ($\geq 89\%$).

Response-time improvement. We calculated the response-time improvement of the app in the edge switching-decisions for neural networks (Fig. 7). We depict at the left-hand side of the x-axis of the chart of Fig. 7, the datasets for which the accuracy of the neural-network model varies/fluctuates. As expected, the response-time improvement varies/fluctuates too. We also depict at the right-hand side of the x-axis, the datasets for which the response-time improvement is stabilised (in accordance with the accuracy of the model). The average response-time improvement in neural networks is 45%.

8.2.2 AI Self-Back-End Uses K-Nearest Neighbours

Experimental configuration. Table 7 presents the statistics of the 810 representative testing-datasets of this experiment. The datasets span all the available ranges of data-points and as expected, a small number of datasets contain a very high number of data-points.

Fig. 7. The response-time improvement of the app in neural networks.

TABLE 7

The statistics of the testing datasets we used in k-nearest neighbours.

Data-Points	Dimension of Data-Points	Number of Datasets
[1, 999]	[3, 12]	231
[1000, 1999]	[3, 8]	137
[2000, 2999]	[3, 8]	129
[3000, 3999]	[3, 14]	132
[4000, 4999]	[3, 12]	124
[5000, 5999]	[3, 8]	10
[6000, 6999]	[3, 7]	13
[7000, 7999]	[3, 11]	10
[8000, 8999]	[3, 11]	13
[9000, 9999]	[3, 8]	9
[10000, 10999]	8	2
Total number of testing datasets:		810

Fig. 8. Overall accuracy of the AI self-back-end in k-nearest neighbours.

Accuracy. Fig. 8 presents the overall accuracy of the k-nearest-neighbours model for each testing dataset (Eq. 12). 810 switching decisions made by the AI self-back-end in the prediction mode. We present in Fig. 8 the accuracy of the model for the datasets with IDs that are greater than 578 because we use the first 578 datasets for training the model. 557 out of 810 switching decisions are correct decisions. We depict at the left-hand side of the x-axis of the chart of Fig. 8, representative datasets for which the accuracy of the model varies/fluctuates. We depict at the right-hand side of the x-axis, approximately 100 representative datasets for which the accuracy of the model is stabilised. The overall accuracy of the model is 69%.

To get an insight in the switching decisions, we present in Table 8 the confusion matrices of the switching decisions and the precision/recall of the model (Eq. 8-11). We observe the AI self-back-end that is configured to use k-nearest neighbours achieves high precision ($\geq 63\%$) and medium recall ($\geq 46\%$).

Response-time improvement. We calculated the response-

TABLE 8

The confusion matrix and the precision/recall of k-nearest neighbours.

	Remote	Edge		Precision	Recall
Correct Remote	813	85	Remote	71%	83%
Correct Edge	168	144	Edge	63%	46%

Fig. 9. The response-time improvement in k-nearest neighbours.

TABLE 9

The statistics of the testing datasets in support-vector machines.

Data-Points	Dimension of Data-Points	Number of Datasets
[1, 999]	[3, 14]	435
[1000, 1999]	[3, 12]	255
[2000, 2999]	[3, 9]	244
[3000, 3999]	[3, 14]	251
[4000, 4999]	[3, 12]	232
[5000, 5999]	[3, 8]	26
[6000, 6999]	[3, 8]	17
[7000, 7999]	[3, 11]	27
[8000, 8999]	[3, 8]	20
[9000, 9999]	[3, 8]	28
[10000, 10999]	8	5
Total number of the testing datasets:		1540

Fig. 10. The overall accuracy in support-vector machines.

time improvement of the app in the edge switching-decisions for k-nearest neighbours (Fig. 9). We depict at the left-hand side of the x-axis of Fig. 9, the datasets for which the accuracy of the model of the k-nearest neighbours varies/fluctuates. We also depict at the right-hand side of the x-axis, the datasets for which the accuracy of the model is stabilised. As expected, the response-time improvement increases as the accuracy is stabilised. The average response-time improvement in k-nearest neighbours is 23%.

8.2.3 AI Self-Back-End Uses Support-Vector Machines

Experimental configuration. Table 9 presents the 1540 representative testing-datasets of this experiment. The datasets span all the available ranges of data-points and as expected, a small number of datasets contain a very high number of data-points.

Accuracy. Fig. 10 presents the overall accuracy of the model of the support-vector machines for each testing dataset (Eq. 12). 1540 switching decisions made by the AI self-back-end in the prediction mode. We present in Fig. 10 the accuracy of the model for the datasets with IDs that are greater than 578 because we use the first 578 datasets for training the model. 1244 out of 1540 switching decisions are correct decisions. We depict at the left-hand side of the x-axis of the chart of Fig. 10, representative datasets for which the accuracy of the model varies/fluctuates. We depict at the right-hand side of the x-axis, approximately 30 representative datasets for which the accuracy of the model is stabilised. The overall accuracy of the model is 81%.

To get an insight in the switching decisions, we present in Table 10 the confusion matrices of the switching decisions for this experiment. We further present in Table 10 the precision and the recall of the model (Eq. 8-11). We observe in the results that the AI self-back-end that is configured to use support-vector machines achieves high precision ($\geq 81\%$) and low recall ($\geq 9\%$).

Response-time improvement. We calculated the response-time improvement of the app in the edge switching-decisions for support-vector machines (Fig. 11). We depict at the left-hand side of the x-axis of Fig. 11, the datasets for which the accuracy

TABLE 10
Confusion matrix and precision/recall of support-vector machines.

	Remote	Edge		Precision	Recall
Correct Remote	1215	3	Remote	81%	100%
Correct Edge	293	29	Edge	91%	9%

Fig. 11. The response-time improvement in support-vector machines.

of the model of support-vector machines varies/fluctuates. As expected, the response-time improvement varies/fluctuates for the first datasets, but it is stabilised for the datasets for which the accuracy is stabilised too. The average response-time improvement in support-vector machines is 20%.

8.2.4 AI Self-Back-End Uses Decision Trees

Experimental configuration. Table 11 presents the 1016 representative testing-datasets of this experiment. The datasets span all the available ranges of data-points and as expected, a small number of datasets contain a very high number of data-points.

Accuracy. Fig. 12 presents the overall accuracy of the decision-tree model for each testing dataset (Eq. 12). 1016 switching decisions made by the AI self-back-end in the prediction mode. We present in Fig. 12 the accuracy of the model for the datasets with IDs that are greater than 578 because we use the first 578 datasets for training the model. 804 out of 1016 switching decisions are correct decisions. We depict at the left-hand side of the x-axis of the chart of Fig. 12, representative datasets for which the accuracy of the model varies/fluctuates. We depict at the right-hand side of the x-axis of the same chart, approximately 200 representative datasets for which the accuracy of the model is stabilised. The overall accuracy of the model is 81%.

To get an insight in the switching decisions, we present in Table 12 the confusion matrices of the switching decisions and the precision/recall of the model (Eq. 8-11). We observe the AI self-back-end that is configured to use decision trees achieves medium precision ($\geq 54\%$) and recall ($\geq 68\%$).

TABLE 11
The statistics for the testing datasets we used in decision trees.

Data-Points	Dimension of Data-Points	Number of Datasets
[1, 999]	[3, 14]	284
[1000, 1999]	[3, 12]	169
[2000, 2999]	[3, 9]	164
[3000, 3999]	[3, 11]	163
[4000, 4999]	[3, 12]	149
[5000, 5999]	[3, 4]	14
[6000, 6999]	[3, 8]	14
[7000, 7999]	[3, 11]	20
[8000, 8999]	[3, 8]	14
[9000, 9999]	[3, 8]	19
[10000, 10999]	8	6
Total number of testing datasets:		1016

Fig. 12. The overall accuracy of the AI self-back-end in decision trees.

TABLE 12
The confusion matrix and the precision/recall of decision trees.

	Remote	Edge		Precision	Recall
Correct Remote	645	136	Remote	89%	83%
Correct Edge	76	159	Edge	54%	68%

Response-time improvement. We calculated the response-time improvement of the app in the edge switching-decisions for decision trees (Fig. 13). We depict at the left-hand side of the x-axis of Fig. 13, the datasets for which the accuracy of the model of decision trees varies/fluctuates. We also depict at the right-hand side of the x-axis, the datasets for which the accuracy of the model is stabilised. As expected, the response-time improvement increases as the accuracy is stabilised. The average response-time improvement in decision trees is 37%.

8.2.5 Overall Comparison of the Experimental Results

We compare in Table 13 the machine-learning models we used with respect to their overall accuracy and the response-time improvement. We observe the k-nearest neighbours model achieves the lowest accuracy, medium accuracy, and medium precision/recall. The neural-network model achieves the highest overall accuracy, the highest precision/recall, and the highest response-time improvement. The second better model is the decision-tree model that achieves high overall accuracy, the second highest response-time improvement, but medium precision/recall. The support-vector-machine model has the same overall accuracy with the decision-tree model, higher precision than the precision of the decision-tree model, but it achieves very low recall and the lowest response-time improvement.

Fig. 13. The response-time improvement in decision trees.

TABLE 13
Overall comparison of the experimental results.

	Accuracy	Precision	Recall	RTI
Neural networks	92%	70%	89%	45%
K-nearest neighbours	69%	63%	46%	23%
Support-vector machines	81%	81%	9%	20%
Decision trees	81%	54%	68%	37%

9 CONCLUSIONS & FUTURE WORK

We contributed with the specification of the architecture and the mechanisms of the AI self-back-end for edge/remote instances. We evaluated the accuracy and efficiency of the AI self-back-end in four machine-learning models. The evaluation results showed that the neural-network model has the highest overall accuracy and the highest precision and recall.

A future research-direction is to extend our approach to use further categories of monitoring data (e.g., memory consumption). Another direction is to consider the synchronisation time (spent by the edge/remote instances for storing datasets) in the training of machine-learning models. A final direction is to enhance the AI self-back-end with the capability to dynamically (de-)register service instances that are (not) available on the Fog and the Cloud.

REFERENCES

- [1] P. Plebani, D. García-Pérez, M. Anderson, D. Bermbach, C. Capiello, R. I. Kat, F. Pallas, B. Pernici, S. Tai, and M. Vitali, "Information logistics and fog computing: The ditas approach," in *International Conference on Advanced Information Systems Engineering*, 2017, pp. 129–136.
- [2] J. Aikat, J. Kaur, F. D. Smith, and K. Jeffay, "Variability in tcp round-trip times," in *ACM SIGCOMM Conference on Internet Measurement*, 2003, pp. 279–284.
- [3] C. Hong and B. Varghese, "Resource management in fog/edge computing: A survey on architectures, infrastructure, and algorithms," *ACM Computing Surveys*, vol. 52, no. 5, pp. 97:1–97:37, 2019.
- [4] D. Athanasopoulos, M. McEwen, and A. Rainer, "Mobile apps with dynamic bindings between the fog and the cloud," in *International Conference on Service-Oriented Computing*, 2019, pp. 539–554.
- [5] E. Saurez, K. Hong, D. Lillethun, U. Ramachandran, and B. Ottenwälder, "Incremental deployment and migration of geo-distributed situation awareness applications in the fog," in *ACM International Conference on Distributed and Event-based Systems*, 2016, pp. 258–269.
- [6] A. Evangelidis, D. Parker, and R. Bahsoon, "Performance modelling and verification of cloud-based auto-scaling policies," *Future Generation Computer Systems*, vol. 87, pp. 629–638, 2018.
- [7] C. Barna, M. Litoiu, M. Fokaefs, M. Shtern, and J. Wigglesworth, "Runtime performance management for cloud applications with adaptive controllers," in *International Conference on Performance Engineering*, 2018, pp. 176–183.
- [8] C. Barna, H. Khazaei, M. Fokaefs, and M. Litoiu, "Delivering elastic containerized cloud applications to enable devops," in *International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2017, pp. 65–75.
- [9] A. Brogi, S. Forti, and A. Ibrahim, "How to best deploy your fog applications, probably," in *International Conference on Fog and Edge Computing*, 2017, pp. 105–114.
- [10] X. Guo, R. Singh, T. Zhao, and Z. Niu, "An index based task assignment policy for achieving optimal power-delay tradeoff in edge cloud systems," in *IEEE International Conference on Communications*, 2016, pp. 1–7.
- [11] T. Chen and R. Bahsoon, "Self-adaptive and online qos modeling for cloud-based software services," *IEEE Transactions on Software Engineering*, vol. 43, no. 5, pp. 453–475, 2017.
- [12] A. Yousefpour, G. Ishigaki, and J. P. Jue, "Fog computing: Towards minimizing delay in the internet of things," in *IEEE International Conference on Edge Computing*, 2017, pp. 17–24.
- [13] A. Yousefpour, A. Patil, G. Ishigaki, I. Kim, X. Wang, H. C. Cankaya, Q. Zhang, W. Xie, and J. P. Jue, "FOGPLAN: A lightweight qos-aware dynamic fog service provisioning framework," *IEEE Internet of Things Journal*, vol. 6, no. 3, pp. 5080–5096, 2019.
- [14] D. H. Tran, N. H. Tran, C. Pham, S. M. A. Kazmi, E. Huh, and C. S. Hong, "Oaas: Offload as a service in fog networks," *Computing*, vol. 99, no. 11, pp. 1081–1104, 2017.
- [15] Q. Zhang, M. Lin, L. T. Yang, Z. Chen, S. U. Khan, and P. Li, "A double deep q-learning model for energy-efficient edge scheduling," *IEEE Transactions on Services Computing*, vol. 12, no. 5, pp. 739–749, 2019.
- [16] W. Chen, D. Wang, and K. Li, "Multi-user multi-task computation offloading in green mobile edge cloud computing," *IEEE Transactions on Services Computing*, vol. 12, no. 5, pp. 726–738, 2019.
- [17] N. Wang, B. Varghese, M. Matthaiou, and D. S. Nikolopoulos, "ENORM: A framework for edge node resource management," *IEEE Transactions on Services Computing*, vol. 13, no. 6, pp. 1086–1099, 2020.
- [18] A. Brogi and S. Forti, "Qos-aware deployment of iot applications through the fog," *IEEE Internet of Things Journal*, vol. 4, no. 5, pp. 1185–1192, 2017.
- [19] R. Deng, R. Lu, C. Lai, T. H. Luan, and H. Liang, "Optimal workload allocation in fog-cloud computing toward balanced delay and power consumption," *IEEE Internet of Things Journal*, vol. 3, no. 6, pp. 1171–1181, 2016.
- [20] H. Gupta, A. V. Dastjerdi, S. K. Ghosh, and R. Buyya, "ifogsim: A toolkit for modeling and simulation of resource management techniques in the internet of things, edge and fog computing environments," *Software: Practice and Experience*, vol. 47, no. 9, pp. 1275–1296, 2017.
- [21] S. Ma, S. Guo, K. Wang, W. Jia, and M. Guo, "A cyclic game for service-oriented resource allocation in edge computing," *IEEE Transactions on Services Computing*, vol. 13, no. 4, pp. 723–734, 2020.
- [22] J. O. Kephart and D. M. Chess, "The vision of autonomic computing," *IEEE Computer*, vol. 36, no. 1, pp. 41–50, 2003.
- [23] L. Richardson and S. Ruby, *Restful Web Services*, 1st ed. O'Reilly, 2007.
- [24] T. Erl, *Service-Oriented Architecture: Analysis and Design for Services and Microservices*, second edition ed. Prentice Hall, 2016.
- [25] P. Kaur, M. Goyal, and J. Lu, "Pricing analysis in online auctions using clustering and regression tree approach," in *Agents and Data Mining Interaction*, 2012, pp. 248–257.
- [26] P. Tan, M. Steinbach, and V. Kumar, *Introduction to Data Mining*. Pearson Addison Wesley, 2006.
- [27] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the internet of things," in *Workshop on Mobile Cloud Computing*. ACM, 2012, pp. 13–16.
- [28] E. Gamma, R. Helm, R. Johnson, and J. M. Vlissides, *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1994.
- [29] R. A. Baeza-Yates and B. A. Ribeiro-Neto, *Modern Information Retrieval*. ACM Press/Addison-Wesley, 1999.

Dionysis Athanasopoulos is a lecturer with the school of Electronics, Electrical Engineering & Computer Science of Queen's University Belfast, UK. His research interests mainly include service-oriented architecture, Fog computing, self-adaptive software, data engineering, and software design principles & patterns. Contact him at D.Athanasopoulos@qub.ac.uk.

Dewei Liu received his BEng with first-class honours in Computer Science (2020) from Queen's University Belfast, UK. He is now pursuing an MSc in machine learning and data science at Imperial College London, UK. His main field of interest is computational statistics and data analysis. Contact him at dliu08@qub.ac.uk.