



Kolomvatsos, K. and Anagnostopoulos, C. (2020) An intelligent edge-centric queries allocation scheme based on ensemble models. *ACM Transactions on Internet Technology*, 20(4), 45.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

© 2020 Copyright is held by the author/owner(s). This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Information Systems* 1,4 <http://dx.doi.org/10.1145/3417297>

<http://eprints.gla.ac.uk/222104/>

Deposited on: 12 August 2020

An Intelligent Edge-Centric Queries Allocation Scheme based on Ensemble Models

Kostas Kolomvatsos¹ Christos Anagnostopoulos²

¹Department of Informatics and Telecommunications, University of Thessaly, 35131 Lamia Greece

²School of Computing Science, University of Glasgow, G12 8RZ Glasgow UK

Abstract

The combination of Internet of Things (IoT) and Edge Computing (EC) can assist in the delivery of novel applications that will facilitate end users activities. Data collected by numerous devices present in the IoT infrastructure can be hosted into a set of EC nodes becoming the subject of processing tasks for the provision of analytics. Analytics are derived as the result of various queries defined by end users or applications. Such queries can be executed in the available EC nodes to limit the latency in the provision of responses. In this paper, we propose a meta-ensemble learning scheme that supports the decision making for the allocation of queries to the appropriate EC nodes. Our learning model decides over queries' and nodes' characteristics. We provide the description of a matching process between queries and nodes after concluding the contextual information for each envisioned characteristic adopted in our meta-ensemble scheme. We rely on widely known ensemble models, combine them and offer an additional processing layer to increase the performance. The aim is to result a subset of EC nodes that will host each incoming query. Apart from the description of the proposed model, we report on its evaluation and the corresponding results. Through a large set of experiments and a numerical analysis, we aim at revealing the pros and cons of the proposed scheme.

Internet of Things, Edge Computing, Query Allocation, Ensemble Classification Models

1 Introduction

Currently, we are in the middle of a data management revolution brought by the numerous devices present in the *Internet of Things* (IoT) infrastructure. These devices are capable of interacting with their environment, perform some processing activities and report to the upper layers either the collected data or

the produced knowledge. On top of the discussed devices, the *Edge Computing* (EC) structure, the *Fog layer* and, finally, the *Cloud* provide multiple storage and processing locations where knowledge can be delivered [52]. In this architecture, finding efficient techniques for data management becomes significant due to the high volumes of data produced by numerous devices [66]. Recent studies show that data offloading to the Cloud requires 100 to 200 ms of additional latency compared to EC technologies [15]. Hence, to reduce the latency, we could store and process the collected data to EC nodes. However, due to their limited computational capabilities only a part of data can be stored and processed locally while the remaining will be sent to Cloud. Currently, local data processing gains more attention as any solution has to build over the limited delay in the provision of the requested analytics. Pushing as much computing workload as close to the edge as possible can bring serious benefits, particularly where communication costs are high or where instant action is needed ¹.

Edge Nodes (ENs) can maintain a local dataset forming a network of distributed data repositories. Data may also be replicated for supporting fault tolerant applications. Queries/tasks generated by end users or applications could request for analytics over the collected data. As data are distributed among ENs, it is imperative to define a mechanism that allocates the incoming queries to the appropriate nodes. Such a decision should be made not only based on queries demand for data but also on various characteristics of the ENs themselves.

Motivating Example. If a query asks for temperature recordings between 10 and 20 degrees in the Celsius scale, it is useless to allocate the query to a dataset where its statistics indicate temperature recordings above 20 degrees or below 10 degrees. Allocating the query to the specific dataset, we will waste resources and time for receiving an empty set. In addition, when multiple datasets exhibit statistics ‘similar’ to the incoming query, a load balancing aspect should be taken into consideration to avoid the creation of bottlenecks in the processing activities of each EN.

Traditionally, the demand for a low latency in the provision of responses is achieved by employing more computing resources and parallelizing the application logic over the datacenter infrastructure [63]. In our case, the incoming queries can be distributed to the available ENs, however, such nodes are equipped with different resources and exhibit a different status. It is worth noticing that a query/task could be allocated in multiple ENs imposing new requirements related to the aggregation of the final responses. We should also keep in mind that the number of ENs is growing while the need for increased decentralization and autonomy is also imperative. To reduce the latency in the provision of responses, it is not enough to perform a parallel execution on top of the available ENs but also to allocate the incoming queries/tasks to appropriate ENs that will conclude the envisioned calculations in short time. The turnaround time is affected by multiple parameters like the load of ENs (the

¹https://www.accenture.com/t20170628T011733Z__w__/us-en/_acnmedia/PDF-54/Accenture-Insight-Mobility-Edge-Analytics.pdf

higher the load is, the higher the response time becomes), the number and the statistics of data present in them (the size of the datasets affects the response time) and so on and so forth. In this paper, we are motivated by the vast EC infrastructure where ENs exhibit a dynamic behaviour concluded by their interactions with IoT devices and the Cloud back end. Any decision for offloading queries/tasks to the available ENs should be mandated by their current status (i.e., their ability to start and conclude the execution in short time) and the matching between the queries/tasks requirements (as dictated by the constraints imposed by end users of applications) and the data present in the EC infrastructure. The baseline solution deals with a query/task allocated to all the available ENs which is governed by the following disadvantages: (i) the requestor should wait to receive the aggregated results from all the available ENs. The aggregation process will start only when the last EN reports its outcomes. Hence, the turnaround time is affected by the longest response time observed in the infrastructure; (ii) in the responses aggregation phase, (near) empty lists will be provided from ENs not having data related to queries/tasks constraints. Resources and time will be spent for processing such data that will, finally, will not add value to the overall processing while requiring time and resources to process them. We should not forget that ENs are usually nodes with limited computational and resources capabilities, thus, we should avoid unnecessary executions of queries/tasks.

In this paper, we propose an intelligent scheme that takes into consideration the characteristics of both, queries and ENs to conclude the most efficient allocations. Such characteristics are related to all aspects of queries and ENs and their matching, i.e., we rely on the ‘burden’ that a query adds to ENs, the data required by the query, the load and speed of ENs and so on and so forth. We consider an entity, the *Query Controller* (QC) that is responsible to receive queries (from this point forward, we refer to queries when we want to depict the execution of a query/task) and proceed with their allocation. The QC can be present at the Cloud being directly connected with a set of *Query Processors* (QPs) placed in every EN. QPs execute the incoming queries and return the final response to the interested QC. Our scheme enhances the behaviour and the decision making capabilities of QCs. We adopt the contextual information of queries and ENs characteristics at the time when the decision should be concluded. We propose the use of multiple contextual vectors, i.e., our basis for ‘matching’ queries with the available nodes. Our mechanism adopts a classification model over the the complexity class of any incoming query and the estimation of the distance between the query constraints and the data present at every EN. In the final decision making, we use a meta-ensemble model that builds on top of multiple ensemble schemes. This way, we build an ‘hierarchy’ of classification models adopted to deliver better results than individual classification modules.

The efficient allocation of queries to a number of nodes is also the subject of our previous efforts presented in [39], [41] and [43]. In [39], every allocation is concluded based on an optimal stopping scheme that examines all the available ENs concerning their ability to host and execute a query. The model proposed in [43] aims at the adoption of a learning scheme that will, finally, decide each

allocation when a query arrives. In [41], we adopt reinforcement learning and clustering to deliver the final allocation. Our motivation for extending our previous work is related to our desire to avoid the drawbacks of adopting a single solution. For instance, when relying on a specific learning model, we can meet the following disadvantages [19]: (i) a single model can be viewed as searching a space for detecting the best possible hypothesis. A statistical problem arises when the amount of training data is too small compared to the size of the hypotheses space; (ii) many algorithms are trapped to local optima. Even if there are enough training data, it may be computational expensive to find the best hypothesis in the available space; (iii) in most applications, the true function cannot be represented by any of the hypotheses in the search space. We depart from our previous work and provide a scheme that avoids the drawbacks of individual models. The contributions of this paper depict the differences from the previous efforts being described in the following list:

- we propose a meta-ensemble learning scheme for queries allocation to a set of ENs. The meta-ensemble learning scheme adopts multiple ensemble learning (sub-)models to provide a powerful and efficient allocation model.
- we propose a decision making mechanism that builds over the contextual data related to queries and ENs characteristics. Through the formulated contextual vectors, our decision making model is capable of ‘matching’ the incoming queries with the appropriate ENs.
- we propose a model for delivering the complexity of any incoming query. The complexity class depicts the ‘burden’ that a query will add in ENs where it will be executed.
- we adopt a model for estimating the distance of a query with datasets present in ENs. Such information is significant for the final allocation as we aim at allocating queries in datasets with which they exhibit the minimum distance.

The paper is organized as follows. Section 2 reports on the prior work in the field by presenting important activities related to our problem. Section 3 presents the problem under consideration and some introductory information. Section 4 discusses how to model the incoming queries and the envisioned ENs while Section 5 presents the proposed meta-ensemble learning scheme. In Section 6, we proceed with our experimental evaluation and in Section 7, we conclude our paper by presenting our future research plans.

2 Related Work

With the advent of IoT, numerous devices can interact with their environment and the network in order to collect and process data. In many application domains, data play a central role towards the generation and the provision of knowledge that will support the efficient decision making. Example domains are

financial services, life sciences, mobile services, etc. Apart from the discussed devices, end users may also generate data like tweets, social networking interactions or photos [59]. Through analytics, one can support the efficient decision making having a view on the ‘meaning’ of the collected data. Analytics aim to discover patterns, especially in the case of unstructured data. Various tools for large scale data analytics have been proposed in the literature. The majority of them concern batch oriented systems and they build over Hadoop [28]. A number of research efforts try to provide performance insights to the discussed framework [1], [20], [37]. Researchers provide new functionalities to increase the performance of legacy systems. For instance, the authors of [33] propose the Starfish, a self-tuning tool for large scale data analytics.

The aforementioned processing can be either centralized or distributed, i.e., the processing takes place where data are initially collected. We can also identify the need for streams processing to facilitate the online, (near) real time provision of responses in a set of analytics queries. Queries/tasks allocation and scheduling are important research subjects in multiple domains. Both subjects have a significant impact in the IoT and EC. IoT/EC nodes have limited computational capabilities while being restricted by various energy constraints. It seems that streams processing is the appropriate methodology for delivering real time analytics. ENs can apply a decision making mechanism to process the incoming tasks/queries. They should take into consideration tasks’/queries’ specific characteristics in combination with their current internal status for any decision making. Additionally, nodes may train and update machine learning models locally to serve the incoming tasks/queries. These local models can be aggregated in an upper layer [18]. This approach is appropriate to detect local data shifts and create collaborative learning and model-sharing environments in which local models can quickly adapt to any changes in the collected data.

A widely studied research subject is task scheduling in *Wireless Sensor Networks* (WSNs). Mapping and scheduling should take into consideration energy constraints to secure an efficient execution [8], [61]. A pre-processor and a scheduler can be responsible for the final allocation. The pre-processor tries to identify the energy requirements of the incoming tasks/queries and, based on energy monitoring activities, decides on the final scheduling. In any case, taking into consideration only a single parameter (e.g., remaining energy) in decision making can limit the reasoning capabilities. Another approach is to study a fair energy balance among sensors while minimizing the delay using a market-based architecture [23]. Taking into consideration the defined constraints, nodes may cooperate to conclude the final allocations [7]. Example algorithms involve tasks/queries clustering and node assignment mechanisms based on tasks/queries duplication and migration schemes. The aim is to minimize the execution time, thus, to deliver the final response with minimized latency. A model that could be adopted for such purposes is to cluster the network and build intra-cluster and inter-cluster scheduling relations. An *Integer Linear Programming* (ILP) formulation and a 3-phase heuristic are also adopted to solve the allocation problem in [67]. In [65], the authors propose a modified version of the binary *Particle Swarm Optimization* (PSO). The method adopts

a different transfer function, a new position updating process and mutation for the task/query allocation problem. Another PSO-based solution is presented in [55] which allocates tasks/queries into a number of robots trying to decrease the communication cost. It is worth noticing that PSO-based models suffer from a low convergence rate while they can be trapped to a local optimal especially in complex problems. In [35], the authors present a task/query allocation mechanism of a dynamic alliance based on a Genetic Algorithm to acquire the balance between energy consumption and accuracy. In [16], the authors discuss three algorithms to solve the discussed problem: a centralized, an auction-based, and a distributed algorithm. The distributed algorithm adopts a spanning tree over the static sensors to assign tasks/queries.

The continuous reporting of queries demanding for immediate responses is also the subject of various research efforts. When large scale data applications involve continuous queries, for having near real-time responses, such applications, usually, involve a large number of data partitions. Obtaining a response in near real-time could be very difficult due to limitations defined by the amount of data and the underlying hardware performance. Querying data samples and the provision of progressive analytics is an efficient solution for the described problem [3]. Specific sampling techniques have been proposed [14], [32]. In progressive analytics, the *Approximate Query Processing* (AQP) technique can secure the accuracy in early results and provide the corresponding confidence [14], [21], [54]. Users defining queries are not involved in the process, however, based on the retrieved confidence, they could rely on an intelligent mechanism for handling early results. When accuracy is at acceptable levels (according the specific application domain), the process could be stopped.

An analytics provisioning system is presented in [13]. The system is based on the Prism framework and allows users to communicate samples to the system. Queries are processed over the defined samples. The authors propose Now!, a progressive data-parallel computation framework for Windows Azure. Now! mainly works with streaming engines to support progressive SQL over large scale data. It is worth noticing that the selection of samples affects the statistical error of the final (sub-)dataset, thus, it may have negative consequences in the final decision making. In [17], the authors present an online MapReduce scheme that supports online aggregation and continuous queries. For decreasing the latency of the system, the authors propose to have the Map task sending early results to the Reduce tasks. This mechanism enables the generation of approximate results, which is particularly useful for interactive analytics. In [47], the authors present a continuous MapReduce model. The execution of Map and Reduce functions is coordinated by a data stream processing platform. Latency is improved through a model where mappers are continually fed by data and the retrieved results are transferred to reducers. CONTROL [32] is an AQP system intended to support progressive analytics. Users have the opportunity to refine answers and have online control of processing. DBO [36] is another AQP system capable of calculating the exact answer to queries over a large relational database. DBO can have an insight on the final response together with specific bounds for the accuracy of early results. As more information is processed, the

DBO can provide provide more accurate results. Users can stop the process at any time, if the accuracy level is sufficient.

3 Rationale & Preliminaries

In this section, we provide a description of the interacting entities towards the efficient allocation of the incoming queries and present the problem under consideration. We also setup the basis for solving the problem and apply the envisioned meta-ensemble classification scheme.

3.1 Edge Network Architecture

Consider the setting presented in Figure 1 where N ENs collect contextual data and locally process them in light of producing knowledge; each EN is indexed in the set $\mathcal{N} = \{n_1, n_2, \dots, n_N\}$. Contextual data are either collected by various sensors/end devices (e.g., users’ smartphones) or are generated through local processing.

Definition 3.1 *Contextual data are depicted by the information that provides context (a value for a specific attribute) to an event, entity or a processing activity.*

Every EN has specific characteristics related to its computational power and a limited storage capacity. Hence, only a subset of contextual data could be stored locally. The remaining data are sent to the back end infrastructure present at Cloud. In this context, by supporting local data processing at the EN, it facilitates the minimization of latency in the provision of responses.

Local data processing involves statistical reasoning, inferential analytics, and real-time data management, e.g., estimation of the top- k lists over the incoming data streams [42]). A dataset DS_i is available over which the local data processing takes place. This dataset is continuously updated as fresh data arrive through streams. The i th EN stores multivariate vectors in DS_i , i.e., $\mathbf{x} = [x_1, x_2, \dots, x_L]^T \in \mathbb{R}^L$ (L is the number of dimensions - contextual attributes). As multiple devices may report vectorial data to multiple ENs, replicates among DS_i and DS_j , with $i \neq j$ may be present. The management of potential replicas is beyond the scope of this paper. The whole data at the network edge form the set $\mathcal{DS} = \{DS_1, \dots, DS_N\}$. In each EN, we introduce a local QP responsible to (i) receive a stream of analytics queries (e.g., estimating the regression plane among contextual variables within a time frame) from end users (e.g., applications, data analysts); (ii) execute them over DS_i and (iii) send the results back to the requestor.

Definition 3.2 *An analytics query is a request for information responded by meaningful patterns found in the available dataset while being extracted based on a scientific process.*

We associate the i th QP with the i th EN, thus, we have N ENs/QPs in the set $\mathcal{QP} = \{qp_1, qp_2, \dots, qp_N\}$. We adopt a queue in every QP which can handle a maximum number of queries. Without loss of generality, we consider that queues adopted in QPs have the same length. Each QP has specific characteristics encoded in the set $\mathcal{C}_i = \{c_{i,1}, c_{i,2}, \dots, c_{i,m}\}$. For instance, $\mathcal{C}_i = \{l, s\}$ with l representing the current load and s depicting the speed of the corresponding QP.

Definition 3.3 *The load is the quantity of analytics queries that can be carried at one time by an EN.*

Definition 3.4 *The processing speed of an EN is its rate of execution of analytics queries.*

l can be easily estimated through the current number of queries waiting in the corresponding queue, while s indicates the throughput of a QP, i.e., the number of queries responded in a given time unit.

In the upper layer, shown in Figure 1 (i.e., Fog/Cloud), there is a federation of QCs. QCs play the role of endpoints for end users or applications and achieve efficient provision of responses to incoming queries. QCs have direct access to the ENs, thus, to their corresponding QPs and they interact with them to get responses for the incoming queries. QCs, after receiving a query, should be able to find the most appropriate subset of QPs for allocating the query. Afterwards, QCs should aggregate the partial results to derive the final aggregated response, which will be delivered to end-users/applications. With the term *appropriate*, we refer to the subset of those QPs that, at the specific time the query is issued, they exhibit characteristics that will facilitate its *efficient execution* and return the result in the minimum turn around time. The efficient execution is realized through the realization of various parameters that should be optimal for any allocation. For instance, the response time should be limited, the outcomes should match against the query constraints, the response time should fulfil the pre-defined time constraints (if set) and so on and so forth. It should be noted that QCs, through their continuous interaction with the ENs and their QPs, can maintain historical performance data as well as the statistics of data present in each EN. Based on this context, QCs obtain a holistic overview on the current status of QPs and the data that they are stored locally in ENs.

The incoming queries are represented via a stream: $Q = \{q_1, q_2, \dots, q_j\}$. At time instance $t \in \mathbb{T} = \{1, 2, \dots\}$ a query q_t arrives to a QC. q_t belongs to a specific *query class*, exhibiting specific characteristics, i.e., $C^q = \{c_1^q, c_2^q, \dots, c_{|C^q|}^q\}$. Suppose that all queries have the same number and types of characteristics. For instance, $C^q = \{p, a\}$ where p stands for the computational complexity and a depicts a deadline for delivering the final result.

Definition 3.5 *The computational complexity of an analytics query represents the amount of the required resources to execute it.*

If we focus on a database management system, we can identify the following query classes: (i) *Selection Queries*: The main representative of such queries is the SELECT command. This type of queries aim to provide data (tuples) that fulfil specific conditions; (ii) *Modification Queries*: Such queries are adopted to perform changes/modifications in the underlying data, e.g., UPDATE. In most of the cases, these queries are demanding and require significant computational time and resources; (iii) *Aggregation Queries*: Usually, they are executed over other queries and apply algebraic operators, e.g., AVG, SUM, MIN, MAX, etc. According to [30], the generic characteristics of queries are: (i) the *type* of the query, e.g., repetitive, ad-hoc; (ii) the query *shape*; and (iii) the *size* of the query, e.g., simple, medium, complex. Based on the characteristics of each query class, specific execution plans could be defined in the form of processing trees [30]. In [62], one can find a study on the calculation of queries complexity, however, with the focus being in the underlying databases. In the database community, the complexity of a query is, usually, measured through the resources required by the database server for executing it. The most significant parameters for doing that is to focus on the space and time required for executing the query during the optimization phase. The query optimizer compares different execution strategies and selects the one with the least expected cost (or the maximum expected reward). In our context, we want to adopt additional parameters not being bounded to the internal processes of any data management system but being aligned with the needs of an EC setting. We focus on the upper layer of the aforementioned architecture and try to limit the time required for initiating the execution of a query. We extend the aforementioned characteristics' list and incorporate new, 'high level', parameters that depict the complexity and the need for instant response.

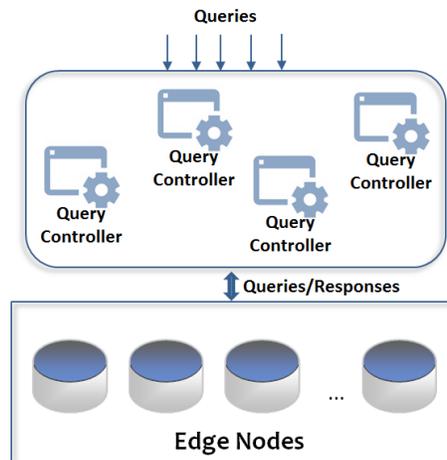


Figure 1: The connection of query controllers and edge nodes.

3.2 Problem Definition

For the description of our problem let us focus on an individual QC attached to a stream of queries. Suppose that at time t , a query q_t arrives at the QC demanding for immediate processing. The QC should be based on q_t 's and ENs characteristics to take the appropriate decision and find the best subset of nodes to engage for the execution of q_t . Among others, we are interested in data constraints as dictated by q_t . For instance, in case of an SQL query, we focus on the WHERE clause. These constraints define the conditions that should be met when we retrieve data to construct the final response for q_t . To easily depict the discussed constraints, we represent q_t as a $2L$ -dimensional vector

$$\mathbf{w} = [\{min_1, max_1\}, \dots, \{min_L, max_L\}]^T \in \mathbb{R}^{2L} \quad (1)$$

such that $\{min_i, max_i\}$ are the minimum and the maximum values defined for the i -th dimension (attribute). Constraints should be matched against data present in ENs where, potentially, q_t is directed to be executed. After the reception of q_t , the QC creates N context vectors (one for each QP) referring to q_t 's characteristics and the current status of ENs/QPs. Context vectors have the following form:

$$\mathbf{v}_i = \langle o, a, r_i, l_i, s_i \rangle \quad (2)$$

In addition, o is the q_t 's expected complexity (elaborated later), a is the deadline set for q_t , r_i is the *information relevance* of q_t with DS_i (the dataset present at the i th EN) and l_i and s_i are the load and the speed of the i th QP, respectively. The aforementioned vectors represent the minimum sufficient statistics for both an incoming query and each EN/QP based on which the QC should decide the final allocation. It should be noted that context vectors are easily concluded in short time either using: (i) pre-defined values (e.g., a is set during the reception of q_t , s_i is defined for each EN beforehand - it depends on the internal characteristics of QPs); (ii) simple calculations (e.g., l_i can be extracted by the number of queries waiting for processing in the corresponding queue); or (iii) our proposed models (e.g., for concluding o and r_i). Overall, context vectors are fed into our proposed ensemble scheme for predicting the appropriateness of each EN/QP to be the host of each query.

Problem: *Given an incoming query q_t represented by \mathbf{w} and the associated context vectors $\{\mathbf{v}\}_{i=1}^N$, predict the most appropriate subset of QPs that should be engaged for the execution of that query.*

In the remainder, we elaborate on the creation of the context vectors. It is worth mentioning that QCs can have: (i) access to the statistical synopses (digests) of data present in ENs (i.e., the QC exploits the minimum sufficient statistics of each dataset); (ii) access to statistical data related to the performance of ENs. Based on the above, our mechanism is, both, *performance-aware* and *data-aware*. We deal with the estimation of the ability of a QP to efficiently execute a query in the minimum time while delivering the outcomes that perfectly match to queries' constraints.

4 Processors Characteristics and Queries Complexity

4.1 Query Processor Speed & Load

As noted, ENs maintain a queue where the incoming queries wait for processing. The size of the queue is adopted such that it can deliver up to l_i queries per time unit, i.e., the percentage of the maximum load that can be afforded by the i th EN. Without loss of generality, we get $l_i \in [0, 1]$ given that a maximum queue size is adopted for such purposes. When $l_i \rightarrow 1$, the EN experiences a high load. The load is also connected with the throughput and the queries reporting rate. Additionally, s_i depicts the speed of processing. A resource demanding query, e.g., a join query, potentially requires more execution time and resources compared with a less resource demanding query like simple select commands. s_i is, then, directly affecting the throughput, i.e., the number of queries executed in a time unit and assumes values in $[0, 1]$. When $s_i \rightarrow 1$, the i th EN exhibits the best possible speed approaching the maximum theoretical performance. When $s_i \rightarrow 0$, the i th EN delays in the delivery responses to the the incoming queries.

4.2 Query Complexity

The classification of the computational complexity o of q_t and its effect on l and s is the key element in our scheme. Various research efforts study the complexity of queries [6], [60], [62]. For realizing o , we rely on our previous effort presented in [40]. We focus on a simple, however, efficient mechanism that will deliver the final outcome in the minimum time. We consider that $|\Theta|$ complexity classes are available, i.e., $\Theta = \{\theta_1, \theta_2, \dots, \theta_{|\Theta|}\}$. θ_i is aligned with the complexity performed by the operations of q_t as required for producing the final result. We assign q_t to a complexity class θ based on a typical classification task. The final complexity is defined based not only on quantitative characteristics, e.g., number of constraints and conditions, but also on qualitative characteristics, e.g., type of operations. For handling this complicated process, we introduce a fuzzy logic based approach and define a *Fuzzy Classification Process* (FCP). The FCP evaluates the membership of q_t in each of the pre-defined complexity classes. Hence, we can obtain an estimate of the computational burden added to a candidate EN. The FCP is executed over a set of historical executed queries along with their corresponding classes. This means that QCs maintain the previous queries set realized by past interactions with ENs being dynamically updated as new requests for processing arrive. The aforementioned set can be concluded by servers' logs reported after each processing activity. In any case, the creation of the historical set of queries and the pairing process with the available complexity classes is beyond of the scope of this work. We deliver a scheme that can be adopted in processing of real, unknown queries providing the final result in limited time. When focusing on a query decomposition process (it is the theoretical ground coming from the database community) to calculate the complexity, we need from 3.0 to 22.65 seconds depending on the query and

the underlying DBMS [48]. This amount of time should be added to the time required for the allocation of a query and the waiting time in the queue of the selected EN. In our work, we specifically focus on the minimization of the allocation time together with the reduction of the waiting time. We try to limit the time at every part of the envisioned processing, i.e., the initial allocation of q_t combined with models proposed by the database community aiming at query execution optimization.

For evaluating θ for every q_t , we adopt widely known similarity techniques instead of relying on a machine learning model that requires a training phase. We try to minimize the complexity of our scheme, proposing the use of a simple, fast, however, efficient model that delivers the final outcome in the minimum time. Let the available dataset of training queries be DS_Q containing tuples in the form $\langle p_k, \theta_k \rangle$, $k \in \{1, 2, \dots, |DS_Q|\}$. p_k represents q_t 's statement along with its complexity class $\theta_k \in \Theta$. An example statement could be: `SELECT NAME, PRICE FROM STOCKS WHERE PRICE < = 100 AND PRICE >= 10`. We, then, develop a function f over q_t and, based on DS_Q , deliver a vector encoding the similarity of q_t with every complexity class in Θ :

$$f(q; DS_Q) \rightarrow \mathbf{q}^s \in [0, 1]^{|\Theta|} \quad (3)$$

\mathbf{q}^s 's components assume values in $[0, 1]$ demonstrating the degree of membership of q_t to each complexity class, thus, forming the basis of our FCP. For instance, consider the vector $\mathbf{q}^s = [0.2, 0.8, 0.3]^\top$ given the complexity classes: $\Theta = \{\theta_1 = O(n \log n), \theta_2 = O(n), \theta_3 = O(n^2)\}$. In this example, \mathbf{q}^s shows that q_t belongs with 20% to θ_1 , 80% to θ_2 and 30% to θ_3 .

For calculating every component of \mathbf{q}^s , we adopt a set of metrics that deliver the similarity between vectors. The interested reader can refer in [45] for more details. We propose an ensemble scheme for evaluating the similarity of q_t with every tuple $\langle p_k, \theta_k \rangle \in DS_Q$. Recall that tuples present in DS_Q depict a set of training queries along with their complexity classes. Our aim is to find the similarity of q_t with every training query merging the results in a subsequent step to deliver the final q_t 's similarity with each of the pre-defined complexity classes. The main theoretical foundation behind the adoption of an ensemble similarity scheme is related to the special requirements of the problem and the need for avoiding any wrong decisions when relying on a single metric. The ensemble scheme 'combines' the 'opinions' of multiple metrics while the aggregated result is the one that will support the final outcome. Our ensemble scheme utilizes a set $E = \{e_1, e_2, \dots, e_{|E|}\}$ of similarity metrics. E can involve the Hamming distance [53], the Jaccard coefficient [51], and the Cosine similarity [53]. Any distance metric available in the literature could be transformed to depict the similarity between q_t and the p_k from the tuple $\langle p_k, \theta_k \rangle$. For instance, if e_d is the Euclidean distance between q_t and p_k , their similarity can be defined as $\frac{1}{1+e_d}$. The adopted similarity metrics are applied on each tuple 'pre-classified' to θ_k aggregated to determine the k -th component q_k^s of the complexity similarity vector. Formally the '2D aggregation' (see Figure 2 - retrieved by [40]) is calculated as follows: $q_k^s = \Omega(\omega \{e_i(q_t, \langle s_k, \theta_k \rangle)\}, \forall i, \forall \langle p_k, \theta_k \rangle$. ω realizes the envisioned ensemble

similarity scheme while the aggregation operator Ω produces the q_k^s over multiple ω values.

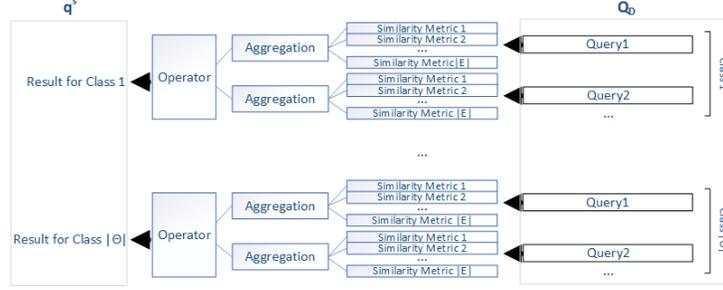


Figure 2: The envisioned similarity aggregation process.

For ω , we consider that every individual result (i.e., $e_i(q_t, \langle p_k, \theta_k \rangle)$) represents the membership of q_t to a ‘virtual’ fuzzy set. We have $|\mathcal{E}|$ membership degrees combined to get the final similarity for each individual tuple. For instance, if we get $e_1 = 0.2$, $e_2 = 0.5$ and $e_3 = 0.3$, q_t ‘belongs’ to the e_1 fuzzy set by 0.2, to the e_2 by 0.5 and to the e_3 by 0.3. ω is a *fuzzy aggregation operator*; an $|\mathcal{E}|$ -place function $\omega : [0, 1]^{|\mathcal{E}|} \rightarrow [0, 1]$ that takes into consideration the membership to every fuzzy set and returns the final value. Aggregation operators are well studied in various research efforts [24], [34]. Through a high number of experiments [24], [34], the Einstein product, the algebraic product, the Hamacher product [34] and the Schweizer-Sklar metric [24] are identified to have the best performance. In the proposed model, we adopt the Hamacher product as it gives us more opportunities to ‘tune’ the result through the parameter $\alpha \geq 0$. The final ω value is defined as:

$$\omega = \frac{\acute{e} \cdot \grave{e}}{a + (1 - a)(\acute{e} + \grave{e} - \acute{e} \cdot \grave{e})}$$

where \acute{e} and \grave{e} are two similarity values. As similarity metrics may ‘disagree’, we propose the use of the top- n similarity values based on their significance level. The *Significance Level* (SL) depicts if a similarity value is ‘representative’ for many other outcomes. We borrow the idea from the density based clustering [31] where centroids of the detected clusters are points that ‘attract’ many other data features around them in close distance. We propose the use of the radius γ and calculate the SL for each similarity result as follows: $SL_{e_i} = \frac{1}{1 + e^{-(\delta_1 |d(e_i, e_k) \leq \gamma| - \delta_2)}}$, $\forall i$, where δ_1 and δ_2 are parameters adopted to smooth the sigmoid function. With the sigmoid function, we eliminate the SL of similarity values with a low number of ‘neighbors’ in the radius γ . The final results are sorted in a descending order of the SL and the top- n of them are processed with the Hamacher product to deliver the final ω .

The Ω operator builds over ω values produced for each tuple in Q_D classified in θ_k . Let $\omega_1, \omega_2, \dots, \omega_m$ are those values. For their aggregation, we rely on a

Quasi-Arithmetic mean, i.e., $q_k^s = \left[\frac{1}{m} \sum_{i=1}^m \omega_i^\alpha \right]^{\frac{1}{\alpha}}$ where α is a parameter that ‘tunes’ the function. When $\alpha = 1$, the function is the arithmetic mean, when $\alpha = 2$, it is the quadratic mean and so on and so forth. After calculating the final values for each θ_k (i.e., realized by Ω_k), we get $\mathbf{q}^s = \langle \Omega_1, \Omega_2, \dots, \Omega_{|\Theta|} \rangle$.

4.3 Distance Between Queries and Datasets

We propose a distance model aiming at concluding r_i adopted in the context vector \mathbf{v} . Recall that r_i is adopted to depict the similarity between \mathbf{w} and data present in ENs. We consider that the dimensions of the collected/stored vectors are not correlated. At pre-defined intervals, ENs send to QCs the statistics of local data expressed by two vectors, i.e., the vector of means $\mu = \langle \mu_1, \mu_2, \dots, \mu_L \rangle$ and the vector containing the standard deviation for each dimension, i.e., $\sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_L \rangle$. Again, we have to recall that \mathbf{w} represents the intervals depicted by the constraints of q_t . Our intention is to find the overlapping between two sets of intervals, i.e., intervals defined by \mathbf{w} and intervals represented by data statistics (the combination of μ and σ). Actually, μ and σ can define the confidence interval for each dataset as follows. We have $\mu \pm z \cdot \frac{\sigma}{|DS_i|}$ for all the adopted dimensions with $|DS_i|$ depicting the cardinality of the corresponding dataset. z represents the z-value retrieved by the standard normal (Z-) distribution for our desired confidence level. The area between $-z$ and $+z$ is, approximately, the confidence percentage. For instance, for $z = 1.28$, the area between -1.28 and $+1.28$ is, approximately, 0.80.

Based on the above analysis, we proceed to calculations related to the similarity between \mathbf{w} and N vectors, i.e.,

$$\mathbf{f}_i = \langle \{ \mu_{i1} - z \cdot \sigma_{i1}, \mu_{i1} + z \cdot \sigma_{i1} \}, \dots, \{ \mu_{iL} - z \cdot \sigma_{iL}, \mu_{iL} + z \cdot \sigma_{iL} \} \rangle, \forall i$$

2. This way, we conclude N similarity values, one for each dataset, thus, we deliver N context vectors \mathbf{v} . For deriving the final r_i , we have to find the final similarity between L intervals, i.e., \mathbf{w} and \mathbf{f}_i . Typical distance/similarity measures (e.g., the Euclidean distance) cannot efficiently manage cases where \mathbf{w} can be completely contained in \mathbf{f}_i . We rely on the research performed in [29], where a study on calculating the distance over interval data is provided. Based on the discussed metrics, we propose the use of the overlapping metric ψ_k to finally deliver r_i as follows: $r_i(\mathbf{w}, \mathbf{f}_i) = h(\psi_{ik}), \forall k \in \{1, 2, \dots, L\}$. In this equation, we propose the use of the aggregation function h responsible to aggregate L distance results. Our calculations are exposed by the following equation (applied for each dimension of the adopted dataset): $\psi_{ik} = 1 - \frac{\|\mathbf{w} \cap \mathbf{f}_i\|}{\min\{\|min_k, max_k\|, \|\mu_{ik} - \sigma_{ik}, \mu_{ik} + \sigma_{ik}\|\}}$ where

$$w_k \cap f_{ik} = \begin{cases} (\max\{min_k, \mu_{ik} - \sigma_{ik}\}, \min\{max_k, \mu_{ik} + \sigma_{ik}\}) & \text{for } \max\{min_k, \mu_{ik} - \sigma_{ik}\} < \min\{max_k, \mu_{ik} + \sigma_{ik}\} \\ 0 & \text{Otherwise} \end{cases} \quad (4)$$

The previous equation defines the interval depicting the ‘area’ between $\{min_k, \mu_{ik} - \sigma_{ik}\}$

²We consider that every σ_i vector depicts the result $z \cdot \frac{\sigma}{|DS_i|}$

and $\{\max_k, \mu_{ik} + \sigma_{ik}\}$. For h , we adopt the Quasi-arithmetic mean [12], i.e.,

$$r_i = \left(\frac{1}{L} \sum_{k=1}^L (\psi_{ik})^\alpha \right)^{1-\alpha}$$

Based on α , we can alter the results derived by the Quasi-arithmetic mean, e.g., if $\alpha = 1$, r_i is calculated based on a ‘simple’ mean.

5 The Ensemble Learning Scheme

We propose the use of a *meta-ensemble learning scheme* that will deliver the final allocation of the incoming queries on top of multiple other ensemble schemes. Our ensemble scheme tries to ‘optimize’ the allocation of every q_t to the available ENs. Let $g()$ be a function that results the matching degree between q_t and an EN. The higher the $g()$ realization is, the higher the matching becomes. Our selection mechanism deals with the maximization of $g()$ based on queries and ENs characteristics. In terms of the description of an optimization problem that fits in our scenario, we can denote that our ensemble mechanism tries to maximize $g(\mathbf{w}, \{\mathbf{v}\}_{i=1}^N)$ subject to our models’ results for delivering l_i , s_i , \mathbf{q}^s and r_i . Actually, our ensemble model is responsible to detect the most efficient allocations over a complicated ‘reasoning’ mechanism that adopts multiple ML schemes.

Any ensemble method aims at the creation of multiple learning models combined to deliver the final decision. The ensemble approach manages to provide better results when adopted for prediction being compared with other incremental techniques [2]. Some approaches in ensemble learning are [9]: (i) Voting; (ii) Stacking; (iii) Bagging; (iv) Boosting. To have the successful application of the ensemble learning model, we have to apply diversity of the obtained results. The required diversity is achieved by [9]: (i) the diversity of the models, i.e., we have to use different algorithms or the same algorithms with different parameters; (ii) the diversity of data, i.e., the training data should be different for each of the adopted models. The most critical issue is the combination of the results. Ensemble schemes attract the attention of researchers working in the machine learning domain based on the view that ensembles are often much more accurate than the individual classifiers [22].

5.1 The Adopted Classifiers

In our research, we rely on a wide range of classifiers. and build on the result of already defined ensemble schemes to provide our meta-ensemble model. We adopt widely known ensemble learning models to provide a more powerful learner. We adopt the three ‘basic’ ensemble models, i.e., (i) the AdaBoost model Y_1 ; (ii) the Stacking model Y_2 ; (iii) the Bagging method Y_3 . Our meta-ensemble scheme receives the results of the aforementioned ‘sub-ensemble’ models (i.e., Y_1, Y_2, Y_3) and delivers the final allocation decision Y . It should be

noted that the aforementioned individual ensemble schemes are already focusing on combining the individual, ‘single’ learners, thus, we deliver another, additional, layer of processing. As the aforementioned schemes are ‘high-level’ ensemble models, they can be applied over any individual learning algorithm. Below, we present the list of the adopted individual learners incorporating into our ‘reasoning’ mechanism different types of schemes, e.g., decision trees, probabilistic models, neural networks. This exhibits the ability of our mechanism to be unbounded of the type of individual learners and its strength to combine various outcomes for the discussed allocations. Our meta-ensemble model performs a two-level aggregation of the outputs retrieved by the individual schemes when applied over the characteristics of queries and ENs trying to derive the most efficient allocations. We have to notice that the outputs of the adopted models should be in the same interval to be easily aggregated for concluding the final decision.

In [56], the interested reader can study the theoretical background behind the adoption of ensemble machine learning schemes. In short, there are three main theories that explain the effectiveness of ensemble models. The first considers ensemble models under the perspective of large margin classifiers [50]. Characteristics of such schemes are the enlargement of margins, enhancements of the generalization capabilities of Output Coding [5] and boosting based ensemble algorithms [58]. The second theoretical approach deals with study of the classical bias/variance decomposition of the error [27], showing that ensembles can reduce variance [10], [46] or both bias and variance [44], [11]. Finally, the third theory is a stochastic discrimination theory focusing on a set-theoretic abstraction to remove all the algorithmic details [38]. Adopting such an abstraction, we are able to ‘see’ the classifiers as a combination of subsets of points of the feature space and their decisions are also point sets. The set of classifiers are, then, just a sample into the power set of the feature space.

The Adaptive Boosting (AdaBoost) model [25] conveys the basic processing of the Boosting model while combining, in an adaptive way, multiple base learners. Initially, the Boosting model builds a first, weak classifier and, accordingly, a succession of models are built iteratively, each one being trained on a dataset in which points misclassified by the previous model are given more weight. All the adopted models are weighted by their success to provide an accurate classification result and their outputs are combined through voting. In any case, the same training dataset is used over and over again. The Stacked Generalization (Stacking) model [64] is another meta-learner that aims at combining models of different types. Initially, it splits the training dataset in two disjoint sets and, then, trains several base learners on the first part of the dataset. Accordingly, it tests the base learners with the second part of the training dataset. Afterwards, it gets predictions in the testing phase as the input and correct responses as the output to train a higher level learner. The Bootstrap Aggregation (Bagging) [10] aims at incorporating multiple versions of the training set through the use of sampling with replacement. Every produced dataset is adopted to train a different learning model. The final output is delivered by averaging the individual outputs or voting when the final results involves a classification process.

Bagging is efficient only when using unstable non-linear models because a small change in the training set can cause a significant change in the model.

In the aforementioned meta-learning schemes, we incorporate the following learning algorithms/base learners (to have the necessary diversity in the selection of algorithms): (i) C4.5 decision tree; (ii) Random tree; (iii) Naive Bayes model; (iv) Bayesian Network; (v) Multinomial Naive Bayes model; (vi) Random Forest; (vii) Logistic Model Tree; (viii) REPTree model; (ix) JRip algorithm; (x) Multilayer Perceptron.

5.2 Multi-Classifer Decision Fusion

Our meta-ensemble model is adopted to realize the aforementioned function $g()$, i.e., to deliver the matching and a ranking of the available ENs for a query q_t . The proposed model gets as input ENs and q_t 's characteristics exposed by \mathbf{w} and $\{\mathbf{v}\}_{i=1}^N$ and results the most efficient allocation. The efficiency of the allocation is realized by the 'appropriateness' of ENs to host q_t and provide the outcome in the minimum time with the best possible performance. The performance is depicted by the matching between q_t 's constraints (i.e., \mathbf{w}) and the available datasets as well as the ability of ENs to minimize the response time (i.e., the execution of q_t should start immediately - we target to a low load - while being completed as soon as possible - we target to a high processing speed). For detecting the most efficient assignment, we adopt the aforementioned ensemble schemes, we get their 'opinion' for every potential allocation and combine them through the proposed meta-ensemble model to conclude the final decision.

The meta-ensemble learning scheme is based on a training dataset \mathcal{TD} where tuples of context vectors are present. Each training tuple is accompanied by the appropriate decision for the 'virtual' node that represents. More formally, every training tuple has the form $v^D = \langle o^{TD}, a^{TD}, r^{TD}, l^{TD}, s^{TD}, B \rangle$. The first part of a training tuple is related to the aforementioned context vectors while the second part (i.e., B) is related to the final classification result. We consider a binary classification setup with two classes, i.e., $B^1 = 1$: *Allocate* and $B^2 = \overline{B^1} = 0$: *Do not allocate*.

The three aforementioned ensemble schemes, i.e., the AdaBoost, the Stacking and the Bagging methods are trained based on \mathcal{TD} and are adopted to generate the envisioned results Y_1, Y_2, Y_3 . When Y_1, Y_2, Y_3 are produced, we have to combine them to get the final result Y based on which, the final decision is concluded. In this effort, we adopt three approaches. The first is a very 'strict' scheme to produce Y relying on the Boolean model originated in Information Retrieval [49]. Based on this model, the final result is delivered through a simple conjunction, i.e., $Y = \prod_{i=1}^3 Y_i$. This means that if at least one of the aforementioned schemes 'votes' against the allocation of q_t to a specific EN/QP, the final result will be the class B^2 . To conclude the class B^1 , all the classifiers should agree upon this decision. The second approach is a majority voting scheme where the majority of classifiers conclude the final result. Hence, when $count\{Y_i = B^1\}$ is at least 2.0, $Y = B^1$ as well. Our future research plans involve the definition of a more complex methodology for combining the

‘opinion’ of the adopted ensemble classification schemes. In [26], the interested reader can find a set of aggregation techniques for classifiers outputs. It should be noted that our decision to incorporate the ‘strict’ technique together with the majority voting scheme deals with our intention to use two representative schemes that are capable of delivering the final result in limited time requiring limited resources (no additional data structures and variables that consume memory).

Each tuple in \mathcal{TD} represents a combination of the adopted parameters involving queries and processors characteristics. For the delivery of the ENs where q_t will be allocated, we apply the one-over-all (OVA) methodology [31]. It consists of a widely adopted technique for multiclass classification with high performance. In [57], a number of optimization algorithms are compared with the OVA model. Although these approaches may have a theoretical interest, it does not appear that they offer any advantage over a simple OVA scheme. In addition, OVA compared to other techniques does not require any additional time for training. In our scenario, N classes are available; one for each EN/QP. Hence, we adopt N binary classifiers built by the above described meta-ensemble model. Actually, we apply N times the same meta-ensemble binary classifier for deciding in which nodes q_t could be allocated. The classification of an unknown vector v^D concerns a voting scheme. If the adopted classifier positively predicts the i th node for q_t , the i th EN gets a vote. If the result is negative (i.e., class B^2) all the remaining nodes except i get a vote. The node(s) with the most votes is selected to host q_t . In the case of ties, we rank the nodes based on their load.

6 Experimental Evaluation

We perform a high number of simulations adopting synthetic traces and datasets found in the literature. In addition, we compare the proposed scheme with our previous work. The aim is to reveal the advantages and the drawbacks of our model when adopted for the allocation of analytics queries. In the upcoming sub-sections, we describe our experimental setup and the assessment of our model.

6.1 Performance Metrics & Experimental Setup

We perform our experiments with a custom software written in Java that contains a number of classes for the simulation of the adopted nodes and datasets present in them. Initially, we deal with the accuracy of the proposed FCP for detecting the correct complexity class of a query. For this, we define the metric v as follows:

$$v = \frac{Q_C}{|Q|} \quad (5)$$

where Q_C is the number of the correct predictions. Obviously, we want $v \rightarrow 1$ to enjoy the best possible performance. We also focus on the time required for

the allocation of each incoming query that affects the final throughput of *QCs*. We define the metric R that represents the discussed throughput, i.e.,

$$R = \frac{|Q|}{\sum_{i=1}^{|Q|} T_i^c} \quad (6)$$

where T_i^c is the conclusion time (in ms) for the i th query. The *conclusion time* is the time required to allocate a query to one of the available nodes just after its reception. R depicts the amount of queries executed by a *QC* in a time unit (i.e., ms). The higher the R is, the higher the performance of the proposed model becomes.

We also adopt parameters l^* and s^* depicting the load and the speed of the selected node, respectively. Based on these parameters, we define the metrics D_l and D_s . Both of them represent the ‘distance’ of l and s of the selected node from the optimal load and speed in the entire group. As the optimal load, we define the lowest load (l_{min}) observed in the available nodes at the decision time. In a similar way, as the optimal speed, we define the highest speed (s_{max}) among the available nodes at the decision time. Through the adopted performance metrics, we try to reveal if the proposed model is capable of selecting the best possible node at the decision time. If this is true, our scheme allocates the incoming queries to nodes that will return the final response in the minimum time. The following equations hold true:

$$D_l = l^* - l_{min} \quad (7)$$

$$D_s = s_{max} - s^* \quad (8)$$

We compare the proposed Conjunctive Scheme *CS* and the Majority Voting Scheme (*MVS*) with our previous work [39], [41], [43] (Model 1 - M1, Model 2 - M2, Model 3 - M3, respectively). We have already discussed the novelty of the current framework in contrast to our previous schemes, thus, for providing a holistic comparison in our experimental evaluation, we deliver numerical results. We evaluate our model for different realizations of N getting $N \in \{10, 50, 100, 500\}$. For each experiment, we consider the execution of 1,000 queries and take results for the aforementioned models. For the type of the incoming queries and the delivery of their complexity class, we rely on the dataset and the methodology presented in [40]. For the remaining parameters, i.e., query vectors, a , l , s and the data reported in nodes, we rely on a set of traces. Our data are retrieved by the following traces:

- Dataset 1. A synthetic trace based on the Uniform distribution defining a for each query. For the generation of each value, we consider an maximum a equal to 10.
- Dataset 2. A synthetic trace based on the Uniform distribution delivering values for query vectors, the data present in each node, l and s .

- Dataset 3. A synthetic trace based on the Gaussian distribution delivering values for query vectors, the data present in each node, l and s .
- Dataset 4. The trace reported by [4]. From this dataset, we adopt the processor utilization dimension that depicts the percent of time that threads are running in a processor. The data are incorporated to our simulator for providing the realizations of l . The remaining parameters take values as described for Datasets 2 & 3.

6.2 Performance Evaluation

Initially, we report on the accuracy of the proposed FCP for delivering \mathbf{q}^s . Recall that values present in \mathbf{q}^s depict the ‘membership’ of each query to the available complexity classes. Executing the FCP for the queries dataset reported in [40], we get $v = 1.0$ for a threshold less than or equal to 0.8. Over this threshold, we consider that the query has the complexity of the class corresponding to the index of the discussed value in \mathbf{q}^s . As already described, it is difficult to assign an individual complexity class for each query, thus, the maximum value of \mathbf{q}^s represents the class that exhibits the highest similarity with q_t . If we get the threshold equal to 0.9, $v \rightarrow 0.55$. It is worth noticing that the maximum value in \mathbf{q}^s in the entire set of our experiments is over 0.8.

We report on the probability density estimate (*pde*) of the adopted performance metrics achieved by the proposed model. In Figure 3, we see our results for T^c (the time required in ms for the allocation of a query). We observe that the uniformity of the data (upper figure) leads to the uniformity of the discussed performance metric. In contrast, the use of the Gaussian distribution (lower figure) leads to results affected by the number of the available ENs. In the latter case, a low number of ENs (e.g., $N \ll 500$) leads to a low T^c . In these scenarios, the throughput of QCs is high managing to allocate multiple queries in a limited time interval. It should be noted that the discussed results refer in the *CS*. If we pay attention on the *MVS* (see Figure 4), we observe the uniformity of T^c as well. However, the higher the N is the higher the T^c becomes. This is natural, as QCs should process multiple context vectors as inputs into the matching process. In any case, if we consider that the provided results are in the scale of ms, we can conclude that the proposed mechanism, even if it involves multiple classification models, manages to deliver the final allocation in a limited time interval.

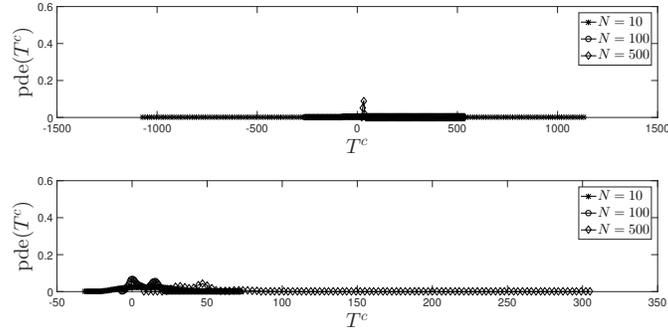


Figure 3: Pde of the allocation time for CS (up: Uniform distribution, down: Gaussian distribution).

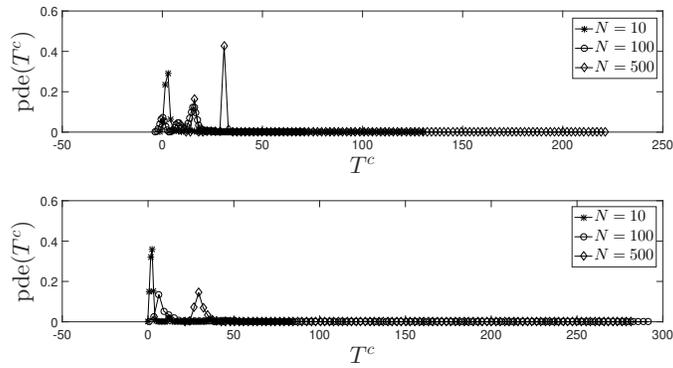


Figure 4: Pde of the allocation time for MVS (up: Uniform distribution, down: Gaussian distribution).

In Figures 5 & 6, we present our results related to l^* for both *CS* and *MVS*. The higher the N is, the lower the l^* becomes. This is more intense in *MVS*. A high number of ENs give us more opportunities to seek for the best possible node. The provided classification scheme and the OVA approach manage to conclude nodes that exhibit low load, thus, it could deliver the final result in a limited time (however, this finally depends on the complexity of the process demanded by the query). A low number of ENs may lead to a higher load compared to the previous scenario in both models, i.e., *CS*, *MVS*.

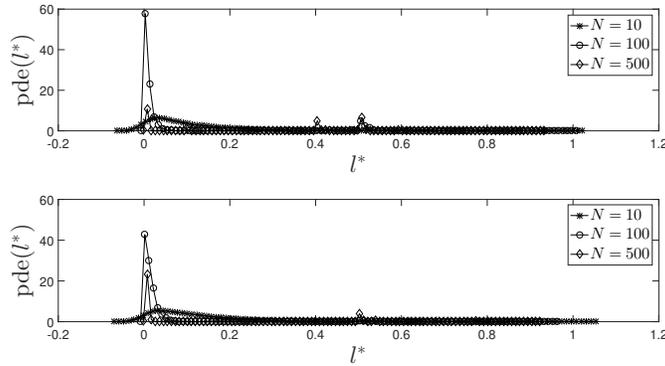


Figure 5: Pde of l^* for CS (up: Uniform distribution, down: Gaussian distribution).

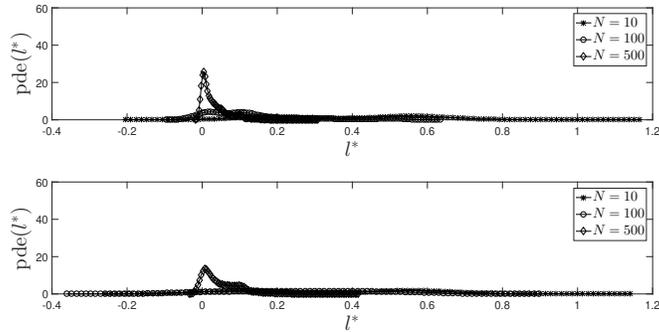


Figure 6: Pde of l^* for MVS (up: Uniform distribution, down: Gaussian distribution).

In Figure 7, we plot the throughput of the proposed mechanisms. The throughput is minimized when we have to deal with multiple nodes; a natural consequence of the involvement of multiple processes until the final decision. The *CS* exhibits the best performance among the models, however, the difference between them is limited when $N \rightarrow 500$. When $N = 10$, the *CS* exhibits the worst performance when the Uniform distribution is adopted to produce the values for our parameters. In terms of the number of queries allocated in a second, we get that the *CS* can allocate 420 to 35 queries while the *MVS* can allocate 430 to 16 queries.

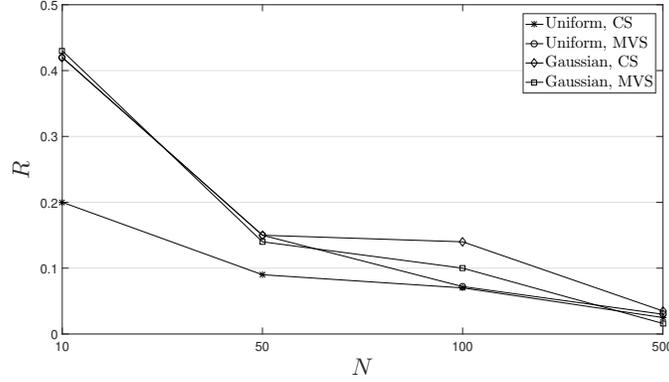


Figure 7: The throughput of the proposed model.

In Figure 8, we plot our results related to the optimality of the proposed model as far as D_l concerns. The *CS* outperforms the *MVS* except one experimental scenario ($N = 10$). Recall that if at least one of the adopted ensemble schemes rejects the allocation, the query is not classified to the corresponding ENs. This means that the *CS* tries to take a decision under unanimity that is not the case in the *MVS*. The significant is that as $N \rightarrow 500$, both models result in a limited distance with the optimal solution. For instance, the *CS* results in a distance equal to 0.016 & 0.022 for the Uniform and the Gaussian distributions, respectively. The *MVS* results in a distance equal to 0.038 & 0.110 for the Uniform and the Gaussian distributions, respectively. In any case, these numbers exhibit the efficiency of the proposed scheme keeping in mind that we plot the average D_l throughout our experiments.

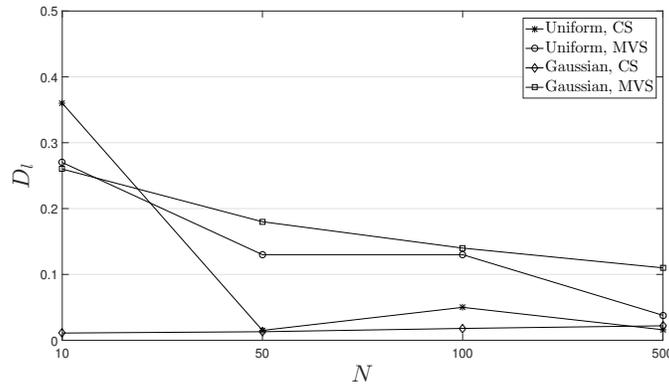


Figure 8: The load of the selected node as delivered by the proposed model compared to the lowest available load.

The next set of our experiments deal with the speed of the selected node. We should note that the speed of processors is randomly selected as depicted by the aforementioned traces with a maximum value equal to 10. In Figure 9, we present the relevant results. The distance from the optimal speed in the group increases as N increases. Especially when $N = 500$, the distance is high. We have to compare these results with the results derived for D_l . The increased number of nodes positively affects the load of the selected node while it negatively affects the speed of the selected EN. The proposed model pays more attention on the minimization of load, however, we should note that this is affected by the training set adopted to prepare the ensemble classifiers for decision making. This dependence on the training set is ‘typical’ for any supervised machine learning algorithm, thus, in our future research plans we want to build an ensemble model based on unsupervised techniques and compare the performance.

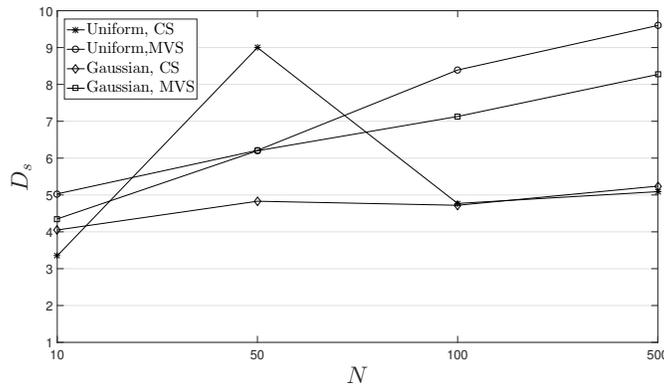


Figure 9: The speed of the selected node as delivered by the proposed model compared with the highest available speed.

In Figure 10, we plot the l^* realizations, i.e., the average load of the selected nodes for each of the 1,000 queries. The load is limited for a high number of nodes, which means that the burden added to ENs is minimized. The interesting is that the statistical error in our measurements is also limited (for the majority of the scenarios) exhibiting the stability of our model. In this set of our experiments, the *CS* exhibits the best performance compared to the *MVS*.

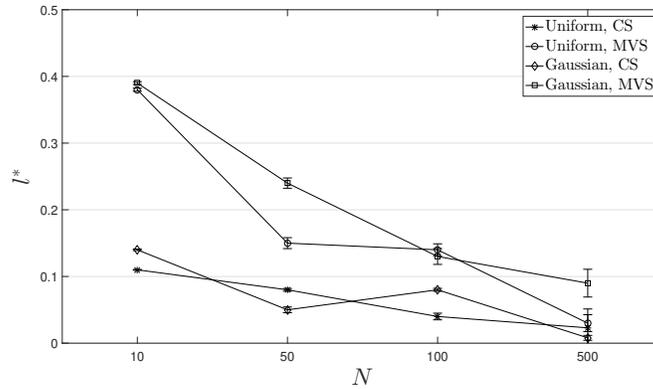


Figure 10: The load of the selected node as delivered by the proposed model.

The next set of experiments deals with the dataset provided by [4]. Our results are presented in Figures 11 & 12. The outcomes confirm our previous observations when adopted the synthetic traces. An increment in the number of ENs positively affects the performance of the model and leads to the selection of nodes with a low load. The average load of the finally selected node is similar for *CS* and *MVS*.

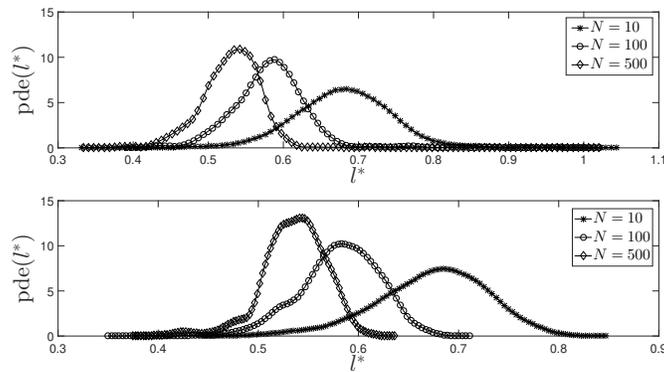


Figure 11: Pde for the load of the selected node based on Dataset 4 (up: CS, down: MVS).

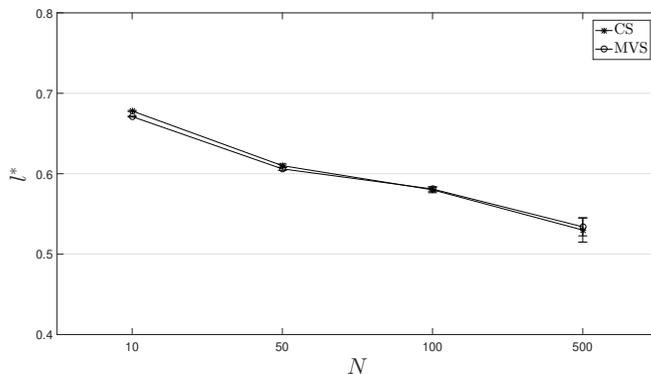


Figure 12: The load of the selected node as delivered by the proposed model based on the real trace.

We also report on our model’s sensitivity in parameters L and α . Recall that L depicts the number of dimensions into the available datasets and α is adopted to ‘smooth’ the result of the Quasi arithmetic mean when we calculate \mathbf{q}^s and r_i . In Figure 14, we consider Dataset 4 and $L \in \{10, 50\}$ to present our results for D_l . We observe that the difference with the optimal (minimum) load increases as N increases, however, D_l is below 0.015 for both, i.e., the *CS* and the *MVS*. Apart from that, we can conclude that L does not heavily affect the performance of the proposed approach. Our mechanism can be adopted for a high number of dimensions to result the best possible node to host the incoming queries. It is worth noticing that an increased L will lead to an increased processing time affecting R . In Figure 13, we present our results for $\alpha \in \{0.5, 1.0, 5.0\}$ and for Dataset 4. Again, the proposed scheme is not heavily affected by the α realization exhibiting efficiency in the detection of the best possible node to host the incoming queries. The majority of results are below 0.01 (the selected node is very close to the node exhibiting the lowest load).

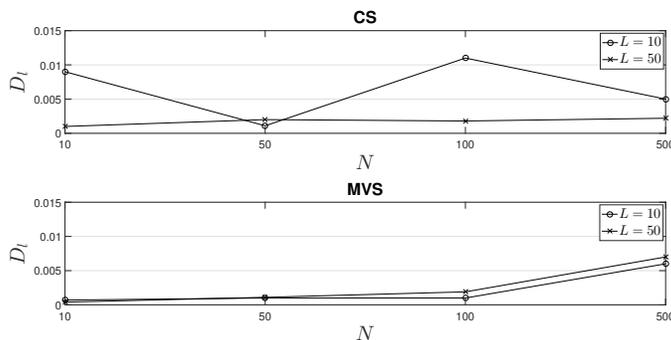


Figure 13: Our sensitivity results for various L realizations.

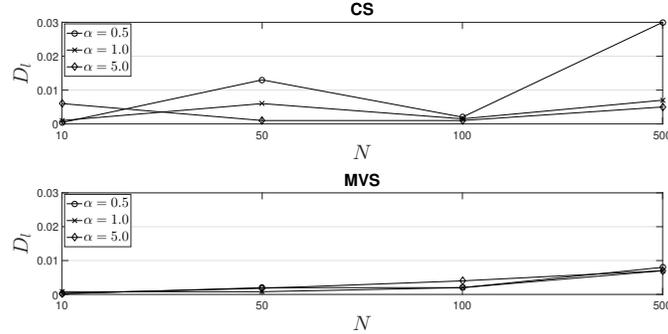


Figure 14: Our sensitivity results for various α realizations.

We also compare the proposed model with our previous efforts in the domain. The comparison refers in the throughput of QCs as well as the load of the selected node. Our model results an R varying from 0.42 to 0.022 for $N \in \{10, 50, 100, 500\}$ and average $l^* \in \{[0.008, 0.14], [0.03, 0.39]\}$, for the *CS* and the *MVS*, respectively. The model M1 discussed in [39] adopts an optimal stopping scheme that sequentially examines a set of processors before it delivers the final allocation. M1 results R values between 0.02 (for $N = 50$) and 0.27 (for $N = 10$). Hence, for a low number of nodes *CS* and *MVS* outperform M1 while in the remaining scenarios we observe a similar performance. The Model M2 presented in [41] adopts a learning mechanism accompanied by a load balancing module. The discussed scheme manages to allocate 58.48 queries for $N = 100$ to 454.55 queries for $N = 20$. The average load of the selected node (the learning scheme) is around 0.10 for $N \in \{2, 5, 20, 50, 100\}$. The average load for the clustering scheme is between 0.160 and 0.670. Our current models outperform the clustering scheme and exhibit a similar performance with the learning model. Finally, the model M3 discussed in [43] adopts a ‘simple’ learning scheme for the decision making. The performance of M3 related to the throughput is 50-100 queries per time unit when $N \rightarrow 500$. As we observe, our current approach outperforms, in the majority of the scenarios, our previous efforts leading to an efficient model. We have to take into consideration that the current scheme is more complex than the previous ones as it involves multiple classification decisions before the final result is in place. In addition, the current model takes every decision based on multiple parameters providing a holistic solution to the discussed problem. The throughput is at acceptable levels which means that the proposed scheme does not delay to provide the final allocation. In addition, as depicted by the comparison for the l parameter, the current model manages to result the nodes with a low load and, in many cases, better than our previous efforts.

7 Conclusions

The present status of Web, IoT and pervasive computing deals with the presence of numerous devices able to collect and process data. All these data are transferred to the Cloud to be the subject of further processing and knowledge production. Usually, such knowledge has the form of analytics that are the response in queries defined by users or applications. In this paper, we focus on an edge computing scenario where the collected data could be processed in edge nodes to reduce the latency in the provision of responses. Knowledge can be produced locally, close to the source of data and end users. A number of edge nodes can serve as aggregation points of data reported by the surrounding devices. As data are geo-distributed, a critical question is raised: in which edge node a query should be executed? Our paper tries to respond to this question and presents a model that efficiently allocates the incoming queries to the appropriate nodes. We provide a decision making mechanism, i.e., a meta-ensemble learning model responsible to select the best possible node to allocate a query. The decision is made on top of multiple parameters related to queries as well as to the available nodes. Our model takes into consideration the form of queries, their complexity, the deadline for having the final response together with the load and the speed of nodes. More importantly, we propose the use of a parameter that represents the distance between queries and datasets present in nodes. The aim is to exclude nodes that do not own data related to the conditions raised by queries as these nodes cannot provide responses that ‘match’ the requirements of the incoming queries. Our meta-ensemble learning scheme adopts multiple ensemble learning models to support a powerful decision making mechanism. We describe the performance of our framework through a large set of simulations adopting synthetic traces and traces found in the literature. Our evaluation results show that our model is capable of efficiently allocating the incoming queries towards the support of (near) real time responses. In the first places of our future research plans is the definition of an adaptive learning model that will be fully aligned with the environment of the edge nodes. This is important as it will save resources devoted to the training process through an ‘incremental’ training approach.

Acknowledgment

This research received funding from the European’s Union Horizon 2020 research and innovation programme under the grant agreement No. 745829.

References

- [1] Abouzeid, A., Bajda-Pawlikowski, K., Abadi, D. J., Rasin, A., Silberschatz, A., ‘HadoopDB: An Architectural Hybrid of MapReduce and DBMS Technologies for Analytical Workloads’, PVLDB, 2(1), 2009.

- [2] Ade, R., Deshmukh, P. R., 'An Incremental Ensemble of Classifiers as a Technique for Prediction of Student's Career Choice', 1st International Conference on Networks & Soft Computing (ICNSC), 2014.
- [3] Agarwal, S., Milner, H., Kleiner, A., Talwalkar, A., Jordan, M., Madden, S., Mozafari, B., Stoica, I., 'Knowing When You're Wrong: Building Fast and Reliable Approximate Query Processing Systems', ACM SIGMOD, USA, 2014.
- [4] Akay, M. Aci, C., Abut, F., 'Predicting the Performance Measures of a 2-Dimensional Message Passing Multiprocessor Architecture by Using Machine Learning Methods', Neural Network World, vol. 71(5), 2015, pp. 1907–1931.
- [5] Allwein, E. L., Schapire, R. E., Singer, Y., 'Reducing multiclass to binary: a unifying approach for margin classifiers', Journal of Machine Learning Research, 1:113–141, 2000.
- [6] Artail, H., El Amine, H., Sakkal, F., 'SQL query space and time complexity estimation for multidimensional queries', International Journal of Intelligent Information and Database Systems, vol. 2(4), 2008, pp. 460—480.
- [7] Awadalla, M. H. A., 'Task Mapping and Scheduling in Wireless Sensor Networks', Int. Journal of Computer Science, 440(4), 2013.
- [8] Bharti, S., Pattanaik, K., 'Task Requirement Aware Pre-processing and Scheduling for IoT Sensory Environments', Ad Hoc Networks, 50, 2016, pp. 102–114.
- [9] Blachnik, M., 'Ensembles of Instance Selection Methods based on Feature Subset', Procedia Computer Science, vol. 35, 2014, pp. 388–396.
- [10] Breiman, L., 'Bagging Predictors', Machine Learning, vol. 24(2), 1996, pp. 123–140.
- [11] Breiman, L., 'Arcing classifiers', The Annals of Statistics, 26(3):801–849, 1998.
- [12] Bullen, P., 'Quasi-Arithmetic Means', Handbook of Means and Their Inequalities', vol. 560, 2003.
- [13] Chandramouli, B., Goldstein, J., Quamar, A., 'Scalable Progressive Analytics on Big Data in the Cloud', in Proceedings of the VLDB Endowment, vol. 6(14), 2013.
- [14] Chaudhuri, S., Das, G., Srivastava, U., 'Effective use of block-level sampling in statistics estimation', In SIGMOD, 2004.
- [15] Chen, Z., Hu, W., Wang, J., Zhao, S., Amos, B., Wu, G., Ha, K., Elgazzar, K., Pillai, P., Klatzky, R., Siewiorek, D., Satyanarayanan, M., 'An Empirical Study of Latency in an Emerging Class of Edge Computing Applications for Wearable Cognitive Assistance', 2nd ACM/IEEE Symposium on Edge Computing, 2017

- [16] Coltin, B., Veloso, N., 'Mobile Robot Task Allocation in Hybrid Wireless Sensors Networks', in 2010 Int. Conf. on Intelligent Robots and Systems, 2010.
- [17] Condie, T., Conway, N., Alvaro, P., Hellerstein, J. M., Elmeleegy, K., Sears, R., 'MapReduce online', 7th Conference on Networked Systems Design and Implementation, 2010.
- [18] Daga, H., Nicholson, P., Gavrilovska, A., Lugones, D., 'Cartel: A System for Collaborative Transfer Learning at the Edge', ACM Symposium on Cloud Computing, 2019. pp. 25–37.
- [19] Dietterich, T., 'Ensemble Methods in Machine Learning', In Multiple Classifier Systems, MCS 2000, Lecture Notes in Computer Science, vol 1857, Springer, Berlin, Heidelberg.
- [20] Dittrich, J., Quiane-Ruiz, J. A., Jindal, A., Kargin, Y., Setty, V., Schad, J., 'Hadoop++: Making a Yellow Elephant Run Like a Cheetah', PVLDB, 3(1), 2010.
- [21] Doucet, A., Briers, M., Senecal, S., 'Efficient block sampling strategies for sequential Monte Carlo methods', Journal of Computational and Graphical Statistics, 2006.
- [22] Dzeroski, S., Zenko, B., 'Is Combining Classifiers with Stacking Better than Selecting the Best One?', Machine Learning, vol. 54, 2004, pp. 255–273.
- [23] Edalat, N., Xiao, W., Tham, C.-K., Keikha, E., Ong, L.-L., 'A price-based adaptive task allocation for Wireless Sensor Network', 6th ICMASS, 2009.
- [24] Farahbod, F., Eftekhari, N., 'Comparison of Different T-Norm Operators in Classification Problems', IJFELS, vol.2(3), 2012.
- [25] Freund, Y., Schapire, R. E., 'A decision-theoretic generalization of on-line learning and an application to boosting', Journal of Computer and System Sciences, vol. 55(1), 1997, pp. 119–139.
- [26] Galar, M., Fernandez, A., Barrenechea, E., Bustince, H., Herrera, F., 'Aggregation schemes for binarization techniques. Methods' description', Research Group on Soft Computing and Intelligent Information Systems, 2011.
- [27] Geman, S., Bienenstock, E., Doursat, R., 'Neural networks and the bias-variance dilemma', Neural Computation, 4(1):1–58, 1992.
- [28] Gualtieri, M., Yuhanna, N., 'The Forrester Wave: Big Data Hadoop Solutions', Technical Report, 2014.
- [29] Haßler, M. Jeschke, S. Meisen, T., 'Similarity Analysis of Time Interval Data Sets Regarding Time Shifts and Rescaling', In Proceedings International work-conference on Time Series, 2017.

- [30] Hameurlain, A., Morvan, F., 'Evolution of Query Optimization Methods', TL-SDK-CS, 2009, pp. 211—242.
- [31] Han, J., Kamber, M., Pei, J., 'Data Mining: Concepts and techniques', Morgan Kaufmann Publishers, 2012.
- [32] Hellerstein, J. M., Avnur, R., 'Informix under control: Online query Processing', Data Mining and Knowledge Discovery Journal, 2000.
- [33] Herodotou, H., Lim, H., Luo, G., Borisov, N., Dong, L., Cetin, F. B., Babu, S., 'Starfish: A Self-tuning System for Big Data Analytics', in CIDR, 2011.
- [34] Hossain, K., Raihan, Z., Hashem, M., 'On Appropriate Selection of Fuzzy Aggregation Operators in Medical Decision Support System', 8th Int. Conf. on Comp. and Inf. Technology, 2005.
- [35] Hu, X., Xu, B., 'Task Allocation Mechanism Based on Genetic Algorithm in Wireless Sensor Networks', in ICAIC, 2011.
- [36] Jermaine, C., Arumugam, S., Pol, A., Dobra, A., 'Scalable approximate query processing with the DBO engine', In SIGMOD, 2007.
- [37] Jiang, D., Ooi, D. C., Shi, L., Wu, S., 'The Performance of MapReduce: An In-depth Study', PVLDB, 3(1), 2010.
- [38] Kleinberg, E., 'A Mathematically Rigorous Foundation for Supervised Learning', In 1st International Workshop MCS 2000, 67—76, Springer-Verlag, 2000.
- [39] Kolomvatsos, K., 'An Intelligent Scheme for Assigning Queries', Springer Applied Intelligence, <https://doi.org/10.1007/s10489-017-1099-5>, 2018.
- [40] Kolomvatsos, K., Anagnostopoulos, C., 'An Edge-Centric Ensemble Scheme for Queries Assignment', 8th International Workshop on Combinations of Intelligent Methods and Applications, 2018.
- [41] Kolomvatsos, K., Anagnostopoulos, C., 'Reinforcement Machine Learning for Predictive Analytics in Smart Cities', Informatics, MDPI, 4, 16, 2017.
- [42] Kolomvatsos, K., Anagnostopoulos, C., Hadjiefthymiades, S., 'A Time Optimized Scheme for Top-k List Maintenance over Incomplete Data Streams', Elsevier Information Sciences (INS), vol. 311, pp. 59-73, 2015.
- [43] Kolomvatsos, K., Hadjiefthymiades, S., 'Learning the Engagement of Query Processors for Intelligent Analytics', Applied Intelligence, vol. 46, 2017, pp. 96–112.
- [44] Kong, E., Dietterich, T. 'Error - correcting output coding correct bias and variance', XII International Conference on Machine Learning, 313–321, 1995.

- [45] Kul, G., Luong, D. T. A., Xie, T., Chandola, V., Kennedy, O., Ypadhyaya, S., 'Similarity Measures for SQL Query Clustering', *IEEE Transactions on Knowledge and Data Engineering*, 2018.
- [46] Lam, L., Sue, C., 'Application of majority voting to pattern recognition: an analysis of its behavior and performance', *IEEE Transactions on Systems, Man and Cybernetics*, 27(5):553–568, 1997.
- [47] Logothetis, D., Yocum, K., 'Ad-hoc Data Processing in the Cloud', *VLDB Endowment*, 1(2), pp. 1472-1475, 2008.
- [48] Lübcke, A., Köppen, V., Saake, G., 'A Query Decomposition Approach for Relational DBMS using Different Storage Architectures', Technical report, *Elektronische Zeitschriftenreihe der Fakultät für Informatik der Otto-von-Guericke-Universität Magdeburg*, 2011.
- [49] Manning, C., Raghavan, P., Schütze, H., 'An Introduction to Information Retrieval', Cambridge University Press, 2009.
- [50] Mason, L., Bartlett, P., Baxter, J., 'Improved generalization through explicit optimization of margins', *Machine Learning*, 2000.
- [51] Moulton, R., Jiang, Y., 'Maximally Consistent Sampling and the Jaccard Index of Probability Distributions', *International Conference on Data Mining, Workshop on High Dimensional Data Mining*, 2018.
- [52] Moysiadis, V., Sarigannidis, P., Moscholios, I., 'Towards Distributed Data Management in Fog Computing', *Wireless Communications and Mobile Computing*, vol. 2018, ID 759686, 2018.
- [53] Pandit, S., Gupta, S., 'A Comparative Study on Distance Measuring Approaches for Clustering', *International Journal of Research in Computer Science*, 2(1), 2011, pp. 29–31.
- [54] Raman, V., Raman, B., Hellerstein, J. M., 'Online dynamic reordering for interactive data processing', *VLDB*, 1999.
- [55] Razavinegad, A., 'Task Allocation In Robot Mobile Wireless Sensor Networks', *Int. Journal of Scientific & Technology Research*, 3(6), 2014.
- [56] Re, M., Valentini, G., 'Ensemble Methods: A Review', *Advances in Machine Learning and Data Mining for Astronomy*, 2012, pp. 563–594.
- [57] Rifkin, R., Klautau, A., 'In Defense of One-Vs-All Classification', *Journal of Machine Learning Research*, vol. 5, 2004, pp. 101–141.
- [58] Schapire, R., Freund, Y., Bartlett, P., Lee, W., 'Boosting the margin: A new explanation for the effectiveness of voting methods', *The Annals of Statistics*, 26(5):1651–1686, 1998.

- [59] Singh, S., Singh, N., 'Big Data Analytics', International Conference on Communication, Information and Computing Technology, 2012.
- [60] Simon, M., Pataki, N., 'SQL Code Complexity Analysis', 8th International Conference of Applied Informatics, 2010.
- [61] Tian, Y., Ekici, E., Ozguner, F., 'Energy-Constrained Task Mapping and Scheduling in Wireless Sensor Networks', IEEE ICMASS, 2005.
- [62] Vashistha, A., Jain, S., 'Measuring Query Complexity in SQLShare Workload', International Conference on Management of Data, 2016.
- [63] Wen, Z., Le Quoc, D., Bhatotia, P., Chen, R., Lee, M., 'ApproxIoT: Approximate Analytics for Edge Computing', 38th International Conference for Edge Computing, 2018.
- [64] Wolpert, D. H., 'Stacked generalization', Neural Networks, vol. 5(2), 1992, pp. 241–259.
- [65] Yang, J., Zhang, H., Ling, Y., Pan, C., Sun, W., 'Task Allocation for Wireless Sensor Network Using Modified Binary Particle Swarm Optimization', IEEE Sensors Journal, 14(3), 2014, pp. 882–892.
- [66] Yongrui, Q., Sheng, Q. Z., Falkner, N. J. G., Dustdar, S., Wang, H., Vasilakos, A. V., 'When things matter: a survey on data-centric internet of things', Journal of Network and Computer Applications, vol. 64, pp. 137–153, 2016.
- [67] Yu, Y., Prasanna, V., 'Energy-Balanced Task Allocation for Collaborative Processing in Wireless Sensor Networks', Mobile Networks and Applications, 10(1-2), 2005, pp. 115–131.