CONTENTS

Automat. Comput. Appl. Math.
Volume 30 (2024), Number 1

COMPUTER SCIENCE

Harnessing Appliance Data for Trend Analysis and Dish Quality Prediction in IoT-enabled Kitchens

Andreea Pastiu
Technical University of Cluj Napoca

Ramona Tolas
Technical University of Cluj Napoca

Raluca Portase
Technical University of Cluj Napoca

Mihaela Dinsoreanu
Technical University of Cluj Napoca

Camelia Lemnaru
Technical University of Cluj Napoca

Rodica Potolea
Technical University of Cluj Napoca

***Abstract*—In smart homes and the Internet of Things (IoT), integrating connected appliances has become increasingly prevalent. This paper delves into the innovative domain of trend analysis and dish quality prediction by harnessing appliance data. By leveraging data generated from experiments with smart cooking devices, we explore patterns, correlations, and trends to gain insights into user behaviors and appliance performance. Understanding data collected from cooking experiments can help identify the evolution of the appliances over time, as well as the impact of different appliance features on the quality of the cooked dish. This paper introduces general pre-processing techniques employed to clean datasets, along with diverse methods for identifying trends and correlations. Additionally, we present two data classification methods, highlighting potential methods for their enhancement. The final step involves the practical application of these methodologies to real-world experimental data. From the results obtained, we can observe general trends, what features are the most important in the output of a dish, and how we can predict the outcome of a dish, given specific cooking settings.**

I. Introduction

Cooking appliances have always played an essential role in people's lifestyles. The recent technological evolution did not bypass these devices, so nowadays, cooking home appliances have evolved into smart devices. Most of them are equipped with sensors that measure different characteristics of the environment and the cooking appliance's internal state. Studying user needs and preferences, manufacturers of these types of appliances introduced many functionalities, making the appliances complex from the usage perspective. Users might need help to deduct the combination of settings that will lead to a successful dish. To tackle this, cooking appliance manufacturers have dedicated laboratories where experiments are performed to gather data about the impact of different settings of cooking appliances on the general success of a dish. Besides the analysis of the dish quality, energy consumption optimization strategies can be inferred from the data generated by these kinds of experiments.

In this work, we present generally applicable methods for knowledge inference in the context of experimental data. Pre-processing steps required in such a context are presented together with specific data preparation phases. The processing step consists of statistical and machine-learning steps. The focus is to define general trend identification methodology and to show how it can be applied to infer valuable knowledge, such as the evolution of cooking time and cooking energy in the monitored period. We use machine learning to predict the success or failure of a recipe, given certain conditions determined by appliance settings. The findings can be used as a data-driven decision by the cooking appliance producer to give setting recommendations to customers for specific recipes, improving the overall user experience.

This paper is structured as follows: Section 2 briefly describes the theoretical aspects needed for developing the presented methodologies and presents related findings in the topics of interest. Section 3 presents our proposed methods, and Section 4 contains the results of applying these methods to the experimental dataset. We conclude this work in Section 5, where we also present future development ideas.

II. Theoretical Background and Related Work

The purpose of this chapter is to explain the main theoretical and technical concepts that were used in developing the analysis and knowledge inference methods. Trends identification and unveiling relations between features are tackled first, followed by prediction-related theoretical aspects.

A. Trends identification

The trend is a pattern in data that shows the movement of a series to relatively higher or lower values over a long period. This can be used to predict future values and extract insights about the data. For example, by analyzing the energy consumed by an appliance in time, given that the appliance suffered changes, a decreasing trend indicates that the changes applied in the appliance building process have a positive outcome.

Trend analysis is one of the classical methods of data analysis. It is present in literature and studied in practical contexts such as water quality analysis [1] [2], where we can see that trend analysis can help identify stations' water quality characteristics are changing significantly or rapidly and require further investigation and distinguish between trends that are caused by natural factors (such as discharge, seasonality or weather) and those that are caused by human factors (such as land use, pollution or management), or from a theoretical perspective [3] [4], where various methods of trend analysis are compared.

Reviewing the existing literature, we identified several steps needed for trend identification. The first is dataset understanding, followed by pre-processing steps such as data cleaning (outliers removal, tackling missing values, ensuring data consistency). The next step is to define the time period on which trend analysis is performed.

Visualizing the data is also vital in this case because it can provide initial insights. Scatter plots (2D or 3D) are a suitable mechanism for constructing meaningful visualizations in this context. The visualization method is the research subject in [5]. The authors conclude that scatter plots are preferred

Donut – Autonomous robot with sentiment analysis for serving in restaurants

Robot that performs food orders at tables

Alexandra Șerban
Technical University of Cluj-Napoca
Cluj-Napoca, România
alexandra.serban.2510@gmail.com

Ana Rednic
Technical University of Cluj-Napoca
Cluj-Napoca, România
Ana.Rednic@cs.utcluj.ro

Radu Gabriel Dănescu
Technical University of Cluj-Napoca
Cluj-Napoca, România
Radu.Danescu@cs.utcluj.ro

***Abstract*— The pandemic context determined our society to enjoy many activities involving technology during lockdown. Recently, technology underwent an increase in many domains such as education, sales or restaurant. One of the biggest challenges for restaurant managers was to find different ways to keep alive their business with as few financial losses as possible. In order to do that, a possible solution to attenuate the lack of staff and increase profit would be a robotic waiter that delivers food at tables. This article offers a low-cost solution to the previous topic and detailed design and implementation of the system. The proposed waiter robot has not only ability to perform delivery services in real-time through an online menu but it also interacts with customers, allowing them to express their enthusiasm or, if necessary, their dissatisfaction through the dedicated feedback section.**

Keywords*— robot, waiter, system, mechanism, autonomous, restaurant, movement

I. INTRODUCTION

The Hotel, Restaurant, and Café/Catering sector (HoReCa) has undergone major changes in recent decades due to the way the Internet of Things (IoT) has contributed to the automation and improvement of service quality. However, the biggest challenge in this industry still remains the management of operational costs which include: supply of food and materials (utensils, equipment), food preparation, wasted food utilities and employee salaries. In 2019 restaurants experienced difficulties in passing on inflationary costs to consumers (at that time the average price increase in a restaurant was 3.1%, the cost of labor increased by 3.8% and that of food with 5%). In 2021 the cost of restaurants increased by up to 4.5% compared to production costs that suffered an increase of up to 8.9%. [1]

In this context, certain estimates have been made that show the introduction of robots on premises can save operational costs in percentages between 30% and 70%. And even if there is only one robot serving food for every 1500 restaurants (reported globally in 2022), many of them are emerging from the prototype stage and can do almost everything a human can do. [2]

Figure I.1 shows the extent to which the jobs of various positions in a restaurant could be replaced by robots becoming fully autonomous from an operational point of view, with minimal or no human intervention [2]:

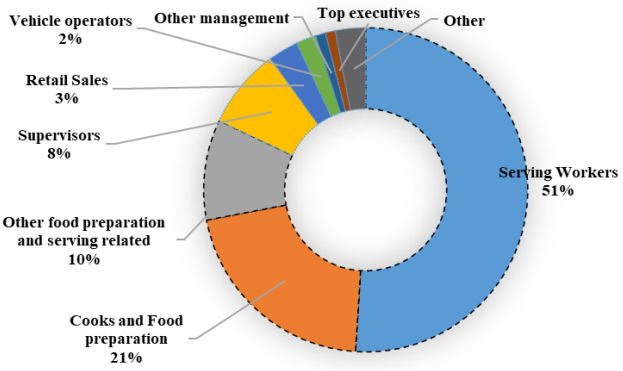

Figure I.1 The ability to replace restaurant jobs with robots (82%)

The goal of this project is to automate the process of serving food in restaurants by introducing a robot to serve the customers. In order to require as little staff as possible on the premises, the product must, first of all, be autonomous and successfully carry out the orders via the mobile phone by as many age groups as possible.

II. RELATED WORK

A. Similar algorithms

Integration of several advanced algorithms are made in the implementation of autonomous robotic waiters to ensure efficient and effective service. One of the critical components in this system is the movement algorithm because restaurants have a dynamic environment in which the robot must navigate seamlessly. Some of the algorithms studied for movement are: **BLE** (Bluetooth Low Energy is a wireless communication technology designed for short-range and has become essential for developing smart, connected devices) [3], **Bluetooth 5.0 Beacon** (the main purpose of this type of BLE is to efficiently broadcast signals that help identify physical objects by nearby portable electronic devices. It is used in many areas such as the parking system of a car.) [4] and **SLAM** (Simultaneous Localizations And Mapping is a method used for autonomous vehicles where a map can be built to locate the system – which is in charge of creating the map, at the same time. It allows devices to map unfamiliar environments specifically for

A High Trust Document Storage System for Public Institutions using Distributed Ledger Technology and IPFS

Sergiu-Emanuel Blaj, Marcel Antal, Cristinel Mihai Mocan and Teodora-Melania Furcovici
Department of Computer Science
Technical University of Cluj-Napoca
Cluj-Napoca, Romania
Email: blaj.co.sergiu@student.utcluj.ro, {marcel.antal, cristi.mocan}@cs.utcluj.ro, teodora.io.furcovici@student.utcluj.ro

***Abstract*—Public institutions possess many documents and signed contracts with various companies. Those documents contain sensitive information but citizens are having a difficult time trying to obtain and analyze them, despite the fact that a public institution's activity is possible because of taxes paid by citizens. The documents being left unsupervised may be one of the causes the fraud is reaching alarming levels. Developing a decentralized platform, intuitive to use, would reduce the fraud and corruption in public domain. This paper presents a Blockchain-based solution that allows public institutions to publish their documents in a decentralized storage (IPFS), allows citizens to easily access the documents and report anything suspicious and allows judges to examine the complaints in order to determine culpability. We performed a set of experiments to evaluate the scalability of the proposed solution on a setup similar to the public Ethereum blockchain, showing that our system using IPFS for document storage can retrieve 100 documents/second and roughly 360,000 documents per hour, achieving a 50% improvement compared to a standard blockchain architecture where all the storage is on-chain.**

Keywords - Blockchain, Ethereum, IPFS, public institution, Govern, document, citizen, complaint, judge, fraud, corruption

I. Introduction

The digital revolution changed the way the public domain of a country acts, starting from the way it is organised, to the way it engages with citizens. However, the digitalisation did not happen "over nigh", but expanded through a couple of decades. Article [1] identifies five essential steps of the bureaucracy digitization that started since the development of internet: Basic Web Presence (1996 - 1999), Interactive Web (1997 - 2000), Transaction Web (1998 - 2003), Integrative and Transformation Web (2000 - 2004) and Smart City Governance Web (2005 - present). These eras that started with small steps, such as simply displaying vital information using static websites on internet, are now using technologies such as *Big Data*, *Internet of Things* and *Artificial Intelligence* to improve the process of governing.

The authors of [2] pointed the difference between two terms involved in digitalisation of public domain, *e-Government* and *e-Governance*. While the first represents the way the Govern exposes its activity to citizens or integrates different useful services, such as online payments, voting and tendering, the latter aims to involve citizens in the governing process, with the use of internet, to increase their trust in public domain.

In Europe's developed countries, *e-Governance* systems helped with improving citizens life by removing the long waiting-times queues and the bureaucracy. In countries such as Finland, Denmark or The Netherlands, where the percentage of citizens using the internet to help with the governing process is close to the percentage of people having access to internet, not only the fraud has recorded lowest limits since digitalisation (figure 1), but the overall quality of life and trust of citizens in public institutions have improved.

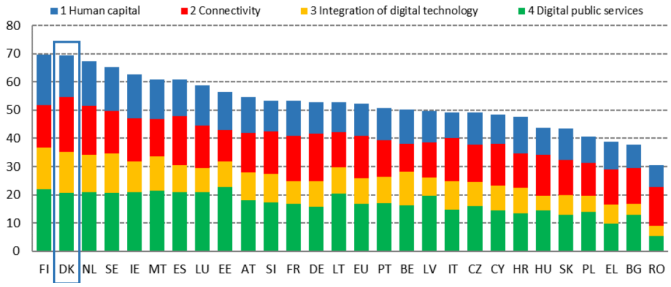

Fig. 1. Digital Economy and Society Index [2]

In Romania, the two government platforms that allows its citizens to analyze the official documents are *Sistemul Electronic de Achiziții Publice*[1] and *Transparența bugetară*[2]. While they are developed under the aegis of the Romanian Government information system, this does not guarantee that the information presented in the documents are 100% accurate and unchanged. Moreover, they do not offer a special *Complaints* section to allow citizens report anything suspicious. They provide, however, a *Contact* section, where citizens can submit their feedback or any other thoughts, but in case of a message written as a complaint nothing ensures that a third party authority, like a judge, would consider analyzing the complaint and, furthermore, taking actions to solve it.

[1]https://www.e-licitatie.ro/pub
[2]https://mfinante.gov.ro/transparenta-bugetara

Governance Architecture for Citizen Developers in the Digitalization Era

Authors: Călin Cenan[(1)], Bogdan Nicuşor Bindea[(2)], Gabriel Dragomir[(3)]
[(1)]calin.cenan@cs.utcluj.ro, [(2)]bogdan.bindea@cs.utcluj.ro, [(3)]gabriel.dragomir@cs.utcluj.ro

***Abstract*— The presented paper aims to propose an architecture for digitalization that is intended for individuals who take on the role of developers—commonly referred to as citizen developers—and create software using low-code tools. Such working models are crucial in the context of continuous and rapidly growing digital transformations observed across various companies, organizations, public administrations, and educational institutions, enabling them to successfully face the challenges within a regional innovation system.**

Keywords — Enterprise Resource Planning; Low-Code tools; Automatic Code Generation; Agile Approaches; Drag-and-drop interfaces; Cloud-based development tools.

I. INTRODUCTION

Traditionally, organizations rely on their Research and Development departments and specialists to help other teams solve business challenges using state-of-the-art technologies. However, a new role is emerging people who decide to tackle their own tech needs without turning to the IT experts. These citizen developers, often without formal technical background, use the available tools on the market, particularly visual integrated development environments, called IDEs, to create their own software solutions through drag-and-drop interfaces. Frequently leveraging cloud-based development tools, which allow them to bypass traditional Software Departments. A study highlighted in [01] presents such stakeholders in the Tyrol-Veneto macro-region, a notable example of an innovative European area.

Low-code and no-code solutions enable users to create software applications using graphical interfaces, often replacing the traditional coding process with drag-and-drop functionality. This approach makes it possible to build entire systems with minimal or no manual coding [02], [03]. By empowering employees to design and develop applications tailored to their specific needs, this method allows domain experts to create solutions that align with their requirements while also reducing the workload on IT departments. As a result, many companies have adopted this approach [04].

Citizen developers can introduce several challenges to a corporation. One challenge is the responsibility for maintaining the software solutions they create, particularly when the software is critical to operations or involves legal risks. Another issue is ensuring effective communication between different teams in a complex software environment, especially when citizen developers introduce additional software that needs to be managed. Typically, administrators, system architects, or product managers are responsible for maintaining and managing the organization's information systems.

Software development typically involves five key activities: defining requirements, designing, developing, testing, and maintaining software, each of them representing a phase in the software development life cycle [05]. Traditionally, building new features requires a deep understanding of the business domain and collaboration between the software development department and software users. However, the rise of low-code/no-code tools has empowered non-programmers, known as citizen developers, to create software with minimal or no coding experience, addressing business needs without relying on IT support.

While this approach can reduce the workload for specific departments and speed up development due to the citizen developer's domain expertise, it also introduces risks. These risks arise when development deviates from the company's standardized processes, including design, documentation, and testing, and when IT tools are used without official approval, a practice known as Shadow IT [06].

II. RELATED WORK

Low-code and no-code tools are recognized for empowering experts without a technical background to create software solutions, often without involving traditional IT departments. Integrating citizen developers into various contexts has been the focus of several studies, including case studies that examine key considerations for implementing this approach. However, there remains a gap in research, particularly concerning the management of citizen developers at the organizational level.

Previous research has primarily focused on the use of low-code technologies and the involvement of citizen developers. For example, the impact analysis of no-code development platforms discussed in [07] highlights how such analysis can provide valuable feedback to users. The study in [08] addresses the adoption of citizen developers within organizations, extending the technology-organization-environment framework by considering various technical, organizational, and inter-organizational factors. This research suggests that centralized IT governance may be beneficial for integrating citizen developers, and it uses cross-case analysis to explore the challenges associated with citizen development.

In [09], a framework is presented for applying low-code tools in the development of Enterprise Resource Planning (ERP) systems, which integrate and automate business functions into a single platform. This study identifies where low-code tools are most effective and discusses the role of

Employing Specific Techniques for Client Data Mining in the Domain of Car Insurance

Delia Mitrea, Paulina Mitrea
Department of Computer Science
Technical University of Cluj-Napoca
Cluj, Romania
delia.mitrea@cs.utcluj.ro

Erik Barna
Life is Hard S. R. L, Cluj-Napoca.
Cluj, Romania
erik.barna@lifeishard.ro

***Abstract*— Segmenting clients considering common characteristics and determining the needs for each group is an important objective in the car insurance domain. We aim to achieve client segmentation through specific computerized techniques, employing and comparing multiple clustering (grouping) methods, considering both classical and deep learning algorithms. Regarding the classical techniques, clustering methods, adequate for both numerical and categorical data were adopted, such as k-means clustering, respectively k-prototype clustering. Regarding the deep-learning methods, a stacked denoising autoencoder, followed by k-means clustering was experimented, the corresponding performance being compared with that resulted after the individual application of the classical k-means clustering. After employing the clustering methods, the relevant attributes that separate among clusters were identified, the dependencies between the policy insurance type and other attributes being also analyzed through both graphical representations and association rules. The experiments were performed considering the data extracted from a relational database specific for the car insurance domain, containing about 1000 instances for the main tables.**

Keywords*— *client segmentation, clustering (grouping) methods, relevant features, deep learning techniques, association rules, car insurance

I. INTRODUCTION

Data mining specific methods, such as the clustering (grouping) techniques, or the association rules, have the potential to improve many areas of life, especially those regarding various insurance types. The main goals are client segmentation considering criteria of interest, for applying the most appropriate policies, next to the exploration of hidden relationships among ttributes. The most important challenges in this domain target the adaptation of the specific techniques to categorical features, respectively to "big data", both conventional and deep learning methods being employed. Thus, techniques such as k-means clustering or fuzzy k-means clustering [1-2] also Minimum Spanning Tree (MSP) [3] and hierarchical clustering [4] have been successfully applied for performing client segmentation in the domains of financial or health insurance, while the association rules based methods, namely the *Apriori* and *FP Growth* algorithms [5-6] have been adopted for attribute dependency mining. However, no systematic comparison was performed between the classical and the deep learning techniques in the field of car insurance, while the relevant attributes that best separate among clusters were less analyzed. In the current approach, we experimented and analyzed several clustering techniques for client segmentation in the field of car insurance, considering groups of attributes that emphasize the financial characteristics or the life stage of the clients, corresponding to the following scenarios: (1.) client segmentation based on financial data, which will allow the insurer to sell insurance policies suitable for each type of client; (2.) grouping customers by life stage, which could contribute to recommending an appropriate insurance policy. We employed various techniques, specific to both numerical data (k-means clustering), as well as to categorical data (k-prototype clustering). We also experimented deep learning methods for this purpose, in the form of Stacked Denoising Autoencoder (SAE) [7] followed by k-means clustering [8], then we compared the performance resulted when applying classical k-means clustering, with the situation when k-means clustering was employed through deep learning techniques. The relevant features that separate between clusters were derived as well using specific techniques, such as Correlation based Feature Selection (CFS)[9]. Further data analysis has been achieved by plotting the attributes against each other, for studying the dependency of some attributes of interest against the policy type.

II. STATE OF THE ART

The clustering techniques have been widely adopted for client segmentation considering various types of insurances. Thus, in [10], the authors compared multiple improved k-means clustering approaches for performing client segmentation in the health insurance field. In this manner, a maximum Root Mean Squared Error (RMSE) of 6.20%, respectively a maximum accuracy of 93.79% was obtained. The superiority of k-means clustering on small datasets was claimed, as well. In [11] the authors evaluate the application of clustering methods such as k-means, Minimum Spanning Tree (MST) and hierarchical clustering in the financial domain, achieving satisfying results. The authors focus upon the implementation of the association rules, respectively on the neuronal segmentation technique that involved Self Organizing Maps (SOM), in the health insurance field. Concerning the experimental dataset, the attributes of interest were the medical services, a *transaction id* associated to each record being considered. The association rules were implemented with the aid of the *Apriori* algorithm, these being applied upon the pathological service database, which contained the patient medical visits. The association rules were derived considering a minimum confidence value of 50%, respectively three different minimum values for the

Leveraging efficient semantic and spatial neighbourhood attention for anchor free 3D object detection

Selma Deac[1] and Sergiu Nedevschi[2]

***Abstract*—In this paper we identify the lack of awareness regarding the semantic and spatial neighbourhood context around detected objects' centers in the late stages of anchor free 3d object detection networks. Given the inherent high sparsity of LiDAR data, these attention-less methods struggle in long-range detection (e.g. more than 50m) which we consider essential for achieving safe automation. Transformer-based detection heads offer neighbourhood attention but under extensive training and high resources demand. To tackle this issues, we introduce Center Context and Class-aware Anchor Free head (aka. C3AF), a detection head that attends to the most relevant spatial and semantic neighbourhood cues of the detected objects, using simple, yet effective, attention mechanisms. Additionally, to address LiDAR sensors' sparsity issue, we use an efficient fully sparse backbone network to be combined with our C3AF and call our standalone network C3AFnet.**

We perform extensive evaluation on Waymo Open Dataset that show C3AFnet achieves higher accuracy metrics with a better trade-off between runtime and memory-footprint than previous state-of-the art methods. This is especially notable when detecting objects that are occluded or at far distances (beyond 50 meters). Additionally, to demonstrate the effectiveness of the C3AF head as an independent detection network, we integrate it with a popular backbone network, Pointpillars, and showcase its superiority compared to other center-based heads.

***Keywords*—3D object detection, Autonomous vehicles, Lidar**

I. Introduction

LiDAR-based 3D object detection methods for autonomous driving, have made remarkable progress in the last years thanks to advances in deep neural network field.

Early works in 3d object detection in the deep learning era are inspired by 2d object detection state of the art methods, which use 2D convolutional backbones to extract high level feature maps and then fit prior anchor boxes to the 2D object proposals in the feature image. These proposals are then classified and refined to generate the final detection results. Nowadays, the majority of 3D object detection methods for autonomous tasks are composed of the following network layers:

1) the input data encoding layer: where the 3d LiDAR point cloud is projected onto a 2d/3d grid. The points from each non-empty cell are then passed into feature space using a feature encoding layer commonly referred as Voxel or Pillar Feature Encoding layer(VFE [31] or PFE [11]). The resulted features are rearranged into a 2d or 3d grid structure to enable the usage of convolutional blocks for high level feature extraction.
2) the backbone and neck layer : The backbone network uses 2d/3d convolutions to learn features across multiple levels most commonly on the Bird's Eye View (BEV) space [11], [18], [27], [31], [28], [7], [29], [32], [9]. It is followed by the neck module which aggregates the backbone's multi-scale feature maps before feeding them to the last layer, the detection head.
3) the detection head layer: predicts final bounding boxes with their associated classification score and regression offsets. The detection is performed on the BEV feature maps due to the fact that same class objects tend to keep their size and the only parameter which changes is the objects' rotation.

All layers of the network have an influence on the quality of the final predictions. If the network does not learn high quality features, the detection head will not have sufficient information to make good predictions. Also, what a network must learn and how the network learns is encoded inside the detection head. So if the detection requirements are not clear and relevant inside the detection head, the network will not converge properly while training, thus it will not learn relevant information.

There are two types of detection heads: anchor-based detection head [11], [10], [27], [31], and anchor free detection head [8], [18], [24], [22], [13], [32].

Anchor based methods use predefined bounding boxes, known as anchors, to classify and localize objects. The idea behind anchor-based 3d object detection is to generate a set of anchor boxes of prior fixed sizes at various rotations in each pixel of the scenes BEV feature image. The network learns the classification scores and the bounding box regression offsets for each anchor box.

Anchor free methods, on the other hand, directly predicts objects bounding box regressions and classification scores without any reference to anchors. The anchor free approach eliminates the need for anchor-related hyperparameters. Because of this, an anchor free detection model is more flexible. This flexibility is important for complex outdoor scenarios where same class objects can have various shapes and sizes.

In this paper, we propose C3AFnet, a robust and efficient 3D

[1]Selma Deac is a PhD candidate at the Computer Science Department, Faculty of Automation and Computer Science, Technical University of Cluj-Napoca, Romania selma.goga@cs.utcluj.ro

[2]Sergiu Nedevschi is a professor at the Computer Science Department, Faculty of Automation and Computer Science, Technical University of Cluj-Napoca, Romania sergiu.nedevschi@cs.utcluj.ro

Time Efficiency in the Encrypted Domain: TFHE Benchmarking

Diana-Elena Petrean*, Rodica Potolea*
*Computer Science Department, Technical University of Cluj-Napoca, Romania
Email: {diana.petrean, rodica.potolea}@cs.utcluj.ro

***Abstract*—In the digital age, the growing importance of data privacy and individuals' personal information protection have led to the development of various privacy-preserving technologies. Some of this technologies are based on advanced cryptographic techniques such as Homomorphic Encryption (HE), that allows computations to be performed directly on encrypted data. While the major benefit of HE is the capability of performing privacy-preserving data processing, the integration of this technique in real-world applications comes at the cost of performance degradation. The time efficiency of traditional operations performed on plaintext data is significantly impacted when the corresponding operations are performed on encrypted data using HE. In this work, we measure the running times of various algorithms implemented using the TFHE scheme (Fully Homomorphic Encryption over the Torus [1]) and we compute the ratio between the running times obtained in the encrypted domain and the running times obtained by evaluating the corresponding algorithms on plaintext data. Our benchmarking demonstrates that, even if this ratio is of the order of thousands or even millions, the considered algorithms are evaluated using HE in reasonable times of the order of seconds. Although the HE operations are considerably slower than the traditional plaintext operations, the use of HE has become reliable in practice for privacy-preserving data processing.**

***Keywords*—Homomorphic encryption, Privacy-preserving data processing, Time efficiency, Benchmarking**

I. Introduction and related work

Nowadays, data privacy is a strong requirement for the vast majority of software applications and information systems. Various regulations have been imposed worldwide in order to enhance the privacy rights that individuals have over their personal information, such as the General Data Protection Regulation (GDPR) in the European Union and the California Consumer Privacy Act (CCPA) in the United States. In this context, various privacy-preserving technologies based on advanced cryptographic techniques have been developed.

Encryption represents a powerful tool for the protection of persistently stored data. However, in many scenarios, data processing should be allowed while also ensuring data privacy. A possible solution is Homomorphic Encryption (HE), as it allows computations to be performed directly on encrypted data. Some scenarios where HE can be used for privacy-preserving data processing are presented in [2] (fog computing for Internet of Things, cloud computing, privacy-preserving machine learning), [3] (medical research and healthcare applications), [4] (biometric security), [5] (electronic voting), [6](the financial services sector).

In the context of HE, addition and multiplication are considered base operations and, using this two operations, complex circuits can be built and various functions can be evaluated. Considering the base operations supported by a scheme and the number of operations that can be chained, the HE schemes are classified in the literature as follows:

- Partially homomorphic encryption (PHE): supports circuits consisting of only one type of base operation (either addition or multiplication);
- Leveled homomorphic encryption (LHE): supports circuits consisting of both base operations, but the maximum number of base operations that can be chained is bounded. Generally, the ciphertexts contain a small amount of noise (or a small error) that increases when performing a base operation. After a critical noise threshold is reached, decryption is no longer possible;
- Fully homomorphic encryption (FHE): supports circuits consisting of both base operations and the maximum number of base operations that can be chained is unbounded. Bootstrapping is used in order to reduce the noise in ciphertexts, but in general this is a time-consuming operation.

The concept of HE dates back to 1978, but it has become increasingly popular during the past decade. Not only multiple HE schemes have been proposed by academic researchers, but various HE-based methods have also been integrated in business solutions. The surveys from [7] and [8] present some of the well-known HE schemes and libraries implementing these schemes. According to underlying techniques on which the HE schemes are based, the history of HE is divided in the pre-FHE period and four FHE generations. Some representative schemes for each generation are:

- pre-FHE: the Paillier cryptosystem [9];
- 1^{st}-generation FHE: Gentry's lattice-based scheme [10];
- 2^{nd}-generation FHE: the BGV (Brakerski-Gentry-Vaikuntantan) scheme [11], the BFV (Brakerski-Fan-Vercauteren) scheme [12] [13];
- 3^{rd}-generation FHE: the GSW (Gentry-Sahai-Waters) scheme [14], the FHEW scheme [15], the TFHE (Torus Fully Homomorphic Encryption) scheme [1];
- 4^{th}-generation FHE: the CKKS (Cheon-Kim-Kim-Song) scheme [16].

The great benefit of HE is the capability of privacy-preserving data processing. However, HE also poses several challenges, including the performance impact and the high complexity of the algorithms. The operations in the encrypted domain are significantly slower compared to the corresponding plaintext operations, which results in a performance degradation for the HE-based software solutions compared to the traditional solutions that work with unencrypted data. Various enhancements and acceleration methods [17] were proposed in order to increase the efficiency of the HE schemes and to achieve reasonable computational times in the encrypted domain. Due to the progress made in the last years in cryptography, the implementation of