

Subject 1

Title: Food recognition and recipe identification

Coordinators/Contacts:

Prof. dr. ing. Adina Magda Florea (adina.florea@cs.pub.ro),

Dr. ing. Mihai Trăscău (mihai.trascau@cs.pub.ro)

Description:

The research topic aims at recognizing basic food elements from pictures taken by the user with a mobile phone and then finding the recipes that can be prepared with these items. A link with a smart refrigerator having the ability to capture in images the foods in the refrigerator is to be integrated in the topic. One of the main challenges is to determine also the amount of food items available for cooking in such a way as to scale the culinary recipe to the available quantity. Moreover, the system has to find the recipes with the least number of missing ingredients that match the current existing items, find on-lien stores that provide these missing items and propose them to the user.

<https://foodai.org/#index>

Subject 2

Title: AI based assistant for smart fridges

Coordinators/Contacts:

dr. ing. Mihai Trăscău (mihai.trascau@cs.pub.ro),

dr. ing. Alexandru Sorici (alexandru.sorici@cs.pub.ro),

Description:

Topic designed in partnership with Infosys Consulting Romania

Along with the expansion of IoT solutions many *smart things* have made their way into people's houses. One such item is the *smart fridge* a connected appliance which is equipped with various sensors: video, temperature, status, door open-close, etc. Even though information becomes easily available given the connectivity and the variety of sensors, the need to interpret the data requires intelligent algorithms.

In this project we will explore two tasks related to a smart fridge:

- Detect products inside and keep track of their quantity and usage by employing computer vision techniques
- Analyze activity sequences regarding the contents of the fridge and its utilization using machine learning methods

The partner company (Infosys) is providing an instance of a smart fridge, equipped with the type of sensors mentioned previously (in particular, an RGB camera). Prospective students have the opportunity/requirement to integrate their work on a real device.

Subject 3

Title: User detection and tracking

Coordinators/Contacts:

Prof. dr. ing. Adina Magda Florea (adina.florea@cs.pub.ro),

dr. ing. Mihai Trăscău (mihai.trascau@cs.pub.ro),

drd. ing. Ștefania Ghiță (stefania.a.ghita@cs.pub.ro)

Description:

People tracking, unlike other recognition and interpretation tasks, is difficult both from the point of view of recognition and prediction of the trajectory, and from the one of the identification of the ground truth. Partially visible, occluded, or cropped targets, reflections in mirrors or windows, and objects that very closely resemble targets, all impose intrinsic ambiguities, such that even humans may not agree on one particular ideal solution. Moreover, establishing evaluation metrics with free parameters and ambiguous definitions often lead to conflicting quantitative results [Lea et.al. 2017].

People detection from videos or images is subsumed by the problem of object detection. During the last years, the problem has received a lot of attention and many robust solutions exist, both based on traditional computer vision processing and based on deep neural networks. Even in this case, the problem is not yet entirely solved if we are to consider real-life situations, fast processing time and limited resources.

On the other hand, people tracking from videos is a more difficult problem and is currently a challenge. Some years ago, the trend on people detection from video sequences was to find strong, preferably optimal methods to solve the data association problem. Linking detections in a set of consistent trajectories (matching two detections based on either simple distances or weak appearance models) was solved by various methods such as DP NMS, by Conditional Random Fields or as a variational Bayesian model; performances were not very good.

Recently, the focus of people tracking from videos is on building robust pairwise similarity costs, mostly based on strong appearance cues, leading to better tracker performances and more complex scenarios. Some good approaches use sparse appearance models or integral channel feature appearance models or aggregated local flow of long-term interest point trajectories to improve detection affinity. Still, most of the available tracking approaches do not include a learning algorithm to determine the set of model parameters for a dataset. Some recent approaches tried to use deep learning, such as Recurrent Neural Networks to encode appearance, motion, and interactions or deep matching to improve the affinity measure. These approaches are rather few and results are promising, however they do not surpass other approaches for the time being.

The research topic implies to use both traditional tracking methods (e.g. Kalman filters) but also deep learning models to achieve better performances than recent approaches, and evaluating the results on the MOTT Challenge [Lea et.al. 2017].

<https://towardsdatascience.com/people-tracking-using-deep-learning-5c90d43774be>

[Lea et.al. 2017] Laura Leal-Taixe, Anton Milan, Konrad Schindler, Daniel Cremers, Ian Reid, Stefan Roth, (2017) Tracking the Trackers: An Analysis of the State of the Art in Multiple Object Tracking, <https://arxiv.org/abs/1704.02781>

Subject 4

Title: Pedestrian trajectory analysis in outdoor environments

Coordinators/Contacts:

**Prof. dr. ing. Adina Magda Florea (adina.florea@cs.pub.ro),
dr. ing. Mihai Trăscău (mihai.trascau@cs.pub.ro),
drd. ing. Ștefania Ghiță (stefania.a.ghita@cs.pub.ro)**

Description:

The trajectory prediction problem is an important part of autonomous systems, in both robot navigation and autonomous driving applications. This problem improves tracking accuracy and prevents collision between autonomous robotic platforms and moving objects (people or other autonomous platforms). This project focuses on analyzing pedestrian trajectories based on visual data extracted from videos. The analysis can include:

- a tracking component applied on individuals or groups of people
- an observation component which computes the general flows of pedestrians and constructs heat maps based on them
- an estimation component to predict possible trajectories of the tracked targets.

The solution will be adapted to outdoor environments, which may imply very crowded places.

Subject 5

Title: Elderly assistance in intelligent environments

Coordinators/Contacts:

**Prof. dr. ing. Adina Magda Florea (adina.florea@cs.pub.ro),
Conf. dr. ing. Irina Mocanu (irina.mocanu@cs.pub.ro),
dr. ing. Alexandru Sorici (alexandru.sorici@cs.pub.ro)**

Description:

The aim of this research is to develop a system for the elderly, which also allows family members or other carers to provide help and assistance when needed. Existing systems dedicated to the elderly focus on either safety (alert buttons or fall detectors) or health (for patients with chronic diseases) or consist only of simplified phones / smartphones. The system will offer all these capabilities, while providing a permanent "direct link" between the elderly and their family. Moreover, while the existing solutions are largely aimed at the environment of the old people's homes, The system aims to offer ubiquitous services both indoors and outdoors (smart homes and cities). It is based on a portable device carried by the elderly and a web application hosted on secure servers for family or other informal caregivers. GW (smartphone, small, light and low power) aims to act as a gateway that allows the collection of data, directly or indirectly, from a large number of sensors and platforms. The Web application server allows caregivers to remotely configure GW, monitor collected data, and receive timely alerts. Thus, the main objective of this research topic is to develop a solution for monitoring and meeting elderly needs for increased security and well-being. The solution aims to provide with a minimum of hardware, a platform that works as an assistant for the elderly, offering, on the one hand, contextual monitoring and adaptable services and, on the other, family support whenever needed. Developed in cooperation with the CITST company.

Subject 6

Title: Detecting drivers drowsiness

Coordinator/Contact:

Prof. dr. ing. Adina Magda Florea (adina.florea@cs.pub.ro)

Description:

The research aims at detecting driver's drowsiness.

<https://arxiv.org/ftp/arxiv/papers/1811/1811.01627.pdf>

Subject 7

Title: Human activity recognition from video sequences

Coordinator/Contact:

Prof. dr. ing. Adina Magda Florea (adina.florea@cs.pub.ro),

dr. ing. Mihai Trăscău (mihai.trascau@cs.pub.ro),

drd. ing. Mihai Nan (mihai.nan.cti@gmail.com)

Description:

The problem of recognizing human activity using the RGB image extracted from video is one of the first problems that computer vision has tried to solve. From the beginning, this problem was considered a challenging one, because there are many variables like that: the height of the person, the scene in which the action takes place, the brightness, the angle from which it is viewed, the fact that an action can be executed in a different manner from one person to another. The information provided by the video cameras has been extensively studied and analyzed as input for systems capable of identifying and recognizing human actions. Models applied initially for action classification used 2D images as features and classifiers such as Support Vector Machines (SVM) and Hidden Markov Models. Recently deep learning methods have emerged as a good candidate for human activity recognition.

The research topic implies developing different deep models for user activity recognition based on RGB images extracted from video sequences.

<https://escholarship.org/uc/item/2mr798mn>

<http://blog.qure.ai/notes/deep-learning-for-videos-action-recognition-review>

Subject 8

Title: User gesture recognition

Coordinator/Contact:

Prof. dr. ing. Adina Magda Florea (adina.florea@cs.pub.ro),

Description:

The purpose of this research is to recognize and track users gestures made with the hands from video sequences.

<https://towardsdatascience.com/tutorial-using-deep-learning-and-cnns-to-make-a-hand-gesturer-cognition-model-371770b63a51> <https://arxiv.org/abs/1901.10323>

<http://blog.qure.ai/notes/deep-learning-for-videos-action-recognition-review>

Subject 9

Title: Car crash detection in video using neural networks

Coordinator/Contact:

dr. ing. Mihai Trăscău (mihai.trascau@cs.pub.ro),

drd. ing. Mihai Nan (mihai.nan.cti@gmail.com)

Description:

Increasing the number of cars and excessive traffic congestion in cities is a major problem in the current time. Statistics show that more and more accidents happen daily, and many of these could be avoided.

The aim of this research topic is to develop a system capable of detecting the possibility of an accident by analyzing a video sequence. Starting from the available datasets (ex. CADP, CarCrashDetector), we propose to implement a system that combines Computer Vision and Machine Learning techniques to identify possible car crash.

Subject 10

Title: Dialogue Management Service supporting the Romanian language in assistive robotics scenarios

Coordinators/Contacts:

dr. ing. Alex Awada (alex.awada@cs.pub.ro),

dr. ing. Alexandru Sorici (alexandru.sorici@cs.pub.ro)

Description:

A dialogue management system, in the context of assistive robotics, is an application that models and guides a conversation carried out between a human user and a robot, by means of multiple modalities (e.g. written text, voice interaction), taking into account external cues (e.g. environment state, user fatigue or emotional state).

The purpose of this project is to develop a dialogue management service that works across different platforms and supports the Romanian language.

The service will be built in a modular, micro-service oriented architectural style, making use of existing web APIs that facilitate text-to-speech and speech-to-text processing for the Romanian language. The dialogue management service is expected to enable scriptable interaction scenarios, whereby the answers given by the robot are informed by both conversational context, as well as emotional state of the user (i.e. the response in the dialogue depends on what the user has said previously and on how he is feeling).

At the same time, the start of a dialogue can be triggered by contextual cues (e.g. the light is turned on in a room, a motion sensor is triggered), apart from direct voice interaction.

Subject 11

Title: Enhancing the functionality of the AMIRO User Interface

Coordinators/Contacts:

dr. ing. Alex Awada (alex.awada@cs.pub.ro),
dr. ing. Alexandru Sorici (alexandru.sorici@cs.pub.ro),
drd. Ing. Alexandru Gavril (alexandru.gavril@cti.pub.ro)

Description:

AMIRO (AMblent RObotics) is a ROS-based system enabling monitoring (e.g. detecting when the door of the lab is opened or closed) and actuation (e.g. raise or lower the blinds in the lab, turn the smart lights on/off or change their color) of an indoor lab environment, as well as access to external context information (e.g. health parameters of lab personnel) by a socially assistive robot.

The service also defines basic robot behaviors (e.g. navigating to a given position, identifying a person, searching for an object, speaking or listening for a voice command) that can be composed in a hierarchical manner to create more elaborate human-robot interaction scenarios (e.g. guiding a user to a location, finding a person in the lab to inform them of a notification).

The purpose of this project is to enhance the current UI (user interface) for the AMIRO system by creating a back-end tool (robot behavior management) that allows the creation of different robot behaviors by dragging, dropping and linking different elements (such as: predefined behaviors, predefined or new animations and dialogs, custom Python codes, etc.) into a scene. The tool output must support the export of constructed behavior compositions into tasks executable through the AMIRO system.

Subject 12

Title: UPB Mobile Application

Coordinators/Contacts:

dr. ing. Alex Awada (alex.awada@cs.pub.ro),
Conf. dr. ing. Irina Mocanu (irina.mocanu@cs.pub.ro)

Description:

he proposed project includes the development of a mobile application for the university, it should integrate three main functionalities:

- Information functionality: that presents different information about the university: faculties, departments, research centers, professors...
- Events functionality: that contains the different events that are organized in or by UPB (workshops, conferences, social events...) and information that are related to those events with the possibility for the user to subscribe to an event and receive related notifications.
- Navigation functionality: that allows the user to navigate easily through the whole campus (google maps & augmented reality - [Link 1](#) | [Link 2](#)).

The proposed project includes also the development of an administrative tool for the application. The tool allows the admin to easily add information and events, to visualize the users and to manage their subscriptions as well to send notifications (to all users or a specific group).

Subject 13 (for students in the ISI Master program)

Title: Context-awareness in a Sensing-as-a-Service Application Paradigm

Coordinators/Contacts:

dr. ing. Alexandru Sorici (alexandru.sorici@cs.pub.ro),

dr. ing. Mihai Trăscău (mihai.trascau@cs.pub.ro),

Conf. dr. ing. Andrei Olaru (cs@andreiolaru.ro)

Description:

In the domain of Ambient Intelligence, one of the most important research directions is the creation of scalable, long-lived information management infrastructures to facilitate search/discovery and querying of context information retrieved from physical or virtual sensors deployed in smart environments (e.g. smart homes/offices, smart cities).

One set of application scenarios, where *scale*, *long-lived interactions* and *openness* come into play as decisive factors is the paradigm of *sensing-as-a-service*.

In this view, the sensor owner, the sensor information aggregator/analyzer and the analysis consumer are all separate entities, with possible competing goals.

This means that dynamic search, discovery, query and interaction protocols have to be set up, by which producers and consumers of sensor information reach the desired level of information sharing.

Starting from a smart refrigerator application (envisaged in partnership with the software consulting company Infosys), where patterns of food consumption can act as primary information for retailers, received in exchange for discounts, the objective of this research topic is the development of a web-enabled context management system featuring the following:

- An explicit semantic representation for the *intended use* of context information (e.g. a retailer must make explicit the fact that he is only interested in information related to observed dairy products in a user fridge, with the *goal* to collect data about global dairy product usage)
- A scalable web application by which *search*, *discovery* and *querying* of context information can be performed by respecting *locality of context consumption* rules (e.g. a retailer should only be able to *find* the fridge sensors of a given user, if explicitly allowed by that user and only for a well defined scope; an assistive social robot can only discover and interact with the sensors in a room, when physically present in that room) [Sorici et al, 2019]
- A means to automatically/dynamically create *semantic social links* between sensing/actuation capable devices, depending on the relations of users/organizations who use them [Ciortea et al, 2016]

The research work in design and implementation of the system will explore the use of several of the most recent semantic web technologies, such as RDF Streams, Linked Data, W3C standards in sensor and actuator descriptions, RESTful web-services, FI-WARE enablers for context information sharing.

[Sorici et al, 2019] Sorici, Alexandru, Andrei Olaru, and Adina Magda Florea. "Towards Enabling Internet-Scale Context-as-a-Service: A Position Paper." In *Companion Proceedings of The 2019 World Wide Web Conference*, pp. 668-671. ACM, 2019.

[Ciortea et al, 2016] Ciortea, Andrei, Antoine Zimmermann, Olivier Boissier, and Adina Magda Florea. "Hypermedia-driven Socio-technical Networks for Goal-driven Discovery in the Web of Things." In *Proceedings of the Seventh International Workshop on the Web of Things*, pp. 25-30. ACM, 2016.

Subject 14

Title: Human Daily Activity Recognition

Coordinator/Contact:

Conf. dr. ing. Irina Mocanu (irina.mocanu@cs.pub.ro)

Description:

The proposed project consists in recognizing daily human activities by analyzing RGB-D videos. The proposed method will analyze videos with many persons, each of them performing the same or other activity.

<https://paperswithcode.com/task/action-recognition-in-videos>

<http://blog.qure.ai/notes/deep-learning-for-videos-action-recognition-review>

Subject 15

Title: Pose Estimation and Tracking Using RGB+D Data

Coordinator/Contact:

Conf. dr. ing. Irina Mocanu (irina.mocanu@cs.pub.ro)

Description:

The proposed project consists in detection and tracking the poses of multiple persons in an unconstrained setting. The method must have to deal with large pose and scale variations, fast motions, and a varying number of persons and visible body parts due to occlusion or truncation.

<https://paperswithcode.com/task/pose-tracking>

<https://posetrack.net/>

Subject 16

Title: Fashion design assistant for creative support using GAN models

Coordinators/Contacts:

dr. ing. Mihai Trăscău (mihai.trascau@cs.pub.ro),

dr. ing. Alexandru Sorici (alexandru.sorici@cs.pub.ro)

Description:

Topic designed in partnership with Infosys Consulting Romania

Generative Adversarial Networks (GANs) have been implemented with great success in many computer vision tasks, like image-to-image translation [1], conditional synthesis [2], data augmentation, style transfer [3], image super-resolution [4], etc.

In this project we aim to use GAN models to support creative efforts for the fashion design task. [5, 6, 7]. Starting from a given dataset of fashion articles, the goal is to generate new items conditioned by style considerations and desired attributes of the target.

The process of garment generation can be further improved/expanded to lead to the generation of 3D meshes of the clothing items, by further conditioning the GAN models on person specific physiological attributes, such as height, waist size, etc.

The physiological attributes are directly determined by a *person scan* procedure, carried out using the Kinect v2/v3 or Intel RealSense RGB-D cameras.

- [1] ZHU, Jun-Yan, et al. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *Proceedings of the IEEE international conference on computer vision*. 2017. p. 2223-2232.
- [2] PARK, Hyojin; YOO, YoungJoon; KWAK, Nojun. Mc-gan: Multi-conditional generative adversarial network for image synthesis. *arXiv preprint arXiv:1805.01123*, 2018.
- [3] ISOLA, Phillip, et al. Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017. p. 1125-1134.
- [4] WANG, Xintao, et al. Esrgan: Enhanced super-resolution generative adversarial networks. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018.
- [5] ZHU, Shizhan, et al. Be your own prada: Fashion synthesis with structural coherence. In: *Proceedings of the IEEE International Conference on Computer Vision*. 2017. p. 1680-1688.
- [6] SBAI, Othman, et al. Design: Design inspiration from generative networks. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018.
- [7] KANG, Wang-Cheng, et al. Visually-aware fashion recommendation and design with generative image models. In: *2017 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2017. p. 207-216.

Subject 17

Title: A Distributed and Flexible Testbed for Agent Learning

Coordinator/Contact:

Conf. dr. ing Andrei Olaru (cs@andreiolaru.ro)

Description:

Almost ever since ML has become a scientific interest, researchers have also become interested in the case where multiple entities learn and act in the same environment. The field of multi-agent learning (MAL) has, however, evolved more slowly, as MAL is a more difficult problem than single-agent learning, adding an additional layer of complexity by multiplying the number of agents. Moreover, the computational resources needed to perform multi-agent learning increase much faster than the number of agents itself. The purpose of the project is to **develop new multi-agent learning techniques** through the creation of a **holistic framework for the design, deployment, and execution of MAL simulations**. The framework will support representative cases and applications in the field and will enable the comparison of MAL applications. The framework will provide the necessary building blocks for enabling rapid development of new multi-agent learning methods and distributed execution of experiments in high-performance computing environments.

Subject 18

Title: Human pose estimation

Coordinators/Contacts:

dr. ing. Mihai Trăscău (mihai.trascau@cs.pub.ro),

drd. ing. Mihai Nan (mihai.nan.cti@gmail.com)

Description:

Human pose estimation is defined as the problem of localizing and identifying anatomical key-points of the human body and it is considered a fundamental and challenging task in Computer Vision. Poses can serve as base features in other vision problems like activity recognition, virtual and augmented reality, human reidentification or in human-computer interaction. This task can be very useful for solving the problem of recognizing human actions using a skeleton-based approach.

This research project aims to create a module capable of extracting human skeleton coordinates from RGB images. There are many challenges which need to be taken into account when inferring human pose. First, the number of people in the image is unknown and usually varies. Tackling this gives the option of either iteratively processing each person instance in the image or to attempt to obtain all human poses simultaneously. Moreover, the image may contain people who are in contact (e.g. shaking hands, carrying one another) making it difficult to assign the correct correspondence between person and key-point or body part. Occlusion, be it of actual key-points or not, also affects accuracy quite drastically.

[Zhe Cao et al.] "OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields". In: CoRR abs/1812.08008 (2018). arXiv: 1812.08008. URL : <http://arxiv.org/abs/1812.08008>.

[Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos.] "Densepose: Dense human pose estimation in the wild". In: arXiv preprint arXiv:1802.00434 (2018).

[Zhe Cao et al.] "Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields". In: CVPR. 2017.

[Shih-En Wei et al.] "Convolutional pose machines". In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2016, pp. 4724–4732.

Subject 19

Title: Human action prediction

Coordinators/Contacts:

dr. ing. Mihai Trăscău (mihai.trascau@cs.pub.ro),

drd. ing. Mihai Nan (mihai.nan.cti@gmail.com)

Description:

Future human action prediction is a probabilistic process that aims to identify ongoing action from video only containing the beginning part of the action. For this task, the goal is to allow early recognition of unfinished action from temporally incomplete video data. In contrast to the task of human action recognition, human action prediction is a before-the-fact video understanding task and is focusing on the future state. This aspect makes this task very useful

for many real-life scenarios where various critical situations could be avoided if they could be predicted quickly enough.

The purpose of this research project is to implement a framework for future human action prediction, integrating several modules (e.g. human pose estimation, object detection) and analyzing various approaches.

Subject 20

Title: Depth estimation for surrounding vehicles in driving scenarios

Coordinators/Contacts:

dr. ing. Alexandru Sorici (alexandru.sorici@cs.pub.ro),

drd. ing. Iancu David (davidiancudvd@gmail.com)

Description:

Autonomous driving is a hot research topic today but there are still a lot of challenges to solve until we can trust a completely autonomous car. One of the biggest problems is how to estimate the distance to the surrounding cars, to approximate the breaking time and distance, to avoid collisions, etc. There are some sensors or even systems that rely on a combination of sensors, for example cars equipped with radar, lidar, maybe even other sensors (GPS, IMU, etc) that can approximate the distance to the cars. However, there are not 100% accurate and also there are very expensive, so the ideal case would be to estimate the depth using only video cameras, which are much less expensive.

There are some algorithms for estimating the depth from two cameras using some mathematical observations, but they require similar cameras that are very close to each other. The most interesting approach regarding Artificial Intelligence is to estimate the depth using deep neural networks (without mathematical approaches) from a single camera.

The purpose of the project is to evaluate the state of the art regarding depth estimation from single camera and to improve the existing models in order to have a better estimation for driving scenarios recorded at University Politehnica of Bucharest.

Subject 21

Title: Context-aware fake news detection

Coordinators:

S.L. dr. ing. Ciprian-Octavian Truică (ciprian.truica@cs.pub.ro)

Conf. dr. ing. Elena-Simona Apostol (elena.apostol@cs.pub.ro)

Description:

Fake news consists of news that is intentionally and verifiably false, and which could mislead readers by presenting alleged, imaginary facts about social, economic and political subjects of interest. This topic aims to design and implement new Natural Language Processing, Machine Learning, and Deep Learning models that use context for detecting fake news using new context features extracted from the textual information.

Implementation Language: Python

Libraries: Scikit-learn, PyTorch, Keras with Tensorflow, etc.

Storage & Distribution Technologies: MongoDB, Spark.

Subject 22

Title: Context-aware harmful speech mitigation on social media

Coordinators:

S.L. dr. ing. Ciprian-Octavian Truică (ciprian.truica@cs.pub.ro)

Conf. dr. ing. Elena-Simona Apostol (elena.apostol@cs.pub.ro)

Description:

The spread of social networks has seen an increase in the use of harmful speech. This type of speech is defined as a form of expression that can increase the risk that its audience will condone or participate in violence against members of another group. This topic aims to design new Natural Language Processing, Machine Learning, Deep Learning, Network Analysis, and Graph Mining models, techniques, and strategies that use context and network information for mitigating the damaging effect of harmful speech on social networks.

Implementation Language: Python

Libraries: Scikit-learn, PyTorch, Keras with Tensorflow, etc.

Storage & Distribution Technologies: MongoDB, Spark.

Subject 23

Title: Network-aware fake news mitigation on social media

Coordinators:

S.L. dr. ing. Ciprian-Octavian Truică (ciprian.truica@cs.pub.ro)

Conf. dr. ing. Elena-Simona Apostol (elena.apostol@cs.pub.ro)

Description:

Fake news consists of news that is intentionally and verifiably false, and which could mislead readers by presenting alleged, imaginary facts about social, economic and political subjects of interest. This topic aims to design and implement new Network Analysis and Graph Mining models and strategies for mitigating the spread of fake news on social media using information extracted from the diffusion medium.

Implementation Language: Python

Libraries: Scikit-learn, PyTorch, Keras with Tensorflow, etc.

Storage & Distribution Technologies: MongoDB, Spark.

Subject 24

Title: Context and network-aware fake news detection and mitigation

Coordinators:

S.L. dr. ing. Ciprian-Octavian Truică (ciprian.truica@cs.pub.ro)

Conf. dr. ing. Elena-Simona Apostol (elena.apostol@cs.pub.ro)

Description:

Fake news consists of news that is intentionally and verifiably false, and which could mislead readers by presenting alleged, imaginary facts about social, economic and political subjects of interest. This topic aims to design new Natural Language Processing, Machine Learning, Deep Learning, Network Analysis, and Graph Mining models, techniques, and strategies that use context and network information for mitigating the spread of fake news on social networks.

Implementation Language: Python

Libraries: Scikit-learn, PyTorch, Keras with Tensorflow, etc.

Storage & Distribution Technologies: MongoDB, Spark.

Subject 25

Title: Gravitational-wave detection

Coordinators:

S.L. dr. ing. Ciprian-Octavian Truică (ciprian.truica@cs.pub.ro)

Conf. dr. ing. Elena-Simona Apostol (elena.apostol@cs.pub.ro)

Description:

Recent research on developing new algorithms to optimize the search for gravitational waves has led to integrating machine learning techniques to the field of astrophysics. The aim of this research is to develop new Anomaly Detection and Change Point Detection models and techniques using Machine Learning and Deep Learning for analyzing time series data for detecting gravitational waves.

Implementation Language: Python

Libraries: Scikit-learn, PyTorch, Keras with Tensorflow, etc.

Storage & Distribution Technologies: MongoDB, Spark.

Subject 26

Title: Domain-specific multi-word extraction and abbreviation disambiguation for e-Health

Coordinators:

S.L. dr. ing. Ciprian-Octavian Truică (ciprian.truica@cs.pub.ro)

Conf. dr. ing. Elena-Simona Apostol (elena.apostol@cs.pub.ro)

Description:

Hospitals collect huge amounts of data about their patients every year, in various ways. Such massive sets of data can provide great knowledge and information which can improve the medical services, and overall the healthcare domain, such as disease prediction by analyzing the patient's symptoms or disease prevention, by facilitating the discovery of behavioral factors that can turn into risk factors for disease. Unfortunately, only a relatively small volume of e-Health data is processed and interpreted, an important factor being the difficulty in efficiently performing Big Data operations, and that often this data, even if anonymous, is hard to obtain. This topic aims to design and implement new Natural Language Processing, Machine Learning, and Deep Learning methods and models for disambiguating abbreviations and extracting domain-specific multi-words to better understand medical documents and correlate medical terms with their meanings.

Implementation Language: Python

Libraries: Scikit-learn, PyTorch, Keras with Tensorflow, etc.

Storage & Distribution Technologies: MongoDB, Spark.