Three sides of AI for Industry 4.0

Ermanno Cordelli, Luisa Francini, Giulio Iannello, Rosa Sicilia, Paolo Soda, Matteo Tortora
Unit of Computer Systems and Bioinformatics, Universita' Campus Bio-Medico di Roma
{e.cordelli,l.francini,g.iannello, r.sicilia,p.soda,m.tortora}@unicampus.it

Abstract

Artificial Intelligence is disrupting industrial markets and forcing enterprises to reevaluate how traditional work is performed such as workforce training, advertising, communication and reputation monitoring, maintenance and repair, operations forecasting and scheduling. In this contribution we overview rumour detection in social media, fleet maintenance operations on the basis of road quality analysis, and representation methods for multivariate time-series. Such topics are useful in industry applications to create business value in different areas, such as marketing, knowledge discovery and energy management.

1 Introduction

The global artificial intelligence (AI) market size was valued at 51.56 billion USD in 2021 and it is expected to expand at a compound annual growth rate of 40.2% from 2021 to 2028 [BusinessWire, 2021]. The continuous research and innovation directed by the tech giants are driving the adoption of advanced technologies in several main areas of industry, such as automotive, healthcare, retail, finance, and manufacturing. However, technology has always been an essential element for these industries, and AI has brought technology to the center of organizations [Grand View Research, 2021].

Among the different industry verticals, in our laboratory we are directing our efforts toward three issues. The first concerns rumour detection, an hot-topic of research since today an increasing number of users inform themselves via social media instead of traditional news sources [L.Wong, 2017]. Social media can be accessed by everyone, so even ordinary citizens can report events as well as their own feelings and experiences, but the absence of systematic control of posts on these platforms may lead to spread misinformation, especially in the context of breaking news and marketing information, which often appear first on social media and only after on traditional media systems [Hamidian e Diab, 2019]. Let us recall that among the different kinds of misinformation, a rumour is defined as an unverified news in circulation with an instrumental value and likely to be dangerous [DiFonzo e Bordia, 2007]. In this respect we are focusing on Twitter platform that, due to how it allows posts to be created, facilitates the real-time dissemination of news, which causes rumour to spread rapidly with many consequences.

The second topic deals with predictive fleet maintenance. In recent years, the Internet of Things (IoT), big data and AI have made possible predictive maintenance applications applied to fleets of different vehicles [Killeen e others, 2019]. Among the various vehicles that can be considered, several work focused on public transport buses, attempting to diagnose faulty buses that deviate from the rest of the bus fleet [Rögnvaldsson e others, 2018]. To this goal, they used many on-board signals sampled from electronic control unit by CAN (Controller Area Network) based protocol, further to the GPS signal. To process the data so collected there exist a pletora of different methods that, in general, are self-organised agents whose outputs are compared against each other to find systems which deviate from the consensus [Rögnvaldsson e others, 2018]. Nevertheless, the quality of the roads impact both the driving style as well as the amount of mechanical stress to which the buses are subjected. This motivates us to mine together sensor data with information related to the quality of the roads in the bus routes, which will be extracted using computer vision techniques from Google Street View images.

The third topic deals with multivariate time series (TS). Indeed, the rising capabilities of storing and registering data has increased the number of temporal datasets, boosting the attention on TS classification and forecasting. In the multivariate case, symbolic methods that try to predict phenomena transform the data into a more compact format to produce a representation of the time series easy to be handled in a machine learning framework. However, these representations do not grasp information on both inter-attribute variability and temporal variability. Hence, we developed an approach that, taking into account the relationships between attributes and their periodicity, reduces the multivariate TS to a collection of symbols, whose distribution is represented by histograms.

2 Rumour detection

The breakthrough of social media has boosted to an increase in the spread of misleading information, with a serious impact on society especially when related to health knowledge. Recently, researchers have been developing AI-based automatic systems to detect rumours in social microblogs and, within this context, we are facing with the use and integration of multimodal data available from the network to develop machine learning-based rumour detection systems.

Work on rumour detection can be divided into *macro-level* and *micro-level*. The first case considers as rumours news carried by a set of microblog posts (e.g. conversational thread, topic, event, etc.), while the second approach aims to discriminate among the single posts which one are rumours from those that are not. This type of study is more effective since in many cases it is possible to find that the subject of a conversation may not be a rumour, but actually several of the posts in the conversation are. Hence, in all our work on this topic we focus on this second rumour detection task tackling different open challenges in this field: (i) how to transfer knowledge between different trending topics in order to build a system robust to of newly rumours on a previously unseen topic; (ii) how to efficiently represent the multimodal data with handcrafted descriptors; (iii) how to exploit the multimodality of the data without employing deep learning strategies, which need a great amount of labelled data and are computationally expensive.

2.1 Methods and main results

We hereby describe three main contributions with the achieved results: in the first we tackle the transfer learning problem in a rumour detection task for the health care domain ("Knowledge Transfer"); in the second, we focus on the comparison between two well-established handcrafted representations that model different properties of the multimodal problem of rumour detection on Twitter ("Comparison of two handcrafted representations"); in the last one we tackle the rumour detection problem from a different angle leveraging the multimodal nature of the task ("multimodal meta-classifier"). The presentation order of our three approaches corresponds to the respective open challenge cited in the previous section. It is worth specifying that in all contributions we tackle the binary task rumours vs. non-rumours.

Knowledge Transfer: in [Sicilia e others, 2021a] we cope with the challenge of exploiting traditional feature-based Transfer Learning (TL) approaches to identify rumour Twitter posts in an unlabelled health-related topic (*#Vaccine* and *#Zikavirus* dataset). To this end we first compare two state-of-the-art representations for micro-level rumour detection (User-Network [Francini e others, 2021] and Social-Content [Zubiaga e others, 2017]) and then we experimentally investigate three TL methods in an unsupervised scenario (Adaptation Regularization-based Transfer Learning, Transfer Kernel Learning and Graph co-regularized Transfer Learning). The comparison with a baseline that does not use any knowledge transfer from the source and target domains reveals that negative transfer occurs probably due to the small dataset size which does not provide enough information to fill the gap between the source and the target domain.

Comparison of two handcrafted representations: in [Francini e others, 2021] we offer two main contributions: (i) we compare the two state-of-the-art handcrafted representations User-Network and Social-Content designed for developing machine learning-based rumour detection systems, considering different machine learning classifiers over a pool of 8 Twitter datasets (5 of them are public) for micro-level rumour detection; (ii) we present two novel datasets (*#RoyalWedding* and *#TrumpRussia*) retrieved from Twitter and manually labelled with a robust methodology. To compare the two representations, we set up an experimental assessment implementing a Leave-One-Topic-Out evaluation on 8 different topics, so the representations would be evaluated for their discriminative power in recognizing rumours in tweets of a topic not present in the training set; this is a solution to model a real environment where a classifier trained on a known set of topics might be challenged to discover rumours regarding unknown new topics. Despite the low results obtained for both representations, our aim is to compare two feature groups in order to see which one is more informative. As main result we find out that the User-Network set of feature is more stable to topic changes.

Multimodal meta-classifier: in this work, still unpublished, we tackle the rumour detection problem considering its intrinsic multimodality. We exploit the robust handcrafted representation analysed in [Francini e others, 2021], i.e. the User-Network representation, breaking it up into the three modalities User, Text and Network. This approach leverages on three unimodal cost-sensitive random forests, which are combined by a a multimodal meta-classifiers trained to select the best modality for each test sample using an estimate of unimodal posterior probabilities. The method was tested in Leave-One-Topic-Out on the same 8 Twitter datasets of the previous work, outperforming state-of-the-art competitors.

2.2 Challenges and perspectives

We identify three main future challenges. First, with reference to knowledge transfer work, we would investigate other TL techniques in unsupervised and semi-supervised scenarios which are not based on feature transfer. Second, with respect to the multimodal meta-classifier, we would study the robustness of the method in the missing-modality scenario, i.e. cases when it is not possible to extract all the modalities from a dataset. Third, we also plan to investigate deep learning methodologies.

2.3 Papers and available resources

We have been working on the this topic since 2017 and, further to work already mentioned, others are [Sicilia e others, 2018; Sicilia e others, 2021b]. Furthermore, we made available upon request several annotated datasets.

3 Road quality detection for fleet maintainance

Advances in AI, Big Data and IoT have made possible predictive maintenance (PdM) applications. Unlike previous approaches, PdM aims to predict the most appropriate time for maintenance operations, by taking into account system health information and/or historical maintenance records [Theissler *et al.*, 2021]. When dealing with vehicles under the same operational conditions, collective information can be used and it can be referred to as predictive fleet maintenance. Several work have focused on buses in public transport, trying to diagnose faulty buses whose behaviour deviates from the rest of the bus fleet [Rögnvaldsson e others, 2018]. Nevertheless, in fleet management, the use of data collected from vehicle sensors should not be enough to develop ad hoc maintenance and repairs. In this context we aim to mine multimodal data collected from vehicle sensors and images of bus routes extracted from Google Street View for better PdM applications. The idea is that multimodal data sources allows the extraction of a complementary, more robust and richer data representation than a stand-alone modality.

Materials This work is part of a PhD project funded by FSTechnology, a company of the Gruppo Ferrovie dello Stato Italiane. FSTechnology provided us with the data of a fleet of 21 buses of the Busitalia group. Each bus is characterized by 123 signals sampled asynchronously and with different sampling frequencies, acquired with the logic *as a new data exists, send it* and collected in a central server. The sampled signals are either discrete signals, such as the door opening status, the cooling air conditioning, and the fuel level status, or continuous signals, such as the steering wheel angle, the brake pedal position, and the engine percent torque. In addition, each bus provides us with GPS signals acquired using the same logic as above. This information is used to reconstruct through the Google Street View Static API query the videos of the bus routes.

3.1 Methods

As already mentioned, the available bus fleet provides us with GPS data for each bus sampled at a variable sampling rate. These GPS signals can be used to reconstruct the route of the buses and thus allow us to add an important element to the predictive maintenance problem. Indeed, by reconstructing the route we can analyse the conditions of the roads the bus has travelled on.

The proposed road quality detection system is depicted in Figure 1 and consists of the following three main blocks:

1. **Map-Matching algorithm & data oversampling**: Due to the built-in error of the GPS sensor, we have to snap the GPS data point to the correct location on the map and reconstruct the bus route. This is performed with a map-matching algortihm followed by an oversampling process. Using the Python library OSMnx we extracted the road network graph from OpenStreetMap, consisting of nodes and edges, where each node represents an intersection and each edge a section of road. To obtain the complete path, the map-matching algorithm finds the edges of the graph closest to the sampled point using the rtree algorithm and the projection of the point onto the edge and then searching for the missing edges using a shortest path search algorithm.
2. **Bus route video extraction**: Once we have aligned and oversampled the geospatial points, we can use this information to query the Google Street View Static API. Since the API provides 360° panoramic images, we use the route oversampling information to set the camera bearing property. Once the route images have been extracted, they can be merged into a video.
3. **Road condition evaluation**: We plan to perform road condition analysis using computer vision techniques applied directly to the extracted video. This block would consist of a first segmentation phase to select only the portion of the road in the image, which is followed by a road condition evaluation. This latter step can be performed through various computer vision approaches. We plan to solve it in an unsupervised way by applying a bag-of-words model-based clustering on patches extracted via sliding window, that may provide a fine characterisation of the regions of the road and is well suited if no labelled dataset is available, as in our case.

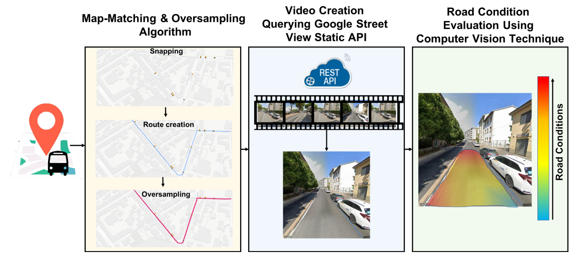

Figure 1: Proposed road quality detection system consisting of three blocks identified by the coloured boxes.

3.2 Challenges and perspectives

We identify three different directions for further research. First, as already mentioned, we can use the road quality detection system outputs in a multimodal learning framework merging them with the bus signals previously described, to improve the capabilities of the predictive maintenance system. Second, we can use the dataset provided by FSTechnology to develop a system to evaluate the driving style of a driver. This work could have several potential applications, such as encouraging drivers to adopt a safer driving style, or it could be used by insurance agencies to reward their customers for their safe behaviour. Finally, the road condition analysis performed through the Google Street View 360° panoramic images can expand and evolve into a stand-alone line of business for the public administration.

3.3 Available resources

The map-matching algorithm implemented is mainly based on the following two Python libraries: OSMnx and NetworkX [Boeing, 2017; Hagberg *et al.*, 2008]. The first is a library that allows to download geospatial data from OpenStreetMap and model, project, visualise and analyse real-world road network graphs. The latter is a library that allows the creation, manipulation and analysis of the structure, dynamics and features of complex graph networks. Instead, the images for analysing road conditions were collected by querying the Google Street View Static API, which provides a set of geolocated 360° panoramic images.

4 Multivariate time series

Representation, classification and forecasting of TS are important tasks with several challenging applications such as speech recognition, financial analysis, environmental modeling and, more recently, Internet of Things (IoT) Univariate time series (UTS) is a sequence of data points, measured typically at successive points in time and spaced at uniform time intervals, whereas a multivariate time series (MTS) is a collection of UTS, each named as an *attribute*. Both for UTS and MTS, either in the case of classification or in the case of forecasting, data representation is a key issue that has attracted research efforts [Fu, 2011].

Methods for MTS classification and regression can use generative models, such as Hidden Markov models or vector autoregressive models, artificial neural networks, distance-based approaches or symbolic representations. Several experience indicates that symbolization can offer a unique potential for computational efficiency, ease of visualization, and connections with AI. Nevertheless, being able to grasp both the inter-attribute and the temporal variabilities in symbolic MTS representations is an open issue that we addressed in [Soda e others, 2019] via a symbolization approach. It first builds a codebook using all time stamp in the MTS; second, it explores the periodicity of the data and, third, it transforms the data into symbols via a two step transformation based on the codebook and on the periodicity. Fourth, histograms of symbols extracted at different locations are concatenated to get the MTS representation. The approach was successfully tested on a publicly available dataset, the Telecom Italia Big Data Challenge 2014 dataset, to forecast energy consumption in the province of Trento. It also favorably compared with results attained by other methods available in the literature.

4.1 Challenges and perspectives

We identify four different directions for further research. First, to study how data granularity may impact the representation method proposed. Second, to apply the proposed representation to other datasets where there is the need to grasp the inter-attribute and temporal variability, experimentally investigating also how the method is affected by TS attributes number. Third, to study if and how the proposed symbolic representation can be applied in stream mining to detect data drifts. Fourth, to investigate how this representation can boost the trustworthy of MTS forecasting.

Acknowledgements

This work is partially supported by two PhD grants offered by FSTechnology SpA and INPS.

References

[Boeing, 2017] Geoff Boeing. Osmnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, 65:126–139, 2017.

[BusinessWire, 2021] BusinessWire. Global Artificial Intelligence Market Report (2021 to 2030) - COVID-19 Growth and Change - ResearchAndMarkets.com. https://www.businesswire.com/news/home/20210816005399/en/Global-Artificial-Intelligence-Market-Report-2021-to-2030---COVID-19-Growth-and-Change---ResearchAndMarkets.com, 2021.

[DiFonzo e Bordia, 2007] Nicholas DiFonzo e Prashant Bordia. *Rumor psychology: social and organizational approaches.* American Psychological Association, 2007.

[Francini e others, 2021] Luisa Francini et al. Describing rumours: a comparative evaluation of two handcrafted representations for rumour detection. In *Int Conf on Information and Digital Technologies*, pages 311–318. IEEE, 2021.

[Fu, 2011] Tak Chung Fu. A review on time series data mining. *Engineering Applications of Artificial Intelligence*, 24(1):164–181, 2011.

[Grand View Research, 2021] Grand View Research. Artificial intelligence market size, share trends analysis report (2021 - 2028). https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-ai-market, 2021.

[Hagberg *et al.*, 2008] Aric Hagberg, Pieter Swart, e Daniel S Chult. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab, 2008.

[Hamidian e Diab, 2019] S. Hamidian e Mona T. Diab. Rumor detection and classification for twitter data. *ArXiv*, abs/1912.08926, 2019.

[Killeen e others, 2019] Patrick Killeen et al. Iot-based predictive maintenance for fleet management. *Procedia Computer Science*, 151:607–613, 2019.

[L.Wong, 2017] J. Burkell L.Wong. Motivation for sharing news on social media. *Proc. 8th Int. Conf. Social Media Soc.*, pages 1–5, 2017.

[Rögnvaldsson e others, 2018] Thorsteinn Rögnvaldsson et al. Self-monitoring for maintenance of vehicle fleets. *Data mining and knowledge discovery*, 32(2):344–384, 2018.

[Sicilia e others, 2018] Rosa Sicilia et al. Twitter rumour detection in the health domain. *Expert Systems with Applications*, 110:33–40, 2018.

[Sicilia e others, 2021a] Rosa Sicilia et al. Representation and knowledge transfer for health-related rumour detection. In *IEEE 34th Int Symp on Computer-Based Medical Systems*, pages 591–596, 2021.

[Sicilia e others, 2021b] Rosa Sicilia et al. Rule-based space characterization for rumour detection in health. *Engineering Applications of Artificial Intelligence*, 105:104389, 2021.

[Soda e others, 2019] Paolo Soda et al. Grasping inter-attribute and temporal variability in multivariate time series. *IEEE Trans on Big Data*, 2019.

[Theissler *et al.*, 2021] Andreas Theissler, Judith Pérez-Velázquez, Marcel Kettelgerdes, e Gordon Elger. Predictive maintenance enabled by machine learning: Use cases and challenges in the automotive industry. *Reliability engineering & system safety*, 215:107864, 2021.

[Zubiaga e others, 2017] Arkaitz Zubiaga et al. Exploiting context for rumour detection in social media. In *Int Conf on Social Informatics*, pages 109–123. Springer, 2017.