Issue 8 / Summer 2022

Social Mining & Big Data Ecosystem

SoBigData RESEARCH INFRASTRUCTURE Magazine

Editorial

The SoBigData RI is consolidating its presence in the European landscape. In 2022, two new projects will become part of the SBD ecosystem, funding the research infrastructure.

On the one hand, the SoBigData RI Preparation Phase Project (PPP) will begin. It will be the first step of the SoBigData RI in the ESFRI RoadMap 2021. The SoBigData RI PPP consortium comprises 33 partners from 17 countries (Bulgaria, Belgium, France, Greece, Spain, Austria, Italy, Germany, Estonia, Finland, The Netherlands, Poland, Sweden, UK, and Switzerland). It aims to enhance the RI from the simple awareness of ethical and legal challenges in social mining to the development of concrete tools that operationalize ethics with value-sensitive design, incorporating values and norms for privacy protection, fairness, transparency, and pluralism.

The main objectives of the SoBigData RI PPP are to define operative strategies for modelling and definition of the legal entity; to produce and review a Business Plan

[Continues on pag. 3]

Inside this issue

3 **EDITORIAL**
Roberto Trasarti & Valerio Grossi (CNR)

5 **NEWS**
Editorial Board

7 **CALL FOR PAPERS**
Editorial Board

8 **RESEARCH HIGHLIGHTS**
Multiple authors

19 **TRANSNATIONAL ACCESS**
Multiple authors

22 **EVENTS HIGHLIGHTS**
Multiple authors

Editorial

The future of the SoBigData Research Infrastructure. 3

News

TransNational Access: a program of Short-Term Scientific Missions to carry forward your own big data project. 5

Call for Paper

IFDaD 2022: International Forum on Digital and Democracy. 7

Research Highlights

The good, the bad, and the ugly of inferring high-resolution poverty maps with multimodal data. 8

A new dataset to assess the long-tail effect of the COVID-19 lockdown on Italians 11

Detecting Mesoscale Structures by Surprise. . 13

Research and Development of Quantification Methods with QuaPy. 15

Human-centred Machine Learning with Visual Analytics. 17

TransNational Access

A TNA experience on Ordinal Quantification. 19

Dynamics of opinion polarization in social networks 20

Events Highlights

Challenge Us 2021: SoBigData meets Industry. 22

Highlights from the Event “Towards a Digital Ecosystem of Trust: Ethical, Legal and Societal Implications” 24

Editorial

The future of the SoBigData Research Infrastructure

Roberto Trasarti, SoBigData RI coordinator, ISTI-CNR, Pisa, Italy | roberto.trasarti@isti.cnr.it

Valerio Grossi, SoBigData RI project manager, ISTI-CNR, Pisa, Italy | valerio.grossi@isti.cnr.it

[Continues from pag. 3]

for long-term sustainability; to engineer, plan, and optimize the technical infrastructure, and define strategies for services design, community involvement and partnerships with third parties. These objectives pose several key challenges that must be addressed within the SoBigData RI PPP.

Due to the multidisciplinary nature of the SoBigData RI, identify and develop sustainable services will be needed, thus responding to the demand for cross-disciplinary data-driven research and innovation.

Particular attention will be devoted to the relationship between the central hub in Italy and the national nodes (distributed in the countries of the consortium) and the involvement of the Member States and Associated Countries in this structure (and the taking in of new partners). The definition of a legal framework implies challenges in defining an agreement regarding the central hub hosting and its relationship to national nodes (and national sites); in preparing the ground to become ERIC, finding financial support from countries involved in the consortium using specific agreements.

On the other hand, considering Italian context, the SoBigData.it project (related to Italian National Recovery and Resilience Plan funds) aims to strengthen the Italian node of the existing RI. The consortium includes 11 research excellence centres in Italy (Consiglio Nazionale delle Ricerche, Università di Pisa, Scuola Superiore di Studi Universitari e di Perfezionamento Sant'Anna, Scuola Normale Superiore, IMT School for Advanced Studies Lucca, Università degli Studi dell'Aquila, Sapienza Università di Roma, Università degli Studi di Genova, Alma Mater Studiorum – Università di Bologna, Università degli Studi di Palermo, Università degli Studi di Catania).

From a technological perspective, SoBigData.it will empower and create new computational nodes through data centres which will be connected with the SoBigData RI network as well as new technologies to enlarge its capability of tackling future social challenges.

Due to the multidisciplinary nature of the SoBigData RI, identify and develop sustainable services will be needed, thus responding to the demand for cross-disciplinary data-driven research and innovation.

Moreover, new kinds of technologies will be integrated into the RI for pervasive computing, edge computing, and communication networks beyond 5G research. From a scientific-research point of view, SoBigData.it will introduce the concept of Virtual Laboratories (VLs) collecting scientific contributions from the operative units involved in the project. A VL is an environment where users share information, execute experiments, and generally find resources related to a specific research line inside SoBigData.

Finally, in a broad vision of promoting responsible data science principles, SoBigData.it pones particular attention to the definition of Ethical aspects and aims at operationalizing the legal and ethical framework of the SoBigData RI in accordance with European and national legislations. This will ensure the development and application of ethical-by-design policies for large-scale dataset collection and analysis, generalize experiences, and create awareness in society and industry by offering a consultancy service.

To promote and incentivize the diffusion of data science, SoBigData.it will devise campaigns to foster awareness in high school students to attract bright, young minds, with particular attention to female students by promoting specific initiatives. Finally, the service provision perspective includes delivering a set of libraries composed of the algorithms developed in the virtual laboratories. All the libraries will be integrated and ready to use in virtual and trans-national access.

The upgraded SoBigData RI resulting from the SoBigData PPP and SoBigData.it projects will produce and review a plan for long-term sustainability, including details describing the nature of our core business (related to RI services), the RI financial projections, and all the strategies aimed at achieving sustainability. Finally, considering ESFRI requirements on long-term sustainability, a robust RI risk and performance management system will be deployed.

The technical infrastructure related to the nodes operationalizing SoBigData RI will be re-engineered, planned, and optimized. The future success of SoBigData RI is related to

the capacity to develop and operationalize specific strategies to support and create new research lines, involve new communities of researchers and new stakeholders, and disseminate (and advise) the RI services beyond our reference community and stakeholders.

Photo Credit: Robynne Hu, Unsplash.com

SoBigData Research Infrastructure (RI), *with its tools and services, empowers researchers and innovators through a platform for the design and execution of large-scale data science and social mining experiments, open to users with diverse backgrounds, cloud accessible, including supercomputing facilities.*

SoBigData RI is rendering social mining experiments *designed to enhance efficacy in an adjusted and repeatable manner by non-data scientists promoting FAIR (Findable, Accessible, Interoperable) and FACT (Fair, Accountable, Confidential, and Transparent) principles.*

The SoBigData RI was established *by the first H2020 SoBigData project in 2015 and is currently consolidated in the SoBigData++ project. The SoBigData RI features resources from multiple sources; e-infrastructures and online services developers; big data analytics and AI; complex systems focussed on modelling social phenomena; Ethical, Legal, Socio-Economic, and Cultural aspects of data protection; privacy-preserving techniques.*

TransNational Access: a program of Short-Term Scientific Missions to carry forward your own big data project

We welcome applications from individuals with a scientific interest, professionals, startups and innovators that may benefit from training in data science and social media analytics.

The SoBigData++ RI manages vertical, thematic environments, called exploratories, on top of the SoBigData infrastructures, for performing cross-disciplinary social mining research. The Transnational Activities offered in this call will be for Short-Term Scientific Missions (STSM), between 3 weeks and 2 months.

Under this call, there will be two kinds of proposals funded: **STSM research proposals** and **STSM tool/data integration proposals.**

Funding is available **up to 5000 euros** per participant (to cover the cost of daily subsistence, accommodation, and economy flights/train).

STSM bursaries are awarded on a competitive basis, according to the procedure described in the application pack and eligibility criteria below, and based upon the quality of the applicant, the scientific merit of the proposed project, and their personal statement.

Applications from female scientists are particularly encouraged.

Visitors are welcome subject to the host country's and host institution's **Covid-19 regulations.** We will consider offering up to a 6 month postponement of an accepted application if travel restrictions are imposed due to Covid-19.

APPLY NOW!

Visit our website
http://www.sobigdata.eu/content/call-2021-22-sobigdata-transnational-access

Pre-requisites for projects to carry out hosted research:
- Good understanding of social data and, ideally, track record of prior social data analysis projects;
- Experience with using at least one of machine learning, natural language processing, and/or complex networks algorithms.

Pre-requisites for projects to integrate new tools/datasets/services:
- An already existing open-source tool for social media mining to be integrated or an already created openly licensed dataset of relevance to SoBigData++, that can be integrated within the infrastructure.

The goal is to provide researchers and professionals with **access to big data computing platforms, big social data resources, and cutting-edge computational methods.**

STSM visitors will be able to:
- Interact with the local experts
- Discuss research questions
- Run experiments on non-public big social datasets and algorithms
- Present results at workshops/seminars

The STSM visits will enable multi-disciplinary social mining experiments with the SoBigData++ Research Infrastructure assets: big data sets, analytical tools, services and skills.

Photo credit: Joshua Woroniecki - Pixabay

IFDaD 2022: International Forum on Digital and Democracy

17-18 November 2022
Rome, Italy
https://www.ifdad.org

CALL FOR PAPERS

Since the first edition of the Forum, in December 2020, the stretching of pandemic control measures constrained furthermore individual freedoms. In some cases, it provided legitimacy to authoritarian shift in frail democratic governance. According to Democracy Index in 2021 global democracy continued its decline and experienced its biggest annual drop since 2010 after the global financial crisis. Whereas even the United States, the stronghold of democracy, in occasion of the Capitol attack, showed to be unshielded to the outburst of illiberal forces.

For fragile democracies difficult times lie ahead. However, if the spirit of democracy and its needs remains strong among people, recovery hopes are also to be placed in the righteous use of digital technologies.

The expected outcome of the Forum is to gain a greater understanding and shared view on how digital technologies can positively affect electoral process and pluralism. As a follow-up of the debate of the first edition, the Forum will explore and develop the interplay between digital technology and data, political participation, and governance. Altogether with two additional focus: the functioning of eGovernment and disinformation.

Internet and algorithms caused a transformation of the centennials democratic processes. When exploited for political purposes digital tools affect the functioning of institutions created to evolve at human speed and not at the exponential pace and massive scale of digital technologies.

In the effort to gain a better understanding on how democracy can be strengthened in an era of rapid technological changes, the Forum engages scholars and experts in a Call for Papers.

For further information, visit
https://www.ifdad.org/

The good, the bad, and the ugly of inferring high-resolution poverty maps with multimodal data

Traditionally, socioeconomic indicators of countries are coarse-grained and inferred through surveys or census data. One limitation of such a process is that it can take years to plan, execute and process the vast amount of collected data. A further limitation is that certain regions are difficult to access, and therefore indicators for such regions end up missing. In this study, we propose machine learning models to produce high-resolution poverty maps in African countries. To this end, we combine various kinds of data that can serve as proxies of poverty and social segregation. These freely available data sources include satellite images, infrastructure, population, and mobility networks. Our models explain more than 80% of the variation of wealth, which outperform previous approaches and help to make better decisions on feature selection and policy.

Lisette Espín-Noboa, Central European University (CEU) | EspinL@ceu.edu

János Kertész, Central European University (CEU) | KerteszJ@ceu.edu

Márton Karsai, Central European University (CEU) | KarsaiM@ceu.edu

Photo Credit: Sergio Rola - Unsplash.com

The first Sustainable Development Goal set by the United Nations is to eradicate poverty by 2030 [L1]. Although the number of people living in extreme poverty has been decreasing since 2018, the decline in poverty rates has slowed down ever since. This is partly due to the COVID-19 pandemic, but the ongoing impacts of conflicts and natural catastrophes set further barriers for progress in this direction [L2]. Traditional data collection methods, such as census or survey, fail to follow the effects of such rapid changes, therefore new techniques are required to capture the aftermath of unexpected global-impact events and processes.

The identification of places-in-need requires rapid, flexible, and precise

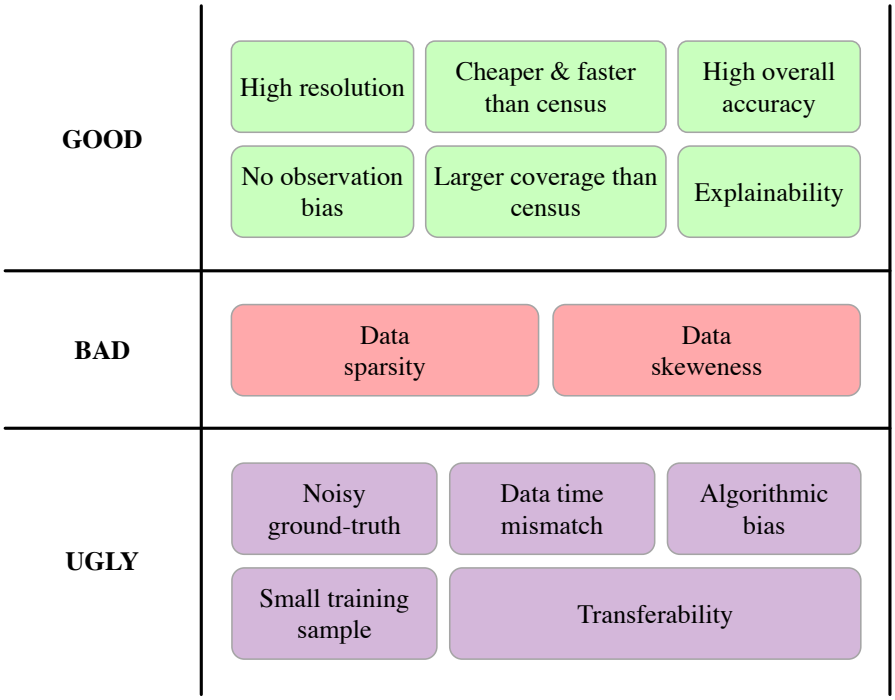

Figure 1: The good, the bad, and the ugly of inferring poverty maps with multimodal data. Machine learning techniques often require a large amount of training data to produce highly accurate generalizable models. However, poverty map inference often relies on survey data which is scarce and not up to date. Moreover, models tend to suffer from the class imbalance problem by over-representing the majority class. By leveraging multiple data sources, we can mitigate these issues and produce scalable and explainable models for inferring high-resolution poverty maps.

inference to inform governments and NGOs to adequately allocate resources, which are often misplaced due to coarse-grained and outdated statistics provided by census and survey data. In this regard, poverty maps are essential tools for authorities and organizations to locate places in need.

Remotely sensed data combined with advanced machine learning methods has provided a recent breakthrough in poverty map inference [R3, R2, R1]. However, these methods are not optimized to produce accountable results that lead to accurate predictions for a broad variety of sub-populations. Here we focus on the challenges and opportunities of predicting high-resolution poverty maps using multimodal data (see Figure 1) and propose three regression models that predict the wealth of populated places from seven independent and freely available data sources obtained through: Google [L3, L4], Facebook [L5, L6, L7], OpenStreetMap [L8], and OpenCelliD [L9].

Our first model is based on a convolutional neural network architecture (CNN) that learns to predict the wealth of places from daylight-satellite images. The second model is an XGBoost regressor (XGB) that learns

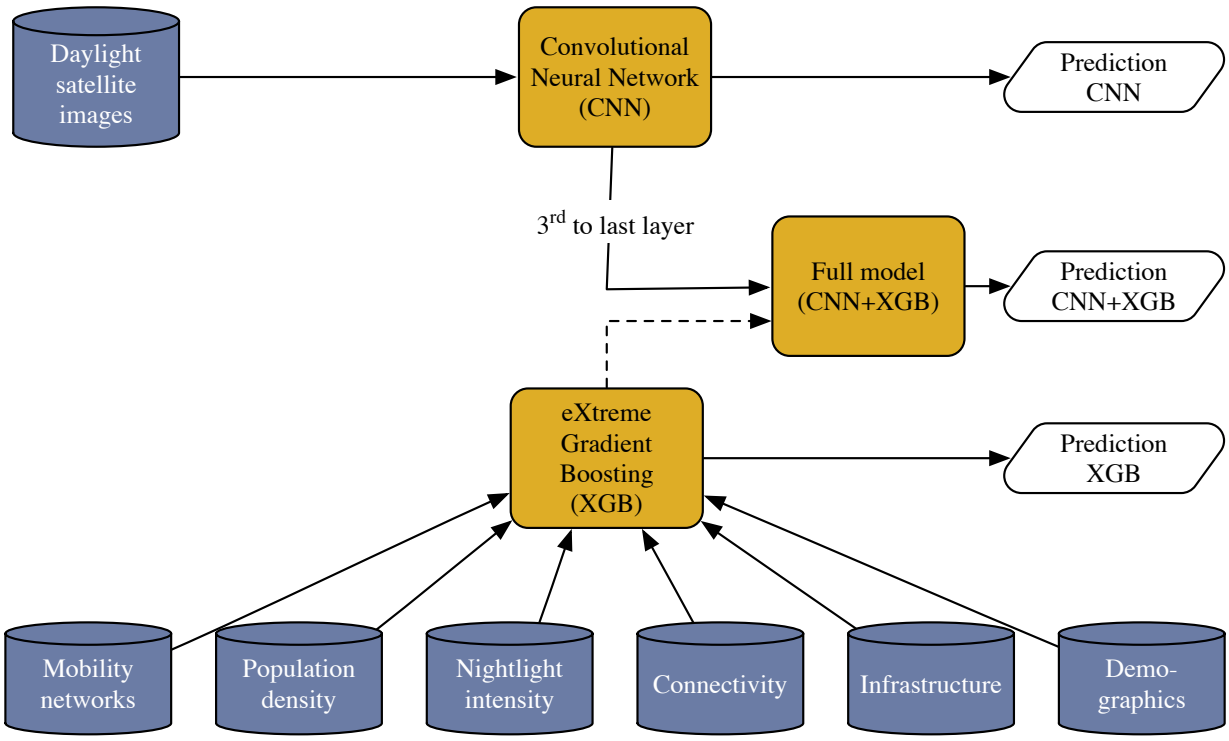

Figure 2: Proposed models. We propose three linear regression models that learn to predict both the mean wealth of places and their respective standard deviation. The former represents the average wealth of households in a given area, and the latter indicates by how much the wealth of those households deviates from each other. The CNN model predicts the wealth of places from daylight satellite images using a convolutional neural network. The XGB model infers the wealth of places using metadata features, e.g., population and mobility. The CNN+XGB model is an extension of the XGB model where in addition to the metadata features, the third to last layer of the CNN model is included as a low-dimension feature vector representation of the satellite images.

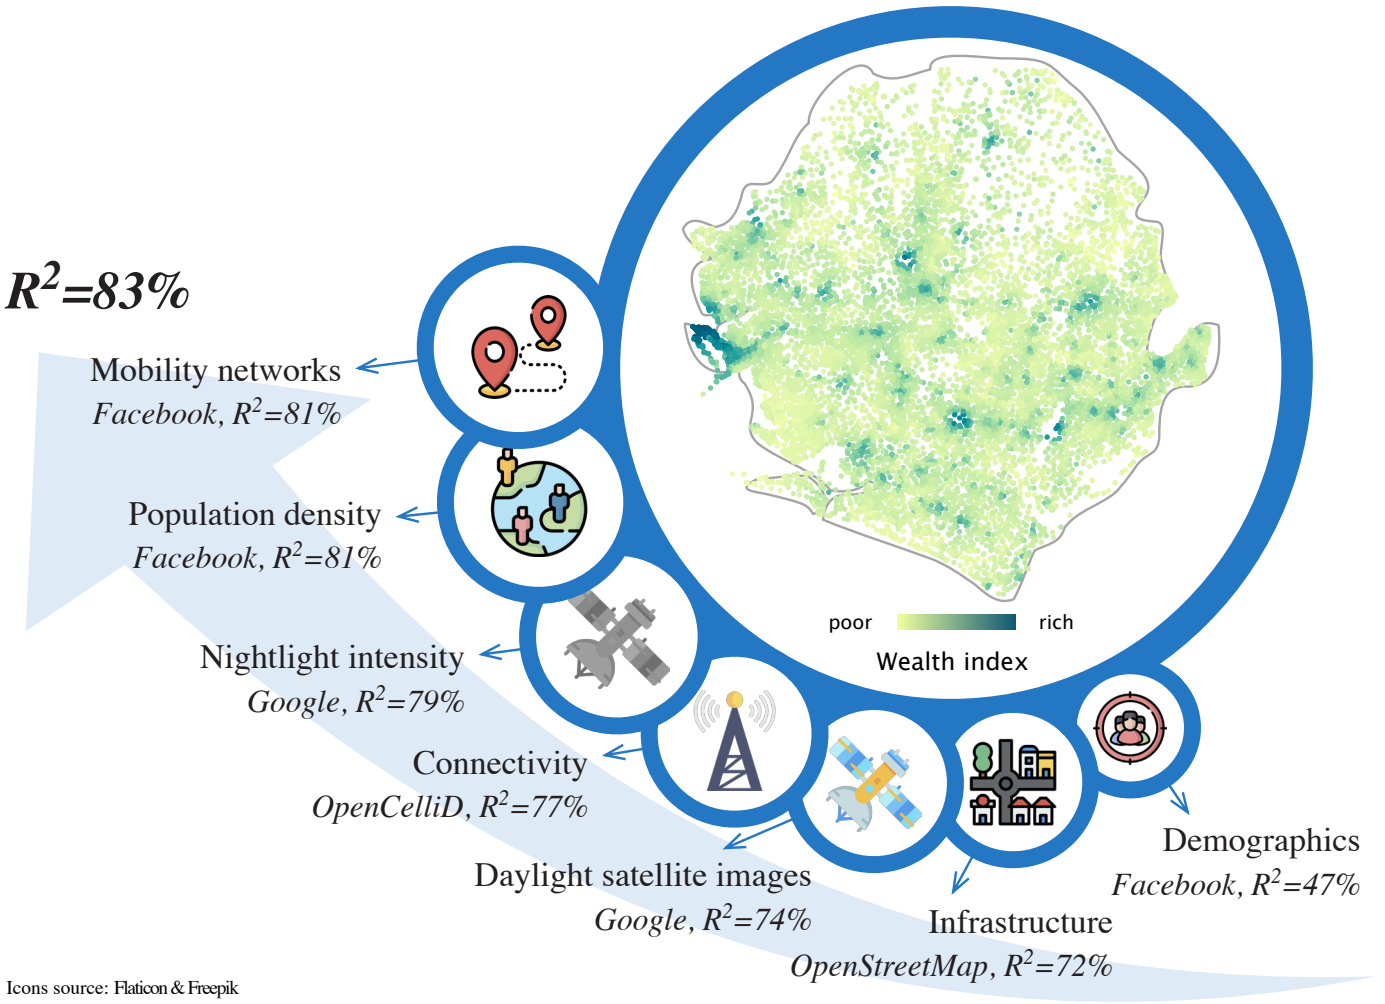

Figure 3: Inferred high-resolution poverty map in Sierra Leone, 2022. By using our proposed models, we found that mobility and population features are the best predictors of wealth, R2=81%. However, this performance may be improved by combining all features, R2=83%.

to predict wealth from metadata features such as mobility and infrastructure. The third and last model extends the second one by including a feature vector extracted from the third to last layer of the CNN model (CNN+XGB), see Figure 2.

Using our models, we found that mobility and population features are the best predictors of wealth. Moreover, the predictive power of our models improves when all features are used at once, see Figure 3. Finally, we also found that satellite images are the best for predicting the poor, while metadata-features are the best for predicting the rural-poor and the urban-middle classes.

As next steps, we aim to explore the spatial and temporal transferability of multi-source models with missing data sources at the target population. Nevertheless, even these preliminary results shed light on the importance of producing accountable models that guarantee accurate predictions to all sub-populations.

Links:
Sustainable development goals - Goal 1: End poverty in all its forms everywhere. https://www.un.org/sustainabledevelopment/poverty
Decline of Global Extreme Poverty Continues but Has Slowed: World Bank. https://www.worldbank.org/en/news/press-release/2018/09/19/decline-of-global-extreme-poverty-continues-but-has-slowed-world-bank
Google Earth Engine. https://earthengine.google.com/
Google Maps static API. https://developers.google.com/maps/documentation/maps-static/overview
High-resolution population density maps by Meta. https://dataforgood.facebook.com/dfg/tools/high-resolution-population-density-maps
Movement Maps by Meta. https://dataforgood.facebook.com/dfg/tools/movement-maps
Facebook Marketing API. https://developers.facebook.com/docs/marketing-apis/
OpenStreetMap. https://www.openstreetmap.org/
OpenCelliD – The world's largest open database od cell towers. https://opencellid.org/

References:
Guanghua Chi, Han Fang, Sourav Chatterjee, and Joshua E. Blumenstock. "Microestimates of wealth for all low-and middle-income countries." Proceedings of the National Academy of Sciences 119, no. 3 (2022).
Kumar Ayush, Burak Uzkent, Marshall Burke, David Lobell, and Stefano Ermon. "Generating interpretable poverty maps using object detection in satellite images." In Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence, pp. 4410-4416 (2021).
Jacob Levy Abitbol, and Marton Karsai. "Interpretable socioeconomic status inference from aerial imagery through urban patterns." Nature Machine Intelligence 2, no. 11 (2020): 684-692.

A new dataset to assess the long-tail effect of the COVID-19 lockdown on Italians

This dataset can be useful to assess the impact of COVID-19 related lockdowns on Italians' psychological and behavioral aspects. This dataset provides accurate insights on changes in the population's general health status permitting practical suggestions at local, regional, and national level to improve infrastructures and services, thus improving Italians' well-being. Understanding the restrictions' consequences may help to better design strategies to contain the virus diffusion and, at the same time, to ensure a reduced impact on the lives of citizens in terms of well-being, sleep quality, and physical activities.

Alessio Rossi, Department of Computer Science, University of Pisa, Pisa, Italy | alessio.rossi2@gmail.com

Michela Natilli, ISTI-CNR, Italy | michela.natilli@isti.cnr.it

From March 2020 to May 2021, several lockdown periods caused by the COVID-19 pandemic have limited people's usual activities and mobility in Italy, as well as around the world. These unprecedented confinement measures dramatically modified citizens' daily lifestyles and behaviors. However, in summer 2021 and thanks to the vaccination campaign (that significantly prevents serious illness and deaths and reduces the risk of contagion) all Italian regions finally returned to regular behaviors and routines. Due to the reduction in citizens life quality caused by long periods of home restriction, it is plausible that the potential consequences of lockdowns are still persistent within the Italian citizens, and whether people's quality of life, sleep and physical activity-related behaviors have returned to the pre-lockdowns level is still unclear.

To answer this question, researchers from the University of Pisa, National Research Council of Pisa, University of Insubria and University of Milan developed a project aimed at gathering subjective data assessing quality of life, sleep and physical activity-related behaviors. This project includes the collection and curation of a database derived from an online survey shared between November and December 2021. Questionnaires quantifying quality of life, sleep, and physical activity were administered, referring to two different moments: before COVID-19 quarantine (i.e., November 2019) and the time frame in which the survey was filled. As the lockdown periods might have a different effect on each participant, contextual information was also collected, regarding socioeconomic status, living space, employment changes. The dataset is publicly available on the SoBigData Catalogue [L1], and an in-depth description of the dataset is being published in Nature Scientific Data.

This dataset may be useful for investigating the impact of the COVID-19 lockdown on psychological and behavioural aspects in Italian citizens. It will be possible to assess whether the quality of life, sleep, and physical

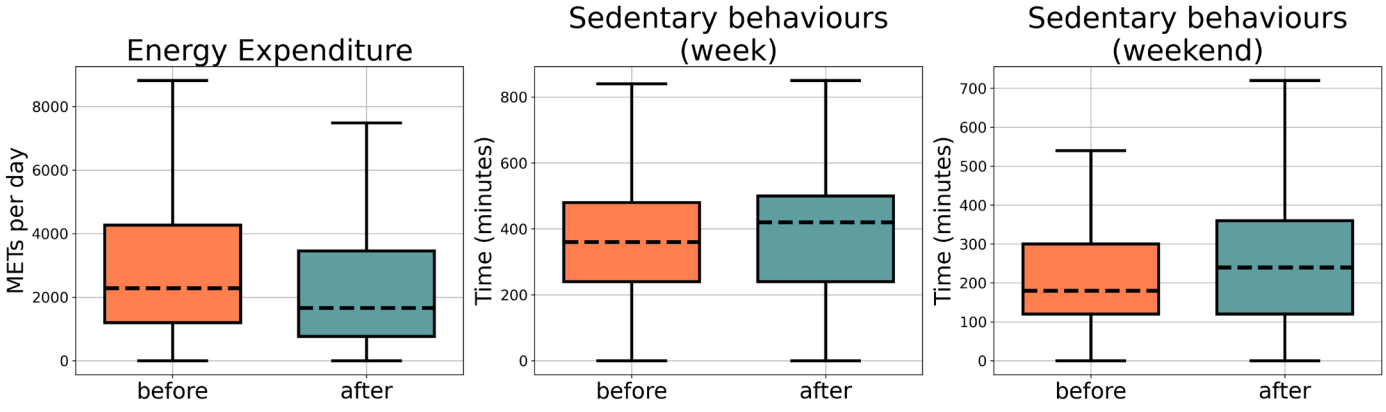

Figure 1. Sankey diagram describing the change in perceived quality of life.

activities may return to pre-isolation levels after a sufficient period without movement restrictions.

The detailed information on each aspect of the respondents' life will permit a complete overview of their habits, perceptions, and well-being status before and after the COVID-19 lockdown. By analysing this dataset, it will be possible to provide practical suggestions at regional, and national level to improve infrastructures and services, thus improving Italians' well-being.

Furthermore, despite the many benefits connected with the vaccination campaign in reducing the transmissibility of the virus, the presence of several COVID-19 variants does not make us safe with respect to possible future actions: understanding the consequences of the restrictions may help to better design strategies to contain virus widespread and, at the same time, to ensure a reduced impact on the lives of citizens. Finally, as this dataset is related to public health, there are many benefits of rendering it open to other researchers' access such as: i) reduce the cost of public research, avoiding the need of launching new studies in this field; ii) replicate/repeat experiments in a trustworthy way, perhaps with somewhat different methods; iii) increase the trust of public opinion toward science, and to encourage 'data altruism' with citizens to facilitate data sharing [R1].

In a preliminary analysis, the researchers found that 45.9% of respondents did not perceive any change in health. However, the health status of 45.2% of our respondents got worse, while only 8.9% of them perceived an increase in their health status. Figure 1 shows that 10.7% of the people changed from a positive health perception to a negative one. In particular, the responses' frequency in excellent and very good health classes were reduced by about 66% and 17% after the lockdown, respectively.

Similarly, poor sleep quality was recorded after returning to regular life from a long time compared to the previous period, when the respondents perceived a better sleep quality. This change in sleep patterns could either affect and be affected by mental and physical health. First, stress disorders induced by movement restriction, and maladaptive coping strategies (e.g., elevated alcohol consumption and time watching television or small screens such as smartphones before sleeping), have been associated with sleep problems [R2]. On the other hand, prolonged periods of irregular sleep routines might affect the sleep-wake circadian rhythmicity, thus causing an impoverishment of the sleep quality inducing an increased risk of mental and physical health disorders.

Finally, the movement restrictions caused by COVID-19 lockdown have limited individuals' possibility to perform physical activity increasing their inactivity tendency and consequently the sedentary behaviours. The respondents of this survey seem to continue having sedentary behaviours also after several months from the end of the first general lockdown when compared to the pre-lockdown period. Indeed, Figure 2 shows that the energy expenditure performing physical tasks per week (Metabolic Equivalent of Task - MET) is still lower compared to the period before the lockdown, despite the return to regular habits. Consequently, the sitting time increased from before to after lockdown period in both week and weekend days.

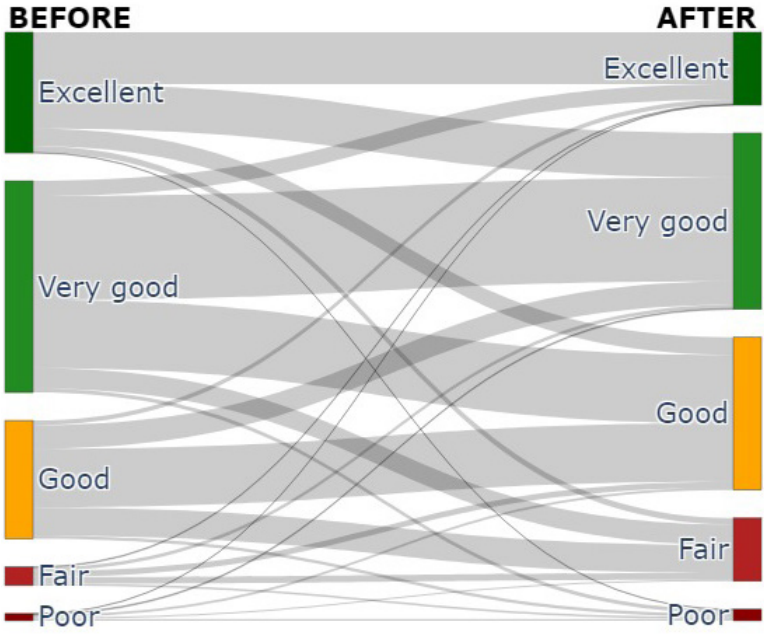

Figure 2. Boxplots describing the differences between the period before and after the lockdown in energy expenditure (METS), and time spent sitting during week and weekend.

Summarising, the dataset could help to deepen the analysis about the effects of lockdown on quality of life, sleep, and physical activity of Italian citizens. The intent of the researchers was to add another piece to the research on the effects of lockdown, which on the one hand has mainly focused on periods of restrictions and on the other hand has not always focused on the perception of individuals.

Links:
[L1]:
https://ckan-sobigdata.d4science.org/dataset/lifequality_lockdown_effect

References:
[R1]: van de Hoven, J. et al. Towards a digital ecosystem of trust: Ethical, legal and societal implications. Opinio Juris In Comp. 131–156 (2021).
[R2]: Robillard, R. et al. Profiles of sleep changes during the covid-19 pandemic: Demographic, behavioural and psychological factors. J sleep research. 30, e13231 (2021).

Detecting Mesoscale Structures by Surprise

In our paper, we devise a unified framework for mesoscale structures detection based upon the score-function called surprise, i.e., a p-value that can be assigned to any given partition of nodes (on undirected/directed, binary/weighted networks). In our paper, we generalize the concept of surprise to tackle the problem of detecting both modular and bimodular structures on weighted networks; besides, we introduce an enhanced version of surprise, capable of simultaneously considering the binary and weighted information of a network to detect mesoscale structures.

Emiliano Marchese, IMT School for Advanced Studies Lucca | emiliano.marchese@imtlucca.it

Tiziano Squartini, IMT School for Advanced Studies Lucca | tiziano.squartini@imtlucca.it

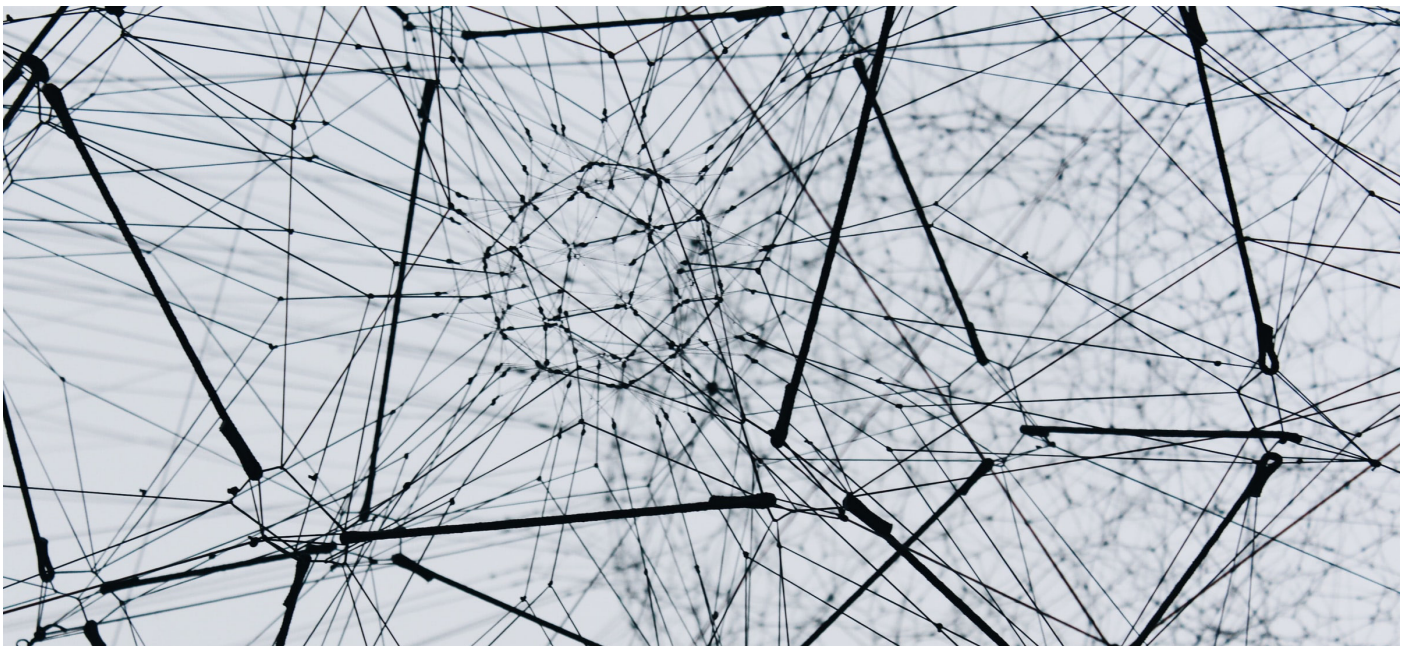

Photo Credit: Alina Grubnyak - Unsplash.com

The importance of identifying the signature of some kind of mesoscopic organization in complex networks (be it due to the presence of communities or bipartite, core-periphery, bow-tie structures) can be hardly overestimated. The best example of complex systems whose behaviour is deeply affected by their mesoscopic structural organization (e.g., resilience to the propagation of shocks, to the failure of nodes, etc.) is provided by financial networks (e.g., the inter-bank network of financial liabilities). So far, much attention has been devoted to the detection of binary mesoscale structures such as communities and (to a far less extent) core-periphery structures.

Generally speaking, the detection of each of the aforementioned structures is carried out by optimizing a score-function that encodes the peculiar character of the pair-wise interactions to be revealed. In our paper, we devise a unified framework for mesoscale structures detection based upon the score-function called surprise, i.e., a p-value that can be assigned to any given partition of nodes (on undirected/directed, binary/weighted networks). Surprise was introduced by Aldecoa et al. for the detection of communities on binary networks and then generalized by Van Lidth de Jeude et al. to handle bimodular structures on the same kind of configurations. In our paper, we generalize the concept of surprise to tackle the problem of detecting both modular and bimodular structures on weighted networks; besides, we introduce an enhanced version of surprise, capable of simultaneously considering the binary and weighted information of a network to detect mesoscale structures.

As a result, our work introduces a statistically grounded framework for detecting mesoscale structures on both binary and weighted networks. To support our findings, we tested surprise-based methods on both synthetic and real-world networks and compared our results with the output of alternative, popular algorithms such as the Louvain and the Clauset's ones – both are modularity based – and Infomap. Our findings show that surprise-based algorithms generally outperform modularity-based ones in detecting mesoscale structures on both binary and weighted networks.

The comparison between surprise-based algorithms and Infomap reveal that the identity of the best-performing algorithm depends on the heterogeneity of the size of the communities characterizing the benchmark planted partitions: a larger heterogeneity seems to favour surprise and vice versa.

This evidence suggests a strategy to carry out mesoscale structure detection on real-world networks, i.e., running different algorithms on the same empirical network, checking the consistency of their outputs and 'combine' them as routinely done in multi-model inference.

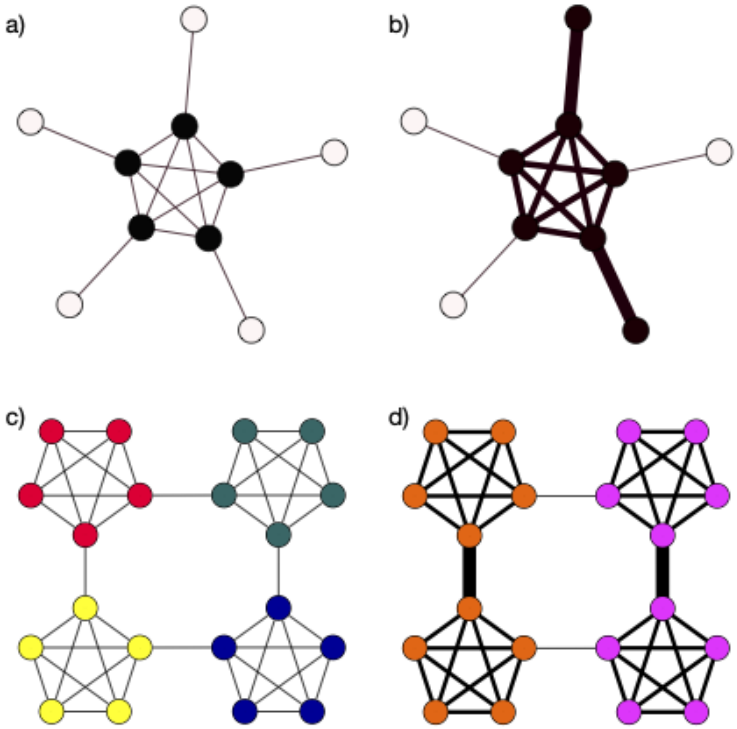

Figure 1: Two examples of how link weights impact on the detection of network mesoscale structures. Top panels: black nodes are detected as the core ones and white nodes are detected as ones belonging to the periphery: if the weight of any two edges linking the core with the periphery is increased, the peripherical ones become part of the core. Bottom panels: similarly, increasing the weight of the edges connecting different communities leads the weighted surprise to detect two modules (Panel d) instead of the four one characterizing the binary case (Panel c).

Research and Development of Quantification Methods with QuaPy

Quantification is an emerging research topic in machine learning aiming at estimating the class prevalence in samples of data items.
QuaPy is an open-source python framework developed at ISTI-CNR for supporting research and development of machine learning methods for quantification.
QuaPy implements many quantification methods, and covers all steps of the quantification pipeline, from data processing to model training, optimization, evaluation, and the analysis of results with the aid of visualization tools.

Andrea Esuli, ISTI-CNR, Pisa, Italy | andrea.esuli@isti.cnr.it

by Alejandro Moreo ISTI-CNR, Pisa, Italy | alejandro.moreo@isti.cnr.it

Photo Credit: Markus Spiske - Unsplash.com

Quantification [R1] (variously called "learning to quantify", "supervised prevalence estimation", or "class prior estimation") is the supervised learning task of training unbiased estimators (or simply "quantifiers") of the relative frequencies (a.k.a. "prevalence values" or "prior probability") of the classes of interest in a sample of unlabeled data items.

For example, in many applications of Sentiment Analysis, the interest often lies in predicting the sentiment distribution in an entire population, rather than in accurately predicting the opinion of every single person. Let us consider, for example, a stream of tweets about a certain debated issue, e.g., the mandate to use face masks in open spaces during the COVID-19 pandemic. Journalists, stakeholders, or the government itself, might certainly be interested in knowing the percentage of people supporting or opposing the measure. Applications like this naturally arise in many fields inherently concerned with aggregated data (in opposition to individual data), such as the social sciences, epidemiology, market research, and ecological modeling, among many others. A characteristic that is common to all these disciplines is that individual decisions are, by and large, unimportant, provided that the ultimate class prevalence estimations are sufficiently accurate.

This problem might, in principle, be cast as a classification task, that is, by training a classifier on labeled data, using it to issue predictions for each of the unlabeled instances in the sample of interest, and then simply counting the number of instances assigned to each class. Such a solution is, however, unsatisfactory whenever the classifier is not a perfect one (as it is indeed expected for any non-trivial classification problem), and particularly so in situations of "distribution shift", that is, in situations where the class prevalence in the deployment set largely differ with respect to the class prevalence observed during training. Needless to say, those cases constitute the most important ones for any discipline targeting the study of distributions over popula-

tions, for the simple fact that those disciplines arise from the need to monitor changing conditions in the population. This naïve approach is called "Classify and Count" in literature and is widely known to deliver biased estimators of class prevalence, i.e., the estimations tend to be close to the distribution observed in the training data (see Figure 1).

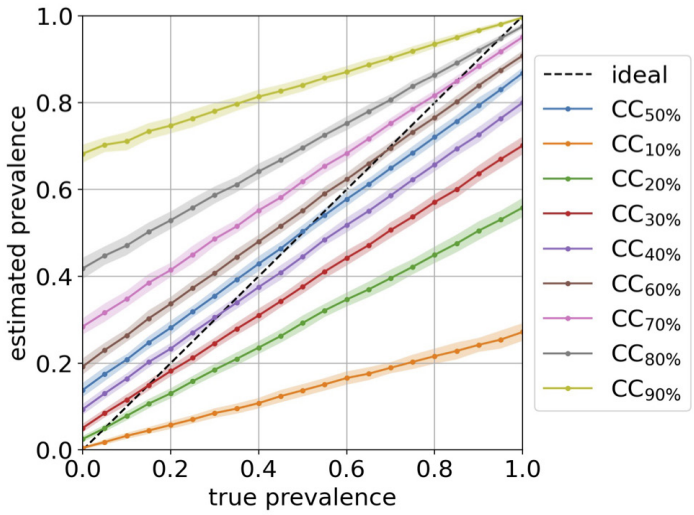

Figure 1: Quantification predictions for a Classify and Count (CC) quantifier on a binary problem. The x-axis reports the true prevalence of the positive class in the test sample, while the y-axis reports the predicted prevalence. Each line represents the performance of a naïve CC when trained on a labeled set with a specific prevalence for the positive class (as indicated in the legend). Note CC tends to deliver biased estimators.

That it is possible to deliver accurate estimations of class prevalence notwithstanding the fact that the classifier is imperfect can be easily shown by the following simple example. Let us assume that we know that every time our classifier provides a positive output, it has a 50% chance of getting it wrong, and, to keep it simple, let us assume that the chance of misclassification for the negative output is 0%. If for an unlabeled sample our classifier predicts a positive label for 50% of its element, we could use our knowledge about the misclassification rates to correct the quantification estimate for the positive label, thus predicting 25% of positives. This is possible since we have characterized the type of systematic error that the classifier exhibits, even though we do not know which instances have been misclassified. Obviously, the error distribution of the classifier is unknown. How to accurately characterize a classifier in terms of the errors it produces, and how to effectively leverage this information in order to provide more accurate estimations are, among other things, examples of typical problems faced in quantification.

Intuitions like this have inspired many of the corrections implemented at the heart of quantification systems, and have given rise to the field of Quantification in 2005 [R3]. Since then, research on quantification has progressed with dedicated learning methods (including methods that do not rely on classification at all), dedicated evaluation measures, and dedicated evaluation protocols. In recent years, quantification research is gaining momentum in the machine learning, information retrieval, and data mining communities, as witnessed by the appearance of the very first dedicated workshops [L2] and competitions [L3]. This growing interest calls attention to the need for specialized frameworks to support and boost research activities on quantification.

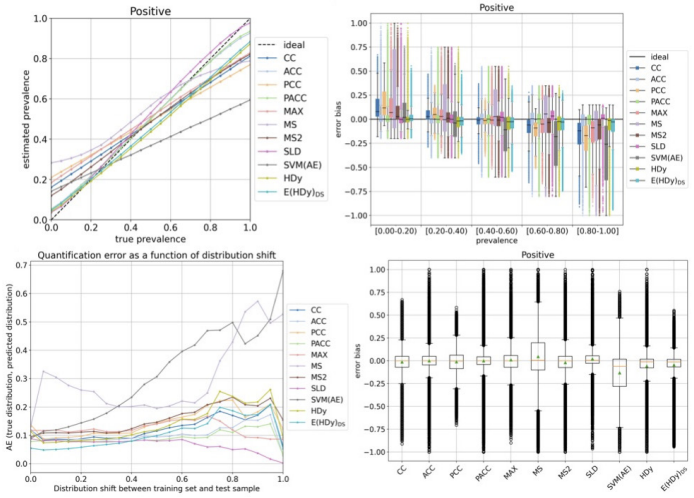

Figure 2: Visualizations produced by QuaPy of results from quantification experiments comparing several quantification methods.

At ISTI-CNR, we have taken a step towards filling this gap by developing QuaPy [R3], a framework entirely written in Python that provides implementations of the most important tools for research, development, and experimentation in quantification. QuaPy implements a rich number of quantification algorithms described in the literature. QuaPy provides abstractions of many central concepts in quantification literature, like the concept of data sample and effective sampling strategies. It supports experimental activity on quantification by providing data loaders for the datasets commonly used in literature, and implementations of the most important metrics and protocols routinely used in the evaluation of quantification systems. QuaPy also provides visualization tools for the analysis of results (see Figure 2). QuaPy is accompanied by rich API documentation and a wiki guide illustrating the most important concepts and features of the framework.

QuaPy is an open-source software distributed under the BSD-3 license; it is available on GitHub [L1] and PyPI (allowing users to quickly install QuaPy using the command "pip install quapy").

Links:
[L1]: https://github.com/HLT-ISTI/QuaPy
[L2]: https://cikmlq2021.github.io/
[L3]: https://lequa2022.github.io/

References:
[R1]: González, P., Castaño, A., Chawla, N. V., & Coz, J. J. D. (2017). A review on quantification learning. ACM Computing Surveys (CSUR), 50(5), 1-40.
[R2]: Forman, G. (2005). Counting Positives Accurately Despite Inaccurate Classification. In: Gama, J., Camacho, R., Brazdil, P.B., Jorge, A.M., Torgo, L. (eds) Machine Learning: ECML 2005. ECML 2005. Lecture Notes in Computer Science, vol 3720. Springer, Berlin, Heidelberg.
[R3]: Moreo, A., Esuli, A., & Sebastiani, F. (2021, October). QuaPy: A Python-Based Framework for Quantification. In Proceedings of the 30th ACM International Conference on Information & Knowledge Management (pp. 4534-4543).

Human-centred Machine Learning with Visual Analytics

The current methods of eXplainable AI (XAI), in fact, do not properly explain behaviour and outputs of machine learning (ML) models to humans. The generated "explanations" do not match human mental models and cognitive capabilities. We call for developing Visual Analytics (VA) approaches that support involvement of human knowledge in both creating and explaining ML models.

Natalia Andrienko and Gennady Andrienko, Fraunhofer Institute IAIS (Germany) and City, University of London (UK) | natalia.andrienko@iais.fraunhofer.de | gennady.andrienko@iais.fraunhofer.de

Photo Credit: Jess Bailey - Unsplash.com

We have recently published a paper setting the research agenda for Visual Analytics to support human-centred ML [R1]. The paper discusses the deficiencies of the current XAI methods and proposes a framework in which VA methods are used to involve human expert knowledge in building of ML models, generation of explanations, and presentation of the explanations to users of the models.

The research field of XAI emerged in response to the pressing need to explain untransparent ML models ("black boxes") to users. The work in this field was boosted by the European Parliament's adoption of the General Data Protection Regulation (GDPR), which introduces the right of individuals to receive explanations of automatically made decisions relevant to them. However, some ML researchers argue that the current XAI approaches fail to provide satisfactory explanations that can be well understood by humans, i.e., linked to their mental models. Consider, for example, an explanation given in the form of a logical rule: "If (x1 > 5) and (x2 < 7) and (x3 > 10) then x belongs to class 1". It may be quite accurate in classifying data instances, and a domain expert can understand what it says if attributes x1 to x3 are meaningful in the domain where the data are taken from. However, the domain expert can say that, despite its high empirical confirmation, it is not clear why this model should work. The model is not explained in the terms of the domain knowledge such as causal relations known in the domain. This problem is referred to as conceptual mismatch between ML models and human mental models. Another problem is that a model interpretable in theory may be incomprehensible in practice due to its size and complexity. Consider, for example, a decision tree containing hundreds of nodes. A human can trace and understand any small part of it, but the whole tree is beyond the human capabilities for tracing and understanding.

To address the first problem, human expert knowledge, including domain concepts and causal relationships, needs to be incorporated in ML models and used in generating explanations. To address the second problem, XAI-generated explanations need to be organised and presented at an appropriate level of abstraction.

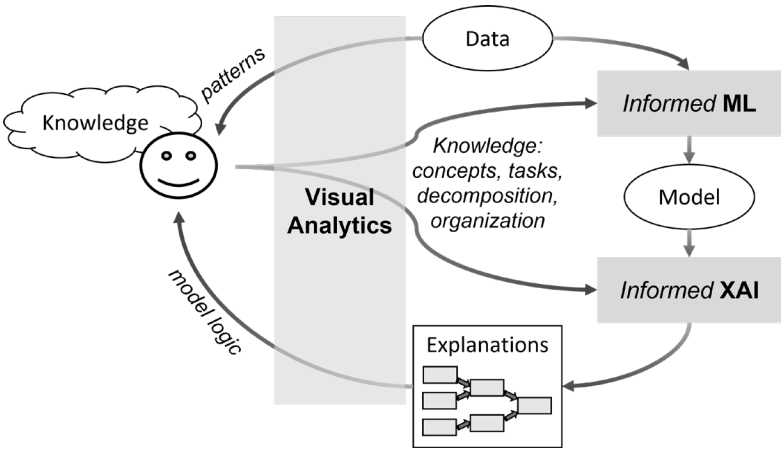

Figure 1. Schematic representation of the research framework for human-centred ML supported by VA. Models are developed in tight interaction of ML algorithms with humans. VA supports the transfer of human knowledge to the algorithms. The knowledge is involved not only in model building but also in generating explanations of the models. Such explanations are produced by XAI methods labelled "Informed XAI", which are yet to be developed. VA techniques are used to present the explanations to users in a comprehensible form at a suitable level of abstraction.

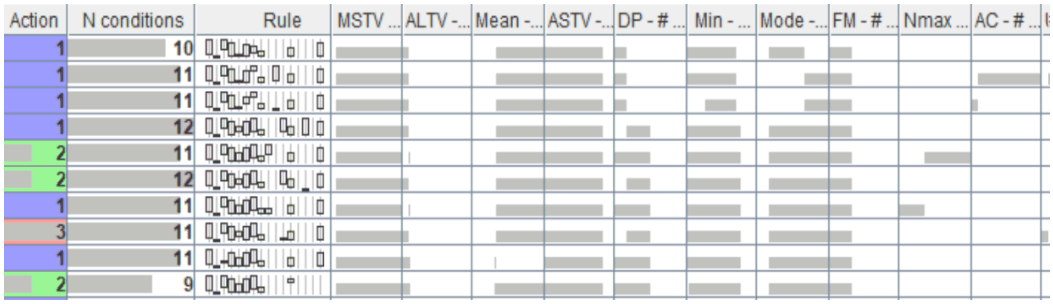

Figure 2. A fragment of a table representing a subset of rules. Each row represents a rule. The columns, starting from column 4, correspond to attributes involved in the rules. The grey bars in the cells represent intervals of attribute values appearing in rule condition. Thus, a condition a ≤ x ≤ b, where x is an attribute and a and b are numbers, is represented in the column corresponding to x by a bar whose left end corresponds to a and right end to b while the cell width represents the whole range of the possible attribute values from minimal to maximal. The same conditions are also represented in column 3 by vertical bars. Each vertical axis corresponds to one of the attributes.

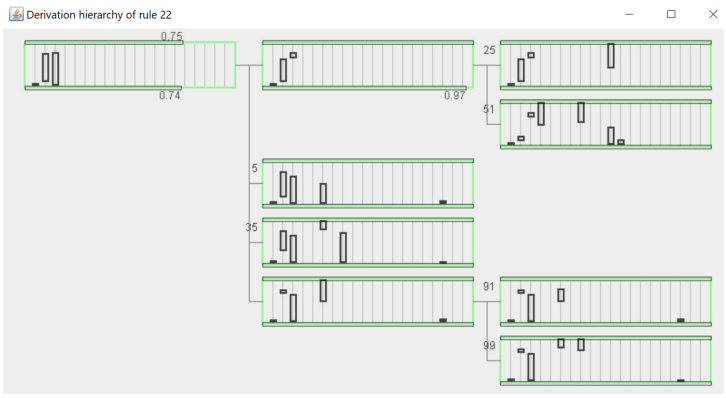

Figure 3. Visualisation of a hierarchically organised group of compatible rules. Each rule is represented by a rectangular glyph in the same way as in column 3 of the table shown in Figure 2. The leaf nodes of the tree represent original rules. The other nodes represent rules derived by merging and generalising two or more rules with a common result and similar conditions.

VA is a natural partner of ML and AI in the research both on involving users in ML processes and on explaining ML to users. Combining human and machine intelligence is the central idea of VA. Complementing the current efforts in XAI, VA can contribute by exploiting the potential of visualization as an effective way of communicating information to humans and a strong trigger of human abstractive perception and thinking. We propose a cross-disciplinary research framework that is schematically represented in Figure 1 and explained in the figure caption.

As an example of abstraction-based simplification of model explanations, we demonstrate how model explanations in the form of an ensemble of decision rules can be lifted to a higher level of organisation and abstraction through hierarchical grouping and generalisation. Figure 2 gives an example of visual representation of rules in a table form. Figure 3 demonstrates a hierarchy of rules, where higher-level rules have been obtained by merging and generalising two or more lower-level rules with a common result (i.e., prediction or decision) and similar conditions.

With this paper, we aim to motivate and trigger research on bridging gaps between machine learning and human mental models using a synergy of approaches from machine learning, artificial intelligence, and visual analytics. We expect it to result in methods and systems for building computer models reinforced by the power of human intelligence and readily accepted by humans as extensions of their mental models and enhancers of their reasoning.

References:
[R1]: N. Andrienko, G. Andrienko, L. Adilova and S. Wrobel, "Visual Analytics for Human-Centered Machine Learning," in IEEE Computer Graphics and Applications, vol. 42, no. 1, pp. 123-133, 1 Jan.-Feb. 2022, doi: 10.1109/MCG.2021.3130314.

A TNA experience on Ordinal Quantification

What follows is the account of a Transnational Access experience by Mirko Bunse, who visited ISTI-CNR Pisa under the SoBigData++ TNA program. The author describes his experience in sharing work with fellow researchers interested in ordinal quantification.

Mirko Bunse, TU Dortmund University | mirko.bunse@cs.tu-dortmund.de

While a typical objective in supervised machine learning is to predict a label for each individual instance, quantification aims instead at predicting the label distribution within a set of instances. It has been widely acknowledged that this problem can't be solved by simply classifying and counting the individual instances in the set; quite to the contrary, quantification deserves consideration as a learning task in its own right. Applications of quantification range from the social sciences over technical support to Astro-particle physics. Just recently, the topic gained momentum from the first international workshop on quantification, which was co-organized by Alejandro Moreo and Fabrizio Sebastiani from ISTI-CNR in Pisa, Italy.

I met for the first time Alejandro and Fabrizio during this workshop. We realized that the three of us had been working on ordinal quantification (OQ) – a quantification task where the classes are totally ordered, and misclassifications are weighted by the distance between classes. For instance, if we need to predict the label distribution of sentiment classes (negative, neutral, positive), we are facing an OQ problem: a misprediction between the class prevalence of negative and positive induces a larger error than a misprediction between negative and neutral or between neutral and positive.

Meeting other researchers who work on OQ is not very common. In fact, the largest part of the existing quantification literature deals with non-ordinal settings, despite the importance of ordinality in many applications. For several years, I myself was not even aware of OQ as a task in its own right. Before I came across quantification literature, I was eagerly working on OQ methods that have been proposed within physics research, under the name "unfolding". Indeed, physicists have devised a large collection of OQ methods, which, however, kept disconnected from quantification literature due to the interdisciplinary gap between computer science and physics. This gap manifests not only in different names for the OQ problem, but also in different notational conventions and different focuses.

Alejandro, Fabrizio, and I wanted to bridge this gap. We wanted to find the commonalities and differences between OQ methods from quantification literature and physics literature. We wanted to find their strengths and weaknesses. So, I applied for a TNA visit at ISTI-CNR. During my stay, we all worked intensively on the topic. After one month, our manuscript was ready for submission. And we are quite happy with the results.

First, we have created two datasets for OQ research that overcome the inadequacies of the previously available ones. Second, we have experimentally compared the most important OQ algorithms proposed in quantification literature and in physics literature. Third, we have proposed three novel OQ algorithms, which are based on the idea of preventing ordinally implausible estimates through regularization, an idea that originates within physics research. In a nutshell, we have indeed bridged the gap between quantification literature and physics literature.

My most vivid memories, however, are all the activities that we undertook after working hours. Meeting the welcoming people who work at ISTI-CNR was an incredible pleasure. The city of Pisa turned out to be not only a beautiful city with lots of cafés, restaurants, and pubs (I am a huge fan of Italian cuisine), but also a beautiful place for outdoor activities (I am also a huge fan of outdoor activities). Monte Pisano, a small mountain chain next to Pisa, is easily reachable by bike or even by foot. These mountains offer several possibilities for hiking, rock climbing, and cycling. I am eternally grateful to SoBigData++ for these memories, and for the honour of working with some of the most renowned quantification researchers. We are planning to continue our collaboration on quantification in the near future, to tackle several exciting issues that we have identified.

Dynamics of opinion polarization in social networks

Why do we witness opinion polarization? This recent research has proved that the Friedkin-Johnsen opinion formation process can explain polarization and has determined the initial opinions that will produce the most polarizing effects.

Elisabetta Biondi, IIT-CNR | elisabetta.biondi@iit.cnr.it

For decades, researchers from different fields have been trying to understand how people form their opinions. With the rise of social media platforms, this quest has become even more significant, especially due to the emergence of some alarming and extreme phenomena, such as the polarization of opinions, online hating, etc.

Understanding how people form their opinions means designing a model that can reproduce the dynamics of opinion formation. To this aim, many models have been proposed, and only a few of them have been already validated in real social networks. One of the most popular is the Friedkin-Johnsen (FJ) model, whereby every person in a network forms his/her opinion depending on an average (weighted on his/her susceptibility attitude) between his/her initial opinion (prejudice) and the opinion of his/her neighbours. Thus, two factors play in it: the anchoring to the original prejudice and the susceptibility towards the neighbourhood's opinion, which are weighted according to the social weights among nodes. Both mechanisms have sociological foundations: the anchoring reproduces the personal capacity or willingness of each person to depart from the prejudice; the susceptibility models the fact that people change their opinions to reduce the discomfort caused by disagreement with others.

A key question to test the validity of the Friedkin-Johnsen model is the following: does this model capture polarization? Is it possible to identify opinions that produce polarization?

In the literature, however, there are different variants of the model (see Figure 1). In the most general form of the FJ model (which we denote as gFJ), each node iteratively updates its opinion balancing the anchoring to its personal prejudice and susceptibility toward the opinions of its neighbours. The balance between the two is controlled by a parameter, called susceptibility i, namely the vulnerability of node i to others' opinion influence. The non-anchoring part considers the opinions of the neighbours and the opinion of the node itself at the step before, and these quantities are weighted by the strengths of its social ties (collected into the matrix W). The variational Friedkin-Johnsen (vFJ) is more restrictive, because when nodes update their opinion they do not take into account their current opinion. Apart from this fact, gFJ and vFJ are very similar. Finally, the restricted FJ (rFJ) is derived from vFJ but here there isn't any way to control the anchoring part of the model, which is assumed to be assigned in the same way for all the nodes of the network.

Table 1. The FJ family of models

gFJ:	$z_i(k+1) = (1-\lambda_i)s_i + \lambda_i \sum_{j \in \{i\} \cup N(i)} w_{ij} z_j(k)$
vFJ:	$z_i(k+1) = \frac{\hat{w}_{ii} s_i + \sum_{j \in N(i)} \hat{w}_{ij} z_j(k)}{\hat{w}_{ii} + \sum_{j \in N(i)} \hat{w}_{ij}}$
rFJ:	$z_i(k+1) = \frac{s_i + \sum_{j \in N(i)} \hat{w}_{ij} z_j(k)}{1 + \sum_{j \in N(i)} \hat{w}_{ij}}$

Figure 1: The FJ Model Family

Each model is said to be polarizing when the polarization in the network increases during the opinion formation process. How is the polarization measured? The metrics of polarization can be classified into four classes. Metrics in the local class measure the total disagreement in the network as a whole, summing up the disagreement perceived by each node weighted by the social ties. The other classes comprise instead global measures and do not consider the social relationships among nodes. The dispersion class considers polarization as the quadratic deviation of opinions at equilibrium from the average. The other two classes consider a slightly different semantic since they distinguish between "neutral" and "extreme" opinions. Since opinions span over [-1,1], the neutral opinion is the opinion 0, whereas the extreme ones are -1 and 1. Thus, the absolute class measures the polarization as the quadratic deviation from the neutrality. The total class measures the polarization as the absolute linear deviation from the neutrality. Besides polarization metrics, we also consider the concept of choice shift. A choice shift occurs

when the sum (let's say the average) of the initial opinion is different from that of the final opinion.

The research question we have tackled is whether, given an underlying social network, the different FJ versions lead to polarization for any of the classes described above. We have proved important results. First, all the variants of the model are locally depolarizing by design. Furthermore, the gFJ and vFJ are globally polarizing in general (if some non-restrictive conditions are met) and this proves that the Friedkin-Johnsen model is realistic, since it is able to reproduce polarization. Next, and even more important, we have identified how to obtain polarizing vectors, i.e., the initial opinions that lead to opinion polarization: we identified three different local maxima and the global maximum for the absolute class of polarization and a local and a global maximum for the total class of polarization. Finally, we have proved that the weaker version of the Friedkin-Johnsen model is also globally depolarizing, and it does not induce a choice shift in the network. This means that this variant of the model is not realistic.

We have tested our results on two real-life social network graphs: the Karate Club graph and a Facebook subgraph. We needed to design a way to assign the susceptibility of each node (which does not come from the network itself) and, since these networks have both very few central nodes, we decided to correlate the susceptibility value according to their centrality, and to compare this configuration with one where all nodes have the same susceptibility. Thus, for the susceptibility, we evaluated three configurations: the individual susceptibility of a node is (i) proportional to its Pagerank centrality Pi, (ii) inversely proportional, and (iii) all susceptibilities are set to 0.8.

In Figure 2 we can see the shape of the polarizing vectors for the absolute class of polarization. Each node is represented by an arrow, whose starting point is the initial prejudice and ending point is the final opinion The three rows represent the three configurations of susceptibility: directly proportional to Pi, inversely proportional, and equal to 0.8. The first three columns are the local maxima, the fourth one is the global maximum.

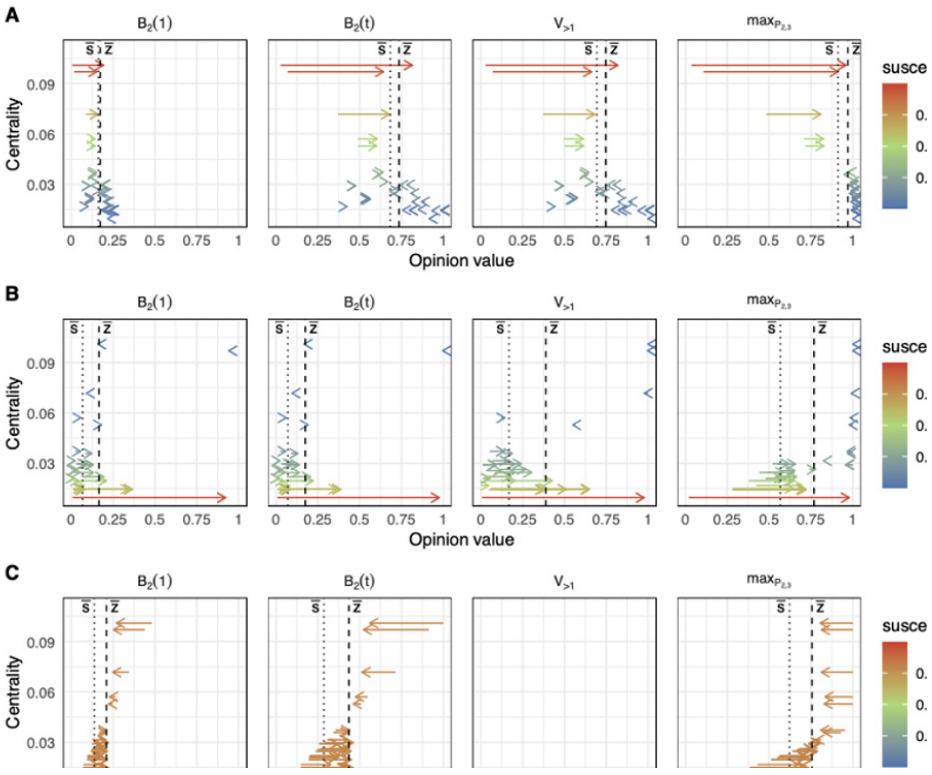

Figure 2: gFJ on the Karate network: polarizing opinions vs nodes centrality. Each arrow corresponds to one node in the Karate graph. An arrow starts at the initial prejudice and ends at the final opinion. The colour of the arrow corresponds to the susceptibility assigned to the node. (A) susc. ∞ P_i (B) susc. ∞ P_i^{-1} (C) susc. = 0.8. With the dotted and dashed line, we denote the average initial and final opinion, respectively.

We can see that the global solution is the one that features the maximum number of nodes with an extreme initial opinion (equal to 1): in this way, the nodes with more extreme opinions work synergistically to push the others' opinions closer to theirs.

For what concerns the centrality, instead, we can see that when the susceptibility is set according to the centrality of nodes, the nodes with larger opinions are those more stubborn (so that they can influence more the neutral nodes toward their opinion). In the third row, where the susceptibility of all nodes is the same (panel C of Figure 2), the polarizing vectors that maximise polarization (and so that are able to influence more the opinions of nodes) are those in which the most central nodes have initial opinions extreme. This confirms that the PageRank centrality is able to capture the ability of nodes to convince the others, and thus it identifies the most influential nodes.

In conclusion, in this work, we have done a comprehensive revision of all the variants of the FJ model and have proven that they are not equivalent. Our findings show that the simplified version of the model, despite being commonly used in the related literature, does not reproduce polarization, while under the original model proposed by Friedkin and Johnsen, polarization can emerge. This is fundamental for an opinion dynamic model to be realistic. Furthermore, we provide the conditions under which polarization happens as well as discuss how to find the initial opinions leading to polarization. We have finally seen that both susceptibility and centrality of nodes play a role in polarization.

Challenge Us 2021: SoBigData meets Industry

The Challenge Us program provided opportunities to companies interested in harvesting their own data they bring, thus entering the "world of Big Data" and exploit its potential. The Challenge Us was designed to build a bridge between industry and academia, offering the free support of SoBigData++ scientists to design solutions and produce proof of concepts.

Nicola Del Sarto, Scuola Superiore Sant'Anna, Pisa, Italy | nicola.delsarto@santannapisa.it

Photo Credit: Joyce McCown - Unsplash.com

The Challenge Us program 2021 selected four companies. The selected companies have been able to benefit from free support of a team of data scientists from the SoBigData++ and have been able to use the methodologies and software tools available on SoBigData++ big-data analysis platform.

The 22nd of November Scuola Superiore Sant'Anna hosted the Challenge Us program final event. During the event, the four selected firms presented the results of their collaboration with SoBigData partners.

At the final event, each company displayed the collaboration's results and how they tried to solve the problem they submitted to the challenge. The companies selected where four:

A2A is actively involved in innovation and research, specifically in the area of energy. A2A S.p.A. is an Italian company that generates, distributes, and markets renewable energy, electricity, gas, integrated water supply, and waste management services. The company asked for the support of Scuola Normale Superiore to handle their big data in order to improve their price system.

BIZPORTAL was founded in Bulgaria in 2014 as a global open data collection and analytics solutions provider. BizPortal's flagship product is TenderAlpha.com – a global government contract database product consisting of aggregated government procurement contract awards and related company information from more than 50 countries worldwide with 10+ years of historical records and various ongoing delivery options. Bizportal asked for the help of SoBigData to explore the possibility to use their data to improve internal processes and operations.

GIUSTA is a food delivery company promoting a new delivery model, which aims to be sustainable for restaurants, ethical for riders and guaranteed for consumers. Giusta requested the support of SoBigData to create a database by connecting different data sources together and to improve forecasts regarding the trend in demand for their products.

OMNITECH S.r.l. is a company active in technology, research, and innovation, engaged in the design of innovative platforms with high technological value, as well as in the production, installation, assistance, maintenance and marketing of software and hardware products. The company has asked SoBigData++ for support to be able to combine data from different sources and be able to make more accurate climate predictions.

Each company had 15 minutes to show the results of the project and indicate possible future developments. At the end of the presentation, an interesting Q&A session took place which made it possible to produce interesting ideas for the various projects. The event was held in the main hall of the Sant'Anna School of Advanced Studies and was attended by more or less twenty people. At the end of the event, the problems relating to the organization of the event were underlined, especially those related to the Covid-19 pandemic and the 2022 edition launched, which will be organized by Tartu University.

Links:
http://www.sobigdata.eu/challenge-us-2021

Highlights from the Event "Towards a Digital Ecosystem of Trust: Ethical, Legal and Societal Implications"

On the 9th of March 2022, Re-Imagine Europa and SoBigData++ organised an event to present the white paper "Towards a Digital Ecosystem of Trust: Ethical, Legal and Societal Implications" authored by Jeroen van de Hoven, Giovanni Comandè, Salvatore Ruggieri, Josep Domingo-Ferrer, Francesca Musiani, Fosca Giannotti, Francesca Pratesi and Marc Stauch as part of the SoBigData++ project. The event saw the participation of over eighty key representatives from academia, policy, industry, CSOs, NGOs, and other stakeholders.

Matteo Nicolosi, ReImagine Europe | matteo.nicolosi@reimagine-europa.eu

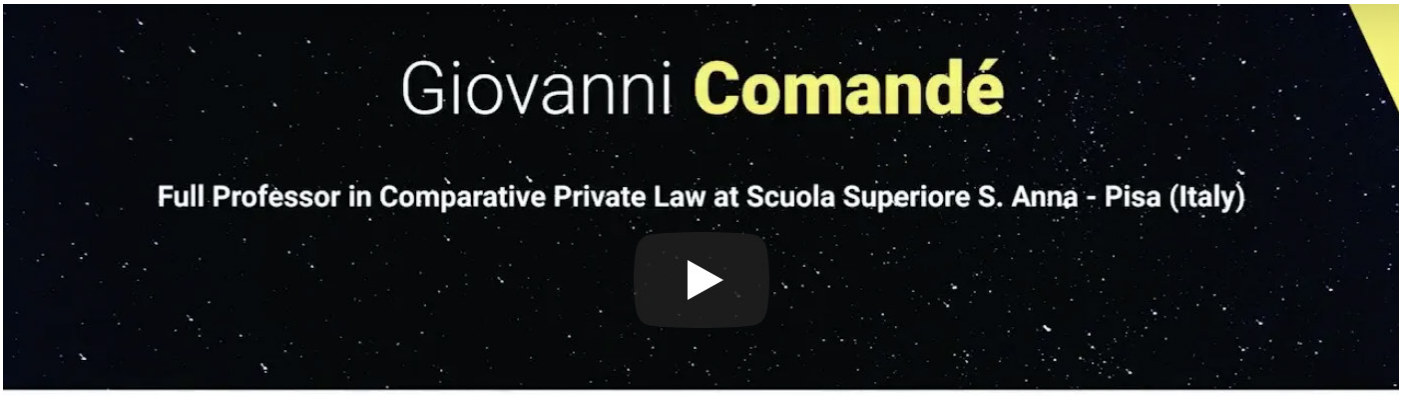

The European vision of a digital ecosystem of trust that the white paper outlines is based on innovative technological solutions, a comprehensive regulatory framework, and respect for the core values and principles of ethics that are shared by the Union's Member States. SoBigData++ is the largest pan-European social science data research project. Establishing a framework where data subjects can feel safe sharing their data is crucial to facilitating the development of data science in Europe. The white paper's aim is precisely to discuss the methods for promoting data sharing, building privacy-preserving technologies, and fostering data decentralisation and data altruism to develop a model for responsible digital innovation in Europe and beyond.

Re-Imagine Europa's Chief Executive Erika Widegren opened the meeting by underlining the timeliness of a discussion on a digital ecosystem of trust within the context of the broader geopolitical challenges. The floor was then passed to Brando Benifei, Member of the European Parliament and Rapporteur for the Special Committee on Artificial Intelligence in a Digital Age (AIDA).

Mr Benifei reiterated that a healthy digital environment has a substantial influence on how the EU is positioning itself in such a difficult geopolitical situation as it is an increasingly important part of the battlefield on which the struggle between democracies and authoritarian regimes is taking place.

He then explained that all the regulations recently issued by the European institutions, from the Data Services Act to the Data Market Act and the AI Act currently in preparation, are conceived to build a concrete frame for a Digital Ecosystem based on safety, transparency, and trust. The European model EU institutions are working to develop will incorporate the shared values of respect for fundamental rights. It shall be based on high levels of protection for data and the principle of human oversight over AI and digital processes. At the same time, a certain level of flexibility

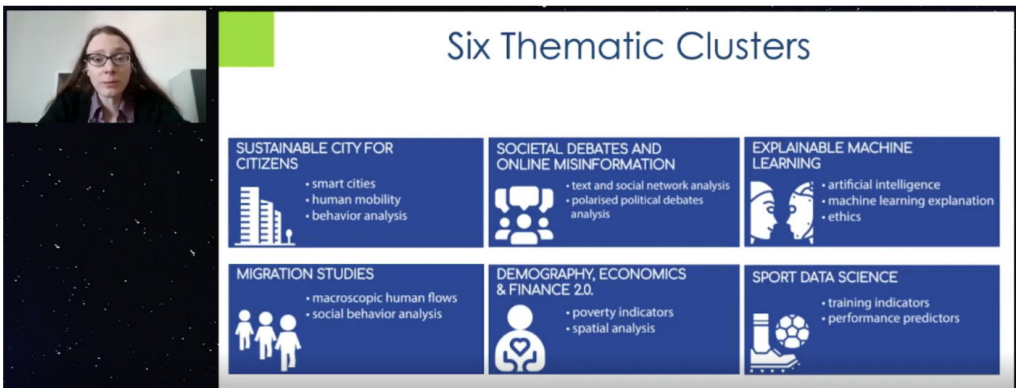

Figure 1 – The SoBigData ++ six thematic clusters

should be allowed to maintain a dialogue with systems based on different approaches.

The floor then passed to the event's Chair and moderator, Luca De Biase, Founder and Chief Editor at Nòva24 – Il Sole 24 Ore and Media Research Director at Re-Imagine Europa. After a few remarks about the fact that distrust is often the rule in the current European digital environment, and the effort to rebuild trust is thus critical, he introduced the participants to the first panel.

The first to take the floor was Giovanni Comandé, Full Professor in Comparative Private Law at Scuola Superiore S. Anna – Pisa (Italy). He briefly presented the topics included in the "Towards a Digital Ecosystem of Trust: Ethical, Legal and Societal Implications" white paper, which he authored along with the other participants of the first panel. One of the most relevant points of his intervention was that the Data Act's text clarified the rules of the GDPR in terms of data portability, addressing the previous legal uncertainty and effectively enabling innovation. He also explained that the same uncertainty is still unaddressed on data standards and interoperability. That issue will have to be solved to achieve a legal framework that fosters a digital ecosystem of trust.

The floor then passed to Josep Domingo-Ferrer, Professor of Computer Science at Universitat Rovira i Virgili – Tarragona (Catalonia). Taking from Professor Comandé's intervention, he stated that we need to balance the need for digital innovation and the protection of our ethical values. In light of this, technologies conceived to provide privacy-by-design could prove very useful. Professor Domingo-Ferrer illustrated a few examples: decentralised anonymisation of data, which exploits the computational capacity of portable devices to anonymise data with a higher security level and incentivised participation by offering services in exchange for data sharing by users.

Francesca Pratesi, Research Member of the Information Science and Technology Institute of the Italian National Research Council (ISTI-CNR) – Pisa (Italy) was next on the floor. She described the SoBigData++ Research Infrastructure and its features more in detail, presenting the six thematic clusters developed and offering a few examples of data micro-aggregation and other ethical and privacy-by-design characteristics.

After a quick one-minute break, the second panel started. Yvo Volman, Head of the Data Policy and Innovation Unit of DG Connect (European Commission), spoke first. Mr Volman said that trust is a critical prerequisite for people willing to make their data publicly available to the wider public. Therefore, clearer rules are needed to define who and in which circumstances can use the various pools of data shared and guarantee that the rule of law is respected in the digital world as they are in the physical world. At the same time, these rules should also ensure that companies are accountable for the innovations they introduce into the market since it would be irresponsible to propose innovations without carefully considering their effects on our society and economy. Europe is making an effort to emphasise the idea of the autonomous citizen with fundamental rights as the central idea of its digital development model.

The floor then passed to Jeroen van den Hoven, Leader of SoBigData++ Ethics Board, Full Professor of Ethics at the Technical University Delft, and Chair of the European Group of Ethics (EGE). Professor van den Hoven started his intervention by making an overview of the various approaches to digital technologies in the different parts of the world. He gave an overview of the most remarkable: the US model, which tends to privilege the big corporate interests over individual rights, the Chinese state-controlled surveillance economy style, and the Russian model, whose features are still unclear. Professor van den Hoven then explained that the next decade would show if this idea of a European "new way" to digital ethics is feasible to defend the liberal-democratic system that is part of our European lifestyle. Furthermore, it will be decisive in defining the fate of Europe: will Europe be the museum of the world or the cradle of a new digital renaissance?

The third panel started with the intervention of Gry Hasselbalch, Senior Key AI Ethics Expert and Research Lead at INTouchAI.eu. Dr Hasselbalch stated that there are different paths to creating data infrastructures that are now competing at the global and local levels in the current development phase of digital technologies. The main stakeholders of this process are the citizens, companies, and institutions supposed to regulate it. It is thus important to consider that while developing a European way to ethics about AI and data, we are also distributing power "by design". Therefore, policymakers should carefully reflect on the kind of society we want to foster for our future. Which type of institution are we empowering? We need

Figure 2 – The major threats to a European Responsible Digital Future

to think carefully about which kind of society we are creating. We also need more dissemination and open conversations since European citizens must have a say about the role of data technology in our communities.

The floor then passed to Maria Savona, Professor of Applied Economics at the Department of Economics at LUISS University, Rome. Professor Savona started her intervention by expressing her idea that, especially when it implies creating value from a social perspective, we should protect data property rights and create other incentives for individuals to share their data. She then presented a few essential questions about how data regulation should be conceived: By making privacy central, are we giving up other fundamental rights? What kind of data architecture can facilitate data collection and stimulate data sharing while protecting individual and collective rights? Among the objectives of the European Commission's Data Governance Act, issued in late 2021, was creating an EU market for data intermediaries, which is, in principle, an excellent idea. But, as recently stated also by the European Data Protection Supervisor, this is also opening new questions: Who is going to guarantee that those data intermediaries will act transparently? Who will ensure they are accountable for how the data shared by citizens is used?

Fosca Giannotti, Professor of Computer Science at Scuola Normale Superiore and Researcher at the Information Science and Technology Institute "A. Faedo" of the National Research Council – Pisa (Italy), opened the fourth panel. Professor Giannotti began her intervention by highlighting the need to support human oversight on algorithms and AI-supported processes to improve citizens' trust in technology. She then introduced the XAI project, whose main goals are empowering individuals over automatised processes and implementing explainable algorithms and human oversight in the field of AI and digital technologies. After a brief description of the two different methods available to provide explainability to AI algorithms and similar processes, professor Giannotti stated that allowing interaction between AI mechanisms and humans could be an excellent methodology to enhance the citizens' trust in such technologies.

Virginia Dignum, Professor of Responsible Artificial Intelligence at Umeå University, Sweden, was next on the floor. She stated her concerns about how AI systems are conceptualised and how the data economy transforms our societies. AI systems are currently regarded as having agency, while they should be considered an ecosystem of different technological, societal, and human elements. Moreover, the data economy pushes our society towards "datification", the tendency to think of data produced as the only way to represent itself. We are much more than our data, both as individuals and society, and we should always keep this in mind. Therefore, we should extend our ex-ante wishful thinking on data and algorithms beyond the concepts of data privacy and security. The same debate should also encompass data governance and the methods we apply to collect and use it.

The floor then passed to the event's Chair Luca De Biase, who gave his final remarks, mentioning the need to balance between more concrete governance to improve the impact of technology on our society and the flexibility required to stimulate innovation in a constantly developing field.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 871042

SoBigData Magazine is published under the

project N° 871042 | Programme: H2020 - INFRAIA

Duration: 01/01/2020 - 31/12/2023