

José Pagán · Mounir Mokhtari ·
Hamdi Aloulou · Bessam Abdulrazak ·
María Fernanda Cabrera (Eds.)

LNCS 11862

How AI Impacts Urban Living and Public Health

17th International Conference, ICOST 2019
New York City, NY, USA, October 14–16, 2019
Proceedings



 Springer Open

Founding Editors

Gerhard Goos

Karlsruhe Institute of Technology, Karlsruhe, Germany

Juris Hartmanis

Cornell University, Ithaca, NY, USA

Editorial Board Members

Elisa Bertino

Purdue University, West Lafayette, IN, USA

Wen Gao

Peking University, Beijing, China

Bernhard Steffen

TU Dortmund University, Dortmund, Germany

Gerhard Woeginger 

RWTH Aachen, Aachen, Germany

Moti Yung

Columbia University, New York, NY, USA

More information about this series at <http://www.springer.com/series/7409>

José Pagán · Mounir Mokhtari ·
Hamdi Aloulou · Bessam Abdulrazak ·
María Fernanda Cabrera (Eds.)

How AI Impacts Urban Living and Public Health

17th International Conference, ICOST 2019
New York City, NY, USA, October 14–16, 2019
Proceedings

Editors

José Pagán
New York University
New York, NY, USA

Hamdi Aloulou
Digital Research Centre of Sfax
Sfax, Tunisia

María Fernanda Cabrera
Technical University of Madrid
Madrid, Spain

Mounir Mokhtari
Institut Mines-Télécom
Paris, France

Bessam Abdulrazak
Université de Sherbrooke
Sherbrooke, QC, Canada



ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-32784-2 ISBN 978-3-030-32785-9 (eBook)
<https://doi.org/10.1007/978-3-030-32785-9>

LNCSSublibrary: SL3 – Information Systems and Applications, incl. Internet/Web, and HCI

© The Editor(s) (if applicable) and The Author(s) 2019. This book is an open access publication.

Open Access This book is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this book are included in the book's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the book's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

This year we organized the 17th ICOST conference, an event which has succeeded in bringing together a community from different continents for over a decade and half and raised awareness about frail and dependent people's quality of life in our societies.

After 16 very successful conferences held in France (2003, 2009, 2017), Singapore (2004, 2013, 2018), Canada (2005, 2011), Northern Ireland (2006), Japan (2007), USA (2008, 2014), Korea (2010), Italy (2012), Switzerland (2015), and China (2016), we decided to open the conference to public health experts and tackle emerging challenges related to AI technologies for health and well-being. This 17th edition of the International Conference on Smart Living and Public Health (ICOST 2019), was hosted by the College of Global Public Health at New York University (NYU) and the New York Academy of Medicine (NYAM) in New York City, USA, during October 14–16, 2019. The theme of the conference this year was “How Does AI Impact Urban Living and Public Health?”

ICOST 2019 provided a premier venue for the presentation and discussion of research in the design, development, deployment, and evaluation of AI for health, smart urban environments, assistive technologies, chronic disease management, and coaching and health telematics systems. ICOST 2019 aimed to understand and assess how research impacts public health policies when facing emerging social and economic challenges. ICOST 2019 brought together stakeholders from health care, public health, academia and industry along with end users and family caregivers to explore how to utilize technologies to foster health prevention, independent living, and offer an enhanced quality of life. The ICOST 2019 conference featured a dynamic program incorporating a range of technical, clinical, and industrial related keynote speakers, oral and poster presentations, along with demonstrations and technical exhibits. Specific use cases were presented, such as the PULSE European funded project aiming at accelerating the use and vision of Big Data by designing a system to exploit Big Data Value (BDV) in the public health sector. PULSE (Participatory Urban Living for Sustainable Environments) is incorporating several major cities as test beds: Barcelona (Spain), New York City (USA), Paris (France), Birmingham (United Kingdom), Singapore, Keelung (Taiwan) and Pavia (Italy).

ICOST 2019 was proud to extend its hospitality to an international community consisting of researchers from major universities and research centers, representatives from industry and users from 13 different countries. We would like to thank the authors for submitting their current research work and the Program Committee members for their commitment to reviewing submitted papers. The ICOST proceedings have now reached over 150,000 downloads, and are in the top 25% of downloads of Springer LNCS.

We were very pleased to host world-renowned keynote speakers from multiple backgrounds coming from all over the world. We are extremely thankful to our sponsors for their commitment and support to the vision and mission of ICOST.

October 2019

José Pagán
Mounir Mokhtari
Hamdi Aloulou
Bessam Abdulrazak
María Fernanda Cabrera

Organization

General Chair

José Pagán New York University and the New York Academy
of Medicine, USA

Conference Co-chairs

Mounir Mokhtari Institut Mines-Télécom, France, and Image
and Pervasive Access Lab., Singapore
Maria Fernanda Cabrera University Politecnica de Madrid, Spain

Steering Committee

Mounir Mokhtari Institut Mines Télécom, France, and Image
and Pervasive Access Lab., Singapore
Sumi Helal Lancaster University (2017–present), UK,
and University of Florida (1998–2017), USA
Bessam Abdulrazak University of Sherbrooke, Canada
Hamdi Aloulou University of Monastir, Tunisia, and Institut
Mines-Télécom, France
Jose Pagan New York University and the New York Academy
of Medicine, USA
Maria Fernanda Cabrera University Politecnica de Madrid, Spain

Scientific Advisory Board

Daqing Zhang Institut Mines Télécom, Télécom SudParis, France
Hisato Kobayashi Hosei University, Japan
Jongbae Kim Science and Technology Research Center, Yonsei
University, South Korea
Christian Roux Institut Mines Télécom, France
Dong Jin Song National University of Singapore, Singapore,
and Griffith University, Australia
Sungyoung Lee College of Electronics and Information, Kyung Hee
University, South Korea
Timo Jämsä EAMBES, University of Oulu, Finland
Daby Sow IBM Research AI, USA

Program Committee

Chairs

Bessam Abdulrazak University of Sherbrooke, AmI-Lab., Canada
Hamdi Aloulou University of Monastir, Tunisia, and Institut
Mines-Télécom, France

Members

Abdallah M'Hamed Télécom SudParis - Evry, France
Belkacem Chikhaoui University of Sherbrooke, Canada
Benoît Encelle Université de Lyon, CNRS Université Lyon 1, LIRIS,
France
Boussada Rihab National School for Computer Science (ENSI), Tunisia
Charles Gouin-Vallerand Télé-Université du Québec, Canada
David Menga EDF R&D, France
Diane Cook Washington State University, USA
Farah Arab Université Paris 8, France
Hisato Kobayashi Hosei University, Japan
Hongbo Ni Northwestern Polytechnical University, China
Houssein Aloulou Faculté des Sciences Économiques et de Gestion de
Sfax, Tunisia
Ibrahim Sadek Singapore University of Technology and Design
(SUTD), Singapore
Jeffrey Soar University of Southern Queensland, Australia
Laurent Billonnet University of Limoges, France
Manfred Wojciechowski University of Applied Sciences Duesseldorf, Germany
Salim Hima ESME-SUDRIA, France
Sofia Ben Jebara Ecole supérieure des communications de Tunis
(SUPCOM), Tunisia
Timo Jämsä University of Oulu, Finland
Victor Manuel Ponce Diaz Université de Sherbrooke, Canada
Vigouroux Nadine Institut de Recherche en Informatique de Toulouse
(IRIT), France
Yves Demazeau Laboratoire d'Informatique de Grenoble (LIG), France

Industry Liaison Committee

Peter Wu ASUS Cloud, Taiwan
Daby M. Sow IBM Research AI, USA
Vladimir Urosevic Belit, Serbia
Suzanne Ballard FCL, UK

Organizing Committee

Chair

Nina Raffio NYU College of Global Public Health, USA

Members

Elisa Fisher The New York Academy of Medicine, USA

Hamdi Aloulou University of Monastir, Tunisia, and Institut
Mines-Télécom, France

Antoine de Marassé Enouf CNRS, Singapore

Valentina Tageo European Connected Health Alliance, Spain

Sponsors

Institut Mines Télécom, Paris, France

National Center for Scientific Research, France

PULSE H2020 Project, Europe

New York University College of Global Public Health, USA

The New York Academy of Medicine, USA

Contents

E-health Technology Design

Privacy and Security of IoT Based Healthcare Systems: Concerns, Solutions, and Recommendations	3
<i>Ibrahim Sadek, Shafiq Ul Rehman, Josué Codjo, and Bessam Abdulrazak</i>	
Designing an ICT Solution for the Empowerment of Functional Independence of People with Mild Cognitive Impairment: Findings from Co-design Sessions with Older People	18
<i>Silvia de los Ríos, Rebeca I. García-Betances, Miguel Páramo, María Fernanda Cabrera-Umpiérrez, Marta Vancells, Maite Garolera, Jakub Kazmierski, and María Teresa Arredondo Waldmeyer</i>	
Deployment of an IoT Solution for Early Behavior Change Detection	27
<i>Hamdi Aloulou, Mounir Mokhtari, and Bessam Abdulrazak</i>	
Long Short Term Memory Based Model for Abnormal Behavior Prediction in Elderly Persons	36
<i>Meriem Zerkouk and Belkacem Chikhaoui</i>	

Well-being Technology

A Deep Learning Method for Automatic Visual Attention Detection in Older Drivers	49
<i>Belkacem Chikhaoui, Perrine Ruer, and Évelyne F. Vallières</i>	
Smart Mat for Respiratory Activity Detection: Study in a Clinical Setting . . .	61
<i>Samuel Otis, Bessam Abdulrazak, Sofia Ben Jebara, Francois Tournoux, and Neila Mezghani</i>	
Non-invasive Classification of Sleep Stages with a Hydraulic Bed Sensor Using Deep Learning.	73
<i>Rayan Gargees, James M. Keller, Mihail Popescu, and Marjorie Skubic</i>	

Biomedical and Health Informatics

A Convolutional Gated Recurrent Neural Network for Epileptic Seizure Prediction	85
<i>Abir Affes, Afef Mdhaffar, Chahnez Triki, Mohamed Jmaiel, and Bernd Freisleben</i>	

Ubiquitous Healthcare Systems and Medical Rules in COPD Domain	97
<i>Hicham Ajami, Hamid Mcheick, and Karam Mustapha</i>	
DL4DED: Deep Learning for Depressive Episode Detection on Mobile Devices	109
<i>Ajef Mdhaftar, Fedi Cherif, Yousri Kessentini, Manel Maalej, Jihen Ben Thabet, Mohamed Maalej, Mohamed Jmaiel, and Bernd Freisleben</i>	
ICT-Based Health Care Services for People with Spinal Cord Injury: A Pilot Study	122
<i>Wanho Jang, Dongwan Kim, Jeonghyun Kim, Seungwan Yang, Yunjeong Uhm, and Jongbae Kim</i>	
Smart Environment Technology	
An Interconnected Smart Technology System for Individuals with Mental Illness Living in the Community and Transitional Hospital Apartments	131
<i>Cheryl Forchuk, Jonathan Serrato, Abraham Rudnick, Deborah Corring, Rupinder Mann, and Barbara Frampton</i>	
Transfer Learning for Urban Landscape Clustering and Correlation with Health Indexes	143
<i>Riccardo Bellazzi, Alessandro Aldo Caldarone, Daniele Pala, Marica Franzini, Alberto Malovini, Cristiana Larizza, and Vittorio Casella</i>	
An IoT Architecture of Microservices for Ambient Assisted Living Environments to Promote Aging in Smart Cities	154
<i>Hubert Kenfack Ngankam, H�el�ene Pigot, Maxime Parenteau, Maxime Lussier, Aline Aboujaoud�e, Catherine Lalibert�e, M�elanie Couture, Nathalie Bier, and Sylvain Giroux</i>	
Designing a Navigation System for Older Adults: A Case Study Under Real Road Condition	168
<i>Perrine Ruer, Damien Brun, Charles Gouin-Vallerand, and �Evelyne F. Valli�eres</i>	
Short Contributions	
User Embeddings Based on Mobile App Behavior Data.	183
<i>Kushal Singla, Satyen Abrol, and Sungdeuk Park</i>	
Design, Development and Initial Validation of a Wearable Particulate Matter Monitoring Solution	190
<i>Jos�e G. Teri�us-Padr�on, Rebeca I. Garc�ia-Betances, Nikolaos Liappas, Mar�ia F. Cabrera-Umpi�errez, and Mar�ia Teresa Arredondo Waldmeyer</i>	

Cooperative System and Scheduling Algorithm for Sustainable Energy-Efficient Communities 197
Esther Palomar, Carlos Cruz, Ignacio Bravo, and Alfredo Gardel

ForeSight - Platform Approach for Enabling AI-based Services for Smart Living 204
Jochen Bauer, Hilko Hoffmann, Thomas Feld, Mathias Runge, Oliver Hinz, Andreas Mayr, Kristina Förster, Franz Teske, Franziska Schäfer, Christoph Konrad, and Jörg Franke

Mobility Application with Semantic Reasoning 212
Martin Kodyš, Antoine de Marassé, and Mounir Mokhtari

Author Index 223



Transfer Learning for Urban Landscape Clustering and Correlation with Health Indexes

Riccardo Bellazzi¹(✉), Alessandro Aldo Caldarone¹, Daniele Pala¹,
Marica Franzini², Alberto Malovini³, Cristiana Larizza¹,
and Vittorio Casella²

¹ Department of Electrical, Computer and Biomedical Engineering,
via Ferrata 5, 27100 Pavia, Italy
{riccardo.bellazzi, cristiana.larizza}@unipv.it,
{alessandroaldo.caldarone01,
daniele.pala02}@universitadipavia.it

² Department of Civil Engineering and Architecture, via Ferrata 3,
27100 Pavia, Italy

{marica.franzini, vittorio.casella}@unipv.it

³ IRCCS ICS Maugeri, via S. Maugeri 2, 27100 Pavia, Italy
alberto.malovini@icsmaugeri.it

Abstract. Within the EU-funded Pulse project, we are implementing a data analytic platform designed to provide public health decision makers with advanced approaches to jointly analyze maps and geospatial information with health care data and air pollution measurements. In this paper we describe a component of such platform, designed to couple deep learning analysis of geospatial images of cities and some healthcare and behavioral indexes collected by the 500 cities US project, showing that, in New York City, urban landscape significantly correlates with the access to healthcare services.

Keywords: Transfer learning · Deep learning · Urban landscape · Health indexes

1 Introduction

Recent advances in machine learning and deep learning enable the design and implementation of novel data analysis pipelines that allow fusing heterogeneous data sources to extract novel insights and predictive patterns. These approaches seem particularly suitable to help increasing our insights in the relationships between the urban landscape of cities and the behavior of their residents, with particular focus on well-being and healthcare indexes. In this context, it can be of interest of health care planners and city decision-makers to have instruments able to find clusters of city areas that share similar urban structures and to analyze some behavioral indexes of their residents, in particular to see potential correlations and to plan similar interventions in the different clusters, even if such clusters contain areas that are geographically far. We have applied such approach in the context of the PULSE (Participatory Urban Living

for Sustainable Environment) EU-project¹. PULSE aims at developing a set of models and technologies to predict and manage public health problems in cities and promote health. It follows a participatory approach where citizens provide data through personal devices that are integrated with information from heterogeneous sources: open city data, health systems, urban sensors and satellites. The project deals with various issues concerning air quality, lifestyle and personal behavior and it aims to investigate the correlations between the exposure to atmospheric pollutants, the citizen habits and the health of the citizen themselves, focusing on asthma and type 2 diabetes. PULSE is being implemented in 5 major cities all over the world. Within PULSE, we are implementing a data analytic platform that will provide public health decision makers with advanced approaches to jointly analyze maps and geospatial information with health care data and air pollution measurements.

In this paper we will describe the results obtained with a prototypical component of such platform, designed to couple deep learning analysis of geospatial images of cities and some healthcare and behavioral indexes, showing that in New York City urban landscape significantly correlates with the access to healthcare services.

2 Deep Learning and Transfer Learning Models

Deep neural models provide flexible instruments to perform non-linear approximation of a variety of multivariate functions and to extract latent variables from a data set. In a nutshell, deep neural models are neural networks with many layers, able to map non-linear functions with a number of parameters that is typically lower than their equivalent models with one layer only. Such models are particularly attractive since they can be used to perform clustering, regression and classification starting from data sets made of images, texts, time series.

In dependence of the nature of the input data set, different architectures can be exploited, ranging from the combination of many Convolutional layers in the case of images to the use of Long-term/Short-term networks in the case of time series and speech/text data.

Recently, an increasing number of papers are using deep learning to examine the relationships between the urban landscape and some environmental or citizens' behavioral data [1–3].

One of the main limits of deep learning models is related to the need of very large data sets in order to be able to gain advantage of their capability of encoding even the finest details that can be important to map input data, without getting trapped into noise and poor parameters estimates.

Rather interestingly, in order to deal with this problem, it is possible to resort to an increasing set of pre-trained deep learning models that can be used for the task of transfer learning [4], i.e. models that are able to represent the input space into a set of latent variables on the basis of a mapping mechanism, usually a deep neural network, learned on a large (external) data set, so that the relationships between such latent

¹ "<http://www.project-pulse.eu>".

variables and the outcomes can be later learned on a specific and smaller data set. A well-known example is Inception-v3, a convolutional neural network trained on more than a million images from the ImageNet database (<http://www.image-net.org>). The network has 48 layers and can classify images into one thousand object categories, including trees and many animals. Another interesting example is represented by the Painters [5] networks, developed to automatically classify Paintings of famous artists. In principle, any of those methods can be used following the transfer learning paradigm to represent images coming from urban landscape of New York City.

3 Data and Methods

Our analysis is based on two data sources: NYC high resolution images and healthcare data coming from the 500 cities project [6]. NYC images have been collected by the “The National Agriculture Imagery Program” (NAIP) that acquires aerial imagery during the agricultural growing seasons in the continental United States. In particular, we have downloaded an image having an original resolution of 0.5 m and have downsampled it to 2 m which allows to have a fine-grained representation of the aerial urban landscape (see Fig. 1).

As it will be explained in the following, the reason for the downsampling is that the big image has been subdivided into tiles and the neural network adopted can accept images having maximum size of 299 pixel; we had to tune the ground resolution in order to have meaningful tiles, embracing a sufficiently-sized area.




Fig. 1. NAIP image of NYC.

Health care data have been extracted from the repository made available by the 500 Cities project. “500 cities” is a collaboration between CDC, the Robert Wood Johnson Foundation, and the CDC Foundation². The project provides city- and census tract-level small area estimates for chronic disease risk factors (unhealthy behaviors), health outcomes, and clinical preventive service use for the largest 500 cities in the United States. NYC is divided in 2166 census tracts and the latest data available concerns 2017. The 27 chronic diseases measures provided by the project are listed in Table 1.

The measures include major risk behaviors that lead to illness, suffering, and early death related to chronic diseases and conditions, as well as the conditions and diseases that are the most common, costly, and preventable of all health problems.

Table 1. 500 cities measures grouped by category. The 27 measures include 13 health outcomes, 9 prevention practices and 5 unhealthy behaviors.

Category	Measure
Health outcomes	Arthritis among adults aged ≥ 18 years
	Current asthma among adults aged ≥ 18 years
	High blood pressure among adults aged ≥ 18 years
	Cancer among adults aged ≥ 18 years
	High cholesterol among adults aged ≥ 18 years who have been screened in the past 5 years
	Chronic kidney disease among adults aged ≥ 18 years
	Chronic obstructive pulmonary disease among adults aged ≥ 18 years
	Coronary heart disease among adults aged ≥ 18 years
	Diagnosed diabetes among adults aged ≥ 18 years
	Mental health not good for ≥ 14 days among adults aged ≥ 18 years
	Physical health not good for ≥ 14 days among adults aged ≥ 18 years
	All teeth lost among adults aged ≥ 65 years
	Stroke among adults aged ≥ 18 years
Prevention	Current lack of health insurance among adults aged 18–64 years
	Visits to doctor for routine checkup within the past year among adults aged ≥ 18 years
	Visits to dentist or dental clinic among adults aged ≥ 18 years
	Taking medicine for high blood pressure control among adults aged ≥ 18 years with high blood pressure
	Cholesterol screening among adults aged ≥ 18 years
	Mammography use among women aged 50–74 years
	Papanicolaou smear use among adult women aged 21–65 years
	Fecal occult blood test, sigmoidoscopy, or colonoscopy among adults aged 50–75 years
	Older adults aged ≥ 65 years who are up to date on a core set of clinical preventive services by age and sex
Unhealthy behaviors	Binge drinking among adults aged ≥ 18 years
	Current smoking among adults aged ≥ 18 years
	No leisure-time physical activity among adults aged ≥ 18 years
	Obesity among adults aged ≥ 18 years
	Sleeping less than 7 h among adults aged ≥ 18 years

² <https://www.cdc.gov/500cities/index.htm>.

3.1 The Data Analysis Pipeline

The pipeline implemented in our work is described in Fig. 2. The NAIP NYC image has been subdivided into image square blocks having size of 256×256 pixels, corresponding to a 512 m edge. Therefore, it was possible to estimate the value of each of the 27 variables collected by “500 Cities” for each block. During this process, blocks out of the tracts or over the sea have been excluded, thus reducing the dataset. The images have been then processed by a pretrained deep model, thus extracting the final features for each image. Images are clustered by resorting to k-means clustering, and the clusters, confirmed with visual inspection, are associated to the healthcare indexes by statistical analysis.

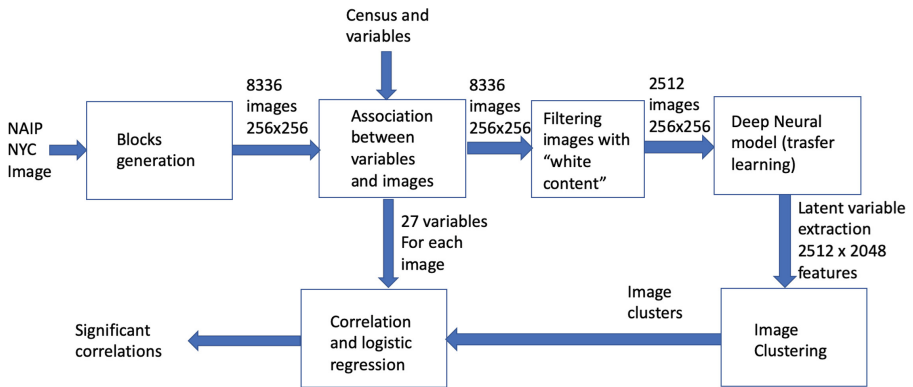


Fig. 2. The data analysis pipeline

3.2 Image Blocks

The NAIP NYC image has been subdivided into 8336 images of 256×256 pixels. Each image is a square with edge equal to 512 m. It must be underlined that the original image is georeferenced, meaning that each single pixel is precisely located in space. The small derived tiles are georeferenced as well and can be effectively overlapped to the health and well-being maps. The images have been processed resorting to the Matlab Image Processing and Mapping toolboxes, which are capable of properly managing georeferenced images. Figure 3 shows some examples of the resulting images.

It is possible to note that some of the squares have white areas, corresponding to unmapped zones, due to the irregular borders of the image and to the presence of sea, rivers. Due to the availability of the vector map of the borders of NYC, we have been able to quantify, for each tile, the amount of its surface lying inside the borders of the city; we then filtered the original tile set and maintained only those having a minimal overlapping of 90%.

3.3 Estimation of the Healthcare Indexes for Each Image Block

The healthcare indexes of the 500 Cities database are collected for census tracts and NYC has, as already reported, 2166 census tracts. In order to carry out our analysis,

we had to determine the value of the considered variables for each image block. In fact, a given tile overlaps, in general, several tracts. Therefore, we had to implement a simple estimator of the healthcare index of the block, as:

$$hci(block) = \frac{\sum_j w_j hci(j)}{\sum_j w_j}$$

where $hci(j)$ is the value of the generic health care index for the j -th census tract and w_j is the percentage of the image block covered by the mentioned tract. An example is shown in Figs. 3 and 4.

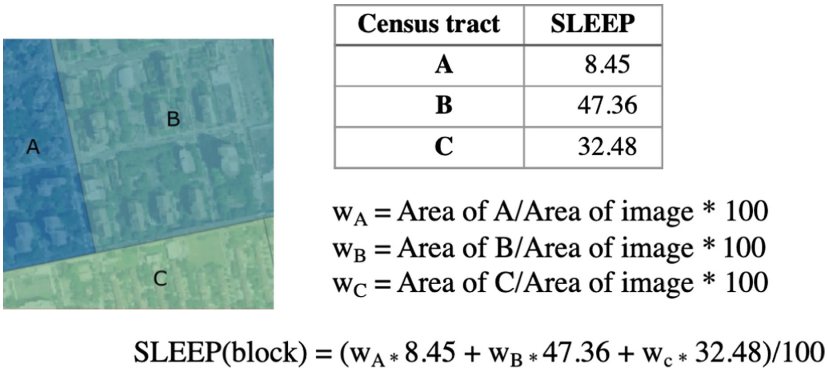


Fig. 3. The quantification of the healthcare index value (SLEEP) of a block.




Fig. 4. Original census tracts with the 500 Cities SLEEP index (left hand side) and derived quantification of the healthcare index values for SLEEP variable (right hand side).

In order to properly quantify $hcis$, the blocks with a white area greater than 10% of the image have been removed. The final number of image blocks used for the following analysis has thus lowered to 2512 images. Each image has been then processed by resorting to a deep neural model to extract a set of latent features.

3.4 Deep Neural Networks Processing and Clustering

As a deep neural network model used for transfer learning, we have selected the network developed for the 2016 Painters by number competition [6]. In such competition the goal was to learn how to discriminate the authors of paintings between 1584 unique painters, starting from a training set of 79433 instances; the test set was composed of 23817 instances. In this case, a deep neural network model was learned, with 23 layers, mostly convolutional layers with some max pooling layer. The Painters network computes a layer of 2048 latent variables before the final discrimination layer implemented with a soft-max non-linear function. Those latent variables can be used as a way to embed generic images in the latent space. Therefore, using the software Orange (<https://orange.biolab.si>) and its Python pipeline, we have processed all image blocks with the Painters model, thus obtaining a final data matrix of 2512 examples with 2048 features.

Such features have been used to cluster the image blocks by resorting to the well-known K-means clustering algorithm. The value of K has been derived with a grid search between 2 and 6 and taking the value that maximize the Silhouette coefficient.

3.5 Correlation and Statistical Analysis

The final step of the data analysis pipeline is represented by the search of statistical correlations between the clusters and their *hcis*. Univariate multinomial logistic regression was applied to estimate the probability to belong to a specific cluster given single variables' values. Multivariate multinomial logistic regression was performed after removal of samples characterized by missing values. A backward stepwise selection procedure based on AIC was applied to identify the most informative set of variables jointly modulating the probability to belong to the clusters. Multinomial logistic regression and the stepwise selection procedure were implemented in the R packages "nnet" and "stats", respectively. Analyses were performed by the R software tool version 3.5.1 (<http://www.r-project.org>).

4 Results

4.1 Clustering

K-means was run on the 2512 instances with Euclidean distance and 10 reruns. 4 clusters were found to maximize Silhouette coefficient. The output of the clustering algorithm has been validated by analyzing the cluster distribution with the tSNE two-dimensional mapping, as reported in Fig. 5. It is easy to see that the four clusters are in general well separated in the two-dimensional space³.

³ It is worthwhile mentioning that this criterion was qualitatively used to assess also other deep neural networks model; Painters turned out to generate the clusters that had the best tSNE spatial distribution of clusters.

Thanks to visual inspection, it is possible to highlight that the four clusters well correspond to different urban landscapes. Cluster C1 corresponds to green areas, Cluster C2 to residential areas with small houses, Cluster C3 to industrial areas and larger buildings, Cluster C4 to residential with larger buildings. Four examples are shown in Fig. 6. Cluster analysis clearly show that the deep neural network model is able to map images in the latent space that share the intuitive notion of similarity that humans may use when they have to classify urban landscape. The method is thus able to automatically cluster similar areas where similar interventions can be planned.




Fig. 5. The tSNE representation of the data with colors identifying the four clusters.




Fig. 6. Four images representing the clusters.

4.2 Correlation and Statistical Analysis

Univariate analysis shows that 22 over 27 variables were significantly correlated with the clusters. This is also confirmed with visual inspection showing the variables distributions after equal frequency discretization against the clusters, as shown in Fig. 7.

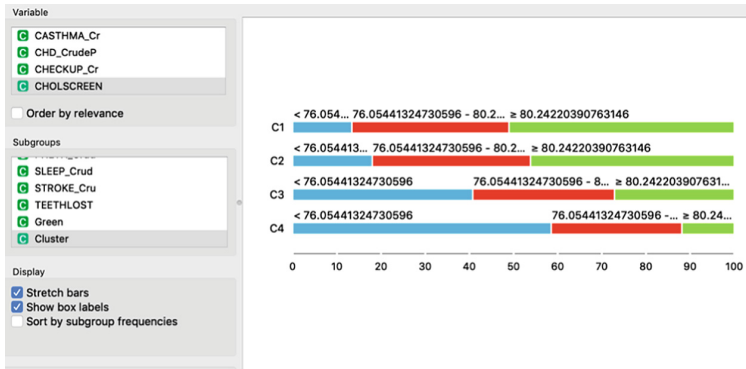


Fig. 7. The different distributions of Cholesterol screening among adults aged ≥ 18 years in the different clusters. Inhabitants of cluster C1 have much higher propensity towards screening than those who live in Cluster C4.

In general, cluster C1, which is the one that groups green areas, has consistently better prevention and health indicators, but worse sleeping indexes and leisure time. Overall, there is a gradient with all indexes moving from cluster C1, to C2, to C3 and finally to C4, which are the residential areas with large buildings.

A multivariate multinomial logistic regression has been performed to assess if significant correlations are present even in the multivariate setting. In this case, after a stepwise feature selection process, 20 variables have been selected. Of those, five variables have been found to be significant ($p \ll 0.01$) in all sub-regressions performed by the multinomial model: Colon screening (Fecal occult blood test, sigmoidoscopy, or colonoscopy among adults aged 50–75 years), Chronic obstructive pulmonary disease among adults aged ≥ 18 years, High cholesterol among adults aged ≥ 18 years who have been screened in the past 5 years, Chronic kidney disease among adults aged ≥ 18 years and finally Stroke among adults aged ≥ 18 years.

4.3 Mapping

The blocks and the clusters have been represented in the original map, confirming the qualitative evaluation of the clusters reported above (Fig. 8).




Fig. 8. The clusters remapped in NYC. (Color figure online)

We can find green areas (C1), residential areas with larger buildings (C2), industrial areas (C3) and finally residential areas (C4).

5 Discussion and Conclusions

The data analysis pipeline described in this paper shows that it is possible to automatically correlate urban landscape with healthcare indicators at the whole city level. In the NYC case, such correlation seems particularly strong, probably because of social factors, which, in the US society, makes health indicators related to the urban areas where people live.

Our work has a number of implications.

First of all, it shows that deep neural networks designed to encode image data can be successfully reused within transfer learning approaches. Their application to represent urban landscape seems very effective.

Second, in the context of the PULSE project, the capability of finding clusters of similar urban landscape may allow to profile city areas, in which health care decision makers may plan similar interventions.

Finally, the combination of urban landscape and healthcare indicators is not only useful to hypothesize the intertwining of these two dimensions, but also to further profile urban areas by finding similar areas with similar behaviors of their inhabitants, thus allowing also life style interventions and more “precise” health care policies.

Of course, the analysis has some limitations. First of all, the “quantification” of the health care indexes in the city blocks have been performed by a weighted averaging of the indexes of the census tracts included in the blocks. The weights are computed taking into account only the spatial overlap and not the actual number of inhabitants of the blocks. Second, the results obtained are probably “proxies” of the wealth of the people living in the different areas. For this reason, results may be representative of specific cities and not generalizable to other ones.

Acknowledgement. The work is part of the project PULSE, H2020 - 727816, funded by the European Union.

References

1. Helbich, M., Yao, Y., Liu, Y., Zhang, J., Liu, P., Wang, R.: Using deep learning to examine street view green and blue spaces and their associations with geriatric depression in Beijing, China. *Environ Int.* **126**, 107–117 (2019)
2. Hong, K.Y., Pinheiro, P.O., Minet, L., Hatzopoulou, M., Weichenthal, S.: Extending the spatial scale of land use regression models for ambient ultrafine particles using satellite images and deep convolutional neural networks. *Environ. Res.* **30**(176), 108513 (2019)
3. Zewdie, G.K., Lary, D.J., Levetin, E., Garuma, G.F.: Applying deep neural networks and ensemble machine learning methods to forecast airborne ambrosia pollen. *Int. J. Environ. Res. Public Health* **16**(11), 1992 (2019)
4. Sharma, S., Ball, J.E., Tang, B., Carruth, D.W., Doude, M., Islam, M.A.: Semantic segmentation with transfer learning for off-road autonomous driving. *Sensors (Basel)*. **19**(11), 2577 (2019)
5. Painters by Numbers. <http://blog.kaggle.com/2016/11/17/painter-by-numbers-competition-1st-place-winners-interview-nejc-ilenic/>. Accessed 21 June 2019
6. Centers for disease control and prevention. National center for chronic disease prevention and health promotion. Division of population health. 500 cities project data (2019). <https://www.cdc.gov/500cities>. Accessed 21 June 2019

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

