BRUNO EMANUEL DA GRAÇA MARTINS
Department of Computer Science and Engineering
Instituto Superior Técnico
University of Lisbon
Email: bruno.g.martins@ist.utl.pt

Bruno Martins is an assistant professor at the Computer Science and Engineering Department of IST, the Engineering school of the University of Lisbon, and a researcher at the Information and Decision Support Systems Lab of INESC-ID. His research focuses on the general areas of information retrieval, text mining, and the geographical information sciences. He has been involved in several research projects related to geospatial aspects in information access and retrieval, and he has accumulated a significant expertise in addressing challenges at the intersection of multi-modal information retrieval and the geographical information sciences (i.e., in an area that is often referred to as geographical information retrieval).

In terms of activities related to the Spatial Data Science Symposium, Martins is currently the principal investigator in a nationally funded project named MIning MUlti-source and MUltimodal geo-referenced information (MIMU), which concerns the study of intelligent services for integrating and processing multi-modal geospatial information from non-traditional sources, e.g. text documents or geo-referenced photos collected from social media. He is also a participating researcher and/or national coordinator in several other research projects (e.g., the RiverCure project on flood monitoring, the DigCH project on spatial humanities, or the MATISSE project on remote sensing and data mining for shellfish safety monitoring) and network activities (e.g., the CYTED IDEAIS network on intelligent assistants supporting access to geospatial data). He is also involved in the organization of ACM SIGSPATIAL workshops on GeoAI and Geospatial Humanities, and a guest-editor for journal special issues also focusing on these areas.

Geospatial Artificial Intelligence and the Combined Analysis of Multimodal Data

Over the last few years, deep learning techniques had a transformative impact in fields such as natural language processing or computer vision, significantly advancing the stateof-the-art in problems like parsing natural language, assigning text or images to predefined classes, or semantic image segmentation. These same techniques can also empower a next generation of geographical information systems, providing the ability to combine spatial analysis with fast and near human-level perception, this way facilitating location-based discovery and analysis of relevant information, originally available in multiple modalities. With colleagues and students from the University of Lisbon, I have been involved in a variety of applied research projects related to this general objective, which I believe are at the core of Spatial Data Science.

In connection to remotely sensed data, or combining remote sensing with volunteered geographic information (e.g., from sources such as the OpenStreetMap), recent studies have for instance shown that deep learning techniques can be used in a variety of applications. Examples include (a) geo-referencing ground-level photos in order to complement remote sensing data, (b) interpolating, downscaling, or fusing multiple remote sensing products, e.g. to circumvent problems of missing data, or (c) segmenting aerial imagery, e.g. to support the extraction of building footprints and/or road segments.

In connection to textual data, or combining text with other data modalities, recent studies have also described the use of deep learning techniques for geo-referencing vague and imprecise place references into unambiguous coordinates, so that traditional spatial analysis methods can be employed over data originally encoded in text. Example applications include (a) automatically geo-referencing Twitter users or their textual posts, (b) matching place records combining names and other attributes, e.g. for gazetteer conflation, or (c) disambiguating place references, presented in the context of textual descriptions, into the corresponding geo-spatial coordinates.

Many of aforementioned applications can already highlight important challenges in the use of machine learning for processing geospatial data, including model interpretability, small sample sizes relative to the complexity of the problems, the lack of ground truth information, or the high degree of noise and uncertainty. Despite the many successful applications, overcoming the aforementioned challenges requires additional developments (e.g., combining deep learning with symbolic and/or reasoning based approaches, or exploring probabilistically sound deep generative models), so that recent deep learning techniques can be more widely and easily applicable to a broader range of tasks. At the UCSB Spatial Data Science Symposium, I specifically plan to present examples of ongoing research at the University of Lisbon, aligned with the perspectives outlined in the previous paragraphs, and which I believe can contribute to a discussion regarding a possible agenda for Spatial Data Science research. Particular emphasis will be given to the combination of language technologies, image processing, and spatial analysis methods, specifically focusing on the following examples:

- Research related to place name resolution in textual documents, combining recent advances in natural language processing (e.g., contextual word embedding methods such as ELMo or BERT, to build effective text representations) together with the idea of exploring external sources of geo-physical information (e.g., raster datasets encoding terrain elevation, land coverage, population density, etc.) to guide model training, under the idea that place names often appear in sentences containing descriptions for the characteristics of the corresponding locations. We have conducted experiments with deep neural network models that take as input sequences of words corresponding to place name occurrences, plus the surrounding textual context. These inputs are transformed into vector representations, which are then used to predict a probability distribution over possible geospatial regions, corresponding to the place name. The probability distribution is used to infer the geospatial coordinates corresponding to the place name, or the geo-physical properties for the location corresponding to the place name, this way better guiding the model towards making correct predictions. We have assessed the advantages of this approach in several corpora used in previous studies, and we also assessed the impact of using distant-supervision (i.e., data collected from Wikipedia) in order to increase the amount of available training data.

- Research related to semantic segmentation of aerial imagery (e.g., delimiting building footprints, or delimiting flooded areas), leveraging fully-convolutional neural networks (e.g., dense U-net neural architectures) together with optimisations regarding model training (e.g., use data pre-processing and augmentation procedures to improve model generalisation, such as slightly changing the colour temperature of the RGB images, together with adversarial training and/or loss functions that account for class imbalance and emphasise correct predictions at boundary regions). Our experiments have given particular emphasis to the processing of historical aerial imagery (e.g., segmenting grayscale imagery, together with a separate model branch that attempts to colorize the images), and also to the processing of RGB images together with other spectral information (i.e., raw near infra-red data, or derived indexes such as NDVI or NDWI for segmenting flooded regions) or additional remote sensing products (e.g., information on terrain elevation or permeability). More recent efforts have also explored the usage of aerial imagery in tasks involving other data modalities (e.g., remote sensing image captioning, for supporting information retrieval applications), or the use of remote sensing data together with information collected on the ground (e.g., using geo-referenced photos, collected from social media, in the context of flood mapping).

Besides advances in fundamental research, we believe the adaptation/customisation of existing methods, as well as the corresponding experimental validation, demands particular emphasis in a research agenda for Spatial Data Science. Application areas such as the digital humanities, or challenges like geographical information retrieval within digital libraries, are also particularly interesting, providing ground-truth data and plentiful of challenging problems to support advances in the field. Examples include assessing the use of deep learning developments in contexts such as (a) the processing of historical maps instead of traditional remote sensing products, (b) the analysis of old grayscale photos instead of modern high-resolution aerial images, or (c) the parsing of information within historical documents in a variety of languages, instead of just modern English text. To accelerate developments in connection to these challenges, I believe that the organisation of joint workshops or journal special issues can assume a particular importance, fostering synergistic discussions and collaborative research. Perhaps even more importantly, the organisation of joint evaluation efforts in connection to the previous forums, considering challenges/tasks and ground-truth datasets with crisp problem statements, can be important to accelerate progress and lower the barriers for more people to get involved. Finally, surveys and curated repositories of learning materials can help educating new researchers, particularly by reducing the challenges involved in understanding advanced cross-disciplinary topics.