Editor's Note

Artificial Intelligence (AI) has become nowadays one of the main relevant technologies that is driven us to a new revolution, a change in society, just as well as other human inventions, such as navigation, steam machines, or electricity did in our past. There are several ways in which AI might be developed, and the European Union has chosen a path, a way to transit through this revolution, in which Artificial Intelligence will be a tool at the service of Humanity. That was precisely the motto of the 2020 European Conference on Artificial Intelligence ("Paving the way towards Human-Centric AI"), of which these special issue is a selection of the best papers selected by the organizers of some of the Workshops in ECAI 2020.

These workshops constitute a diverse list of different subjects that are relevant to AI at present, some that envision the future, and finally also multidisciplinary topics in an increasingly transversal discipline. This selection aims at proposing, discussing and finding ways to confront the many challenges that lie ahead and for which solutions need to be found. Designing the correct strategy is crucial to be able embrace a future in which AI ensures empowering people, making true the conference motto.

The selected papers belong to the following Workshops: Hybrid Intelligence for Natural Language Processing Tasks (HI4NLP); Applied Deep Generative Networks (ADGN); Declarative Problem Solving (DPSW); Advancing towards the Sustainable Development Goals: AI for a fair, just and equitable world (AI4EQ); Evaluating Progress in AI (epAI); Singular Problems for Health Care (SP4HC); Intelligent Information Processing and Natural Language Generation (IntelLanG); and Data Fusion for Artificial Intelligence (DAFUSAI).

With the workshop "Advancing Towards the Sustainable Development Goals (SDGs): Artificial Intelligence for a Fair, Just and Equitable World (AI4Eq)" we aimed to illustrate the R&D path that would confer a decisive role to AI in achieving the UN Agenda 2030. Eradicating poverty is a central objective of the SDGs so the emphasis is on AI benefits for low and middle income countries, and the growing pockets of underdevelopment in high income countries.

There is a growing interest in the role that AI can play in achieving SDGs on the part of international organisations, such as UN Global Pulse [1], UNHCR [2], the UNICEF Global Innovation Centre [3], the World Wide Web Foundation [4], the International Telecommunications Union [5], and even the World Economic Forum [6]. In order for AI to catalyse the necessary transformation promoted by the 2030 Agenda, a research agenda that is practice-oriented and that goes beyond cataloging AI risks and potentialities is required, in part as a counter-weight to the heavily-plugged corporate sector view on AI ethics, which is often little more than "ethicswash" for a program in which the effect of AI/S development and deployment will most likely be to increase inequality [7] [8]. The three papers selected from the submissions to AI4Eq for this special issue describe research on SDG-oriented AI applications, as well as AI tools conceived to support the development of AI respectful of, and even actively committed to, fundamental human rights, focusing particularly on protecting and empowering the most vulnerable and marginalized.

The paper "Achieving fair inference using error-prone outcomes" focusses on a field that is attracting increasing interest: the assessment of fairness criteria in supervised learning. The authors demonstrate that existing methods to assess and calibrate fairness criteria do not extend to the true target variables of interest, when error-prone proxy targets are used. They propose a framework that combines fair machine-learning methods, such as those found in the fairness literature, and measurement models found in the statistical literature; and illustrate their approach in a healthcare decision problem showing how a latent variable model to account for measurement error removes the unfairness detected previously.

The paper "Attesting Digital Discrimination Using Norms" also addresses the problem of digital discrimination arising from bias in machine-learning algorithms. In this case, the authors point to the need to provide non-expert users of machine-learning algorithms with simple tools to determine if a machine-learning system is potentially discriminatory, and to make explicit under which assumptions the systems are discrimination free. The authors suggest using "norms" as an abstraction to represent different situations that may lead to digital discrimination. In particular, they formalise non-discrimination norms in the context of machine-learning systems and propose a digital-discrimination attesting algorithm to check whether the systems violate these norms, illustrating its performance in three case studies where, in particular, gender and racial biases are identified.

"No App is an Island: Collective Action and Sustainable Development Goal-Sensitive Design" deals with the challenges of engineering ever more complex socio-technical systems to address "wicked" societal problems, with respect to satisfying qualitative human values and to assessing their impact on global challenges. The authors present a set of sets of design principles and an associated meta-platform, which focusses the design of socio-technical systems on the potential interaction of human and artificial intelligence with respect to three aspects: firstly, decision support regarding the codification of deep social knowledge; secondly, visualisation of community contribution to successful collective action; and thirdly, systemic improvement with respect to the SDGs through impact assessment and measurement. This SDG-sensitive design methodology is illustrated through the design of two collective action apps, one for encouraging plastic re-use and reducing plastic waste, and the other for addressing redistribution of surplus food.

The Workshop on Hybrid Intelligence for Natural Language Processing Tasks (HI4NLP) has provided a forum to discuss exciting research on hybrid technologies for NLP. In particular, our interest lies in those methodologies and architectures which combine and integrate symbolic information into statistical methods, including neural networks, thus allowing for building more transparent and interpretable models.

The paper "Assessing Lexical-Semantic Regularities in Portuguese Word Embeddings" introduces TALES, a new dataset with lexical-semantic word analogies. The authors use this resource to perform a detailed analysis of various word embeddings for Portuguese, including static representations such as GloVe [9] and word2vec [10], and current contextualized models (e.g., BERT [11]). Interestingly, this paper also discusses how distributional models can be used to enlarge lexical-semantic knowledge bases, which can be beneficial to various natural language processing tasks.

On "The Semantics of History. Interdisciplinary Categories and Methods for Digital Historical Research", Travé et al. present a conceptual framework for interdisciplinary research in History, focusing on data modeling and labeling methods. In particular, they propose identifying minimum units of information (units of topography, units of stratigraphy, and actors) and their relations, for which methodological aspects are described. A detailed case study on landscape archaeology shows the usefulness of the proposed framework, which takes advantage of knowledge obtained from several sources.

ECAI 2020 also hosted the first edition of the Declarative Problem Solving Workshop (DPSW), gathering researchers from different AI disciplines with a common interest in solving computational problems via their explicit representation in some declarative language. This covers, for instance, the solution of combinatorial problems, optimization, numerical constraints, planning, scheduling, temporal constraints, etc, or combinations of these categories provided that their specification is made in terms of some declarative formal language. The workshop spanned two days, August 29th and 30th, and included the presentations corresponding to thirteen accepted regular papers. The average audience was around 25 participants, reaching a maximum of 40 during the invited talk by Vladimir Lifschitz, that closed the event. Although the papers covered different disciplines like constraints, planning, natural language or pattern mining, perhaps the most frequent topic was the problem solving paradigm of Answer Set Programming (ASP). Two contributions obtained the workshop best paper award in a tie, and were extended into full journal papers included in this volume. The first best paper, "Smoke Test Planning using Answer Set Programming" by Tobias Philipp, Valentin Roland and Lukas Schweizer, presented a declarative method for optimizing the automated generation of smoke tests for hardware devices, that is, quick tests of main functionalities that may spot a large error in the early stages of hardware design. The second best paper, "An Application of Declarative Languages in Distributed Architectures: ASP and DALI microservices", by Stefania Costantini, Giovanni de Gasperis and Lorenzo de Lauretis, introduced an innovative combination of the microservices architecture with a modular variant of ASP, showing potential applications of declarative problem solving to Multi-Agent Systems, Internet of Things (IoT) or Cloud Computing.

The workshop on Intelligent Information Processing and Natural Language Generation (IntelLang) aimed to identify challenges and explore current results that arise from the interaction of Intelligent Information Processing techniques and research in Natural Language Generation (NLG), both at the level of models and applications. The use of intelligent data and information processing techniques can help in many relevant aspects of the NLG problem, for example in the contribution of formalisms for knowledge modeling and management, or in the development of models for the evaluation of the quality of the proposals, among many others. The workshop provided a forum for discussion of these new research directions and attracted a broad spectrum of contributions, emphasising either or both of the workshop's main themes - NLG and Information Processing. Our hope is that these contributions will serve to enhance the sharing of ideas among the two communities.

The paper "Improving Asynchronous Interview Interaction with Follow-up Question Generation" proposes a follow-up question generation model (followQG) capable of generating relevant and diverse follow-up questions based on the previously asked questions, and its answers. This model is integrated in a 3D virtual interviewing system, Maya, with capability of follow-up question generation, taking advantage of the implicit knowledge from deep pre-trained language models to generate rich and varied natural language follow-up questions. Empirical results suggest that followQG generates questions that humans rate as high quality, achieving 77% relevance, and a comparison with strong baselines of neural network and rule-based systems shows that it produces better quality questions.

The paper "Neural Scoring of Logical Inferences from Data using Feedback" proposes a neural network model that generates personalised lifestyle insights based on a model of their significance, and feedback from the user. These insights are derived from wearable sensors in smartwatches or sleep trackers, and their generation should adapt automatically to the preferences and goals of the user. Simulated analysis of the presented model shows its ability to assign high scores to a) insights with statistically significant behaviour patterns and b) topics related to simple or complex user preferences at any given time. The authors believe that the proposed neural networks model could be adapted for any application that needs user feedback to score logical inferences from data.

The first workshop on Evaluating Progress in Artificial Intelligence (EPAI) took place on September 4th and comprised 13 presentations and one invited talk from Professor Barry O'Sullivan, President of the European AI Association. There were over 30 attendants and a reasonably good number of (very active) attendees. EPAI 2020 served not only as a meeting point for people from different backgrounds and goals, but also to identify the most challenging/urgent needs for AI evaluation [12]. In this regard, it is very well-known that AI capabilities are growing at an unprecedented rate. Countless AI approaches and applications are being developed and can be expected over the long term. In hindsight, one would say that progress certainly has taken place just looking at the range of tasks that AI are able to solve autonomously today (according to the benchmarks, challenges, and competitions [13]) and were not solvable a few years ago, from machine translation to medical image analysis or self-driving vehicles [14]. Moreover, progress in AI is widely believed to have substantial social and economic benefits, and possibly to create unprecedented challenges. In order to properly prepare policy initiatives for the arrival of such technologies, accurate forecasts and timelines are necessary to enable timely action among policymakers and other stakeholders. However, there is still much uncertainty over how to assess and monitor the state, development, uptake, and impact of AI as a whole, including its future evolution, progress and benchmarking capabilities. While measuring the performance of state-of-the-art AI systems on narrow tasks is useful and fairly easy to do, where the assessment really becomes difficult, though, is in trying to map these narrow-task performances onto more general AI and how it can have an impact on society in terms of benefits, risks, interactions, values, ethics, oversight into these systems, etc.

EPAI papers covered different formalisations, methodologies and testbenches for the evaluation of AI systems with the final goal of measuring the field's rates of development, progress, and impact. Two contributions obtained, respectively, the workshop best paper award and runner-up award, and have been extended into full journal papers included in this volume. The best paper, "Artificial Canaries: Early Warning Signs for Anticipatory and Democratic Governance of AI" by Carla Zoe Cremer and Jess Whittlestone, propose a method for identifying early warning signs of transformative progress in AI, and discuss how these can support the anticipatory and democratic governance of AI. Their method combines expert elicitation and collaborative causal graphs to identify key mile-stones and identify the relationships between them.

The runner-up award paper, "Efficient and Robust Model Benchmarks with Item Response Theory and Adaptive Testing", by Hao Song and Peter Flach, investigate adaptive approaches to achieve better efficiency in model benchmarking. In this regard, they propose and analyse methods that allow machine learning practitioners to pick only a few representative datasets to quantify the quality of a technique, from which to extrapolate the performance on other datasets. To this end, they adapt existing approaches from psychometrics, Item Response Theory and Adaptive Testing specifically, by implementing certain modifications following the requirements of machine learning experiments, and present experimental results to validate the approach.

The workshop "Singular Problems for Healthcare (SP4HC)" was devoted to advances in Artificial Intelligence applied to Healthcare and Well-Being, with an active interest in frontier-of-knowledge Machine Learning subjects, sometimes named as singular Machine Learning problems. In particular, imbalanced classification, ordinal classification

or multi-label classification, which are pervasive in important practical problems in healthcare, have consequently generated a tremendous interest. From the scientific point of view, the workshop intended to serve as a basis for proposing and discussing advances in the artificial intelligence arena, with a range of applications. Some contributions have dealt with common challenges in healthcare applications, as imbalanced classes and feature selection using simple interpretable classifiers like logistic models and decision trees. Others used feature extraction from images and artificial neural networks approaches. Natural language processing, reinforcement learning and model-based reinforcement techniques, recommender systems or echo state networks, which are an alternative to standard recurrent neural networks, have also been examined. Finally, several temporal modelling approaches to manage the concept drift phenomenon have been applied for identification and classification tasks. From the medical perspective, the papers of this workshop had coped with different medical topics like melanoma skin cancer detection, wellness application for providing personalized health activities, Type 1 diabetes blood glucose control for insulin dose decisions, and antimicrobial multidrug resistance in Intensive Care Units (ICUs) for their characterization and prediction.

The paper selected as best of this workshop, "Antimicrobial Resistance Prediction in Intensive Care Unit for Pseudomonas Aeruginosa using Temporal Data-Driven Models" proposes new paradigms to address the problem of the increasing bacterial resistance to antibiotics, a particularly serious problem in hospital's ICUs because of the vulnerability of these patients. Knowing in advance whether a concrete bacterium is resistant or susceptible to an antibiotic is a crux step for clinicians to determine an effective antibiotic treatment. This article focuses on cultures of the Pseudomonas Aeruginosa bacterium because is one of the most frequent, dangerous and difficult to treat in the ICU. Several temporal data-driven models are proposed and analysed to predict the resistance or susceptibility to a specific antibiotic family previously to obtain the result of the antimicrobial susceptibility test and only using historical data registered in the electronic health system. The approach provides reasonably accurate results for some antimicrobial families, and could be used by clinicians as an early-warning system to support the election of the antibiotic therapy. This early prediction can save valuable time to start the adequate treatment for an ICU patient.

The Workshop on Data Fusion for Artificial Intelligence (DAFUSAI) was dedicated to discuss this crucial problem from both theoretical and applied point of views. Classiffication, image processing, decision-making, big data or deep learning require collecting data and fusing them in appropriate ways in order to solve specific problems. For this reason, a huge effort is devoted to the developments and analysis of data fusion methods [15]. Aggregation functions are one of the most widely used methods in this sense. They are defined as monotone functions with appropriate boundary conditions and include, among others, most of the means or functions such as the product, the minimum or the maximum. However, in recent years it has been shown that the concept of aggregation function can be too restrictive, as it does not cover some examples which can provide good results in particular applications, as it is the case of the mode. Furthermore, some data fusion functions more general than aggregations,- the so called pre-aggregation functions [16]-, have been proposed to deal with problems ranging from classiffication [17] to the computational brain [18], with very promising results.

From the papers presented at the DAFUSAI Workshop, we have selected the paper by Sicui Zhang et al., entitled "Towards Multi-perspective Conformance Checking with Fuzzy Sets". The paper points out the problem faced by the organizations concerning the necessity to employ data-driven techniques to audit their business processes and ensure they comply with laws and internal/external regulations, which can be inefficient and subject to frauds or abuses. An increasingly popular approach to automatically assess the compliance of the executions of organization processes is represented by alignment-based conformance checking. These techniques present several advantages, e.g., by comparing real process executions with the models, they can show possible discrepancies. However, there are also some drawbacks, e.g., as they perform a crisp evaluation of process compliance, a behavior process is classiffied as either compliant or d-viant (even if such deviation is not severe). In this paper, the authors discuss about these problems, proposing a novel conformance checking approach aimed at representing actors' tolerance with respect to process deviations, taking it into account when assessing the severity of the deviations. Additionally, as a proof of concept, the authors performed a set of synthetic experiments to assess the approach. The obtained results clearly show the advantages of considering a more flexible evaluation of process deviations, and the impact on the quality and the interpretation of the obtained diagnostics.

With this special issue, we have attempted to describe progress in various areas of Artificial Intelligence by outlining some of the most interesting papers presented in several selected workshops of the European Conference in AI in 2020. These papers reveal a mature discipline, stating novel aspects, and opening questions and problems that require equally novel approaches. We hope that its reading can inspire new directions and solutions that can lead us to both theoretical and practical developments, helping us to advance to a future in which humans undoubtedly will be at the center of the technology.

Amparo Alonso-Betanzos[1], Pedro Cabalar[2],
Gracaliz P. Dimuro[10,11], Marcos Garcia[3], José Hernández-Orallo[4],
Raquel Hervás[5], Ángeles Manjarés[6], Fernando Martínez-Plumed[7,4],
Inmaculada Mora-Jiménez[8], Miquel Sànchez-Marré[9]

[1] Research Center on Information and Communication Technologies (CITIC). Universidade da Coruña (Spain)

[2] Universidade da Coruña (Spain)

[3] Research Center in Intelligent Technologies (CiTIUS). Universidade de Santiago de Compostela, Galicia (Spain)

[4] Valencian Research Institute for Artificial Intelligence (VRAIN). Universitat Politècnica de València (Spain)

[5] Facultad de Informática, Instituto de Tecnología del Conocimiento

Universidad Complutense de Madrid (Spain)

[6] Dept. Inteligencia Artificial ETSI, UNED (Spain)

[7] European Commission Joint Research Centre, Seville (Spain)

[8] Department of Signal Theory and Communications, Telematics and Computing Systems. Universidad Rey Juan Carlos, Madrid (Spain)

[9] Intelligent Data Science and Artificial Intelligence Research Centre (IDEAI-UPC) Dept. of Computer Science, Universitat Politècnica de Catalunya, Barcelona (Spain)

[10] Centro de Ciências Computacionais (C3) Universidade Federal do Rio Grande (Brazil)

[11] Departamento de Estadística, Informática y Matemáticas Universidad Pública de Navarra, Pamplona (Spain)

References

[1] UN Global Pulse, "Big data for development: Challenges and opportunities," 2012. [Online].

[2] UN, "Human Rights Based Approach to development," 2003. [Online].

[3] UNICEF Global Innovation Centre, "Generation AI," 2003. [Online].

[4] World Wide Web Foundation, "Artificial Intelligence: The Road Ahead in Low and Middle-Income Countries," World Wide Web Foundation, 2017. [Online].

[5] International Telecommunications Union, "AI for good global summit 2018," International Telecommunication Union, Geneva, 2018. [Online].

[6] World Economic Forum, "Harnessing artificial intelligence for the earth," 2018. [Online].

[7] R. Ochigame, "The invention of "Ethical AI"," 2019. [Online].

[8] Y. Benkler, "Don't let industry write the rules for AI," *Nature*, vol. 569, no. 7754, pp. 161–162, 2019.

[9] J. Pennington, R. Socher, C. D. Manning, "Glove: Global vectors for word representation," in *Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543.

[10] T. Mikolov, K. Chen, G. Corrado, J. Dean, "Efficient estimation of word representations in vector space," in *Workshop Proceedings of the International Conference on Learning Representations (ICLR) 2013*, 2013. arXiv preprint arXiv:1301.3781.

[11] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Minneapolis, Minnesota, June 2019, pp. 4171–4186, Association for Computational Linguistics.

[12] F. Martínez-Plumed, J. Hernández-Orallo, E. Gómez, "Tracking AI: The capability is (not) near," *Frontiers in Artificial Intelligence and Applications*, vol. 320: ECAI 2020, pp. 2915–2916, 2020.

[13] J. Hernández-Orallo, M. Baroni, J. Bieger, N. Chmait, D. L. Dowe, K. Hofmann, F. Martínez-Plumed, C. Strannegård, K. R. Thórisson, "A new AI evaluation cosmos: Ready to play the game?," *AI Magazine*, vol. 38, no. 3, pp. 66–69, 2017.

[14] F. Martínez-Plumed, E. Gómez, J. Hernández-Orallo, "Futures of artificial intelligence through technology readiness levels," *Telematics and Informatics*, vol. 58, p. 101525, 2021.

[15] G. Pereira Dimuro, J. Fernández, B. in Bedregal, R. Mesiar, J. Antonio Sanz, G. Lucca, H. Bustince, "The state-of-art of the generalizations of the Choquet integral: From aggregation and pre-aggregation to ordered directionally monotone functions," *Information Fusion*, vol. 57, pp. 27 – 43, 2020, doi: https://doi.org/10.1016/j.inffus.2019.10.005.

[16] G. Lucca, J. Sanz, G. Pereira Dimuro, B. Bedregal, R. Mesiar, A. Kolesárová, H. Bustince Sola, "Pre-aggregation functions: construction and an application," *IEEE Transactions on Fuzzy Systems*, vol. 24, pp. 260–272, April 2016, doi: 10.1109/TFUZZ.2015.2453020.

[17] M. Elkano, M. Galar, J. A. Sanz, P. F. Schiavo, S. Pereira, G. caliz Pereira Dimuro, E. N. Borges, H. Bustince, "Consensus via penalty functions for decision making in ensembles in fuzzy rule-based classification systems," *Applied Soft Computing*, vol. 67, pp. 728 –740, 2018, doi: https://doi.org/10.1016/j.asoc.2017.05.050.

[18] L. Ko, Y. Lu, H. Bustince, Y. Chang, Y. Chang, J. Ferandez, Y. Wang, J. A. Sanz, G. Pereira Dimuro, C. Lin, "Multimodal fuzzy fusion for enhancing the motor-imagery-based brain computer interface," *IEEE Computational Intelligence Magazine*, vol. 14, no. 1, pp. 96–106, 2019, doi: 10.1109/MCI.2018.2881647.