Guest Editorial:
Special Section on New Frontiers in Computing for Next-Generation Healthcare Systems

Advances in information and communication technology (ICT) are driving a profound transformation in the healthcare sector. The traditional healthcare has evolved to smart healthcare. The prominent role played by technology in the fight against the novel Coronavirus disease 2019 (COVID-19) pandemic, from artificial intelligence to telehealth services, from contact tracing to 3D printing, is a driving force towards increased digitalization of healthcare services. This Special Section aims at providing a comprehensive picture of the technological advancements in computing that are underpinning the current revolution in healthcare. It is dedicated to emerging technological trends and their potential implications for improving healthcare interventions and tackling the challenges of an increasingly aging population.

The increasing use of mobile devices, wearable sensors, smart devices, etc. are shifting the healthcare provision model from a centralized, standardized, hospital-centered, episodic model, to a decentralized, home-based, personalized, patient-centered, and continuous paradigm, pushing and redefining the boundary of traditional Health Informatics systems. The availability of pervasive and detailed health information brings forth new opportunities for the treatment and the management of chronic illnesses and the prevention of lifestyle-related diseases. Social media enables extensive mining of healthcare-related data, from the study of social networks in relation to epidemics and infectious diseases, to the promotion of mental healthcare and wellbeing. Advances in mHealth (mobile Health), pHealth (pervasive Health), sHealth (smart Health) have led to an exponentially increasing amount of healthcare-related data, both within and outside of hospital boundaries. Several submissions to this Special Section reflect the expanding use of artificial intelligence and data analytics enabled by far-reaching data collection.

The adoption of assistive and immersive technologies including personal assistants, chatbots, recommender systems, virtual and augmented reality open new opportunities to improve the life and well-being of the elderly, the disable and those suffering from chronic and mental illnesses and empower people to take charge of their own health and well-being.

It must be noted that applying computing technologies in healthcare introduces specific and strict requirements. Ensuring the proper, meaningful, trusted, and safe use of devices and computing technologies by the users is needed in order to minimize risks for their health, and requires novel and creative approaches to technology engineering, testing and validation. Another major concern is the protection of patients' privacy and more generally, the potential misuse of health-related personal data. In this context, research is needed to ensure an appropriate balance between protecting individuals' rights and reaping the benefits of data sharing.

With the above thoughts, the current Special Issue was proposed for the *IEEE Transactions on Emerging Topics in Computing (TETC)*. The call received a significant number of submissions, which were reviewed rigorously by a large set of reviewers across the globe. Thus, it represents a collective effort from the research community and industry participants on an international scale. From the many excellent submissions received, a set of manuscripts were selected to appear in this Special Section. The works address different issues related to computing applications in the healthcare, and can be grouped in 3 broad categories:

(1) computing aspects of pervasive and ubiquitous healthcare systems
(2) semantic computing applications and methods, and
(3) technologies, methods, and applications empowering user privacy, security and safety.

Considering *computing aspects of pervasive and ubiquitous healthcare systems*, in "IoTility: Architectural Requirements for Enabling Health IoT Ecosystems" by Wyatt Lindquist, Sumi Helal, Ahmed Khaled and Wesley Hutchinson, the architectural design of IoT applications tailored to the healthcare scenario is analyzed. Indeed, Health IoT presents specific requirements that are not addressed in traditional IoT systems. These include combining data on patients' health acquired by multiple sensors, detecting sensors' malfunctions or improper use by the patient, or the necessity of sharing the same device by multiple users. To address these specific needs, the authors propose the IoTility architecture especially targeted to Health IoT. This includes the Atlas Thing architecture, focusing on communication between *things*, the IoT Device Description Language (IoT-DDL), and the Inter-Thing Relationships framework, defining a concrete set of relationships between *things*. The proposed architecture supports developers in analyzing existing services and inferring new opportunities they may provide. Furthermore, it simplifies the development of applications interacting with data provided by Health IoT devices.

The massive deployment of IoT devices opens new opportunity for the management of chronic diseases by collecting far-reaching and punctual data in real-life conditions. In "Collaborative Trajectory Mining in Smart-homes to Support Early Diagnosis of Cognitive Decline" by Elham Khodabandehloo and Daniele Riboni, a system to recognize symptoms of cognitive decline by relying on smart-home sensors is presented. Moving from the assumption that cognitive issues are frequently associated to locomotion anomalies, the authors propose a system to detect wandering episodes in a challenging indoor environment, where several

movement patterns associated to everyday activities may resemble wandering.

The proposed collaborative trajectory mining system collects anonymous trajectories from multiple instrumented smart-homes and trains a personalized model, fine-tuned to the characteristics of each individual and their home environments. Experimental results on more than 100 subjects from the CASAS study shows promise in the early recognition of cognitive decline.

The paper "DiabDeep: Pervasive Diabetes Diagnosis based on Wearable Medical Sensors and Efficient Neural Networks" by Hongxu Yin, Bilal Mukadam, Xiaoliang Dai and Niraj K. Jha, tackles the important issue of how to train and deploy deep learning (DL) networks on resource-constrained edge devices. Although the state-of-the-art in most pattern recognition problems, large and over-parameterized deep neural networks are computationally intensive and require expensive GPUs for fast inference. Besides advances in dedicated hardware, from a computing perspective many strategies are being explored to reduce the computational burden. Leveraging an efficient grow-and-prune training algorithm, the system by Yin *et al.* achieves a sparse architecture that supports both server-side and edge-side inference with minimal loss of accuracy. The proposed methodology was successfully applied to the detection of diabetes from off-the-shelf wearable medical sensors. The accuracy of the system in classifying diabetics against healthy individuals is 96.3% (server) and 95.3% (edge), whereas, for distinguishing among type-1 diabetic, type-2 diabetic and healthy individual accuracy values are 95.7% (server) and 94.6% (edge).

Concerning *semantic computing applications and methods*, in "A Mental Health Chatbot for Regulating Emotions (SERMO) - Concept and Usability Test" Kerstin Denecke, Sayan Vaaheesan, and Aaganya Arulnathan present a conversational agent designed to support people suffering from depression and mood disorders in regulating emotions. Users can describe in natural language events that occurred during the day. The system, leveraging natural language processing (NLP) techniques, recognizes the users' emotions and recommends beneficial activities such as mindfulness or relaxation exercises.

The work entitled "Recommending Activities for Mental Health and Well-being: Insights from Two User Studies" by Darius A. Rohani, Aaron Springer, Victoria Hollis, Jakob E. Bardram and Steve Whittaker also presents a recommender system to support emotional well-being of patients. Different from the previous work whose focus was on the detection of users' emotions, this work focuses on the identification of those activities that could have the higher impact on a user's mood. The authors collected activities from two populations (a clinical sample diagnosed with mood disorder and a non-clinical sample) and labeled them based on category and impact on mood. To predict the outcome on a user's mood, the authors exploited a probabilistic multinomial Naïve Bayes (NB) Classifier and a Support Vector Machine (SVM). The evaluation shows that the personalized model provided by the system outperformed the generalized model in both populations. Collectively, both works highlighted difficulties in evaluating the effectiveness of the proposed technologies and assessing how they can effectively empower the user and achieve a positive impact on relevant health outcomes, especially in susceptible and at-risk populations.

An application of semantic computing in the more traditional hospital setting is presented in "FarSight: Long-Term Disease Prediction Using Unstructured Clinical Nursing Notes", by Tushaar Gangavarapu, Gokul S Krishnan, Sowmya Kamath S, and Jayakumar Jeganathan. The authors move from the assumption that unstructured clinical nursing notes contain valuable information on patients' state, complementary or even richer than traditional structured Electronic Health Records (HER). Hence, they present an aggregation mechanism intended to recognize the onset of a disease with the earliest detected symptoms from clinical nursing notes. To model patients' information, the authors rely on document embeddings, topic modeling and multiple state-of-the-art DL models for classification. On the clinical task of ICD-9 code group prediction in the MIMIC-III benchmark, the proposed model achieves 0.72 AUPRC and 0.82 AUROC, improving the EHR based state-of-the-art models by 19.34% and 5.41%, respectively.

The last three works of this Special Section deal with *technologies, methods, and applications empowering user privacy, security, and safety*. The role of federated learning in preserving patients' privacy is explored in "Privacy-Preserving Deep Learning NLP Models for Cancer Registries", by Mohammed Alawad, Hong-Jun Yoon, Shang Gao, Brent Mumphrey, Xiao-Cheng Wu, Eric B. Durbin, Jong Cheol Jeong, Isaac Hands, David Rust, Linda Coyle, Lynne Penberthy, and Georgia Tourassi. Their goal is to extract characteristics from pathology reports using deep learning. In this context, although collaboration among cancer registries is essential, data privacy and confidentiality concerns prevent cancer registries from sharing patient data and benefiting from each other's knowledge. The premise of federated learning is that each site trains the model on the local, private dataset, and then only the trained model parameters are shared. In spite of that, the authors show how private information may still be accidentally leaked, e.g., by embedding them in the source data dictionary. To overcome this issue, they rely on public, anonymized data to build a shared data dictionary. The proposed DL NLP approach is used to distribute a multi-task DL model among cancer registries, to extract cancer characteristics from pathology reports. Results show how federated learning outperforms significantly the single registry model, and achieves a comparable performance as centralized training.

Another work dealing with the issue of privacy preservation is "Differential Privacy-based Genetic Matching in Personalized Medicine", by Jianhao Wei, Yaping Lin, Xin Yao, Jin Zhang, Xinbo Liu. In this paper, the authors propose a differential privacy-based genetic matching (DPGM) scheme. Genetic matching, i.e., searching for genomes that have segments that are identical or nearly identical to segments of the user's genome, has important applications in personalized medicine, such as detecting mutated genes or genetic diseases. The genetic testing industry, especially cloud-based services, has grown rapidly in the past few years, holding a considerable amount of private genetic data. Outsourced genetic data has important privacy implications as gene sequences contain sensitive information related to the users as well as their

relatives. Differential privacy (DP) is a methodology that leverages added noise to protect individuals' privacy when publishing or analyzing aggregated data from a population, with a minimal loss in accuracy. The scheme proposed by the authors first exploits a DP-based algorithm to construct a published sequence that contains significantly noisy Single-Nucleotide Polymorphisms (SNPs) associated with diseases, thereby ensuring outsourced genetic data privacy. Then, a DP-based N-order Markov algorithm is exploited to generate a noisy query sequence. Finally, the longest common subsequence based on dynamic programming algorithm is calculated. The proposed scheme achieves ε-differential privacy and high efficiency and data utility.

Finally, in "Access Control for Implantable Medical Devices", by Carmen Camara, Pedro Peris-Lopez, Jose Maria de Fuentes, and Samuel Marchal, the authors address security aspects of cardiac implants. Modern implantable medical devices allow remote access and re-programming, which is highly efficient but opens the door to unauthorized accesses and possible attacks. An emerging solution to this problem involves the exploitation of physiological signals, such as the ECG, to ensure accurate identity verification. Current solutions usually exploit only a single feature, such as the R peak, and are still vulnerable to attacks. This work combines two security mechanisms, namely, identity verification, based on the entire physiological signal, and proximity verification, by measuring the similarity between an internal and external physiological signal. The proposed system was evaluated using ECG signals of 199 individuals recorded over 24 hours while considering three adversary strategies, resulting in 92.92% accuracy.

The guest editors thank the reviewers for their valued time, expertise, and constructive feedback in their reviews. We also thank all the authors for their submissions and their accommodation of the publication deadlines and constraints. Finally, we would also like to thank the two consecutive Editors-in-Chief of IEEE TETC, Prof. Paolo Montuschi and Prof. Cecilia Metra, whose help made this Special Section possible.

LIA MORRA, *Guest Editor*
Politecnico di Torino, Italy.

VALENTINA GATTESCHI, *Guest Editor*
Politecnico di Torino, Italy.

SARAJU P. MOHANTY, *Guest Editor*
University of North Texas, Denton, USA.

YUAN-HAO CHANG, *Corresponding Editor*
Academia Sinica, Taipei, Taiwan.

Lia Morra (M'18-SM'19) received the M.Sc. and the Ph.D. degrees in computer engineering from Politecnico di Torino in 2002 and 2006, respectively. Currently, she is post-doctoral Associate at the Department of Control and Computer Engineering of Politecnico di Torino, Italy. From 2006 to 2016 she joined im3D, where she served as Chief Scientific Officer from 2014 to 2017; she guided the team developing the core medical image analysis and artificial intelligence technology components from design to regulatory approval. Her research interests include computer vision, pattern recognition, machine learning and computer-assisted medical image interpretation. She has authored numerous papers in international journals and conferences and was awarded three patents. She is a Senior Member of IEEE, and member of the AAPM Computer Aided Image Analysis Subcommittee.

Valentina Gatteschi (M'19-SM'19) received her B.Sc. and M.Sc. degrees in management engineering and the Ph.D. degree in computer engineering in 2005, 2008 and 2013, respectively. She currently is an Assistant Professor with time contract at the Department of Control and Computer Engineering of Politecnico di Torino, Italy. Her main interests are natural language processing, human-computer interaction, intelligent systems and blockchain. She has authored numerous papers in international journals and conferences. She has been involved in a number of European projects. She has been an invited speaker at several conferences/workshops on blockchain and is currently serving as a member of the IEEE Blockchain Education Subcommittee.

Saraju P. Mohanty (M'99-SM'08) received the bachelor's degree (Honors) in electrical engineering from the Orissa University of Agriculture and Technology, Bhubaneswar, in 1995, the masters degree in Systems Science and Automation from the Indian Institute of Science, Bengaluru, in 1999, and the Ph.D. degree in Computer Science and Engineering from the University of South Florida, Tampa, in 2003. He is a Professor with the University of North Texas. His research is in "Smart Electronic Systems" which has been funded by National Science Foundations (NSF), Semiconductor Research Corporation (SRC), U.S. Air Force, IUSSTF, and Mission Innovation. He has authored 300 research articles, 4 books, and invented 4 U.S. patents. His Google Scholar h-index is 35 and i10-index is 130 with 5600+ citations. He has over 20 years of research experience on security and protection of media, hardware, and system. He introduced the Secure Digital Camera (SDC) in 2004 with built-in security features designed using Hardware-Assisted Security (HAS) or Security by Design (SbD) principle. He is widely credited as the designer for the first digital watermarking chip in 2004 and first the low-power digital watermarking chip in 2006. He is a recipient of 12 best paper awards, IEEE Consumer Electronics Society Outstanding Service Award in 2020, the IEEE-CS-TCVLSI Distinguished Leadership Award in 2018, and the PROSE Award for Best Textbook in Physical Sciences and Mathematics category from the Association of American Publishers in 2016 for his Mixed-Signal System Design book published by McGraw-Hill. He has delivered 9 keynotes and served on 5 panels at various International Conferences. He has been serving on the editorial board of several peer-reviewed international journals, including IEEE Transactions on Consumer Electronics (TCE), and IEEE Transactions on Big Data (TBD). He is the Editor-in-Chief (EiC) of the IEEE Consumer Electronics Magazine (MCE). He has mentored 2 post-doctoral researchers and supervised 12 Ph.D. dissertations and 26 M.S. theses.