

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 subseries at <http://www.springer.com/series/7409>


K. Selçuk Candan · Bogdan Ionescu ·
Lorraine Goeuriot · Birger Larsen ·
Henning Müller · Alexis Joly ·
Maria Maistro · Florina Piroi ·
Guglielmo Faggioli · Nicola Ferro (Eds.)


Experimental IR Meets Multilinguality, Multimodality, and Interaction

12th International Conference of the CLEF Association, CLEF 2021
Virtual Event, September 21–24, 2021
Proceedings


Editors

K. Selçuk Candan 
Arizona State University
Tempe, AZ, USA

Lorraine Goeuriot 
Université Grenoble Alpes
Saint-Martin-d'Hères, France


Henning Müller 
HES-SO Valais-Wallis
Sierre, Switzerland


Maria Maistro 
University of Copenhagen
Copenhagen, Denmark


Guglielmo Faggioli 
University of Padua
Padova, Italy

Bogdan Ionescu
Politehnica University of Bucharest
Bucharest, Romania

Birger Larsen 
Aalborg University Copenhagen
Copenhagen, Denmark

Alexis Joly 
University of Montpellier
Montpellier, France

Florina Piroi 
TU Wien
Vienna, Austria

Nicola Ferro 
University of Padua
Padova, Italy

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-85250-4 ISBN 978-3-030-85251-1 (eBook)
<https://doi.org/10.1007/978-3-030-85251-1>

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

© Springer Nature Switzerland AG 2021

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

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

Since 2000, the Conference and Labs of the Evaluation Forum (CLEF) has played a leading role in stimulating research and innovation in the domain of multimodal and multilingual information access. Initially founded as the Cross-Language Evaluation Forum and running in conjunction with the European Conference on Digital Libraries (ECDL/TPDL), CLEF became a standalone event in 2010 combining a peer-reviewed conference with a multi-track evaluation forum. The combination of the scientific program and the track-based evaluations at the CLEF conference creates a unique platform to explore information access from different perspectives, in any modality and language.

The CLEF conference has a clear focus on experimental information retrieval (IR) as seen in evaluation forums (like CLEF Labs, TREC, NTCIR, FIRE, MediaEval, RomIP, TAC) with special attention to the challenges of multimodality, multilinguality, and interactive search ranging from unstructured to semi-structured and structured data. The CLEF conference invites submissions on new insights demonstrated by the use of innovative IR evaluation tasks or in the analysis of IR test collections and evaluation measures, as well as on concrete proposals to push the boundaries of the Cranfield/TREC/CLEF paradigm.

CLEF 2021¹ was organized by the University “Politehnica” of Bucharest, Romania, during September 21–24, 2021. The continued outbreak of the COVID-19 pandemic affected the organization of CLEF 2021. The CLEF steering committee along with the organizers of CLEF 2021, after detailed discussions, decided to run the conference fully virtually. The conference format remained the same as in past years and consisted of keynotes, contributed papers, lab sessions, and poster sessions, including reports from other benchmarking initiatives from around the world. All sessions were organized and run online.

CLEF 2021 continued the initiative introduced in the 2019 edition during which the European Conference for Information Retrieval (ECIR) and CLEF joined forces: ECIR 2021 hosted a special session dedicated to CLEF Labs where lab organizers present the major outcomes of their labs and their plans for ongoing activities, followed by a poster session to favour discussion during the conference. This was reflected in the ECIR 2021 proceedings, where CLEF Lab activities and results were reported as short papers. The goal was not only to engage the ECIR community in CLEF activities but also to disseminate the research results achieved during CLEF evaluation cycles as submission of papers to ECIR.

¹ <http://clef2021.clef-initiative.eu/>.

Naila Murray (Facebook AI Research) was invited to deliver a keynote talk about video understanding and multi-modal search. Further keynote talks were in the process of being confirmed and were not final at the writing of this editorial.

CLEF 2021 received a total of 21 scientific submissions, of which a total of 11 papers (10 long, one short) were accepted. Each submission was reviewed by three Program Committee members, and the program chairs oversaw the reviewing and follow-up discussions. Twelve countries are represented in the accepted papers, as several of them were a product of international collaboration. This year, researchers addressed the following important challenges in the community: application of neural methods for entity recognition as well as misinformation detection in the health area, skills extraction in job-match databases, stock market prediction using financial news, and extraction of audio features for podcast retrieval. Evaluation remains a strong interest with papers on the evaluation of 1) pseudo-relevance feedback based on web-based data enrichment, 2) evolving datasets using pivot systems, and 3) multitask learning models for relevance assessment. Creating shareable open datasets is also a strong focus this year with datasets or frameworks created for 1) linguistic uncertainty in NLP, 2) an Italian corpus for subjectivity detection in newspapers, and 3) personalized type-based facet ranking tasks.

Like in previous editions since 2015, CLEF 2021 invited CLEF lab organizers to nominate a “best of the labs” paper that was reviewed as a full paper submission to the CLEF 2021 conference according to the same review criteria and PC. Five full papers were accepted for this “best of the labs” section.

The conference integrated a series of workshops presenting the results of lab-based comparative evaluations. CLEF 2021 was the 12th year of the CLEF conference and the 22nd year of the CLEF initiative as a forum for IR evaluation. A total of 15 lab proposals were received and evaluated in peer review based on their innovation potential and the quality of the resources created. The 12 selected labs represented scientific challenges based on new datasets and real world problems in multimodal and multilingual information access. These datasets provide unique opportunities for scientists to explore collections, to develop solutions for these problems, to receive feedback on the performance of their solutions, and to discuss the challenges with peers at the workshops. In addition to these workshops, the labs reported results of their year long activities in overview talks and lab sessions. Overview papers describing each of the labs are provided in this volume. The full details for each lab are contained in a separate publication, the Working Notes².

The 12 labs running as part of CLEF 2021 comprised mainly labs that continued from previous editions at CLEF (ARQMath, BioASQ, CheckThat!, CheMU, CLEF eHealth, eRisk, ImageCLEF, LifeCLEF, Lilas, PAN, and Touché) along with a new pilot/workshop activity (SimpleText). In the following we give a few details for each of the labs organized at CLEF 2021 (ordered alphabetically):

² Faggioli, G., Ferro, N., Joly, A., Maistro, M., and Piroi, F. editors (2021). *CLEF 2021 Working Notes*. CEUR Workshop Proceedings (CEUR-WS.org), ISSN 1613-0073.

ARQMath: Answer Retrieval for Mathematical Questions³ considers the problem of finding answers to new mathematical questions among posted answers on the community question answering site Math Stack Exchange. The goals of the lab were to develop methods for mathematical information retrieval based on both text and formula analysis. Objectives to reach these goals include creating test collections for training and evaluating Math IR systems, establishing a state-of-the-art set of retrieval solutions on these test collections to be used as future baselines, and promoting Math IR to the research community. Compared to the 2020 ARQMath edition, this year the test collection size has doubled, the same being observed in the number of participants.

BioASQ⁴ challenges researchers with large-scale biomedical semantic indexing and question answering (QA). The challenges include tasks relevant to hierarchical text classification, machine learning, information retrieval, QA from texts and structured data, multi-document summarization, and many other areas. The aim of the BioASQ workshop is to push the research frontier towards systems that use the diverse and voluminous information available online to respond directly to the information needs of biomedical scientists. Four tasks were organized in 2021, two on biomedical semantic indexing, one on QA, and a new task on COVID-19 QA. For the Spanish semantic indexing task, a continuation of a task introduced in previous years, new data was added that contained Spanish clinical trials and Spanish patents. Participant systems generally outperformed strong baselines, with participant solution clearly shifting towards the use of deep neural approaches. As such approaches, however, necessitate large amounts of training data, which were not available for BioASQ tasks, participants were investigating knowledge and model transfer from other resources.

CheckThat!: Detecting Check-Worthy Claims, Previously Fact-Checked Claims, and Fake News⁵ aims to foster the development of technologies capable of both spotting and verifying check-worthy claims in short messages and political debates in various languages. This year there were three main shared tasks where participants were to estimate the check-worthiness of a claim in a short message (tweet) and in political texts (debates/speeches), to check if a detected claim was previously verified, retrieve evidence to fact-check a claim, and verify the factuality of a claim. The data in 2021 included more languages than in the previous year, with some teams addressing the challenges for all languages, while others tackling one language only.

ChEMU: Cheminformatics Elsevier Melbourne University⁶ proposes two key information extraction tasks over chemical reactions from patent texts. The ChEMU corpus builds on the one used in the previous lab edition, being extended to provide data for two distinct 2021 tasks: reference resolution for chemical reactions and anaphora resolution to identify relationships (i.e. coreference and bridging relationships) between expressions in descriptions of chemical reactions. Out of 19 originally registered teams, only two managed to submit experiments. The tasks proved to be complex, the submitted experiments barely over-performing the baseline results.

³ <https://www.cs.rit.edu/dprl/ARQMath>.

⁴ <http://www.bioasq.org/workshop2021>.

⁵ <https://sites.google.com/view/clef2021-checkthat>.

⁶ <http://chemu2021.eng.unimelb.edu.au/>.

CLEF eHealth⁷ aims to support the development of techniques to aid laypeople, clinicians, and policy-makers in easily retrieving and making sense of medical content to support their decision making. The goals of the lab are to develop processing methods and resources in a multilingual setting to enrich difficult-to-understand eHealth texts and provide valuable documentation. Organized since 2012, the CLEF eHealth labs have provided a recurring contribution to the creation and dissemination of text analytics resources, methods, test collection, and evaluation benchmarks that support both medical professionals and laypersons when dealing with health-related information. The 2021 CLEF eHealth edition organized two tasks. The first one was a multilingual Information Extraction task, focusing on Spanish language ultrasound reports. The second task, Consumer Health Search, was a continuation of previous CLEF eHealth IR tasks with a new representative web corpus and layperson medical queries. From the 67 teams that originally registered, 11 of them submitted runs to the two tasks.

eRisk: Early Risk Prediction on the Internet⁸ explores challenges of evaluation methodology, effectiveness metrics, and other processes related to early mental health risk detection. Early detection technologies can be employed in different areas, particularly those related to health and safety. Over the years these evaluations have taken place, it has become evident that the interplay between psychological disorders and the users' expression through language is a very challenging task, with currently available solutions not reaching satisfactory performance levels. The 2021 edition of the lab contained three tasks, two being continuations of tasks organized in the previous years (self-harm and depression severity detection), and a new one on the topic of pathological gambling. The data provided to campaign participants consisted of texts written in social media. From the 76 teams that originally registered to this lab, 18 had submitted experiments, with a total of 117 runs (26 for Task 1, 55 for Task 2, and 36 for Task 3).

ImageCLEF: Multimedia Retrieval⁹ provides an evaluation forum for visual media analysis, indexing, classification/learning, and retrieval in medical, nature, security, and lifelogging applications with a focus on multimodal data, that is data from a variety of sources and media. The 2021 ImageCLEF edition consisted of four main tasks dedicated to multimedia retrieval in four areas: medical, nature, identification of hand-drawn components, and social media, with the latter being newly introduced this year. The first task consisted of three subtasks related to radiology images (visual question answering, CT-based tuberculosis evaluation, and captioning concepts across radiology images). The nature-related task contained training and test data to form 3D reconstructions of coral environments. The task on hand-drawn images focused on user interface drawings as well as screenshot images (new this year) which, by segmentation and labeling steps, are to provide additional support for code developers. The social media-related task aimed to assess the vulnerability potential and real-life effects of users sharing personal visual data. 42 participating groups submitted over 250

⁷ <https://clefehealth.imag.fr/>.

⁸ <https://erisk.irlab.org/>.

⁹ <https://www.imageclef.org/2021>.

experiments to these tasks, with results varying in performance improvements over previous task editions.

LifeCLEF: Multimedia Life Species Identification¹⁰ aims at boosting research on the identification and prediction of living organisms in order to solve the taxonomic gap and improve our knowledge of biodiversity. Through its biodiversity informatics related challenges, LifeCLEF is intended to push the boundaries of the state of the art in several research directions at the frontier of multimedia information retrieval, machine learning, and knowledge engineering. LifeCLEF in 2021 organized four challenges (PlantCLEF, BirdCLEF, GeoLifeCLEF, and SnakeCLEF) involving image data, audio data, and geolocations. In terms of participating teams that submitted runs, the BirdCLEF task (a bird sound recognition task) stands out with over 800 teams submitting experimental results. The main LifeCLEF outcome is that, taken together, the solutions used by the participants to solve the lab tasks provide a new snapshot of state-of-the-art systems' performances in computer vision, audio analysis techniques, and machine learning algorithms that can be part of a real-world biodiversity monitoring system.

LiLAS: Living Labs for Academic Search¹¹ aims to bring together researchers interested in the online evaluation of academic search systems. The long term goal is to foster knowledge on improving the search for academic resources like literature, research data, and the connections between these resources in fields from the life sciences and the social sciences. The immediate goal of this lab is to develop ideas, best practices, and guidelines for a full online evaluation campaign at CLEF 2021. The first LiLAS iteration as a workshop-lab provided participants exclusive access to real-world academic data search systems, LIVIVO for scientific literature search and GESIS Search for data sets and open access publication search, for each of which a use case was defined. STELLA was introduced as the living lab framework to assess participant submissions which were provided either as static search results sets or as Docker images to be integrated in the live search systems. Nine experimental systems were evaluated with metrics designed for assessing interleaved results, combining results from the participants with baseline results provided by the search systems.

PAN: Digital Text Forensics and Stylometry¹² is a networking initiative for digital text forensics, where researchers and practitioners study technologies that analyze texts with regard to originality, authorship, and trustworthiness. PAN provides evaluation resources consisting of large-scale corpora, performance measures, and web services that allow for meaningful evaluations. The main goal is to provide for sustainable and reproducible evaluations, to get a clear view of the capabilities of state-of-the-art algorithms. This year, PAN organized three shared tasks: detecting authors of hate speech spreaders, authorship verification, and multi-author writing style analysis. Each of the tasks made use of its own specifically designed collection of documents. For the first task, focusing on profiling hate speech spreaders, a data set of social media postings (i.e. Twitter) was created, with manually annotated tweets as hater/not-hater labels. For the second task, authorship verification, the lab organizers

¹⁰ <https://www.imageclef.org/LifeCLEF2021>.

¹¹ <https://clef-lilas.github.io/>.

¹² <http://pan.webis.de/>.

aimed for a scaled up benchmark setting using fan-fiction literature. Finally, the multi-author style analysis task, a task that evolved over the years PAN was organized as a lab, used a collection of Q&A postings from StackExchange where paragraphs from different answers were joined into one text, thus creating a document with multiple authors.

SimpleText: (Re)Telling Scientific Stories to Non-specialists via Text Simplification¹³ aims to create a community interested in generating a simplified summary of scientific documents and to contribute in making the science really open and accessible for everyone. The goal is to generate a simplified abstract of multiple scientific documents based on a given query. SimpleText was organized as a workshop which discussed three pilot tasks on text simplification for scientific information access, all contributing steps towards arriving at a simplified text summary of an input scientific text. The first pilot task addressed the passage selection challenge, i.e. which parts of a document are appropriate for inclusion into a simplified summary. The second pilot task aimed to decide which terms in a selected passage require a simplifying explanation and contextualisation. Finally, the last pilot task discussed aimed to obtain simplified text passages derived from input scientific text passages. The document collection used for these tasks was compiled from preprint and open access repositories, Wikipedia, and science journalism article resources.

Touché: Argument Retrieval¹⁴ is the first shared task on the topic of argument retrieval. Decision making processes, be it at the societal or at the personal level, eventually come to a point where one side will challenge the other with a “why” question, which is a prompt to justify one’s stance. Thus, technologies for argument mining and argumentation processing are maturing at a rapid pace, giving rise for the first time to argument retrieval. In its second year, Touché has organized two shared tasks: an argument retrieval for controversial questions task and an argument retrieval for comparative questions task. The two tasks used different document collections: for the first task the args.me corpus was provided, while for the second one argument retrieval was performed on the ClueWeb12 collection. Out of 36 registered teams, 27 sent in their retrieval experiments, where relevance judgements from the 2020 lab edition could be used for training.

As a group, the 152 lab organizers were based in 22 countries, with Germany, and France leading the distribution. Despite CLEF’s traditionally Europe-based audience, 44 (28.9%) organizers were affiliated with international institutions outside of Europe. The gender distribution was biased towards 75% male organizers.

The success of CLEF 2021 would not have been possible without the huge effort of several people and organizations, including the CLEF Association¹⁵, the Program

¹³ <https://www.irit.fr/simpleText/>.

¹⁴ <https://touche.webis.de/>.

¹⁵ <http://www.clef-initiative.eu/association>.

Committee, the Lab Organizing Committee, the reviewers, and the many students and volunteers who contributed.

Finally, we thank the generous support of the H2020 AI4Media project, Facebook AI Research, Keysight Technologies Romania, and Siemens Romania R&D, who provided general funding support.

July 2021

K. Selçuk Candan
Bogdan Ionescu
Lorraine Goeuriot
Birger Larsen
Henning Müller
Alexis Joly
Maria Maistro
Florina Piroi
Guglielmo Faggioli
Nicola Ferro

Organization

CLEF 2021, Conference and Labs of the Evaluation Forum – Experimental IR meets Multilinguality, Multimodality, and Interaction, was hosted (online) by the University “Politehnica” of Bucharest, Romania.

General Chairs

K. Selçuk Candan	Arizona State University, USA
Bogdan Ionescu	University “Politehnica” of Bucharest, Romania

Program Chairs

Lorraine Goeuriot	Université Grenoble Alpes, France
Birger Larsen	Aalborg University, Denmark
Henning Müller	University of Applied Sciences Western Switzerland (HES-SO), Switzerland

Lab Chairs

Alexis Joly	Inria Sophia-Antipolis, France
Maria Maistro	University of Copenhagen, Denmark
Florina Piroi	Vienna University of Technology, Austria

Lab Mentorship Chair

Lorraine Goeuriot	Université Grenoble Alpes, France
-------------------	-----------------------------------

Publicity Chairs

Liviu-Daniel Ștefan	University “Politehnica” of Bucharest, Romania
Mihai Dogariu	University “Politehnica” of Bucharest, Romania

Outreach Program Chairs

Yu-Gang Jiang	Fudan University, China - Asian Liaison
Hugo Jair Escalante	Instituto Nacional de Astrofisica, Optica y Electronica, Mexico - Central American Liaison
Fabio A. Gonzalez	National University of Colombia, Colombia - South American Liaison

Ben Herbst	Praelexis, South Africa - African Liaison
Abdulmotaleb El Saddik	University of Ottawa, Canada - North American Liaison

Industry and Sponsorship Chairs

Şeila Abdulamit	Vodafone, Romania
Mihai-Gabriel Constantin	University “Politehnica” of Bucharest, Romania
Bogdan Boteanu	University “Politehnica” of Bucharest, Romania

Website and Social Media Chair

Denisa Ionaşcu	University “Politehnica” of Bucharest, Romania
----------------	--

Finance Chair

Ion Marghescu	University “Politehnica” of Bucharest, Romania
---------------	--

Proceedings Chairs

Guglielmo Faggioli	University of Padua, Italy
Nicola Ferro	University of Padua, Italy

Program Committee

Martin Braschler	ZHAW, Switzerland
Fabio Crestani	University of Lugano, Switzerland
Elöd Egyed-Zsigmond	LIRIS, France
Sebastien Fournier	LSIS, France
Norbert Fuhr	University of Duisburg-Essen, Germany
Teresa Goncalves	University of Evora, Portugal
Gareth Jones	Dublin City University, Ireland
Jaap Kamps	University of Amsterdam, The Netherlands
Jussi Karlgren	Spotify, UK
Liadh Kelly	Maynooth University, Ireland
Christina Lioma	University of Copenhagen, Denmark
David Losada	University of Santiago de Compostela, Spain
Josiane Mothe	IRIT, France
Irina Ovchinnikov	Haifa University, Israel
Martin Potthast	Leipzig University, Germany
Paolo Rosso	Universitat Politècnica de València, Spain
Eric Sanjuan	Université d’Avignon, France
Benno Stein	Bauhaus-Universität Weimar, Germany
Theodora Tsikrika	Information Technologies Institute, CERTH, Greece
Christa Womser-Hacker	University of Hildesheim, Germany
Md Zia Ullah	CNRS, France

CLEF Steering Committee

Steering Committee Chair

Nicola Ferro University of Padua, Italy

Deputy Steering Committee Chair for the Conference

Paolo Rosso Universitat Politècnica de València, Spain

Deputy Steering Committee Chair for the Evaluation Labs

Martin Braschler Zurich University of Applied Sciences, Switzerland

Members

Khalid Choukri	Evaluations and Language resources Distribution Agency (ELDA), France
Paul Clough	University of Sheffield, UK
Fabio Crestani	Università della Svizzera italiana, Switzerland
Carsten Eickhoff	Brown University, USA
Norbert Fuhr	University of Duisburg-Essen, Germany
Lorraine Goeuriot	Université Grenoble Alpes, France
Julio Gonzalo	National Distance Education University (UNED), Spain
Donna Harman	National Institute for Standards and Technology (NIST), USA
Evangelos Kanoulas	University of Amsterdam, The Netherlands
Birger Larsen	University of Aalborg, Denmark
David E. Losada	Universidade de Santiago de Compostela, Spain
Mihai Lupu	Vienna University of Technology, Austria
Josiane Mothe	IRIT, Université de Toulouse, France
Henning Müller	University of Applied Sciences Western Switzerland (HES-SO), Switzerland
Jian-Yun Nie	Université de Montréal, Canada
Eric SanJuan	University of Avignon, France
Giuseppe Santucci	Sapienza University of Rome, Italy
Jacques Savoy	University of Neuchâtel, Switzerland
Laure Soulier	Pierre and Marie Curie University (Paris 6), France
Theodora Tsikrika	Information Technologies Institute (ITI), Centre for Research and Technology Hellas (CERTH), Greece
Christa Womser-Hacker	University of Hildesheim, Germany

Past Members

Djoerd Hiemstra	Radboud University, The Netherlands
Jaana Kekäläinen	University of Tampere, Finland

Séamus Lawless	Trinity College Dublin, Ireland
Carol Peters	ISTI, National Council of Research (CNR), Italy (Steering Committee Chair 2000–2009)
Emanuele Pianta	Centre for the Evaluation of Language and Communication Technologies (CELCT), Italy
Maarten de Rijke	University of Amsterdam, The Netherlands
Alan Smeaton	Dublin City University, Ireland

Supporters and Sponsors

Supporter



Funded under the European Union’s Horizon 2020 research and innovation programme, ICT-48-2020, grant #951911, the AI4Media project—A European Excellence Centre for Media, Society and Democracy—aspires to become a Centre of Excellence engaging a wide network of researchers across Europe and beyond, focusing on delivering the next generation of core AI advances and training to serve the media sector, while ensuring that the European values of ethical and trustworthy AI are embedded in future AI deployments (<https://www.ai4media.eu/>).

Gold Sponsor

FACEBOOK AI

Facebook AI Research seeks to further our fundamental understanding in both new and existing domains, covering the full spectrum of topics related to AI, with the mission of advancing the state of the art of AI through open research for the benefit of all. Along with the key principles of Facebook AI - openness, collaboration, excellence, and scale - we believe FAIR researchers also need to have the freedom and autonomy to design and follow their own research agendas so they can take on the most impactful work and develop the most disruptive projects, all while sharing their results with the community (<http://ai.facebook.com>).

Silver Sponsor

Keysight Technologies Romania is a leading test and measurement equipment provider for electronic design, e-mobility, network monitoring, 5G, LTE, IoT, connected cars, and more. The company's nearly 12,600 employees serve customers in more than 100 countries (<https://www.keysight.com/>).

Silver Sponsor

Located in Braşov, Cluj-Napoca, and Bucureşti, the Siemens R&D activity is represented in Romania by the Siemens Advanta Development and Siemens Technology teams. The programmers, engineers, and researchers in our teams work with Siemens divisions and business units to identify and implement tomorrow's technologies. The world is constantly changing and we're always thinking about what's next. Siemens Advanta Development is involved in the development of hardware, software, and automation products and solutions in the fields of energy, industry, healthcare, mobility, building technologies, and e-business. The IT Services and Solutions teams provide internal support services for various Siemens platforms. It is also involved in the automated development and testing of applications on Web, mobile, or SAP platforms. The teams from the Technology group carry out their research activity on new design methodologies and advanced verification of the implementations made in ASIC and FPGA, analytics and Big Data, IT Security Software and Systems Innovation, applied research based on computational models, and Artificial Intelligence (<https://new.siemens.com/ro/ro.html>).

Contents

Full Papers

Audio Features, Precomputed for Podcast Retrieval and Information Access Experiments.	3
<i>Abigail Alexander, Matthijs Mars, Josh C. Tingey, Haoyue Yu, Chris Backhouse, Sravana Reddy, and Jussi Karlgren</i>	
A Hybrid Approach for Stock Market Prediction Using Financial News and Stocktwits	15
<i>Alaa Alhamzeh, Saptarshi Mukhopadhyaya, Salim Hafid, Alexandre Bremard, Előd Egyed-Zsigmond, Harald Kosch, and Lionel Brunie</i>	
Dataset Creation Framework for Personalized Type-Based Facet Ranking Tasks Evaluation.	27
<i>Esraa Ali, Annalina Caputo, Séamus Lawless, and Owen Conlan</i>	
SubjectivITA: An Italian Corpus for Subjectivity Detection in Newspapers	40
<i>Francesco Antici, Luca Bolognini, Matteo Antonio Inajetovic, Bogdan Ivasiuk, Andrea Galassi, and Federico Ruggeri</i>	
Evaluating Elements of Web-Based Data Enrichment for Pseudo-relevance Feedback Retrieval	53
<i>Timo Breuer, Melanie Pest, and Philipp Schaer</i>	
End-to-End Fine-Grained Neural Entity Recognition of Patients, Interventions, Outcomes.	65
<i>Anjani Dhrangadhariya, Gustavo Aguilar, Thamar Solorio, Roger Hilfiker, and Henning Müller</i>	
Comparing Traditional and Neural Approaches for Detecting Health-Related Misinformation	78
<i>Marcos Fernández-Pichel, David E. Losada, Juan C. Pichel, and David Elsweller</i>	
Towards the Evaluation of Information Retrieval Systems on Evolving Datasets with Pivot Systems.	91
<i>Gabriela Nicole González-Sáez, Philippe Mulhem, and Lorraine Goeuriot</i>	

A Multi-Task Learning Model for Multidimensional Relevance Assessment	103
<i>Divi Galih Prasetyo Putri, Marco Viviani, and Gabriella Pasi</i>	
Skill Extraction for Domain-Specific Text Retrieval in a Job-Matching Platform	116
<i>Ellery Smith, Andreas Weiler, and Martin Braschler</i>	
Linguistic Uncertainty in Clinical NLP: A Taxonomy, Dataset and Approach	129
<i>Mark Turner, Julia Ive, and Sumithra Velupillai</i>	
Best of 2020 Labs	
Self-calibrating Neural-Probabilistic Model for Authorship Verification Under Covariate Shift	145
<i>Benedikt Boenninghoff, Dorothea Kolossa, and Robert M. Nickel</i>	
Priberam at MESINESP Multi-label Classification of Medical Texts Task . . .	159
<i>Ruben Cardoso, Zita Marinho, Afonso Mendes, and Sebastião Miranda</i>	
Herbarium-Field Triplet Network for Cross-Domain Plant Identification	173
<i>Sophia Chulif and Yang Loong Chang</i>	
BERT-Based Transformers for Early Detection of Mental Health Illnesses . . .	189
<i>Rodrigo Martínez-Castaño, Amal Htait, Leif Azzopardi, and Yashar Moshfeghi</i>	
Dowsing for Math Answers	201
<i>Yin Ki Ng, Dallas J. Fraser, Besat Kassaie, and Frank Wm. Tompa</i>	
Overviews 2021 Labs	
Overview of ARQMath-2 (2021): Second CLEF Lab on Answer Retrieval for Questions on Math	215
<i>Behrooz Mansouri, Richard Zanibbi, Douglas W. Oard, and Anurag Agarwal</i>	
Overview of BioASQ 2021: The Ninth BioASQ Challenge on Large-Scale Biomedical Semantic Indexing and Question Answering	239
<i>Anastasios Nentidis, Georgios Katsimpras, Eirini VANDOROU, Anastasia Kriihara, Luis Gasco, Martin Krallinger, and Georgios Paliouras</i>	

Overview of the CLEF–2021 CheckThat! Lab on Detecting Check-Worthy Claims, Previously Fact-Checked Claims, and Fake News	264
<i>Preslav Nakov, Giovanni Da San Martino, Tamer Elsayed, Alberto Barrón-Cedeño, Rubén Míguez, Shaden Shaar, Firoj Alam, Fatima Haouari, Maram Hasanain, Watheq Mansour, Bayan Hamdan, Zien Sheikh Ali, Nikolay Babulkov, Alex Nikolov, Gautam Kishore Shahi, Julia Maria Struß, Thomas Mandl, Mucahid Kutlu, and Yavuz Selim Kartal</i>	
Overview of ChEMU 2021: Reaction Reference Resolution and Anaphora Resolution in Chemical Patents.	292
<i>Yuan Li, Biaoyan Fang, Jiayuan He, Hiyori Yoshikawa, Saber A. Akhondi, Christian Druckenbrodt, Camilo Thorne, Zubair Afzal, Zenan Zhai, Timothy Baldwin, and Karin Verspoor</i>	
Overview of the CLEF eHealth Evaluation Lab 2021	308
<i>Hanna Suominen, Lorraine Goeuriot, Liadh Kelly, Laura Alonso Alemany, Elias Bassani, Nicola Brew-Sam, Viviana Cotik, Darío Filippa, Gabriela González-Sáez, Franco Luque, Philippe Mulhem, Gabriella Pasi, Roland Roller, Sandaru Seneviratne, Rishabh Upadhyay, Jorge Vivaldi, Marco Viviani, and Chenchen Xu</i>	
Overview of eRisk 2021: Early Risk Prediction on the Internet.	324
<i>Javier Parapar, Patricia Martín-Rodilla, David E. Losada, and Fabio Crestani</i>	
Overview of the ImageCLEF 2021: Multimedia Retrieval in Medical, Nature, Internet and Social Media Applications.	345
<i>Bogdan Ionescu, Henning Müller, Renaud Péteri, Asma Ben Abacha, Mourad Sarrouiti, Dina Demner-Fushman, Sadid A. Hasan, Serge Kozlovski, Vitali Liauchuk, Yashin Dicente Cid, Vassili Kovalev, Obioma Pelka, Alba García Seco de Herrera, Janadhip Jacutprakart, Christoph M. Friedrich, Raul Berari, Andrei Tauteanu, Dimitri Fichou, Paul Brie, Mihai Dogariu, Liviu Daniel Ştefan, Mihai Gabriel Constantin, Jon Chamberlain, Antonio Campello, Adrian Clark, Thomas A. Oliver, Hassan Moustahfid, Adrian Popescu, and Jérôme Deshayes-Chossart</i>	
Overview of LifeCLEF 2021: An Evaluation of Machine-Learning Based Species Identification and Species Distribution Prediction.	371
<i>Alexis Joly, Hervé Goëau, Stefan Kahl, Lukáš Picek, Titouan Lorieul, Elijah Cole, Benjamin Deneu, Maximilien Servajean, Andrew Durso, Isabelle Bolon, Hervé Glotin, Robert Planqué, Rafael Ruiz de Castañeda, Willem-Pier Vellinga, Holger Klinck, Tom Denton, Ivan Eggel, Pierre Bonnet, and Henning Müller</i>	

Overview of LiLAS 2021 – Living Labs for Academic Search 394
*Philipp Schaer, Timo Breuer, Leyla Jael Castro, Benjamin Wolff,
Johann Schaible, and Narges Tavakolpoursaleh*

Overview of PAN 2021: Authorship Verification, Profiling Hate Speech
Spreaders on Twitter, and Style Change Detection. 419
*Janek Bevendorff, Berta Chulvi, Gretel Liz De La Peña Sarracén,
Mike Kestemont, Enrique Manjavacas, Ilia Markov, Maximilian Mayerl,
Martin Potthast, Francisco Rangel, Paolo Rosso, Efstathios Stamatatos,
Benno Stein, Matti Wiegmann, Magdalena Wolska, and Eva Zangerle*

Overview of SimpleText 2021 - CLEF Workshop on Text Simplification
for Scientific Information Access 432
*Liana Ermakova, Patrice Bellot, Pavel Braslavski, Jaap Kamps,
Josiane Mothe, Diana Nurbakova, Irina Ovchinnikova,
and Eric SanJuan*

Overview of Touché 2021: Argument Retrieval 450
*Alexander Bondarenko, Lukas Gienapp, Maik Fröbe, Meriem Beloucif,
Yamen Ajjour, Alexander Panchenko, Chris Biemann, Benno Stein,
Henning Wachsmuth, Martin Potthast, and Matthias Hagen*

Author Index 469