Report on the 5th International Workshop on Narrative Extraction from Texts (Text2Story 2022) at ECIR 2022

Ricardo Campos
INESC TEC & IPT
Portugal
ricardo.campos@ipt.pt

Alípio M. Jorge
INESC TEC & FCUP
Portugal
amjorge@fc.up.pt

Adam Jatowt
Univ. of Innsbruck
Austria
jatowt@gmail.com

Sumit Bhatia
Media and Data Science Research Lab, Adobe
India
sumit.summit@gmail.com

Marina Litvak
Shamoon College of Engineering
Israel
litvak.marina@gmail.com

João Paulo Cordeiro
INESC TEC & Univ. of Beira Interior
Portugal

Conceição Rocha
INESC TEC
Portugal

Hugo Sousa
INESC TEC
Portugal

Behrooz Mansouri
University of Southern Maine
USA

Abstract

The Fifth International Workshop on Narrative Extraction from Texts (Text2Story'22) was held on the April 10^{th}, 2022, in conjunction with the 44^{th} European Conference on Information Retrieval (ECIR 2022) in Stavanger, Norway. Due to the COVID-19 restrictions that are still active in some countries, the workshop was held as an hybrid event, combining a "live" in-person and a "virtual" online participation. The online participation was allowed using the Zoom platform. During the course of the day, more than 50 attendees - about 40 in person and 10 in zoom - had the opportunity to follow-up and discuss the recent advances in topics related to representation, extraction, and generation of narratives. The workshop program included two invited keynotes and twelve paper presentations. The proceedings of the workshop are available online at http://ceur-ws.org/Vol-3117/

Date: 10 April, 2022.

Website: https://text2story22.inesctec.pt/.

Figure 1: Text2Story2022 Workshop: Best Paper Rewarding

1 Introduction

The Text2Story series of Workshops started with the aims of bringing together researchers from diverse but related fields such as Information Retrieval, Natural Language Processing, Computational Linguistics, Artificial Intelligence, Human-Computer Interaction, and Data Visualization to share the recent advances in their respective fields towards narrative understanding. Building upon the success of the past editions [Jorge et al., 2018, 2019a; Campos et al., 2020, 2021] and on the Text2Story Special Issue at IPM Journal [Jorge et al., 2019b], this year, we organized the fifth edition of the Text2Story workshop, held as a hybrid event under the umbrella of ECIR 2022. Figure 1 depicts the Best Paper Rewarding.

This report summarizes the key activities at the workshop. The enthusiasm and active participation of the attendees ensured that the hybrid edition of the workshop was as intellectually stimulating as the past physical and virtual (in 2020 and 2021) editions. To summarize,there were 21 papers submitted for peer-review to the workshop this year. Out of these, 12 papers were accepted for this volume, 7 as regular papers, 1 short paper, 1 demo paper, 1 position paper, and 2 work in progress papers. The workshop program included a total of 12 research papers from 40 authors and two invited keynote talks. The presentation slides and the videos of the part of talks at the workshop can be found at https://text2story22.inesctec.pt/.

2 The Program

The workshop program consisted of two invited keynote talks, seven regular papers, one short paper, and one demo paper. In addition, we also had one position paper and two works-in-progress that allowed the respective authors to present their ongoing work and get useful feedback from the workshop attendees. The papers presented at the workshop covered diverse aspects of the narrative analysis including alternative text representations, understanding, simplifying, and visualization of news, summarization, data annotation and augmentation, and more. The presented methods covered both extraction and generation techniques from supervised and unsupervised approaches. The research works were applied on a range of datasets—from social media to fiction stories—while their domains varied from COVID-19 to climate change.

2.1 Keynotes

Andreas Spitz, Head of the Data and Information Mining lab at the University of Konstanz, presented Quobert, a transformer-based model that exploits the parallelism in news reporting for the extraction and attribution of quotes from news. Andreas demonstrated how a comprehensive corpus of 235 million unique quotations extracted with Quobert from a decade of news, can be used to quantify trends in political language. He focused on the uptick in negativity in U.S. politicians' language after the end of Obama's tenure, quantified the shifts in language tone, and unraveled to whom these shifts could feasibly be attributed. [Spitz, 2022].

The second keynote by **Antoine Doucet**, a tenured full Professor in computer science at the L3i laboratory of the University of La Rochelle, presented recent advances in AI and natural language understanding that enable historical and handwritten documents to be analyzed in a way that is robust to digitization. Antoine showed how the research group of the H2020 NewsEye project created state-of-the-art results for the cross-lingual recognition and disambiguation of named entities in a large corpora of historical newspapers written in 4 languages between 1850 and 1950. These results pave the way to a large-scale analysis of digitized documents, notably able to cross linguistic borders [Doucet, 2022].

2.2 Tools and Frameworks for Narrative Analysis

Sriharsh Bhyravajjula and Shrivastava [2022] presented MARCUS (Modelling Arcs for Understanding Stories), an NLP pipeline that extracts events, participant characters, implied emotion, and sentiment to model inter-character relations. MARCUS addresses the novel task of computationally generating event-centric, relation-based character arcs from narratives. Character arcs are important theoretical devices employed in literary studies to understand character journeys, identify tropes across literary genres, and establish similarities between narratives.

João Santos and Miranda [2022] discuss a task of organizing and clustering multilingual news articles for media monitoring in real time. The authors presented an online system that is able to cluster an incoming stream of documents without depending on language-specific features. The proposed system achieves state-of-the-art results on a multilingual news stream clustering dataset.

2.3 Interpreting and Understanding Narratives

Jonas Rieger and Jentsch [2022] addressed a task of automatic monitoring of an evolution of topics identified in textual data. The authors proposed a dynamic change detection method that relies on a rolling version of the classical LDA. The changes are detected by comparison of the actual intensity of word change in the LDA's topics over time to the expected intensity of word change under stable conditions. Jonas Rieger and Jentsch [2022] applied their method to topics obtained by the RollingLDA to COVID-19 related news data from CNN and illustrated that the detected changes in these topics are well interpretable.

Margaret Meehan and Piper [2022] addressed a problem of understanding causality in narrative. They established the relative performance of baseline causal mining models on examples drawn from fictional narratives and compare them with standard NLP benchmarks drawn from SemEval data. With a new labeled dataset that the authors introduce, they train models to detect causality within and between sentences, and uncover linguistic features that indicate a causal relationship between phrases.

Chen and Bunescu [2022] discussed a problem of changing the narrative perspective and its NLP solution that requires the capability to select mention strings that refer to the character in a natural and non-ambiguous manner. The authors introduced and evaluated three mention selection architectures: LSTM with attention over frozen BERT embeddings, fine-tuned BERT with coreference-modulated self-attention, and prompt-based tuning over either frozen or fine-tuned T5. They experimentally showed that the prompt-tuning approach over frozen T5 obtains the best performance, also outperforming the previous state-of-the-art on this task.

Piyush Kumar Garg and Dandapat [2022] proposed an entropy and diversity based summarizer, named EnDSUM, for disaster tweet summarization. EnDSUM can automatically generate summary of all the tweets related to a particular disaster, which is a necessary to ensure quick crisis response and provide situational updates from government agencies and humanitarian organizations. The authors performed a comprehensive analysis on six datasets which indicates the effectiveness of EnDSUM.

2.4 Applications and Case Studies

Ishrat Rahman Sami and Soldatova [2022] demonstrated visualizations designed to provide cognitive guidance for planning editorial news stories. Their approach is based on five basic questions "Who", "Where", "What", "When" and "Why" which are fundamental for any readers'understanding. The visualizations are contextual: global (considering the whole news archive), relative (considering topic-based news collection) and local (considering single news).

Ali Emre Varol and Talby [2022] analyzed more than 36,000 COVID-19 related news articles from CNN and The Guardian (Pasquali et al. [2021]), using the clinical and biomedical NLP models. The analysis covers key entities and phrases, observed biases, and change over time in news coverage by correlating mined medical symptoms, procedures, drugs, and guidance with commonly mentioned demographic and occupational groups. They also performed another analysis about drug and vaccine manufacturers, which when reported by major news outlets has an impact on vaccine hesitancy.

Bart Gajderowicz and Fox [2022] addressed the problem of matching client needs to the social purpose organizations (SPO) services. They presented an approach, based on Named Entity

Recognition, to extracting an SPO's impact model, i.e., "narrative", from various text sources. The authors showed how to use it for the sequencing of services and impacted clients with information extracted from unstructured text. They define each service as a "story" describing what they offer, to whom, how, and when. The narrative, then, forms a system of services that guides clients through various programs towards achieving their goals.

2.5 Ongoing Research Efforts

The workshop also included special tracks for work-in-progress and position papers to offer a platform for authors to share initial results of their research efforts and gather useful feedback from the community. Joana Valente and Nunes [2022] presented a position paper, where a new narrative visualization approach is proposed. This approach uses icons to represent important narrative elements and is integrated in Brat2Viz, a narrative annotation visualization tool that implements a pipeline transforming text annotations into formal representations leading to narrative visualizations. Authors implemented a narrative element extraction process that uses automatic sentence extraction, automatic translation methods, and an algorithm that determines the actors' most adequate descriptions. They also introduced a method for an icon dictionary collection, with the ability to automatically search for icons. The reported human evaluation shows positive and promising results.

Satya Almasian and Gertz [2022] explored transfer learning strategies to improve the quality of a German temporal tagger. They employ a weakly-supervised pre-training strategy to stabilize the convergence of Transformer-based taggers. Also, the authors augmented data with automatically translated English resources, which served as an alternative to commonly used alignments of latent embedding spaces. The authors provided empirical evidence of the suitability of transfer approaches to other low-resourced languages: a small number of gold data coupled with an existing data set in a resource-rich language and a weak labeling baseline system may be sufficient to boost performance.

Jakub Piskorski and Linge [2022] reported on preliminary experiments of exploiting Data Augmentation techniques for improving climate change denial classification. The authors focused on a selection of known techniques, and also augmentation transformations not reported elsewhere that replace certain type of named entities with high probability of preserving labels. They also introduced a new benchmark dataset consisting of text snippets extracted from online news labeled with fine-grained climate change denial types.

3 Key Contributors

The Program Chairs were Ricardo Campos, Alípio Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak. The proceedings were setup by our Proceedings Chairs João Paulo Cordeiro and Conceição Rocha. Hugo Sousa and Behrooz Mansouri were the web and dissemination chairs. Vasco Campos (INESC TEC & University of Porto) chaired the last technical session.

Finally, we would like to acknowledge the effort and valuable contribution of the researchers and industry experts that have served on the Program Committee of the Text2Story'22 workshop. Our thanks go to:

- Álvaro Figueira (INESC TEC & University of Porto)
- Andreas Spitz (University of Konstanz)
- António Horta Branco (University of Lisbon)
- Arian Pasquali (CitizenLab)
- Brenda Santana (Federal University of Rio Grande do Sul)
- Bruno Martins (IST and INESC-ID - Instituto Superior Técnico, University of Lisbon)
- Demian Gholipour (University College Dublin)
- Daniel Gomes (FCT/Arquivo.pt)
- Daniel Loureiro (University of Porto)
- Deya Banisakher (Defense Threat Reduction Agency (DTRA), Ft. Belvior, VA, USA.)
- Dhruv Gupta (Norwegian University of Science and Technology (NTNU), Trondheim, Norway)
- Dwaipayan Roy (ISI Kolkata, India)
- Dyaa Albakour (Signal)
- Evelin Amorim (INESC TEC)
- Florian Boudin (Université de Nantes)
- Grigorios Tsoumakas (Aristotle University of Thessaloniki)
- Henrique Lopes Cardoso (University of Porto) Hugo Sousa (INESC TEC)
- Ismail Sengor Altingovde (Middle East Technical University)
- Jeffery Ansah (BHP)
- João Paulo Cordeiro (INESC TEC & University of Beira Interior)
- Kiran Kumar Bandeli (Walmart Inc.)
- Ludovic Moncla (INSA Lyon)
- Marc Spaniol (Université de Caen Normandie)
- Mark Finlayson (Florida International University)
- Nina Tahmasebi (University of Gothenburg)
- Pablo Gamallo (University of Santiago de Compostela)
- Paulo Quaresma (Universidade de Évora)
- Pablo Gervás (Universidad Complutense de Madrid)
- Paul Rayson (Lancaster University)
- Satya Almasian (Heidelberg University)
- Sérgio Nunes (INESC TEC & University of Porto)
- Udo Kruschwitz (University of Regensburg)
- Yihong Zhang (Kyoto University)

4 Recognition Awards

The workshop ended with the announcement of the recognition awards. The Text2Story 2022 Best Paper Award went to *Sriharsh Bhyravajjula*, *Ujwal Narayan*, and *Manish Shrivastava* for their paper entitled *MARCUS: An Event-Centric NLP Pipeline that generates Character Arcs from Narratives*. In addition to this, the workshop chairs have recognized the following researchers as recipients of the 2022 reviewer award for their insightful and valuable reviews - *Andreas Spitz*

(University of Konstanz), *Deya Banisakher* (Defense Threat Reduction Agency (DTRA)), *Pablo Gervás (Universidad Complutense de Madrid)* and *Satya Almasian* (Heidelberg University).

5 Concluding Thoughts

This was the fifth edition of the Text2Story workshop series. We started with an objective to bring together the interested participants from different geographies and research expertise to collectively set the agenda for the emerging multi-disciplinary area of *narrative extraction from texts*. We are overwhelmed from the increasingly positive response from the research community and the growing participation in the workshops underlines the relevance of the topic.

Acknowledgements

We would like to thank the ECIR 2022 organizers for helping us to organize the workshop, our wonderful keynote speakers, members of the program committee, the paper authors, and the participants who made the workshop a grand success.

References

Hasham Ul Haq Ali Emre Varol, Veysel Kocaman and David Talby. CUnderstanding COVID-19 News Coverage using Medical NLP. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 45–53. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper5.pdf.

Daniela Rosu Bart Gajderowicz and Mark S. Fox. Extracting Impact Model Narratives from Social Services' Text. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 55–65. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper6.pdf.

Ricardo Campos, Alípio Jorge, Adam Jatowt, and Sumit Bhatia. The 3rd international workshop on narrative extraction from texts: Text2story 2020. In Joemon M. Jose, Emine Yilmaz, João Magalhães, Pablo Castells, Nicola Ferro, Mário J. Silva, and Flávio Martins, editors, *Advances in Information Retrieval - 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14-17, 2020, Proceedings, Part II*, volume 12036 of *Lecture Notes in Computer Science*, pages 648–653. Springer, 2020. doi: 10.1007/978-3-030-45442-5_86. URL https://doi.org/10.1007/978-3-030-45442-5_86.

Ricardo Campos, Alípio Jorge, Adam Jatowt, Sumit Bhatia, and Mark Finlayson. The 4th international workshop on narrative extraction from texts: Text2story 2021. In Djoerd Hiemstra, Marie-Francine Moens, Josiane Mothe, Raffaele Perego, Martin Potthast, and Fabrizio Sebastiani, editors, *Advances in Information Retrieval*, pages 701–704, Cham, 2021. Springer International Publishing. ISBN 978-3-030-72240-1.

Mike Chen and Razvan Bunescu. Changing the Narrative Perspective: From Ranking to Prompt-Based Generation of Entity Mentions. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 35–43. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper4.pdf.

Antoine Doucet. Robust and multilingual analysis of historical documents. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 3–4. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/keynote2.pdf.

Tony Russell-Rose Ishrat Rahman Sami and Larisa Soldatova. Comprehensive contextual visualization of a news archive for aiding story planning. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 75–81. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper8.pdf.

Nicolas Stefanovitch Bonka Kotseva Irene Vianini Sopho Kharazi Jakub Piskorski, Nikolaos Nikolaidis and Jens P. Linge. Exploring Data Augmentation for Classification of Climate Change Denial: Preliminary Study. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 97–109. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper11.pdf.

Alípio Jorge Joana Valente and Sérgio Nunes. Text2Icons: linking icons to narrative participants. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 111–116. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper12.pdf.

Jonathan Flossdorf Jonas Rieger, Kai-Robin Lange and Carsten Jentsch. Dynamic change detection in topics based on rolling LDAs. In Ricardo Campos, Alípio Mário Jorge, Adam

Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 5–13. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper1.pdf.

Alípio M. Jorge, Ricardo Campos, Adam Jatowt, and Sérgio Nunes. First international workshop on narrative extraction from text (text2story'18). In *Pasi G., Piwowarski B., Azzopardi L., Hanbury A. (eds). Advances in Information Retrieval. ECIR 2018 (Grenoble, France. March 26 – 29). Lecture Notes in Computer Science, vol 10772*, pages 833–834, 2018.

Alípio M. Jorge, Ricardo Campos, Adam Jatowt, and Sumit Bhatia. Second international workshop on narrative extraction from text (text2story'19). In *Azzopardi L., Stein B., Fuhr N., Mayr P., Hauff C., Hiemstra D. (eds), Advances in Information Retrieval. ECIR'19 (Cologne, Germany. April 14 – 18). Lecture Notes in Computer Science, vol 11438*, pages 389–393, 2019a.

Alípio M. Jorge, Ricardo Campos, Adam Jatowt, and Sérgio Nunes. Special issue on narrative extraction from texts (text2story): Preface. In *Information Processing Management an International Journal. Elsevier, Vol 56(5)*, pages 1771–1774, 2019b. URL https://www.sciencedirect.com/science/article/abs/pii/S0306457319304455.

Afonso Mendes João Santos and Sebastião Miranda. Simplifying Multilingual News Clustering Through Projection From a Shared Space. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 15–24. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper2.pdf.

Dane Malenfant Margaret Meehan and Andrew Piper. Causality Mining in Fiction. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 25–34. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper3.pdf.

Arian Pasquali, Ricardo Campos, Alexandre Ribeiro, Brenda Santana, Alípio Jorge, and Adam Jatowt. Tls-covid19: A new annotated corpus for timeline summarization. In Djoerd Hiemstra, Marie-Francine Moens, Josiane Mothe, Raffaele Perego, Martin Potthast, and Fabrizio Sebastiani, editors, *Advances in Information Retrieval*, pages 497–512, Cham, 2021. Springer International Publishing. ISBN 978-3-030-72113-8.

Roshni Chakraborty Piyush Kumar Garg and Sourav Kumar Dandapat. EnDSUM: Entropy and Diversity based Disaster Tweet Summarization. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference*

on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022, volume 3117 of *CEUR Workshop Proceedings*, pages 91–96. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper10.pdf.

Dennis Aumiller Satya Almasian and Michael Gertz. Time for some German? Pre-Training a Transformer-based Temporal Tagger for German. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 83–90. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper9.pdf.

Andreas Spitz. We have the best words: From the web-scale extraction and attribution of quotes to analyzing negativity in u.s. political language. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, page 1. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/keynote1.pdf.

Ujwal Narayan Sriharsh Bhyravajjula and Manish Shrivastava. MARCUS: An Event-Centric NLP Pipeline that generates Character Arcs from Narratives. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway, April 10th, 2022*, volume 3117 of *CEUR Workshop Proceedings*, pages 67–74. CEUR-WS.org, 2022. URL http://ceur-ws.org/Vol-3117/paper7.pdf.