

Report on the 9th International Workshop on Narrative Extraction from Texts (Text2Story 2026) at ECIR 2026

Ricardo Campos

Univ. of Beira Interior & INESC TEC & Ci2@IPT
Portugal
ricardo.campos@ubi.pt

Alípio M. Jorge

Univ. of Porto & INESC TEC
Portugal
amjorge@fc.up.pt

Adam Jatowt

Univ. of Innsbruck
Austria
jatowt@acm.org

Sumit Bhatia

Media and Data Science Research Lab, Adobe
India
Sumit.Bhatia@adobe.com

Marina Litvak

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

João Paulo Cordeiro

(NOVA LINC'S & Univ. of Beira Interior
Portugal

Conceição Rocha

INESC TEC
Portugal

Luis Filipe Cunha

Univ. of Porto & INESC TEC
Portugal

Behrooz Mansouri

University of Southern Maine
USA

Abstract

The Ninth International Workshop on Narrative Extraction from Texts (Text2Story'26) was held on March 29th, 2026, in conjunction with the 48th European Conference on Information Retrieval (ECIR 2026) in Delft, The Netherlands. During this full-day event, attendees engaged in discussions and presentations focused on recent advancements in narrative representation, extraction, generation, and evaluation, with particular emphasis on the growing convergence of Information Retrieval with Transformers, Large Language Models (LLMs), and agentic AI. The workshop featured two keynote addresses and a mix of oral presentations and a poster session covering fourteen accepted papers. The workshop proceedings are available online¹.

Date: 29 March 2026.

Website: <https://text2story26.inesctec.pt/>.

¹<https://ceur-ws.org/Vol-4202/>

1 Introduction

The Text2Story workshop series was initiated to foster collaboration between researchers from various disciplines with a shared interest in narrative understanding. These fields include Information Retrieval, Natural Language Processing, Computational Linguistics, Artificial Intelligence, Human-Computer Interaction, and Data Visualization. Building on the momentum of previous successful editions [Jorge et al., 2018, 2019a; Campos et al., 2020, 2021, 2022, 2023, 2024, 2025] and the Text2Story Special Issue at IPM Journal [Jorge et al., 2019b], we held this year the ninth iteration of the Text2Story workshop event at ECIR 2026.



Figure 1. Text2Story 2026 Workshop Participants.

This year’s edition placed particular emphasis on the growing impact of LLMs and agentic AI on the field of narrative understanding. While these technologies have driven significant progress in the field, important challenges persist around fine-grained narrative structure, dialogue and multi-modal narratives, robustness to domain shift, and evaluation beyond surface accuracy. Recognizing these issues, the call for papers explicitly invited contributions on narrative information extraction and representation; narrative analysis and generation; narrative analysis of LLM-generated text; data creation and augmentation with LLMs; evaluation of narrative extraction with LLMs and the role of narratives in human-agent communication; datasets, resources, and evaluation; ethics, bias, and misinformation; and a wide range of narrative applications such as narrative-focused search, summarization, question answering, timeline generation, and multimodal narratives.

This report provides an overview of the main activities that took place during the workshop. The attendees’ enthusiastic engagement ensured that it was just as intellectually stimulating as previous editions. A total of nineteen submissions were received and each paper was reviewed by three PC members. The review process resulted in a total of fourteen papers being accepted for the workshop under various tracks, providing a platform for researchers to share ongoing work and disseminate previously published works from related venues. The workshop also featured two

invited keynotes and a vibrant poster session. The presentation slides for papers and the keynote talks can be found at the workshop website².

2 The Program

The workshop program featured two keynote addresses alongside fourteen research papers. These papers were presented in a mix of oral sessions and during a dedicated poster session, fostering in-depth discussions among workshop participants. The day was organized in six sessions chaired by Ricardo Campos, Alípio Jorge, Liana Ermakova, Adam Jatowt, Jaap Kamps, Purificação Silvano, and Nuno Guimarães. In the following sections, we provide a concise overview of the research presented at the workshop.

2.1 Keynotes

Professor Federico Pianzola, Assistant Professor of Computational Humanities at the University of Groningen and Principal Investigator of the ERC Starting Grant project GOLEM (“Graphs and Ontologies for Literary Evolution Models”), delivered the keynote titled “*Knowledge Extraction from Fiction Using Generative Large Language Models and Multilingual Semantic Parsing*” [Pianzola, 2026]. The talk addressed the well-known challenge of extracting event knowledge from unstructured text, which is particularly difficult in fiction due to the prevalence of subtext and figurative style. Professor Pianzola presented an evaluation of different techniques for extracting narrative information from fiction, including the prompting of open-weights LLMs and semantic parsing via Abstract Meaning Representation (AMR), framed within the context of the GOLEM project.

Arian Pasquali, Applied AI Researcher at Orq.ai (Amsterdam), delivered the second keynote titled “*Evaluating Generative AI and Agentic Systems: Methodologies for Systematic Assessment and Continuous Improvement*” [Pasquali, 2026]. Drawing on real-world industry practice, the talk reviewed evaluation methodologies for generative AI and agentic systems, with particular emphasis on the gap between academic benchmarks and production realities. The talk introduced Evaluation-Driven Development as a framework for continuous, iterative improvement of AI systems in production, and discussed the operationalization of LLM-as-a-Judge approaches – how to align automated judges with human judgment, avoid common calibration pitfalls, and scale evaluation through panels of judges. The keynote also covered agentic systems specifically, including agent simulation as a method for stress-testing agent behavior before deployment, and trace intelligence as a practice for extracting structured meaning from agent execution traces.

2.2 Full Research Papers

The full papers presented at the workshop spanned a wide variety of topics, covering foundational issues in narrative representation, extraction, and similarity, narrative mining in sociopolitical and environmental contexts, and the safety and evaluation of LLMs in clinical settings.

²<https://text2story26.inesctec.pt/>

2.2.1 Narrative Extraction, Visualization, and Similarity

Several papers presented at the workshop advanced computational approaches to extract, visualize, and compare narratives. [Keith-Norambuena et al. \[2026a\]](#) presented a user study evaluating semantic interaction (SI) for narrative map sensemaking. With 33 participants examining narratives from the 2021 Cuban protests under three conditions, a timeline baseline, a basic narrative map, and an interactive narrative map with SI capabilities, the study found that map-based prototypes yielded more insights than the timeline baseline, with the SI-enabled condition reaching statistical significance. Qualitative analysis identified two distinct strategies, corrective and additive, that analysts employ to impose quality judgments and organizational structure on extracted narratives, and SI users achieved comparable exploration breadth with less parameter manipulation.

In a complementary contribution to the narrative maps research line, [Keith-Norambuena et al. \[2026b\]](#) introduced agenda-based narrative extraction, a method that integrates LLMs into the Narrative Trails pathfinding process to steer storyline construction toward user-specified perspectives expressed in natural language. Evaluated on a news article corpus using LLM judges (Claude Opus 4.5 and GPT-5.1), their approach achieved 9.9% higher alignment with semantic agendas than keyword matching, with only a 2.2% reduction in coherence relative to an agenda-agnostic baseline, while counter-agendas that contradicted the source material scored uniformly low across all methods, confirming that steering cannot fabricate unsupported narratives.

[Barancová and Sileno \[2026\]](#) examined how structured narrative elements extracted with LLMs can complement document-level embeddings for narrative similarity detection. Focusing on the role of characters, they extended a baseline approach for similar-story retrieval using a taxonomy-free character role schema and tested the pipeline on the Tell me Again! dataset of movie summaries. Their results show a small but consistent improvement over an embedding-only baseline, providing a proof of concept for integrating narratological abstractions into otherwise embedding-driven similarity detection.

2.2.2 Narrative Mining in Sociopolitical and Environmental Contexts

Two full papers explored the use of narrative-aware methods to study sociopolitical and environmental discourse. [Agarwal and Subedi \[2026\]](#) analyzed how social, cultural, economic, and political (SCEP) symbols shaped the spread of narratives surrounding the 2025 U.S. tariff expansion across X/Twitter, Instagram, and TikTok. Combining LLM-based clustering to identify dominant narratives, vision-language models to detect symbolic cues, and the SEIZ epidemiological diffusion model to characterise propagation, the study found that economic symbols were more prominent in Instagram narratives focused on markets and finance, while political symbols played a stronger role on X/Twitter and TikTok, and that posts containing multiple symbolic cues exhibited higher transmission rates on Instagram and TikTok.

[Rieger et al. \[2026\]](#) applied a spatio-temporal topic modelling framework to a corpus of 1,996 validated news articles covering socio-environmental conflicts associated with Chile’s energy transition between 2011 and 2025. Using RollingLDA, a dynamic adaptation of latent Dirichlet allocation that prevents information leakage from future documents, the authors identified twelve topics and traced how specific conflicts – such as the HidroAysén dam project, the Dominga mining controversy, and pollution in the Quintero–Puchuncaví sacrifice zone – evolved over time, alongside the emergence of new concerns around green hydrogen development and lithium extraction.

The analysis is complemented by an interactive dashboard that maps the prevalence of topics across Chilean regions, providing a reproducible framework for tracking how socio-environmental narratives emerge, evolve, and fade over time.

2.2.3 Safety and Evaluation of Medical Large Language Models

Continuing the line of clinical applications that has featured prominently in recent editions of the workshop, [Feier et al. \[2026\]](#) presented a multi-domain red teaming framework for the safety, robustness, and fairness evaluation of medical LLMs. The framework evaluated eleven contemporary models across 690 clinically grounded scenarios spanning nine domains and over 150 subcategories, using adversarial transformations and a seven-dimension scoring rubric combined with an LLM-assisted scoring pipeline and human-in-the-loop validation. Mean scores ranged from 0.791 to 0.984, but several high-performing systems produced complete failures in individual safety-critical scenarios, demonstrating that aggregate accuracy can mask clinically meaningful risk. Equity-related tasks showed 10–20% error amplification under demographic modifications, and human reviewers identified clinically relevant failures missed by automated evaluation, underscoring that performance variance and worst-case failures provide more clinically meaningful reliability indicators than mean accuracy alone.

2.3 Short Papers

Continuing the practice of previous years, the workshop also featured a dedicated track for short papers covering work-in-progress and position papers. These papers provide a space for researchers to present early-stage results, articulate research visions, and receive community feedback on emerging ideas.

Two papers focused on the evaluation and generation of stories produced by AI systems. [Argese et al. \[2026\]](#) proposed StoryScore, a composite metric for evaluating AI-generated scientific stories that integrates semantic alignment, lexical grounding, narrative control, structural fidelity, redundancy avoidance, and entity-level hallucination detection. Their analysis highlights why many existing hallucination detection methods conflate legitimate pedagogical creativity with factual errors, exposing a key limitation of similarity-based metrics for evaluating audience-adapted scientific narratives. [Palma and Jager \[2026\]](#) revisited TRIZ – a systematic theory of inventive problem-solving – as a framework for AI-driven narrative generation. They proposed a 24×24 narrative contradiction matrix grounded in classical narratological theory, encoded as a complement to Narrative Ontologies using RDF-Star, and outlined a human-in-the-loop workflow in which an agentic system proactively prompts the writer at identified narrative bottlenecks within the ArchiveMyLegacy photo-to-story pipeline.

Three contributions explored the analysis of narratives in literary and cultural texts. [Genoni et al. \[2026\]](#) evaluated the applicability of LLMs to character identification in the low-resource narrative domain of Italian Renaissance epic poetry, using Ludovico Ariosto’s *Orlando Furioso* as a case study. Their pipeline integrates recursive narrative memory generation with character extraction, achieving an F1 score of 0.758 with Qwen-2.5-72B; qualitative error analysis nonetheless revealed persistent failures in cataphoric resolution, where character identities are disclosed non-linearly. [Kaplan et al. \[2026\]](#) investigated whether five literary patterns (Comedy, Mystery,

Romance, Satire, and Tragedy) introduced in prior work are prevalent in award-winning novellas, by manually annotating 35 novellas and comparing the results with automated LLM annotation. Their feature selection analysis identified Closure, Final Confrontation, and Return as the most structurally central elements across genres, and reported high recall but low precision of the LLM with respect to human annotation. [López Calle and Gervás \[2026\]](#) examined whether a sufficiently detailed description of a structured narrative representation can support zero-prompt segmentation of movie dialogues into Action Units. Using *The Princess Bride* as a reference, the study showed that ChatGPT provides reasonable responses on small fragments of screenplay dialogue but tends to omit relevant information when prompted with longer fragments.

Three further short papers contributed methodological reflections and position pieces. [Harshina et al. \[2026\]](#) presented a work-in-progress participatory storytelling pipeline for sensitive narrative elicitation in contexts of forced migration and displacement. The system uses a locally hosted LLM as a neutral chatbot interviewer to support self-paced narration without cloud processing, prioritizing privacy, data sovereignty, and participant control, and transforms interview transcripts into a structured scene representation that can be replayed through a prototype interface. A small formative expert evaluation reported that participants described the interviewer as low-pressure and reflective, while highlighting limitations such as weak acknowledgement of long answers and occasional forced turns. [Huang and Usbeck \[2026\]](#) contributed a position paper on Human Label Variation (HLV) in computational narrative understanding, providing a distance-focused analysis of the machine learning pipeline. They examined how the single ground truth assumption shapes the interpretation of distance in annotation reliability, loss computation, and model evaluation metrics, arguing that the pluralistic nature of reading and interpreting narratives is suppressed by majority-voting label aggregation, and called on the community to provide unaggregated datasets and to interpret reliability scores, training losses, and evaluation metrics with HLV in mind. [Munna et al. \[2026\]](#) argued, in a position paper, that annotation-free localization of findings in chest X-ray images can be strengthened by integrating explicit domain knowledge in a narrative-aware approach. Rather than relying solely on raw report text, they enrich free-text radiology reports with disease-centric narrative descriptions encoding clinically meaningful attributes such as definition, anatomical location, visual appearance, shape, and extent. Initial empirical observations on chest X-ray benchmarks suggest progressive – though not yet statistically significant – improvements in localization performance over free-text supervision alone.

3 Key Contributors

The Program Chairs were Ricardo Campos, Alípio Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak. The proceedings were set up by our Proceedings Chairs Conceição Rocha and João Cordeiro. Luís Filipe Cunha and Behrooz Mansouri served as web and dissemination chairs.

Finally, we would like to acknowledge the effort and valuable contribution of the researchers and industry experts who served on the Program Committee of the Text2Story’26 workshop. Our thanks go to:

- Alvaro Figueira (CRACS / INESC TEC and University of Porto)
- Andreas Spitz (University of Konstanz)
- Antoine Doucet (University of La Rochelle)

-
- António Branco (Faculdade de Ciências da Universidade de Lisboa)
 - Arian Pasquali (Orq.ai)
 - Behrooz Mansouri (University of Southern Maine)
 - Bruce Santos (Laboratório de Inteligência Artificial)
 - Bruno Martins (IST and INESC-ID, Instituto Superior Técnico, University of Lisbon)
 - Dhruv Gupta (Norwegian University of Science and Technology)
 - Evelin Amorim (INESC TEC)
 - Federico Pianzola (University of Groningen)
 - Florian Boudin (Université de Nantes)
 - Henrique Lopes Cardoso (University of Porto)
 - Hugo Sousa (University of Porto)
 - Irina Rabaev (Shamoon College of Engineering, SCE, Beer Sheva, Israel)
 - Ismail Sengor Altinogvde (Middle East Technical University)
 - Jakub Piskorski (Joint Research Centre of the European Commission)
 - Joao Cordeiro (Universidade Beira Interior)
 - Joaquim Silva (Nova School of Science and Technology – NOVA LINCS)
 - Junbo Huang (University of Hamburg)
 - Lise Stork (University of Amsterdam)
 - Luca Cagliero (Politecnico di Torino)
 - Ludovic Moncla (LIRIS (UMR 5205 CNRS), INSA Lyon)
 - Luís Filipe Cunha (University of Minho)
 - Marc Spaniol (Université de Caen Normandie)
 - Marcelo Finger (University of São Paulo)
 - Mark Finlayson (Florida International University)
 - Natalia Vanetik (Sami Shamoon College of Engineering)
 - Nuno Guimarães (LIAAD – INESC TEC)
 - Pablo Gervás (Universidad Complutense de Madrid)
 - Paul Rayson (Lancaster University)
 - Paulo Quaresma (Universidade de Évora)
 - Purificação Silvano (Faculdade de Letras da Universidade do Porto)
 - Ricardo Marcacini (ICMC/USP)
 - Ross Purves (University of Zurich)
 - Sara Tonelli (Fondazione Bruno Kessler)
 - Sérgio Nunes (INESC TEC and Faculty of Engineering, University of Porto)
 - Sérgio Matos (IEETA, Universidade de Aveiro)
 - Thiago Pardo (University of São Paulo)
 - Udo Kruschwitz (University of Regensburg)
 - Valentina Bartalesi (ISTI-CNR)
 - Veysel Kocaman (John Snow Labs)
 - Yangyang Chen (Brandeis University)

4 Recognition Awards

The workshop ended with the announcement of the recognition awards. The Text2Story 2026 Best Paper Award went to the paper entitled *Agenda-based Narrative Extraction: Steering Pathfinding Algorithms with Large Language Models* by Brian Keith-Norambuena, Carolina Rojas-Córdova, Claudio Meneses-Villegas, Elizabeth LamEsquenazi, Angélica Flores-Bustos, Ignacio Molina-Villablanca, and Joshua Leyton-Vallejos. A runner-up Best Paper Award was given to the paper entitled *Semantic Interaction for Narrative Map Sensemaking: An Insight-based Evaluation* by Brian Keith, Fausto German, Eric Krokos, Sarah Joseph, and Chris North. In addition, the workshop chairs recognized the following researchers as recipients of the 2026 reviewer award for their insightful and valuable reviews:

- Álvaro Figueira
- Junbo Huang
- Veysel Kocaman
- Pablo Gervás
- Ricardo Marcacini

5 Concluding Thoughts

The Text2Story workshop series has successfully concluded its ninth iteration. Started with the aim of promoting interdisciplinary collaboration and defining the research agenda for the emerging field of narrative extraction from text, the series has continued to attract growing participation from the research community, indicating significant and sustained interest in this domain. The increasing prominence of generative AI and agentic systems, and their pervasive role in producing, summarising, and reasoning over text, has made the questions raised by the Text2Story community more important than ever. The contributions presented this year, ranging from interactive narrative extraction with LLM-guided pathfinding, evaluation methodologies for AI-generated stories and clinical AI, to position papers on human label variation and narrative-aware domain knowledge, illustrate the breadth and depth of these challenges. The workshop series will continue to provide a crucial platform for researchers to collaboratively shape the future direction of this evolving field.

Acknowledgments

The success of this workshop would not have been possible without the invaluable support of the ECIR 2026 organizers. We also extend our gratitude to the esteemed keynote speakers, Program Committee members, paper authors, and participants whose contributions significantly enriched the workshop experience.

References

- Nitin Agarwal and Bishwa Prakash Subedi. Symbolic framing meets diffusion modeling: How tariff narratives spread on social media. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper2.pdf>.
- Alex Argese, Pasquale Lisena, and Raphaël Troncy. Hallucination or creativity: How to evaluate AI-generated scientific stories? In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper7.pdf>.
- Alexandra Barancová and Giovanni Sileno. Narrative similarity detection using character roles. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper6.pdf>.
- Ricardo Campos, Alípio Jorge, Adam Jatowt, and Sumit Bhatia. The 3rd International Workshop on Narrative Extraction from Texts: Text2Story 2020. In *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.
- Ricardo Campos, Alípio Jorge, Adam Jatowt, Sumit Bhatia, and Mark Finlayson. The 4th International Workshop on Narrative Extraction from Texts: Text2Story 2021. In *Advances in Information Retrieval – 43rd European Conference on IR Research, ECIR 2021*, pages 701–704. Springer International Publishing, 2021.
- Ricardo Campos, Alípio Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak. The 5th International Workshop on Narrative Extraction from Texts: Text2Story 2022. In *Advances in Information Retrieval – 44th European Conference on IR Research, ECIR 2022, Stavanger, Norway, April 10–14, 2022, Proceedings, Part II*, volume 13186 of *Lecture Notes in Computer Science*, pages 552–556. Springer, 2022. doi: 10.1007/978-3-030-99739-7_68.
- Ricardo Campos, Alípio Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak. The 6th International Workshop on Narrative Extraction from Texts: Text2Story 2023. In *Advances in Information Retrieval – 45th European Conference on Information Retrieval, ECIR 2023, Dublin, Ireland, April 2–6, 2023, Proceedings, Part III*, volume 13982 of *Lecture Notes in Computer Science*, pages 377–383. Springer, 2023. doi: 10.1007/978-3-031-28241-6_40.
- Ricardo Campos, Alípio Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak. The 7th International Workshop on Narrative Extraction from Texts: Text2Story 2024. In *Advances in Information Retrieval – 46th European Conference on Information Retrieval, ECIR 2024, Glasgow, UK, March 24–28, 2024, Proceedings, Part V*, volume 14612 of *Lecture Notes in Computer Science*, pages 391–397. Springer, 2024. doi: 10.1007/978-3-031-56069-9_52.

-
- Ricardo Campos, Alípio Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak. The 8th international workshop on narrative extraction from texts: Text2Story 2025. In Claudia Hauff, Craig Macdonald, Dietmar Jannach, Gabriella Kazai, Franco Maria Nardini, Fabio Pinelli, Fabrizio Silvestri, and Nicola Tonellotto, editors, *Advances in Information Retrieval*, pages 345–351, Cham, 2025. Springer Nature Switzerland. ISBN 978-3-031-88720-8.
- Andrei Marian Feier, Veysel Kocaman, Yigit Gul, Ahmet Korkmaz, Alexander Thomas, Aleksei Zakharov, Jay Gil, Mehmet Butgul, and David Talby. A multi-domain red teaming framework for safety, robustness, and fairness evaluation of medical large language models. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper3.pdf>.
- Gaudenzia Genoni, Alfio Ferrara, and Stefano Montanelli. Evaluating large language models for character identification in Italian Renaissance epics: A case study on “Orlando Furioso”. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper8.pdf>.
- Kseniia Harshina, Tom Tucek, and Mathias Lux. Dynamic participatory game design with local AI: From interviews to trauma-aware interactive narratives. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper10.pdf>.
- Junbo Huang and Ricardo Usbeck. Collapsing distance: The curse of ground truth in computational narrative understanding. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper13.pdf>.
- Alípio M. Jorge, Ricardo Campos, Adam Jatowt, and Sérgio Nunes. First International Workshop on Narrative Extraction from Text (Text2Story’18). In *Advances in Information Retrieval – 40th European Conference on IR Research, ECIR 2018*, volume 10772 of *Lecture Notes in Computer Science*, pages 833–834. Springer, 2018.
- Alípio M. Jorge, Ricardo Campos, Adam Jatowt, and Sumit Bhatia. Second International Workshop on Narrative Extraction from Text (Text2Story’19). In *Advances in Information Retrieval – 41st European Conference on IR Research, ECIR 2019*, volume 11438 of *Lecture Notes in Computer Science*, pages 389–393. Springer, 2019a.
- Alípio M. Jorge, Ricardo Campos, Adam Jatowt, and Sérgio Nunes. Special issue on narrative extraction from texts (Text2Story): Preface. *Information Processing & Management*, 56(5): 1771–1774, 2019b.
- Milka Kaplan, Armin Shmilovici, and Mark Last. Structural patterns in award-winning novellas. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper9.pdf>.

-
- Brian Felipe Keith-Norambuena, Fausto German, Eric Krokos, Sarah Joseph, and Chris North. Semantic interaction for narrative map sensemaking: An insight-based evaluation. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026a. URL <https://ceur-ws.org/Vol-4202/paper4.pdf>.
- Brian Felipe Keith-Norambuena, Carolina Inés Rojas-Córdova, Claudio Juvenal Meneses-Villegas, Elizabeth Johanna Lam-Esquenazi, Angélica María Flores-Bustos, Ignacio Alejandro Molina-Villablanca, and Joshua Emanuel Leyton-Vallejos. Agenda-based narrative extraction: Steering pathfinding algorithms with large language models. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026b. URL <https://ceur-ws.org/Vol-4202/paper1.pdf>.
- José Luis López Calle and Pablo Gervás. Zero prompting segmentation of movie dialogues into narrative units. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper11.pdf>.
- Tahsir Ahmed Munna, Nuno Guimarães, and Alípio Jorge. Narrative-aware domain knowledge integration for annotation-free medical image localisation. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper14.pdf>.
- Cosimo Palma and Bertrand Jager. Rethinking TRIZ for AI-driven narrative generation. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper12.pdf>.
- Arian Pasquali. Evaluating generative AI and agentic systems: Methodologies for systematic assessment and continuous improvement. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/keynote2.pdf>.
- Federico Pianzola. Knowledge extraction from fiction using generative large language models and multilingual semantic parsing. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/keynote1.pdf>.
- Jonas Rieger, Felipe Muñoz, Lars Grönberg, Kai-Robin Lange, Iván Ojeda-Pereira, Darío Briceño, Christian Nass, Carsten Stahl, José Cassola, Carolina Rojas-Córdova, Brian Keith-Norambuena, Marcelo Lufin, Fernando Campos-Medina, and Sebastián Herrera-León. Following socio-environmental conflict narratives about energy transition in Chile: A spatio-temporal analysis using dynamic topic modeling. In *Proceedings of Text2Story – Ninth Workshop on Narrative Extraction From Texts*, volume 4202 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2026. URL <https://ceur-ws.org/Vol-4202/paper5.pdf>.