

NLLP 2026

Natural Legal Language Processing Workshop 2026

Proceedings of the Workshop

October 28, 2026

The NLLP organizers gratefully acknowledge the support from the following sponsors.

Gold

Bloomberg

Engineering



©2026 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
317 Sidney Baker St. S
Suite 400 - 134
Kerrville, TX 78028
USA
Tel: +1-855-225-1962
acl@aclweb.org

ISBN 979-8-89176-494-1

Introduction

Welcome to the eighth edition of the NLLP (Natural Legal Language Processing) Workshop, co-located with the 2026 Conference on Empirical Methods in Natural Language Processing (EMNLP). The NLLP workshop series explores methods and applications of NLP for the Legal Domain by focusing on original research on legal data and data with legal relevance. The rapid advancement of NLP is fundamentally transforming the legal domain, powering new capabilities in legal reasoning, information access, and decision support. Automation has become essential in the practice of law, and by its very nature, the practice of law (including its enforcement) necessarily involves the analysis and interpretation of language.

We organized this workshop to bring together researchers and practitioners from NLP, machine learning and other artificial intelligence disciplines together with legal practitioners, researchers and policy makers who develop or use NLP techniques. This is an exciting opportunity to expand the boundaries of our field by identifying new problems and exploring new data as it interacts with the full inventory of NLP and machine learning approaches. In this spirit, the Organizing and Program Committees were assembled to include researchers from both academia and industry, from NLP and legal backgrounds.

We were interested in the following types of papers: (1) Applications of NLP methods to tasks specific to the legal domain; (2) Legal Reasoning and Generation; (3) Tasks, Resources, Evaluation and Benchmarking for Legal NLP; (4) Methods for applying Large Language Models (LLMs) to the legal domain; (5) Methodological innovations for legal tasks; (6) Tasks, Resources and Evaluation for NLP in the Legal domain; (7) NLP for Online Platforms, Social Media and Regulations; (8) Systems, Demos and Industry Applications; and (9) Interdisciplinary position papers.

We received a record number of submissions. In total, we received 161 submissions, out of which we accepted for publication 51 papers for an overall acceptance rate of 31.6%. This includes 9 submissions from ARR (ACL Rolling Review) – papers that have received reviews and meta-review via the ARR system – of which 6 were accepted for publication. In keeping with past year's, we allowed authors to submit their papers for review for non-archival publication, which allows later publication in an archival venue and 9 papers selected this option.

This edition additionally featured a shared task: ChildSafeAds — Commercial Content in Child-Facing YouTube Videos that attracted 59 participants. The proceedings contain 4 shared task system descriptions and the shared task report. This task is about building components for a platform focused on commercial content that reaches minors on video platforms. The components of the platform classify what is being sold, categorize it, and reason about the advertising-law concerns it raises.

The papers in this year's edition of the workshop cover a total of 15 languages and jurisdictions, showcasing the broad reach and diversity of the workshop's topic. This includes new data sets for legal NLP in lower resource languages such as Nepali, Thai, Greek or Vietnamese.

Continuing the trend from past years, multiple papers covered the creation of new benchmarks for legal tasks including case/information retrieval, creating motion response pairs, obligation extraction, summarization and reasoning.

Application areas include information retrieval and search (case, evidence or statutory retrieval, compliance validation, grounding), information extraction, citation verification and inference, question answering, generation or classification of documents or clauses.

We note a specific focus this year on reasoning models and applications of LLMs for reasoning tasks. Additionally, several papers incorporate knowledge graphs or neuro-symbolic reasoning in LLMs for adapting them to legal tasks.

We remain excited to highlight the diverse array of legal texts examined in the workshop's papers including case law, complaints, statutory texts, acts, regulations, official gazettes, administrative law, contracts, NDAs, terms and conditions, licenses or legal essays.

We thank everyone who expressed interest in the workshop, all authors of submitted papers, the 86 members of the Program Committee who did an excellent job at reviewing papers given a short turnaround time, everyone attending the workshop and the EMNLP 2026 conference for hosting the

workshop. We thank our sponsors: Bloomberg and the European Research Council Starting Grant project HUMANads.

The NLLP Workshop organizers.

<http://nllpw.org>

Organizing Committee

Organizers

Nikolaos Aletras, The University of Sheffield

Ilias Chalkidis, University of Copenhagen

Leslie Barrett, Bloomberg Law

Cătălina Goanță, Utrecht University

Daniel Preoțiuc-Pietro, Bloomberg

Gerasimos Spanakis, Maastricht University

Shanshan Xu, University of Copenhagen

Program Committee

Reviewers

Sallam Abualhaisa, Tomaso Agnoloni, Ion Androutsopoulos, Tom Ault

Ilayda Balaban, Claire Barale, Thales Bertaglia, Floris Bex, Baldwin Breck

Luca Cagliero, Jiahong Chen, Odysseas Spyridon Chlapanis, Ashish Chouhan

Bram Duivenvoorde, Dominik Dworniczak, Arthur Dyeve

Dimitrios Galanis, Piyush Ghai

Ivan Habernal, Ben Hagag, Nils Holzenberger, Abe Hou

Abderrahmane Issam

Baaij Jaap

Constantinos Karouzos, Aykut Koc, Alice Kwak

Tong Liang, Ruta Liepina, Chu Luo

Megan Ma, Pawel Maka, Adam Meyers, Jelena Mitrović

Rohan Nanda, Joel Niklaus

Henrik Palmer Olsen, Katsikouli Panagiota, Ioannis Panagis, Anu Pradhan

Paulo Quaresma

T.Y.S.S. Santosh, Vageesh Saxena, Yusuf Can Semerci, Gil Semo, Madhavan Seshadri, Samyak Sheth, Dan Simonson, Jerrold Soh, Christoph Sorge, Alexandru Sotropa, Ieva Staliunaite

Dimitrios Tsarapatsanis, Kalpana Tyagi

Gijs Van Dijck, Helen van der Kroef

Jianqiang Wang, Hannes Westermann

ShanShan Xu, Huiyin Xue

Marcos Zampieri, Miri Zilka, Frederike Zufall

Table of Contents

<i>thylinao at ChildSafeAds 2026: A Six-Member Ensemble and What Each Level of Data Access Provided</i> Maksim Silchenko	1
<i>Can LLMs Time Travel? Enhancing Temporal Consistency in Legal Agentic Search through Reinforcement Learning</i> Wei Fan, ZHOU Yining, Mufan Zhang, WENG Yanbing, Yiran HU, Tianshi Zheng, Baixuan Xu, Chunyang Li, Jianhui Yang, Haoran Li and Yangqiu Song	12
<i>GRACE: Grounded Adversarial Reasoning over Canadian Law</i> Jiakang Xu, Wantong Huo, Udom Silparcha and Jonathan H. Chan	13
<i>OblEx26: A Multi-Jurisdictional Benchmark for Structured Obligation Extraction from Anti-Money Laundering Regulations</i> Philip Moore, Giovanni Bandi, Tim Luka Horstmann, Russell Moore, Robert H. Wardrop and Paula Buttery	14
<i>LegalPincite: Multi-level Legal Information Retrieval Dataset</i> Theresia Veronika Rampisela, Henrik Palmer Olsen and Giovanni Colavizza	29
<i>Persuading the Judge: Adversarial Robustness and RAG Mitigation in LLM-Based Legal Citation Verification</i> Rohann Rahul Shah, Rishabh Kartik Jawagal and Gowri Srinivasa	42
<i>Evidence-Grounded Contract Clause Classification with Small Language Models</i> Dina Mohamed Ali, hager hamed abd elrady allam, Hayam Abdelbaset Rezk, Rahatara Ferdousi and Mohammed Anber AbdElhamed	55
<i>Alphastream at ChildSafeAds 2026: Gated Per-Label Merging and an Advertiser-Precedent Rule for Rare Compliance Classes</i> Raghavendra Prasad Narayan	64
<i>Structuring license permissiveness from pairwise comparisons</i> Hamidah Oderinwale, Rachel Hong, David Atkinson, Art Abal and Benjamin Laufer	76
<i>Privacy Later: Reflections on Design Decisions, Lessons Learned, and Open Problems of the OPP-115 Corpus a Decade After Release</i> Shomir Wilson, Florian Schaub and Norman Sadeh	94
<i>The EUR-Lex Retrieval Dataset: A Section-Aware and Human-Annotated Dataset for Long-Form Similar Case Retrieval</i> Christof Koolen	101
<i>Verdict Without the Rule: Diagnosing and Auditing Regulatory Rule Sensitivity in LLM Compliance Systems</i> Saisab Sadhu, Aadit Sengupta, Vinay kumar Sankarapu and Pratinav Seth	119
<i>What the Public Asked to Change, and What It Would Touch? Complaint-to-Statute Grounding for Thailand's Regulatory Reform</i> Pawitsapak Akarajaradwong, Wuttikrai Lertprasertphakorn, Chompakorn Chaksangchaichot and Sarana Nutanong	133
<i>ARGUS: Role-Aware Event Knowledge Graphs for U.S. Employment-Discrimination Complaints</i> Sriram Kannan, Swetha Saseendran, Vishnu Vardhan Reddy Kandi, Leslie Barrett, Madhavan Seshadri and Enrico Santus	142

<i>Where Experts Disagree, Models Fail: Detecting Implicit Legal Citations in French Court Decisions</i> Avrile Floro, Tamara Dhorasoo, Soline Pellez and Nils Holzenberger	156
<i>Tensions Between Copyright, Contract, Accessibility and Technology in Australia: A Case Study on Legal Information Access</i> Hunter Briegel, Sumedh Vemuganti and J. Shane Culpepper	176
<i>A Strong Harness Narrows the Model Gap: A Comparative Evaluation for Legislative Document-to-XML Conversion</i> Palvika Bansal, Rittika Jindal and Ruichong Wang	183
<i>A Dataset of Motion-Response Text Pairs from US District Court Cases for Legal Judgment Prediction</i> Damian Curran, Eduard Hovy and Lea Frermann	199
<i>An Empirical Study of Evidence Retrieval for Retrieval Augmented Generation over the Thai Royal Gazette</i> Pantai Suyasri, Pawitsapak Akarajaradwong and Chompakorn Chaksangchaichot	212
<i>Terms and Conditions Apply: Evaluating Consumer LLM Reliability in NDA Interpretation to Investigate the Non-Linear Role of Reasoning</i> Simeon Sayer and Shaurya Jain Mathur	237
<i>Cross-Lingual Legal QA for Vietnamese Labour Law: Retrieval, Translation, and Verifier-Guided Correction</i> Nguyen Minh Chi, Mo El-Haj, Nguyen Ha Thanh, Dawn Knight and Paul Rayson	257
<i>Recall Is Not Completeness: Retrieving All the Evidence a Question Needs</i> Sebastien Kawada, Adhiban Arulselvan and Manolis Kellis	258
<i>AI-Monitors at ChildSafeAds 2026: A Routed Hybrid for Commercial Content in Child-Facing YouTube Videos</i> Pragyananda sahuo, Vishnu Tripathi, Prakhar joshi, Gaurav Kumar and Neel Mani	259
<i>Do Small Models Use the Law You Give Them? Measuring Context Use on a Bilingual Bangladesh Legal Benchmark</i> Moniruzzaman Mahadi, Abrar Mohammed Tanzim Alam, Mir Mohammad Asif Abdullah, Sayma Siddika Monalisa, Swakkhar Shatabda and Md Adnan Arefeen	271
<i>Beyond the Single-Expert Ceiling: Localizing LLM-as-Judge Stability on Free-Form Legal Essays</i> Pawitsapak Akarajaradwong, Wuttikrai Lertprasertphakorn, Chompakorn Chaksangchaichot and Sarana Nutanong	290
<i>What Drives an LLM's Legal Reasoning? Measuring Faithfulness at the Concept Level</i> Yassine Thlija, Lena Held and Ivan Habernal	303
<i>CLARE: A Criminal-Law DSL for Neuro-Symbolic Long-Form Legal Reasoning</i> Jaehoon Jeong and Jay-Yoon Lee	321
<i>CANLEGALRAGBENCH: Evaluating Retrieval-Augmented Generation on Canadian Case Law</i> Ethan Zhao, Maksym Taranukhin, Wei Cui, Moira Aikenhead and Vered Shwartz	322
<i>Nürnberg NLP at ChildSafeAds 2026: Structurally Dissimilar Voter Ensembles under Four Levels of Data Access</i> Philipp Steigerwald, Eric Rudolph and Jens Albrecht	323
<i>LegalBench.PT: A Benchmark for Portuguese Law</i> Beatriz Canaverde, Telmo Pires and Andre Martins	334

<i>Retrieval for Legal AI: Complementarity and Cross-Collection Interference in Brazilian Legal Collections</i>	
Flavio Soriano, Guilherme Henrique Gomes Evangelista, Celso França, Gisele L. Pappa, Wagner Meira Jr. and Marcos André Gonçalves	350
<i>GreekBarRetrieval: A Benchmark for Greek Statutory Retrieval</i>	
Ernest Beta, Odysseas S. Chlapanis, Dimitris Galanis and Ion Androutsopoulos	361
<i>Imitation Thresholds: Quantifying Influence Propagation in Large Language Models</i>	
Gil Semo, Dor Bernsohn, Lior Magen, Kfir Bar and Gila Hayat	381
<i>Discovering Hierarchical Semantic Structures in EU Regulatory Embeddings: A Comparison of Two K-means Solutions and Their Alignment with EuroVoc</i>	
Kayoko Fukasawa	382
<i>Deka Retrieval: Component-Level Criteria for Mining Legal Case Pairs in Thai Supreme Court Decisions</i>	
Pasin Buakhaw, Pawitsapak Akarajaradwong and Chompakorn Chaksangchaichot	399
<i>FIRAC-APPEALS: A Structured Indian Appellate Judgments Dataset for Summarisation and Reasoning-Aware Legal NLP</i>	
Martin Borčín and Joemon M. Jose	420
<i>Outside the Four Corners: Charting the Missing Evidence in Korean AI Training-Data License Screening</i>	
Jiwoo Lee and Sungwon Chae	432
<i>GapToCompliance: A Taxonomy-Driven Agentic AI Framework for Automated AI Act Gap Analysis and Remediation</i>	
Ilias Driouich, Hongliu Cao and Eoin Thomas	456
<i>The Bluebook as Code: A Deterministic Engine for Legal Citation</i>	
Eric H Martínez and Matthew Dahl	468
<i>OPENER: A Frugal Open-World NER Pipeline for the Legal Domain</i>	
Thibault GAREL, Faiza BELBACHIR and Assia Soukane	469
<i>Building Structured Sentencing Data from Federal Criminal Docket Entries: An Ontology and Entity Extraction Study</i>	
David Guecha, Anan Biswas and Charlotte S. Alexander	484
<i>Cited but Not Consulted: A Counterfactual Audit of Legal Chain-of-Thought Faithfulness</i>	
Saisab Sadhu, Shreeyans Arora and Pratinav Seth	508
<i>NELSR: A Nepali Legal Benchmark for Case-to-Statute Retrieval</i>	
Anjan Parajuli and Darshan Kunwar	524
<i>Evaluating LLMs on Explicit Knowledge of Italian Legal Concepts</i>	
Mattia Proietti, Lucia Passaro and Alessandro Lenci	535
<i>Enhancing Legal Reasoning in Pre-trained Language Models via Knowledge Graph-Guided Multi-Task Pre-training</i>	
Ibrahim Siddig and Munir Georges	548
<i>Legal Experts Disagree With Rationale Extraction Techniques for Explaining ECtHR Case Outcome Classification</i>	
Mahammad Namazov, Tomáš Koref and Ivan Habernal	562

<i>Protecting De-identified Documents from Search-based Linkage Attacks</i>	
Pierre Lison and Mark Anderson	587
<i>UCVE: Graph Structured Retrieval for Compliance Validation</i>	
Emily Chen, Michelle Menachery, Rahul Nenavath, Himanshu Kumar and Amirfarrokh Iranitalab	
609	
<i>Right Code, Wrong Reasons: An Audit Trail Benchmark for Customs Classification</i>	
Kancharla Aditya Hari, Sai Khadloya, Arghya Bhattacharya and Kumarmanas Nethil	628
<i>From Legal Text to Executable Process: LLM-Based Extraction of DCR Determinations from Danish Administrative Law</i>	
Ali Al-Laith and Manex Agirrezabal	640
<i>ChildSafeAds Shared Task 2026: Commercial Content in Child-Facing YouTube Videos</i>	
Thales Bertaglia, Catalina Goanta, Gerasimos Spanakis and Gunes Acar	641

Program

Wednesday, October 28, 2026

09:00 - 17:00 *Papers*

Privacy Later: Reflections on Design Decisions, Lessons Learned, and Open Problems of the OPP-115 Corpus a Decade After Release

Shomir Wilson, Florian Schaub and Norman Sadeh

Imitation Thresholds: Quantifying Influence Propagation in Large Language Models

Gil Semo, Dor Bernsohn, Lior Magen, Kfir Bar and Gila Hayat

Where Experts Disagree, Models Fail: Detecting Implicit Legal Citations in French Court Decisions

Avrile Floro, Tamara Dhorasoo, Soline Pellez and Nils Holzenberger

NELSR: A Nepali Legal Benchmark for Case-to-Statute Retrieval

Anjan Parajuli and Darshan Kunwar

LegalBench.PT: A Benchmark for Portuguese Law

Beatriz Canaverde, Telmo Pires and Andre Martins

Can LLMs Time Travel? Enhancing Temporal Consistency in Legal Agentic Search through Reinforcement Learning

Wei Fan, ZHOU Yining, Mufan Zhang, WENG Yanbing, Yiran HU, Tianshi Zheng, Baixuan Xu, Chunyang Li, Jianhui Yang, Haoran Li and Yangqiu Song

The EUR-Lex Retrieval Dataset: A Section-Aware and Human-Annotated Dataset for Long-Form Similar Case Retrieval

Christof Koolen

A Strong Harness Narrows the Model Gap: A Comparative Evaluation for Legislative Document-to-XML Conversion

Palvika Bansal, Rittika Jindal and Ruichong Wang

Tensions Between Copyright, Contract, Accessibility and Technology in Australia: A Case Study on Legal Information Access

Hunter Briegel, Sumedh Vemuganti and J. Shane Culpepper

From Legal Text to Executable Process: LLM-Based Extraction of DCR Determinations from Danish Administrative Law

Ali Al-Laith and Manex Agirrezabal

Wednesday, October 28, 2026 (continued)

Structuring license permissiveness from pairwise comparisons

Hamidah Oderinwale, Rachel Hong, David Atkinson, Art Abal and Benjamin Laufer

OPENER: A Frugal Open-World NER Pipeline for the Legal Domain

Thibault GAREL, Faiza BELBACHIR and Assia Soukane

LegalPincite: Multi-level Legal Information Retrieval Dataset

Theresia Veronika Rampisela, Henrik Palmer Olsen and Giovanni Colavizza

A Dataset of Motion-Response Text Pairs from US District Court Cases for Legal Judgment Prediction

Damian Curran, Eduard Hovy and Lea Frermann

An Empirical Study of Evidence Retrieval for Retrieval Augmented Generation over the Thai Royal Gazette

Pantai Suyasri, Pawitsapak Akarajardwong and Chompakorn Chaksangchaichot

Persuading the Judge: Adversarial Robustness and RAG Mitigation in LLM-Based Legal Citation Verification

Rohann Rahul Shah, Rishabh Kartik Jawagal and Gowri Srinivasa

Cross-Lingual Legal QA for Vietnamese Labour Law: Retrieval, Translation, and Verifier-Guided Correction

Nguyen Minh Chi, Mo El-Haj, Nguyen Ha Thanh, Dawn Knight and Paul Rayson

Do Small Models Use the Law You Give Them? Measuring Context Use on a Bilingual Bangladesh Legal Benchmark

Moniruzzaman Mahadi, Abrar Mohammed Tanzim Alam, Mir Mohammad Asif Abdullah, Sayma Siddika Monalisa, Swakkhar Shatabda and Md Adnan Arefeen

GapToCompliance: A Taxonomy-Driven Agentic AI Framework for Automated AI Act Gap Analysis and Remediation

Ilias Driouich, Hongliu Cao and Eoin Thomas

GreekBarRetrieval: A Benchmark for Greek Statutory Retrieval

Ernest Beta, Odysseas S. Chlapanis, Dimitris Galanis and Ion Androutsopoulos

Recall Is Not Completeness: Retrieving All the Evidence a Question Needs

Sebastien Kawada, Adhiban Arulselvan and Manolis Kellis

Wednesday, October 28, 2026 (continued)

Discovering Hierarchical Semantic Structures in EU Regulatory Embeddings: A Comparison of Two K-means Solutions and Their Alignment with EuroVoc

Kayoko Fukasawa

What Drives an LLM's Legal Reasoning? Measuring Faithfulness at the Concept Level

Yassine Thlija, Lena Held and Ivan Habernal

Terms and Conditions Apply: Evaluating Consumer LLM Reliability in NDA Interpretation to Investigate the Non-Linear Role of Reasoning

Simeon Sayer and Shaurya Jain Mathur

OblEx26: A Multi-Jurisdictional Benchmark for Structured Obligation Extraction from Anti-Money Laundering Regulations

Philip Moore, Giovanni Bandi, Tim Luka Horstmann, Russell Moore, Robert H. Wardrop and Paula Buttery

Evidence-Grounded Contract Clause Classification with Small Language Models

Dina Mohamed Ali, hager hamed abd elrady allam, Hayam Abdelbaset Rezk, Rahatara Ferdousi and Mohammed Anber AbdElhamed

CANLEGALRAGBENCH: Evaluating Retrieval-Augmented Generation on Canadian Case Law

Ethan Zhao, Maksym Taranukhin, Wei Cui, Moira Aikenhead and Vered Shwartz

ARGUS: Role-Aware Event Knowledge Graphs for U.S. Employment-Discrimination Complaints

Sriram Kannan, Swetha Saseendran, Vishnu Vardhan Reddy Kandi, Leslie Barrett, Madhavan Seshadri and Enrico Santus

Building Structured Sentencing Data from Federal Criminal Docket Entries: An Ontology and Entity Extraction Study

David Guecha, Anan Biswas and Charlotte S. Alexander

Retrieval for Legal AI: Complementarity and Cross-Collection Interference in Brazilian Legal Collections

Flavio Soriano, Guilherme Henrique Gomes Evangelista, Celso França, Gisele L. Pappa, Wagner Meira Jr. and Marcos André Gonçalves

The Bluebook as Code: A Deterministic Engine for Legal Citation

Eric H Martínez and Matthew Dahl

Cited but Not Consulted: A Counterfactual Audit of Legal Chain-of-Thought Faithfulness

Saisab Sadhu, Shreeyans Arora and Pratinav Seth

Wednesday, October 28, 2026 (continued)

Verdict Without the Rule: Diagnosing and Auditing Regulatory Rule Sensitivity in LLM Compliance Systems

Saisab Sadhu, Aadit Sengupta, Vinay kumar Sankarapu and Pratinav Seth

CLARE: A Criminal-Law DSL for Neuro-Symbolic Long-Form Legal Reasoning

Jaehoon Jeong and Jay-Yoon Lee

GRACE: Grounded Adversarial Reasoning over Canadian Law

Jiakang Xu, Wantong Huo, Udom Silparcha and Jonathan H. Chan

What the Public Asked to Change, and What It Would Touch? Complaint-to-Statute Grounding for Thailand's Regulatory Reform

Pawitsapak Akarajaradwong, Wuttikrai Lertprasertphakorn, Chompakorn Chaksangchaichot and Sarana Nutanong

Beyond the Single-Expert Ceiling: Localizing LLM-as-Judge Stability on Free-Form Legal Essays

Pawitsapak Akarajaradwong, Wuttikrai Lertprasertphakorn, Chompakorn Chaksangchaichot and Sarana Nutanong

Deka Retrieval: Component-Level Criteria for Mining Legal Case Pairs in Thai Supreme Court Decisions

Pasin Buakhaw, Pawitsapak Akarajaradwong and Chompakorn Chaksangchaichot

Outside the Four Corners: Charting the Missing Evidence in Korean AI Training-Data License Screening

Jiwoo Lee and Sungwon Chae

FIRAC-APPEALS: A Structured Indian Appellate Judgments Dataset for Summarisation and Reasoning-Aware Legal NLP

Martin Borčín and Joemon M. Jose

thylinao at ChildSafeAds 2026: A Six-Member Ensemble and What Each Level of Data Access Provided

Maksim Silchenko

AI-Monitors at ChildSafeAds 2026: A Routed Hybrid for Commercial Content in Child-Facing YouTube Videos

Pragyananda sahu, Vishnu Tripathi, Prakhar joshi, Gaurav Kumar and Neel Mani

Nürnberg NLP at ChildSafeAds 2026: Structurally Dissimilar Voter Ensembles under Four Levels of Data Access

Philipp Steigerwald, Eric Rudolph and Jens Albrecht

Wednesday, October 28, 2026 (continued)

Alphastream at ChildSafeAds 2026: Gated Per-Label Merging and an Advertiser-Precedent Rule for Rare Compliance Classes

Raghavendra Prasad Narayan

ChildSafeAds Shared Task 2026: Commercial Content in Child-Facing YouTube Videos

Thales Bertaglia, Catalina Goanta, Gerasimos Spanakis and Gunes Acar

Evaluating LLMs on Explicit Knowledge of Italian Legal Concepts

Mattia Proietti, Lucia Passaro and Alessandro Lenci

Enhancing Legal Reasoning in Pre-trained Language Models via Knowledge Graph-Guided Multi-Task Pre-training

Ibrahim Siddig and Munir Georges

Legal Experts Disagree With Rationale Extraction Techniques for Explaining EC-tHR Case Outcome Classification

Mahammad Namazov, Tomáš Koref and Ivan Habernal

Protecting De-identified Documents from Search-based Linkage Attacks

Pierre Lison and Mark Anderson

UCVE: Graph Structured Retrieval for Compliance Validation

Emily Chen, Michelle Menachery, Rahul Nenavath, Himanshu Kumar and Amir-farrokh Iranitalab

Right Code, Wrong Reasons: An Audit Trail Benchmark for Customs Classification

Kancharla Aditya Hari, Sai Khadloya, Arghya Bhattacharya and Kumarmanas Nethil