

PaperID	Title	Proceedings	Format
2	Structured Pruning for Large Language Models Using Coupled Components Elimination and Minor Fine-tuning	Findings	Virtual Poster
3	Named Entity Recognition Under Domain Shift via Metric Learning for Life Sciences	Main Conference	Oral
4	Weight-Inherited Distillation for Task-Agnostic BERT Compression	Findings	Virtual Poster
5	Text Diffusion Model with Encoder-Decoder Transformers for Sequence-to-Sequence Generation	Main Conference	Oral
6	An Interactive Framework for Profiling News Media Sources	Main Conference	Poster
7	Ignore Me But Don't Replace Me: Utilizing Non-Linguistic Elements for Pretraining on the Cybersecurity Domain	Findings	Poster
8	Assessing Logical Puzzle Solving in Large Language Models: Insights from a Minesweeper Case Study	Main Conference	Poster
9	TelME: Teacher-leading Multimodal Fusion Network for Emotion Recognition in Conversation	Main Conference	Poster
14	Effective and Efficient Conversation Retrieval for Dialogue State Tracking with Implicit Text Summaries	Main Conference	Poster
15	Promptly Predicting Structures: The Return of Inference	Main Conference	Poster
16	On Linearizing Structured Data in Encoder-Decoder Language Models: Insights from Text-to-SQL	Main Conference	Poster
17	Extremely efficient online query encoding for dense retrieval	Findings	Poster
18	Extractive Summarization with Text Generator	Main Conference	Poster
19	Self-generated Replay Memories for Continual Neural Machine Translation	Main Conference	Oral
20	Measuring and Improving Chain-of-Thought Reasoning in Vision-Language Models	Main Conference	Virtual Poster
21	Revisiting Zero-Shot Abstractive Summarization in the Era of Large Language Models from the Perspective of Position Bias	Main Conference	Oral
22	Building Knowledge-Guided Lexica to Model Cultural Variation	Main Conference	Oral
23	Adaptive Rank Selections for Low-Rank Approximation of Language Models	Main Conference	Virtual Poster
25	An Empirical Study of Consistency Regularization for End-to-End Speech-to-Text Translation	Main Conference	Virtual Poster
26	Unleashing the Emergent Cognitive Synergy in Large Language Models: A Task-Solving Agent through Multi-Persona Self-Collaboration	Main Conference	Virtual Poster
27	FPT: Feature Prompt Tuning for Few-shot Readability Assessment	Main Conference	Oral
28	Struc-Bench: Are Large Language Models Really Good at Generating Complex Structured Data?	Main Conference	Poster
29	Self-Prompting Large Language Models for Zero-Shot Open-Domain QA	Main Conference	Virtual Poster
30	Head-to-Tail: How Knowledgeable are Large Language Models (LLMs)? A.K.A. Will LLMs Replace Knowledge Graphs?	Main Conference	Poster
31	\$\$k\$N-ICL: Compositional Task-Oriented Parsing Generalization with Nearest Neighbor In-Context Learning	Main Conference	Virtual Poster
32	DIVKNOWQA: Assessing the Reasoning Ability of LLMs via Open-Domain Question Answering over Knowledge Base and Text	Findings	Virtual Poster
33	ARES: An Automated Evaluation Framework for Retrieval-Augmented Generation Systems	Main Conference	Oral
34	DEMO: A Statistical Perspective for Efficient Image-Text Matching	Main Conference	Virtual Poster
35	SeaEval for Multilingual Foundation Models: From Cross-Lingual Alignment to Cultural Reasoning	Main Conference	Virtual Poster
36	Volcano: Mitigating Multimodal Hallucination through Self-Feedback Guided Revision	Main Conference	Poster
37	LLMs Are Few-Shot In-Context Low-Resource Language Learners	Main Conference	Poster
38	SpeedE: Euclidean Geometric Knowledge Graph Embedding Strikes Back	Findings	Poster
40	Simple and effective data augmentation for compositional generalization	Main Conference	Poster
41	Language Guided Exploration for RL Agents in Text Environments	Findings	Poster
42	GPT-who: An Information Density-based Machine-Generated Text Detector	Findings	Poster
43	Rethinking Tabular Data Understanding with Large Language Models	Main Conference	Poster
44	From Shortcuts to Triggers: Backdoor Defense with Denoised PoE	Main Conference	Poster
45	BookSQL: A Large Scale Text-to-SQL Dataset for Accounting Domain	Main Conference	Poster
46	Improving Toponym Resolution by Predicting Attributes to Constrain Geographical Ontology Entries	Main Conference	Virtual Poster
47	FLAP: Flow-Adhering Planning with Constrained Decoding in LLMs	Main Conference	Oral
48	DuRE: Dual Contrastive Self Training for Semi-Supervised Relation Extraction	Main Conference	Poster

PaperID	Title	Proceedings	Format
49	Query-Efficient Textual Adversarial Example Generation for Black-Box Attacks	Main Conference	Virtual Poster
50	Embrace Divergence for Richer Insights: A Multi-document Summarization Benchmark and a Case Study on Summarizing Diverse Information from News Articles	Main Conference	Poster
51	AMRFact: Enhancing Summarization Factuality Evaluation with AMR-Driven Negative Samples Generation	Main Conference	Oral
52	PILOT: Legal Case Outcome Prediction with Case Law	Main Conference	Oral
53	ALoRA: Allocating Low-Rank Adaptation for Fine-tuning Large Language Models	Main Conference	Virtual Poster
55	R-Spin: Efficient Speaker and Noise-invariant Representation Learning with Acoustic Pieces	Main Conference	Poster
57	DEED: Dynamic Early Exit on Decoder for Accelerating Encoder-Decoder Transformer Models	Findings	Poster
58	InsCL: A Data-efficient Continual Learning Paradigm for Fine-tuning Large Language Models with Instructions	Main Conference	Virtual Poster
60	Language Agnostic Code Embeddings	Main Conference	Virtual Poster
61	Advancing Regular Language Reasoning in Linear Recurrent Neural Networks	Main Conference	Virtual Poster
63	Attention Alignment and Flexible Positional Embeddings Improve Transformer Length Extrapolation	Findings	Virtual Poster
64	Automatic Pair Construction for Contrastive Post-training	Findings	Poster
65	An Examination of the Compositionality of Large Generative Vision-Language Models	Main Conference	Virtual Poster
66	Extracting Lexical Features from Dialects via Interpretable Dialect Classifiers	Main Conference	Poster
68	Two Heads are Better than One: Nested PoE for Robust Defense Against Multi-Backdoors	Main Conference	Oral
69	Clear Up Confusion: Advancing Cross-Domain Few-Shot Relation Extraction through Relation-Aware Prompt Learning	Main Conference	Virtual Poster
70	Self-Checker: Plug-and-Play Modules for Fact-Checking with Large Language Models	Findings	Virtual Poster
71	VertAttack: Taking Advantage of Text Classifiers' Horizontal Vision	Main Conference	Oral
72	Fusion Makes Perfection: An Efficient Multi-Grained Matching Approach for Zero-Shot Relation Extraction	Main Conference	Virtual Poster
74	Personalized Review Recommendation based on Implicit dimension mining	Main Conference	Virtual Poster
75	KDMCSE: Knowledge Distillation Multimodal Sentence Embeddings with Adaptive Angular margin Contrastive Learning	Main Conference	Virtual Poster
77	Low-resource neural machine translation with morphological modeling	Findings	Virtual Poster
78	Self-Cleaning: Improving a Named Entity Recognizer Trained on Noisy Data with a Few Clean Instances	Findings	Virtual Poster
80	The taste of IPA: Towards open-vocabulary keyword spotting and forced alignment in any language	Main Conference	Poster
82	Think Before You Act: A Two-Stage Framework for Mitigating Gender Bias Towards Vision-Language Tasks	Main Conference	Virtual Poster
83	VLUE: A New Benchmark and Multi-task Knowledge Transfer Learning for Vietnamese Natural Language Understanding	Findings	Virtual Poster
84	LETI: Learning to Generate from Textual Interactions	Findings	Poster
85	BeLLM: Backward Dependency Enhanced Large Language Model for Sentence Embeddings	Main Conference	Oral
86	Bilateral Masking with prompt for Knowledge Graph Completion	Findings	Poster
88	Assessing Factual Reliability of Large Language Model Knowledge	Main Conference	Oral
91	Dial-MAE: ConTextual Masked Auto-Encoder for Retrieval-based Dialogue Systems	Main Conference	Virtual Poster
92	MiLe Loss: a New Loss for Mitigating the Bias of Learning Difficulties in Generative Language Models	Findings	Virtual Poster
93	Toolink: Linking Toolkit Creation and Using through Chain-of-Solving on Open-Source Model	Main Conference	Virtual Poster
95	Create! Don't Repeat: A Paradigm Shift in Multi-Label Augmentation through Label Creative Generation	Main Conference	Virtual Poster
96	Neurocache: Efficient Vector Retrieval for Long-range Language Modeling	Main Conference	Poster
97	Unveiling the Generalization Power of Fine-Tuned Large Language Models	Main Conference	Poster
98	A Closer Look at the Self-Verification Abilities of Large Language Models in Logical Reasoning	Main Conference	Virtual Poster
99	Exploring Self-supervised Logic-enhanced Training for Large Language Models	Main Conference	Virtual Poster
100	MATHSENSEI: A Tool-Augmented Large Language Model for Mathematical Reasoning	Main Conference	Poster
101	GOLD: Geometry Problem Solver with Natural Language Description	Findings	Virtual Poster

PaperID	Title	Proceedings	Format
102	RoDia: A New Dataset for Romanian Dialect Identification from Speech	Findings	Virtual Poster
103	CoUDA: Coherence Evaluation via Unified Data Augmentation	Main Conference	Virtual Poster
104	Examining Modularity in Multilingual LMs via Language-Specialized Subnetworks	Findings	Poster
106	mEdit: Multilingual Text Editing via Instruction Tuning	Main Conference	Poster
107	Navigation as Attackers Wish? Towards Building Robust Embodied Agents under Federated Learning	Main Conference	Poster
108	In-context Learning and Gradient Descent Revisited	Main Conference	Poster
109	Corpus Considerations for Annotator Modeling and Scaling	Main Conference	Poster
110	Reverse Chain: A Generic-Rule for LLMs to Master Multi-API Planning	Findings	Virtual Poster
111	Incorporating Exponential Smoothing into MLP: a Simple but Effective Sequence Model	Findings	Poster
113	OpenFMNav: Towards Open-Set Zero-Shot Object Navigation via Vision-Language Foundation Models	Findings	Virtual Poster
114	On Large Language Models' Hallucination with Regard to Known Facts	Main Conference	Poster
116	One-Size-Fits-All? Examining Expectations around What Constitute "Fair" or "Good" NLG System Behaviors	Main Conference	Oral
117	Language Models Hallucinate, but May Excel at Fact Verification	Main Conference	Poster
118	Comparing Two Model Designs for Clinical Note Generation; Is an LLM a Useful Evaluator of Consistency?	Findings	Virtual Poster
119	VOLTA: Improving Generative Diversity by Variational Mutual Information Maximizing Autoencoder	Findings	Poster
120	EcoSpeak: Cost-Efficient Bias Mitigation for Partially Cross-Lingual Speaker Verification	Findings	Virtual Poster
121	A Rationale-centric Counterfactual Data Augmentation Method for Cross-Document Event Coreference Resolution	Main Conference	Virtual Poster
122	TrojFSP: Trojan Insertion in Few-shot Prompt Tuning	Main Conference	Oral
123	Ensuring Safe and High-Quality Outputs: A Guideline Library Approach for Language Models	Main Conference	Virtual Poster
124	Leveraging Contextual Information for Effective Entity Salience Detection	Findings	Poster
125	X-PARADE: Cross-Lingual Textual Entailment and Information Divergence across Paragraphs	Main Conference	Oral
126	Topics, Authors, and Institutions in Large Language Model Research: Trends from 17K arXiv Papers	Main Conference	Poster
127	Unlocking Structure Measuring: Introducing PDD, an Automatic Metric for Positional Discourse Coherence	Main Conference	Oral
128	\$E^5\$: Zero-shot Hierarchical Table Analysis using Augmented LLMs via Explain, Extract, Execute, Exhibit and Extrapolate	Main Conference	Poster
130	S3Eval: A Synthetic, Scalable, Systematic Evaluation Suite for Large Language Model	Main Conference	Poster
131	MMC: Advancing Multimodal Chart Understanding with Large-scale Instruction Tuning	Main Conference	Virtual Poster
132	Visual Grounding Helps Learn Word Meanings in Low-Data Regimes	Main Conference	Oral
133	Accurate Knowledge Distillation via n-best Reranking	Main Conference	Poster
134	AutoPRM: Automating Procedural Supervision for Multi-Step Reasoning via Controllable Question Decomposition	Main Conference	Virtual Poster
135	SEMQA: Semi-Extractive Multi-Source Question Answering	Main Conference	Oral
136	Fine-Tuning Language Models with Reward Learning on Policy	Main Conference	Virtual Poster
137	LLM-as-a-Coach: Can Mixed Human-Written and Machine-Generated Text Be Detected?	Findings	Virtual Poster
138	A (More) Realistic Evaluation Setup for Generalisation of Community Models on Malicious Content Detection	Findings	Virtual Poster
139	Citation: A Key to Building Responsible and Accountable Large Language Models	Findings	Virtual Poster
140	A Universal Dependencies Treebank for Highland Puebla Nahuatl	Main Conference	Oral
141	COPAL-ID: Indonesian Language Reasoning with Local Culture and Nuances	Main Conference	Poster
142	IterAlign: Iterative Constitutional Alignment of Large Language Models	Main Conference	Poster
143	OrchestraLLM: Efficient Orchestration of Language Models for Dialogue State Tracking	Main Conference	Poster
144	Multi-Operational Mathematical Derivations in Latent Space	Main Conference	Poster
145	Graph-Induced Syntactic-Semantic Spaces in Transformer-Based Variational AutoEncoders	Findings	Poster
146	Large Language Models Help Humans Verify Truthfulness – Except When They Are Convincingly Wrong	Main Conference	Poster

PaperID	Title	Proceedings	Format
147	XferBench: a Data-Driven Benchmark for Emergent Language	Main Conference	Poster
148	Narrowing the Gap between Zero- and Few-shot Machine Translation by Matching Styles	Findings	Poster
149	Which Modality should I use - Text, Motif, or Image? : Understanding Graphs with Large Language Models	Findings	Poster
150	Evaluating Large Language Models as Generative User Simulators for Conversational Recommendation	Main Conference	Poster
151	A Symbolic Framework for Evaluating Mathematical Reasoning and Generalisation with Transformers	Main Conference	Poster
154	On-the-Fly Fusion of Large Language Models and Machine Translation	Findings	Poster
155	READ: Improving Relation Extraction from an ADversarial Perspective	Findings	Virtual Poster
156	Identifying Linear Relational Concepts in Large Language Models	Main Conference	Oral
157	REQUAL-LM: Reliability and Equity through Aggregation in Large Language Models	Findings	Poster
158	Benchmark Transparency: Measuring the Impact of Data on Evaluation	Main Conference	Poster
159	JAMDEC: Unsupervised Authorship Obfuscation using Constrained Decoding over Small Language Models	Main Conference	Oral
160	Returning to the Start: Generating Narratives with Related Endpoints	Main Conference	Poster
161	REST: Retrieval-Based Speculative Decoding	Main Conference	Virtual Poster
162	Sub-Sentence Encoder: Contrastive Learning of Propositional Semantic Representations	Main Conference	Poster
163	MSciNLI: A Diverse Benchmark for Scientific Natural Language Inference	Main Conference	Virtual Poster
164	Causal Inference for Human-Language Model Collaboration	Main Conference	Poster
165	SELF-GUARD: Empower the LLM to Safeguard Itself	Main Conference	Poster
166	COSIGN: Contextual Facts Guided Generation for Knowledge Graph Completion	Main Conference	Virtual Poster
167	Toward Informal Language Processing: Knowledge of Slang in Large Language Models	Main Conference	Oral
168	Ghostbuster: Detecting Text Ghostwritten by Large Language Models	Main Conference	Poster
170	End-to-End Beam Retrieval for Multi-Hop Question Answering	Main Conference	Virtual Poster
171	Unified Examination of Entity Linking in Absence of Candidate Sets	Main Conference	Virtual Poster
172	Leveraging Generative Large Language Models with Visual Instruction and Demonstration Retrieval for Multimodal Sarcasm Detection	Main Conference	Virtual Poster
173	Multi-Scale Prompt Memory-Augmented Model for Black-Box Scenarios	Main Conference	Virtual Poster
174	Ungrammatical-syntax-based In-context Example Selection for Grammatical Error Correction	Main Conference	Virtual Poster
175	Addressing Both Statistical and Causal Gender Fairness in NLP Models	Findings	Poster
177	LLM-Rec: Personalized Recommendation via Prompting Large Language Models	Findings	Virtual Poster
178	A Robust Semantics-based Watermark for Large Language Model against Paraphrasing	Findings	Poster
179	BUFFET: Benchmarking Large Language Models for Few-shot Cross-lingual Transfer	Main Conference	Oral
180	TISE: A Tripartite In-context Selection Method for Event Argument Extraction	Main Conference	Virtual Poster
182	Solving Data-centric Tasks using Large Language Models	Findings	Poster
183	Reasoning or Reciting? Exploring the Capabilities and Limitations of Language Models Through Counterfactual Tasks	Main Conference	Oral
185	TRUE-UIE: Two Universal Relations Unify Information Extraction Tasks	Main Conference	Poster
186	zrLLM: Zero-Shot Relational Learning on Temporal Knowledge Graphs with Large Language Models	Main Conference	Oral
187	Embodied Executable Policy Learning with Language-based Scene Summarization	Main Conference	Poster
188	Metacognitive Prompting Improves Understanding in Large Language Models	Main Conference	Virtual Poster
190	A Novel Paradigm Boosting Translation Capabilities of Large Language Models	Findings	Poster
191	MART: Improving LLM Safety with Multi-round Automatic Red-Teaming	Main Conference	Virtual Poster
192	DialogCC: An Automated Pipeline for Creating High-Quality Multi-Modal Dialogue Dataset	Main Conference	Poster
194	Routing to the Expert: Efficient Reward-guided Ensemble of Large Language Models	Main Conference	Virtual Poster
195	Measuring Social Norms of Large Language Models	Findings	Virtual Poster

PaperID	Title	Proceedings	Format
196	Automatic Generation of Model and Data Cards: A Step Towards Responsible AI	Main Conference	Poster
197	FUN with Fisher: Improving Generalization of Adapter-Based Cross-lingual Transfer with Scheduled Unfreezing	Main Conference	Poster
198	Are Multilingual LLMs Culturally-Diverse Reasoners? An Investigation into Multicultural Proverbs and Sayings	Main Conference	Poster
200	The Colorful Future of LLMs: Evaluating and Improving LLMs as Emotional Supporters for Queer Youth	Main Conference	Poster
201	IPED: An Implicit Perspective for Relational Triple Extraction based on Diffusion Model	Main Conference	Poster
202	QualEval: Qualitative Evaluation for Model Improvement	Main Conference	Oral
203	Quantum-inspired Language Model with Lindblad Master Equation and Interference Measurement for Sentiment Analysis	Main Conference	Virtual Poster
204	MultiParaDetox: Extending Text Detoxification with Parallel Data to New Languages	Main Conference	Poster
205	VisLingInstruct: Elevating Zero-Shot Learning in Multi-Modal Language Models with Autonomous Instruction Optimization	Main Conference	Virtual Poster
206	A Wolf in Sheep's Clothing: Generalized Nested Jailbreak Prompts can Fool Large Language Models Easily	Main Conference	Oral
207	P ³ Sum: Preserving Author's Perspective in News Summarization with Diffusion Language Models	Main Conference	Virtual Poster
208	Bridging the Novice-Expert Gap via Models of Decision-Making: A Case Study on Remediating Math Mistakes	Main Conference	Poster
209	Source-Free Unsupervised Domain Adaptation for Question Answering via Prompt-Assisted Self-learning	Findings	Poster
210	SKICSE: Sentence Knowable Information Prompted by LLMs Improves Contrastive Sentence Embeddings	Main Conference	Virtual Poster
211	RST-LoRA: A Discourse-Aware Low-Rank Adaptation for Long Document Abstractive Summarization	Main Conference	Oral
212	Hierarchical Attention Graph for Scientific Document Summarization in Global and Local Level	Findings	Virtual Poster
213	LEETs-Dial: Linguistic Entrainment in End-to-End Task-oriented Dialogue systems	Findings	Poster
214	Strings from the Library of Babel: Random Sampling as a Strong Baseline for Prompt Optimisation	Main Conference	Virtual Poster
215	ReTA: Recursively Thinking Ahead to Improve the Strategic Reasoning of Large Language Models	Main Conference	Poster
217	Fact Checking Beyond Training Set	Main Conference	Virtual Poster
218	Program-Aided Reasoners (Better) Know What They Know	Main Conference	Poster
219	The Perspectivist Paradigm Shift: Assumptions and Challenges of Capturing Human Labels	Main Conference	Poster
220	A Multi-Aspect Framework for Counter Narrative Evaluation using Large Language Models	Main Conference	Poster
221	Principles from Clinical Research for NLP Model Generalization	Main Conference	Oral
222	Efficient Dependency Tree Sampling Without Replacement	Findings	Virtual Poster
223	First Tragedy, then Parse: History Repeats Itself in the New Era of Large Language Models	Main Conference	Poster
224	Found in the Middle: Permutation Self-Consistency Improves Listwise Ranking in Large Language Models	Main Conference	Virtual Poster
225	From Language Modeling to Instruction Following: Understanding the Behavior Shift in LLMs after Instruction Tuning	Main Conference	Oral
226	Towards Better Generalization in Open-Domain Question Answering by Mitigating Context Memorization	Findings	Poster
229	GEE! Grammar Error Explanation with Large Language Models	Findings	Poster
230	POLYIE: A Dataset of Information Extraction from Polymer Material Scientific Literature	Main Conference	Virtual Poster
231	LLM-based Medical Assistant Personalization with Short- and Long-Term Memory Coordination	Main Conference	Oral
232	SumTra: A Differentiable Pipeline for Few-Shot Cross-Lingual Summarization	Main Conference	Poster
233	How does Multi-Task Training Affect Transformer In-Context Capabilities? Investigations with Function Classes	Main Conference	Virtual Poster
234	KTRL+F: Knowledge-Augmented In-Document Search	Main Conference	Poster
235	How Well Do Large Language Models Truly Ground?	Main Conference	Oral
236	CELL: Simple yet Effective Approach to Enhance Out-of-Domain Generalization of Cross-Encoders.	Main Conference	Oral
238	ALBA: Adaptive Language-Based Assessments for Mental Health	Main Conference	Poster
239	FREB-TQA: A Fine-Grained Robustness Evaluation Benchmark for Table Question Answering	Main Conference	Poster
240	MILL: Mutual Verification with Large Language Models for Zero-Shot Query Expansion	Main Conference	Poster
242	Efficient Benchmarking (of Language Models)	Main Conference	Poster

PaperID	Title	Proceedings	Format
243	ReFACT: Updating Text-to-Image Models by Editing the Text Encoder	Main Conference	Poster
244	ContrastiveMix: Overcoming Code-Mixing Dilemma in Cross-Lingual Transfer for Information Retrieval	Main Conference	Oral
245	A Likelihood Ratio Test of Genetic Relationship among Languages	Main Conference	Poster
246	AdaRefiner: Refining Decisions of Language Models with Adaptive Feedback	Findings	Virtual Poster
248	PaD: Program-aided Distillation Can Teach Small Models Reasoning Better than Chain-of-thought Fine-tuning	Main Conference	Oral
249	MEGAVERSE: Benchmarking Large Language Models Across Languages, Modalities, Models and Tasks	Main Conference	Poster
250	Unlocking Emergent Modularity in Large Language Models	Main Conference	Virtual Poster
251	A School Student Essay Corpus for Analyzing Interactions of Argumentative Structure and Quality	Main Conference	Poster
253	Adjusting Interpretable Dimensions in Embedding Space with Human Judgments	Main Conference	Poster
254	PatentEval: Understanding Errors in Patent Generation	Main Conference	Poster
255	Contextual Refinement of Translations: Large Language Models for Sentence and Document-Level Post-Editing	Main Conference	Poster
256	DivTOD: Unleashing the Power of LLMs for Diversifying Task-Oriented Dialogue Representations	Findings	Virtual Poster
257	Metaphor Detection with Context Enhancement and Curriculum Learning	Main Conference	Virtual Poster
258	Teaching a Multilingual Large Language Model to Understand Multilingual Speech via Multi-Instructional Training	Findings	Virtual Poster
259	What Causes the Failure of Explicit to Implicit Discourse Relation Recognition?	Main Conference	Oral
261	CLEAN-EVAL: Clean Evaluation on Contaminated Large Language Models	Findings	Virtual Poster
262	UniverSLU: Universal Spoken Language Understanding for Diverse Tasks with Natural Language Instructions	Main Conference	Oral
263	How Trustworthy are Open-Source LLMs? An Assessment under Malicious Demonstrations Shows their Vulnerabilities	Main Conference	Poster
264	R-BASS : Relevance-aided Block-wise Adaptation for Speech Summarization	Findings	Poster
265	OVM, Outcome-supervised Value Models for Planning in Mathematical Reasoning	Findings	Poster
267	Paraphrase and Solve: Exploring and Exploiting the Impact of Surface Form on Mathematical Reasoning in Large Language Models	Main Conference	Virtual Poster
268	The Whole is Better than the Sum: Using Aggregated Demonstrations in In-Context Learning for Sequential Recommendation	Findings	Virtual Poster
269	TriSum: Learning Summarization Ability from Large Language Models with Structured Rationale	Main Conference	Poster
270	GenRES: Rethinking Evaluation for Generative Relation Extraction in the Era of Large Language Models	Main Conference	Poster
271	Curated Datasets and Neural Models for Machine Translation of Informal Registers between Mayan and Spanish Vernaculars	Main Conference	Virtual Poster
272	Bring Your Own KG: Self-Supervised Program Synthesis for Zero-Shot KGQA	Findings	Poster
273	The Effect of Data Partitioning Strategy on Model Generalizability: A Case Study of Morphological Segmentation	Main Conference	Poster
274	Measuring Entrainment in Spontaneous Code-switched Speech	Main Conference	Poster
275	A Survey of Meaning Representations – From Theory to Practical Utility	Main Conference	Poster
276	Mitigating Language-Level Performance Disparity in mPLMs via Teacher Language Selection and Cross-lingual Self-Distillation	Main Conference	Poster
277	GraSAME: Injecting Token-Level Structural Information to Pretrained Language Models via Graph-guided Self-Attention Mechanism	Findings	Poster
278	Evaluating In-Context Learning of Libraries for Code Generation	Main Conference	Poster
279	Visually-Aware Context Modeling for News Image Captioning	Main Conference	Poster
280	Regularized Conventions: Equilibrium Computation as a Model of Pragmatic Reasoning	Main Conference	Poster
281	SLIDE: Reference-free Evaluation for Machine Translation using a Sliding Document Window	Main Conference	Poster
282	TopicGPT: A Prompt-based Topic Modeling Framework	Main Conference	Poster
283	ChatGPT as an Attack Tool: Stealthy Textual Backdoor Attack via Blackbox Generative Model Trigger	Main Conference	Poster
284	Social Meme-ing: Measuring Linguistic Variation in Memes	Main Conference	Poster
285	ExpertQA: Expert-Curated Questions and Attributed Answers	Main Conference	Oral
286	What if you said that differently?: How Explanation Formats Affect Human Feedback Efficacy and User Perception	Main Conference	Poster
287	Can Public Large Language Models Help Private Cross-device Federated Learning?	Findings	Virtual Poster

PaperID	Title	Proceedings	Format
288	When Life Gives You Lemons, Make Cherryade: Converting Feedback from Bad Responses into Good Labels	Main Conference	Poster
289	Kreyòl-MT: Building MT for Latin American, Caribbean and Colonial African Creole Languages	Main Conference	Oral
290	LangNav: Language as a Perceptual Representation for Navigation	Findings	Poster
291	Instructions as Backdoors: Backdoor Vulnerabilities of Instruction Tuning for Large Language Models	Main Conference	Poster
292	Modeling Empathetic Alignment in Conversation	Main Conference	Oral
293	Separately Parameterizing Singleton Detection Improves End-to-end Neural Coreference Resolution	Main Conference	Oral
294	Native Language Identification in Texts: A Survey	Main Conference	Virtual Poster
295	LoRETTA: Low-Rank Economic Tensor-Train Adaptation for Ultra-Low-Parameter Fine-Tuning of Large Language Models	Main Conference	Oral
296	Which One? Leveraging Context Between Objects and Multiple Views for Language Grounding	Main Conference	Poster
297	Do Localization Methods Actually Localize Memorized Data in LLMs? A Tale of Two Benchmarks	Main Conference	Oral
298	PromptFix: Few-shot Backdoor Removal via Adversarial Prompt Tuning	Main Conference	Oral
299	Comparing Explanation Faithfulness between Multilingual and Monolingual Fine-tuned Language Models	Main Conference	Oral
300	Planning and Editing What You Retrieve for Enhanced Tool Learning	Findings	Poster
301	A Pretrainer's Guide to Training Data: Measuring the Effects of Data Age, Domain Coverage, Quality, & Toxicity	Main Conference	Poster
302	Chart-based Reasoning: Transferring Capabilities from LLMs to VLMs	Findings	Poster
303	Unveiling Divergent Inductive Biases of LLMs on Temporal Data	Main Conference	Poster
305	Instructional Fingerprinting of Large Language Models	Main Conference	Oral
306	Reinforced Multiple Instance Selection for Speaker Attribute Prediction	Main Conference	Virtual Poster
307	SLiM: Speculative Decoding with Hypothesis Reduction	Findings	Virtual Poster
308	DynaMo: Accelerating Language Model Inference with Dynamic Multi-Token Sampling	Main Conference	Poster
309	REMATCH: Robust and Efficient Matching of Local Knowledge Graphs to Improve Structural and Semantic Similarity	Findings	Poster
310	On Retrieval Augmentation and the Limitations of Language Model Training	Main Conference	Poster
311	Few-shot Knowledge Graph Relational Reasoning via Subgraph Adaptation	Main Conference	Virtual Poster
313	Modeling the Sacred: Considerations when Using Religious Texts in Natural Language Processing	Findings	Poster
314	Uncertainty Quantification for In-Context Learning of Large Language Models	Main Conference	Poster
316	HelpSteer: Multi-attribute Helpfulness Dataset for SteerLM	Main Conference	Poster
318	A Preference-driven Paradigm for Enhanced Translation with Large Language Models	Main Conference	Oral
319	GenDecider: Integrating "None of the Candidates" Judgments in Zero-Shot Entity Linking Re-ranking	Main Conference	Virtual Poster
320	Fair Abstractive Summarization of Diverse Perspectives	Main Conference	Virtual Poster
321	What Are We Measuring When We Evaluate Large Vision-Language Models? An Analysis of Latent Factors and Biases	Main Conference	Virtual Poster
322	Show Your Work with Confidence: Confidence Bands for Tuning Curves	Main Conference	Poster
325	GRASP: A Disagreement Analysis Framework to Assess Group Associations in Perspectives	Main Conference	Poster
326	Event Causality Is Key to Computational Story Understanding	Main Conference	Oral
327	Testing the Effect of Code Documentation on Large Language Model Code Understanding	Findings	Poster
328	Subspace Representations for Soft Set Operations and Sentence Similarities	Main Conference	Virtual Poster
329	Aligning Large Language Models with Recommendation Knowledge	Findings	Poster
330	OFA: A Framework of Initializing Unseen Subword Embeddings for Efficient Large-scale Multilingual Continued Pretraining	Findings	Poster
331	Advancing the Robustness of Large Language Models through Self-Denoised Smoothing	Main Conference	Virtual Poster
334	SELF-EXPERTISE: Knowledge-based Instruction Dataset Augmentation for a Legal Expert Language Model	Findings	Poster
335	Can LLM's Generate Human-Like Wayfinding Instructions? Towards Platform-Agnostic Embodied Instruction Synthesis	Main Conference	Poster
336	Re-evaluating the Need for Visual Signals in Unsupervised Grammar Induction	Findings	Poster

PaperID	Title	Proceedings	Format
337	My Heart Skipped a Beat! Recognizing Expressions of Embodied Emotion in Natural Language	Main Conference	Oral
338	Low-Cost Generation and Evaluation of Dictionary Example Sentences	Main Conference	Virtual Poster
339	Making Language Models Better Tool Learners with Execution Feedback	Main Conference	Virtual Poster
340	Complex Claim Verification with Evidence Retrieved in the Wild	Main Conference	Oral
341	Multimodal Multi-loss Fusion Network for Sentiment Analysis	Main Conference	Oral
342	Confronting LLMs with Traditional ML: Rethinking the Fairness of Large Language Models in Tabular Classifications	Main Conference	Oral
344	EDEntail: An Entailment-based Few-shot Text Classification with Extensional Definition	Findings	Virtual Poster
346	What Makes Math Word Problems Challenging for LLMs?	Findings	Poster
347	Analyzing the Use of Metaphors in News Editorials for Political Framing	Main Conference	Oral
348	SharpSeq: Empowering Continual Event Detection through Sharpness-Aware Sequential-task Learning	Main Conference	Virtual Poster
350	Dissecting Paraphrases: The Impact of Prompt Syntax and supplementary Information on Knowledge Retrieval from Pretrained Language Models	Main Conference	Poster
351	SMILE: Multimodal Dataset for Understanding Laughter in Video with Language Models	Findings	Virtual Poster
352	Know When To Stop: A Study of Semantic Drift in Text Generation	Main Conference	Poster
353	T3M: Text Guided 3D Human Motion Synthesis from Speech	Findings	Virtual Poster
355	Curriculum Masking in Vision-Language Pretraining to Maximize Cross Modal Interaction	Main Conference	Poster
356	Elote, Choclo and Mazorca: on the Varieties of Spanish	Main Conference	Poster
357	On the Role of Summary Content Units in Text Summarization Evaluation	Main Conference	Poster
359	Ada-LEval: Evaluating long-context LLMs with length-adaptable benchmarks	Main Conference	Poster
361	More room for language: Investigating the effect of retrieval on language models	Main Conference	Poster
362	Deja vu: Contrastive Historical Modeling with Prefix-tuning for Temporal Knowledge Graph Reasoning	Findings	Poster
363	A Zero-Shot Monolingual Dual Stage Information Retrieval System for Spanish Biomedical Systematic Literature Reviews	Main Conference	Poster
364	LayoutPointer: A Spatial-Context Adaptive Pointer Network for Visual Information Extraction	Main Conference	Virtual Poster
365	Long-form evaluation of model editing	Main Conference	Oral
366	Explanation Extraction from Hierarchical Classification Frameworks for Long Legal Documents	Findings	Poster
367	Analyzing the Role of Semantic Representations in the Era of Large Language Models	Main Conference	Poster
368	Low-Rank Adaptation for Multilingual Summarization: An Empirical Study	Findings	Poster
371	Empowering Multi-step Reasoning across Languages via Tree-of-Thoughts	Findings	Virtual Poster
372	TRAQ: Trustworthy Retrieval Augmented Question Answering via Conformal Prediction	Main Conference	Poster
375	Emergent Abilities in Reduced-Scale Generative Language Models	Findings	Virtual Poster
376	MapGuide: A Simple yet Effective Method to Reconstruct Continuous Language from Brain Activities	Main Conference	Virtual Poster
377	Discourse-Aware In-Context Learning for Temporal Expression Normalization	Main Conference	Virtual Poster
378	On-the-fly Definition Augmentation of LLMs for Biomedical NER	Main Conference	Poster
379	Context Does Matter: Implications for Crowdsourced Evaluation Labels in Task-Oriented Dialogue Systems	Findings	Poster
380	Matching Varying-Length Texts via Topic-Informed and Decoupled Sentence Embeddings	Findings	Virtual Poster
382	This Land is {Your, My} Land: Evaluating Geopolitical Bias in Language Models through Territorial Disputes	Main Conference	Oral
384	Instruction Tuning with Human Curriculum	Findings	Virtual Poster
385	Contextualizing Argument Quality Assessment with Relevant Knowledge	Main Conference	Virtual Poster
387	Set-Aligning Framework for Auto-Regressive Event Temporal Graph Generation	Main Conference	Poster
388	LanguageFlow: Advancing Diffusion Language Generation with Probabilistic Flows	Main Conference	Virtual Poster
389	Natural Language-based State Representation in Deep Reinforcement Learning	Findings	Poster
390	Selective Perception: Learning Concise State Descriptions for Language Model Actors	Main Conference	Poster

PaperID	Title	Proceedings	Format
391	Towards Improved Multi-Source Attribution for Long-Form Answer Generation	Main Conference	Poster
392	Learning Cross-Architecture Instruction Embeddings for Binary Code Analysis in Low-Resource Architectures	Findings	Virtual Poster
393	ALOHa: A New Measure for Hallucination in Captioning Models	Main Conference	Oral
394	Synthetic Query Generation for Privacy-Preserving Deep Retrieval Systems using Differentially Private Language Models	Main Conference	Poster
395	Okay, Let's Do This! Modeling Event Coreference with Generated Rationales and Knowledge Distillation	Main Conference	Oral
396	Can Knowledge Graphs Reduce Hallucinations in LLMs? : A Survey	Main Conference	Oral
397	Pedagogically aligned objectives create reliable automatic cloze tests	Main Conference	Poster
398	ReEval: Automatic Hallucination Evaluation for Retrieval-Augmented Large Language Models via Transferable Adversarial Attacks	Findings	Poster
399	Beyond Yes and No: Improving Zero-Shot LLM Rankers via Scoring Fine-Grained Relevance Labels	Main Conference	Poster
401	Take One Step at a Time to Know Incremental Utility of Demonstration: An Analysis on Reranking for Few-Shot In-Context Learning	Main Conference	Virtual Poster
402	LM-Infinite: Zero-Shot Extreme Length Generalization for Large Language Models	Main Conference	Poster
403	CONSCENDI: A Contrastive and Scenario-Guided Distillation Approach to Guardrail Models for Virtual Assistants	Main Conference	Poster
404	An Effective Automated Speaking Assessment Approach to Mitigating Data Scarcity and Imbalanced Distribution	Findings	Poster
405	LLM-Driven Knowledge Injection Advances Zero-Shot and Cross-Target Stance Detection	Main Conference	Virtual Poster
406	GPT-Fathom: Benchmarking Large Language Models to Decipher the Evolutionary Path towards GPT-4 and Beyond	Findings	Poster
407	Advancing Beyond Identification: Multi-bit Watermark for Large Language Models	Main Conference	Oral
408	HTCCN: Temporal Causal Convolutional Networks with Hawkes Process for Extrapolation Reasoning in Temporal Knowledge Graphs	Main Conference	Poster
409	SemStamp: A Semantic Watermark with Paraphrastic Robustness for Text Generation	Main Conference	Poster
410	Media Bias Detection Across Families of Language Models	Main Conference	Poster
411	Better Zero-Shot Reasoning with Role-Play Prompting	Main Conference	Poster
412	Subword Attention and Post-Processing for Rare and Unknown Contextualized Embeddings	Findings	Virtual Poster
413	Event-Content-Oriented Dialogue Generation in Short Video	Main Conference	Poster
414	UGIF-DataSet: A New Dataset for Cross-lingual, Cross-modal Sequential actions on the UI	Findings	Virtual Poster
416	SimSCOOD: Systematic Analysis of Out-of-Distribution Generalization in Fine-tuned Source Code Models	Findings	Poster
417	Pruning as a Domain-specific LLM Extractor	Findings	Virtual Poster
419	DoG-Instruct: Towards Premium Instruction-Tuning Data via Text-Grounded Instruction Wrapping	Main Conference	Virtual Poster
420	Leveraging Prototypical Representations for Mitigating Social Bias without Demographic Information	Main Conference	Poster
421	Beyond Borders: Investigating Cross-Jurisdiction Transfer in Legal Case Summarization	Main Conference	Poster
423	EDC: Effective and Efficient Dialog Comprehension For Dialog State Tracking	Main Conference	Virtual Poster
424	Automatic Restoration of Diacritics for Speech Data Sets	Main Conference	Poster
425	XNLeu: a dataset for cross-lingual NLI in Basque	Main Conference	Poster
426	MDR: Model-Specific Demonstration Retrieval at Inference Time for In-Context Learning	Main Conference	Poster
427	LLMRefine: Pinpointing and Refining Large Language Models via Fine-Grained Actionable Feedback	Findings	Poster
428	Exploring Cross-Cultural Differences in English Hate Speech Annotations: From Dataset Construction to Analysis	Main Conference	Oral
429	Enhancing Contextual Understanding in Large Language Models through Contrastive Decoding	Main Conference	Oral
430	Noisy Multi-Label Text Classification via Instance-Label Pair Correction	Findings	Virtual Poster
431	Generalizable Sarcasm Detection is Just Around the Corner, of Course!	Main Conference	Oral
432	Composite Backdoor Attacks Against Large Language Models	Findings	Virtual Poster
433	Adapting Fake News Detection to the Era of Large Language Models	Findings	Poster
434	Encoding of lexical tone in self-supervised models of spoken language	Main Conference	Oral
435	A Systematic Comparison of Contextualized Word Embeddings for Lexical Semantic Change	Main Conference	Poster

PaperID	Title	Proceedings	Format
436	iACOS: Advancing Implicit Sentiment Extraction with Informative and Adaptive Negative Examples	Main Conference	Oral
437	Rectifying Demonstration Shortcut in In-Context Learning	Main Conference	Poster
438	MCAD: Multi-teacher Cross-modal Alignment Distillation for efficient image-text retrieval	Findings	Virtual Poster
439	Large Language Models are Effective Text Rankers with Pairwise Ranking Prompting	Findings	Virtual Poster
441	FedLFC: Towards Efficient Federated Multilingual Modeling with LoRA-based Language Family Clustering	Findings	Poster
442	Universal NER: A Gold-Standard Multilingual Named Entity Recognition Benchmark	Main Conference	Poster
443	Direct Preference Optimization for Neural Machine Translation with Minimum Bayes Risk Decoding	Main Conference	Poster
444	Gaussian Process Optimization for Adaptable Multi-Objective Text Generation using Linearly-Weighted Language Models	Findings	Virtual Poster
446	ODD: A Benchmark Dataset for the Natural Language Processing Based Opioid Related Aberrant Behavior Detection	Main Conference	Poster
447	A Comprehensive Study of Gender Bias in Chemical Named Entity Recognition Models	Main Conference	Poster
449	Groundedness in Retrieval-augmented Long-form Generation: An Empirical Study	Findings	Poster
450	The Promises and Pitfalls of Using Language Models to Measure Instruction Quality in Education	Main Conference	Poster
451	TagDebias: Entity and Concept Tagging for Social Bias Mitigation in Pretrained Language Models	Findings	Poster
452	EchoPrompt: Instructing the Model to Rephrase Queries for Improved In-context Learning	Main Conference	Poster
453	Improving Absent Keyphrase Generation with Diversity Heads	Findings	Poster
454	Differentially Private Next-Token Prediction of Large Language Models	Main Conference	Poster
455	Improving Adversarial Data Collection by Supporting Annotators: Lessons from GAHD, a German Hate Speech Dataset	Main Conference	Poster
456	mOthello: When Do Cross-Lingual Representation Alignment and Cross-Lingual Transfer Emerge in Multilingual Models?	Findings	Poster
457	Memory Augmented Language Models through Mixture of Word Experts	Main Conference	Poster
458	LEAF: Language Learners' English Essays and Feedback Corpus	Main Conference	Virtual Poster
459	Impossible Distillation for Paraphrasing and Summarization: How to Make High-quality Lemonade out of Small, Low-quality Model	Main Conference	Poster
460	TofuEval: Evaluating Hallucinations of LLMs on Topic-Focused Dialogue Summarization	Main Conference	Virtual Poster
461	MOKA: Moral Knowledge Augmentation for Moral Event Extraction	Main Conference	Poster
462	Discovering and Mitigating Indirect Bias in Attention-Based Model Explanations	Findings	Poster
463	Fixing Rogue Memorization in Many-to-One Multilingual Translators of Extremely-Low-Resource Languages by Rephrasing Training Samples	Main Conference	Poster
464	Zero-Shot vs. Translation-Based Cross-Lingual Transfer: The Case of Lexical Gaps	Main Conference	Oral
465	i-Code V2: An Autoregressive Generation Framework over Vision, Language, and Speech Data	Findings	Virtual Poster
466	Think While You Write: Hypothesis Verification Promotes Faithful Knowledge-to-Text Generation	Findings	Poster
467	It's All Relative! -- A Synthetic Query Generation Approach for Improving Zero-Shot Relevance Prediction	Findings	Poster
468	RS-DPO: A Hybrid Rejection Sampling and Direct Preference Optimization Method for Alignment of Large Language Models	Findings	Poster
470	Backdoor Attacks on Multilingual Machine Translation	Main Conference	Oral
471	Personalized Jargon Identification for Enhanced Interdisciplinary Communication	Main Conference	Virtual Poster
472	Hypernetwork-Assisted Parameter-Efficient Fine-Tuning with Meta-Knowledge Distillation for Domain Knowledge Disentanglement	Findings	Virtual Poster
473	Flames: Benchmarking Value Alignment of LLMs in Chinese	Main Conference	Poster
474	MICo: Preventative Detoxification of Large Language Models through Inhibition Control	Findings	Poster
475	Reinforcement Learning with Token-level Feedback for Controllable Text Generation	Findings	Virtual Poster
476	CoMM: Collaborative Multi-Agent, Multi-Reasoning-Path Prompting for Complex Problem Solving	Findings	Virtual Poster
477	Tokenization Matters: Navigating Data-Scarce Tokenization for Gender Inclusive Language Technologies	Findings	Poster
478	AdaPT: A Set of Guidelines for Hyperbolic Multimodal Multilingual NLP	Findings	Poster
479	More Samples or More Prompts? Exploring Effective Few-Shot In-Context Learning for LLMs with In-Context Sampling	Findings	Poster
480	Mitigating Bias for Question Answering Models by Tracking Bias Influence	Main Conference	Poster

PaperID	Title	Proceedings	Format
481	Extending CLIP's Image-Text Alignment to Referring Image Segmentation	Main Conference	Poster
482	Generating Attractive and Authentic Copywriting from Customer Reviews	Main Conference	Poster
483	Effective Long-Context Scaling of Foundation Models	Main Conference	Virtual Poster
485	Empowering Diffusion Models on the Embedding Space for Text Generation	Main Conference	Virtual Poster
486	On the True Distribution Approximation of Minimum Bayes-Risk Decoding	Main Conference	Poster
487	ZSEE: A Dataset based on Zeolite Synthesis Event Extraction for Automated Synthesis Platform	Findings	Virtual Poster
488	Mitigating Hallucination in Abstractive Summarization with Domain-Conditional Mutual Information	Findings	Poster
489	Adversarial DPO: Harnessing Harmful Data for Reducing Toxicity with Minimal Impact on Coherence and Evasiveness in Dialogue Agents	Findings	Poster
490	Rehearsal-Free Modular and Compositional Continual Learning for Language Models	Main Conference	Poster
491	Aligning as Debiasing: Causality-Aware Alignment via Reinforcement Learning with Interventional Feedback	Main Conference	Virtual Poster
492	Prompt Space Optimizing Few-shot Reasoning Success with Large Language Models	Findings	Virtual Poster
493	Fake Alignment: Are LLMs Really Aligned Well?	Main Conference	Oral
494	Visually Guided Generative Text-Layout Pre-training for Document Intelligence	Main Conference	Poster
495	DAGCN: Distance-based and Aspect-oriented Graph Convolutional Network for Aspect-based Sentiment Analysis	Findings	Virtual Poster
496	Llama meets EU: Investigating the European political spectrum through the lens of LLMs	Main Conference	Poster
498	HILL: Hierarchy-aware Information Lossless Contrastive Learning for Hierarchical Text Classification	Main Conference	Virtual Poster
499	Connecting the Dots: Inferring Patent Phrase Similarity with Retrieved Phrase Graphs	Findings	Poster
500	Self-Regulated Sample Diversity in Large Language Models	Findings	Poster
502	Investigating the Emergent Audio Classification Ability of ASR Foundation Models	Main Conference	Poster
503	Methods, Applications, and Directions of Learning-to-Rank in NLP Research	Findings	Poster
505	When Quantization Affects Confidence of Large Language Models?	Findings	Poster
509	MedCycle: Unpaired Medical Report Generation via Cycle-Consistency	Findings	Poster
510	In-context Learning Generalizes, But Not Always Robustly: The Case of Syntax	Main Conference	Poster
511	Prompt-Singer: Controllable Singing-Voice-Synthesis with Natural Language Prompt	Main Conference	Virtual Poster
512	Beta-LR: Interpretable Logical Reasoning based on Beta Distribution	Findings	Virtual Poster
513	Lost in Transcription: Identifying and Quantifying the Accuracy Biases of Automatic Speech Recognition Systems Against Disfluent Speech	Main Conference	Oral
514	MAFALDA: A Benchmark and Comprehensive Study of Fallacy Detection and Classification	Main Conference	Oral
517	Diffusion Glancing Transformer for Parallel Sequence-to-Sequence Learning	Main Conference	Virtual Poster
519	Applications of BERT Models Towards Automation of Clinical Coding in Icelandic	Findings	Poster
520	Tell me who you are and I tell you how you argue: Predicting Stances and Arguments for Stakeholder Groups	Findings	Virtual Poster
523	Psychometric Predictive Power of Large Language Models	Findings	Poster
525	No Context Needed: Contextual Quandary In Idiomatic Reasoning With Pre-Trained Language Models	Main Conference	Virtual Poster
526	M3T: A New Benchmark Dataset for Multi-Modal Document-Level Machine Translation	Main Conference	Poster
527	Multi-stage Retrieve and Re-rank Model for Automatic Medical Coding Recommendation	Main Conference	Oral
528	Anisotropy is Not Inherent to Transformers	Main Conference	Virtual Poster
531	Finding Replicable Human Evaluations via Stable Ranking Probability	Main Conference	Poster
532	Large Language Models Sensitivity to The Order of Options in Multiple-Choice Questions	Findings	Poster
533	Stealthy and Persistent Unalignment on Large Language Models via Backdoor Injections	Main Conference	Poster
535	PEEB: Part-based Image Classifiers with an Explainable and Editable Language Bottleneck	Findings	Poster
537	Generalizable and Stable Finetuning of Pretrained Language Models on Low-Resource Texts	Main Conference	Poster
538	Detecting Bipolar Disorder from Misdiagnosed Major Depressive Disorder with Mood-Aware Multi-Task Learning	Main Conference	Poster

PaperID	Title	Proceedings	Format
539	Leveraging Code to Improve In-Context Learning for Semantic Parsing	Main Conference	Oral
540	Improving Pre-trained Language Model Sensitivity via Mask Specific losses: A case study on Biomedical NER	Main Conference	Poster
541	Language Models Implement Simple Word2Vec-style Vector Arithmetic	Main Conference	Poster
542	Control-DAG: Constrained Decoding for Non-Autoregressive Directed Acyclic T5 using Weighted Finite State Automata	Main Conference	Poster
543	AutoLoRA: Automatically Tuning Matrix Ranks in Low-Rank Adaptation Based on Meta Learning	Main Conference	Poster
544	Do Vision-Language Models Understand Compound Nouns?	Main Conference	Poster
546	SportQA: A Benchmark for Sports Understanding in Large Language Models	Main Conference	Oral
547	Revisiting subword tokenization: A case study on affixal negation in large language models	Main Conference	Poster
548	Generating Mental Health Transcripts with SAPE (Spanish Adaptive Prompt Engineering)	Main Conference	Oral
549	Ethos: Rectifying Language Models in Orthogonal Parameter Space	Findings	Poster
550	Crafting In-context Examples according to LMs' Parametric Knowledge	Findings	Poster
552	Where are you from? Geolocating Speech and Applications to Language Identification	Main Conference	Poster
553	Teaching Language Models to Self-Improve through Interactive Demonstrations	Main Conference	Oral
554	MAGID: An Automated Pipeline for Generating Synthetic Multi-modal Datasets	Main Conference	Oral
555	ICXML: An In-Context Learning Framework for Zero-Shot Extreme Multi-Label Classification	Findings	Virtual Poster
556	Zero-shot Generative Linguistic Steganography	Main Conference	Virtual Poster
557	Does GPT-4 pass the Turing test?	Main Conference	Oral
558	CLGSI: A Multimodal Sentiment Analysis Framework based on Contrastive Learning Guided by Sentiment Intensity	Findings	Virtual Poster
559	Polarity Calibration for Opinion Summarization	Main Conference	Virtual Poster
560	Sentence-level Media Bias Analysis with Event Relation Graph	Main Conference	Virtual Poster
561	Interpreting Answers to Yes-No Questions in Dialogues from Multiple Domains	Findings	Poster
562	EMONA: Event-level Moral Opinions in News Articles	Main Conference	Virtual Poster
563	DLM: A Decoupled Learning Model for Long-tailed Polyphone Disambiguation in Mandarin	Main Conference	Oral
565	Enhancing Perception: Refining Explanations of News Claims with LLM Conversations	Findings	Virtual Poster
566	How Interpretable are Reasoning Explanations from Prompting Large Language Models?	Findings	Virtual Poster
567	You don't need a personality test to know these models are unreliable: Assessing the Reliability of Large Language Models on Psychometric Instruments	Main Conference	Oral
568	CASA: Causality-driven Argument Sufficiency Assessment	Main Conference	Poster
569	Plug-in Language Model: Controlling Text Generation with a Simple Regression Model	Findings	Poster
571	MacGyver: Are Large Language Models Creative Problem Solvers?	Main Conference	Poster
572	Signer Diversity-driven Data Augmentation for Signer-Independent Sign Language Translation	Findings	Virtual Poster
573	A Systematic Analysis of Subwords and Cross-Lingual Transfer in Multilingual Translation	Findings	Poster
574	Multi-Granularity Guided Fusion-in-Decoder	Findings	Poster
575	Group Fairness in Multilingual Speech Recognition Models	Findings	Poster
576	Rethinking Machine Ethics -- Can LLMs Perform Moral Reasoning through the Lens of Moral Theories?	Findings	Poster
577	To Translate or Not to Translate: A Systematic Investigation of Translation-Based Cross-Lingual Transfer to Low-Resource Languages	Main Conference	Poster
578	Enhancing Large Language Models Against Inductive Instructions with Dual-critique Prompting	Main Conference	Virtual Poster
579	Role Prompting Guided Domain Adaptation with General Capability Preserve for Large Language Models	Findings	Virtual Poster
580	GLiNER: Generalist Model for Named Entity Recognition using Bidirectional Transformer	Main Conference	Oral
581	XSTest: A Test Suite for Identifying Exaggerated Safety Behaviours in Large Language Models	Main Conference	Poster
582	BERTweet's TACO Fiesta: Contrasting Flavors On The Path Of Inference And Information-Driven Argument Mining On Twitter	Findings	Poster
583	Testing the limits of logical reasoning in neural and hybrid models	Findings	Poster

PaperID	Title	Proceedings	Format
584	Is Prompt Transfer Always Effective? An Empirical Study of Prompt Transfer for Question Answering	Main Conference	Poster
585	METAL: Towards Multilingual Meta-Evaluation	Findings	Poster
586	AGIEval: A Human-Centric Benchmark for Evaluating Foundation Models	Findings	Poster
587	Carpe diem: On the Evaluation of World Knowledge in Lifelong Language Models	Main Conference	Poster
588	Lost in Space: Probing Fine-grained Spatial Understanding in Vision and Language Resamplers	Main Conference	Oral
589	Product Description and QA Assisted Self-Supervised Opinion Summarization	Findings	Virtual Poster
590	COMEM: In-Context Retrieval-Augmented Mass-Editing Memory in Large Language Models	Findings	Poster
592	Fine-grained Gender Control in Machine Translation with Large Language Models	Main Conference	Poster
593	DialogVCS: Robust Natural Language Understanding in Dialogue System Upgrade	Main Conference	Poster
595	LLattribution: LLM-Verified Retrieval for Verifiable Generation	Main Conference	Virtual Poster
598	Content-Specific Humorous Image Captioning Using Incongruity Resolution Chain-of-Thought	Findings	Poster
599	Mapping Long-term Causalities in Psychiatric Symptomatology and Life Events from Social Media	Main Conference	Poster
601	Do Multilingual Language Models Think Better in English?	Main Conference	Poster
602	Multimodal Chart Retrieval: A Comparison of Text, Table and Image Based Approaches	Main Conference	Poster
603	Denoising Attention for Query-aware User Modeling	Findings	Virtual Poster
604	A Lightweight Mixture-of-Experts Neural Machine Translation Model With Stage-wise Training Strategy	Findings	Virtual Poster
605	BEAR: A Unified Framework for Evaluating Relational Knowledge in Causal and Masked Language Models	Findings	Poster
606	Conformal Intent Classification and Clarification for Fast and Accurate Intent Recognition	Findings	Poster
609	Anonymity at Risk? Assessing Re-Identification Capabilities of Large Language Models in Court Decisions	Findings	Poster
610	X-LLaVA: Optimizing Bilingual Large Vision-Language Alignment	Findings	Poster
611	Why So Gullible? Enhancing the Robustness of Retrieval-Augmented Models against Counterfactual Noise	Findings	Virtual Poster
612	Heterogeneity over Homogeneity: Investigating Multilingual Speech Pre-Trained Models for Detecting Audio Deepfake	Findings	Virtual Poster
613	Identifying Self-Disclosures of Use, Misuse and Addiction in Community-based Social Media Posts	Findings	Poster
614	Self-Adaptive Sampling for Accurate Video Question Answering on Image Text Models	Findings	Virtual Poster
616	Retrieval Helps or Hurts? A Deeper Dive into the Efficacy of Retrieval Augmentation to Language Models	Main Conference	Oral
617	AudioChatLlama: Towards General-Purpose Speech Abilities for LLMs	Main Conference	Virtual Poster
619	Towards an On-device Agent for Text Rewriting	Findings	Poster
620	Tailoring Vaccine Messaging with Common-Ground Opinions	Findings	Poster
621	Whispers of Doubt Amidst Echoes of Triumph in NLP Robustness	Main Conference	Poster
625	Sequential Compositional Generalization in Multimodal Models	Main Conference	Oral
626	Best of Both Worlds: A Pliable and Generalizable Neuro-Symbolic Approach for Relation Classification	Findings	Poster
628	Q-Tuning: Queue-based Prompt Tuning for Lifelong Few-shot Language Learning	Findings	Poster
630	Generating Uncontextualized and Contextualized Questions for Document-Level Event Argument Extraction	Main Conference	Poster
632	Evidence-Driven Retrieval Augmented Response Generation for Online Misinformation	Main Conference	Virtual Poster
633	A Continued Pretrained LLM Approach for Automatic Medical Note Generation	Main Conference	Poster
634	In-Context Example Ordering Guided by Label Distributions	Findings	Poster
635	Open-Vocabulary Federated Learning with Multimodal Prototyping	Main Conference	Virtual Poster
636	Beyond Surface Similarity: Detecting Subtle Semantic Shifts in Financial Narratives	Findings	Virtual Poster
638	Exploring Key Point Analysis with Pairwise Generation and Graph Partitioning	Main Conference	Virtual Poster
640	Laying Anchors: Semantically Priming Numerals in Language Modeling	Findings	Poster
641	UEGP: Unified Expert-Guided Pre-training for Knowledge Rekindle	Findings	Virtual Poster

PaperID	Title	Proceedings	Format
642	Understanding the Capabilities and Limitations of Large Language Models for Cultural Commonsense	Main Conference	Poster
643	LatticeGen: Hiding Generated Text in a Lattice for Privacy-Aware Large Language Model Generation on Cloud	Findings	Virtual Poster
644	HateModerate: Testing Hate Speech Detectors against Content Moderation Policies	Findings	Virtual Poster
645	Code Models are Zero-shot Precondition Reasoners	Main Conference	Oral
646	Contrastive and Consistency Learning for Neural Noisy-Channel Model in Spoken Language Understanding	Main Conference	Oral
648	Do Large Language Models Rank Fairly? An Empirical Study on the Fairness of LLMs as Rankers	Main Conference	Poster
651	TabSQLify: Enhancing Reasoning Capabilities of LLMs Through Table Decomposition	Main Conference	Poster
654	Compensate Quantization Errors: Make Weights Hierarchical to Compensate Each Other	Findings	Virtual Poster
657	Contextual Label Projection for Cross-Lingual Structured Prediction	Main Conference	Poster
658	Event Detection from Social Media for Epidemic Prediction	Main Conference	Oral
659	RESPROMPT: Residual Connection Prompting Advances Multi-Step Reasoning in Large Language Models	Main Conference	Virtual Poster
660	Contrastive Preference Learning for Neural Machine Translation	Findings	Virtual Poster
663	BPE-knockout: Pruning Pre-existing BPE Tokenisers with Backwards-compatible Morphological Semi-supervision	Main Conference	Poster
665	How are Prompts Different in Terms of Sensitivity?	Main Conference	Poster
666	LSTDial: Enhancing Dialogue Generation via Long- and Short-Term Measurement Feedback	Main Conference	Virtual Poster
667	SocREval: Large Language Models with the Socratic Method for Reference-free Reasoning Evaluation	Findings	Virtual Poster
669	The ART of LLM Refinement: Ask, Refine, and Trust	Main Conference	Oral
671	Modularized Multilingual NMT with Fine-grained Interlingua	Main Conference	Oral
672	Multilingual Machine Translation with Large Language Models: Empirical Results and Analysis	Findings	Poster
673	Unleashing the Power of LLMs in Court View Generation by Stimulating Internal Knowledge and Incorporating External Knowledge	Findings	Virtual Poster
674	Prompting Vision-Language Models For Aspect-Controlled Generation of Referring Expressions	Findings	Poster
677	ParallelPARC: A Scalable Pipeline for Generating Natural-Language Analogies	Main Conference	Poster
678	AWESOME: GPU Memory-constrained Long Document Summarization using Memory Mechanism and Global Salient Content	Main Conference	Virtual Poster
679	Task-Agnostic Detector for Insertion-Based Backdoor Attacks	Findings	Poster
680	Uncertainty Estimation on Sequential Labeling via Uncertainty Transmission	Findings	Virtual Poster
681	NLP Systems That Can't Tell Use from Mention Censor Counterspeech, but Teaching the Distinction Helps	Main Conference	Poster
682	Debiasing with Sufficient Projection: A General Theoretical Framework for Vector Representations	Main Conference	Poster
684	Semi-Supervised Dialogue Abstractive Summarization via High-Quality Pseudolabel Selection	Main Conference	Oral
686	AfriMTE and AfriCOMET: Enhancing COMET to Embrace Under-resourced African Languages	Main Conference	Poster
687	Exploring Language Model's Code Generation Ability with Auxiliary Functions	Findings	Poster
688	Lost in Translation? Translation Errors and Challenges for Fair Assessment of Text-to-Image Models on Multilingual Concepts	Main Conference	Oral
689	Crossing Linguistic Horizons: Finetuning and Comprehensive Evaluation of Vietnamese Large Language Models	Findings	Poster
690	GoT: Effective Graph-of-Thought Reasoning in Language Models	Findings	Virtual Poster
693	Enhancing the General Agent Capabilities of Low-Paramter LLMs through Tuning and Multi-Branch Reasoning	Findings	Virtual Poster
694	MuMath: Multi-perspective Data Augmentation for Mathematical Reasoning in Large Language Models	Findings	Virtual Poster
695	Self-Improving for Zero-Shot Named Entity Recognition with Large Language Models	Main Conference	Virtual Poster
696	TableLlama: Towards Open Large Generalist Models for Tables	Main Conference	Oral
697	PEMA: An Offsite-Tunable Plug-in External Memory Adaptation for Language Models	Main Conference	Poster
699	Tram: A Token-level Retrieval-augmented Mechanism for Source Code Summarization	Findings	Poster
700	Backdooring Instruction-Tuned Large Language Models with Virtual Prompt Injection	Main Conference	Virtual Poster
701	UNO-DST: Leveraging Unlabelled Data in Zero-Shot Dialogue State Tracking	Findings	Poster

PaperID	Title	Proceedings	Format
702	Lifelong Event Detection with Embedding Space Separation and Compaction	Main Conference	Virtual Poster
706	Evaluating Step-by-Step Reasoning through Symbolic Verification	Findings	Virtual Poster
707	Exploring the Factual Consistency in Dialogue Comprehension of Large Language Models	Main Conference	Poster
709	Multilingual Pretraining and Instruction Tuning Improve Cross-Lingual Knowledge Alignment, But Only Shallowly	Main Conference	Virtual Poster
710	A Study on the Calibration of In-context Learning	Main Conference	Virtual Poster
711	DialogBench: Evaluating LLMs as Human-like Dialogue Systems	Main Conference	Virtual Poster
712	GINopic: Topic Modeling with Graph Isomorphism Network	Main Conference	Poster
713	CMB: A Comprehensive Medical Benchmark in Chinese	Main Conference	Virtual Poster
717	Multi-Review Fusion-in-Context	Findings	Poster
718	Retrieving Examples from Memory for Retrieval Augmented Neural Machine Translation: A Systematic Comparison	Findings	Poster
719	Extending Input Contexts of Language Models through Training on Segmented Sequences	Findings	Poster
720	Reason from Fallacy: Enhancing Large Language Models' Logical Reasoning through Logical Fallacy Understanding	Findings	Poster
722	Massive End-to-end Speech Recognition Models with Time Reduction	Main Conference	Poster
723	Language Models (Mostly) Do Not Consider Emotion Triggers When Predicting Emotion	Main Conference	Poster
724	SlimFit: Memory-Efficient Fine-Tuning of Transformer-based Models Using Training Dynamics	Main Conference	Virtual Poster
726	Exploring Automated Distractor Generation for Math Multiple-choice Questions via Large Language Models	Findings	Poster
727	Effective Large Language Model Adaptation for Improved Grounding and Citation Generation	Main Conference	Poster
730	Assisting in Writing Wikipedia-like Articles From Scratch with Large Language Models	Main Conference	Poster
731	Grounding Gaps in Language Model Generations	Main Conference	Oral
733	When Does Monolingual Data Help Multilingual Translation: The Role of Domain and Model Scale	Main Conference	Poster
735	ContraSim – Analyzing Neural Representations Based on Contrastive Learning	Main Conference	Poster
736	Universal Prompt Optimizer for Safe Text-to-Image Generation	Main Conference	Poster
738	Language Model Based Unsupervised Dependency Parsing with Conditional Mutual Information and Grammatical Constraints	Main Conference	Poster
739	The Bias Amplification Paradox in Text-to-Image Generation	Main Conference	Poster
740	Grammar-based Data Augmentation for Low-Resource Languages: The Case of Guarani-Spanish Neural Machine Translation	Main Conference	Poster
741	Aspect-based Sentiment Analysis with Context Denoising	Findings	Virtual Poster
742	IruMozhi: Automatically classifying diglossia in Tamil	Findings	Poster
743	Global Gallery: The Fine Art of Painting Culture Portraits through Multilingual Instruction Tuning	Main Conference	Poster
744	CPopQA: Ranking Cultural Concept Popularity by LLMs	Main Conference	Poster
747	Toward Interactive Regional Understanding in Vision-Large Language Models	Main Conference	Poster
748	RENOVI: A Benchmark Towards Remediating Norm Violations in Socio-Cultural Conversations	Findings	Virtual Poster
749	The Impact of Language on Arithmetic Proficiency: A Multilingual Investigation with Cross-Agent Checking Computation	Main Conference	Oral
750	Human-in-the-Loop Synthetic Text Data Inspection with Provenance Tracking	Findings	Poster
753	COMMIT: Code-Mixing English-Centric Large Language Model for Multilingual Instruction Tuning	Findings	Poster
754	Script-mix: Mixing Scripts for Low-resource Language Parsing	Main Conference	Poster
756	MT-PATCHER: Selective and Extendable Knowledge Distillation from Large Language Models for Machine Translation	Main Conference	Poster
757	ToXCL: A Unified Framework for Toxic Speech Detection and Explanation	Main Conference	Virtual Poster
758	\$\textit{LinkPrompt}\$: Natural and Universal Adversarial Attacks on Prompt-based Language Models	Main Conference	Poster
759	CoE-SQL: In-Context Learning for Multi-Turn Text-to-SQL with Chain-of-Editions	Main Conference	Poster
761	ContraDoc: Understanding Self-Contradictions in Documents with Large Language Models	Main Conference	Poster
762	Entity Disambiguation via Fusion Entity Decoding	Main Conference	Poster

PaperID	Title	Proceedings	Format
763	DiLM: Distilling Dataset into Language Model for Text-level Dataset Distillation	Findings	Poster
764	PlanRAG: A Plan-then-Retrieval Augmented Generation for Generative Large Language Models as Decision Makers	Main Conference	Poster
765	GPTScore: Evaluate as You Desire	Main Conference	Poster
767	MindAgent: Emergent Gaming Interaction	Findings	Virtual Poster
768	A Survey of Confidence Estimation and Calibration in Large Language Models	Main Conference	Poster
769	Not All Metrics Are Guilty: Improving NLG Evaluation by Diversifying References	Main Conference	Virtual Poster
770	Separation and Fusion: A Novel Multiple Token Linking Model for Event Argument Extraction	Main Conference	Virtual Poster
771	BotChat: Evaluating LLMs' Capabilities of Having Multi-Turn Dialogues	Findings	Poster
772	Learning Mutually Informed Representations for Characters and Subwords	Findings	Poster
773	A Novel Two-step Fine-tuning Framework for Transfer Learning in Low-Resource Neural Machine Translation	Findings	Virtual Poster
774	The Integration of Semantic and Structural Knowledge in Knowledge Graph Entity Typing	Main Conference	Oral
775	ComCLIP: Training-Free Compositional Image and Text Matching	Main Conference	Poster
777	Enhancing Cross-lingual Sentence Embedding for Low-resource Languages with Word Alignment	Findings	Virtual Poster
778	C ³ LPGCN: Integrating Contrastive Learning and Cooperative Learning with Prompt into Graph Convolutional Network for Aspect-based Sentiment Analysis	Findings	Virtual Poster
779	ACLSum: A New Dataset for Aspect-based Summarization of Scientific Publications	Main Conference	Poster
780	Visual Enhanced Entity-Level Interaction Network for Multimodal Summarization	Findings	Virtual Poster
781	XAL: EXplainable Active Learning Makes Classifiers Better Low-resource Learners	Main Conference	Virtual Poster
783	Knowledgeable In-Context Tuning: Exploring and Exploiting Factual Knowledge for In-Context Learning	Findings	Virtual Poster
784	LaDiC: Are Diffusion Models Really Inferior to Autoregressive Counterparts for Image-to-Text Generation?	Main Conference	Virtual Poster
786	Efficient Information Extraction in Few-Shot Relation Classification through Contrastive Representation Learning	Main Conference	Oral
788	Intent-conditioned and Non-toxic Counterspeech Generation using Multi-Task Instruction Tuning with RLAIIF	Main Conference	Virtual Poster
789	Attacks, Defenses and Evaluations for LLM Conversation Safety: A Survey	Main Conference	Virtual Poster
790	Mind's Mirror: Distilling Self-Evaluation Capability and Comprehensive Thinking from Large Language Models	Main Conference	Virtual Poster
791	Divergent Token Metrics: Measuring degradation to prune away LLM components – and optimize quantization	Main Conference	Virtual Poster
792	Beyond Performance: Quantifying and Mitigating Label Bias in LLMs	Main Conference	Poster
793	Instructing Large Language Models to Identify and Ignore Irrelevant Conditions	Main Conference	Poster
794	Time Machine GPT	Findings	Virtual Poster
795	An End-to-End Submodular Framework for Data-Efficient In-Context Learning	Findings	Poster
796	On the Relationship Between Non-deterministic FSLMs and RNN LMs	Main Conference	Virtual Poster
797	Transformer LMs Can Represent Probabilistic Strictly Local Languages	Main Conference	Virtual Poster
798	The Role of n -gram Smoothing in the Age of Neural Networks	Main Conference	Virtual Poster
799	Reliability Estimation of News Media Sources: Birds of a Feather Flock Together	Main Conference	Poster
800	On the Multilingual Ability of Decoder-based Pre-trained Language Models: Finding and Controlling Language-Specific Neurons	Main Conference	Poster
801	NLP Progress in Indigenous Latin American Languages	Main Conference	Oral
802	Teaching Llama a New Language Through Cross-Lingual Knowledge Transfer	Findings	Poster
804	On the Effectiveness of Adversarial Robustness for Abuse Mitigation with Counterspeech	Main Conference	Virtual Poster
805	Leveraging the Structure of Pre-trained Embeddings to Minimize Annotation Effort	Main Conference	Oral
807	UniArk: Improving Generalisation and Consistency for Factual Knowledge Extraction through Debiasing	Main Conference	Poster
810	Adaptive-RAG: Learning to Adapt Retrieval-Augmented Large Language Models through Question Complexity	Main Conference	Poster
811	Knowing What LLMs DO NOT Know: A Simple Yet Effective Self-Detection Method	Main Conference	Virtual Poster

PaperID	Title	Proceedings	Format
812	Are Large Language Model Temporally Grounded?	Main Conference	Poster
813	Simulating Opinion Dynamics with Networks of LLM-based Agents	Findings	Poster
814	A diverse Multilingual News Headlines Dataset from around the World	Main Conference	Poster
815	Probing the Category of Verbal Aspect in Transformer Language Models	Findings	Poster
816	Document Image Machine Translation with Dynamic Multi-pre-trained Models Assembling	Main Conference	Virtual Poster
817	A Measure for Transparent Comparison of Linguistic Diversity in Multilingual NLP Data Sets	Findings	Poster
818	Beyond Read-Only: Crafting a Comprehensive Chinese Text-to-SQL Dataset for Database Manipulation and Query	Findings	Virtual Poster
819	The Unreasonable Effectiveness of Random Target Embeddings for Continuous-Output Neural Machine Translation	Main Conference	Oral
820	Efficient Sample-Specific Encoder Perturbations	Main Conference	Virtual Poster
822	Normalizing without Modernizing: Keeping Historical Wordforms of Middle French while Reducing Spelling Variants	Findings	Virtual Poster
823	Elastic Weight Removal for Faithful and Abstractive Dialogue Generation	Main Conference	Poster
825	R-Tuning: Instructing Large Language Models to Say 'I Don't Know'	Main Conference	Oral
826	Bridging the Gap between Different Vocabularies for LLM Ensemble	Main Conference	Virtual Poster
828	KnowLA: Enhancing Parameter-efficient Finetuning with Knowledgeable Adaptation	Main Conference	Virtual Poster
829	Anti-LM Decoding for Zero-shot In-context Machine Translation	Findings	Virtual Poster
830	Extremely Weakly-supervised Text Classification with Wordsets Mining and Sync-Denoising	Main Conference	Virtual Poster
831	Defending Against Weight-Poisoning Backdoor Attacks for Parameter-Efficient Fine-Tuning	Findings	Virtual Poster
832	F-MALLOC: Feed-forward Memory Allocation for Continual Learning in Neural Machine Translation	Main Conference	Virtual Poster
833	Towards Reducing Diagnostic Errors with Interpretable Risk Prediction	Main Conference	Poster
834	Select and Summarize: Scene Saliency for Movie Script Summarization	Findings	Poster
836	Generalizable Multilingual Hate Speech Detection on Low Resource Indian Languages using Fair Selection in Federated Learning	Main Conference	Poster
837	Key ingredients for effective zero-shot cross-lingual knowledge transfer in generative tasks	Main Conference	Poster
839	Don't be a Fool: Pooling Strategies in Offensive Language Detection from User-Intended Adversarial Attacks	Findings	Poster
840	The Impact of Depth on Compositional Generalization in Transformer Language Models	Main Conference	Poster
842	Z-GMOT: Zero-shot Generic Multiple Object Tracking	Findings	Virtual Poster
843	Pregnant Questions: The Importance of Pragmatic Awareness in Maternal Health Question Answering	Main Conference	Poster
846	Towards Explainability in Legal Outcome Prediction Models	Main Conference	Virtual Poster
849	NLP for Counterspeech against Hate: A Survey and How-To Guide	Findings	Poster
850	PRODIGy: a PROFILE-based Dialogue Generation dataset	Findings	Poster
852	The steerability of large language models toward data-driven personas	Main Conference	Oral
853	WaterJudge: Quality-Detection Trade-off when Watermarking Large Language Models	Findings	Poster
854	Cognitive Overload: Jailbreaking Large Language Models with Overloaded Logical Thinking	Findings	Poster
856	CCSum: A Large-Scale and High-Quality Dataset for Abstractive News Summarization	Main Conference	Oral
857	Capturing Perspectives of Crowdsourced Annotators in Subjective Learning Tasks	Main Conference	Poster
858	PAELLA: Parameter-Efficient Lightweight Language-Agnostic Captioning Model	Findings	Poster
859	Improving Factual Accuracy of Neural Table-to-Text Output by Addressing Input Problems in ToTTo	Main Conference	Poster
861	CERET: Cost-Effective Extrinsic Refinement for Text Generation	Main Conference	Virtual Poster
863	Parameter-Efficient Instruction Tuning of Large Language Models For Extreme Financial Numeral Labelling	Main Conference	Oral
865	Diverse Perspectives, Divergent Models: Cross-Cultural Evaluation of Depression Detection on Twitter	Main Conference	Poster
866	OSCaR: Object State Captioning and State Change Representation	Findings	Virtual Poster
867	SumCSE: Summary as a transformation for Contrastive Learning	Findings	Poster

PaperID	Title	Proceedings	Format
868	The Curious Decline of Linguistic Diversity: Training Language Models on Synthetic Text	Findings	Poster
869	Analysis of State-Level Legislative Process in Enhanced Linguistic and Nationwide Network Contexts	Main Conference	Virtual Poster
870	PersonaLLM: Investigating the Ability of Large Language Models to Express Personality Traits	Findings	Poster
872	DeMuX: Data-efficient Multilingual Learning	Main Conference	Poster
873	Removing RLHF Protections in GPT-4 via Fine-Tuning	Main Conference	Oral
875	MoE-SLU: Towards ASR-Robust Spoken Language Understanding via Mixture-of-Experts	Findings	Poster
876	FIRE: A Dataset for Financial Relation Extraction	Findings	Virtual Poster
880	MusiLingo: Bridging Music and Text with Pre-trained Language Models for Music Captioning and Query Response	Findings	Poster
882	Investigating Acceleration of LLaMA Inference by Enabling Intermediate Layer Decoding via Instruction Tuning with 'LITE'	Findings	Virtual Poster
885	DUQGen: Effective Unsupervised Domain Adaptation of Neural Rankers by Diversifying Synthetic Query Generation	Main Conference	Oral
887	How did we get here? Summarizing conversation dynamics	Main Conference	Poster
888	Unveiling Implicit Toxicity in Life Advice	Main Conference	Oral
890	Instruction-following Evaluation through Verbalizer Manipulation	Findings	Virtual Poster
891	WebWISE: Unlocking Web Interface Control for LLMs via Sequential Exploration	Findings	Virtual Poster
892	Can Language Model Moderators Improve the Health of Online Discourse?	Main Conference	Poster
893	CodeLM: Aligning Language Models with Tailored Synthetic Data	Findings	Virtual Poster
898	LeanReasoner: Boosting Complex Logical Reasoning with Lean	Main Conference	Poster
900	Prompting Few-shot Multi-hop Question Generation via Comprehending Type-aware Semantics	Findings	Virtual Poster
904	When Hindsight is Not 20/20: Testing Limits on Reflective Thinking in Large Language Models	Findings	Poster
905	CoDa: Constrained Generation based Data Augmentation for Low-Resource NLP	Findings	Poster
906	UICoder: Finetuning Large Language Models to Generate User Interface Code through Automated Feedback	Main Conference	Poster
907	Measuring Cross-lingual Transfer in Bytes	Main Conference	Poster
908	Synonym relations affect object detection learned on vision-language data	Findings	Poster
909	CM-TTS: Enhancing Real Time Text-to-Speech Synthesis Efficiency through Weighted Samplers and Consistency Models	Findings	Virtual Poster
910	Arithmetic Reasoning with LLM: Prolog Generation & Permutation	Main Conference	Poster
911	MisgenderMender: A Community-Informed Approach to Interventions for Misgendering	Main Conference	Poster
912	Interplay of Machine Translation, Diacritics, and Diacritization	Main Conference	Poster
913	From Quantity to Quality: Boosting LLM Performance with Self-Guided Data Selection for Instruction Tuning	Main Conference	Poster
916	RobustSentEmbed: Robust Sentence Embeddings Using Adversarial Self-Supervised Contrastive Learning	Findings	Virtual Poster
917	Safer-Instruct: Aligning Language Models with Automated Preference Data	Main Conference	Oral
918	PELMS: Pre-training for Effective Low-Shot Multi-Document Summarization	Main Conference	Poster
919	Characterizing Human and Zero-Shot GPT-3.5 Object-Similarity Judgments	Findings	Poster
920	Deceptive Semantic Shortcuts on Reasoning Chains: How Far Can Models Go without Hallucination?	Main Conference	Poster
921	IndiSentiment140: Sentiment Analysis Dataset for Indian Languages with Emphasis on Low-Resource Languages using Machine Translation	Main Conference	Poster
922	Leveraging LLMs for Synthesizing Training Data Across Many Languages in Multilingual Dense Retrieval	Main Conference	Virtual Poster
924	Verifying Claims About Metaphors with Large-Scale Automatic Metaphor Identification	Main Conference	Oral
926	SCANNER: Knowledge-Enhanced Approach for Robust Multi-modal Named Entity Recognition of Unseen Entities	Main Conference	Poster
928	A Theory Guided Scaffolding Instruction Framework for LLM-Enabled Metaphor Reasoning	Main Conference	Virtual Poster
929	Self-Demos: Eliciting Out-of-Demonstration Generalizability in Large Language Models	Findings	Virtual Poster
932	Learning to Compress Prompt in Natural Language Formats	Main Conference	Virtual Poster
933	Getting Sick After Seeing a Doctor? Diagnosing and Mitigating Knowledge Conflicts in Event Temporal Reasoning	Findings	Virtual Poster

PaperID	Title	Proceedings	Format
935	InstructABSA: Instruction Learning for Aspect Based Sentiment Analysis	Main Conference	Oral
936	MCECR: A Novel Dataset for Multilingual Cross-Document Event Coreference Resolution	Findings	Virtual Poster
937	Automatic, Meta and Human Evaluation for Multimodal Summarization with Multimodal Output	Main Conference	Poster
940	Naive Bayes-based Context Extension for Large Language Models	Main Conference	Oral
942	Leitner-Guided Memory Replay for Cross-lingual Continual Learning	Main Conference	Oral
944	Multilingual Nonce Dependency Treebanks: Understanding how Language Models Represent and Process Syntactic Structure	Main Conference	Oral
945	Actively Learn from LLMs with Uncertainty Propagation for Generalized Category Discovery	Main Conference	Virtual Poster
948	Sentiment Analysis in the Era of Large Language Models: A Reality Check	Findings	Virtual Poster
949	Explaining Text Similarity in Transformer Models	Main Conference	Poster
951	Tokenizer Choice For LLM Training: Negligible or Crucial?	Findings	Poster
955	Think Before You Speak: Cultivating Communication Skills of Large Language Models via Inner Monologue	Findings	Poster
956	Large Language Models can Contrastively Refine their Generation for Better Sentence Representation Learning	Main Conference	Virtual Poster
957	MEMORY-VQ: Compression for Tractable Internet-Scale Memory	Main Conference	Poster
959	HIL: Hybrid Isotropy Learning for Zero-shot Performance in Dense retrieval	Main Conference	Virtual Poster
962	SuperGLEBer: German Language Understanding Evaluation Benchmark	Main Conference	Poster
963	You are an expert annotator: Automatic Best-Worst-Scaling Annotations for Emotion Intensity Modeling	Main Conference	Poster
964	The Impact of Differential Privacy on Group Disparity Mitigation	Findings	Virtual Poster
965	Isometric Neural Machine Translation using Phoneme Count Ratio Reward-based Reinforcement Learning	Findings	Poster
966	What Matters in Training a GPT4-Style Language Model with Multimodal Inputs?	Main Conference	Oral
968	Defining and Detecting Vulnerability in Human Evaluation Guidelines: A Preliminary Study Towards Reliable NLG Evaluation	Main Conference	Poster
969	Read between the lines - Functionality Extraction From READMEs	Findings	Virtual Poster
970	AbsPyramid: Benchmarking the Abstraction Ability of Language Models with a Unified Entailment Graph	Findings	Poster
971	MOSAICo: a Multilingual Open-text Semantically Annotated Interlinked Corpus	Main Conference	Oral
972	Few-TK: A Dataset for Few-shot Scientific Typed Keyphrase Recognition	Findings	Poster
973	SemRoDe: Macro Adversarial Training to Learn Representations that are Robust to Word-Level Attacks	Main Conference	Oral
976	BUST: Benchmark for the evaluation of detectors of LLM-Generated Text	Main Conference	Poster
977	Improving In-context Learning of Multilingual Generative Language Models with Cross-lingual Alignment	Main Conference	Virtual Poster
978	MaCSC: Towards Multimodal-augmented Pre-trained Language Models via Conceptual Prototypes and Self-balancing Calibration	Main Conference	Oral
979	Does Pre-trained Language Model Actually Infer Unseen Links in Knowledge Graph Completion?	Main Conference	Oral
982	Language Models can be Deductive Solvers	Findings	Poster
984	Interpreting User Requests in the Context of Natural Language Standing Instructions	Findings	Virtual Poster
986	Secure Your Model: An Effective Key Prompt Protection Mechanism for Large Language Models	Findings	Virtual Poster
987	Discovering Lobby-Parliamentarian Alignments through NLP	Main Conference	Poster
988	IterCQR: Iterative Conversational Query Reformulation with Retrieval Guidance	Main Conference	Poster
989	Unveiling the Magic: Investigating Attention Distillation in Retrieval-Augmented Generation	Main Conference	Virtual Poster
991	Improving Factuality in Clinical Abstractive Multi-Document Summarization by Guided Continued Pre-training	Main Conference	Oral
993	Enhancing Chain-of-Thoughts Prompting with Iterative Bootstrapping in Large Language Models	Findings	Virtual Poster
994	AceGPT, Localizing Large Language Models in Arabic	Main Conference	Poster
996	Improving Machine Translation with Human Feedback: An Exploration of Quality Estimation as a Reward Model	Main Conference	Poster
1000	Depression Detection in Clinical Interviews with LLM-Empowered Structural Element Graph	Main Conference	Virtual Poster
1002	Do Prompt Positions Really Matter?	Findings	Poster

PaperID	Title	Proceedings	Format
1005	Natural Language Embedded Programs for Hybrid Language Symbolic Reasoning	Findings	Poster
1007	A Study on Scaling Up Multilingual News Framing Analysis	Findings	Poster
1009	SQATIN: Supervised Instruction Tuning Meets Question Answering for Improved Dialogue NLU	Main Conference	Poster
1010	Enhancing Argument Summarization: Prioritizing Exhaustiveness in Key Point Generation and Introducing an Automatic Coverage Evaluation Metric	Main Conference	Oral
1012	ARM: Alignment with Residual Energy-Based Model	Main Conference	Oral
1013	ViGLUE: A Vietnamese General Language Understanding Benchmark and Analysis of Vietnamese Language Models	Findings	Virtual Poster
1014	HumanRankEval: Automatic Evaluation of LMs as Conversational Assistants	Main Conference	Poster
1015	FAMuS: Frames Across Multiple Sources	Main Conference	Virtual Poster
1016	Rationale-based Opinion Summarization	Main Conference	Virtual Poster
1018	Mustango: Toward Controllable Text-to-Music Generation	Main Conference	Oral
1019	Exploring the Trade-off Between Model Performance and Explanation Plausibility of Text Classifiers Using Human Rationales	Findings	Poster
1020	Adaptive Cross-lingual Text Classification through In-Context One-Shot Demonstrations	Main Conference	Oral
1021	CNER: Concept and Named Entity Recognition	Main Conference	Poster
1022	MuLan: A Study of Fact Mutability in Language Models	Main Conference	Poster
1028	Unlocking Parameter-Efficient Fine-Tuning for Low-Resource Language Translation	Findings	Poster
1029	ADaPT: As-Needed Decomposition and Planning with Language Models	Findings	Poster
1031	Language-Independent Representations Improve Zero-Shot Summarization	Main Conference	Poster
1033	Branch-Solve-Merge Improves Large Language Model Evaluation and Generation	Main Conference	Virtual Poster
1034	Guiding Large Language Models to Post-Edit Machine Translation with Error Annotations	Findings	Poster
1036	Trusting Your Evidence: Hallucinate Less with Context-aware Decoding	Main Conference	Poster
1037	REPLUG: Retrieval-Augmented Black-Box Language Models	Main Conference	Oral
1038	Non-contrastive sentence representations via self-supervision	Findings	Poster
1039	Semantically-Prompted Language Models Improve Visual Descriptions	Findings	Poster
1040	David helps Goliath: Inference-Time Collaboration Between Small Specialized and Large General Diffusion LMs	Main Conference	Oral
1041	GenTKG: Generative Forecasting on Temporal Knowledge Graph	Findings	Poster
1042	Efficient End-to-End Visual Document Understanding with Rationale Distillation	Main Conference	Poster
1044	A Systematic Comparison of Syllogistic Reasoning in Humans and Language Models	Main Conference	Oral
1046	Computationally Efficient Active Learning for Large Imbalanced Datasets	Main Conference	Virtual Poster
1049	A Transformer Language Model with Stack Attention	Findings	Virtual Poster
1050	InstructEval: Systematic Evaluation of Instruction Selection Methods	Findings	Virtual Poster
1051	ICLE++: Modeling Fine-Grained Traits for Holistic Essay Scoring	Main Conference	Poster
1052	RecMind: Large Language Model Powered Agent For Recommendation	Findings	Poster
1053	GOLD: Generalized Knowledge Distillation via Out-of-Distribution-Guided Language Data Generation	Findings	Poster
1055	How Lexical is Bilingual Lexicon Induction?	Findings	Poster
1057	Fumbling in Babel: An Investigation into ChatGPT's Language Identification Ability	Findings	Poster
1058	Targeted Augmentation for Low-Resource Event Extraction	Findings	Virtual Poster
1059	Asking More Informative Questions for Grounded Retrieval	Findings	Poster
1060	GuyLingo: The Republic of Guyana Creole Corpora	Main Conference	Oral
1064	UNcommonsense Reasoning: Abductive Reasoning about Uncommon Situations	Main Conference	Poster
1065	To Tell The Truth: Language of Deception and Language Models	Main Conference	Oral
1066	Multilingual Models for ASR in Chibchan Languages	Main Conference	Poster

PaperID	Title	Proceedings	Format
1067	LegalDiscourse: Interpreting When Laws Apply and To Whom	Main Conference	Oral
1071	X-Eval: Generalizable Multi-aspect Text Evaluation via Augmented Instruction Tuning with Auxiliary Evaluation Aspects	Main Conference	Poster
1073	Is Reference Necessary in the Evaluation of NLG Systems? When and Where?	Main Conference	Virtual Poster
1074	Efficient Citer: Tuning Large Language Models for Enhanced Answer Quality and Verification	Findings	Virtual Poster
1075	Addressing Healthcare-related Racial and LGBTQ+ Biases in Pretrained Language Models	Findings	Virtual Poster
1078	Semi-Structured Chain-of-Thought: Integrating Multiple Sources of Knowledge for Improved Language Model Reasoning	Main Conference	Virtual Poster
1079	Evaluating the Deductive Competence of Large Language Models	Main Conference	Poster
1080	DoubleLingo: Causal Estimation with Large Language Models	Main Conference	Poster
1082	ATG: Benchmarking Automated Theorem Generation for Generative Language Models	Findings	Virtual Poster
1084	Large Human Language Models: A Need and the Challenges	Main Conference	Oral
1087	Benchmarking Generation and Evaluation Capabilities of Large Language Models for Instruction Controllable Summarization	Findings	Virtual Poster
1088	On Learning to Summarize with Large Language Models as References	Main Conference	Virtual Poster
1091	NeuroComparatives: Neuro-Symbolic Distillation of Comparative Knowledge	Findings	Virtual Poster
1094	Emotion-Anchored Contrastive Learning Framework for Emotion Recognition in Conversation	Findings	Virtual Poster
1097	SUQL: Conversational Search over Structured and Unstructured Data with Large Language Models	Findings	Virtual Poster
1099	On Evaluating the Integration of Reasoning and Action in LLM Agents with Database Question Answering	Findings	Virtual Poster
1100	Hallucination Diversity-Aware Active Learning for Text Summarization	Main Conference	Virtual Poster
1101	CARE: Extracting Experimental Findings From Clinical Literature	Findings	Poster
1102	Personalized Federated Learning for Text Classification with Gradient-Free Prompt Tuning	Findings	Virtual Poster
1103	SGSH: Stimulate Large Language Models with Skeleton Heuristics for Knowledge Base Question Generation	Findings	Virtual Poster
1104	Biomedical Entity Representation with Graph-Augmented Multi-Objective Transformer	Findings	Poster
1105	{K}keep it {P}rivate: Unsupervised Privatization of Online Text	Main Conference	Poster
1108	Improved Text Emotion Prediction Using Combined Valence and Arousal Ordinal Classification	Main Conference	Virtual Poster
1109	Cross-Lingual Summarization with Pseudo-Label Regularization	Findings	Virtual Poster
1110	Tied-LoRA: Enhancing parameter efficiency of LoRA with Weight Tying	Main Conference	Virtual Poster
1114	On the Way to Gentle AI Counselor: Politeness Cause Elicitation and Intensity Tagging in Code-mixed Hinglish Conversations for Social Good	Findings	Virtual Poster
1116	Investigating Data Contamination in Modern Benchmarks for Large Language Models	Main Conference	Poster
1119	Pre-trained Language Models for Entity Blocking: A Reproducibility Study	Main Conference	Virtual Poster
1120	On Narrative Question Answering Skills	Main Conference	Poster
1122	Leveraging Summarization for Unsupervised Dialogue Topic Segmentation	Findings	Poster
1123	\$RE^2\$: Region-Aware Relation Extraction from Visually Rich Documents	Main Conference	Poster
1124	Mix-Initiative Response Generation with Dynamic Prefix Tuning	Main Conference	Virtual Poster
1126	Value FULCRA: Mapping Large Language Models to the Multidimensional Spectrum of Basic Human Value	Main Conference	Poster
1127	Order-Based Pre-training Strategies for Procedural Text Understanding	Main Conference	Poster
1129	LLaMA-Rider: Spurring Large Language Models to Explore the Open World	Findings	Virtual Poster
1130	Contrastive Learning as a Polarizer: Mitigating Gender Bias by Fair and Biased sentences	Findings	Poster
1133	PolLMgraph: Unraveling Hallucinations in Large Language Models via State Transition Dynamics	Findings	Poster
1137	Breaking the Language Barrier: Can Direct Inference Outperform Pre-Translation in Multilingual LLM Applications?	Main Conference	Oral
1143	Improving Open-Domain Health Question Answering with Evidence Quality Assessment	Findings	Virtual Poster
1146	DecoderLens: Layerwise Interpretation of Encoder-Decoder Transformers	Findings	Virtual Poster
1154	IndiBias: A Benchmark Dataset to Measure Social Biases in Language Models for Indian Context	Main Conference	Poster

PaperID	Title	Proceedings	Format
CL-2374	Can Large Language Models Transform Computational Social Science?	Comp. Linguistics	Poster
TACL-4673	Evaluating a Century of Progress on the Cognitive Science of Adjective Ordering	TACL	Virtual Poster
TACL-4805	MIRACL: A Multilingual Retrieval Dataset Covering 18 Diverse Languages	TACL	Poster
TACL-5433	Data-driven Parsing Evaluation for Child-Parent Interactions	TACL	Virtual Poster
TACL-5449	Explicitly Representing Syntax Improves Sentence-to-layout Prediction of Unexpected Situations	TACL	Poster
TACL-5555	Addressing the Binning Problem in Calibration Assessment through Scalar Annotations	TACL	Poster
TACL-5663	To Diverge or Not to Diverge: A Morphosyntactic Perspective on Machine Translation vs Human Translation	TACL	Virtual Poster
TACL-5691	Large Language Models Enable Few-Shot Clustering	TACL	Poster
TACL-5721	Federated Learning for Exploiting Annotators' Disagreements in Natural Language Processing	TACL	Poster
TACL-5739	Evaluating the Ripple Effects of Knowledge Editing in Language Models	TACL	Poster
TACL-5757	Lost in the Middle: How Language Models Use Long Contexts	TACL	Poster
TACL-5857	Exploring Human-Like Translation Strategy with Large Language Models	TACL	Poster
TACL-5887	Source-Free Domain Adaptation for Question Answering with Masked Self-training	TACL	Poster
TACL-6031	Scope Ambiguities in Large Language Models	TACL	Poster
TACL-6039	What Formal Languages Can Transformers Express? A Survey	TACL	Poster