

The First Workshop on Large Language Model Memorization (L2M2 2025)

Vienna, Austria
1 August 2025

ISBN: 979-8-3313-2411-7

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2025) by the Association for Computational Linguistics
All rights reserved.

Printed with permission by Curran Associates, Inc. (2025)

For permission requests, please contact the Association for Computational Linguistics
at the address below.

Association for Computational Linguistics
209 N. Eighth Street
Stroudsburg, Pennsylvania 18360

Phone: 1-570-476-8006
Fax: 1-570-476-0860

acl@aclweb.org

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

Table of Contents

<i>Factual Knowledge in Language Models: Robustness and Anomalies under Simple Temporal Context Variations</i>	
Hichem Ammar Khodja, Frederic Bechet, Quentin Brabant, Alexis Nasr and Gwénolé Lecorvé .	1
<i>Memorization in Language Models through the Lens of Intrinsic Dimension</i>	
Stefan Arnold	23
<i>From Data to Knowledge: Evaluating How Efficiently Language Models Learn Facts</i>	
Daniel Christoph, Max Ploner, Patrick Haller and Alan Akbik	29
<i>Towards a Principled Evaluation of Knowledge Editors</i>	
Sebastian Pohl, Max Ploner and Alan Akbik	47
<i>On the Way to LLM Personalization: Learning to Remember User Conversations</i>	
Lucie Charlotte Magister, Katherine Metcalf, Yizhe Zhang and Maartje Ter Hoeve	61
<i>From Teacher to Student: Tracking Memorization Through Model Distillation</i>	
Simardeep Singh	78
<i>Understanding Verbatim Memorization in LLMs Through Circuit Discovery</i>	
Ilya Lasy, Peter Knees and Stefan Woltran	83
<i>Quantifying Memorization in Continual Pre-training with Japanese General or Industry-Specific Corpora</i>	
Hiromu Takahashi and Shotaro Ishihara	95
<i>Memorization is Language-Sensitive: Analyzing Memorization and Inference Risks of LLMs in a Multilingual Setting</i>	
Ali Satvaty, Anna Visman, Dan Seidel, Suzan Verberne and Fatih Turkmen	106
<i>Quantifying Memorization and Parametric Response Rates in Retrieval-Augmented Vision-Language Models</i>	
Peter Carragher, Abhinand Jha, Raghav R and Kathleen M. Carley	127
<i>Empirical Evaluation of Loss Masking to Selectively Prevent Memorization</i>	
Tagore Rao Kosireddy and Evan Lucas	142
<i>Bring Your Own Knowledge: A Survey of Methods for LLM Knowledge Expansion</i>	
Mingyang Wang, Alisa Stoll, Lukas Lange, Heike Adel, Hinrich Schuetze and Jannik Strötgen	150
<i>Memorization: A Close Look at Books</i>	
Iris Ma, Ian Domingo, Alberto Krone-Martins, Pierre Baldi and Cristina Lopes	169
<i>Memory Tokens: Large Language Models Can Generate Reversible Sentence Embeddings</i>	
Ignacio Sastre and Aiala Rosá	183
<i>Robust Data Watermarking in Language Models by Injecting Fictitious Knowledge</i>	
Xinyue Cui, Johnny Wei, Swabha Swayamdipta and Robin Jia	190
<i>Better Aligned with Survey Respondents or Training Data? Unveiling Political Leanings of LLMs on U.S. Supreme Court Cases</i>	
Shanshan Xu, Santosh T.y.s.s, Yanai Elazar, Quirin Vogel, Barbara Plank and Matthias Grabmair	205
<i>Capacity Matters: a Proof-of-Concept for Transformer Memorization on Real-World Data</i>	
Anton Changelidis and Aki Härmä	227