

2nd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures (DeeLIO 2021)

Deep Learning Inside Out

Online
10 June 2021

ISBN: 978-1-7138-3022-1

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© (2021) by the Association for Computational Linguistics
All rights reserved.

Printed with permission by Curran Associates, Inc. (2021)

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>Transformer visualization via dictionary learning: contextualized embedding as a linear superposition of transformer factors</i> Zeyu Yun, Yubei Chen, Bruno Olshausen and Yann LeCun	1
<i>Reconstructing Implicit Knowledge with Language Models</i> Maria Becker, Siting Liang and Anette Frank	11
<i>Investigating the Effect of Background Knowledge on Natural Questions</i> Vidhisha Balachandran, Bhuwan Dhingra, Haitian Sun, Michael Collins and William Cohen . . .	25
<i>Augmenting Topic Aware Knowledge-Grounded Conversations with Dynamic Built Knowledge Graphs</i> Junjie Wu and Hao Zhou	31
<i>What Makes My Model Perplexed? A Linguistic Investigation on Neural Language Models Perplexity</i> Alessio Miaschi, Dominique Brunato, Felice Dell’Orletta and Giulia Venturi	40
<i>How Do BERT Embeddings Organize Linguistic Knowledge?</i> Giovanni Puccetti, Alessio Miaschi and Felice Dell’Orletta	48
<i>ERNIE-NLI: Analyzing the Impact of Domain-Specific External Knowledge on Enhanced Representations for NLI</i> Lisa Bauer, Lingjia Deng and Mohit Bansal	58
<i>Enhancing Multiple-Choice Question Answering with Causal Knowledge</i> Dhairya Dalal, Mihael Arcan and Paul Buitelaar	70
<i>Low Anisotropy Sense Retrofitting (LAsER) : Towards Isotropic and Sense Enriched Representations</i> Geetanjali Bihani and Julia Rayz	81
<i>KW-ATTN: Knowledge Infused Attention for Accurate and Interpretable Text Classification</i> Hyeju Jang, Seojin Bang, Wen Xiao, Giuseppe Carenini, Raymond Ng and Young ji Lee	96
<i>Multi-input Recurrent Independent Mechanisms for leveraging knowledge sources: Case studies on sentiment analysis and health text mining</i> Parsa Bagherzadeh and Sabine Bergler	108
<i>What BERTs and GPTs know about your brand? Probing contextual language models for affect associations</i> Vivek Srivastava, Stephen Pilli, Savita Bhat, Niranjan Pedanekar and Shirish Karande	119
<i>Attention vs non-attention for a Shapley-based explanation method</i> Tom Kersten, Hugh Mee Wong, Jaap Jumelet and Dieuwke Hupkes	129
<i>Predicting Numerals in Natural Language Text Using a Language Model Considering the Quantitative Aspects of Numerals</i> Taku Sakamoto and Akiko Aizawa	140