

NAACL-HLT 2022

**The 2022 Conference of the North American Chapter of the
Association for Computational Linguistics: Human Language
Technologies**

Industry Papers

July 10-15, 2022

©2022 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
acl@aclweb.org

ISBN 978-1-955917-72-8

Introduction

We welcome you to NAACL-HLT Industry Track 2022.

The industry track was first introduced to a major NLP conference at NAACL-HLT 2018 in New Orleans and has since become a standard track at the NAACL conferences. The Industry track provides a forum for researchers, engineers, and application developers to exchange ideas, share results and discuss use cases of successful deployment of language technologies in real-world settings. It also inspired industry tracks at other conferences such as COLING and EMNLP.

Each conference year is unique, and 2022 was no exception. The roll-out of COVID vaccines made it possible to return to in-person or at least hybrid conferences. Yet the unpredictability of this pandemic meant that even at the time of writing this preface, we do not know how many people will be comfortable attending in-person and who will prefer remote attendance.

Another big change in 2022 was the transition of the main track to ACL rolling review (ARR). Since the criteria for Industry track papers differ from those used in ARR, we needed to organize a separate review process. However, we used the OpenReview platform to streamline the author experience. This decision came with a steep learning curve for us, and we are grateful to all authors and reviewers for being flexible and working with us as we were learning the ropes. We are also very thankful to the OpenReview support, who were always available to answer our questions.

Finally, we introduced senior area chairs. We selected ten experienced researchers with a broad range of industry and academic experience who wrote meta-reviews for each paper and helped us make the final decisions.

This year we received 128 paper submissions. Five submissions were rejected without review due to incompleteness, non-compliance with format requirements, or submission policies (such as the double submission policy). Our program committee reviewed the remaining 123 papers with a rich representation of the present spectrum of NLP researchers and professionals. Each submission was reviewed by at least three members of the program committee. Reviews solicited committee opinions along two primary aspects: Focus on real-world applications and lessons offered by the paper. Reviews also considered clarity, methodological rigor, ethical use of datasets, and compliance with conference guidelines. The area chairs then reviewed each paper and the reviews and provided their recommendation along with a short metareview. Finally, we accepted 40 papers based on committee recommendations as well as alignment of the papers with the goals of the industry track (acceptance rate of 32%).

This year, the Industry Track program will consist of two oral sessions (10 papers in total) and one poster session (30 posters). Each oral session will have a diverse set of talks covering the areas of Text Mining, Question Answering, Interactive or Dialog Systems, Summarization, Translation, Speech Technologies, Green NLP, Bias, Fairness, and Ethics. The work presented in the poster session paints a rich picture of the many real-world applications of language technologies and the challenges associated with these applications.

NAACL-HLT 2022 Industry Track also features the now traditional “Careers in NLP” panel discussion. This year we are expanding the range of careers discussed in the panel by introducing a career in product. The panel will be moderated by Yunyao Li, and we expect the conversation to include trends in NLP careers, emerging skills, main challenges and opportunities for cross-functional collaboration as NLP professionals in today’s organizations, and more.

It has been a privilege chairing the Industry Track this year. We thank the conference general chair, Dan Roth, for inviting us to the organizing committee. Thanks also to Program Chairs Marine Carpuat, Marie-Catherine de Marneffe and Ivan Vladimir Meza Ruiz, and all members of the organizing committee. Yunayo Li and Owen Rambow served as advisors for the Industry Track to provide continuity for this new track. We were generously helped by every member of this committee over the past year, and organizing this track was possible only with their advice and efforts. We once again recognize and thank every member of the industry track program committee for volunteering their time. Finally, thanks to the authors and attendees of the industry track for embracing this initiative and offering a reason to continue the industry track at NAACL-HLT conferences.

Anastassia Loukina, Bonan Min, Rashmi Gangadharaiah

Program Committee

Industry Track Co-chairs

Anastassia Loukina, Grammarly
Rashmi Gangadharaiah, AWS AI, Amazon
Bonan Min, Raytheon BBN Technologies

Senior Area Chairs

Alborz Geramifard, Meta
Marjorie Freedman, USC/ISI
Debanjan Ghosh, Educational Testing Service
Siddharth Patwardhan, Apple
Mathias Lambert, Meta
Weiwei Guo, Pinterest
Yang Liu, Amazon
William Hartmann, Raytheon BBN Technologies
Aoife Cahill, Dataminr
Roy Schwartz, Hebrew University of Jerusalem

Program Committee

Tong Guo
Rylan Conway
Jangwon Kim
Shihao Ran
Ralph M. Weischedel
Abhijit Mishra
Alex Acero
Alex Marin
Alexey Romanov
Alok Baikadi
Amar Prakash Azad
Ankush Gupta
Anmol Goel
Annemarie Friedrich
Antonie Lin
Aoife Cahill
Aparna Elangovan
Avinesh P.V.S
Byung-Hak Kim
Beata Beigman Klebanov
Benjamin Han
Benjamin Vladimir Rozonoyer
Brian Lester
Budhaditya Deb
Chanjun Park
Charith Peris

Chinnadhurai Sankar
Chris Brew
Christine Doran
Christopher Hidey
Ciprian Chelba
Constantine Lignos
Damianos Karakos
Daniel Dickinson
David Uthus
Deborah A. Dahl
Deepak Muralidharan
Derrick Higgins
Deyi Xiong
Ehud Reiter
Emre Barut
Fabio Mercurio
Feifei Pan
Frederic Mailhot
Geewook Kim
Heba Elfardy
Honglei Guo
Hyeonseok Moon
Hyung Won Chung
Igor Shalyminov
Ioannis Partalas
J. William Murdock
Jaegul Choo
Jaime Lorenzo-Trueba
Jangwon Kim
Jean-Philippe Fauconnier
Jiangning Chen
Jinseok Nam
Jinyeong Yim
Joel R. Tetreault
John Chen
Joo-Kyung Kim
Judith Gaspers
Jun Seok Kang
Juntao Li
Kai Yu
Kalpa Gunaratna
Kartik Mehta
Kasturi Bhattacharjee
Keith Trnka
Kevin Stowe
Kyle Williams
Laura Chiticariu
Lee Becker
Lei Chen
Lei Shu
Li Dong

Liangming Pan
Ling Tsou
Lingjia Deng
Long Qin
Lynette Hirschman
Manuel Rafael Ciosici
Marek Suppa
Margot Mieskes
Marina Danilevsky
Matthew Mulholland
Matthew T. Dunn
Michael Flor
Michal Shmueli-Scheuer
Mingyue Shang
Minhua Chen
Minkai Xu
Mohammad Kachuee
Olga Golovneva
Oliver Ferschke
Pablo Duboue
Petr Lorenc
Prasanna Muthukumar
Radhika Gaonkar
Radityo Eko Prasajo
Rahul Divekar
Rajasekar Krishnamurthy
Ramy Eskander
Ryan Georgi
Rylan Conway
Sachin Agarwal
Saleh Soltan
Saloni Potdar
Sangameshwar Patil
Sanjeev Kumar
Sarah C Campbell
Satwik Kottur
Shang-Yu Su
Shashank Gupta
Shuangyong Song
Shuyan Dong
Song Feng
Sopan Khosla
Sourish Chaudhuri
Sudarshan R. Thitte
Sugyeong Eo
Tariq Alhindi
Tilman Becker
Tong Guo
Tong Wang
Trung Bui
Varun Kumar

Varun Nagaraj Rao
Vinayshekhar Bannihatti Kumar
Wei Zhu
Wenhao Yu
Weiwen Xu
Wonseok Hwang
Yannis Katsis
Yoav Katz
Young-Suk Lee
Youngja Park
Yuta Koreeda
Yuval Marton

Careers in NLP Panel

Yun Yao Li, Apple Knowledge Platform
Yang Liu, Amazon, Alexa AI
Timo Mertens, Grammarly
Thamar Solorio, University of Houston and Bloomberg LP
Luke Zettlemoyer, University of Washington and Meta

Table of Contents

<i>Scalable and Robust Self-Learning for Skill Routing in Large-Scale Conversational AI Systems</i> Mohammad Kachuee, Jinseok Nam, Sarthak Ahuja, Jin-Myung Won and Sungjin Lee	1
<i>CREATER: CTR-driven Advertising Text Generation with Controlled Pre-Training and Contrastive Fine-Tuning</i> Penghui Wei, Xuanhua Yang, ShaoGuo Liu, Liang Wang and Bo Zheng	9
<i>Augmenting Poetry Composition with Verse by Verse</i> David Uthus, Maria Voitovich and R.J. Mical	18
<i>AB/BA analysis: A framework for estimating keyword spotting recall improvement while maintaining audio privacy</i> Raphael Petegrosso, VasistaKrishna Baderdinni, Thibaud Senechal and Benjamin Bullough .	27
<i>Temporal Generalization for Spoken Language Understanding</i> Judith Gaspers, Anoop Kumar, Greg Ver Steeg and Aram Galstyan	37
<i>An End-to-End Dialogue Summarization System for Sales Calls</i> Abdelkadir Asi, Song Wang, Roy Eisenstadt, Dean Geckt, Yarin Kuper, Yi Mao and Royi Ronen	45
<i>Controlled Data Generation via Insertion Operations for NLU</i> Manoj Kumar, Yuval Merhav, Haidar Khan, Rahul Gupta, Anna Rumshisky and Wael Hamza	54
<i>Easy and Efficient Transformer: Scalable Inference Solution For Large NLP Model</i> Gongzheng Li, Yadong Xi, Jingzhen Ding, Duan Wang, Ziyang Luo, Rongsheng Zhang, Bai Liu, Changjie Fan, Xiaoxi Mao and Zeng Zhao	62
<i>Aspect-based Analysis of Advertising Appeals for Search Engine Advertising</i> Soichiro Murakami, Peinan Zhang, Sho Hoshino, Hidetaka Kamigaito, Hiroya Takamura and Manabu Okumura	69
<i>Self-supervised Product Title Rewrite for Product Listing Ads</i> Xue Zhao, Dayiheng Liu, Junwei Ding, Liang Yao, Mahone Yan, Huibo Wang and Wenqing Yao	79
<i>Efficient Semi-supervised Consistency Training for Natural Language Understanding</i> George Leung and Joshua Tan	86
<i>Distantly Supervised Aspect Clustering And Naming For E-Commerce Reviews</i> Prateek Sircar, Aniket Chakrabarti, Deepak Gupta and Anirban Majumdar	94
<i>Local-to-global learning for iterative training of production SLU models on new features</i> Yulia Grishina and Daniil Sorokin	103
<i>CULG: Commercial Universal Language Generation</i> Haonan Li, Yameng Huang, Yeyun Gong, Jian Jiao, Ruofei Zhang, Timothy Baldwin and Nan Duan	112
<i>Constraining word alignments with posterior regularization for label transfer</i> Thomas Gueudre and Kevin Martin Jose	121

<i>Explaining the Effectiveness of Multi-Task Learning for Efficient Knowledge Extraction from Spine MRI Reports</i>	
Arijit Sehanobish, McCullen Sandora, Nabila Abraham, Jayashri Pawar, Danielle Torres, Anasuya Das, Murray Becker, Richard Herzog, Benjamin Odry and Ron Vianu	130
<i>FPI: Failure Point Isolation in Large-scale Conversational Assistants</i>	
Rinat Khaziev, Usman Shahid, Tobias Röding, Rakesh Chada, Emir Kapanci and Pradeep Natarajan	141
<i>Asynchronous Convergence in Multi-Task Learning via Knowledge Distillation from Converged Tasks</i>	
Weiyi Lu, Sunny Rajagopalan, Priyanka Nigam, Jaspreet Singh, Xiaodi Sun, Yi Xu, Belinda Zeng and Trishul Chilimbi	149
<i>Augmenting Training Data for Massive Semantic Matching Models in Low-Traffic E-commerce Stores</i>	
Ashutosh Joshi, Shankar Vishwanath, Choon Hui Teo, Vaclav Petricek, Vishy Vishwanathan, Rahul Bhagat and Jonathan May	160
<i>Retrieval Based Response Letter Generation For a Customer Care Setting</i>	
Biplob Biswas, Renhao Cui and Rajiv Ramnath	168
<i>Medical Coding with Biomedical Transformer Ensembles and Zero/Few-shot Learning</i>	
Angelo Ziletti, Alan Akbik, Christoph Berns, Thomas Herold, Marion Legler and Martina Viell	176
<i>Knowledge extraction from aeronautical messages (NOTAMs) with self-supervised language models for aircraft pilots</i>	
Alexandre Arnold, Fares Ernez, Catherine Kobus and Marion-Cécile Martin	188
<i>Intent Discovery for Enterprise Virtual Assistants: Applications of Utterance Embedding and Clustering to Intent Mining</i>	
Minhua Chen, Badrinath Jayakumar, Michael Johnston, S. Eman Mahmoodi and Daniel Pressel	197
<i>ReFinED: An Efficient Zero-shot-capable Approach to End-to-End Entity Linking</i>	
Tom Ayoola, Shubhi Tyagi, Joseph Fisher, Christos Christodoulopoulos and Andrea Pierleoni	209
<i>Lightweight Transformers for Conversational AI</i>	
Daniel Pressel, Wenshuo Liu, Michael Johnston and Minhua Chen	221
<i>NER-MQMRC: Formulating Named Entity Recognition as Multi Question Machine Reading Comprehension</i>	
Anubhav Shrimal, Avi Jain, Kartik Mehta and Promod Yenigalla	230
<i>What Do Users Care About? Detecting Actionable Insights from User Feedback</i>	
Kasturi Bhattacharjee, Rashmi Gangadharaiyah, Kathleen McKeown and Dan Roth	239
<i>CTM - A Model for Large-Scale Multi-View Tweet Topic Classification</i>	
Vivek Kulkarni, Kenny Leung and Aria Haghighi	247
<i>Developing a Production System for Purpose of Call Detection in Business Phone Conversations</i>	
Elena Khasanova, Pooja Hiranandani, Shayna Gardiner, Cheng Chen, Simon Corston-Oliver and Xue-Yong Fu	259
<i>Adversarial Text Normalization</i>	
Joanna Bitton, Maya Pavlova and Ivan Evtimov	268

<i>Constraint-based Multi-hop Question Answering with Knowledge Graph</i> Sayantan Mitra, Roshni Ramnani and Shubhashis Sengupta	280
<i>Fast Bilingual Grapheme-To-Phoneme Conversion</i> Hwa-Yeon Kim, Jong-Hwan Kim and Jae-Min Kim	289
<i>Knowledge Extraction From Texts Based on Wikidata</i> Anastasia Shimorina, Johannes Heinecke and Frédéric Herledan	297
<i>AIT-QA: Question Answering Dataset over Complex Tables in the Airline Industry</i> Yannis Katsis, Saneem Ahmed Chemmengath, Vishwajeet Kumar, Samarth Bharadwaj, Mustafa Canim, Michael Glass, Alfio Gliozzo, Feifei Pan, Jaydeep Sen, Karthik Sankaranarayanan and Soumen Chakrabarti	305
<i>Parameter-efficient Continual Learning Framework in Industrial Real-time Text Classification System</i> Tao Zhu, Zhe Zhao, Weijie Liu, Jiachi Liu, Yiren Chen, Weiyan Mao, Haoyan Liu, Kunbo Ding, Yudong Li and Xuefeng Yang	315
<i>Self-Aware Feedback-Based Self-Learning in Large-Scale Conversational AI</i> Pragaash Ponnusamy, Clint Solomon Mathialagan, Gustavo Aguilar, Chengyuan Ma and Chenlei Guo	324
<i>Fast and Light-Weight Answer Text Retrieval in Dialogue Systems</i> Hui Wan, Siva Sankalp Patel, J William Murdock, Saloni Potdar and Sachindra Joshi	334
<i>BLINK with Elasticsearch for Efficient Entity Linking in Business Conversations</i> Md Tahmid Rahman Laskar, Cheng Chen, Aliaksandr Martsinovich, Jonathan Johnston, Xue-Yong Fu, Shashi Bhushan TN and Simon Corston-Oliver	344
<i>Q2R: A Query-to-Resolution System for Natural-Language Queries</i> Shiau Hong Lim and Laura Wynter	353
<i>Identifying Corporate Credit Risk Sentiments from Financial News</i> Noujoud Ahbali, Xinyuan Liu, Albert Aristotle Nanda, Jamie Stark, Ashit Talukder and Rupinder Paul Khandpur	362

Program

Monday, July 11, 2022

13:15 - 14:15 *Careers in NLP Panel*

Tuesday, July 12, 2022

08:00 - 09:00 *Oral session 1*

CREATER: CTR-driven Advertising Text Generation with Controlled Pre-Training and Contrastive Fine-Tuning

Penghui Wei, Xuanhua Yang, ShaoGuo Liu, Liang Wang and Bo Zheng

Augmenting Poetry Composition with Verse by Verse

David Uthus, Maria Voitovich and R.J. Mical

FPI: Failure Point Isolation in Large-scale Conversational Assistants

Rinat Khaziev, Usman Shahid, Tobias Röding, Rakesh Chada, Emir Kapanci and Pradeep Natarajan

ReFinED: An Efficient Zero-shot-capable Approach to End-to-End Entity Linking

Tom Ayoola, Shubhi Tyagi, Joseph Fisher, Christos Christodoulopoulos and Andrea Pierleoni

15:45 - 14:15 *Oral session 2*

Self-supervised Product Title Rewrite for Product Listing Ads

Xue Zhao, Dayiheng Liu, Junwei Ding, Liang Yao, Mahone Yan, Huibo Wang and Wenqing Yao

Local-to-global learning for iterative training of production SLU models on new features

Yulia Grishina and Daniil Sorokin

Medical Coding with Biomedical Transformer Ensembles and Zero/Few-shot Learning

Angelo Ziletti, Alan Akbik, Christoph Berns, Thomas Herold, Marion Legler and Martina Viell

CTM - A Model for Large-Scale Multi-View Tweet Topic Classification

Vivek Kulkarni, Kenny Leung and Aria Haghighi

Self-Aware Feedback-Based Self-Learning in Large-Scale Conversational AI

Pragaash Ponnusamy, Clint Solomon Mathialagan, Gustavo Aguilar, Chengyuan Ma and Chenlei Guo

Aspect-based Analysis of Advertising Appeals for Search Engine Advertising

Soichiro Murakami, Peinan Zhang, Sho Hoshino, Hidetaka Kamigaito, Hiroya Takamura and Manabu Okumura

Tuesday, July 12, 2022 (continued)

17:45 - 16:15 *Poster session*

Scalable and Robust Self-Learning for Skill Routing in Large-Scale Conversational AI Systems

Mohammad Kachuee, Jinseok Nam, Sarthak Ahuja, Jin-Myung Won and Sungjin Lee

AB/BA analysis: A framework for estimating keyword spotting recall improvement while maintaining audio privacy

Raphael Petegrosso, Vasista Krishna Baderdinni, Thibaud Senechal and Benjamin Bullough

Temporal Generalization for Spoken Language Understanding

Judith Gaspers, Anoop Kumar, Greg Ver Steeg and Aram Galstyan

An End-to-End Dialogue Summarization System for Sales Calls

Abdelkadir Asi, Song Wang, Roy Eisenstadt, Dean Geckt, Yarin Kuper, Yi Mao and Royi Ronen

Controlled Data Generation via Insertion Operations for NLU

Manoj Kumar, Yuval Merhav, Haidar Khan, Rahul Gupta, Anna Rumshisky and Wael Hamza

Easy and Efficient Transformer: Scalable Inference Solution For Large NLP Model

Gongzheng Li, Yadong Xi, Jingzhen Ding, Duan Wang, Ziyang Luo, Rongsheng Zhang, Bai Liu, Changjie Fan, Xiaoxi Mao and Zeng Zhao

Efficient Semi-supervised Consistency Training for Natural Language Understanding

George Leung and Joshua Tan

Distantly Supervised Aspect Clustering And Naming For E-Commerce Reviews

Prateek Sircar, Aniket Chakrabarti, Deepak Gupta and Anirban Majumdar

CULG: Commercial Universal Language Generation

Haonan Li, Yameng Huang, Yeyun Gong, Jian Jiao, Ruofei Zhang, Timothy Baldwin and Nan Duan

Constraining word alignments with posterior regularization for label transfer

Thomas Gueudre and Kevin Martin Jose

Tuesday, July 12, 2022 (continued)

Explaining the Effectiveness of Multi-Task Learning for Efficient Knowledge Extraction from Spine MRI Reports

Arijit Sehanobish, McCullen Sandora, Nabila Abraham, Jayashri Pawar, Danielle Torres, Anasuya Das, Murray Becker, Richard Herzog, Benjamin Odry and Ron Vianu

Asynchronous Convergence in Multi-Task Learning via Knowledge Distillation from Converged Tasks

Weiyi Lu, Sunny Rajagopalan, Priyanka Nigam, Jaspreet Singh, Xiaodi Sun, Yi Xu, Belinda Zeng and Trishul Chilimbi

Augmenting Training Data for Massive Semantic Matching Models in Low-Traffic E-commerce Stores

Ashutosh Joshi, Shankar Vishwanath, Choon Hui Teo, Vaclav Petricek, Vishy Vishwanathan, Rahul Bhagat and Jonathan May

Retrieval Based Response Letter Generation For a Customer Care Setting

Biplob Biswas, Renhao Cui and Rajiv Ramnath

Knowledge extraction from aeronautical messages (NOTAMs) with self-supervised language models for aircraft pilots

Alexandre Arnold, Fares Ernez, Catherine Kobus and Marion-Cécile Martin

Intent Discovery for Enterprise Virtual Assistants: Applications of Utterance Embedding and Clustering to Intent Mining

Minhua Chen, Badrinath Jayakumar, Michael Johnston, S. Eman Mahmoodi and Daniel Pressel

Lightweight Transformers for Conversational AI

Daniel Pressel, Wenshuo Liu, Michael Johnston and Minhua Chen

NER-MQMRC: Formulating Named Entity Recognition as Multi Question Machine Reading Comprehension

Anubhav Shrimal, Avi Jain, Kartik Mehta and Promod Yenigalla

What Do Users Care About? Detecting Actionable Insights from User Feedback

Kasturi Bhattacharjee, Rashmi Gangadharaiah, Kathleen McKeown and Dan Roth

Developing a Production System for Purpose of Call Detection in Business Phone Conversations

Elena Khasanova, Pooja Hiranandani, Shayna Gardiner, Cheng Chen, Simon Corston-Oliver and Xue-Yong Fu

Adversarial Text Normalization

Joanna Bitton, Maya Pavlova and Ivan Evtimov

Tuesday, July 12, 2022 (continued)

Constraint-based Multi-hop Question Answering with Knowledge Graph

Sayantana Mitra, Roshni Ramnani and Shubhashis Sengupta

Fast Bilingual Grapheme-To-Phoneme Conversion

Hwa-Yeon Kim, Jong-Hwan Kim and Jae-Min Kim

Knowledge Extraction From Texts Based on Wikidata

Anastasia Shimorina, Johannes Heinecke and Frédéric Herledan

AIT-QA: Question Answering Dataset over Complex Tables in the Airline Industry

Yannis Katsis, Saneem Ahmed Chemmengath, Vishwajeet Kumar, Samarth Bhadraraj, Mustafa Canim, Michael Glass, Alfio Gliozzo, Feifei Pan, Jaydeep Sen, Karthik Sankaranarayanan and Soumen Chakrabarti

Parameter-efficient Continual Learning Framework in Industrial Real-time Text Classification System

Tao Zhu, Zhe Zhao, Weijie Liu, Jiachi Liu, Yiren Chen, Weiquan Mao, Haoyan Liu, Kunbo Ding, Yudong Li and Xuefeng Yang

Fast and Light-Weight Answer Text Retrieval in Dialogue Systems

Hui Wan, Siva Sankalp Patel, J William Murdock, Saloni Potdar and Sachindra Joshi

BLINK with Elasticsearch for Efficient Entity Linking in Business Conversations

Md Tahmid Rahman Laskar, Cheng Chen, Aliaksandr Martsinovich, Jonathan Johnston, Xue-Yong Fu, Shashi Bhushan TN and Simon Corston-Oliver

Q2R: A Query-to-Resolution System for Natural-Language Queries

Shiau Hong Lim and Laura Wynter

Identifying Corporate Credit Risk Sentiments from Financial News

Noujoud Ahbali, Xinyuan Liu, Albert Aristotle Nanda, Jamie Stark, Ashit Talukder and Rupinder Paul Khandpur