Hypertext 2020 Tutorial: Knowledge-infused Deep Learning

Manas Gaur, Ugur Kursuncu, Amit Sheth, Ruwan Wickramarachchi, Shweta Yadav
Artificial Intelligence Institute, University of South Carolina, SC, USA
mgaur@email.sc.edu, kursuncu@mailbox.sc.edu, amit@sc.edu, ruwan@email.sc.edu, shweta@knoesis.org

ABSTRACT

Deep Learning has shown remarkable success during the last decade for essential tasks in computer vision and natural language processing. Yet, challenges remain in the development and deployment of artificial intelligence (AI) models in real-world cases, such as dependence on extensive data and trust, explainability, traceability, and interactivity. These challenges are amplified in high-risk fields, including healthcare, cyber threats, crisis response, autonomous driving, and future manufacturing. On the other hand, symbolic computing with knowledge graphs has shown significant growth in specific tasks with reliable performance. This tutorial (a) discusses the novel paradigm of knowledge-infused deep learning to synthesize neural computing with symbolic computing (b) describes different forms of knowledge and infusion methods in deep learning, and (c) discusses application-specific evaluation methods to assure explainability and reasoning using benchmark datasets and knowledge-resources. The resulting paradigm of "knowledge-infused learning" combines knowledge from both domain expertise and physical models. A wide variety of techniques involving shallow, semi-deep, and deep infusion will be discussed along with the corresponding intuitions, limitations, use cases, and applications. More details can be found http://kidl2020.aiisc.ai/.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence**; *Scene understanding*; • **Human-centered computing** → *Collaborative and social computing*; • **Applied computing** → *Health informatics.*

KEYWORDS

Knowledge-infused Learning, Knowledge Graphs, Deep Learning, Neuro-symbolic Computing, Public Health, Disaster Resilience, Cyber-Social Threats, Autonomous Driving

ACM Reference Format:
Manas Gaur, Ugur Kursuncu, Amit Sheth, Ruwan Wickramarachchi, Shweta Yadav. 2020. Hypertext 2020 Tutorial: Knowledge-infused Deep Learning. In *Proceedings of the 31st ACM Conference on Hypertext and Social Media (HT '20), July 13–15, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3372923.3404862

1 TUTORIAL INFORMATION

Recent advances in statistical and data-driven deep learning demonstrate significant success in natural language understanding (NLU) without using prior knowledge, especially in generic unstructured domains, where data is abundant (e.g., BERT, GPT-3). On the other hand, in problems concerning text mining that are dynamic and impact society at large, existing data-dependent, state-of-the-art deep learning methods remain vulnerable to veracity considerations, especially when high-volume masks small, emergent signals [10]. Statistical natural language processing (NLP) techniques have shown poor performance in capturing: (1) Human well being online especially in evolving events (e.g., mental health communications on Reddit [2, 3]), (2) Culture and context-specific discussion on the web (e.g., sarcasm and humor detection, extremism on social media [6]), (3) Social Network Analysis during pandemic or disaster scenarios [7], and (4) Explainable methods of learning that drive technological innovations and inventions for social good [8, 9, 11]. In such multimodal, social content, leveraging the semantic-web concept of knowledge graphs is a promising approach to enhancing deep learning and NLP [5]. Examples of the desiderata include compensating for data limitations, improving inductive bias, generating explainable outcomes, and enabling trust. These are particularly useful for data-limited but evolving problems in domains such as mental healthcare, online social threats, and pandemics. Despite the general agreement that structured prior knowledge and tacit knowledge resulting from deep learning should be combined, there has been limited progress. Recent debates on neuro-symbolic AI[1], inclusion of innate priors in deep learning[2], and AI fireside chat[3] have discussed incorporating knowledge as a promising approach to improve explainability, interpretability, and trust in AI systems.

2 DESCRIPTION OF THE TUTORIAL

This tutorial includes three broad modules:

Knowledge-infused Deep Learning (KiDL): First, we explain the fundamental role of a relation-preserving knowledge representation in learning from social media, scientific articles, blog posts and scene graphs. Second, we describe Neural, Symbolic, and Semantic computing as the foundational piece for knowledge-infusion. Lastly, we explain the KiDL paradigm that models prior knowledge and cognitive theories, which constitute our understanding of context and relations in the social good domain. We will present a use case driven exposition of the modern aspect of hypertext using Knowledge Graphs (KGs) [4]. Key use cases include social good applications (Mental Health [2], Radicalization [6]) and multimodal aspects of social media (e.g. scene understanding from images, video and text (hypermedia/hypertext) [9]) often found in documentation of critical events. Subsequently, we describe the utility of KiDL for interpretable and explainable, multi-modal learning for text, video, images, and graph data on the web.[4] [5]

[1] https://www.zdnet.com/article/devils-in-the-details-in-bengio-marcus-ai-debate/
[2] https://bit.ly/Manning_LeCun
[3] https://vimeo.com/390814190
[4] http://bit.ly/iSEMANTICS2020, http://bit.ly/k-CCKS
[5] https://bit.ly/KnowledgeGraphs, https://bit.ly/OK_talk

HT '20, July 13–15, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-7098-1/20/07.
https://doi.org/10.1145/3372923.3404862

All about Knowledge Graphs: We elaborate on different forms of external knowledge and explain: (1) What is a Knowledge Graph, (2) Knowledge Graphs and Domain-Specific Search Problems, (3) Knowledge Graph Construction and Evolution, (4) Knowledge Graph Completion and Sub-graph Creation. Further, we introduce KG-driven unsupervised, semi-supervised, and supervised methods of representation learning over unstructured multimodal content.
Forms of Knowledge-Infusion and its Applications: We describe three approaches to knowledge-infusion [10]:
(a) ***Shallow Infusion:*** Both the external knowledge and the method of knowledge infusion is shallow, utilizing syntactic and lexical knowledge in the form of word embedding models. **(b)** ***Semi-Deep Infusion:*** External knowledge is involved through attention mechanisms or learnable knowledge constraints acting as a sentinel to guide model learning. **(c)** ***Deep Infusion:*** Employs a stratified representation of knowledge representing different levels of abstractions in different layers of a deep learning model, to transfer knowledge that aligns with the corresponding layer in the layered learning process. **(d)** ***Domain-specific Applications:*** Considering problems in the real-world applications (e.g., epidemic (e.g., COVID-19 and Ebola), disaster (e.g., Hurricanes) scenarios [1]), Mental Health (e.g. Depression), we explain the utility of KiDL, highlighting key problems:(i)*Context understanding:* Understanding current context with respect to observable objects and events, given learned experience (from past behavior) and external knowledge. (ii)*Abstraction:* A technique with utilizes domain-specific KG to map and associates raw data to action-related information for high-order stakeholders (e.g., Psychiatric clinicians, Emergency responders).

3 TARGET AUDIENCE AND PREREQUISITES

This tutorial will bring researchers in academic, industry, humanitarian organizations, and healthcare practitioners at the confluence of knowledge representation, reasoning, semantic linking, NLP, and deep learning. There are no prerequisites for attending the tutorial. We will cover basics and advanced techniques with sufficient examples. Newcomers in the area will learn the basic principles of data science and fundamentals of the semantic web. Expert attendees will appreciate KiDL options as promising, reliable, and practical approaches to overcoming familiar technical obstacles in social good domains.

4 PRESENTERS' BIOGRAPHIES

Manas Gaur[6] is a Ph.D. Student in AIISC. He has been a Data Science and AI for Social Good Fellow with the University of Chicago and Dataminr Inc. His interdisciplinary research funded by NIH and NSF operationalizes the use of KGs, NLU, and AI in the domain of Mental Health Informatics[7]. His work has appeared in premier AI and Data Science conferences.
Ugur Kursuncu is a Postdoctoral Research Associate at AIISC. He received his Ph.D. from The University of Georgia with awards for excellence in Teaching and Research in 2015 and 2016. His research has focused on context-aware and knowledge-infused learning systems spanning the areas of Cyber Social Threats and Health informatics. His research has been published in top-tier conferences, journals and books, such as CSCW, Springer-Nature.

[6]https://manasgaur.github.io
[7]https://bit.ly/PyDataSal

Ruwan Wickramarachchi is a Ph.D. student at AIISC. His primary research interest is in neuro-symbolic AI for context understanding with applications in Autonomous Driving and Healthcare. Prior to joining the Ph.D. program, he was a Senior Software Engineer at the Machine Learning Research Group of LSEG Technology.
Shweta Yadav is a Postdoctoral Research Associate at AIISC. She completed her Ph.D. from IIT Patna. Her research interests span biomedical and healthcare informatics. Her work has appeared in top-tier conferences and journals including ACL, WWW, Knowledge based Systems, Soft Computing, and ACM TOMM.
Amit Sheth [8] is an Educator, Researcher, and Entrepreneur. He is the founding director of the university-wide Artificial Intelligence Institute at the University of South Carolina (AIISC). Previously, he was the LexisNexis Ohio Eminent Scholar and the executive director of Ohio Center of Excellence in Knowledge-enabled Computing (Kno.e.sis). He is a Fellow of IEEE, AAAI, and AAAS. He has organized 75+ international events (general/program chair, organization committee chair), given 65+ keynotes and many well-attended tutorials and is among the well-cited computer scientists. He has founded three companies by licensing his university research outcomes, including the first Semantic Web company in 1999 that pioneered technology similar to what is found today in Google Semantic Search and Knowledge Graph. Several commercial products and deployed systems have resulted from his research.

ACKNOWLEDGEMENT

The authors would like to thank Dr. Krishnaprasad Thirunarayan and Dr. Valerie L. Shalin for insightful suggestions in structuring the Tutorial. We acknowledge support from NSF awards CNS 1513721, EAR 1520870, and NIH award 1R01MH105384-01A1.

REFERENCES

[1] Chidubem Arachie, Manas Gaur, Sam Anzaroot, William Groves, Ke Zhang, and Alejandro Jaimes. 2019. Unsupervised Detection of Sub-events in Large Scale Disasters. *AAAI* (2019).
[2] Manas Gaur, Amanuel Alambo, Joy Prakash Sain, Ugur Kursuncu, Krishnaprasad Thirunarayan, Ramakanth Kavuluru, Amit Sheth, Randy Welton, and Jyotishman Pathak. 2019. Knowledge-aware assessment of severity of suicide risk for early intervention. In *ACM WWW*.
[3] Manas Gaur, Ugur Kursuncu, Amanuel Alambo, Amit Sheth, Raminta Daniulaityte, Krishnaprasad Thirunarayan, and Jyotishman Pathak. 2018. " Let Me Tell You About Your Mental Health!" Contextualized Classification of Reddit Posts to DSM-5 for Web-based Intervention. In *ACM CIKM*.
[4] Amelie Gyrard, Manas Gaur, Saeedeh Shekarpour, Krishnaprasad Thirunarayan, and Amit Sheth. 2018. Personalized Health Knowledge Graph. In *ISWC*.
[5] Ramnath Kumar, Shweta Yadav, Raminta Daniulaityte, Francois Lamy, Krishnaprasad Thirunarayan, Usha Lokala, and Amit Sheth. 2020. eDarkFind: Unsupervised Multi-view Learning for Sybil Account Detection. In *ACM WWW*.
[6] Ugur Kursuncu, Manas Gaur, Carlos Castillo, Amanuel Alambo, Krishnaprasad Thirunarayan, Valerie Shalin, Dilshod Achilov, I Budak Arpinar, and Amit Sheth. 2019. Modeling Islamist Extremist Communications on Social Media using Contextual Dimensions: Religion, Ideology, and Hate. *ACM CSCW* (2019).
[7] Ugur Kursuncu, Manas Gaur, Usha Lokala, Krishnaprasad Thirunarayan, Amit Sheth, and I Budak Arpinar. 2019. Predictive analysis on Twitter: Techniques and applications. In *Springer Nature*.
[8] Ugur Kursuncu, Manas Gaur, and Amit Sheth. 2020. Knowledge Infused Learning (K-IL): Towards Deep Incorporation of Knowledge in Deep Learning. *AAAI-MAKE Symposium* (2020).
[9] Alessandro Oltramari, Jonathan Francis, Cory Henson, Kaixin Ma, and Ruwan Wickramarachchi. 2020. Neuro-symbolic Architectures for Context Understanding. *arXiv preprint arXiv:2003.04707* (2020).
[10] Amit Sheth, Swati Padhee, and Amelie Gyrard. 2019. Knowledge Graphs and Knowledge Networks: The Story in Brief. *IEEE Internet Computing* (2019).
[11] Ruwan Wickramarachchi, Cory Henson, and Amit Sheth. 2020. An evaluation of knowledge graph embeddings for autonomous driving data: Experience and practice. *AAAI-MAKE Symposium* (2020).

[8]http://aiisc.ai/amit