Neural Language Models for the Multilingual, Transcultural, and Multimodal Semantic Web

Dagmar Gromann [a]
[a] *Center for Translation Studies, University of Vienna, Vienna, Austria*
E-mail: dagmar.gromann@univie.ac.at

Editors: Pascal Hitzler, Kansas State University, Manhattan, KS, USA; Krzysztof Janowicz, University of California, Santa Barbara, USA
Solicited reviews: Claudia d'Amato, Università degli Studi di Bari, Italy; Philipp Cimiano, Bielefeld University, Germany; Freddy Lecue, CortAIx Thales, Montreal, Canada and INRIA, France

Abstract. A vision of a truly multilingual Semantic Web has found strong support with the Linguistic Linked Open Data community. Standards, such as OntoLex-Lemon, highlight the importance of explicit linguistic modeling in relation to ontologies and knowledge graphs. Nevertheless, there is room for improvement in terms of automation, usability, and interoperability. Neural Language Models have achieved several breakthroughs and successes considerably beyond Natural Language Processing (NLP) tasks and recently also in terms of multimodal representations. Several paths naturally open up to port these successes to the Semantic Web, from automatically translating linguistic information associated with structured knowledge resources to multimodal question-answering with machine translation. Language is also an important vehicle for culture, an aspect that deserves considerably more attention. Building on existing approaches, this article envisions joint forces between Neural Language Models and Semantic Web technologies for multilingual, transcultural, and multimodal information access and presents open challenges and opportunities in this direction.

Keywords: Neural Networks, Multilingual Representations, Cross-Linguistic Modeling

1. Introduction

One central endeavor of the Semantic Web (SW) [1] is intelligent access to heterogeneous and distributed sources of knowledge. However, limiting this access to natural languages predominant in the world inevitably creates biases and hegemonies. Supporters of a multilingual SW can account for several successes to overcome the language barrier, from multilingual structured knowledge resources, such as BabelNet [2] and Framester [3], to multilingual methods and applications (cf. e.g. [4]). Nevertheless, approaches that further improve the level of automation, usability, and interoperability are required.

This article proposes a vision that is based on Neural Language Models (NLMs) to foster a multilingual, transcultural, and multimodal Semantic Web. Its contribution is a detailed exploration of this vision based on existing approaches and an outline of currently valid challenges and envisioned opportunities, which provides a solid starting point for the Semantic Web and NLP community to initiate and/or advance such interdisciplinary research.

A language model is designed to assign probabilities to an input sequence, i.e., learn a joint probability function of sequences of signs. Based on this idea, powerful Natural Language Processing (NLP) applications from machine translation (e.g. [5]) and natural language generation (e.g. [6]) to textual entailment (e.g. [7, 8]) have been proposed.

NLMs learn implicit semantic representations of sequences on their hidden layer(s), resulting in a dense

1570-0844/0-1900/$35.00

real-valued vector for each word, phrase, sentence, document, or knowledge base triple, which turned out to be a powerful representation. Such embeddings have been applied to a large variety of traditional SW tasks, from link prediction to ontology alignment [9]. Recent NLMs have provided a strong backbone to many Artificial Intelligence (AI) applications that go beyond traditional NLP tasks, see for instance [10] for a wide range of tasks, including a new best performance on the Winograd Schema Challenge. Resolving pronouns in such schemas requires world knowledge, such as spatio-temporal relations and mental states.

Regarding automation and usability of SW technologies, NLMs have successfully been applied to translating from natural language to natural language but also to ontology representation [11] and structured query [12] languages. Automatically translating natural language questions to queries can improve the usability of SW query interfaces. However, the usage of NLMs goes considerably beyond translating languages, structured or unstructured. Neural Machine Translation (NMT) based on NLMs has even been applied to noise-tolerant RDFS reasoning [13].

Language enables communication and at the same time serves as a vehicle for cultural and social identity. This function of natural language should find consideration in approaches to the multilingual Semantic Web by building on decades of research on cross-cultural and transcultural communication (e.g. [14]). NLMs potentially provide interesting methods to port information learned for one language and culture to another in form of domain adaptation and transfer learning [15, 16]. Nevertheless, a more thorough basis is required to capture cultural aspects, such as cognitive principles guiding our communication.

Communication in natural language is by no means confined to textual boundaries and can be signed, spoken, or written. This calls for multimodal representations of language in relation to SW technologies, which finds strong support in state-of-the-art language modeling. Recent advances of NLMs provide powerful approaches that allow flexible alignments between text and video [17] and translate directly from speech to speech without a need for textual transcriptions [18].

In short, this vision goes beyond plurality of language and envisions multilingual, transcultural, and multimodal information access backed by NLMs and the Semantic Web. As preliminaries, this article first briefly defines language models and the Multilingual Semantic Web. The sections Multilingual, Transcultural, and Multimodal detail existing joint approaches on different SW tasks, each of which is followed by a description of the challenges and opportunities for joining language modeling and SW approaches. Neither of these can be fully accounted for in this article, but are detailed to the point of grounding envisioned future research directions.

2. Language Model: A Brief Definition

Language modeling has been key to the success of NLP applications and tasks, such as machine translation, speech recognition, question-answering, spelling correction and many more. A language model (LM) predicts a probability of a previously unseen sequence of words based on a preceding learned probability distribution over the whole vocabulary of the training corpus. In general, the joint probability of a sequence is decomposed as the product of conditional probabilities of co-occurring words, two in the case of bigrams, three in the case of trigrams, and so on. Smoothing is a procedure to avoid zero probabilities due to unknown words (cf. e.g. [19] for more details).

Neural language models (NLMs) can learn distributed representations without smoothing and generalize well across contexts. Training tasks generally consist in predicting a center word of a sequence given its context words (skip-gram) or predicting a context given a center word (CBOW), which where popularly introduced by [20]. Both tasks train word embeddings as values of their hidden layers and the base methods have been extended to train vector representations of knowledge graph triples (e.g. [21]).

A common architectural style is that of encoder-decoder, either of which can be independent models. One recent best-performer in machine translation, the Transformer model [5], has soon been propagated to many NLP tasks, from multi-task approaches [10] to text-video combinations [17]. A Transformer combines multi-head attention, that is, a mechanism to single out central words in sequences for given queries, and feedfoward layers in an encoder-decoder architecture. Frequently, this architecture is combined with Byte Pair Encodings (PBE), a form of data compression [22] that iteratively merges most frequent characters or character sequences with a single, unused byte. Since it evaluates words on character-level it strongly mitigates the problem of unknown words.

3. Multilingual Semantic Web: A Long-Standing Endeavor

For several decades multiple research endeavors [4, 23] have made it their mission to provide a truly multilingual SW. To this end, algorithms and systems are required that help overcome linguistic and national boundaries, to grant information access to users of different cultures and languages. Limiting such access to languages spoken by majorities inevitably creates a bias. The SW, with its language-independent representation of knowledge, provides an excellent anchor point for multilingual, transcultural, and multimodal information access.

As a first step towards a multilingual SW, several mediation mechanisms to translate between abstract conceptual layers and lexical manifestations, which frequently are different across languages and cultures, have been proposed. In fact, concepts might exist in one language but not in another, so called lexical gaps, such as the German "Schadenfreude" (joy when something bad happens to someone else) that has been readily adopted in English due to a lack of an equivalence.

Knoweldge representation needs to be able to accommodate such differences. First, the OntoLex-Lemon model that provides an ontology-lexicon interface has found broad uptake by the community and has recently been published as a W3C report [24]. Second, similar models have been proposed to interchange domain-specific terminological information grounded in ontological resources [25]. Combined representations of linguistic, terminological, and ontological knoweldge have been modeled [26]. As a final example, the NLP Interchange Format (NIF) [27] based on Linked Data principles serves to improve the interoperability of NLP tools.

Rich combinations of structured knowledge and linguistic information can be applied to a variety of tasks, such as ontology-based information extraction [28], completing and correcting natural language information [29], translating from knowledge resource to natural language and/or vice versa [30], and ontology learning from text [31].

Over the past few decades the Linked Open Data (LOD) cloud and resources published in the Resource Description Framework (RDF) and Web Ontology Language (OWL) have experiences a tremendous growth, however, predominantly in English with several notable exceptions, such as BabelNet [2] and WikiData [32]. To foster this endeavor, automated means, such as NLMs, can improve and fasten approaches. While this section detailed initiatives by the multilingual SW community, Section 4 focuses on the utilization of NLMs towards a multilingual SW.

4. Multilingual

Within the context of this article, multilingual refers to this aspect of the presented vision, that is, how NLMs can contribute to multilingual SW tasks and technologies.

Machine Translating the SW: One most immediate application scenario of NLMs is the translation of natural language contents of the SW. Ontology labels, especially in domain ontologies, provide a rich terminological layer, but are still predominantly in English. To overcome this problem, Neural Machine Translation (NMT) and Statistical Machine Translation (SMT) models have been applied to translate ontology labels [15]. As an interesting side-aspect, the impact of injection approaches of domain-specific terminological knowledge to NMT and SMT on the translation quality are evaluated. The most promising knowledge augmentation method is domain adaptation of a trained model with terminological expressions, which has been utilized before to translate ontology labels [33] and fine-tune machine translation [34].

Challenges and Opportunities: As concluded in a recent survey on machine translation and SW technologies [35], this combination is still in its infancy. SW technologies have the potential to aid NMT models for disambiguating senses and targeting NMT to particular domains of discourse, which in turn can be applied to produce multilingual domain-specific ontology descriptions.

A most promising direction for such combinations lies in the injection of domain, lexical, and terminological knowledge into NMT systems. While some preliminary evaluations, such as the impact of linguistic processing results injected into a neural question-answering system, are available [36], a systematic investigation is yet to be performed. Knowledge injection or augmentation holds the potential to help bridge the neural-symbolic gap (cf. e.g. [37]) and support Explainable AI (cf. e.g. [38]). It refers to the task of actively including external (knowledge) resources in the training process of NLMs, e.g. by continuing training on a pre-trained model with such a resource or by adjusting the attention mechanism during training with external knowledge.

This topic of injecting knowledge to NLMs implicitly raises an important challenge posed to the Semantic Web community. Current semantic representations, such as RDF and OWL, while readily embraced by the multilingual Semantic Web and Linguistic Linked Open Data (LLOD) community, might require an adaptation towards the more lightweight end in order to be readily adopted by the machine learning and NLP community.

Finally, a fully automated and NLM-based translation of existing ontology labels to rich linguistic representations in form of ontology-lexicon or ontology-terminology models would be a very interesting application of NMT, which brings us to the next topic of learning structured languages.

Machine Translating to Structured Languages: NMT can not only translate natural languages. Early neural approaches utilized joint knowledge base and language embeddings to extract relations [39]. [40] utilize multilingual natural-language patterns to learn RDF predicates, which are refined by way of a feedforward neural network. Recent approaches treat the entire problem of structure learning as a machine translation task and utilized an NMT system to learn a specific subset of Description Logic formulas from definitions [11]. For instance, from the input *A bee is an animal that produces honey* the model produces $bee \sqsubseteq animal \sqcap \exists produces.honey$.

A long-standing endeavor in Semantic Web research has been the automated translation of natural language questions to SPARQL queries. Since SPARQL requires syntactic and semantic expertise, a translation from natural language could considerably boost its uptake and make Semantic Web resources broadly available without any prior knowledge of representation and query languages. A broad test of existing NMT models to the task of translating from natural language to SPARQL has been proposed [12].

Challenges and Opportunities: One substantial future application scenario of NLMs is that of learning structured knowledge resources. Ontology learning experiments with NLMs focus on a subset of logical expressions and on English only. However, automating the process of extracting structured knowledge from natural languages, holds the promise of obtaining conceptualizations specific to the language and culture.

This joining of both technologies is not only attractive for its promised speed of creating resources, but also for the ability to adapt trained models to new domains and languages, such as zero-shot translation, the ability to translate from one language to another without ever explicitly training the language model on this particular language pair. Thus, these approaches have the potential to consider low-resource languages. As a central problem in NLP is the predominance of certain widespread language varieties in applications, this could boost the uptake of SW technologies in machine learning.

Machine Translation for Reasoning: A very recent approach is that of tailoring embeddings to accommodate RDFS reasoning in an NMT task [13]. To this end, RDF graphs are layered and encoded as adjacency matrices, where each layer layout represents a graph word. Input graph and entailments are then represented as sequences of graph words, which enables treating RDFS reasoning as a machine translation task.

Challenges and Opportunities: Deductive reasoning as a machine translation task is attractive due to its potential reasoning speed, a major challenge for reasoning engines. Encoding information as input to NLM-based reasoning engines is an open research topic. [13] suggest layered graph word embeddings as a first approach. However, there is a lot of room for experimentation and further proposals in this regard.

It would be interesting to evaluate whether embeddings learned for the purpose of reasoning in one formalism might allow for transitioning to or similarity measures of elements or statements in different representation languages. In other words, this could be a potential approach for tackling diversity of ontology languages, as proposed by [41] in form of a meta language. For instance, distributed vector representations learned for an ontology represented in RDF might be comparable to embeddings trained on an OWL ontology or on information encoded in the Unified Modeling Language (UML).

The challenge in combining machine learning and logic lies in conciliating advantages of both without aggravating their limitations. The advantages of symbolic approaches are the provision of sound and explainable reasoning, while neural approaches with NLMs have the potential of providing fast and robust learning. The difficulty in combining them, in fact, is the current low-level representation of information in neural approaches instead of symbolic representations as done in logic. An additional opportunity here is a hybrid interaction of both methodologies, as proposed for image segmentation by [42]. The ability of a fully integrated or hybrid solution to support explainability of NLMs might currently be the most trendy opportunity. A discussion specific to neural-symbolic systems can be found in [37] in this issue.

NLM-based ontology alignment: has been successfully applied to matching knowledge bases. For instance, utilizing multilingual pretrained embeddings, domain-specific industry classification standards could be aligned [9]. The task of aligning large ontologies has been subdivided into smaller, tractable tasks utilizing a lexical index, neural embeddings, and locality models [43].

A broader alignment strategy is that of bringing together a multitude of resources from the Linguistic Linked Open Data (LLOD) cloud with ontology resources in Framester [3]. Based on this resource, frame-based embeddings are trained and utilized for knowledge reconciliation purposes [44], but could also be applied to a wide range of NLP tasks.

Challenges and Opportunities: NLM-based alignment strategies could benefit from the previous tasks in form of using neural-symbolic reasoning to align large multilingual, transcultural, and multimodal ontologies. In addition, the substantial surge of knowledge graph embedding approaches could be joint with the multitude of word embedding models, building on and contributing to the tradition of modeling at the ontology-natural language interface of the SW community. Joint NLM- and SW-based alignment approaches can potentially also foster the transitioning from knowledge represented in one cultural context to another.

5. Transcultural

When it comes to culture, a multitude of prefixes is commonplace: cross-cultural, intercultural, multicultural, and transcultural. Cross-cultural refers to analytic comparative approaches of different cultures. Intercultural generally establishes a certain understanding for different cultures. Multicultural refers to a plurality of cultures even within a society. And finally transcultural refers to a social concept that denotes a joint shared culture irrespective of origin or nationality. With an ever-growing global connectivity, this last prefix best denotes what this vision entails. Rather than a mere coexisting alignment of cultural representations, a capacity to move between and within cultural and social identity is foreseen. Importance of differences in semantic modeling across cultures finds support in cross-cultural neuro-scientific findings that show differences in categorization and in processing semantic relationships across cultures [45].

Cultural Evolution: Cultural evolution is closely tied to evolutionary biology science and Darwinian evolutionary principles [46]. A set of algorithms based on evolution by natural selection, that is, variation, hereditiy, and selection, has been put forward and recently extended by fission, fusion, and cooperation in their application to cultural phenomena [47]. As a basic assumption, biological concepts for the origin of living beings can be mapped to the cultural and linguistic domain, which have then been combined in a theory of cultural language evolution [48]. An application of the SW tests this assumption in terms of ontology alignments and evolutionary alignment repair in cultural environments utilizing a multi-agent system [49].

Challenges and Opportunities: A theory and experiments for the cultural evolution of human language has been thoroughly investigated [48]. It studies, for instance, how linguistic variants are generated in a population and on which basis some variants survive and become dominant. As a social phenomenon, language features cooperative interaction patterns, such as open-ended questions. While many of these phenomena are language-specific, it is highly unlikely that vocabularies and grammars are stored as different language systems by users. Instead, a widely accepted theory is that of storing such knowledge in form of constructions, associating meaning with form. One construction stores many constraints for efficient parsing and is presumed to incorporate several different language systems [48].

Various grammars have been proposed to capture such linguistic constructions. NLMs potentially complement tested construction grammar approaches, increasing the level of automation and potentially domain coverage by means of transfer learning. In particular, NLM-based multi-agent system negotiations of meaning could foster transcultural modeling of cultural evolution. An alternative way of modeling cultural language evolution is that of formal game-theoretic semantics. Bringing NLMs and formal semantics of constructions spanning across language systems together could potentially boost a transcultural SW. In fact, a SW that facilitates cultural exchange cannot only support human users, but building on evolutionary theories can foster a dynamically evolving SW that embraces diversity rather than rigidity, which again raises the challenge of less focus on formalization and more focus on usability and robustness.

Cultural Heritage: denotes physical artifacts as much as intangible attributes of a culture or society from the past. Several SW approaches can be found, from cultural heritage modeling (e.g, [50]) to creating ontology-based lexicographic tools for the study of ancient culture to enable object-multilingual links [51]. Culture-specific knowledge graphs of cultural heritage have been proposed, such as for Italy [52]. While there is a multitude of NLM-based approaches, little overlap could be detected between NLM- and SW-based research on cultural heritage.

Challenges and Opportunities: The range of possible joint approaches of NLM and SW technologies to model cultural heritage includes all of the multilingual approaches presented above and most of the multimodal approaches presented below. For instance, based on knowledge graphs, NLMs can be utilized to analyze similarities and differences across cultural heritages as well as refine technologies to share and analyze cultural data. Neural-symbolic reasoning could be particularly powerful for such alignments.

One of the most central challenges in terms of cultural heritage is the linking and representing of cultural data in a harmonized way across individual data collections. Here symbolic ontology-based integration methods could strongly benefit from NLMs and their ability to detect similarities even with noisy data, bringing together rich semantic representations and noise-tolerant, robust learners.

On a more local level, organizations procuring cultural heritage data are frequently interested in a possibility to provide highly personalized user experiences. For instance, a virtual tour through a museum or archaeological site should be fully in the control of each user. In this direction it would be interesting to explore the joint power of NLM-based SW technologies or SW-empowered NLMs to suggest or predict interesting paths or items for each individual user.

Culture-specific Modeling: Another important transcultural SW connection is that of utilizing ontologies for culture-specific modeling. For instance, [53] explore Australian Indigenous knowledge systems utilizing SW technologies. When utilizing SW technologies for cross-cultural modeling, lexical gaps rapidly become unavoidable. Cross-language information retrieval (CLIR) tasks equally encounter this problem, and have come up with NLM-based methods to bridge such lexical gaps [54]. Embedding spaces have also been analyzed for their ability to represent culture-specific association [55] and their ability for macro-cultural modeling.

Challenges and Opportunities: Going from modeling individual culture-specific knowledge representations to a transcultural one represents the biggest challenge in this task. Domain ontologies potentially provide a language-independent anchor for transcultural knowledge modeling, joint with NLM-based cross-language information retrieval and analysis approaches. Bringing both together enables transcultural question-answering and potentially automated localization approaches. By transcultural question-answering, this article refers to the phenomenon of not only foreseeing multilingual answers to queries posed to a language-agnostic knowledge base, but the ability to verbalize responses from such a knowledge base in a variety of cultural spheres and states of cultural language evolution. Localization differs from translation in that it focuses on a regional adaptation of contents more than their transformation to a different language or linguistic variation. As such it takes cultural preferences into account.

One powerful aspect that could potentially boost transcultural modeling is a solid cognitive basis, such as multilingual knowledge extraction and modeling related to embodied cognition [56, 57]. Such a cognitive framework can be utilized to analyze and model cultural differences on a cognitive-conceptual basis rather than a primarily data-driven approach.

One important aspect of culture are regional linguistic differences. Considering dialects and linguistic variations in machine translation and semantic speech technologies is still an open field of research. Rich variation-aware linguistic representation models in connection to ontologies, that is, extensions of ontology-lexicon and ontology-terminology models, injected into NLMs are promising for this task. Especially in this regard a connection to other modalities, such as speech synthesis approaches, could bring significant benefits.

6. Multimodal

NLMs promise to boost not only the SW's multilinguality but are capable of contributing to its multimodality. For the sake of the vision, a broad perspective will be adopted also considering multisensory approaches, from vision to tactile. Such multimodal representations can be utilized in intelligent conversational agents, multimodal information extraction, robotics, among many more.

Semantic Speech Technologies: Speech technologies building on SW resources and NLM systems promise to support important present-day applications, such as assisted living. Google registered a patent on utilizing language models for understanding conversations based on SW resources [58]. A speech interface for question-answering systems has been proposed [59], which, in combination with the above multilingual strategies for NLM-based question answering, could provide broad access to SW resources. Another Google patent for reformulations of speech queries has been registered [60], providing alternative queries if the submitted one returns no results. Most of these systems rely on text transcriptions utilizing automated speech recognition (ASR) systems. The recently published Translatotron [18] omits this step and translates directly from speech to speech in the speaker's voice.

Challenges and Opportunities: Intelligent voice interaction is a booming business model as much as vibrant research field. Building on neural-symbolic reasoning, such systems could enable a multilingual, multimodal query-answering system on formally structured resources. Major challenges here are similar to those of transcultural modeling. Local contexts and linguistic variations need to be taken into account to grant broad information access and a high usability. Nevertheless, achieving a speech-empowered SW technologies strongly furthers the endeavor to break down access barriers to represented information.

Semantic Video Technologies: In order to include the visual-manual modality to convey meaning in form of sign language, knowledge needs to be conveyed by video. [61] combine speech synthesis, machine translation, and SW technologies to create a machine-readable knowledge representation for the Turkish sign language. In consequence, NMT can be utilized to translate between natural language and sign language, as has been suggested utilizing the above sign language representation for Turkish [62].

Challenges and Opportunities: Semantic video technologies still suffer from a lack of broad coverage in terms of language and visual-manual modality. Establishing such a coverage potentially provides all users, including users with special needs or illiterate users, with access to information truly breaking down information barriers. In fact, very few SW sign language approaches can be found. Latest NLM advances can contribute to automating and improving existing approaches. For instance, VideoBERT [17] treats video frames as "video words" utilizing a vector quantization approach and an off-the-shelf speech recognition system to transcribe audio. Resulting representations allow for a seamless transition between text and video. Further adding cross-modal reasoning approaches could boost the interface between video and natural language [63]. A video-enabled NLM-based SW can strongly support barrier-free online communication, especially if transformations between different modalities are provided.

Semantic Sensor Web: A Semantic Sensor Web in the Internet of Things vision [64] is probably the closest corpus of related approaches. Building on SW enablement or Linked Data standards, sensor data are linked and annotated. Thus, SW query technologies can be applied to sensor data [65]. Its link to language models comes from the necessity of connecting sensor data to human communication means, the human-robot interface, such as natural language understanding of robot instructions which has been shown to benefit from ontology-natural language groundings [66].

Challenges and Opportunities: Linking sensory data and language can boost human-robot interactions, as sensory information, their semantic representation, and neural-symbolic reasoning could be highly beneficial to the task of explainable robotics and AI [38]. Recent advances in terms of cross-modal predictions [67] connected to SW technologies can potentially boost cognitive AI systems.

One major challenge of connecting NLMs with semantic sensor data is that of magnifying biases. NLMs have been shown in multiple studies to easily suffer from biases, which unfortunately is also true for sensor data, thereby bearing the risk of multiplying and intensifying biases across modalities. For example, sensors in self-driving cars have been shown to detect lighter skin tones better than darker ones [68].

Multisensor semantic data might also relate to neural patterns and the ability to automatically decode them. A recent approach managed to reconstruct a word from neural activation patterns from auditory inputs [69]. Thus, one future scenario of this neural-symbolic vision is the application of SW and NLM models in the brain-computer interface.

7. Conclusion

Building on selected existing approaches, this article laid out a vision of a multilingual, transcultural, and multimodal Semantic Web (SW) utilizing Neural Language Models (NLMs). Joining forces of SW and NLMs can boost a wide variety of tasks, such as extracting data from different languages and channels, formally interlinking them, and verbalizing logical answers in natural language or sign language in response to multimodal queries.

The biggest challenge and at the same time opportunity is a seamless connection of and transition between multilingual, transcultural, and multimodal knowledge representations. Individual bridges across this big gap have been built, such as transitioning from natural language text to video and sign language. However, integrating transcultural representations requires further investigations. In fact, cultural modeling and transcultural alignments might be the pillar that requires most further construction work to provide a footing for this vision, which targets a truly unbiased and fully accessible Semantic Web.

One central opportunity of this vision is the fact that seamlessly accessibly and dynamic SW technologies can foster not only cultural language evolution but simultaneously knowledge evolution. One key to this vision is thus an easily accessible representation mechanism - one that can easily be adopted by other communities, such as machine learning - that strongly embraces diversity and boosts diversity-aware AI, which in the end will foster robustness.

While the focus here was on benefits of SW technologies,some advantages that NLMs can obtain by joining forces with SW technologies have been braced upon. For instance, injecting structured and formal knowledge into NLM architectures have shown improvements for machine translation and textual entailment task. Further investigations in the ability of SW technologies to support NLM tasks and applications would be interesting.

References

[1] T. Berners-Lee, J. Hendler and O. Lassila, The Semantic Web, *Scientific American* **284**(5) (2001), 34–43. http://www.sciam.com/article.cfm?articleID=00048144-10D2-1C70-84A9809EC588EF21.

[2] R. Navigli and S.P. Ponzetto, BabelNet: Building a Very Large Multilingual Semantic Network, in: *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Uppsala, Sweden, 2010, pp. 216–225. https://www.aclweb.org/anthology/P10-1023.

[3] A. Gangemi, M. Alam, L. Asprino, V. Presutti and D.R. Recupero, Framester: A Wide Coverage Linguistic Linked Data Hub, in: *20th International Conference on Knowledge Engineering and Knowledge Management - Volume 10024*, EKAW 2016, Springer-Verlag New York, Inc., New York, NY, USA, 2016, pp. 239–254. ISBN 978-3-319-49003-8. doi:10.1007/978-3-319-49004-5_16.

[4] P. Buitelaar and P. Cimiano, *Towards the Multilingual Semantic Web: Principles, Methods and Applications*, Springer Publishing Company, 2014. ISBN 3662435845, 9783662435847.

[5] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser and I. Polosukhin, Attention is All You Need, in: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, Curran Associates Inc., USA, 2017, pp. 6000–6010. ISBN 978-1-5108-6096-4. http://dl.acm.org/citation.cfm?id=3295222.3295349.

[6] A. Gatt and E. Krahmer, Survey of the State of the Art in Natural Language Generation: Core Tasks, Applications and Evaluation, *Journal of Artificial Intelligence Research* **61**(1) (2018), 65–170. doi:10.1613/jair.5714.

[7] I. Androutsopoulos and P. Malakasiotis, A survey of paraphrasing and textual entailment methods, *Journal of Artificial Intelligence Research* **38** (2010), 135–187.

[8] A. Conneau, D. Kiela, H. Schwenk, L. Barrault and A. Bordes, Supervised Learning of Universal Sentence Representations from Natural Language Inference Data, in: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Copenhagen, Denmark, 2017, pp. 670–680. doi:10.18653/v1/D17-1070. https://www.aclweb.org/anthology/D17-1070.

[9] D. Gromann and T. Declerck, Comparing Pretrained Multilingual Word Embeddings on an Ontology Alignment Task, in: *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*, N. Calzolari, K. Choukri, C. Cieri, T. Declerck, S. Goggi, K. Hasida, H. Isahara, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odijk, S. Piperidis and T. Tokunaga, eds, European Language Resources Association (ELRA), Miyazaki, Japan, 2018, pp. 230–236. ISBN 979-10-95546-00-9. https://www.aclweb.org/anthology/L18-1034.

[10] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei and I. Sutskever, Language models are unsupervised multitask learners, *OpenAI Blog* (2019). https://d4mucfpksywv.cloudfront.net/better-language-models/language_models_are_unsupervised_multitask_learners.pdf.

[11] G. Petrucci, M. Rospocher and C. Ghidini, Expressive Ontology Learning as Neural Machine Translation, *Journal of Web Semantics* **52** (2018), 66–82. doi:10.1016/j.websem.2018.10.002.

[12] X. Yin, D. Gromann and S. Rudolph, Neural Machine Translating from Natural Language to SPARQL, *CoRR* **abs/2740229** (2010). http://arxiv.org/abs/1806.07941.

[13] B. Makni and J. Hendler, Deep learning for noise-tolerant RDFS reasoning, *Semantic Web* **10**(5) (2019), 823–862. doi:10.3233/SW-190363.

[14] A. Hepp, *Transcultural communication*, John Wiley & Sons, 2015. ISBN 978-0-470-67393-5.

[15] M. Arcan and P. Buitelaar, Translating Domain-Specific Expressions in Knowledge Bases with Neural Machine Translation, *CoRR* **abs/1709.02184** (2017). http://arxiv.org/abs/1709.02184.

[16] A. Liu, J. Du and V. Stoyanov, Knowledge-Augmented Language Model and Its Application to Unsupervised Named-Entity Recognition, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran and T. Solorio, eds, Association for Computational Linguistics, 2019, pp. 1142–1150. ISBN 978-1-950737-13-0. https://www.aclweb.org/anthology/N19-1117/.

[17] C. Sun, A. Myers, C. Vondrick, K. Murphy and C. Schmid, VideoBERT: A Joint Model for Video and Language Representation Learning, *CoRR* **abs/1904.01766** (2019). http://arxiv.org/abs/1904.01766.

[18] Y. Jia, R.J. Weiss, F. Biadsy, W. Macherey, M. Johnson, Z. Chen and Y. Wu, Direct speech-to-speech translation with a sequence-to-sequence model, *CoRR* **abs/1904.06037** (2019). http://arxiv.org/abs/1904.06037.

[19] J.T. Goodman, A Bit of Progress in Language Modeling, *Computer Speech and Language* **15**(4) (2001), 403–434. doi:10.1006/csla.2001.0174.

[20] T. Mikolov, I. Sutskever, K. Chen, G.S. Corrado and J. Dean, Distributed Representations of Words and Phrases and their Compositionality, in: *Advances in Neural Information Processing Systems 26*, C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani and K.Q. Weinberger, eds, Curran Associates, Inc., 2013, pp. 3111–3119. http://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf.

[21] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston and O. Yakhnenko, Translating Embeddings for Modeling Multi-relational Data, in: *Advances in Neural Information Processing Systems 26*, C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani and K.Q. Weinberger, eds, Curran Associates, Inc., 2013, pp. 2787–2795. http://papers.nips.cc/paper/5071-translating-embeddings-for-modeling-multi-relational-data.pdf.

[22] P. Gage, A New Algorithm for Data Compression, *C Users J.* **12**(2) (1994), 23–38. http://dl.acm.org/citation.cfm?id=177910.177914.

[23] J.P. McCrae and J. Gracia, Foreword to the Special Issue: "Towards the Multilingual Web of Data", *Information* **10**(2) (2019), 56. doi:10.3390/info10020056.

[24] P. Cimiano, J.P. McCrae and P. Buitelaar, Lexicon Model for Ontologies: Community Report, 2016. https://www.w3.org/2016/05/ontolex/.

[25] D. Gromann, A Model and Method to Terminologize Existing Domain Ontologies, in: *Terminology and Knowledge Engineering 2014*, Berlin, Germany, 2014. https://hal.archives-ouvertes.fr/ghal-01005867.

[26] T. Declerck and D. Gromann, Combining three Ways of Conveying Knowledge: Modularization of Domain, Terminological, and Linguistic Knowledge in Ontologies, in: *Proceedings of the 6th International Workshop on Modular Ontologies, Graz, Austria, July 24, 2012*, T. Schneider and D. Walther, eds, CEUR Workshop Proceedings, Vol. 875, CEUR-WS.org, 2012. http://ceur-ws.org/Vol-875/regular_paper_3.pdf.

[27] S. Hellmann, J. Lehmann, S. Auer and M. Brümmer, Integrating NLP using Linked Data, in: *12th International Semantic Web Conference, 21-25 October 2013, Sydney, Australia*, 2013. doi:10.1007/978-3-642-41338-4_7.

[28] C. Federmann, D. Gromann, T. Declerck, S. Hunsicker, H. Krieger and G. Budin, Multilingual terminology acquisition for ontology-based information extraction, in: *Proceedings of the 10th Terminology and Knowledge Engineering Conference (TKE 2012)*, G.A. de Cea, M.C. Suarez-Figueroa, R. Garcia-Castro and E. Montiel-Ponsoda, eds, Madrid, Spain, 2012, pp. 166–175.

[29] D. Gromann and T. Declerck, A Cross-Lingual Correcting and Completive Method for Multilingual Ontology Labels, in: *Towards the Multilingual Semantic Web, Principles, Methods and Applications*, P. Buitelaar and P. Cimiano, eds, Springer, 2014, pp. 227–242. ISBN 978-3-662-43584-7. doi:10.1007/978-3-662-43585-4_14.

[30] D. Gerber and A.N. Ngomo, From RDF to Natural Language and Back, in: *Towards the Multilingual Semantic Web, Principles, Methods and Applications*, P. Buitelaar and P. Cimiano, eds, Springer, 2014, pp. 193–209. ISBN 978-3-662-43584-7. doi:10.1007/978-3-662-43585-4_12.

[31] S. Rudolph, J. Völker and P. Hitzler, Supporting Lexical Ontology Learning by Relational Exploration, in: *Conceptual Structures: Knowledge Architectures for Smart Applications*, U. Priss, S. Polovina and R. Hill, eds, Springer Berlin Heidelberg, Berlin, Heidelberg, 2007, pp. 488–491. ISBN 978-3-540-73681-3. doi:10.1007/978-3-540-73681-3_41.

[32] D. Vrandečić and M. Krötzsch, Wikidata: A Free Collaborative Knowledge Base, *Communications of the ACM* **57** (2014), 78–85. http://cacm.acm.org/magazines/2014/10/178785-wikidata/fulltext.

[33] Domain Adaptation for Ontology Localization, *Web Semantics* **36**(C) (2016), 23–31. doi:10.1016/j.websem.2015.12.001.

[34] R. Sennrich, B. Haddow and A. Birch, Improving Neural Machine Translation Models with Monolingual Data (2016), 86–96. doi:10.18653/v1/P16-1009. https://www.aclweb.org/anthology/P16-1009.

[35] D. Moussallem, M. Wauer and A.N. Ngomo, Machine Translation using Semantic Web Technologies: A Survey, *Journal of Web Semantics* **51** (2018), 1–19. doi:https://doi.org/10.1016/j.websem.2018.07.001. https://www.sciencedirect.com/science/article/pii/S1570826818300301.

[36] F. Hommel, P. Cimiano, M. Orlikowski and M. Hartung, Extending Neural Question Answering with Linguistic Input Features, in: *Proceedings of the 5th Workshop on Semantic Deep Learning (SemDeep-5)*, Association for Computational Linguistics, Macau, China, 2019, pp. 31–39. https://www.aclweb.org/anthology/W19-5806.

[37] P. Hitzler, F. Bianchi, M. Ebrahimi and K.M. Sarker, Neural-Symbolic Integration and the Semantic Web, *Semantic Web Journal* **this issue** (forthcoming).

[38] F. Lecue, On the Role of Knolwedge Graphs in Explainable AI, *Semantic Web Journal* **this issue** (forthcoming).

[39] J. Weston, A. Bordes, O. Yakhnenko and N. Usunier, Connecting Language and Knowledge Bases with Embedding Models for Relation Extraction, in: *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest*

Group of the ACL, ACL, 2013, pp. 1366–1371. ISBN 978-1-937284-97-8. https://www.aclweb.org/anthology/D13-1136/.

[40] D. Gerber and A.-C.N. Ngomo, Extracting Multilingual Natural-Language Patterns for RDF Predicates, in: *Knowledge Engineering and Knowledge Management*, A. ten Teije, J. Völker, S. Handschuh, H. Stuckenschmidt, M. d'Acquin, A. Nikolov, N. Aussenac-Gilles and N. Hernandez, eds, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 87–96. ISBN 978-3-642-33876-2.

[41] T. Mossakowski, M. Codescu, F. Neuhaus and O. Kutz, *The Distributed Ontology, Modeling and Specification Language – DOL*, in: *The Road to Universal Logic: Festschrift for the 50th Birthday of Jean-Yves Béziau Volume II*, A. Koslow and A. Buchsbaum, eds, Springer International Publishing, Cham, 2015, pp. 489–520. ISBN 978-3-319-15368-1.

[42] M. Alirezaie, M. Längkvist, M. Sioutis and A. Loutfi, Semantic referee: A neural-symbolic framework for enhancing geospatial semantic segmentation, *Semantic Web* **10**(5) (2019), 863–880. doi:10.3233/SW-190362.

[43] E. Jiménez-Ruiz, A. Agibetov, M. Samwald and V. Cross, We Divide, You Conquer: From Large-scale Ontology Alignment to Manageable Subtasks with a Lexical Index and Neural Embeddings, in: *Proceedings of the ISWC 2018 Posters & Demonstrations, Industry and Blue Sky Ideas Tracks co-located with 17th International Semantic Web Conference (ISWC 2018), Monterey, USA, October 8th - to - 12th, 2018.*, M. van Erp, M. Atre, V. López, K. Srinivas and C. Fortuna, eds, CEUR Workshop Proceedings, Vol. 2180, CEUR-WS.org, 2018. http://ceur-ws.org/Vol-2180/paper-28.pdf.

[44] M. Alam, D. Reforgiato Recupero, M. Mongiovi, A. Gangemi and P. Ristoski, Event-based Knowledge Reconciliation Using Frame Embeddings and Frame Similarity, *Know.-Based Syst.* **135**(C) (2017), 192–203. doi:10.1016/j.knosys.2017.08.014.

[45] A.H. Gutchess, R.C. Welsh, A. Boduroĝlu and D.C. Park, Cultural differences in neural function associated with object processing, *Cognitive, Affective, & Behavioral Neuroscience* **6**(2) (2006), 102–109. doi:10.3758/CABN.6.2.102.

[46] A. Mesoudi, A. Whiten and K.N. Laland, Towards a unified science of cultural evolution, *Behavioral and Brain Sciences* **29**(4) (2006), 329–347. doi:10.1017/S0140525X06009083.

[47] P.D. Turney, Conditions for Major Transitions in Biological and Cultural Evolution, *CoRR* **abs/1806.07941** (2018). http://arxiv.org/abs/1806.07941.

[48] L. Steels, *Experiments in cultural language evolution*, Vol. 3, John Benjamins Publishing, 2012. https://benjamins.com/catalog/ais.3.

[49] J. Euzenat, First Experiments in Cultural Alignment Repair, in: *Proceedings of the Third International Workshop on Debugging Ontologies and Ontology Mappings, WoDOOM 2014, co-located with 11th Extended Semantic Web Conference (ESWC 2014), Anissaras/Hersonissou, Greece, May 26, 2014.*, P. Lambrix, G. Qi, M. Horridge and B. Parsia, eds, CEUR Workshop Proceedings, Vol. 1162, CEUR-WS.org, 2014, pp. 3–14. http://ceur-ws.org/Vol-1162/paper1.pdf.

[50] E. Iadanza, F. Maietti, A. Ziri, R. Di Giulio, M. Medici, F. Ferrari, P. Bonsma, B. Turillazzi et al., Semantic Web Technologies Meet BIM for Accessing and Understanding Cultural Heritage, in: *8th International Workshop 3D-ARCH 3D Virtual Reconstruction and Visualization of Complex Architectures"*, Vol. 42, Copenicus, 2019, pp. 381–388.

[51] M. Papadopoulou and C. Roche, Twinning Classics and AI: Building the new generation of ontology-based lexicographical tools and resources for Humanists on the Semantic Web, in: *TwinTalks@ DHN*, Vol. 2365, CEUR Workshop Proceedings, 2019, pp. 67–81. http://ceur-ws.org/Vol-2365/08-TwinTalks-DHN2019_paper_8.pdf.

[52] V.A. Carriero, A. Gangemi, M.L. Mancinelli, L. Marinucci, A.G. Nuzzolese, V. Presutti and C. Veninata, ArCo: the Italian Cultural Heritage Knowledge Graph, *CoRR* **abs/1905.02840** (2019). http://arxiv.org/abs/1905.02840.

[53] A. Corn and w.S.W.J. Patrick, Exploring the applicability of the Semantic Web for discovering and navigating Australian Indigenous knowledge resources, *Archives and Manuscripts* **47**(1) (2019), 131–152. doi:10.1080/01576895.2019.1575248.

[54] B. Li and P. Cheng, Learning Neural Representation for CLIR with Adversarial Framework, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, E. Riloff, D. Chiang, J. Hockenmaier and J. Tsujii, eds, Association for Computational Linguistics, 2018, pp. 1861–1870. ISBN 978-1-948087-84-1. https://www.aclweb.org/anthology/D18-1212/.

[55] A.C. Kozlowski, M. Taddy and J.A. Evans, The Geometry of Culture: Analyzing Meaning through Word Embeddings, *CoRR* **abs/1803.09288** (2018). http://arxiv.org/abs/1803.09288.

[56] D. Gromann and M.M. Hedblom, Body-Mind-Language: Multilingual Knowledge Extraction Based on Embodied Cognition, in: *Proceedings of the 5th International Workshop on Artificial Intelligence and Cognition, AIC 2017, Larnaca, Cyprus, November 1-3, 2017.*, I. Diakidoy, A.C. Kakas, A. Lieto and L. Michael, eds, CEUR Workshop Proceedings, Vol. 2090, CEUR-WS.org, 2017, pp. 20–33. http://ceur-ws.org/Vol-2090/paper2.pdf.

[57] D. Gromann and M.M. Hedblom, Kinesthetic Mind Reader: A Method to Identify Image Schemas in Natural Language, in: *Proceedings of Advancements in Cogntivie Systems*, 2017.

[58] M. Akbacak, D.Z. Hakkani-Tur, G. Tur, L.P. Heck and B. Dumoulin, Language modeling for conversational understanding domains using semantic web resources, Google Patents, 2017, US Patent 9,679,558. https://patents.google.com/patent/US9478145.

[59] A.J. Kumar, C. Schmidt and J. Köhler, A knowledge graph based speech interface for question answering systems, *Speech Communication* **92** (2017), 1–12. doi:10.1016/j.specom.2017.05.001.

[60] B. Dumoulin, A. Ahmadi, S. Parthasarathy, N. Craswell, U. Ozertem, M. Shokouhi, K. Raghunathan and R. Jones, Interactive reformulation of speech queries, Google Patents, 2019, US Patent App. 10/176,219. https://patents.google.com/patent/US8521526.

[61] C. Eryiğit, H. Köse, M. Kelepir and G. Eryiğit, Building machine-readable knowledge representations for Turkish sign language generation, *Knowledge-Based Systems* **108** (2016), 179–194. doi:10.1016/j.knosys.2016.04.014.

[62] M. Selcuk-Simsek and I. Cicekli, Bidirectional Machine Translation Between Turkish and Turkish Sign Language : A Data-Driven Approach, *International Journal on Natural Language Computing (IJNLC)* **6**(3) (2017). doi:10.5121/ijnlc.2017.6303.

[63] S. Ghosh, A. Agarwal, Z. Parekh and A.G. Hauptmann, ExCL: Extractive Clip Localization Using Natural Language Descriptions, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran and T. Solorio, eds, Association for Computational Linguistics, 2019, pp. 1984–1990. ISBN 978-1-950737-13-0. https://www.aclweb.org/anthology/N19-1198/.

[64] A. Sheth, C. Henson and S. S. Sahoo, Semantic Sensor Web, *Internet Computing, IEEE* **12** (2008), 78–83. doi:10.1109/MIC.2008.87.

[65] J. Calbimonte, H. Jeung, Ó. Corcho and K. Aberer, Enabling Query Technologies for the Semantic Sensor Web, *Int. J. Semantic Web Inf. Syst.* **8**(1) (2012), 43–63. doi:10.4018/jswis.2012010103.

[66] S. Patki, A.F. Daniele, M.R. Walter and T.M. Howard, Inferring Compact Representations for Efficient Natural Language Understanding of Robot Instructions, *CoRR* **abs/1903.09243** (2019). http://arxiv.org/abs/1903.09243.

[67] Y. Li, J.-Y. Zhu, R. Tedrake and A. Torralba, Connecting Touch and Vision via Cross-Modal Prediction, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10609–10618.

[68] B. Wilson, J. Hoffman and J. Morgenstern, Predictive Inequity in Object Detection, *CoRR* **abs/1902.11097** (2019). http://arxiv.org/abs/1902.11097.

[69] H. Akbari, B. Khalighinejad, J.L. Herrero, A.D. Mehta and N. Mesgarani, Towards reconstructing intelligible speech from the human auditory cortex, *Scientific reports* **9**(1) (2019). doi:10.1101/350124.

[70] J. Burstein, C. Doran and T. Solorio (eds), Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers), Association for Computational Linguistics, 2019. ISBN 978-1-950737-13-0. https://www.aclweb.org/anthology/volumes/N19-1/.

[71] P. Buitelaar and P. Cimiano (eds), *Towards the Multilingual Semantic Web, Principles, Methods and Applications*, Springer, 2014. ISBN 978-3-662-43584-7. doi:10.1007/978-3-662-43585-4.