Semantic Web 0 (0) 1
IOS Press

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

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 and multilingual text-video knowledge representation with embeddings. 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.

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

1. Introduction

One central endeavor of the Semantic Web (SW) 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 knolwedge resources, such as BabelNet [45] and Framester [16], to multilingual methods and applications (see e.g. [5]). Nevertheless, approaches that further improve the level of automation, usability, and interoperability are needed.

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 and language generation to textual entailment have been proposed.

Neural Language Models (NLMs) learn an implicit semantic representations of sequences on their hidden layer, resulting in a dense 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 [23]. Recent NLMs have provided a strong backbone to many Artificial Intelligence (AI) applications that go beyond traditional NLP tasks, see for instance Radford et al. [49] 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 spaio-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 [48] and structured query [61] 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, structure or unstructured. Neural Machine Translation (NMT) based on NLMs has even been applied to noise-tolerant RDFS reasoning [39].

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. [28]). In terms of culture, NLMs potentially provide interesting methods to port information learned for one language and culture to another in form of domain adaptation and transfer learning [4, 38]. Nevertheless, further principles are 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 [55] and translate directly from speech to speech without a need for textual transcriptions [31].

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. Goodman [20] for more LM 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 Mikolov et al. [43]. 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.

A common architectural style is that of encoder-decoder, either of which can be independent models. The encoder maps an input sequence $(x_1, x_2, ..., x_n)$ to a continuous representation $z = (z_1, z_2, ..., z_n)$, which is given to the decoder, which generates an output sequence $(y_1, y_2, ..., y_n)$ one symbol at a time. At each step previously generated symbols are considered as additional input to generate the next, in form of beam search or similar algorithms.

One recent best-performer in machine translation, the Transformer model [57], has soon been propagated to many NLP tasks, from multi-task approaches [49] to text-video combinations [55]. 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 [15] 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 [5, 40] 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 [8]. Second, similar models have been proposed to interchange domain-specific terminological information grounded in ontological resources [21]. Combined representations of linguistic, terminological, and ontological knoweldge have been modeled [10]. 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 [14], completing and correcting natural language information [22], translating from knowledge resource to natural language and/or vice versa [18], and ontology learning from text [50].

Over the past few decades the Linked Open Data (LOD) cloud and resources published in RDF and OWL have experiences a tremendous growth, however, predominantly in English with several notable exceptions, such as BabelNet [45] and WikiData [58]. To foster this endeavor, automated means, such as NLMs, can improve and fasten approaches.

4. Multilingual

In contrast to multimodal, multilingual focuses on the treatment of written text. Within the context of this article, the focus is on SW technologies and contributes of NLMs to SW tasks.

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 [4]. 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 [41] and fine-tune machine translation [52].

Challenges and Opportunities: As concluded in a recent survey on machine translation and SW technologies [44], 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 and lexical knowledge into NMT systems, which could benefit from further systematic experiments since its exact nature potentially evolves with constantly developing new architectures. Furthermore, a 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 [59]. Gerber and Ngomo [17] 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 [48]. 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 [61].

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 under-resourced languages.

Machine Translation for Reasoning: A very recent approach is that of tailoring embeddings to accommodate RDFS reasoning in an NMT task [39]. 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. Makni and Hendler [39] suggest embeddings tailored to reasoning as a first approach. Furthermore, learning distributed vector representations of embeddings in one formal language might allow for a transition to or similarity measure of different formal languages. A discussion specific to neural-symbolic systems can be found in Hitzler et al. [29] 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 [23]. The task of aligning large ontologies has been subdivided into smaller, tractable tasks utilizing a lexical index, neural embeddings, and locality models [32].

A broader alignment strategy is that of bringing together a multitude of resource from the Linguistic Linked Open Data (LLOD) cloud with ontology resources in Framester [16]. Based on this resource, frame-based embeddings are trained and utilized for knowledge reconciliation purposes [3], 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.

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 withing 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 [26].

Cultural Evolution: Cultural evolution is closely tied to evolutionary biology science and Darwinian evolutionary principles [42]. 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 [56]. 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 [54]. 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 [13].

Challenges and Opportunities: A connection of cultural evolution and language as well as connections to knowledge representation have been thoroughly investigated [54]. 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.

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, [30]) to creating ontology-based lexicographic tools for the study of ancient culture to enable object-multilingual links [46]. Culture-specific knowledge graphs of cultural heritage have been proposed, such as for Italy [7]. 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. Neural-symbolic reasoning could be particularly powerful for such alignments.

Culture-specific Modeling: Another important transcultural SW connection is that of utilizing ontologies for culture-specific modeling. For instance, Corn and Patrick [9] 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 [36]. Embedding spaces have also been analyzed for their ability to represent culture-specific association [33] 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 query-answering and potentially automated localization approaches. 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 [24, 25]. Such a cognitive frameworks 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 mutisensory 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 [1]. A speech interface for question-answering systems has been proposed [34], 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 [11], 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 [31] 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.

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. Eryiğit et al. [12] 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 [51].

Challenges and Opportunities: Semantic video technologies still suffer from a lack of broad coverage in terms of language and visual-manual modality. 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 [55] 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 [19].

Semantic Sensor Web: A Semantic Sensor Web in the Internet of Things vision [53] 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 [6]. 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 [47].

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 [35]. Recent advances in terms of cross-modal predictions [37] 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 [60].

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 [2]. 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.

While the focus here was on benefits of SW technologies, in future it would be interesting to discuss all advantages NLMs can obtain by joining forces with

SW technologies. For instance, injecting structured and formal knowledge into NLM architectures have shown improvements for machine translation and textual entailment task.

References

[1] Murat Akbacak, Dilek Z Hakkani-Tur, Gokhan Tur, Larry P Heck, and Benoit Dumoulin. Language modeling for conversational understanding domains using semantic web resources, June 13 2017. US Patent 9,679,558.

[2] Hassan Akbari, Bahar Khalighinejad, Jose L Herrero, Ashesh D Mehta, and Nima Mesgarani. Towards reconstructing intelligible speech from the human auditory cortex. *Scientific reports*, 9(1):874, 2019.

[3] Mehwish Alam, Diego Reforgiato Recupero, Misael Mongiovi, Aldo Gangemi, and Petar Ristoski. Event-based knowledge reconciliation using frame embeddings and frame similarity. *Knowledge-Based Systems*, 135:192–203, 2017.

[4] Mihael Arcan and Paul Buitelaar. Translating domain-specific expressions in knowledge bases with neural machine translation. *arXiv preprint arXiv:1709.02184*, 2017.

[5] Paul Buitelaar and Philipp Cimiano. *Towards the multilingual semantic web: principles, methods and applications*. Springer, 2014.

[6] Jean-Paul Calbimonte, Hoyoung Jeung, Oscar Corcho, and Karl Aberer. Enabling query technologies for the semantic sensor web. *International Journal On Semantic Web and Information Systems (IJSWIS)*, 8(1):43–63, 2012.

[7] Valentina Anita Carriero, Aldo Gangemi, Maria Letizia Mancinelli, Ludovica Marinucci, Andrea Giovanni Nuzzolese, Valentina Presutti, and Chiara Veninata. Arco: the italian cultural heritage knowledge graph. *CoRR*, abs/1905.02840, 2019. URL http://arxiv.org/abs/1905.02840.

[8] Philipp Cimiano, John P. McCrae, and Paul Buitelaar. Lexicon model for ontolo-gies: Community report, 2016.

[9] Aaron Corn and with Steven Wantarri Jampijinpa Patrick. Exploring the applicability of the semantic web for discovering and navigating australian indigenous knowledge resources. *Archives and Manuscripts*, 47(1):131–152, 2019.

[10] Thierry Declerck and Dagmar 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*, volume 875, pages 28–40, 2012.

[11] Benoit Dumoulin, Ali Ahmadi, Sarangarajan Parthasarathy, Nick Craswell, Umut Ozertem, Milad Shokouhi, Karthik Raghunathan, and Rosie Jones. Interactive reformulation of speech queries, January 8 2019. US Patent App. 10/176,219.

[12] Cihat Eryiğit, Hatice Köse, Meltem Kelepir, and Gülşen Eryiğit. Building machine-readable knowledge representations for turkish sign language generation. *Knowledge-Based Systems*, 108:179–194, 2016.

[13] Jérôme Euzenat. First experiments in cultural alignment repair (extended version). In *European Semantic Web Conference*, pages 115–130. Springer, 2014.

[14] Christian Federmann, Dagmar Gromann, Thierry Declerck, Sabine Hunsicker, H Krieger, and Gerhard Budin. Multilingual terminology acquisition for ontology-based information extraction. In *Proceedings of the 10th Terminology and Knowledge Engineering Conference (TKE 2012)*, pages 166–175, 2012.

[15] Philip Gage. A new algorithm for data compression. *The C Users Journal*, 12(2):23–38, 1994.

[16] Aldo Gangemi, Mehwish Alam, Luigi Asprino, Valentina Presutti, and Diego Reforgiato Recupero. Framester: a wide coverage linguistic linked data hub. In *European Knowledge Acquisition Workshop*, pages 239–254. Springer, 2016.

[17] Daniel Gerber and Axel-Cyrille Ngonga Ngomo. Extracting multilingual natural-language patterns for rdf predicates. In *International Conference on Knowledge Engineering and Knowledge Management*, pages 87–96. Springer, 2012.

[18] Daniel Gerber and Axel-Cyrille Ngonga Ngomo. From rdf to natural language and back. In *Towards the Multilingual Semantic Web*, pages 193–209. Springer, 2014.

[19] Soham Ghosh, Anuva Agarwal, Zarana Parekh, and Alexander G. Hauptmann. Excl: Extractive clip localization using natural language descriptions. 2019.

[20] Joshua T Goodman. A bit of progress in language modeling. *Computer Speech & Language*, 15(4):403–434, 2001.

[21] Dagmar Gromann. A model and method to terminologize existing domain ontologies. In *Terminology and Knowledge Engineering 2014*, pages 10–p, 2014.

[22] Dagmar Gromann and Thierry Declerck. A cross-lingual correcting and completive method for multilingual ontology labels. In *Towards the Multilingual Semantic Web*, pages 227–242. Springer, 2014.

[23] Dagmar Gromann and Thierry 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)*, 2018.

[24] Dagmar Gromann and Maria M. Hedblom. Kinesthetic mind reader: A method to identify image schemas in natural language. In *Proceedings of Advancements in Cogntivie Systems*, 2017.

[25] Dagmar Gromann and Maria M Hedblom. Body-mind-language: Multilingual knowledge extraction based on embodied cognition. In *AIC*, pages 20–33, 2017.

[26] Angela H Gutchess, Trey Hedden, Sarah Ketay, Arthur Aron, and John DE Gabrieli. Neural differences in the processing of semantic relationships across cultures. *Social cognitive and affective neuroscience*, 5(2-3):254–263, 2010.

[27] Sebastian Hellmann, Jens Lehmann, Sören Auer, and Martin Brümmer. Integrating nlp using linked data. In *International semantic web conference*, pages 98–113. Springer, 2013.

[28] Andreas Hepp. *Transcultural communication*. John Wiley & Sons, 2015.

[29] Pascal Hitzler, Frederico Bianchi, Monireh Ebrahimi, and Kamruzzaman Md Sarker. Neural-symbolic integration and the semantic web. *this issue*, forthcoming.

[30] E Iadanza, F Maietti, AE 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"*, volume 42, pages 381–388. Copenicus, 2019.

[31] Ye Jia, Ron J Weiss, Fadi Biadsy, Wolfgang Macherey, Melvin Johnson, Zhifeng Chen, and Yonghui Wu. Direct speech-to-speech translation with a sequence-to-sequence model. *arXiv preprint arXiv:1904.06037*, 2019.

[32] Ernesto Jiménez-Ruiz, Asan Agibetov, Matthias Samwald, and Valerie Cross. We divide, you conquer: From large-scale ontology alignment to manageable subtasks. In *Ontology Matching: OM-2018: Proceedings of the ISWC Workshop*, page 13, 2018.

[33] Austin C. Kozlowski, Matt Taddy, and James A. Evans. The geometry of culture: Analyzing meaning through word embeddings. *CoRR*, abs/1803.09288, 2018. URL http://arxiv.org/abs/1803.09288.

[34] Ashwini Jaya Kumar, Christoph Schmidt, and Joachim Köhler. A knowledge graph based speech interface for question answering systems. *Speech Communication*, 92:1–12, 2017.

[35] Freddy Lecue. On the role of knolwedge graphs in explainable ai. *this issue*, forthcoming.

[36] Bo Li and Ping Cheng. Learning neural representation for clir with adversarial framework. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1861–1870, 2018.

[37] Yunzhu Li, Jun-Yan Zhu, Russ Tedrake, and Antonio Torralba. Connecting touch and vision via cross-modal prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10609–10618, 2019.

[38] Angli Liu, Jingfei Du, and Veselin Stoyanov. Knowledge-augmented language model and its application to unsupervised named-entity recognition. *arXiv preprint arXiv:1904.04458*, 2019.

[39] Bassem Makni and James Hendler. Deep learning for noise-tolerant rdfs reasoning. *Semantic Web: Special Issue on Semantic Deep Learning*, (Preprint), 2019.

[40] John P McCrae and Jorge Gracia. Foreword to the special issue:âĂIJtowards the multilingual web of dataâĂİ. *Information*, 10, 2019.

[41] John P McCrae, Mihael Arcan, Kartik Asooja, Jorge Gracia, Paul Buitelaar, and Philipp Cimiano. Domain adaptation for ontology localization. *Web Semantics: Science, Services and Agents on the World Wide Web*, 36:23–31, 2016.

[42] Alex Mesoudi, Andrew Whiten, and Kevin N Laland. Towards a unified science of cultural evolution. *Behavioral and Brain Sciences*, 29(4):329–347, 2006.

[43] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.

[44] Diego Moussallem, Matthias Wauer, and Axel-Cyrille Ngonga Ngomo. Machine translation using semantic web technologies: A survey. *Journal of Web Semantics*, 51:1–19, 2018.

[45] Roberto Navigli and Simone Paolo Ponzetto. Babelnet: Building a very large multilingual semantic network. In *Proceedings of the 48th annual meeting of the association for computational linguistics*, pages 216–225. Association for Computational Linguistics, 2010.

[46] Maria Papadopoulou and Christophe 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*, pages 67–81, 2019.

[47] Siddharth Patki, Andrea F Daniele, Matthew R Walter, and Thomas M Howard. Inferring compact representations for efficient natural language understanding of robot instructions. *arXiv preprint arXiv:1903.09243*, 2019.

[48] Giulio Petrucci, Marco Rospocher, and Chiara Ghidini. Expressive ontology learning as neural machine translation. *Journal of Web Semantics*, 52:66–82, 2018.

[49] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 1:8, 2019.

[50] Sebastian Rudolph, Johanna Völker, and Pascal Hitzler. Supporting lexical ontology learning by relational exploration. In *International Conference on Conceptual Structures*, pages 488–491. Springer, 2007.

[51] Merve Selcuk-Simsek and Ilyas ÃĞiÃğekli. Bidirectional machine translation between turkish and turkish sign language : A data-driven approach. *International Journal on Natural Language Computing (IJNLC)*, 6(3), 2017.

[52] Rico Sennrich, Barry Haddow, and Alexandra Birch. Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709*, 2015.

[53] Amit Sheth, Cory Henson, and Satya S. Sahoo. Semantic sensor web. *Internet Computing, IEEE*, 12:78–83, 08 2008. .

[54] Luc Steels. *Experiments in cultural language evolution*, volume 3. John Benjamins Publishing, 2012.

[55] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. *arXiv preprint arXiv:1904.01766*, 2019.

[56] Peter D. Turney. Conditions for major transitions in biological and cultural evolution. *CoRR*, abs/1806.07941, 2018. URL http://arxiv.org/abs/1806.07941.

[57] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.

[58] Denny Vrandečić and Markus Krötzsch. Wikidata: a free collaborative knowledge base. 2014.

[59] Jason Weston, Antoine Bordes, Oksana Yakhnenko, and Nicolas Usunier. Connecting language and knowledge bases with embedding models for relation extraction. *arXiv preprint arXiv:1307.7973*, 2013.

[60] Benjamin Wilson, Judy Hoffman, and Jamie Morgenstern. Predictive inequity in object detection. *arXiv preprint arXiv:1902.11097*, 2019.

[61] Xiaoyu Yin, Dagmar Gromann, and Sebstian Rudolph. Neural machine translating from natural language to sparql. *CoRR*, abs/2740229, 2010. URL http://arxiv.org/abs/1806.07941.