

Abstract Brochure

Abstracts for the LCT Annual Meeting Poster Session

June 23–24, 2022

This brochure contains abstracts for posters being presented at the LCT annual meeting of 2022, hosted by the University of the Basque Country. The section numbers in this brochure correspond to the respective poster slots in the poster session.



With the support of the
Erasmus+ Programme
of the European Union



Contents

1	Investigating Adapters Effectiveness in Machine Translation	4
2	Multimodal Information Extraction in visually rich documents: Enriching linguistic features	5
3	Salient entity detection for event extraction	6
4	Grounding Negation in Visual Dialogue Games	7
5	Can neural networks write movie plots?	8
6	MEME-R: Meme Retrieval system	9
7	Interpreting the Task-Dependent Changes in Multimodal Transformers with Multimodal Brains	10
8	Highlight detection in live streams using audience reactions with transformer language models	11
9	Automatic detection of SWR in EEG signal to study correlation of SWR with memory performance in humans. (Preliminar Results)	12
10	Functional Data Analysis of Non-manual Marking of Questions in Kazakh-Russian Sign Language	13
11	On the Cusp of Comprehensibility: Can Models Distinguish Between Metaphors and Nonsense?	14
12	Patterns of human judgments of the complexity, acceptability, and grammaticality in Italian sentences.	15
13	Cross-lingual argument mining in medical domain	16

14 Neural Chart Summarization	17
15 Improving Automatic Speech Recognition for Regionally-Accented French	18
16 Cross-lingual transfer learning with Persian	19
17 Accent stigma of nonnative accents	20
18 Shrinking Knowledge Base Size: Dimension Reduction, Splitting & Filtering	21
19 Investigating LLMs' Processing of Nominal Plurality using Probing and Causal Interventions	22
20 Measuring readability of technical texts	23
21 Research Review and Direction for Dialogue-based Computer Assisted Language Learning	24
22 Subword-based Cross-lingual Transfer of Embeddings from Hindi to Marathi and Nepali	25
23 Fighting hate speech by understanding implied bias and selecting proper counter-narrative	26
24 Entity-centric Event Sequences to Highlight Perspectives	27
25 Towards an Analysis of Ordering Restrictions within Noun Phrases: A Top-Down Minimalist Approach.	30

1 Investigating Adapters Effectiveness in Machine Translation

Author: Aditya Kurniawan

Presented by: Aditya Kurniawan

Abstract

Pre-trained language models received extensive attention in recent years. However, it is still challenging to incorporate a pre-trained model such as BERT into Natural Language Generation tasks. Adapters are a promising approach that allows to fine-tune only a very small fraction of a pretrained network. In this work, we study the effectiveness of adapters in the Transformer model for machine translation task. We put adapters either in the encoder or the decoder only and we also attempt to down-scale the pre-trained model size to decrease GPU memory demands. We found that incorporating adapters in the encoder alone matches the performance of the setup when we include the adapters on both the encoder and decoder. In our down-scaling study, we found that using only a half of the original pre-trained weights can positively impact the performance when fine-tuned with adapters. Our experiments show that we can get almost the same performance as the original BERT model after fine-tuning the cross-attention layer.

2 Multimodal Information Extraction in visually rich documents: Enriching linguistic features

Authors: Sharmila Upadhyaya

Presented by: Sharmila Upadhyaya

Abstract

Information Extraction (IE) is a crucial sub-field of Natural Language Processing. It comprises various business applications obtainable in natural language and this encoded information plays a significant role in business decisions. IE in distinction to these documents provides extracted features to carry out groundwork for further research: in making efficient and feasible business solutions. For years, various multimodal pretraining techniques have been proposed, and this work attempts to implement transfer learning using those models to smaller neural network layers. Furthermore, this thesis incorporates the impact of linguistic features when trained explicitly in IE.

3 Salient entity detection for event extraction

Authors: Antonia Claésia da Costa Souza

Presented by: Antonia Claésia da Costa Souza

Abstract

Events can be explained as dynamic data structures that play an important role in understanding phenomena that occur in the real world. As a result, an event is a way to explain complex relations between people, locations, actions, and objects. This work explores approaches for event extracting and introduces the task of salient named entity detection for event extraction. In this paper, we have as the goal to try to answer and give the primary approach of the following question: “how do we know an entity is relevant within the context of a document for the purpose of identifying news events?”.

4 Grounding Negation in Visual Dialogue Games

Author: Alex Lucassen

Presented by: Alex Lucassen

Abstract

It is crucial for problem solving tasks involving visual dialogue to learn to distribute the probabilities among the candidates, but little work has been done on this topic. For this study, a Visual Dialogue referential game has been used to examine the probabilities assigned by the Guesser model, focusing specifically on cases of negation. An initial model based on the GDSE architecture was found to perform relatively well for positively-answered questions, but to assign probabilities incorrectly when the answer was negative. A new Guesser model trained using KL-divergence loss significantly improved the model's performance for category questions, but not for spatial questions. We believe this may be because unlike category questions, spatial questions require a soft-labelling approach.

5 Can neural networks write movie plots?

Authors: Jelena Sarajlić

Presented by: Jelena Sarajlić

Abstract

In the past few years, deep learning, neural networks, and artificial intelligence have been becoming more and more known and interesting even to the wider public. The reason for this is the fact that neural networks can approach human performance or even outperform humans in some tasks. While this might be true for tasks that can be somehow brought down to linear algebra, vectors and matrices, can it also be true for, for example, arts? Art is something that is very unique and sacred to humans and one of the most important ways of expressing ourselves. In this work, we will try to train or fine tune different neural networks and architectures for (conditioned) natural language using movie plots and test their performance, as well as review some related work that has been carried out in this field.

6 MEME-R: Meme Retrieval system

Author: Murali Kondragunta, Xabier Lahuerta and Ruan Chaves

Presented by: Murali Kondragunta

Abstract

With the rise in social media, memes became an integral part of the culture reflecting current political and social scenarios. Understanding memes requires a perfect multimodal system that can understand humour and map a described situation in text to the theme represented in the image. To this extent, we built a meme retrieval system that can extract relevant images when provided a theme/situation in the textual format and vice-versa. Our system relies on an existing database of memes, which are filtered based on the input scenario mentioned in the query text/image. Qualitative results show that the multimodal models like CLIP are good at capturing the intent behind the memes in comparison to the text-based models like Sentence Transformers and BM25.

7 Interpreting the Task-Dependent Changes in Multimodal Transformers with Multimodal Brains

Authors: Tianai Dong (Dota), Mariya Toneva

Presented by: Tianai Dong (Dota)

Abstract

Fine-tuning pre-trained multimodal transformers have achieved state-of-the-art performance in many language-and-vision tasks. However, little is known about how learned representations in these models change through fine-tuning. In this work, we introduce a novel approach to account for task-dependent variations in artificial neural networks, using the prior knowledge of multimodal information processing in the human brain. We exploit functional MRI recordings of human subjects while watching TV shows, to interpret the changes in a SOTA model within a multimodal video understanding task. We aim to uncover differences between tuned and untuned models in terms of layer depth, context length, and modality type. Our proposal can be seen as an interesting exploration of the interaction between multimodal NLP and cognitive neuroscience.

8 Highlight detection in live streams using audience reactions with transformer language models

Authors: Maximilian Gutsche

Presented by: Maximilian Gutsche

Abstract

With the rise of live streaming of eSports events, highlight videos have become very popular. Their manual creation is resource intensive. Thus, we explore automatic highlight creation with transformer language models. Although this task is inherently multimodal, we focus on the audience reactions in form of chat messages which inform a transformer model for highlight detection. We present approaches for understanding live stream chat specific language, linking video content to chat messages and splitting a continuous stream of messages into units for classification. In this context, we collect a corpus of live stream chat from League of Legends audiences, train a language model on it and compare fine-tuning for highlight detection with different models.

9 Automatic detection of SWR in EEG signal to study correlation of SWR with memory performance in humans. (Preliminar Results)

Author: Edu Vallejo

Presented by: Edu Vallejo

Abstract

Sharp-Wave Ripple (SWR) complexes are brain signal patterns that occur during deep-sleep stages in the hippocampal region of the brain (Buzsáki, 2015). SWR have been object of many studies due to their connection to important brain functions such as memory consolidation (i.e. the transference from medium term memory into long term memory), (Buzsáki, 2015 & Ramadan, Eschenko, and Sara, 2009). We already know that SWR correlate (and even have a causal relation) with memory consolidation performance in mice (Ramadan, Eschenko, and Sara, 2009 & Maingret et al., 2016), however there are not any conclusive results in humans yet. In this context, the Neurorhythms lab is entrusted with data from a few patients who suffer from epilepsy and are being studied because of this condition. The data consists of deep-brain EEG recordings and MRI snapshots during sleep, and performance metrics on a memory task before and after sleep. Although the amount of samples is small and biased (because they all had epilepsy), it was enough to conceive an experiment which yielded positive results indicating this might be a worthwhile avenue of research.

10 Functional Data Analysis of Non-manual Marking of Questions in Kazakh-Russian Sign Language

Author: Anna Kuznetsova, Alfarabi Imashev, Medet Mukushev, Anara Sandygulova, Vadim Kimmelman

Presented by: Anna Kuznetsova

Abstract

This paper is a continuation of Kuznetsova et al. (2021), which described non-manual markers of polar and wh-questions in comparison with statements in an NLP dataset of Kazakh-Russian Sign Language (KRSL) using Computer Vision. One of the limitations of the previous work was the distortion of the 3D face landmarks when the head was rotated. The proposed solution was to train a simple linear regression model to predict the distortion and then subtract it from the original output. We improve this technique with a multilayer perceptron. Another limitation that we intend to address in this paper is the discrete analysis of the continuous movement of non-manuals. In Kuznetsova et al. (2021) we averaged the value of the non-manual over its scope for statistical analysis. To preserve information on the shape of the movement, in this study we use a statistical tool that is often used in speech research, Functional Data Analysis, specifically Functional PCA.

11 On the Cusp of Comprehensibility: Can Models Distinguish Between Metaphors and Nonsense?

Author: Bernadeta Griciūtė

Presented by: Bernadeta Griciūtė

Abstract

Linguistic metaphors can be defined as expressions of an understanding of one concept in terms of another, where there is some similarity between the two. Phrases can vary in metaphoricity, ranging from ubiquitous conventional metaphors, which can be hard to distinguish from their literal meaning to creative metaphors, where a novel meaning emerges in a sentence. However, even when it comes to novel metaphors, language users should still be able to infer the meaning - otherwise, they are just nonsense. In the current work, I extend the traditional task of metaphor detection from the binary classification between literal and metaphorical sentences into a more fine-grained classification with four categories: literal senses, conventional metaphors, unconventional (novel) metaphors, and nonsense. For testing, a new dataset containing sentence pairs from all four categories is created.

12 Patterns of human judgments of the complexity, acceptability, and grammaticality in Italian sentences.

Author: Ioanna Tsani

Presented by: Ioanna Tsani

Abstract

Acceptability and complexity constitute useful metrics for the evaluation of any type of text, with regards to its felicity and to its level of difficulty. For the further investigation of this topic, the EVALITA 2020 dataset has been annotated with human ratings of acceptability and complexity of sentences which present interesting syntactic and semantic phenomena in Italian. The dataset includes ill-formed and well-formed sentences that showcase phenomena such as wh-extraction, negative polarity items, and wh-interrogatives with gap positions inside or outside of a VP conjunction. In this project, we are interested in patterns that such ratings might reveal, such as the correlation of grammaticality, acceptability, and negative complexity (i.e. whether more complex sentences have the tendency to be judged as ungrammatical or unacceptable), the effects of a specific phenomenon in human judgments, and whether specific phenomena allow for more variance among the annotators.

13 Cross-lingual argument mining in medical domain

Author: Anar Yeginbergenova

Presented by: Anar Yeginbergenova

Abstract

Medicine, nowadays, is receiving more and more attention, particularly in the applications involving artificial intelligence to help clinicians with decision-making in their everyday life. Argument mining is addressed as a method to help achieve this goal by detecting argumentative components in the text and classifying the relations between them. The annotated medical data of abstracts of Randomized Controlled Trials (RCT) from the MEDLINE database was designed specifically for this task. This data set consists of abstracts of 5 different diseases: neoplasm, glaucoma, diabetes, hepatitis B, and hypertension. Besides argumentation, we address one of the major problems in NLP which is the lack of datasets in languages other than English. Therefore, in this project, we automatically create a corpus in Spanish by translating the data and projecting the labels from source to target language. Then we perform cross-lingual, multilingual, and monolingual argument mining experiments using the model that was used for English data and compare the performance under each experimental setup to see how multilingual models handle multilingual tests.

14 Neural Chart Summarization

Author: Saad Obaid ul Islam

Presented by: Saad Obaid ul Islam

Abstract

Information visualizations like bar charts, line charts, and pie charts are a popular way of communicating quantitative data. They are used to get important insights and make well informed decisions. However, for some people it can be challenging to understand these charts, especially for people who are visually impaired. Automatic Chart Summarization is the task to explain and summarize the key takeaways from the chart. This task comes under the umbrella of data-to-text. Data is extracted from charts in a form of table and that table is converted to textual summary. In this research, we leverage the power of pre-trained models like T5 to generate summaries. Even though T5 is good at generating summaries, it often leads to hallucinations. To tackle this, we optimize input data representation to the model. Initial results show that the generated output depends highly on input representation. Our model performs better than existing chart summarization models by several BLEU points.

15 Improving Automatic Speech Recognition for Regionally-Accented French

Author: Rasul Jasir Dent

Presented by: Rasul Jasir Dent

Abstract

For widely-spoken European languages such as French, English, and Spanish, the accuracy of automatic speech recognition (ASR) for speakers of prestige dialects has improved tremendously in recent years, while performance for other dialects and pronunciations shows more considerable variation. In this study, we first seek to characterize the performance of a hybrid Time-Delay Neural Network system on speakers from across Francophone Europe, Africa, and Canada, using samples from the Phonologie du Français Contemporain corpus. We then explore how changes to the number of phonetic classes in the acoustic model and alternate pronunciations in the lexicon impact ASR accuracy, with a particular emphasis on those speakers deemed to have stronger accents. Preliminary results suggest that maintaining distinct phonetic classes for mid-vowels may considerably reduce WER for more accented speakers from multiple regions, and that trimming the pronunciation lexicon may provide additional gains when ambiguity is relatively low. In future experiments, we will examine the impact of adding region-specific alternate pronunciations and acoustic models.

16 Cross-lingual transfer learning with Persian

Author: Sepideh Mollanorozy

Presented by: Sepideh Mollanorozy

Abstract

Cross-lingual transfer learning is a popular approach when working with low-resource languages. In this project, we focus on transfer learning with the Persian language. We explore if Persian as a source language can be beneficial for other target languages, and which languages are a good source for Persian as a target. We use XLM-ROBERTA and ParsBERT pre-trained models and experiment by fine-tuning them for POS tagging and sentiment analysis as low-level and high-level tasks. We will discuss whether there are global rules underlying cross-lingual transfer learning or not.

17 Accent stigma of nonnative accents

Author: Allison Keith

Presented by: Allison Keith

Abstract

Adult learners of a second language frequently retain phonetic characteristics of their native language when speaking in their second language which can indicate their identity as an outsider or a member of a stigmatized group. Stigma against nonnative accents can lead to a number of negative professional, social, and emotional outcomes for both native speakers and nonnative speakers. This paper discusses how listeners perceive accents, their thoughts and ideas about accented speech, and the possible outcomes of accent stigma, as well as potential interventions that address accent discrimination.

18 Shrinking Knowledge Base Size: Dimension Reduction, Splitting & Filtering

Author: Vilém Zouhar

Presented by: Vilém Zouhar

Abstract

Recently neural network based approaches to knowledge-intensive NLP tasks, such as question answering, started to rely heavily on the combination of neural retrievers and readers. Retrieval is typically performed over a large textual knowledge base which requires significant memory and compute resources, especially when scaled up. On HotpotQA we explore various filtering & splitting criteria. Primarily, we systematically investigate reducing the size of the KB index by means of dimensionality (sparse random projections, PCA, autoencoders) and numerical precision reduction.

Our results show that PCA is an easy solution that requires very little data and is only slightly worse than autoencoders, which are less stable. All methods are sensitive to pre- and post-processing and data should always be centered and normalized both before and after dimension reduction. Finally, we show that it is possible to combine PCA with using 1bit per dimension. Overall we achieve (1) 100× compression with 75%, and (2) 24× compression with 92% original retrieval performance.

19 Investigating LLMs’ Processing of Nominal Plurality using Probing and Causal Interventions

Author: Michael Hanna

Presented by: Michael Hanna

Abstract

With the proliferation of large pretrained language models (LLMs) in NLP, methods for interpreting such models have similarly multiplied. One particularly popular interpretability technique is probing, which analyzes the linguistic features learned by LLMs by training auxiliary models to extract such features from LLMs’ internal representations. However, probing has recently been criticized for its inability to demonstrate that LLMs actually use the information that probes have found. Recent studies have proposed various causal interventions as a solution. These techniques solve the aforementioned issue by altering LLMs’ internal representations and observing how their behavior changes, thus eliciting causal connections between internal mechanisms and external production. In this paper, we bring together multiple causal intervention techniques to pinpoint the degree to which probes capture functionally relevant information encoded in LLM representations. We find that probes trained to predict noun plurality do capture information relevant to subject-verb agreement. Despite this, the magnitude of the behavioral changes they produce lags behind an upper bound given by other interventions in some scenarios.

20 Measuring readability of technical texts

Author: Anna Kriukova

Presented by: Anna Kriukova

Abstract

Readability measures try to capture how difficult it is for a text's audience to understand its content. Text clarity is especially important in educational materials that introduce and explain new concepts. In this research, I explore the ways to capture the readability of Computer Science texts from an online educational platform. I experiment with the classical readability scores to evaluate how applicable they are for the task at hand, concluding that none of them are able to differentiate between well- and poorly-written texts. I also suggest a way of customizing one of the scores, taking into consideration the specific nature of texts on the platform. Finally, I build a machine learning system that makes use of linguistic, meta-text, and statistical features, as well as those based on the knowledge graph of topics on the platform.

21 Research Review and Direction for Dialogue-based Computer Assisted Language Learning

Author: Trevor Atkins

Presented by: Trevor Atkins

Abstract

Education and language learning in particular has been significantly influenced by technology and globalization. The awareness and approaches of technological systems for language learning have grown with the pandemic as the demand for remote learning and unconventional implementations increased. For language learning, that means Computer Assisted Language Learning (CALL) would be a preferable method. As defined in the work by Levy (1997, p.1) CALL is “the search for and study of applications of the computer in language teaching and learning”. Approaches for these applications can be top-down centered around a language theory or bottom-up centered around the implementation of the computer addressing an issue. Since the conceptualization of CALL is determined mostly by hardware and software use, the design choice makes a significant influence on its creation. For this presentation, I will present both a brief overview of a research review and directions for planned or possible research with the central idea being that the computer is implemented as an interlocutor that allows for the user to practice their target language.

22 Subword-based Cross-lingual Transfer of Embeddings from Hindi to Marathi and Nepali

Author: Niyati Bafna

Presented by: Niyati Bafna

Abstract

Word embeddings are growing to be a crucial resource in the field of NLP for any language. This work introduces a novel technique for static subword embeddings transfer for Indic languages from a relatively higher resource language to a genealogically related low resource language. We primarily work with Hindi-Marathi, simulating a low-resource scenario for Marathi, and confirm observed trends on Nepali. We demonstrate the consistent benefits of unsupervised morphemic segmentation on both source and target sides over the treatment performed by fastText. Our best-performing approach uses an EM-style approach to learning bilingual subword embeddings; we also show, for the first time, that a trivial “copy-and-paste” embeddings transfer based on even perfect bilingual lexicons is inadequate in capturing language-specific relationships. We find that our approach substantially outperforms the fastText baselines for both Marathi and Nepali on the Word Similarity task as well as WordNet-Based Synonymy Tests; on the former task, its performance for Marathi is close to that of pretrained fastText embeddings that use three orders of magnitude more Marathi data.

23 Fighting hate speech by understanding implied bias and selecting proper counter-narrative

Author: Nami Akazawa

Presented by: Nami Akazawa

Abstract

Social bias exists in our day-to-day language and social media has a great danger of spreading hate against specific groups or individuals. For this project we focus on implied biases in hate speech: while many researchers had focused on the binary task of hate/non-hate detection, we want to obtain a deeper understanding of the implied meaning to select possible textual responses. To our knowledge, this is the first attempt to build an automatic selection tool of counter-narratives based on implied bias in the hateful content. In this task, the first step is to distill the underlying biases and to do so, we will utilize transformer-based sequence-to-sequence models such as BART and T5 to map the existing implied statement from the biased comments that especially target different minority groups including migrants, Muslim, Jews, POC, LGBT+, women and disabled.

24 Entity-centric Event Sequences to Highlight Perspectives

Author: Shantanu Nath, Gosse Minnena, Tommaso Caselli

Presented by: Shantanu Nath

Abstract

People talk about the same events in different ways. Narrations of the same facts can vary according to who is telling them, and the perspective they take on it. In this contribution, we use football matches to extract entity-centric event sequences to identify and highlight the presence of perspectives. We argue that football is particularly suited for this purpose because of three defining characteristics: (i) football is a clearly defined semantic domain; (ii) news reports are usually short and simple to process; (iii) it is straightforward to define perspectives, as there are always exactly two teams that compete against each other, and which may be reported on more or less favorably by different sources.

We model events by means of Fillmorean semantic frames [1]. Entity-centric event sequences are based on [2] & [3] and represent coherent sequences of sets of events sharing a single protagonist. Each team in a match is a protagonist of an event sequence. For instance, given the text (1a), we can extract the event sequence (1b) for one of the teams:

- *The Flying Dutchmen* will *advance* to the semi-finals. *Frank de Boer* is in *possession* of the ball, [he] *passes* the ball to *Dennis Bergkamp*, *Bergkamp* *picks up* the ball and *scores*.
- PROGRESSION (advance) \Rightarrow POSSESSION (possession) \Rightarrow PASS (passes) \Rightarrow CONTROL (pick up) \Rightarrow SCORE_GOAL (scores)

To obtain the event sequence in (1b), we first extract events using an end-to-end frame semantic parser [4] trained on Kicktionary [5], a multilingual (German, French, English) football domain-specific framenet resource. Minnema (2021) showed that despite limited training material,

the model achieves good performance (0.83 F1-score on frame detection, and 0.81 F1-score on semantic role detection given gold predicates) on a held-out test set. Since the model is based on XLM-R [6], it can also be applied in a zero-shot cross-lingual transfer setting. Once frames and roles have been identified, protagonists must be aggregated at the team level. In our case, all mentions of individuals (“Frank de Boer”, “he”, “Dennis Bergkamp”, “Bergkamp”) as well as mentions of the targeted team (“The Flying Dutchmen”) are abstracted to a single protagonist representing the team. In this stage, the temporal order of events is assumed to be iconic with the order of mention in the text, although it may not always be the case. By comparing sequences of events for each of the teams, we can highlight different perspectives taken on the same match by different texts. For example, in (1b), the perspective is clearly on the Dutch team: we have a contiguous sequence of 5 events in which only the Dutch team is mentioned, and always in an Agent or Beneficiary role. Currently, we are applying this approach on a large multilingual corpus (Dutch, German, and English) of football matches composed of 312 documents from different sources and languages, allowing us to extensively investigate perspectives associated with entity-centric event sequences.

References

- [1] C.J. Fillmore. Frame semantics. In *Geeraerts, D. (ed.), Cognitive Linguistics: Basic Readings*, 2006.
- [2] Nathanael Chambers and Dan Jurafsky. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, pages 789–797, Columbus, Ohio, June 2008. Association for Computational Linguistics.
- [3] Nathanael Chambers and Dan Jurafsky. Unsupervised learning of narrative schemas and their participants. In *ACL*, 2009.
- [4] Patrick Xia, Guanghui Qin, Siddharth Vashishtha, Yunmo Chen, Tongfei Chen, Chandler May, Craig Harman, Kyle Rawlins, Aaron Steven White, and Benjamin Van Durme. LOME: Large on-

tology multilingual extraction. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, Online, April 2021. Association for Computational Linguistics.

- [5] Thomas Schmidt. *The Kicktionary revisited*, pages 239–252. De Gruyter Mouton, 2008.
- [6] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online, July 2020. Association for Computational Linguistics.

25 Towards an Analysis of Ordering Restrictions within Noun Phrases: A Top-Down Minimalist Approach.

Author: Daniel Gutiérrez

Presented by: Daniel Gutiérrez

Abstract

Greenberg's Universal 20 (U20) puzzle for unattested orders of {Dem, Num, A, N} has been a difficult one to solve in syntax. On this, Cinque [7] shows that there are 24 possible combinations, of which only 14 are attested. To derive (only) attested orders, he proposes a strict, universal merging order (Dem > Num > A > N), and postulates conditions on phrasal movement. Cinque [8] also extends the analysis to adjectives and accounts for their relative order, especially in terms of pre- and post-nominal positions. However, the precise derivational mechanisms that generate cartographic structures must be postulated. Here we argue that Cinque's insights on U20 and adjectival syntax still do not account for some fine-grained, agreement-related distributional phenomena, and we explore the possibility to match structures attested in language by means of a top-down, left to right derivation [9, 10], paired with modifications on the standard mechanism of agreement [11] that would allow the correct generation of phenomena such as asymmetric agreement and other agreement-related distributional effects.

References

- [7] Guglielmo Cinque. Deriving Greenberg's universal 20 and its exceptions. *Linguistic Inquiry*, 36(3), 2005.
- [8] Guglielmo Cinque. *The Syntax of Adjectives: A comparative study*. Linguistic Inquiry Monographs 57. The MIT Press, Cambridge, MA, 2010.

- [9] Cristiano Chesi. *Phases and Cartography in Linguistic Computation: Toward a cognitively motivated computational model of linguistic competence*. PhD thesis, Università degli Studi di Siena, Siena, 2004.
- [10] Cristiano Chesi. Expectation-based minimalist grammars. *arXiv preprint arXiv:2109.13871*, 2021.
- [11] Kunio Kinjo. A probe-goal approach to impoverished agreement. *Proceedings of Seoul International Circle of Generative Grammar*, 17:225, 2015.