L98: Introduction to Computational Semantics

Lecture 10: Meaning Representations in Natural Language Generation

Weiwei Sun and Simone Teufel

Natural Language and Information Processing Research Group
Department of Computer Science and Technology
University of Cambridge

Lent 2021/22

> *[...] you can get away with incomplete semantics when you are doing parsing, but when you're doing generation, you have to specify everything in semantics. And we don't know how to do that. At least we don't know how to do that completely or properly.*
>
> *Mark Steedman*
>
> ACL lifetime archievement award lecture (`vimeo.com/288152682`)

Lecture 10: Meaning Representations in Natural Language Generation

1. Several NLG tasks
2. Parsing a Graph
3. The constructivist hypothesis
4. Small clause

Several NLG Tasks

Example 1: Question generation

- X. Du, J. Shao and C. Cardie. *Learning to Ask: Neural Question Generation for Reading Comprehension*.

Sentence:

Oxygen is used in cellular respiration and released by photosynthesis, which uses the energy of sunlight to produce oxygen from water.

Questions:

– What life process produces oxygen in the presence of light?

photosynthesis

– Photosynthesis uses which energy to form oxygen from water?

sunlight

– From what does photosynthesis get oxygen?

water

Example 2: Knowledge graph to string generation

- R. Koncel-Kedziorski, D. Bekal, Y. Luan, M. Lapata, and H. Hajishirzi. *Text Generation from Knowledge Graphs with Graph Transformers*.

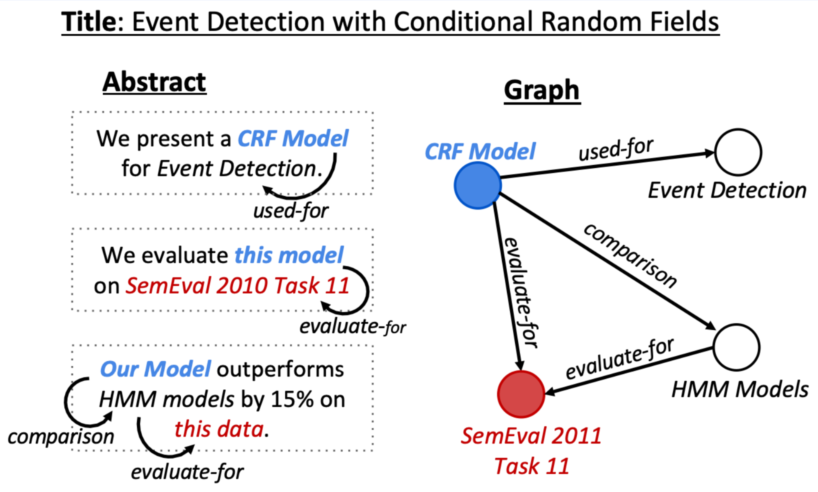

Alchemy

a type of chemistry, especially in the Middle Ages, that dealt with trying to find a way to change ordinary metals into gold and with trying to find a medicine that would cure any disease

https://dictionary.cambridge.org

Surface realisation

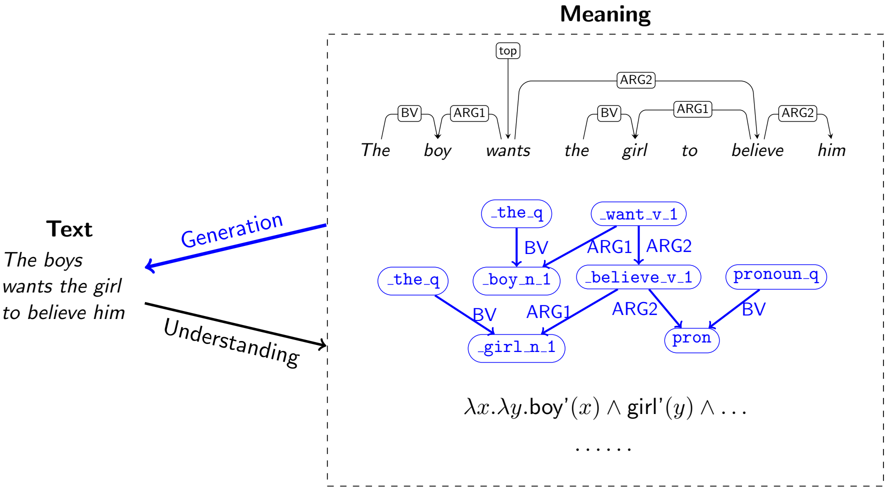

The task

Map meaning representations to sentences.

A very brief history

GraphBank Symbolic + statistical/neural Neural end-to-end

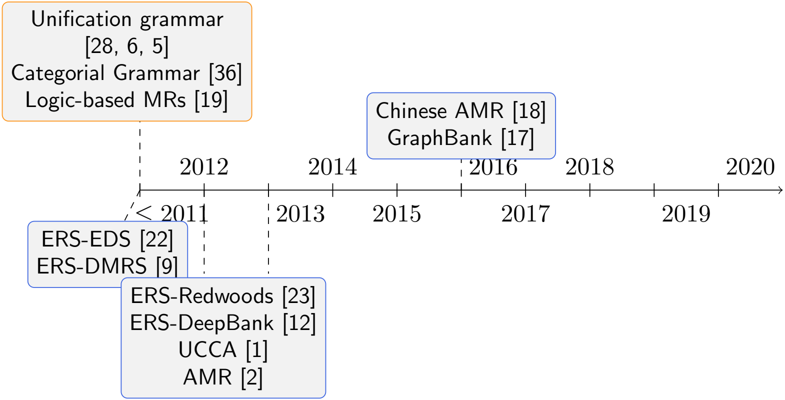

Based on symbolic system:

- Theoretical: modeling syntax–semantics correspondence
- Empirical: building comprehensive rules, improving efficiency. . .

A very brief history

GraphBank   Symbolic + statistical/neural   Neural end-to-end

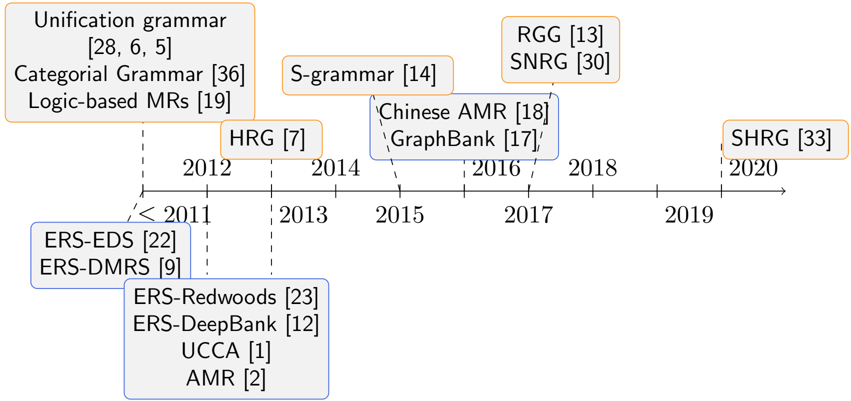

Based on symbolic system:

- Theoretical: modeling syntax–semantics correspondence
- Empirical: building comprehensive rules, improving efficiency. . .

A very brief history

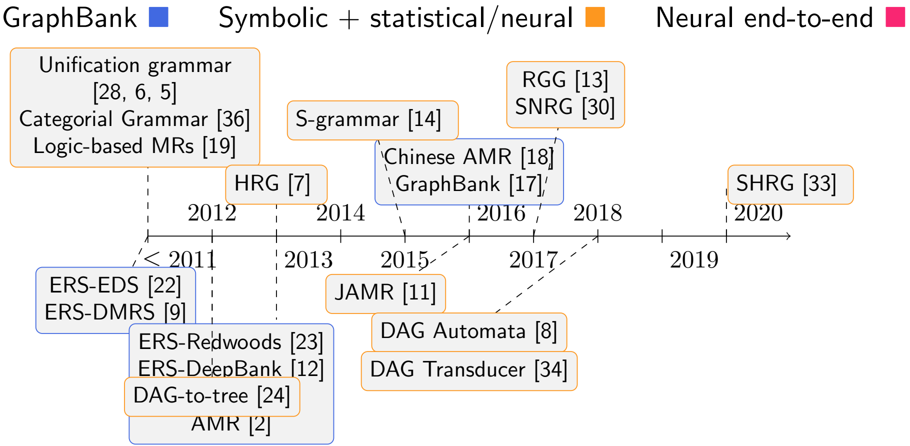

Based on symbolic system:

- Theoretical: modeling syntax–semantics correspondence
- Empirical: building comprehensive rules, improving efficiency. . .

A very brief history

GraphBank Symbolic + statistical/neural Neural end-to-end

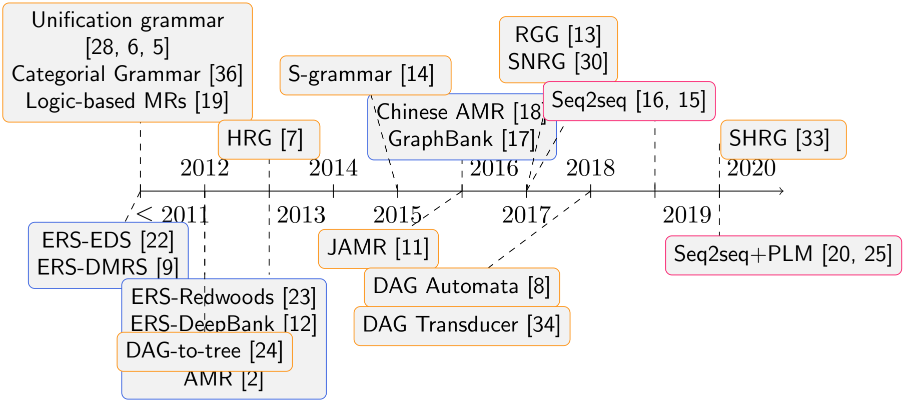

Neural end-to-end:

- Encoder: encoding the structures of MR
- Decoder: generating strings

A very brief history

GraphBank   Symbolic + statistical/neural   Neural end-to-end

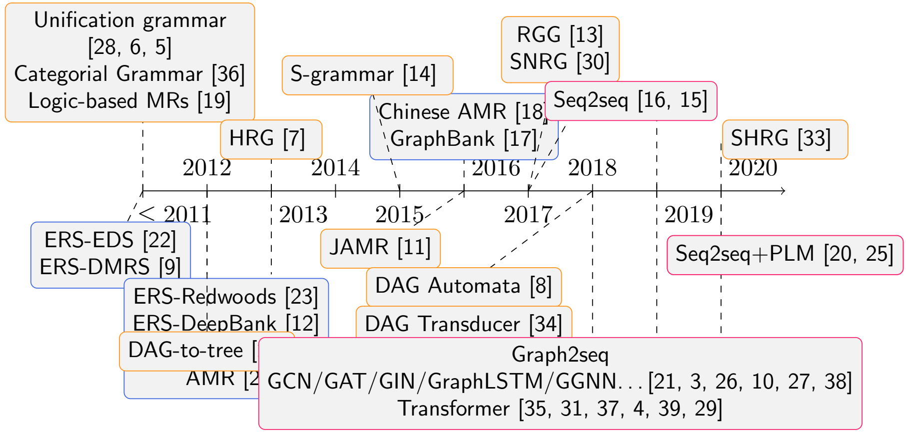

Neural end-to-end:

- Encoder: encoding the structures of MR
- Decoder: generating strings

Semantic relation to syntactic relation

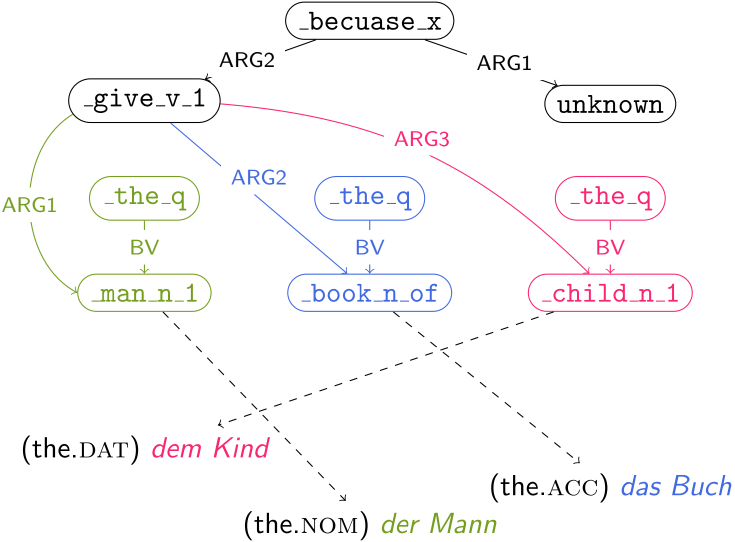

Semantic relation to syntactic construction

The Principle of Compositionality

The meaning of an expression is a function of **the meanings of its parts** *and of* **the way they are syntactically combined**.

B. Partee

Semantic relation to syntactic construction

The Principle of Compositionality

The meaning of an expression is a function of **the meanings of its parts** *and of* **the way they are syntactically combined**.

B. Partee

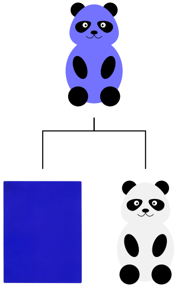

Semantic relation to syntactic construction

The Principle of Compositionality

> *The meaning of an expression is a function of* **the meanings of its parts** *and of* **the way they are syntactically combined**.
>
> *B. Partee*

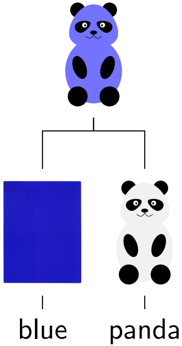

Parsing a Semantic Graph

Recursively decompose a graph

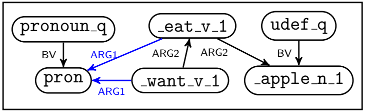

Recursively decompose a graph

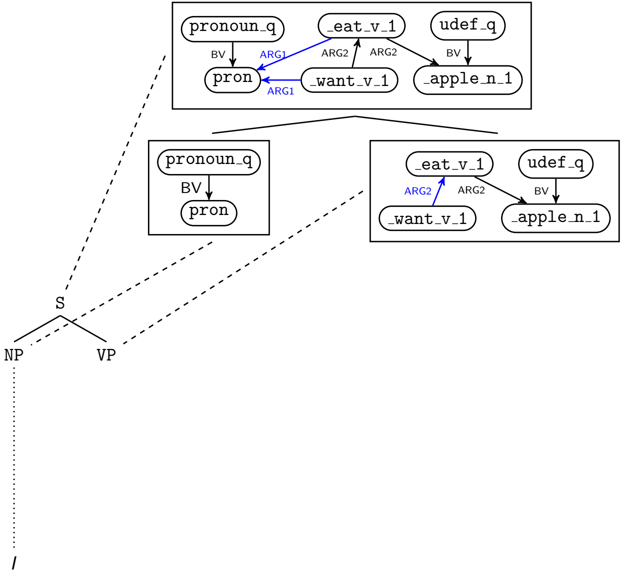

Recursively decompose a graph

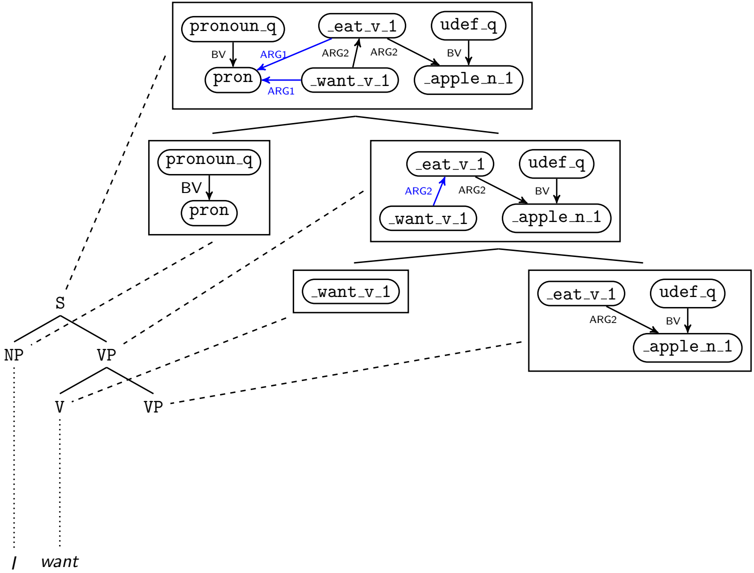

Recursively decompose a graph

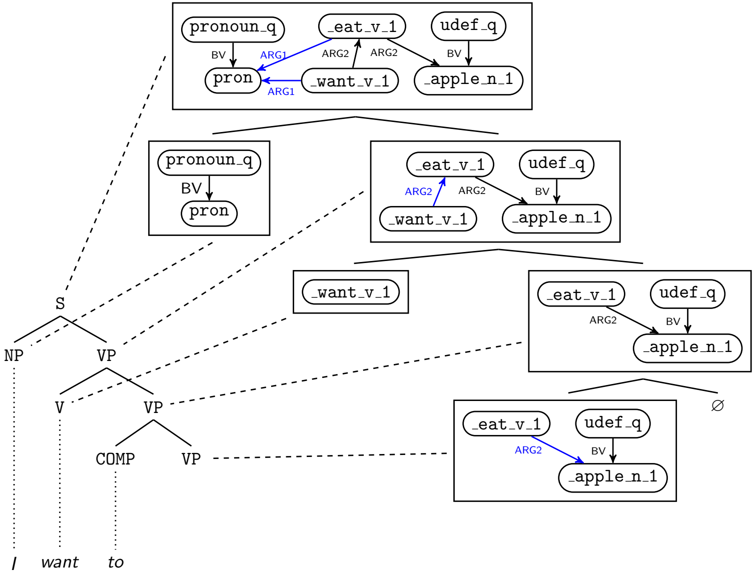

Recursively decompose a graph

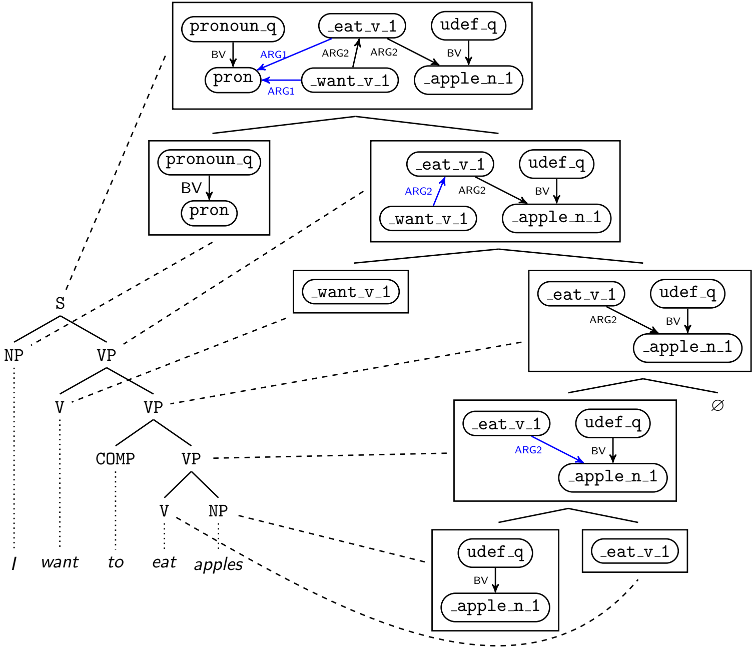

Graph parsing with an HRG

A dynamic programming algorithm [7]

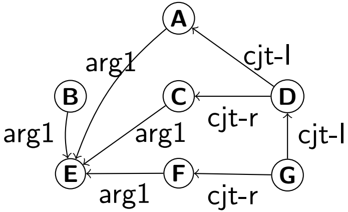

Graph parsing with an HRG

A dynamic programming algorithm [7]

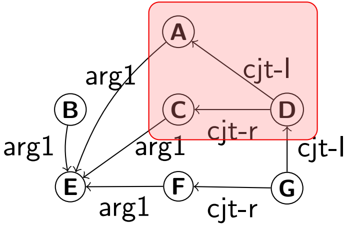

Graph parsing with an HRG

A dynamic programming algorithm [7]

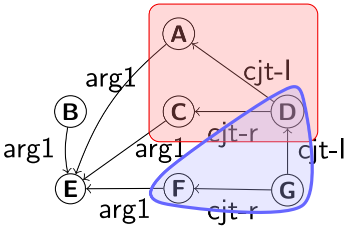

Graph parsing with an HRG

A dynamic programming algorithm [7]

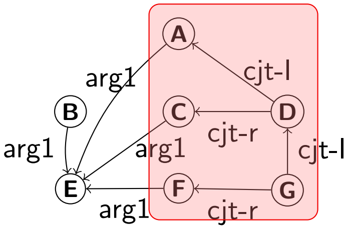

Graph parsing with an HRG

A dynamic programming algorithm [7]

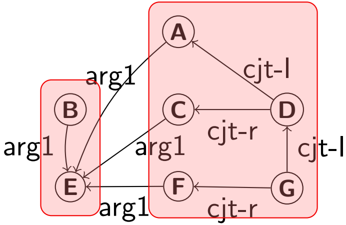

Graph parsing with an HRG

A dynamic programming algorithm [7]

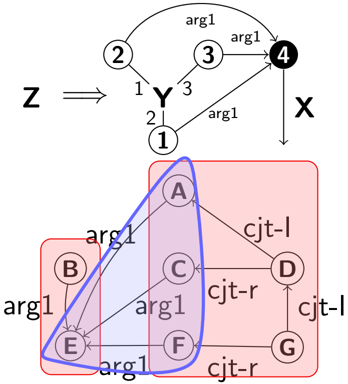

Graph parsing with an HRG

A dynamic programming algorithm [7]

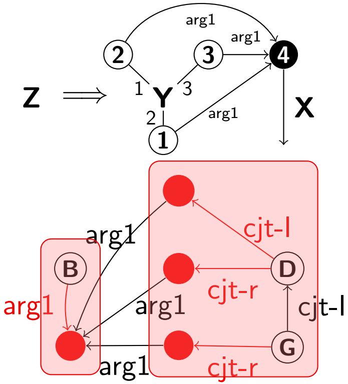

Graph parsing with an HRG

A dynamic programming algorithm [7]

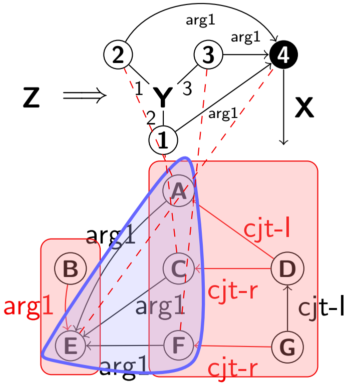

Graph parsing with an HRG

A dynamic programming algorithm [7]

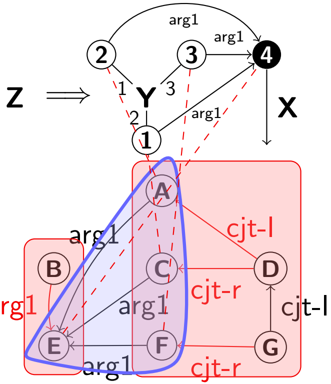

Graph-to-tree parsing

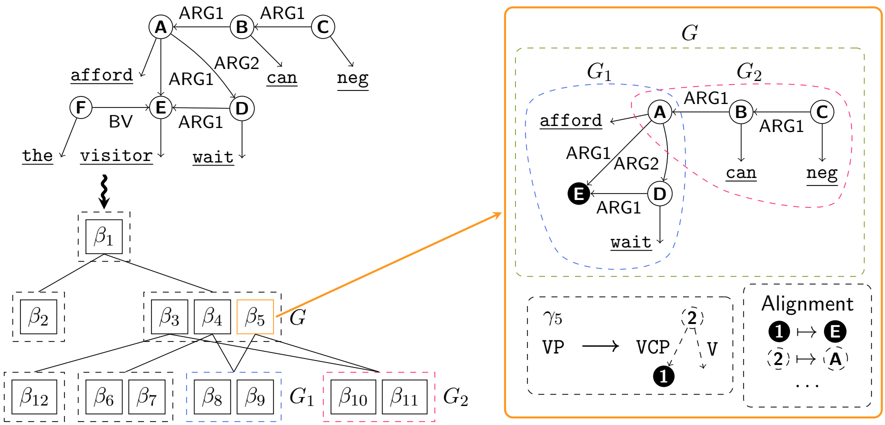

Scoring a derivation tree step-by-step

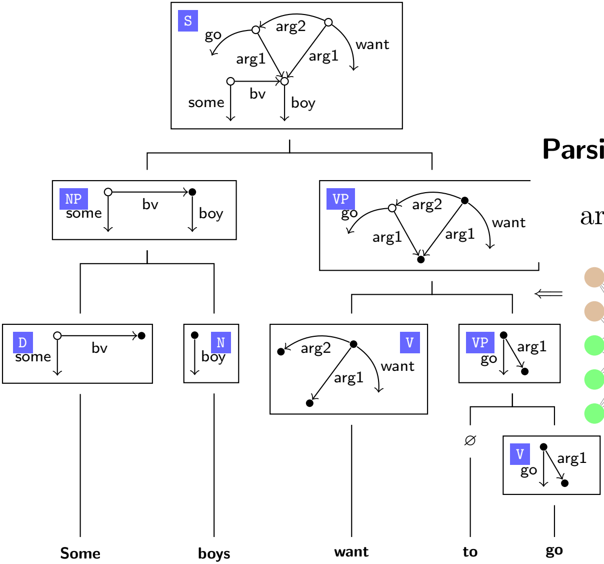

Parsing Semantic Graphs

$$\arg\max_{T \in \mathcal{T}(G)} \mathrm{SCORE}(T)$$

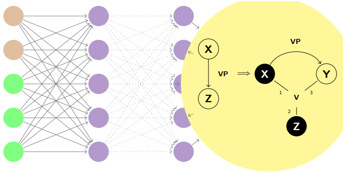

The Constructivist Hypothesis

Locality as terminal edge–adjacency

The Principle of Adjacency [32]

Combinatory rules may only apply to finitely many phonologically realized and string-adjacent entities.

A graph-based view of string-adjacency

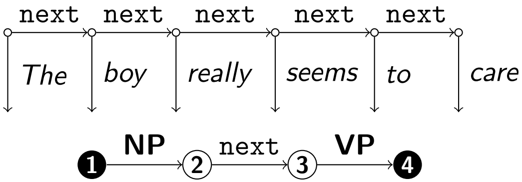

Locality as terminal edge–adjacency

The Principle of Adjacency [32]

Combinatory rules may only apply to finitely many phonologically realized and string-adjacent entities.

A graph-based view of string-adjacency

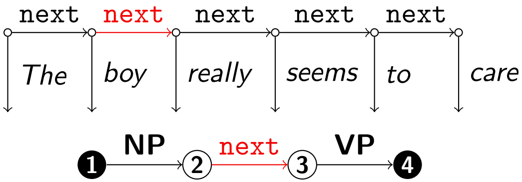

Locality as terminal edge–adjacency

The Principle of Adjacency [32]

> *Combinatory rules may only apply to finitely many phonologically realized and string-adjacent entities.*

A graph-based view of string-adjacency

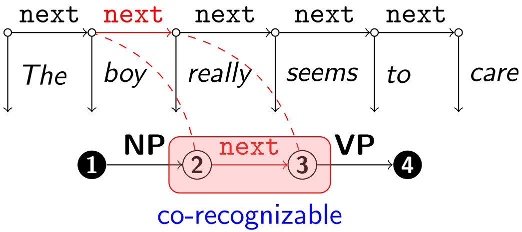

Locality as terminal edge–adjacency

In a to-be-recognized subgraph which consists of only terminal edges, if one node is identified in an input graph, the possible positions of the other nodes are highly restricted.

Example

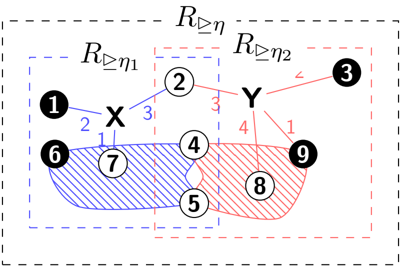

Shadowed areas mean subgraphs which consists of only terminal edges. If ④ is identified, then the cost to recognize ⑤ ❻ ❾ is highly restricted.

Distributed Argument-Structure

Lexicalized Grammar

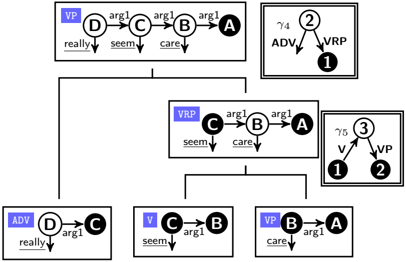

Construction Grammar

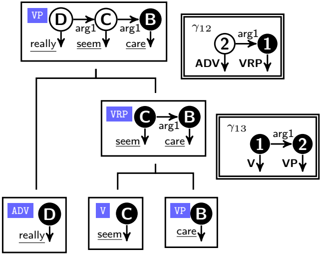

Lexicalism vs. Constructivism

- Lexicalist approaches were dominant in theoretical linguistics.
- Lexicalist approaches are dominant in computational linguistics: HPSG, LFG, CCG, . . .
- HRG-based parsing (as a derivational model) favors the constructivist approach. Parsing a graph can be as fast as parsing a string.
- Roughly speaking, we lexicalise concepts and constructionalise relations of concepts.

Small Clause

Small Clause

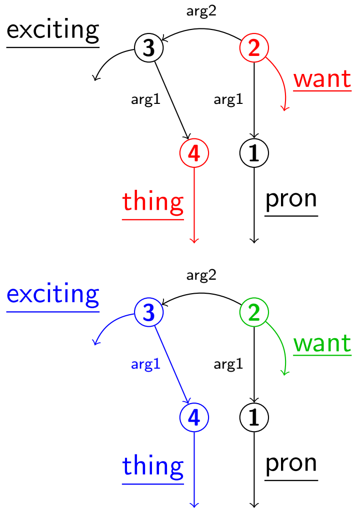

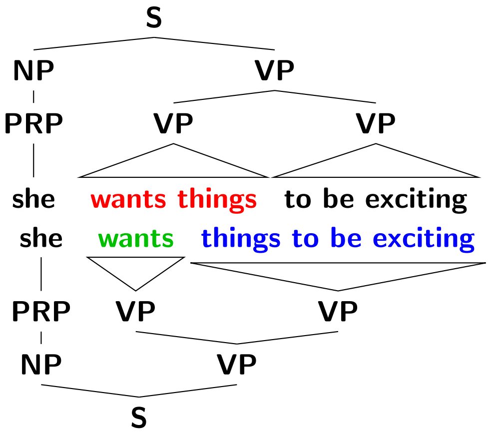

References I

O. Abend and A. Rappoport.
Universal conceptual cognitive annotation (UCCA).
In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 228–238, Sofia, Bulgaria, August 2013. Association for Computational Linguistics.

L. Banarescu, C. Bonial, S. Cai, M. Georgescu, K. Griffitt, U. Hermjakob, K. Knight, P. Koehn, M. Palmer, and N. Schneider.
Abstract Meaning Representation for Sembanking.
In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria, August 2013. Association for Computational Linguistics.

D. Beck, G. Haffari, and T. Cohn.
Graph-to-sequence learning using gated graph neural networks.
In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 273–283, Melbourne, Australia, July 2018. Association for Computational Linguistics.

References II

D. Cai and W. Lam.
Graph Transformer for Graph-to-Sequence Learning.
Proceedings of the AAAI Conference on Artificial Intelligence, 34(05):7464–7471, Apr. 2020.

J. Carroll, A. Copestake, D. Flickinger, and V. Nski.
An efficient chart generator for (semi-)lexicalist grammars.
09 2001.

J. Carroll and S. Oepen.
High efficiency realization for a wide-coverage unification grammar.
In *Second International Joint Conference on Natural Language Processing: Full Papers*, 2005.

D. Chiang, J. Andreas, D. Bauer, K. M. Hermann, B. Jones, and K. Knight.
Parsing graphs with Hyperedge Replacement Grammars.
In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 924–932, Sofia, Bulgaria, August 2013. Association for Computational Linguistics.

References III

D. Chiang, F. Drewes, D. Gildea, A. Lopez, and G. Satta.
Weighted dag automata for semantic graphs.
Computational Linguistics, 44(1):119–186, 2018.

A. Copestake.
Invited Talk: slacker semantics: Why superficiality, dependency and avoidance of commitment can be the right way to go.
In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009)*, pages 1–9, Athens, Greece, March 2009. Association for Computational Linguistics.

M. Damonte and S. B. Cohen.
Structural neural encoders for AMR-to-text generation.
In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3649–3658, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.

References IV

J. Flanigan, C. Dyer, N. A. Smith, and J. Carbonell.
Generation from Abstract Meaning Representation using tree transducers.
In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 731–739, San Diego, California, June 2016. Association for Computational Linguistics.

D. Flickinger, Y. Zhang, and V. Kordoni.
Deepbank: A dynamically annotated treebank of the wall street journal.
In *Proceedings of the Eleventh International Workshop on Treebanks and Linguistic Theories*, pages 85–96, 2012.

S. Gilroy, A. Lopez, and S. Maneth.
Parsing graphs with regular graph grammars.
In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (* SEM 2017)*, pages 199–208, 2017.

References V

J. Groschwitz, A. Koller, and C. Teichmann.
Graph parsing with s-graph grammars.
In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1481–1490, 2015.

V. Hajdik, J. Buys, M. W. Goodman, and E. M. Bender.
Neural text generation from rich semantic representations.
In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2259–2266, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.

I. Konstas, S. Iyer, M. Yatskar, Y. Choi, and L. Zettlemoyer.
Neural amr: Sequence-to-sequence models for parsing and generation.
In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 146–157, Vancouver, Canada, July 2017. Association for Computational Linguistics.

References VI

M. Kuhlmann and S. Oepen.
Towards a catalogue of linguistic graph banks.
Computational Linguistics, 42(4):819–827, 2016.

B. Li, Y. Wen, W. QU, L. Bu, and N. Xue.
Annotating the little prince with chinese amrs.
In *Proceedings of the 10th Linguistic Annotation Workshop held in conjunction with ACL 2016 (LAW-X 2016)*, pages 7–15. Association for Computational Linguistics, 2016.

W. Lu and H. T. Ng.
A probabilistic forest-to-string model for language generation from typed lambda calculus expressions.
In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1611–1622, Edinburgh, Scotland, UK., July 2011. Association for Computational Linguistics.

References VII

M. Mager, R. Fernandez Astudillo, T. Naseem, M. A. Sultan, Y.-S. Lee, R. Florian, and S. Roukos.
GPT-too: A language-model-first approach for AMR-to-text generation.
In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1846–1852, Online, July 2020. Association for Computational Linguistics.

D. Marcheggiani and L. Perez-Beltrachini.
Deep graph convolutional encoders for structured data to text generation.
In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 1–9, Tilburg University, The Netherlands, Nov. 2018. Association for Computational Linguistics.

S. Oepen and J. T. Lønning.
Discriminant-based mrs banking.
In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC-2006)*, Genoa, Italy, May 2006. European Language Resources Association (ELRA).
ACL Anthology Identifier: L06-1214.

References VIII

S. Oepen, K. Toutanova, S. Shieber, C. Manning, D. Flickinger, and T. Brants.
The lingo redwoods treebank motivation and preliminary applications.
In *Proceedings of the 19th International Conference on Computational Linguistics - Volume 2*, COLING '02, pages 1–5, Stroudsburg, PA, USA, 2002. Association for Computational Linguistics.

D. Quernheim and K. Knight.
Towards probabilistic acceptors and transducers for feature structures.
In *Proceedings of the Sixth Workshop on Syntax, Semantics and Structure in Statistical Translation*, SSST-6 '12, pages 76–85, Stroudsburg, PA, USA, 2012. Association for Computational Linguistics.

L. F. Ribeiro, M. Schmitt, H. Schütze, and I. Gurevych.
Investigating pretrained language models for graph-to-text generation.
arXiv preprint arXiv:2007.08426, 2020.

References IX

L. F. R. Ribeiro, C. Gardent, and I. Gurevych.
Enhancing AMR-to-text generation with dual graph representations.
In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3183–3194, Hong Kong, China, Nov. 2019. Association for Computational Linguistics.

L. F. R. Ribeiro, Y. Zhang, C. Gardent, and I. Gurevych.
Modeling global and local node contexts for text generation from knowledge graphs.
Transactions of the Association for Computational Linguistics, 8:589–604, 2020.

G. Russell, J. Carroll, and S. Warwick.
Asymmetry in parsing and generating with unification grammars: Case studies from ELU.
In *28th Annual Meeting of the Association for Computational Linguistics*, pages 205–211, Pittsburgh, Pennsylvania, USA, June 1990. Association for Computational Linguistics.

References X

M. Schmitt, L. F. R. Ribeiro, P. Dufter, I. Gurevych, and H. Sch ü tze.
Modeling graph structure via relative position for text generation from knowledge graphs, 2020.

L. Song, X. Peng, Y. Zhang, Z. Wang, and D. Gildea.
Amr-to-text generation with synchronous node replacement grammar.
In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 7–13, Vancouver, Canada, July 2017. Association for Computational Linguistics.

L. Song, A. Wang, J. Su, Y. Zhang, K. Xu, Y. Ge, and D. Yu.
Structural information preserving for graph-to-text generation.
In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7987–7998, Online, July 2020. Association for Computational Linguistics.

M. Steedman.
The syntactic process.
MIT Press, Cambridge, MA, USA, 2000.

References XI

Yajie Ye and W. Sun.
Exact yet efficient graph parsing, bi-directional locality and the constructivist hypothesis.
In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4100–4110, Online, July 2020. Association for Computational Linguistics.

Yajie Ye, W. Sun, and X. Wan.
In *2021 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 2021.

T. Wang, X. Wan, and H. Jin.
Amr-to-text generation with graph transformer.
Transactions of the Association for Computational Linguistics, 8:19–33, 2020.

M. White, R. Rajkumar, and S. Martin.
Towards broad coverage surface realization with ccg.
In *Proc. of the Workshop on Using Corpora for NLG: Language Generation and Machine Translation (UCNLG+ MT)*, 2007.

References XII

S. Yao, T. Wang, and X. Wan.
Heterogeneous graph transformer for graph-to-sequence learning.
In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7145–7154, Online, July 2020. Association for Computational Linguistics.

C. Zhao, M. Walker, and S. Chaturvedi.
Bridging the structural gap between encoding and decoding for data-to-text generation.
In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2481–2491, Online, July 2020. Association for Computational Linguistics.

J. Zhu, J. Li, M. Zhu, L. Qian, M. Zhang, and G. Zhou.
Modeling graph structure in transformer for better AMR-to-text generation.
In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5459–5468, Hong Kong, China, Nov. 2019. Association for Computational Linguistics.