A structured syntax-semantics interface for English-AMR alignment

Ida Szubert **Adam Lopez** **Nathan Schneider**

Edinburgh NLP
University of Edinburgh
Natural Language Processing

nert
Georgetown University

Abstract Meaning Representation (AMR)

Broad-coverage scheme for **scalable human annotation of English sentences** [Banarescu et al., 2013]

- Unified, readable graph representation
- “Semantics from scratch”: annotation does not use/specify syntax or align words
- **60k sentences** gold-annotated

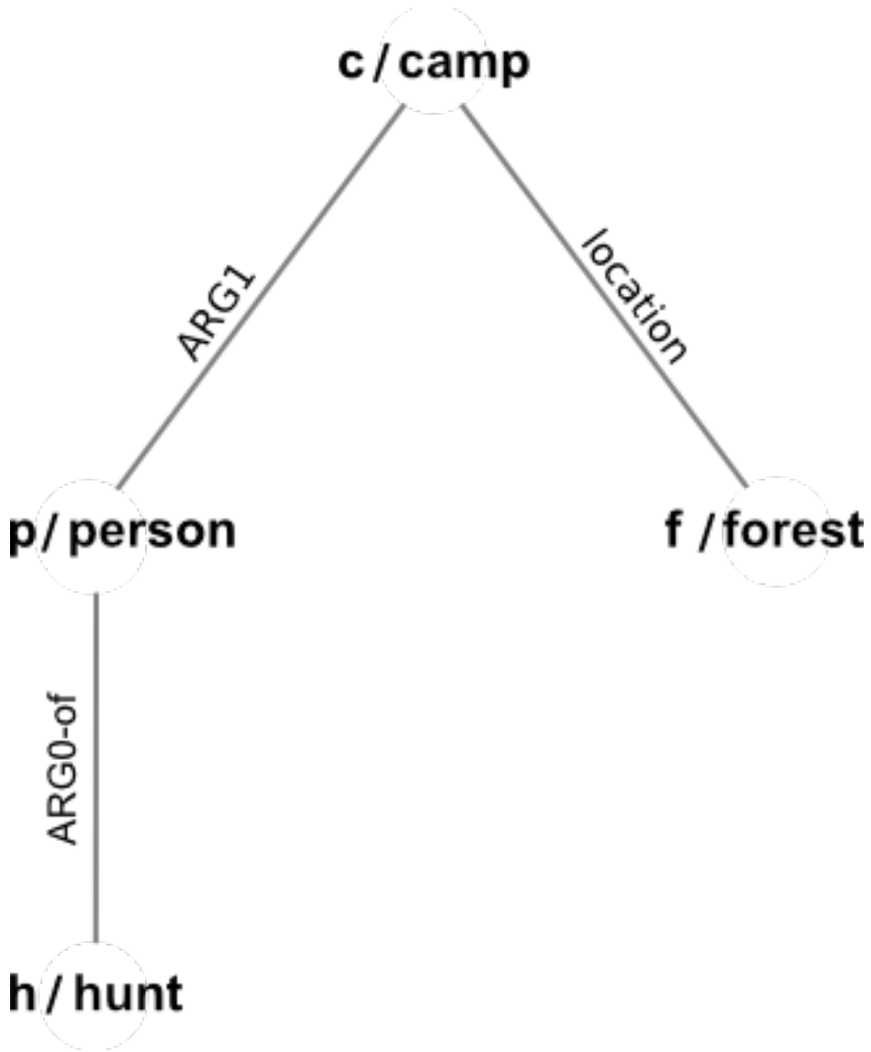

The hunters camp in the forest

Abstract Meaning Representation (AMR)

Broad-coverage scheme for **scalable human annotation of English sentences** [Banarescu et al., 2013]

- Unified, readable graph representation
- “Semantics from scratch”: annotation does not use/specify syntax or align words
- **60k sentences** gold-annotated

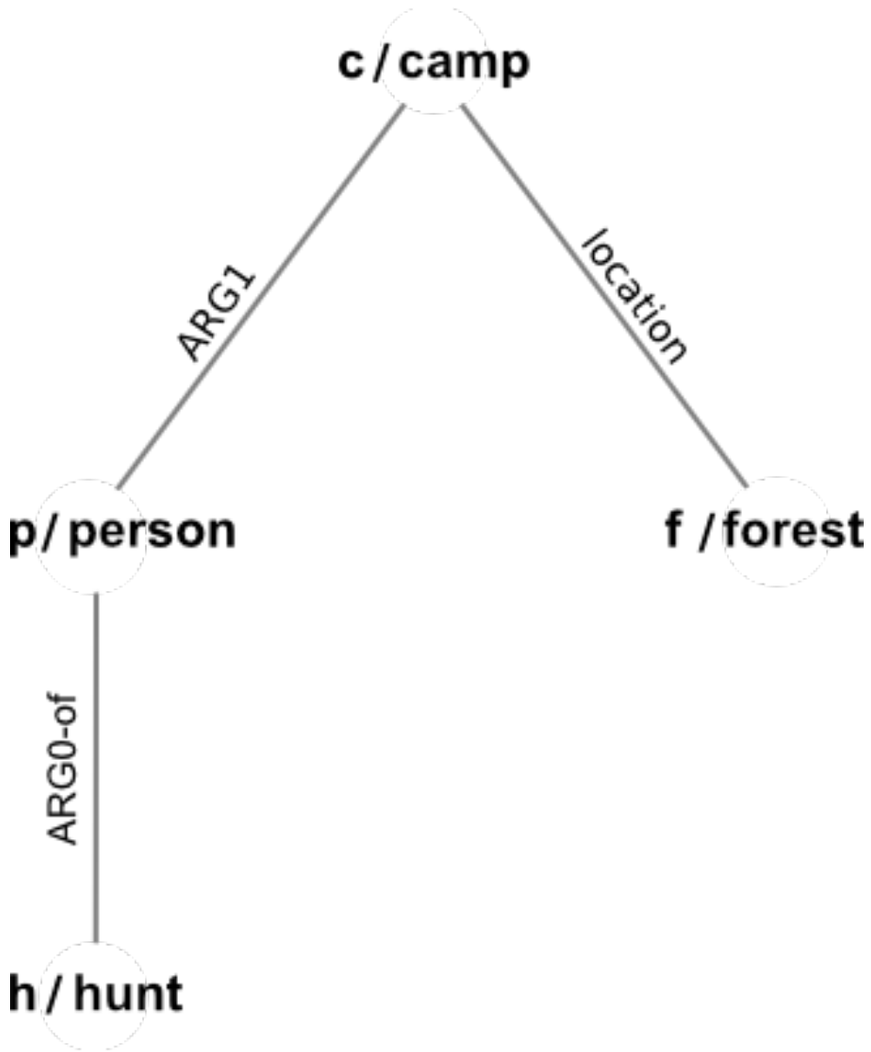

The hunters camp in the forest

AMR in NLP

- Most approaches to AMR parsing/ generation require explicit **alignments** in the training data to learn generalizations [Flanigan et al., 2014; Wang et al., 2015; Artzi et al., 2015; Flanigan et al., 2016; Pourdamghani et al., 2016; Misra and Artzi, 2016; Damonte et al., 2017; Peng et al., 2017; …]

- 2 main alignment flavors/datasets & systems:
  - JAMR [Flanigan et al., 2014]
  - ISI [Pourdamghani et al., 2014]

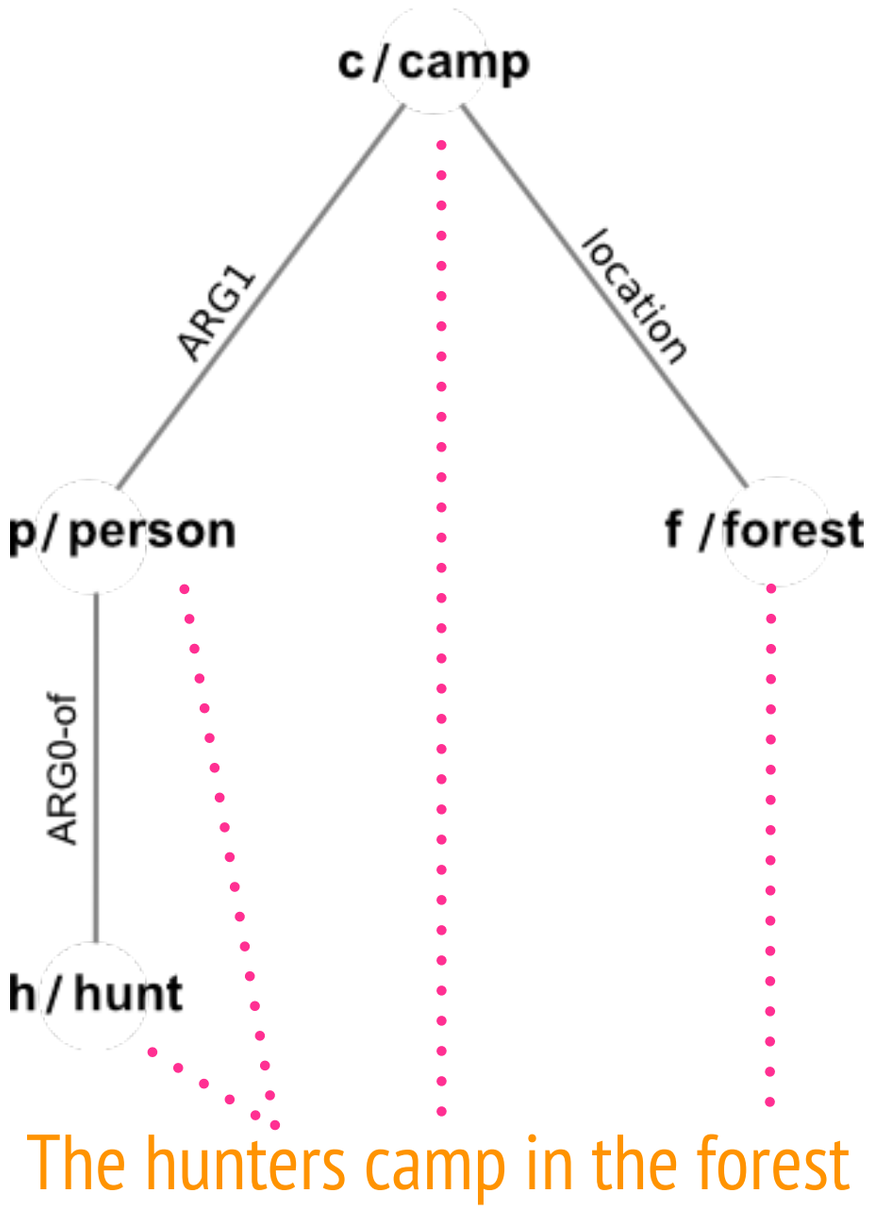

Reactions to Current AMR Alignments

"Wrong alignments between the word tokens in the sentence and the concepts in the AMR graph account for a significant proportion of our AMR parsing errors" [Wang et al., 2015]

"Improvements in the quality of the alignment in training data would improve parsing results." [Foland & Martin, 2017]

"More accurate alignments are therefore crucial in order to achieve better parsing results." [Damonte & Cohen, 2018—4:24 in Empire B!]

"A standard semantics and annotation guideline for AMR alignment is left for future work" [Werling et al., 2015]

This Talk: UD 💖 AMR

- ✓ A **new, more expressive flavor** of AMR alignment that captures the syntax–semantics interface
  - ▸ UD parse nodes and subgraphs ↔ AMR nodes and subgraphs
  - ▸ Annotation guidelines, new dataset of 200 hand-aligned sentences
- ✓ Quantify coverage and similarity of AMR to dependency syntax (97% of AMR aligns)
- ✓ Baseline algorithms for lexical (node–node) and structural (subgraph) alignment

(String, AMR) alignments

The hunters camp in the forest

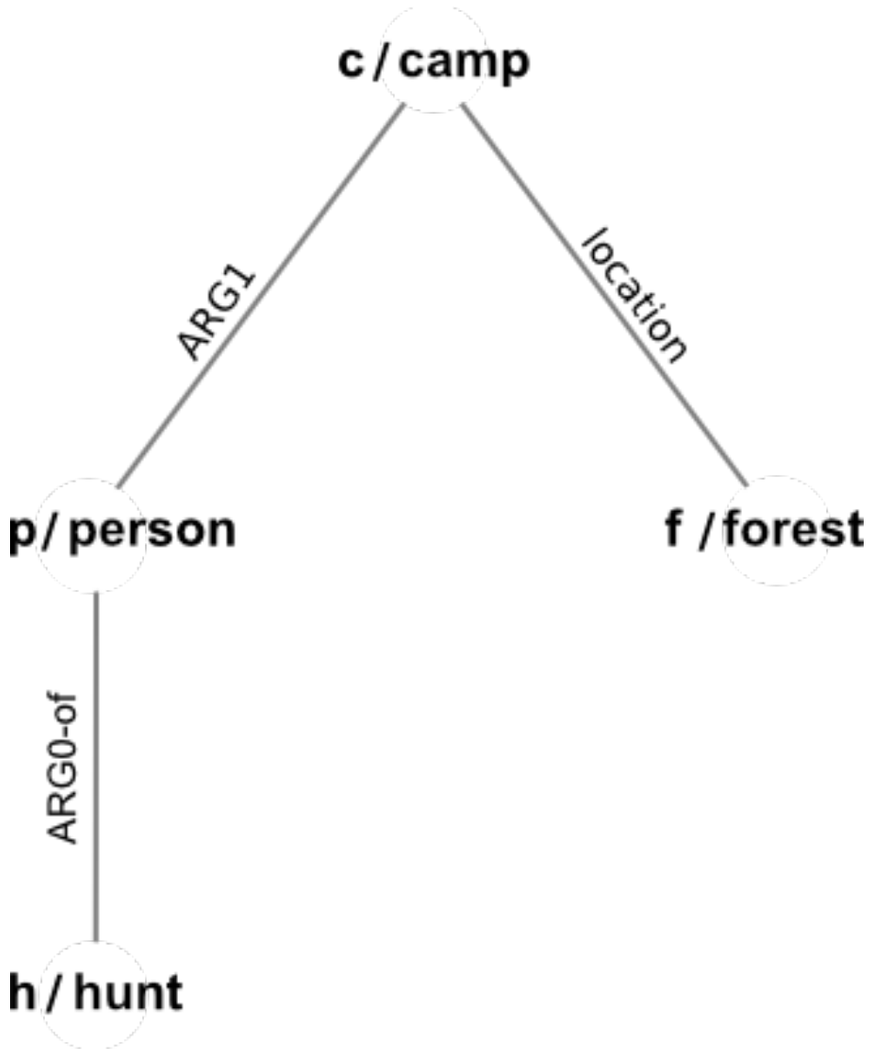

JAMR-style [Flanigan et al., 2014]

- (Word span, AMR node), (Word span, Connected AMR subgraph) alignments
- each AMR node is in 0 or 1 alignments

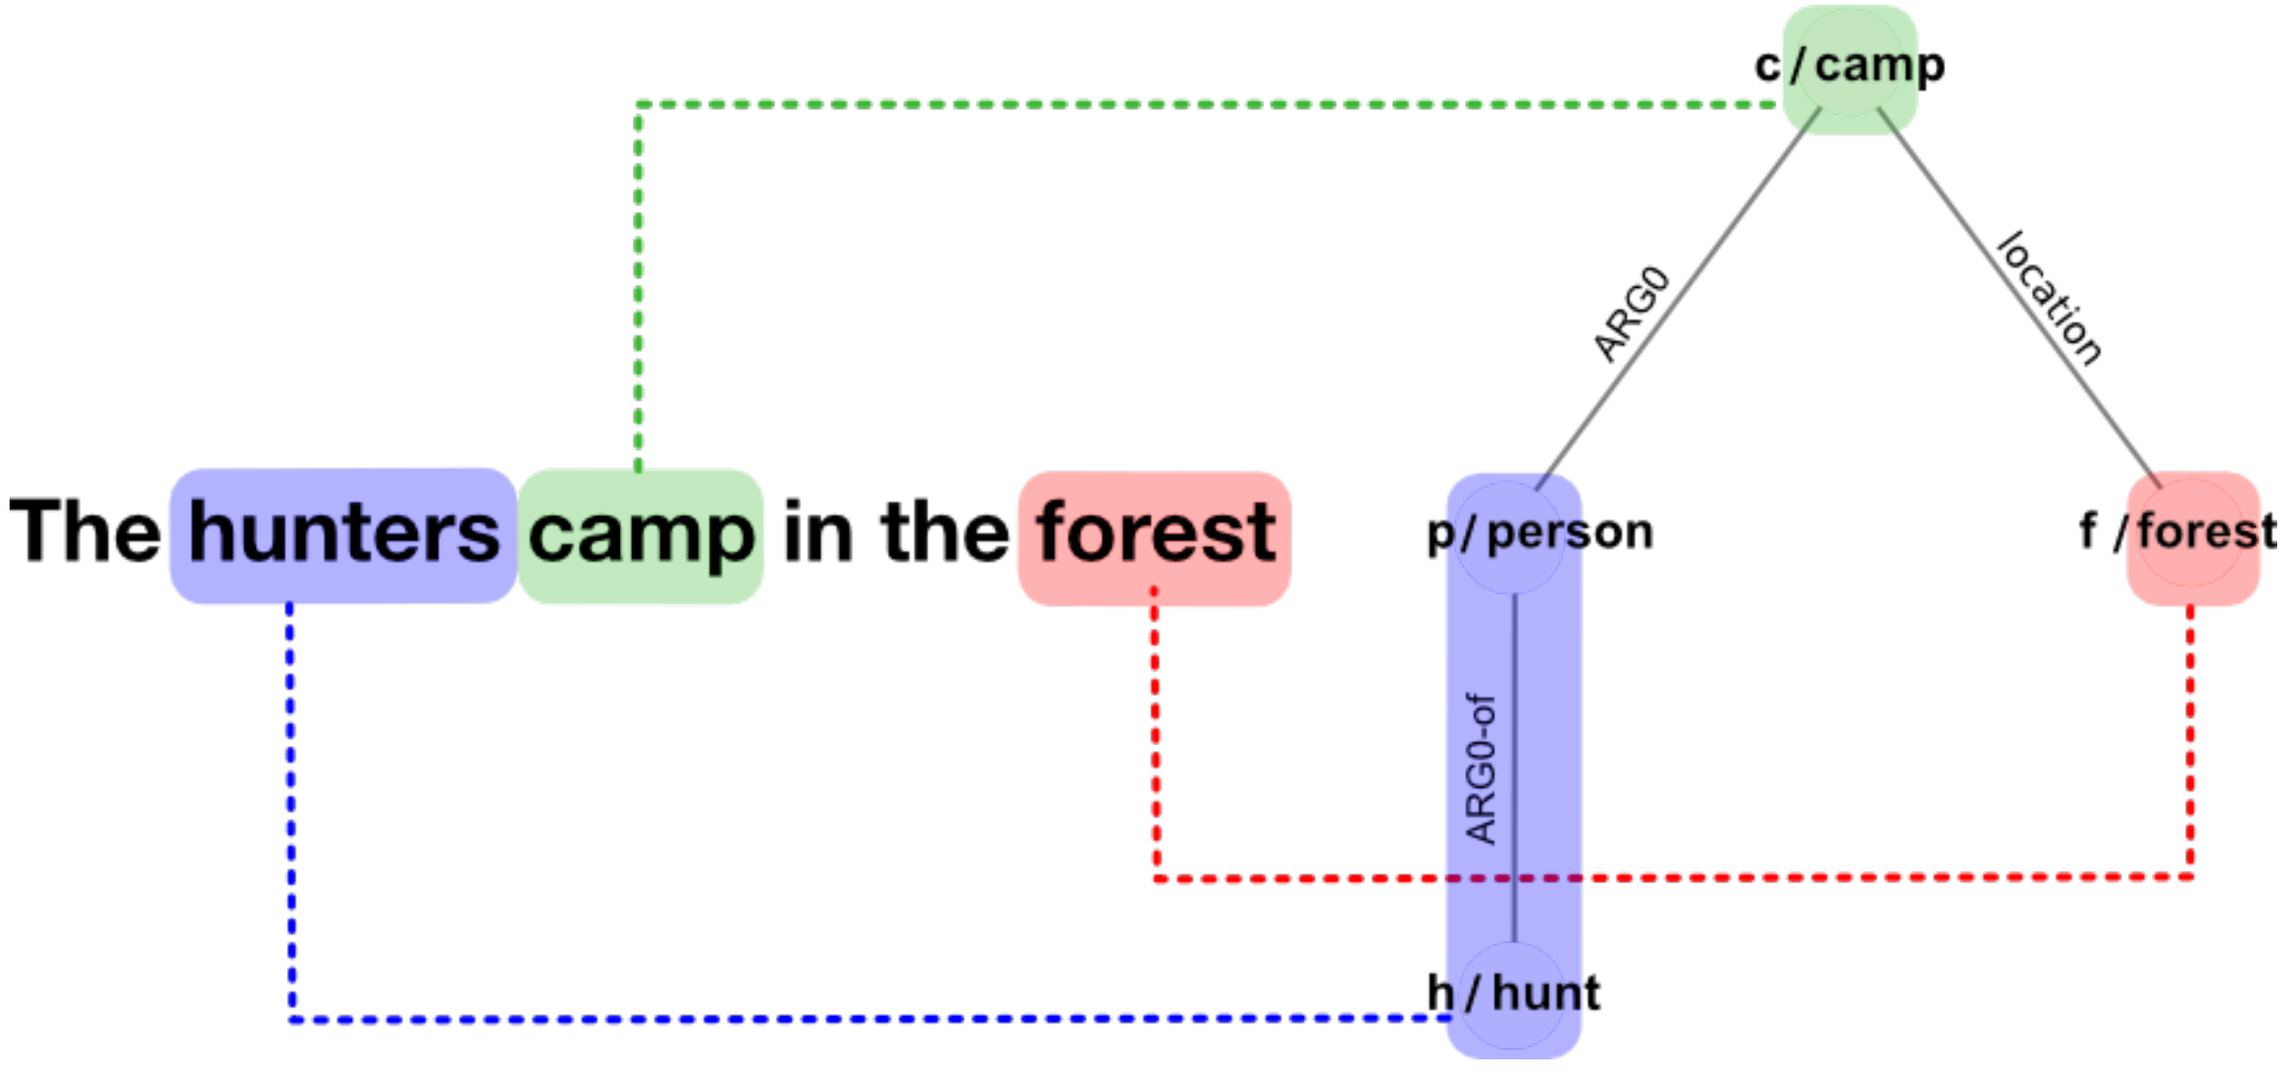

ISI-style [Pourdamghani et al., 2014]

- (Word, AMR node), (Word, AMR edge) alignments
- many-to-many

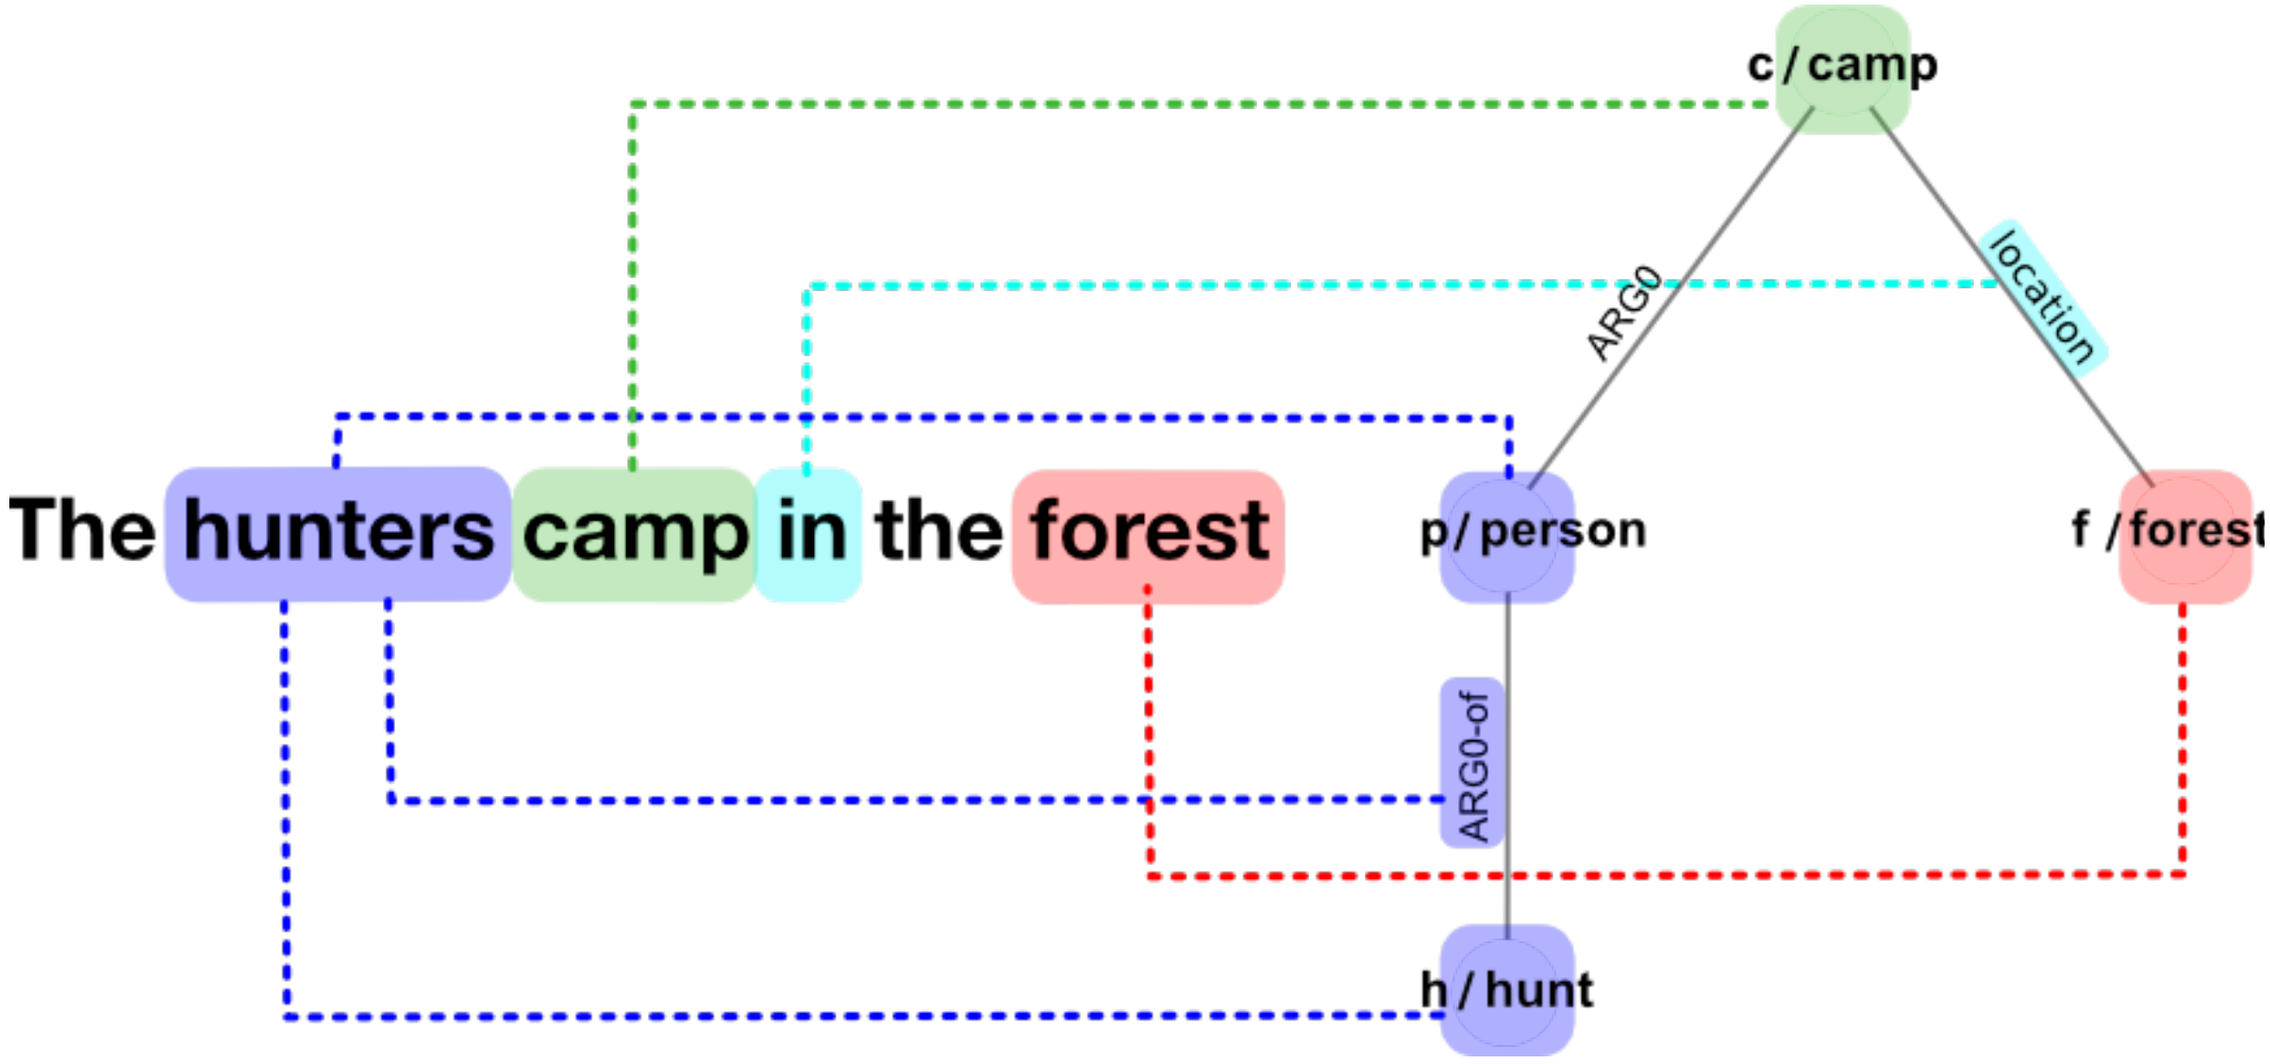

Relative to JAMR: lower level,
+ Compositional relations marked by **function words** (but only 23% of AMR edges covered),
– Distinguishing coreference from multiword expression

Why syntax?

- To explain all (or nearly all) of the AMR in terms of the sentence, we need more than string alignment.
  - Not every AMR edge is marked by a word—some reflected in word order.
- Syntax = grammatical conventions above the word level that give rise to semantic compositionality.
  - Alignments to syntax give a better picture of the derivational structure of the AMR.

Universal Dependencies (UD)

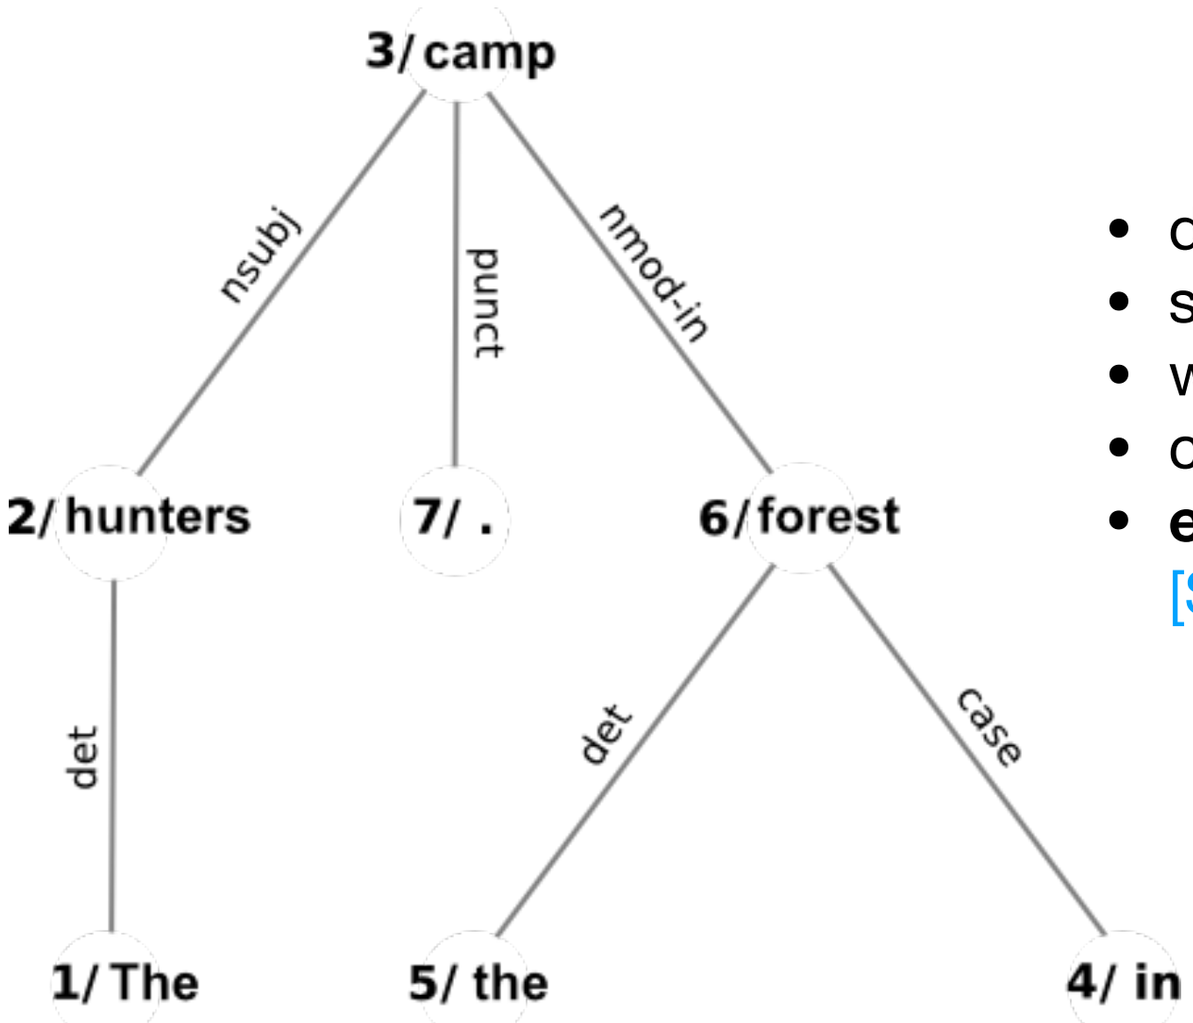

- directed, rooted graphs
- semantics-oriented, surface syntax
- widespread usage
- corpora in many languages
- **enhanced++ variant**
  [Schuster & Manning, 2016]

Syntax ↔ AMR

- Prior AMR work has modeled various kinds of syntax–semantics mappings [Wang et al., 2015; Artzi et al., 2015, Misra and Artzi, 2016, Chu and Kurohashi, 2016, Chen and Palmer, 2017].

- We are the first to
  - present a detailed linguistic **annotation scheme** for syntactic alignments, and
  - release a hand-annotated **dataset** with dependency syntax.

- AMR and dependency syntax are often assumed to be **similar**, but this claim has never been evaluated.

UD ↔ AMR

UD

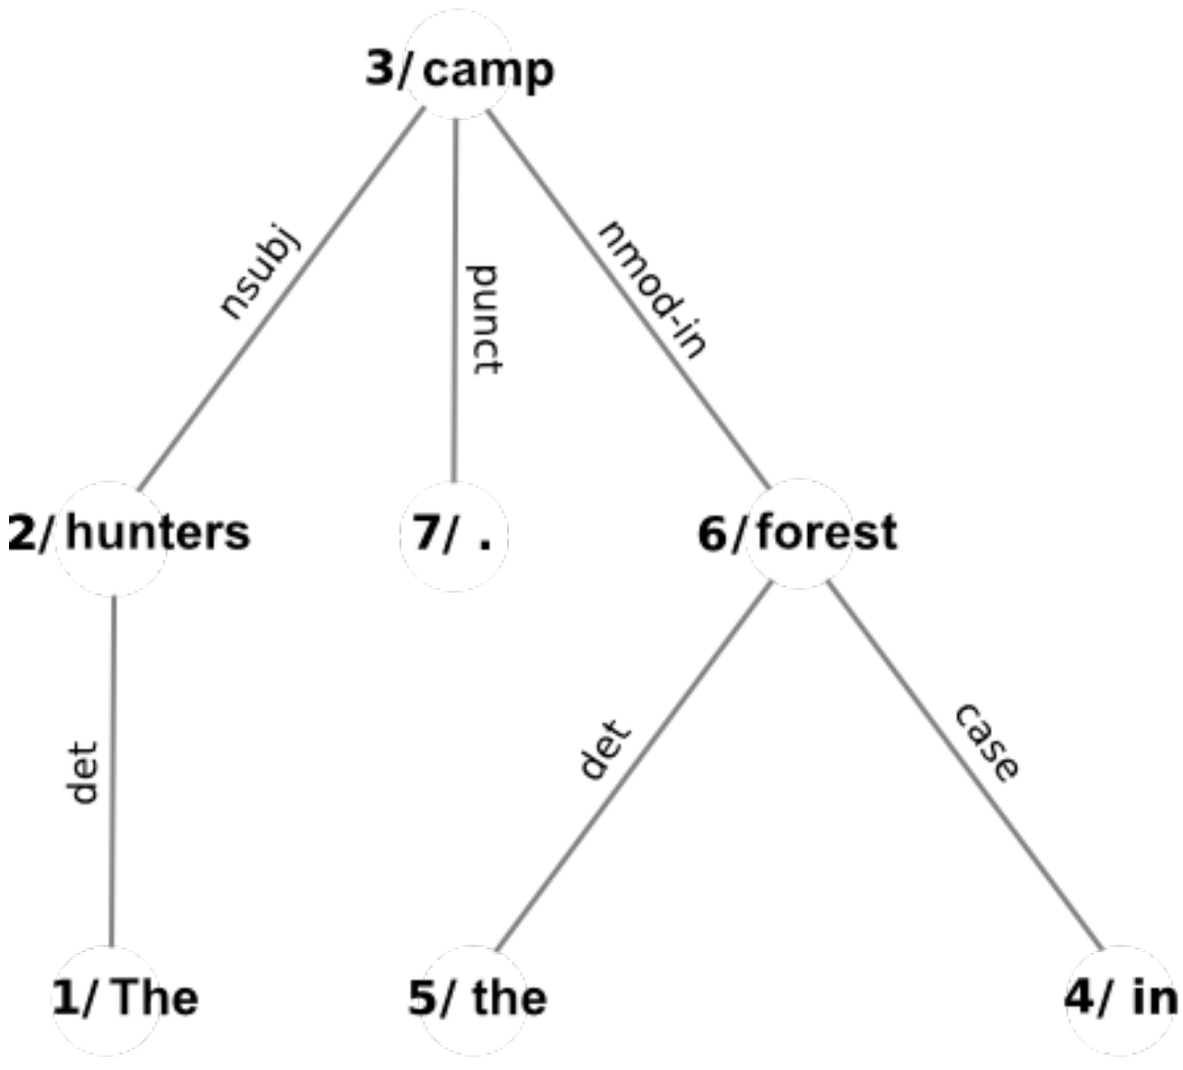

AMR

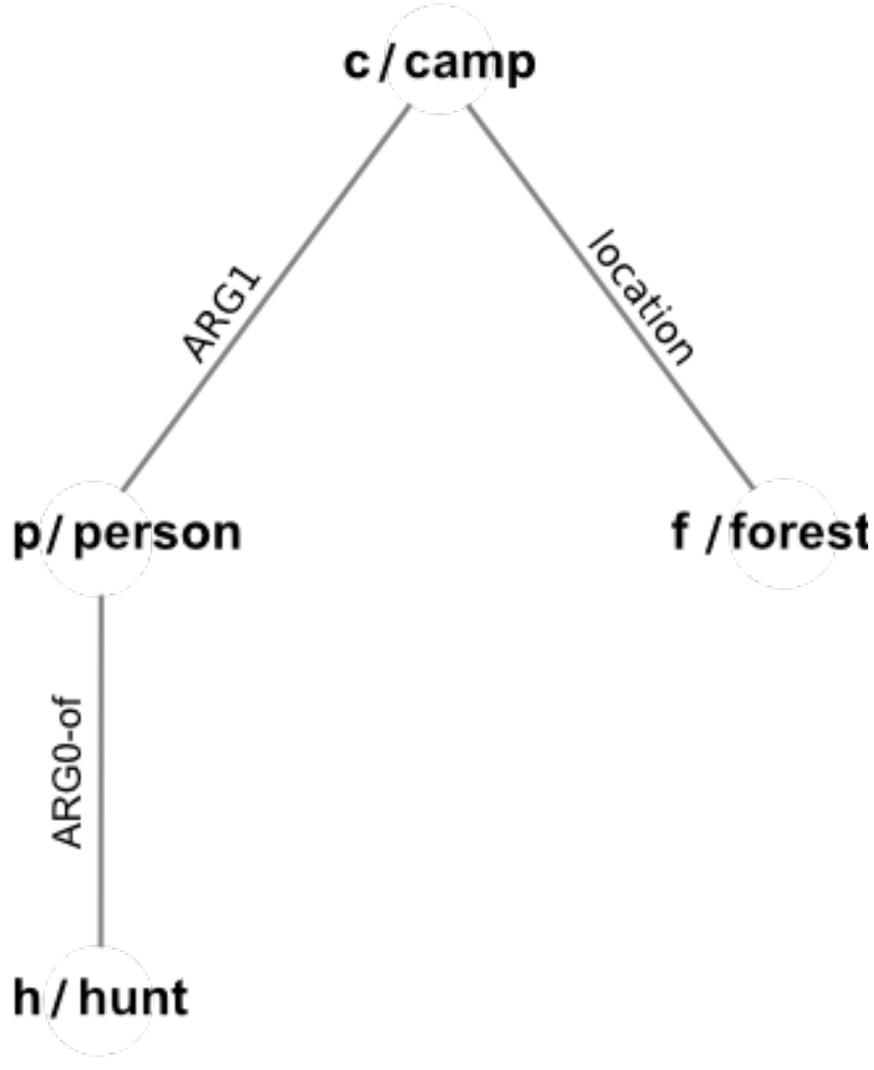

The hunters camp in the forest

Lexical alignments: (Node, Node)

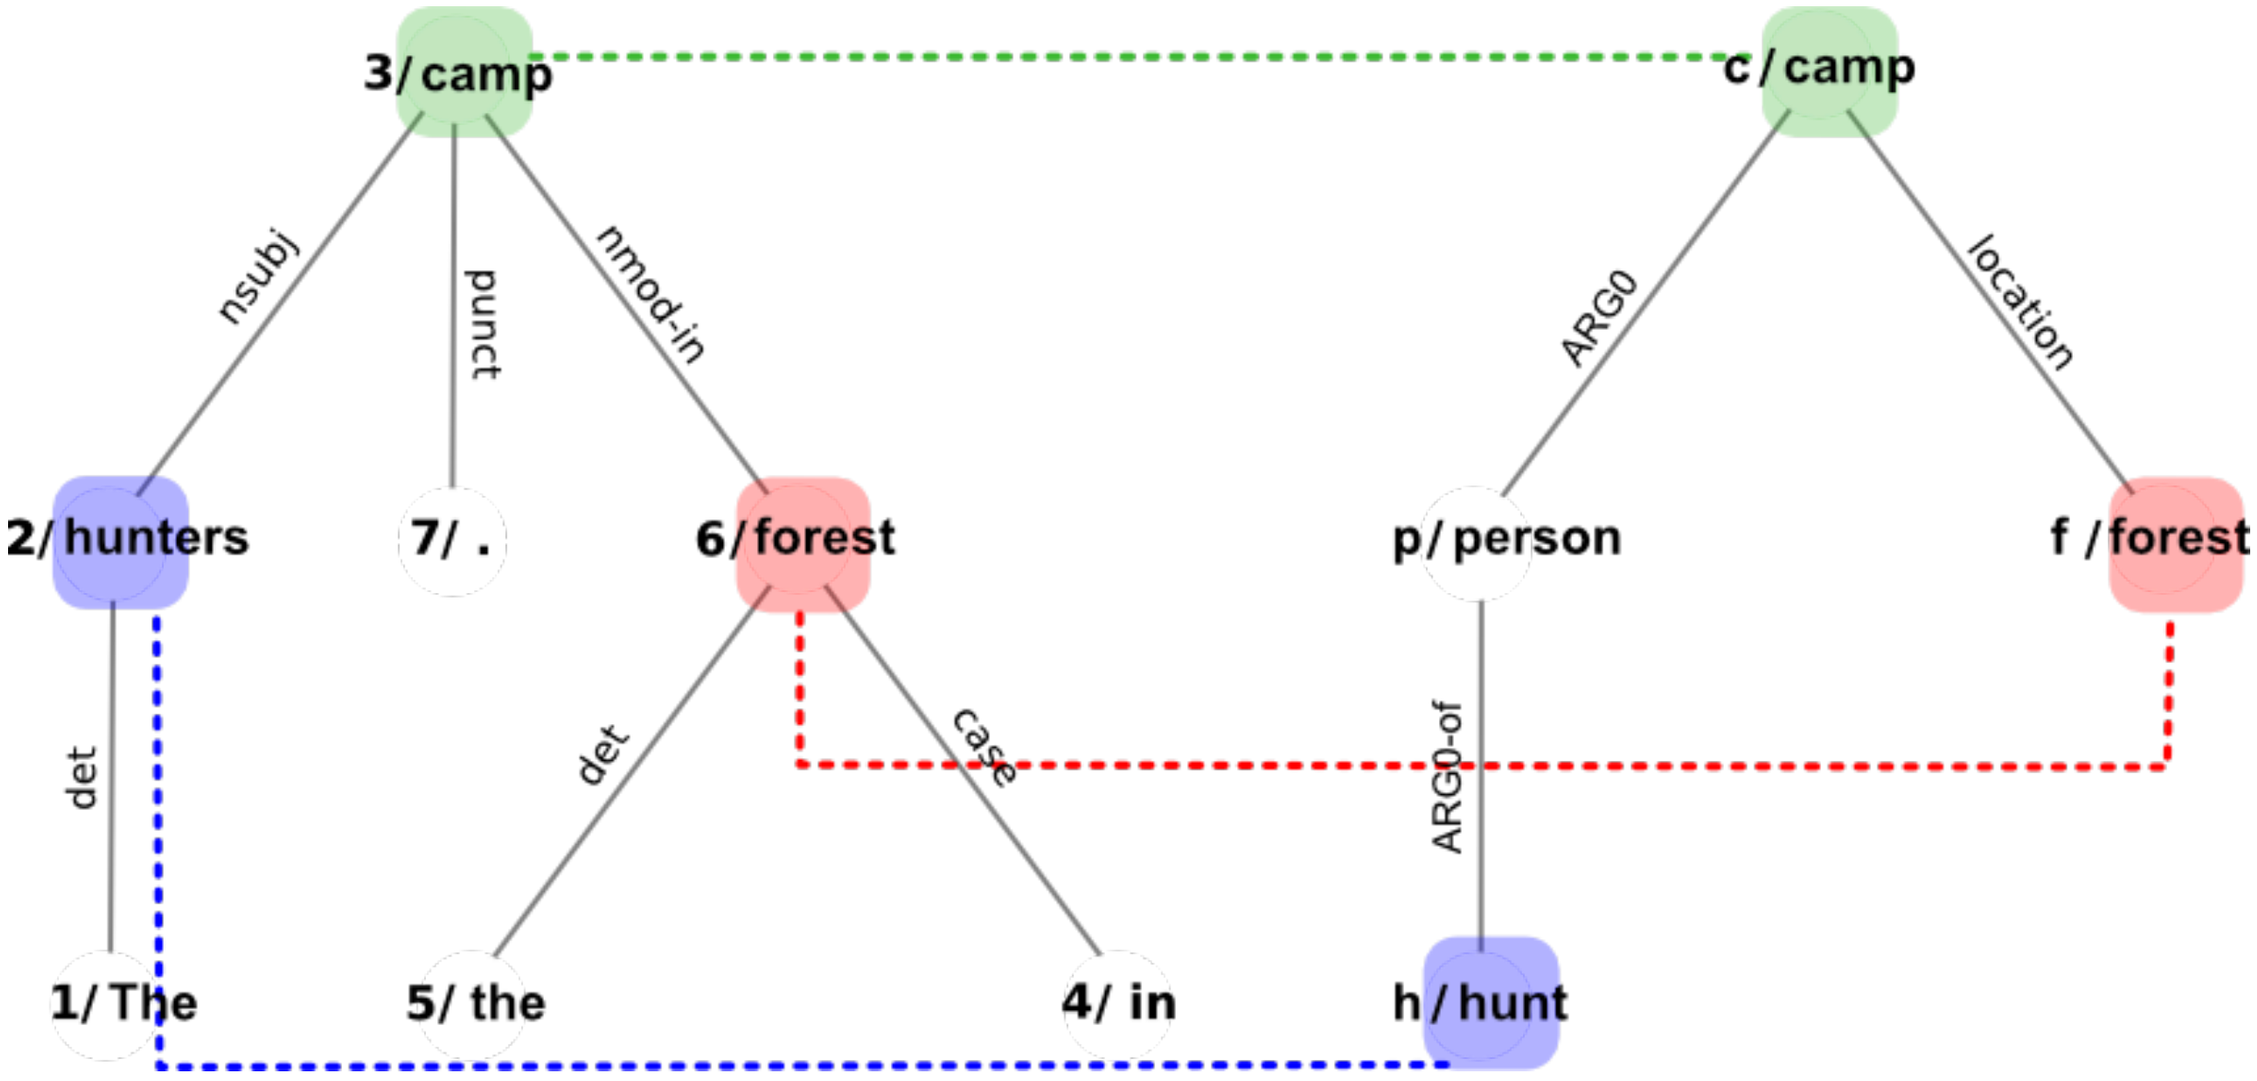

The hunters camp in the forest

Structural alignments

Connected subgraphs on both sides,
at least one of which is larger than 1 node

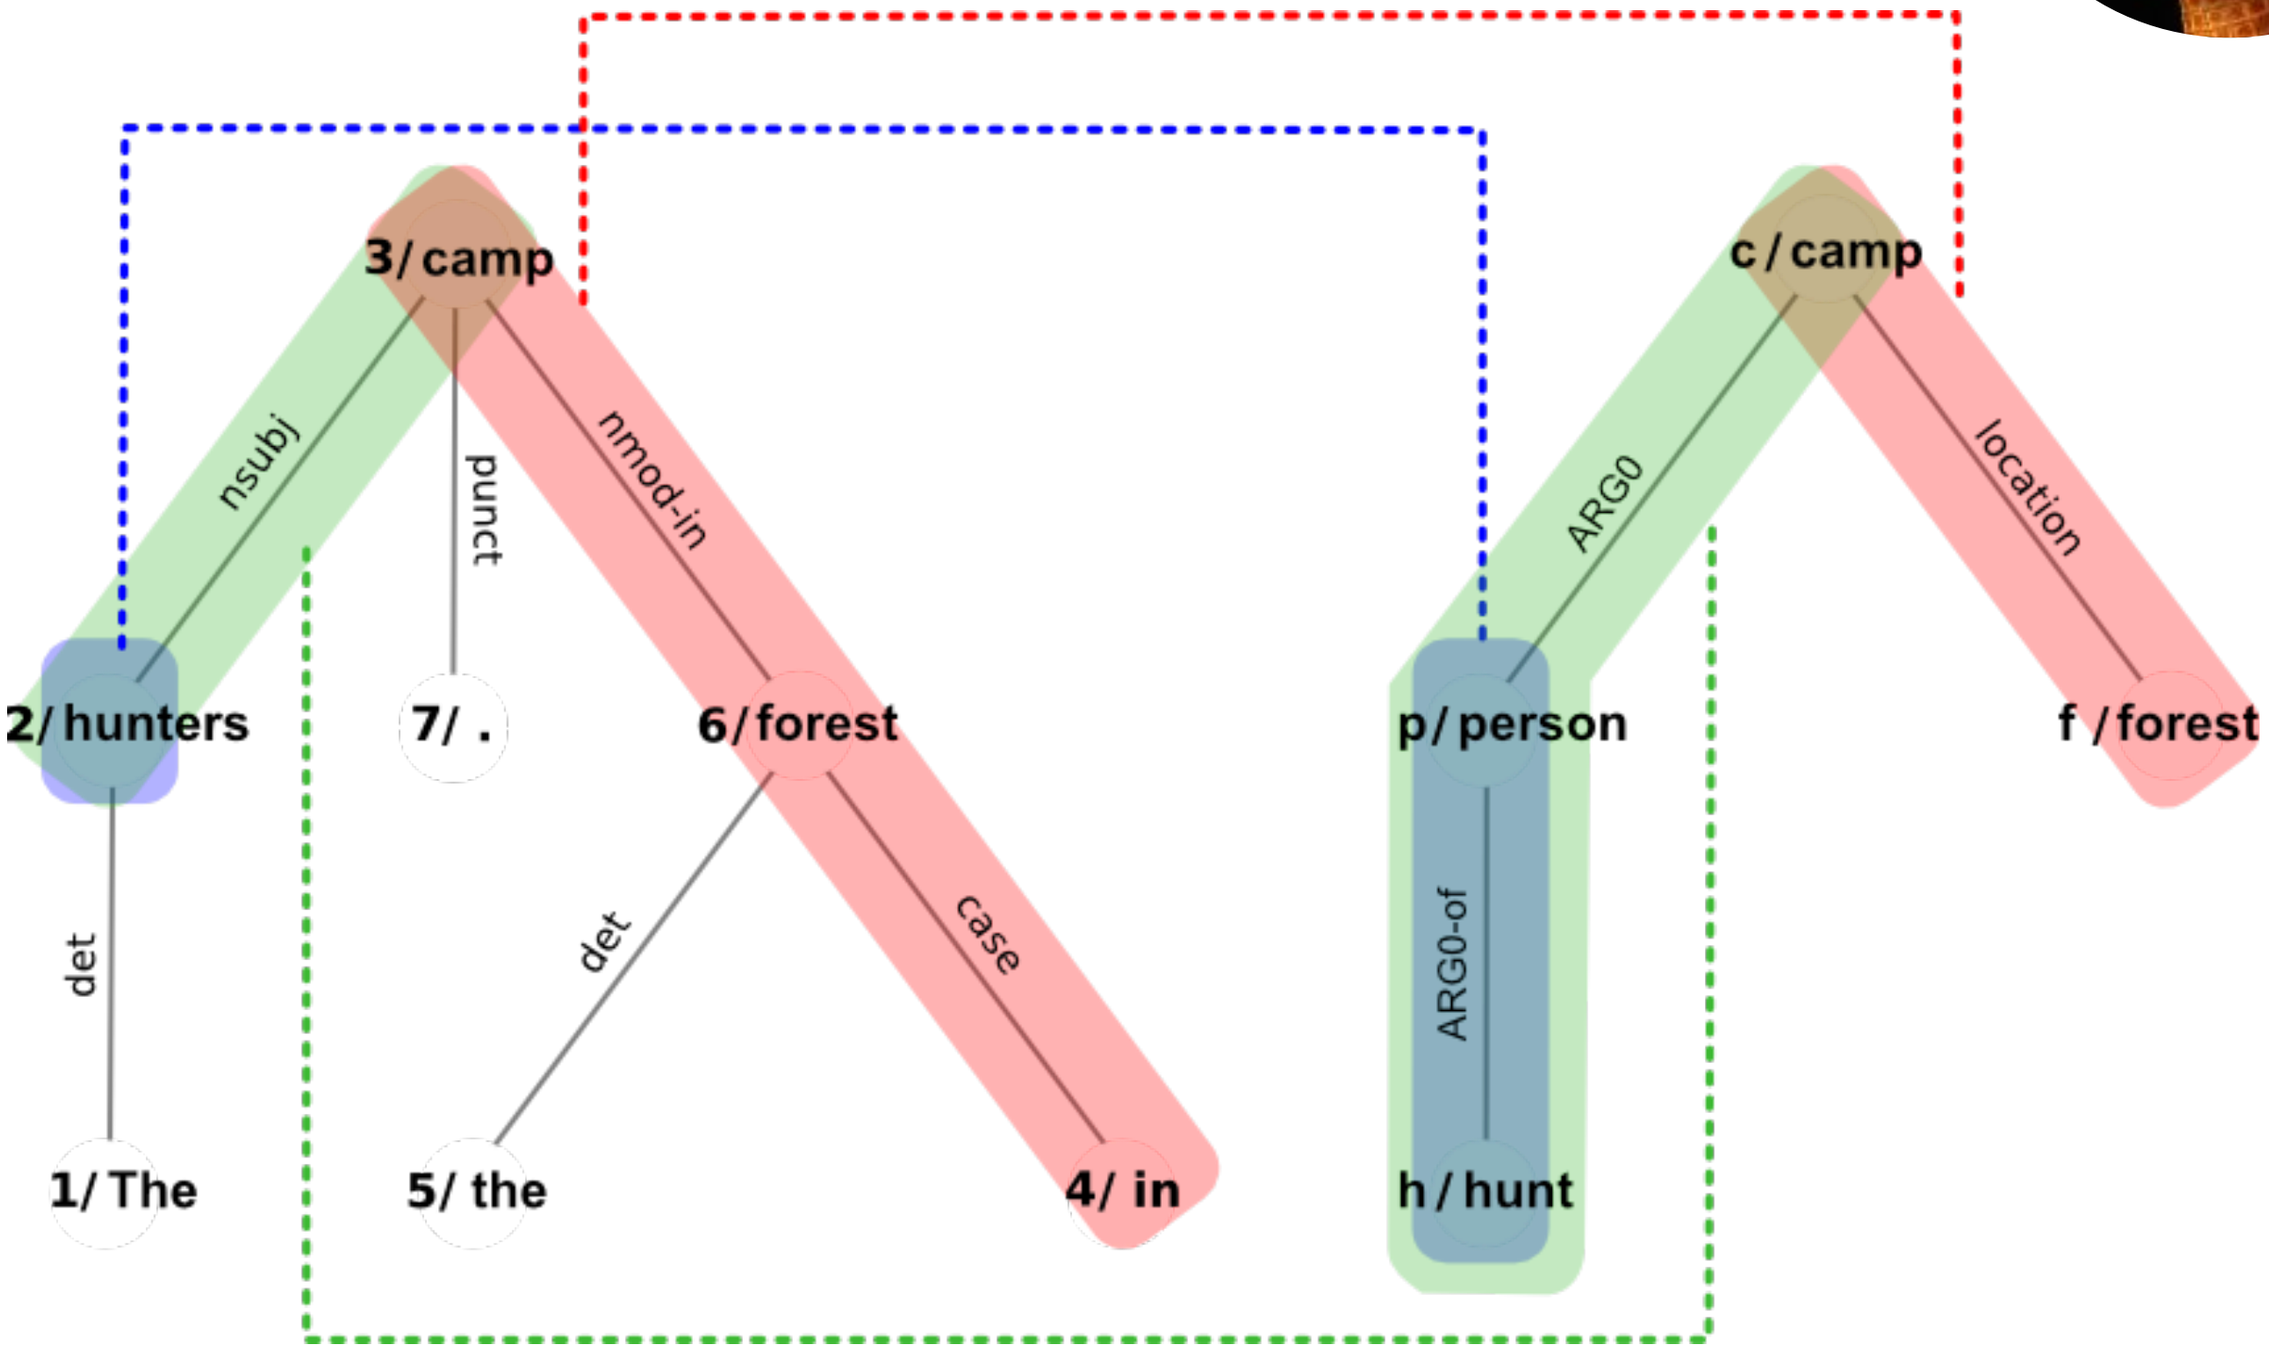

The hunters camp in the forest

Adverbial PP

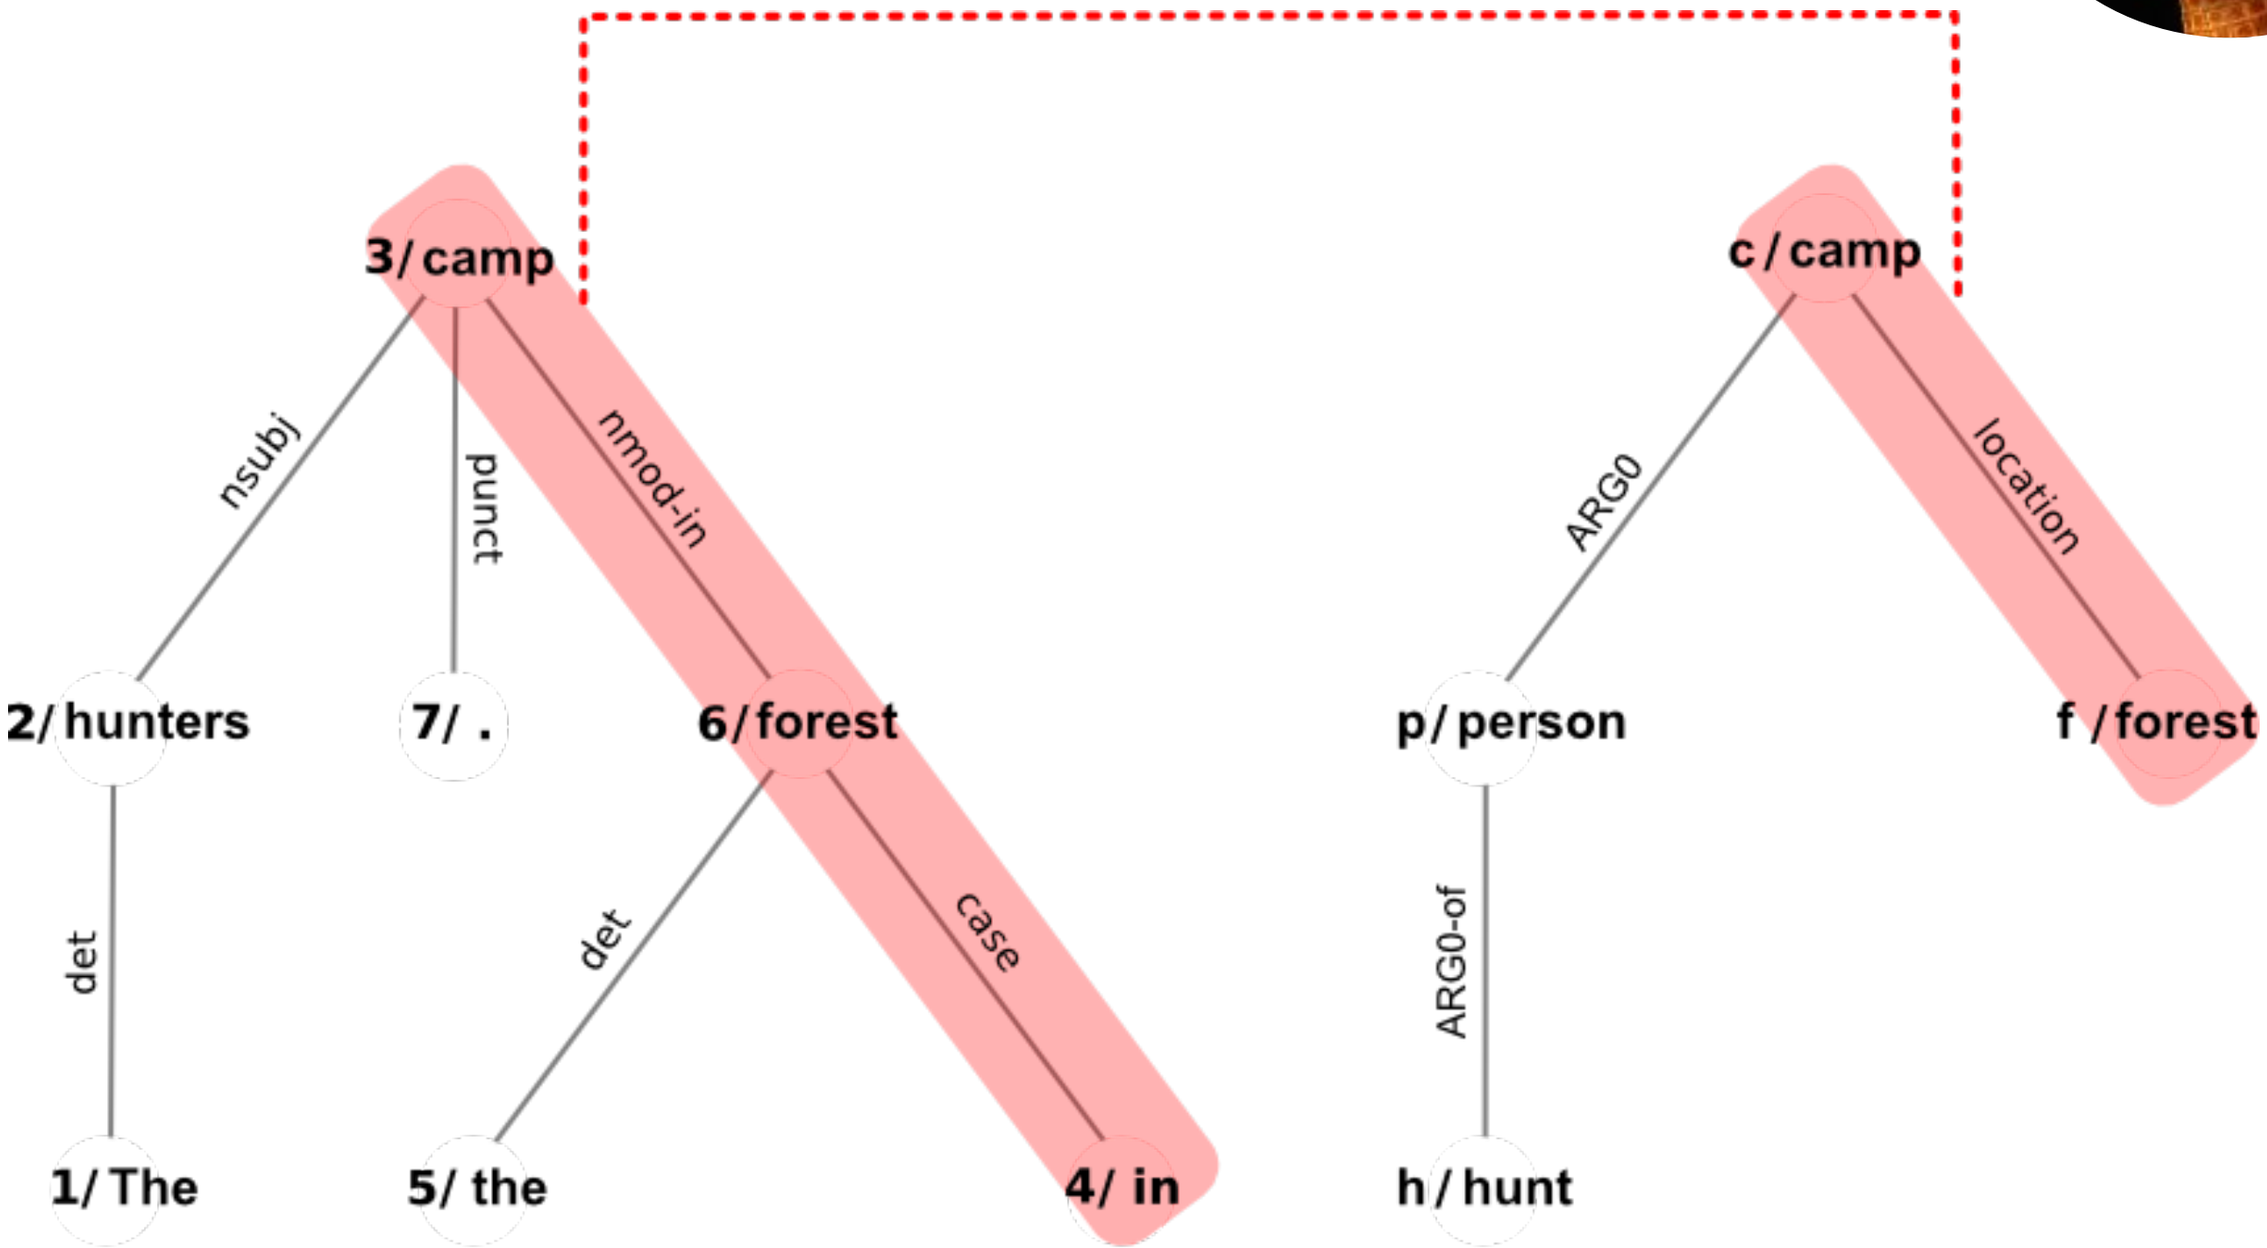

The hunters camp in the forest

Derived Noun

Similar treatment for **named entities.**

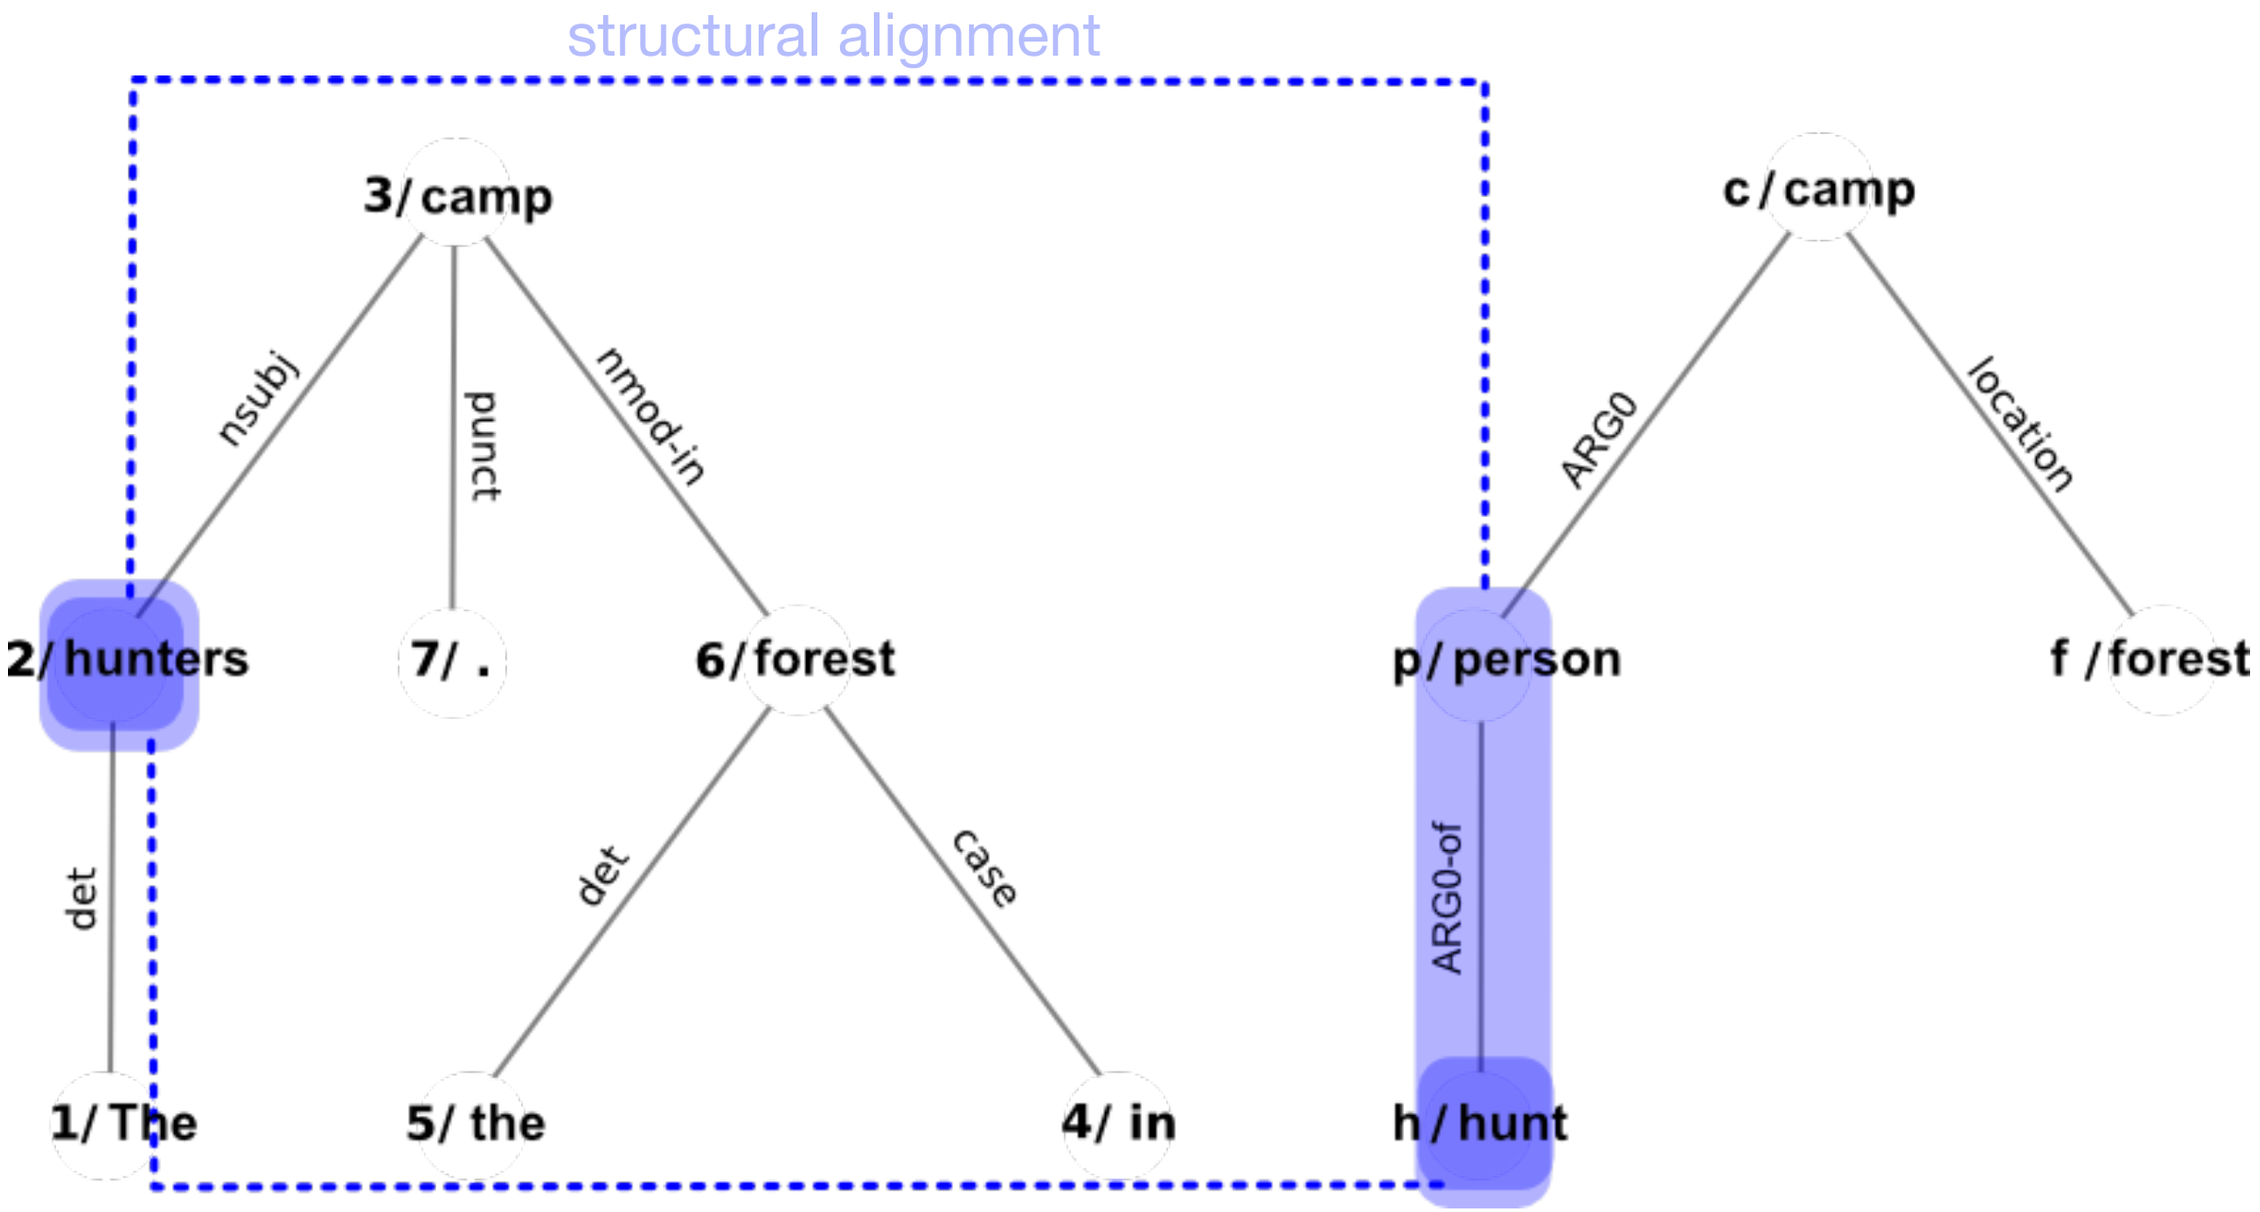

The hunters camp in the forest

Subject

Subsumption Principle for hierarchical alignments: Because the ‘hunters’ node aligns to `person :ARG0-of hunt`, any structural alignment containing ‘hunters’ must contain that AMR subgraph.

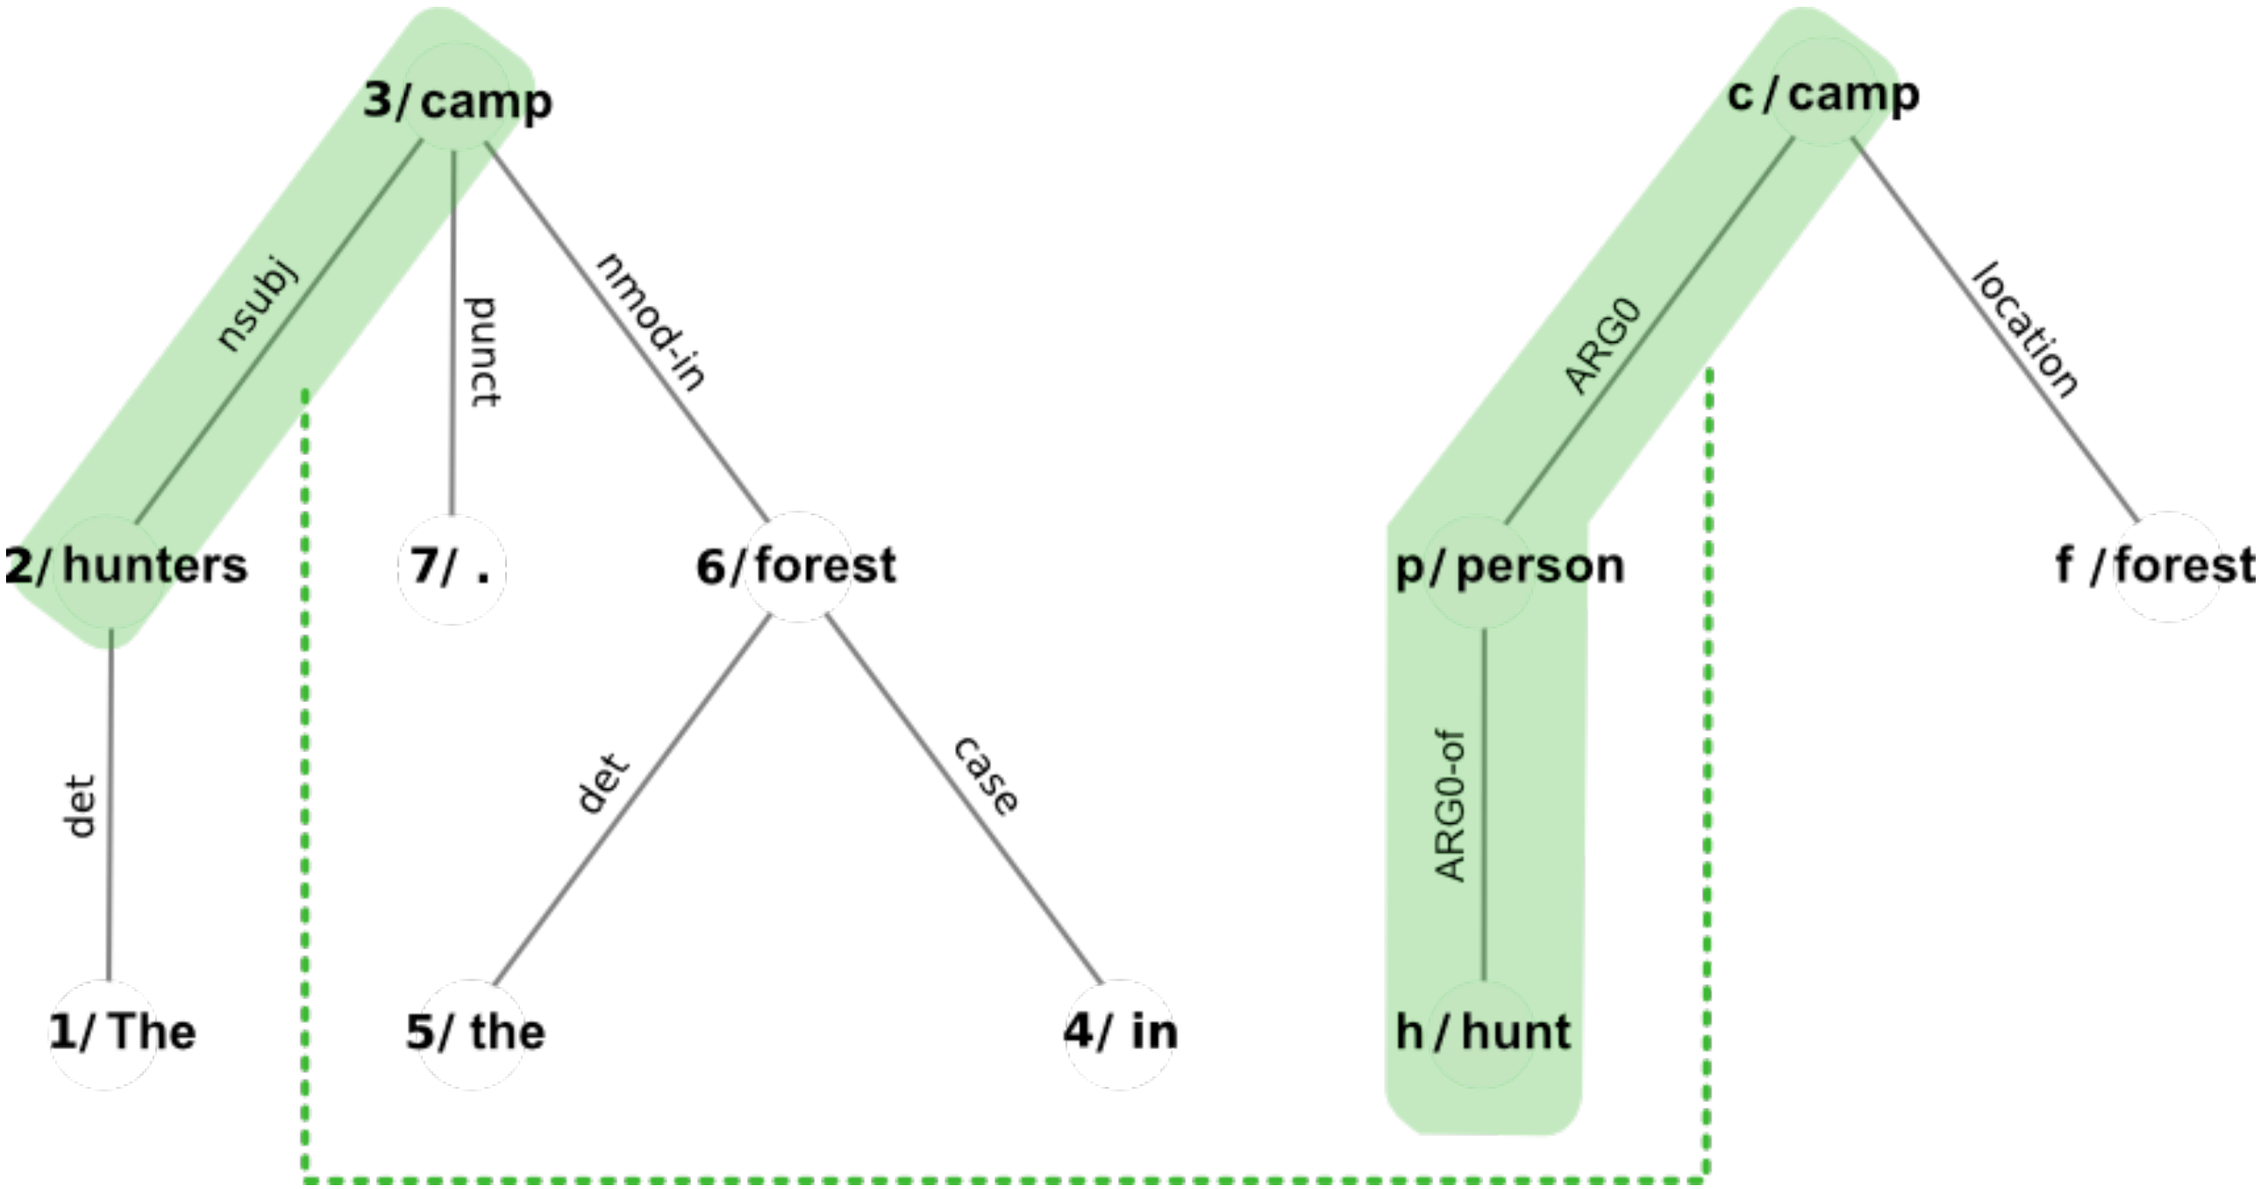

The hunters camp in the forest

Structural alignments

Connected subgraphs on both sides,
at least one of which is larger than 1 node

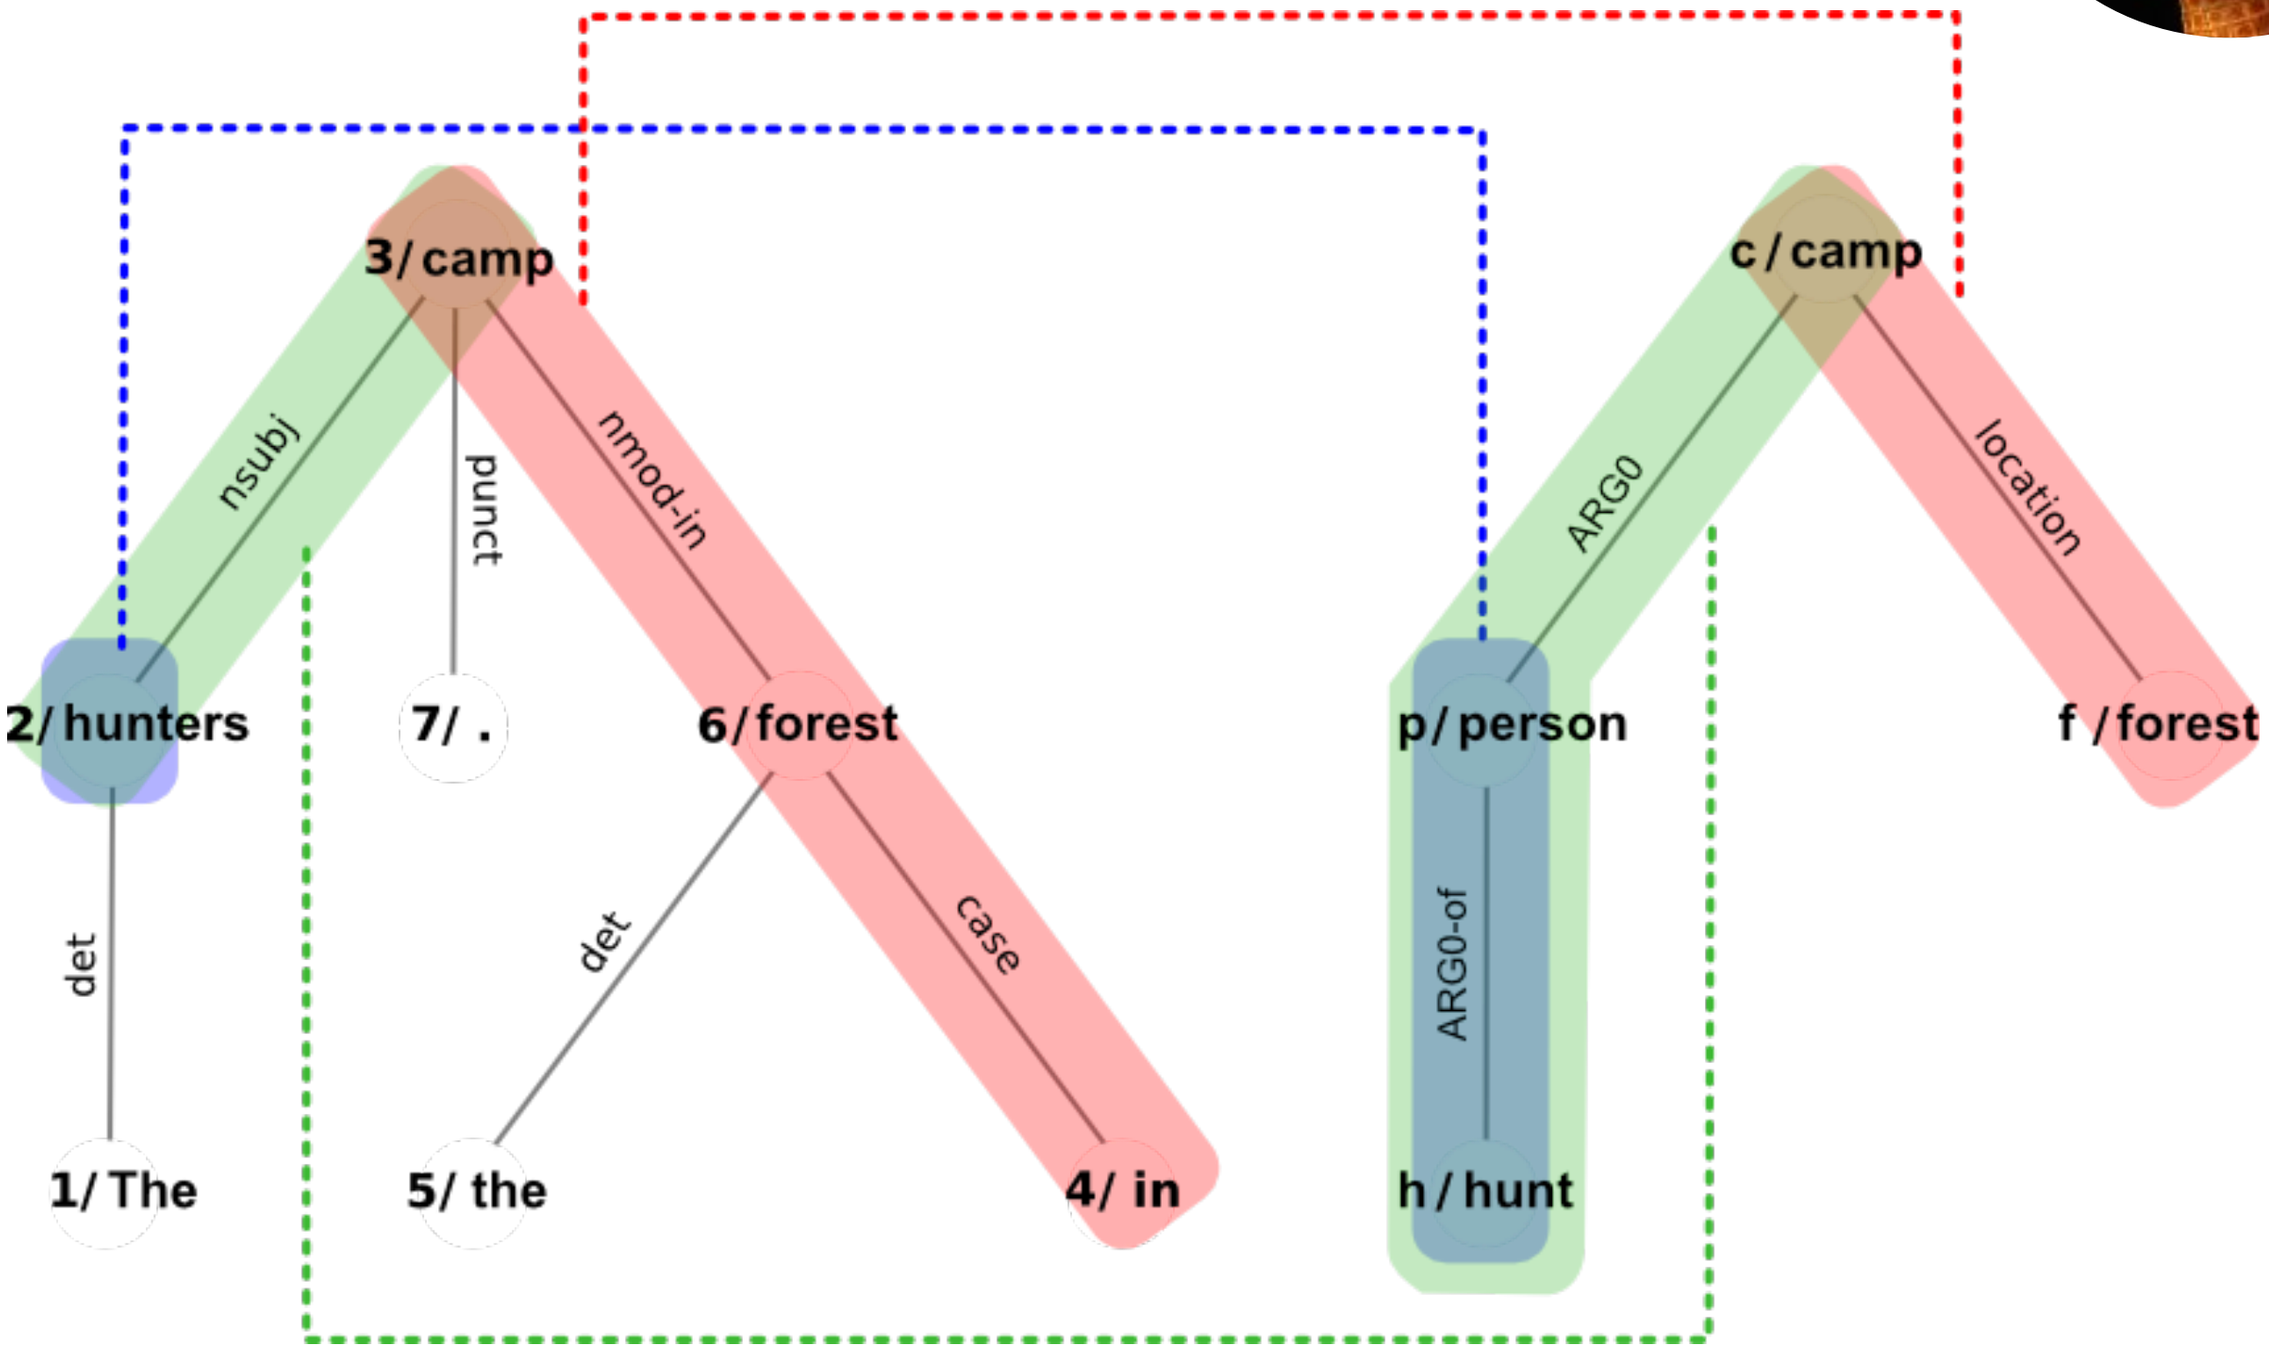

The hunters camp in the forest

Hierarchical alignments

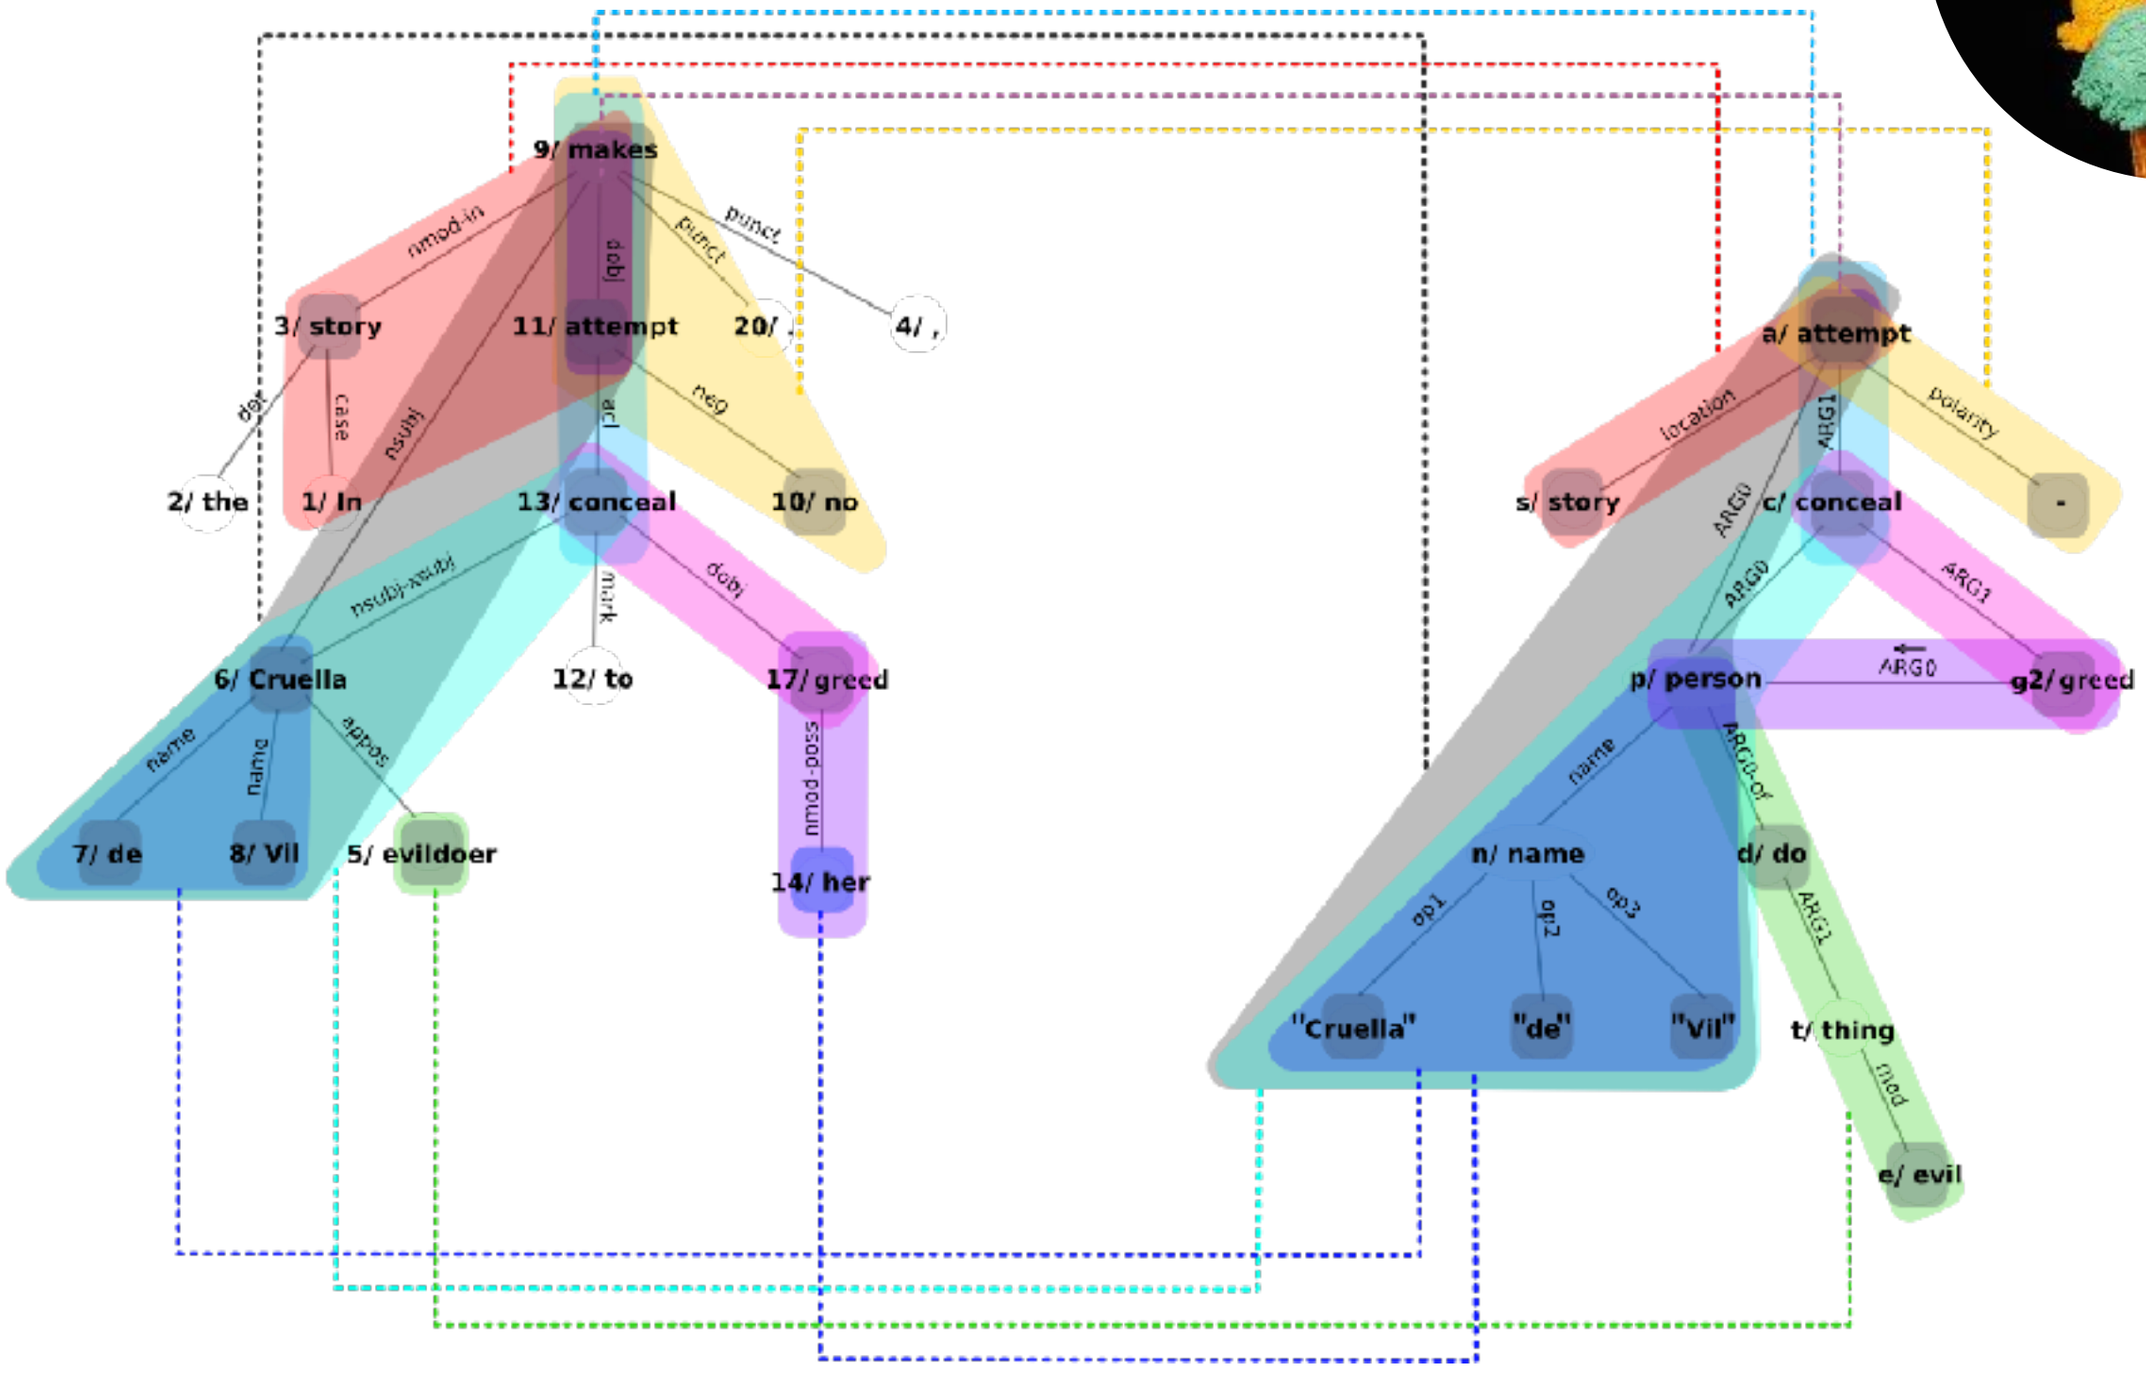

In the story, evildoer Cruella de Vil makes no attempt to conceal her greed.

200 hand-aligned sentences

UD: hand-corrected CoreNLP parses

IAA: 96% for lexical, 80% for structural

http://tiny.cc/amrud

Coverage

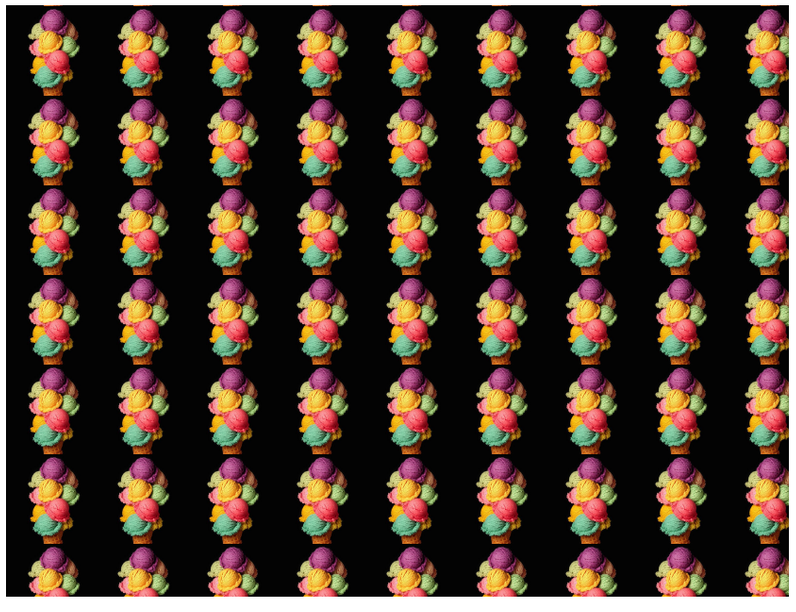

Perhaps from-scratch AMR annotation gives too much flexibility, and annotators incorporate inferences from **beyond the sentence** [Bender et al., 2015]

99.3% of AMR nodes
97.2% of AMR edges
} are part of at least 1 alignment

81.5% of AMRs are fully covered

Thus, **nearly all** information in an AMR is evoked by **lexical items** and **syntax**.

AMR–UD Similarity

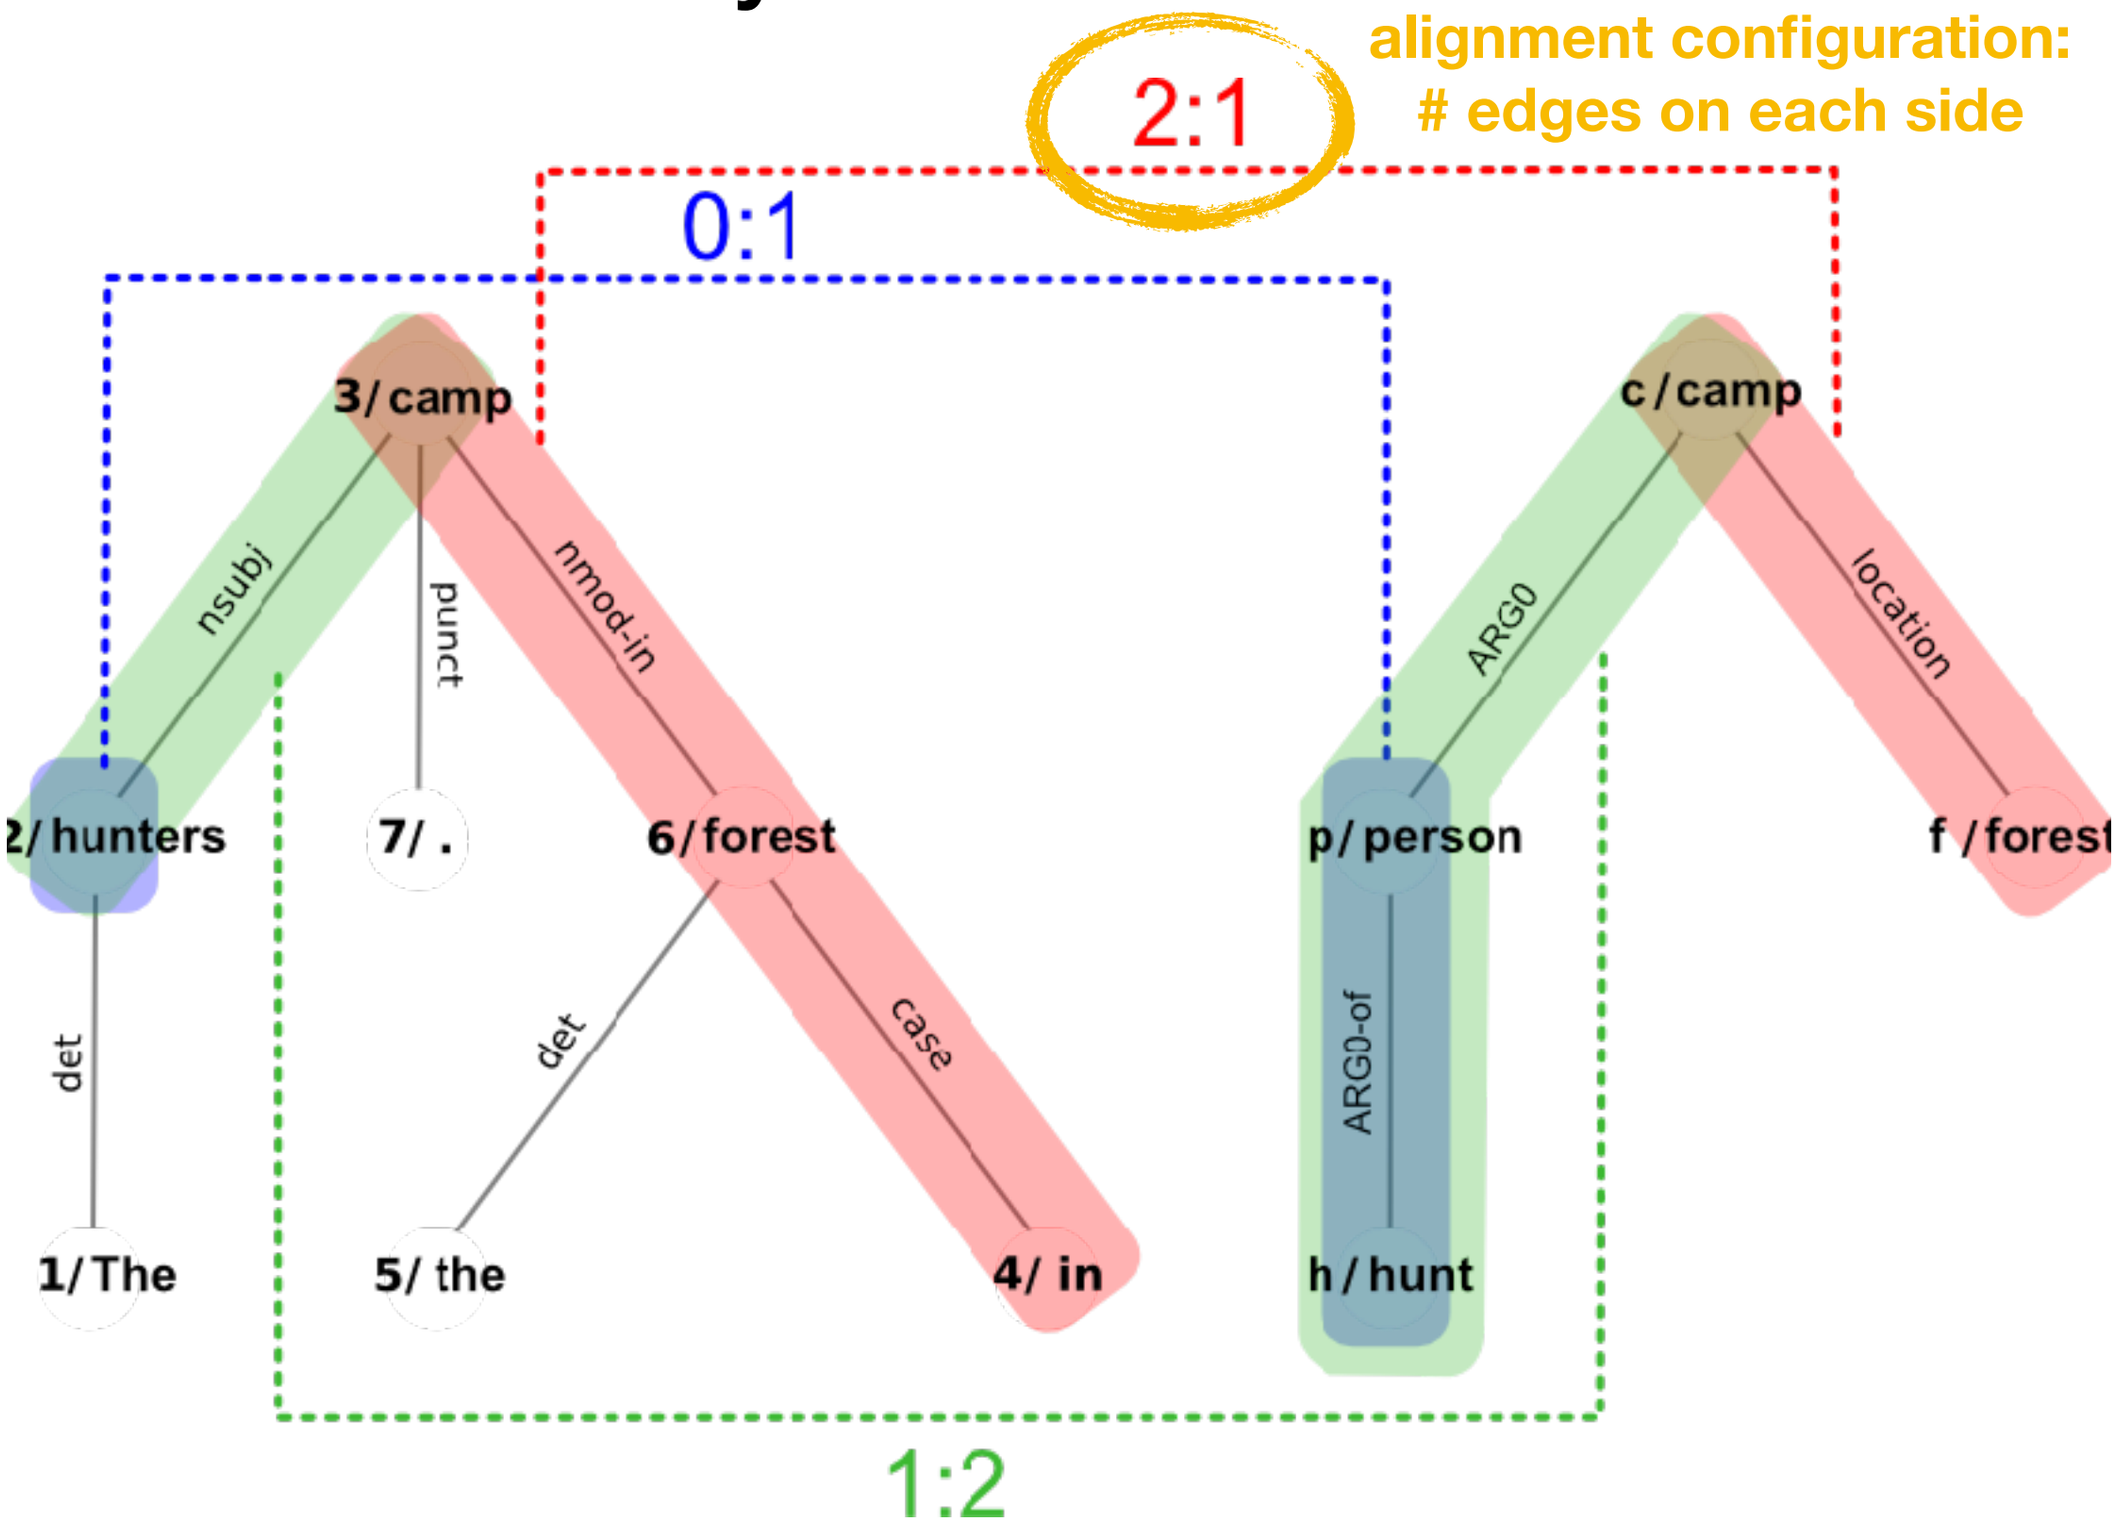

Distribution of alignment configurations

10% complex: multiple UD edges & multiple AMR edges

90% simple

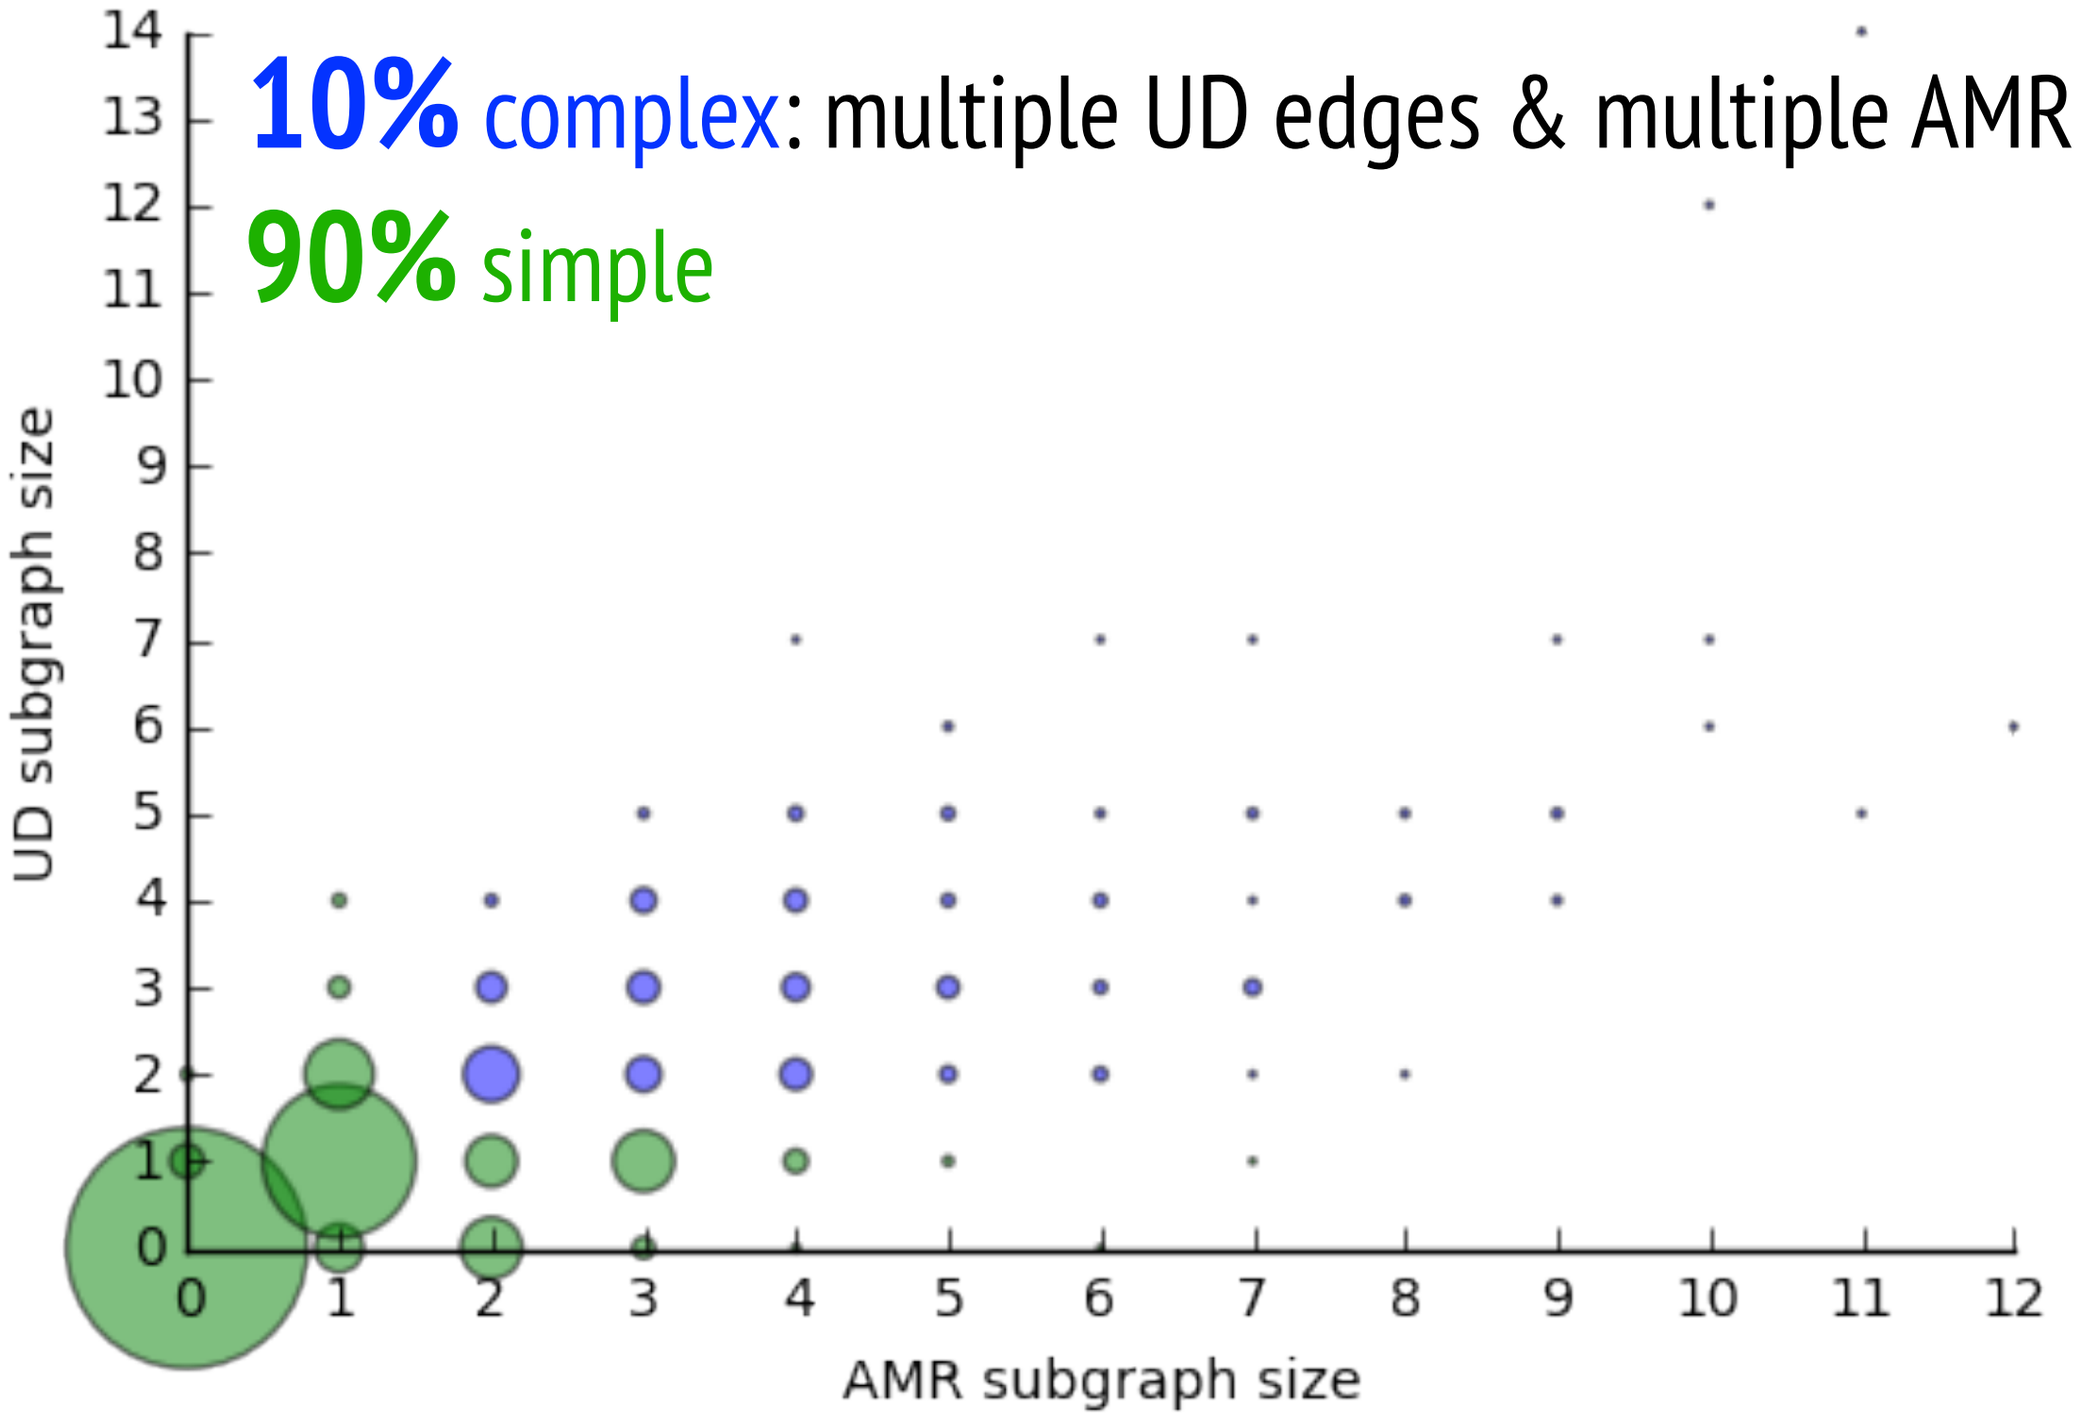

Complex configurations are frequently due to

coordination: **28%** (different head rules)

named entities: **10%** (MWE with each part of name in AMR)

semantic decomposition: **6%**

quantities/dates: **5%**

How similar are AMR and UD?

10% complex alignments

66% of sentences have at least 1 complex alignment

Thus, most AMRs have some
local structural dissimilarity.

Automatic alignment: **lexical**

F1

Our rule-based algorithm: **87%** (mainly string match; no syntax)

Automatic alignment: **structural**

Simple algorithm that infers structural alignments from lexical alignments via path search

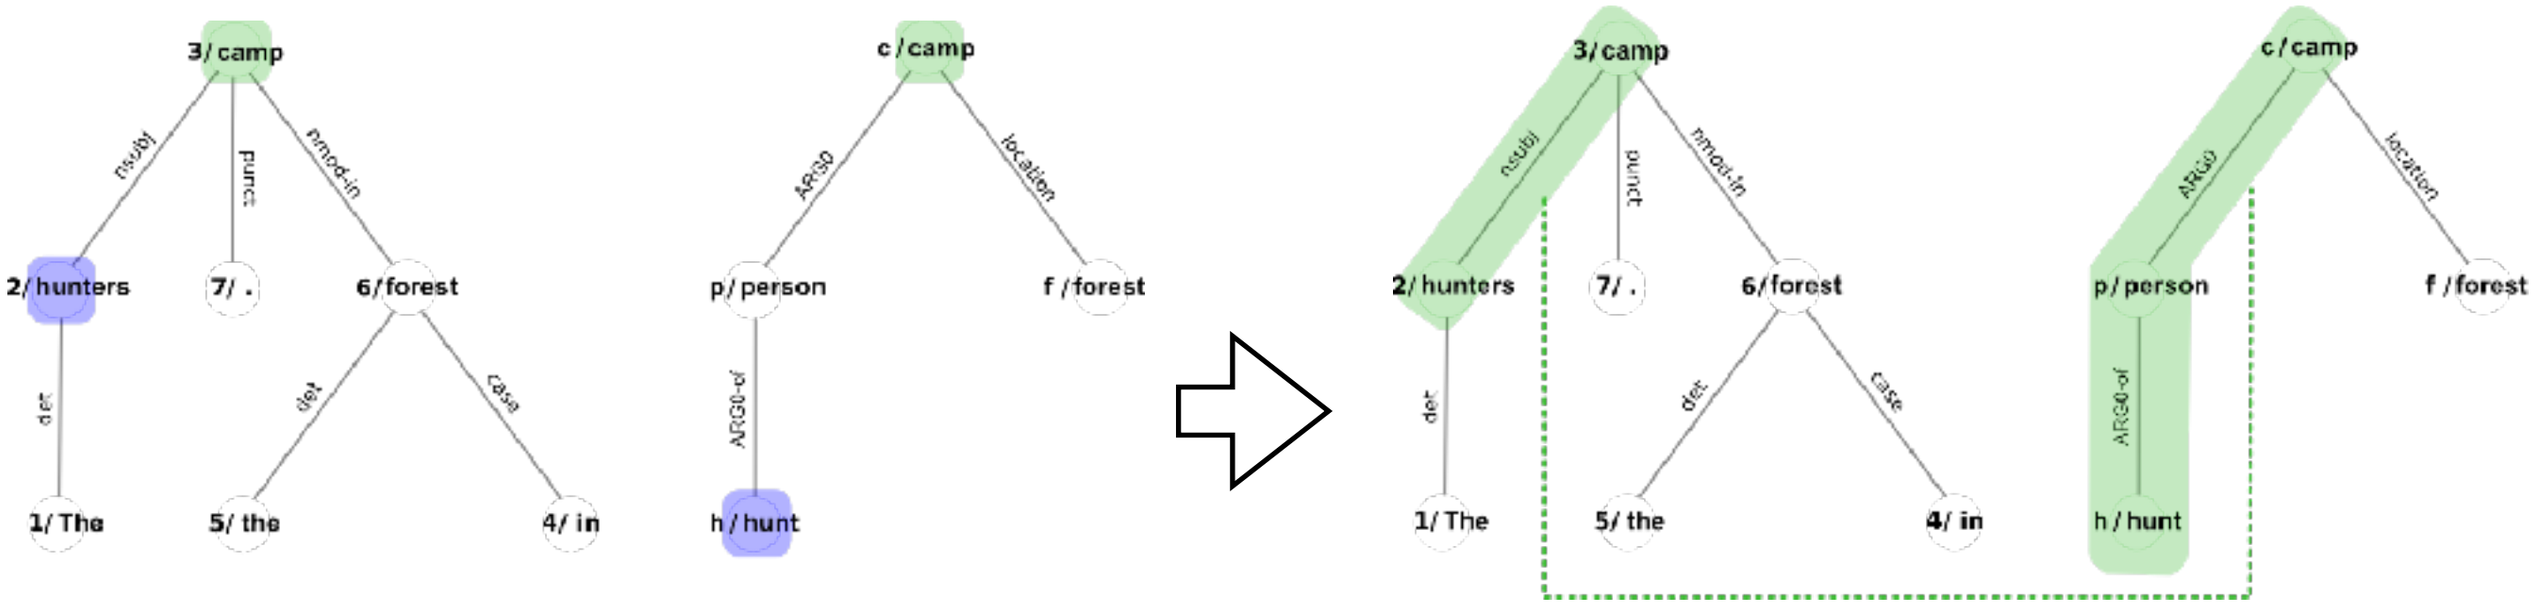

F1

Gold UD & lexical alignments: **76%**

Gold UD, auto lexical alignments: **61%**

Auto UD & lexical alignments: **55%**

Conclusions

- Aligning AMRs to **dependency parses** (rather than strings) accounts for nearly all of the AMR nodes and edges
- AMR and UD are **broadly similar**, but many sources of local dissimilarity
- **Lexical alignment** can be largely automated, but **structural alignment** is harder
- We **release** our guidelines, data, and code

More in the paper

- Linguistic annotation guidelines
- Constraints on structural alignments
- Rule-based algorithms for lexical and structural alignment
- Syntactic error analysis of an AMR parser

Future Work

- Better alignment algorithms
  - Adjust alignment scheme as AMR standard evolves
    [Bonial et al., 2018, …]
- Richer alignments ⇒ better AMR parsers & generators?
  - By feeding the alignments into the system, or
  - Evaluating attention in neural systems

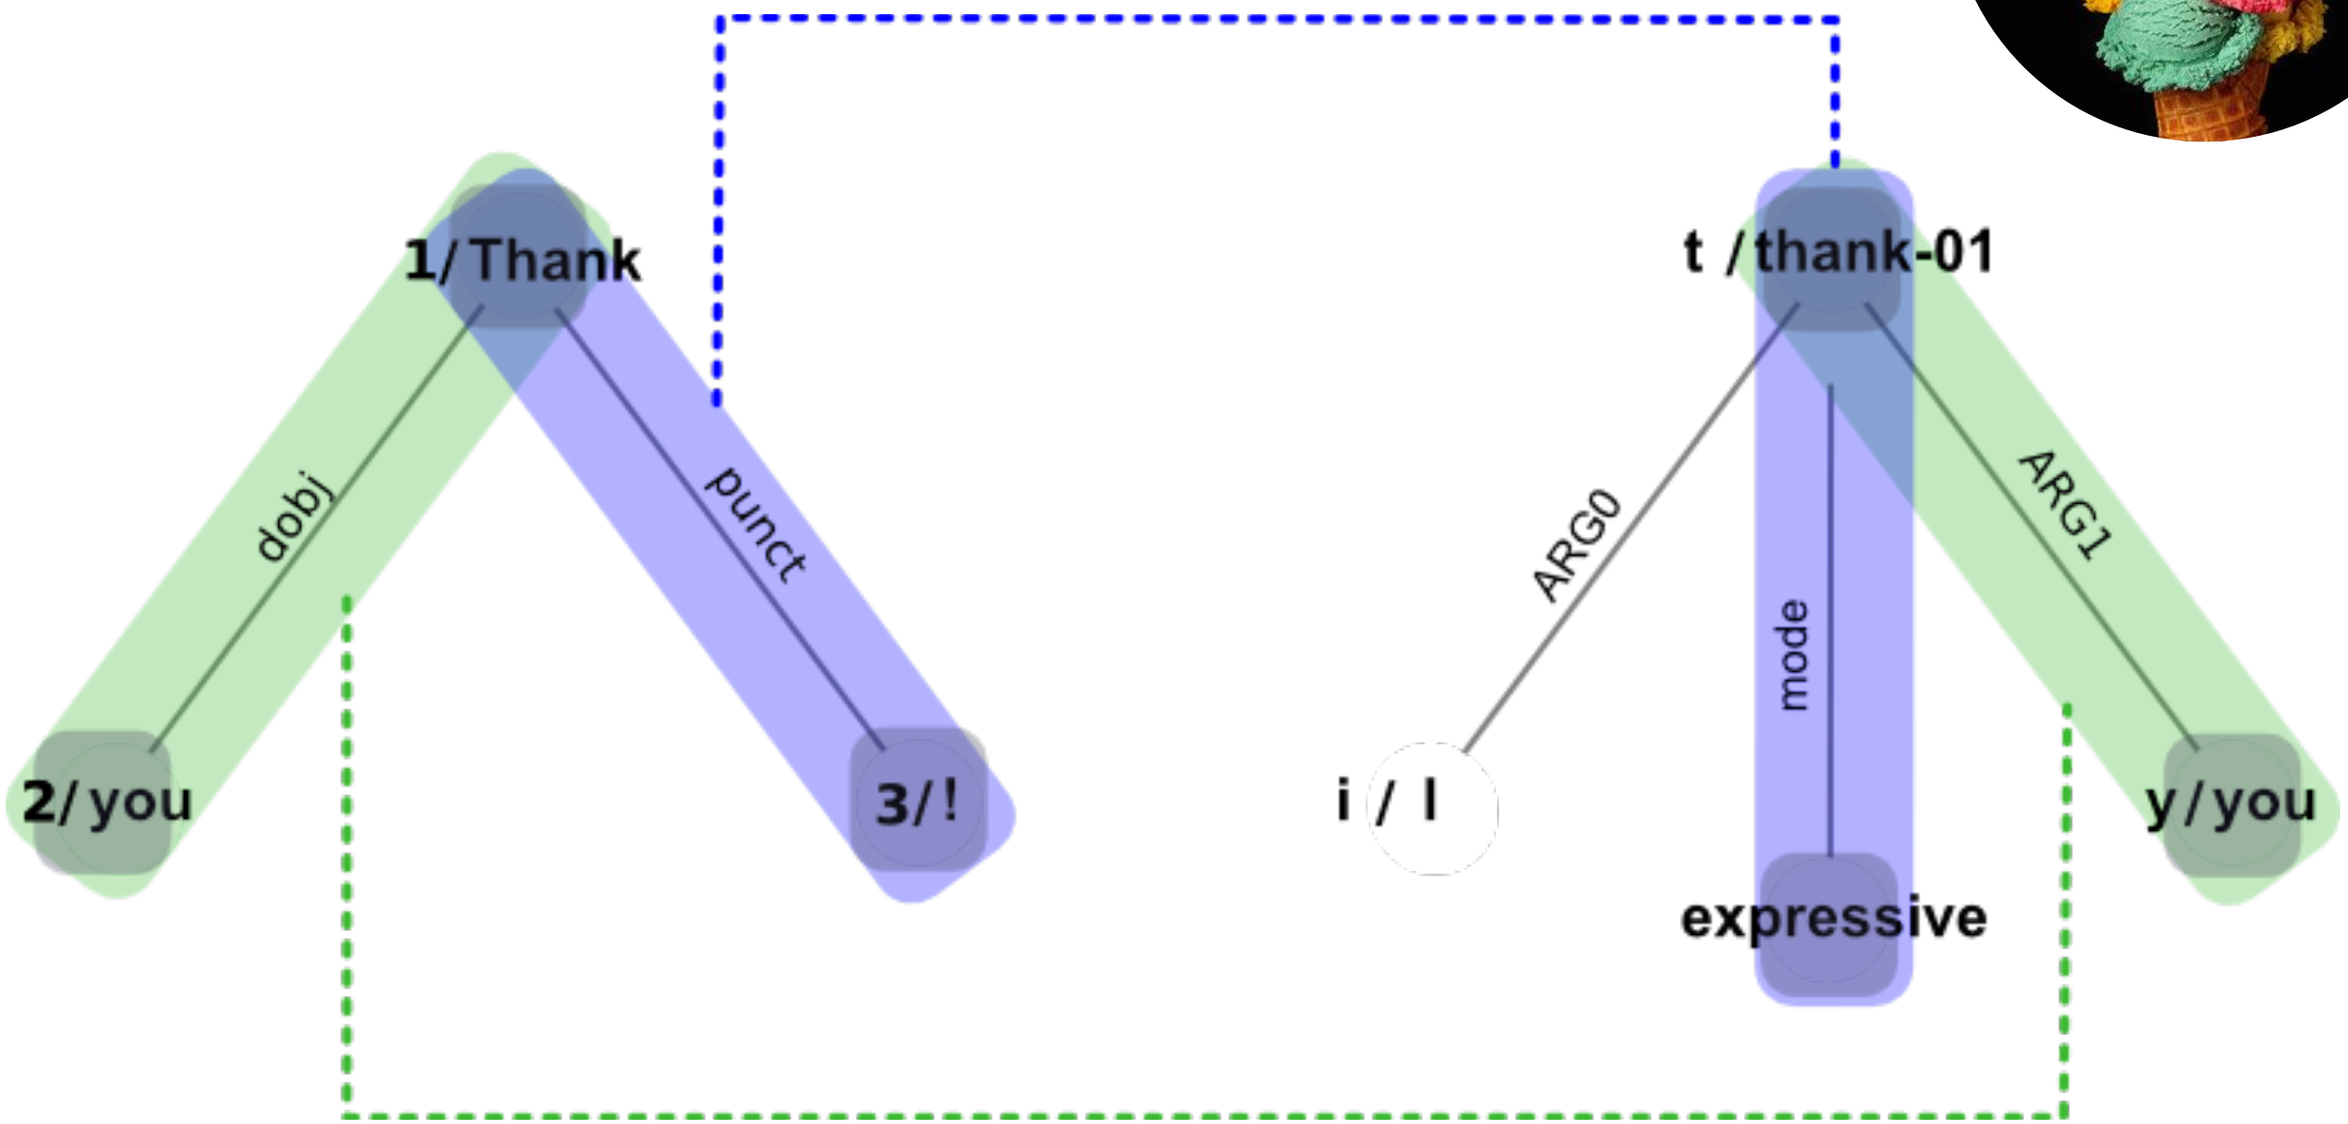

http://tiny.cc/amrud

Advantages of our approach

- **Compositional** syntactic relations between lexical expressions, even if not marked by a function word (subject, object, amod, advmod, compound, …)

- **Subgraphs** preserve contiguity of multiword expressions/morphologically complex expressions (as in JAMR, though we don't require string contiguity)
  - Distinguish from coreference

- Lexical alignments are where to look for spelling overlap; non-lexically-aligned concepts are implicit

- A syntactic edge may attach to different parts of an AMR-complex expression (*tall hunter* vs. *careful hunter*; *bad hunter* is ambiguous). The lexical alignment gives us the *hunt* predicate, while the structural alignment gives us the `person`-rooted subgraph.

Complex configurations indicate structural differences

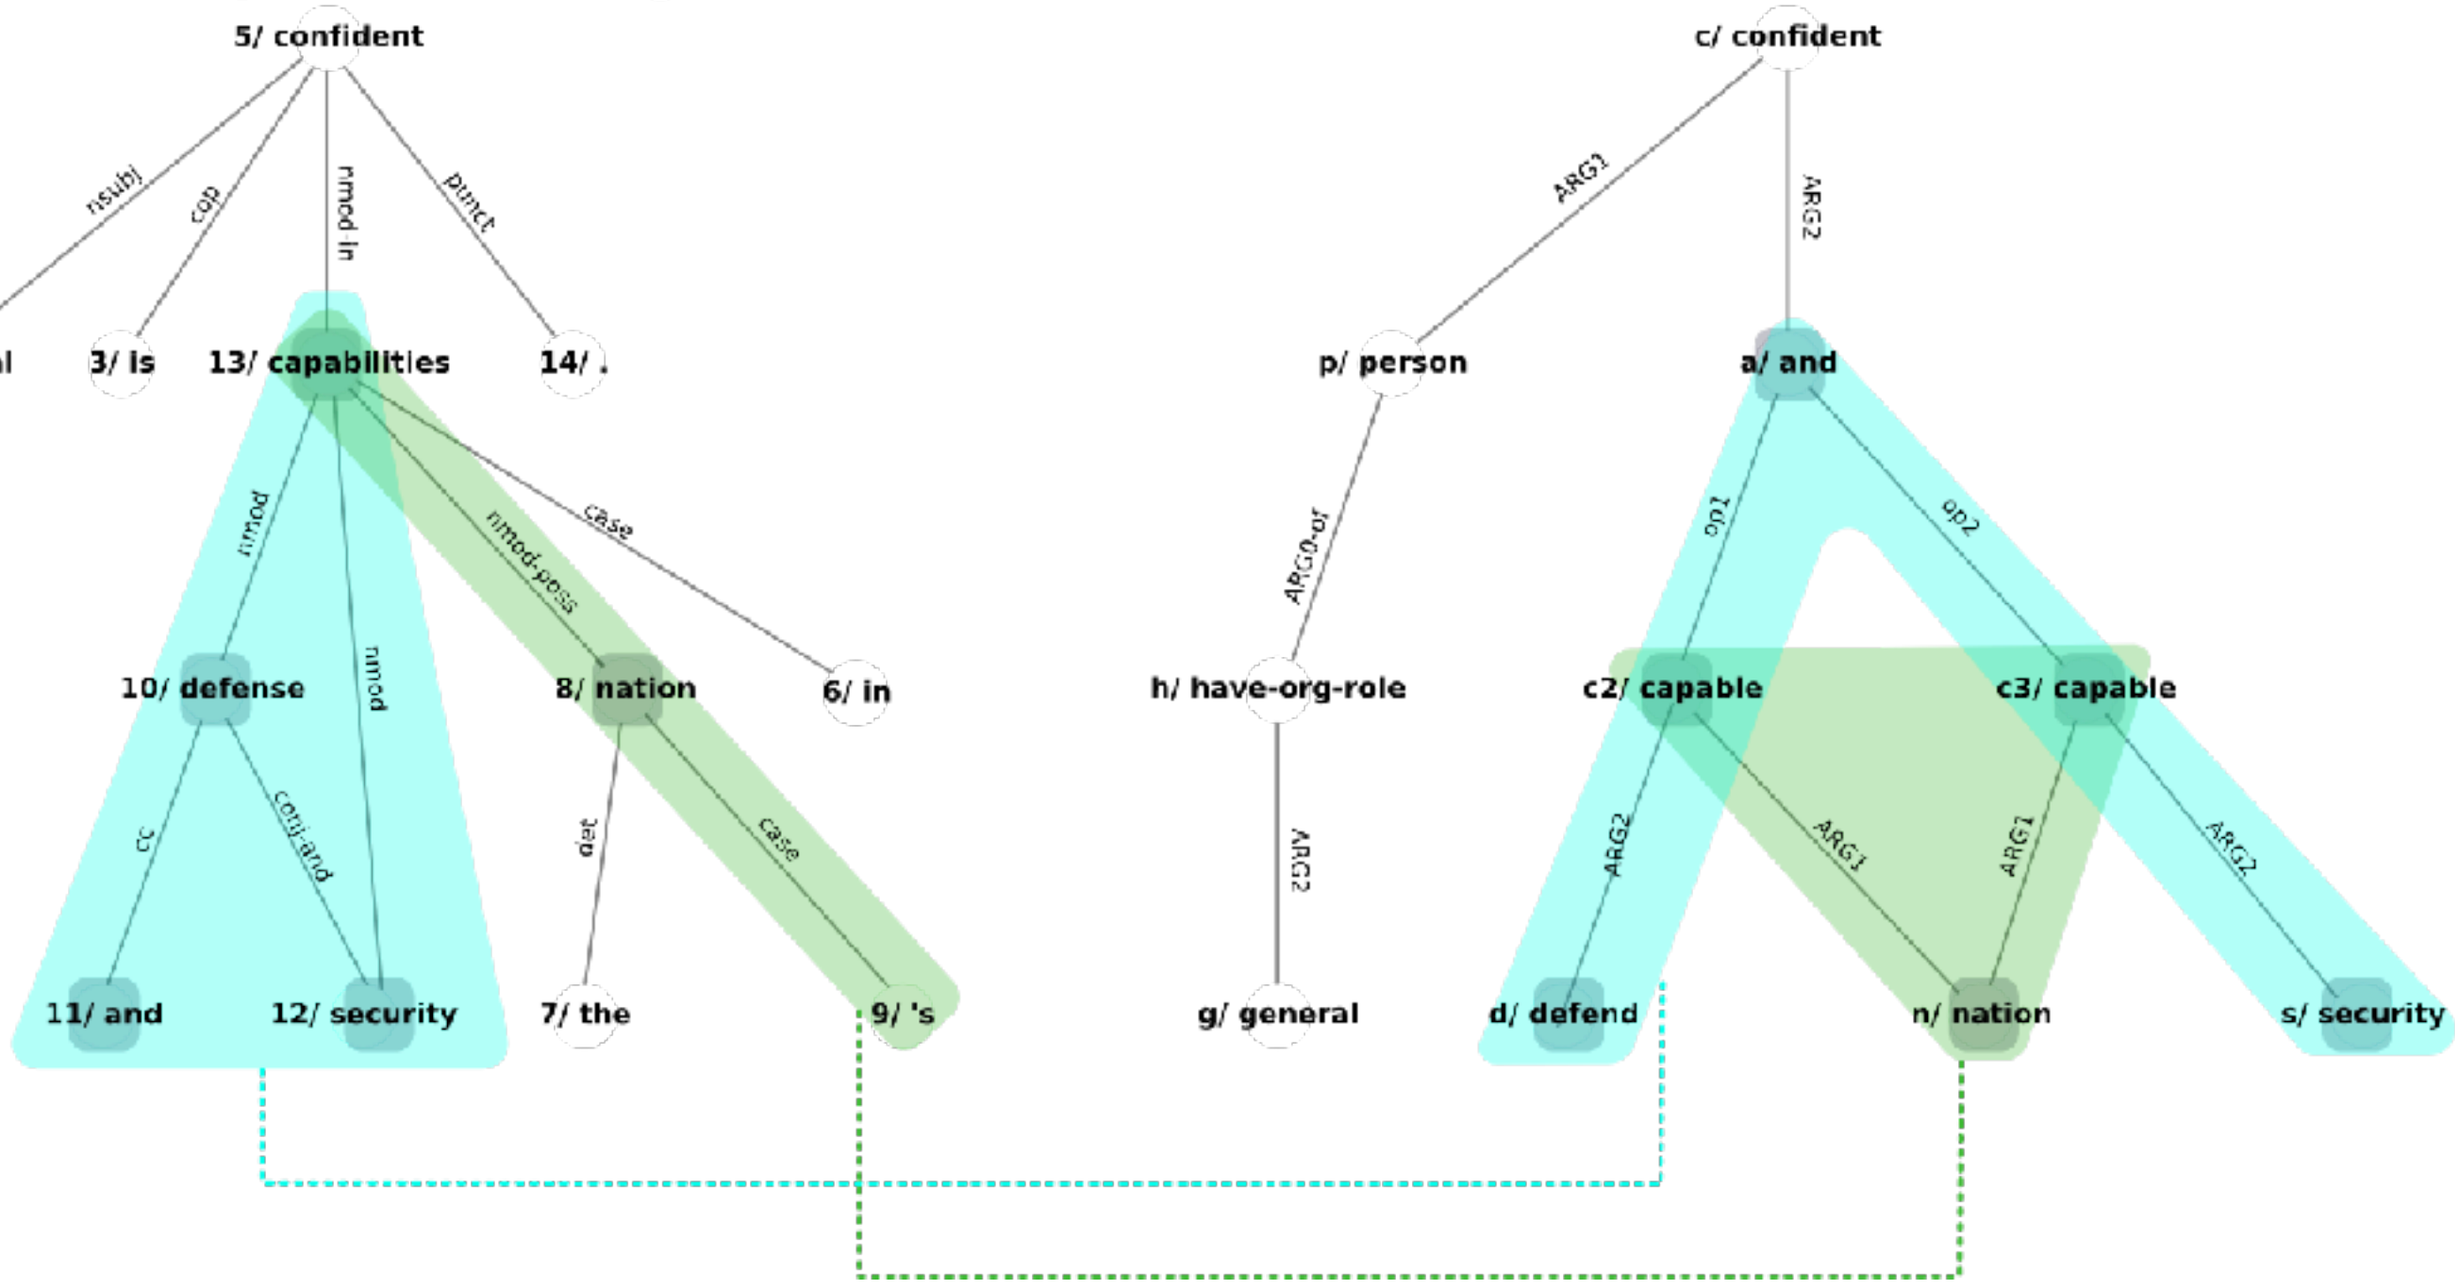

nation's defense and security capabilities

⇒ *nation's defense capabilities and its security capabilities*

Hierarchical alignments

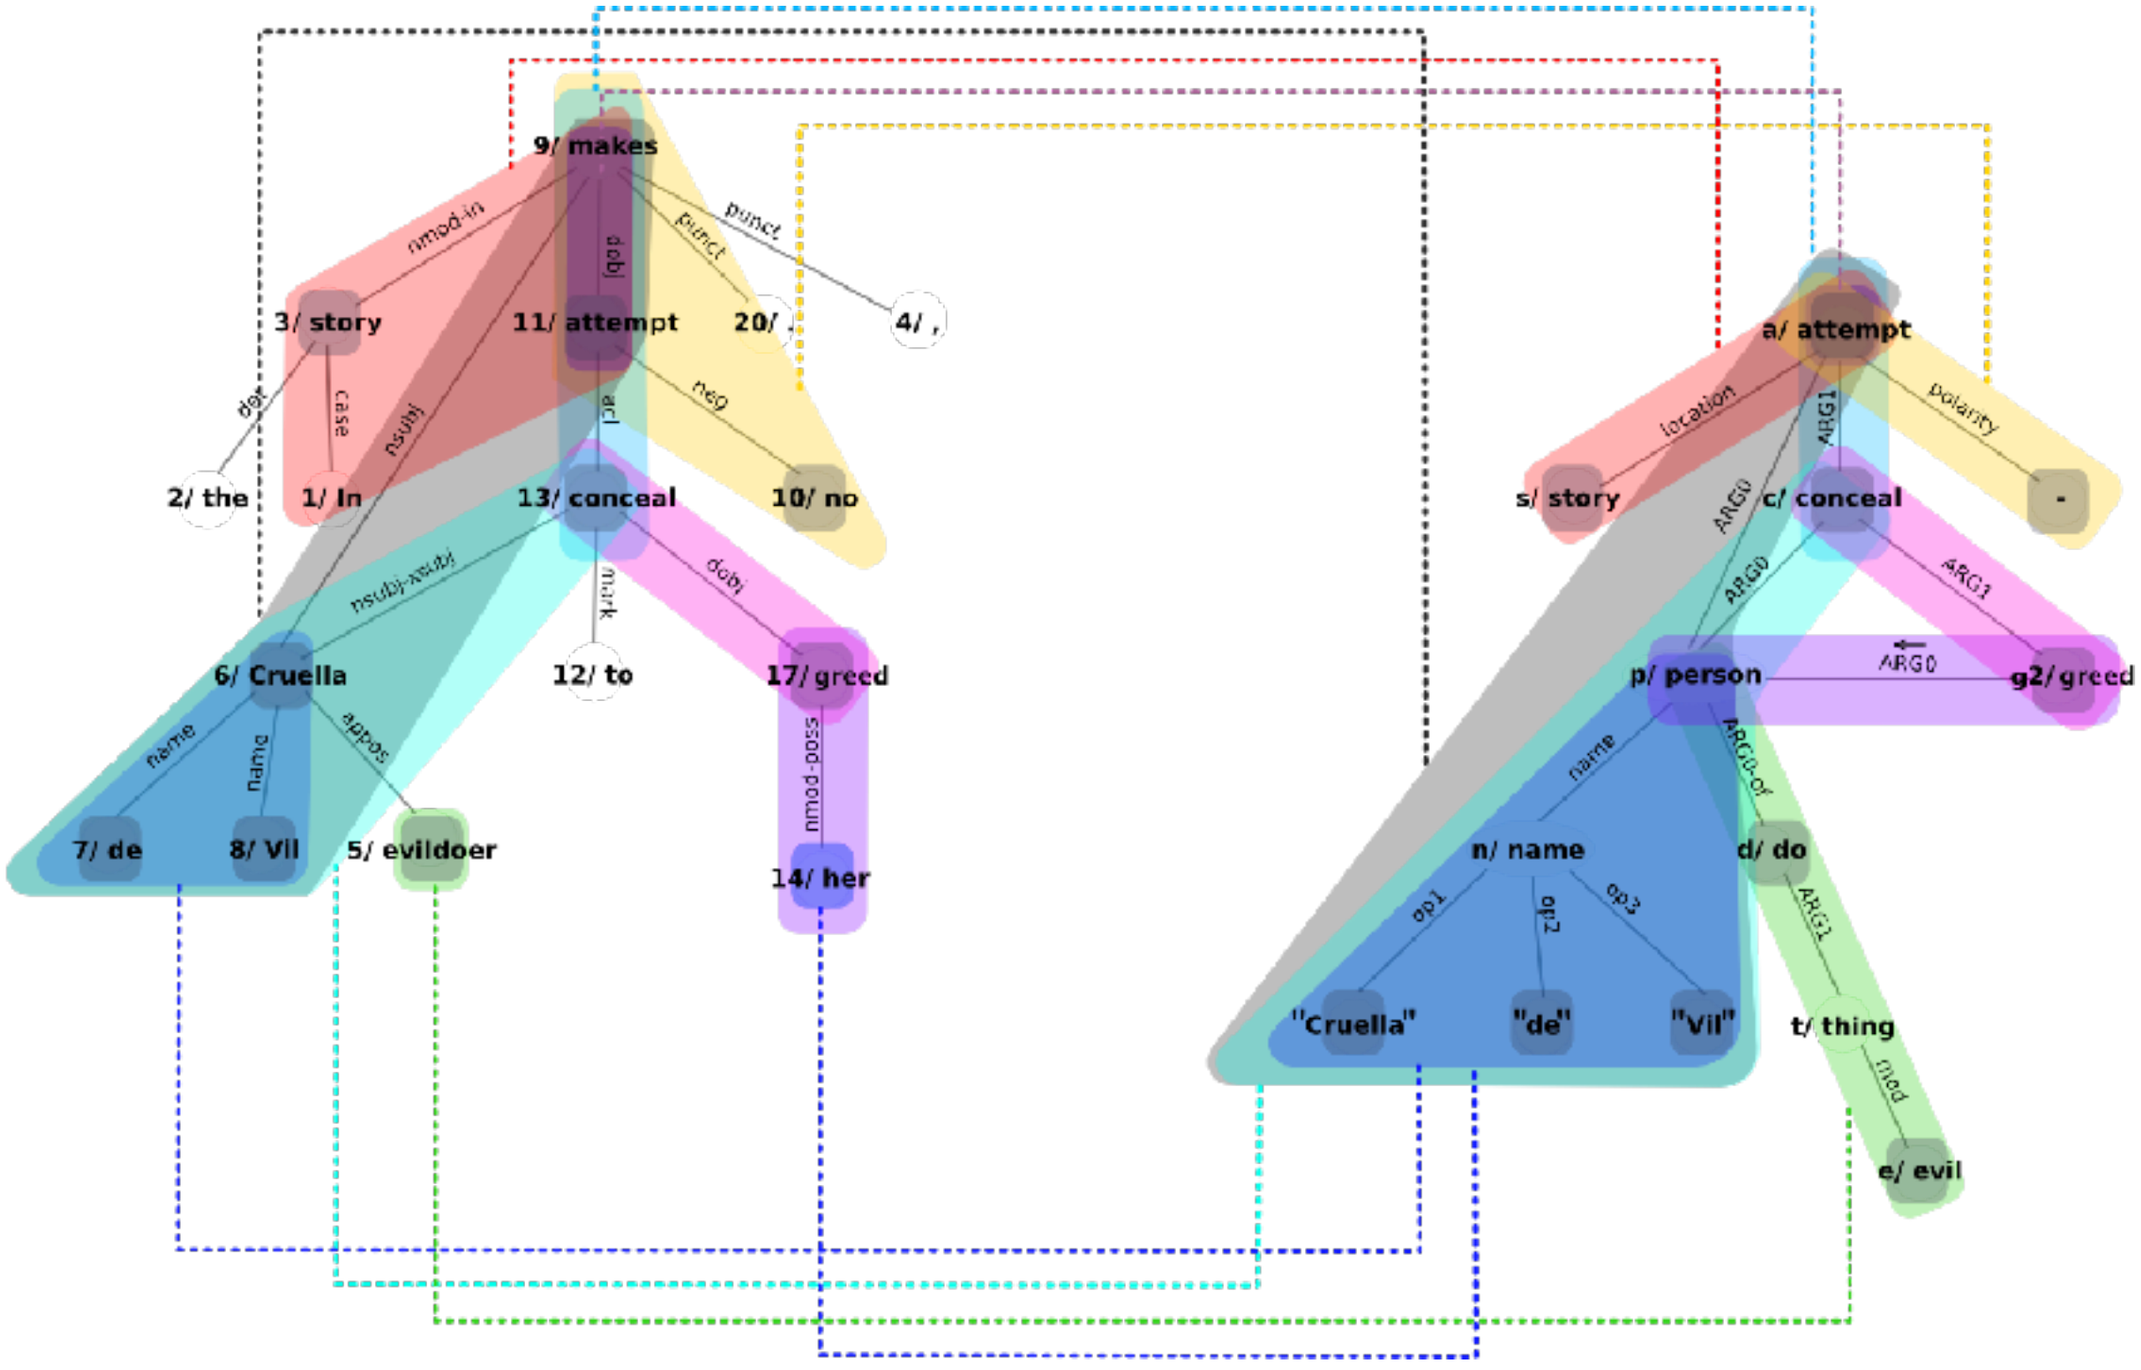

In the story, evildoer Cruella de Vil makes no attempt to conceal her greed.

Named entities + Coreference

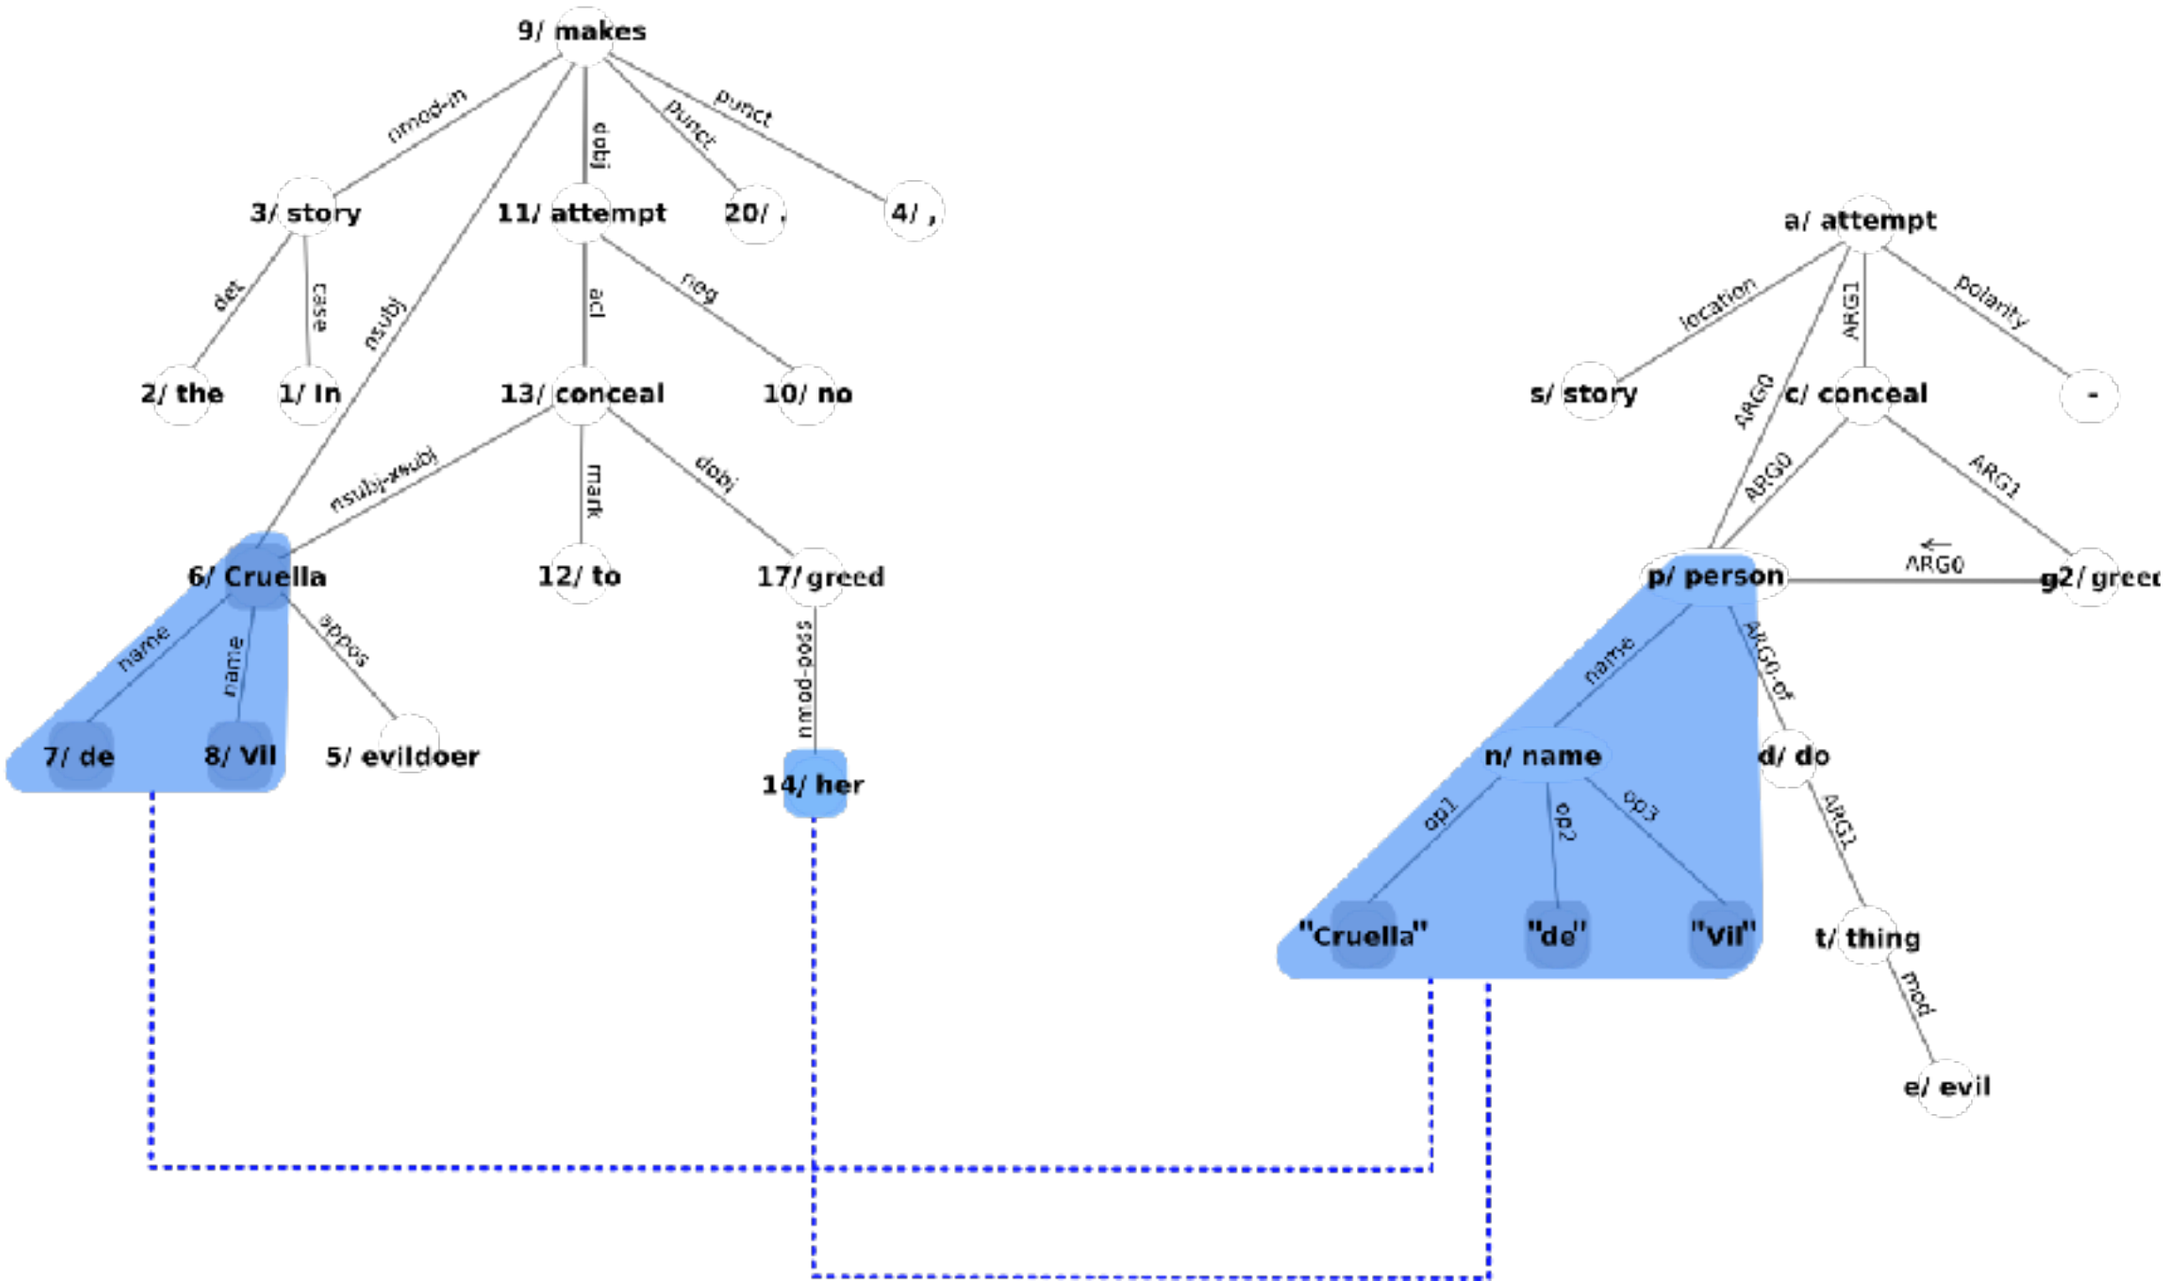

In the story, evildoer Cruella de Vil makes no attempt to conceal her greed.

Light verbs

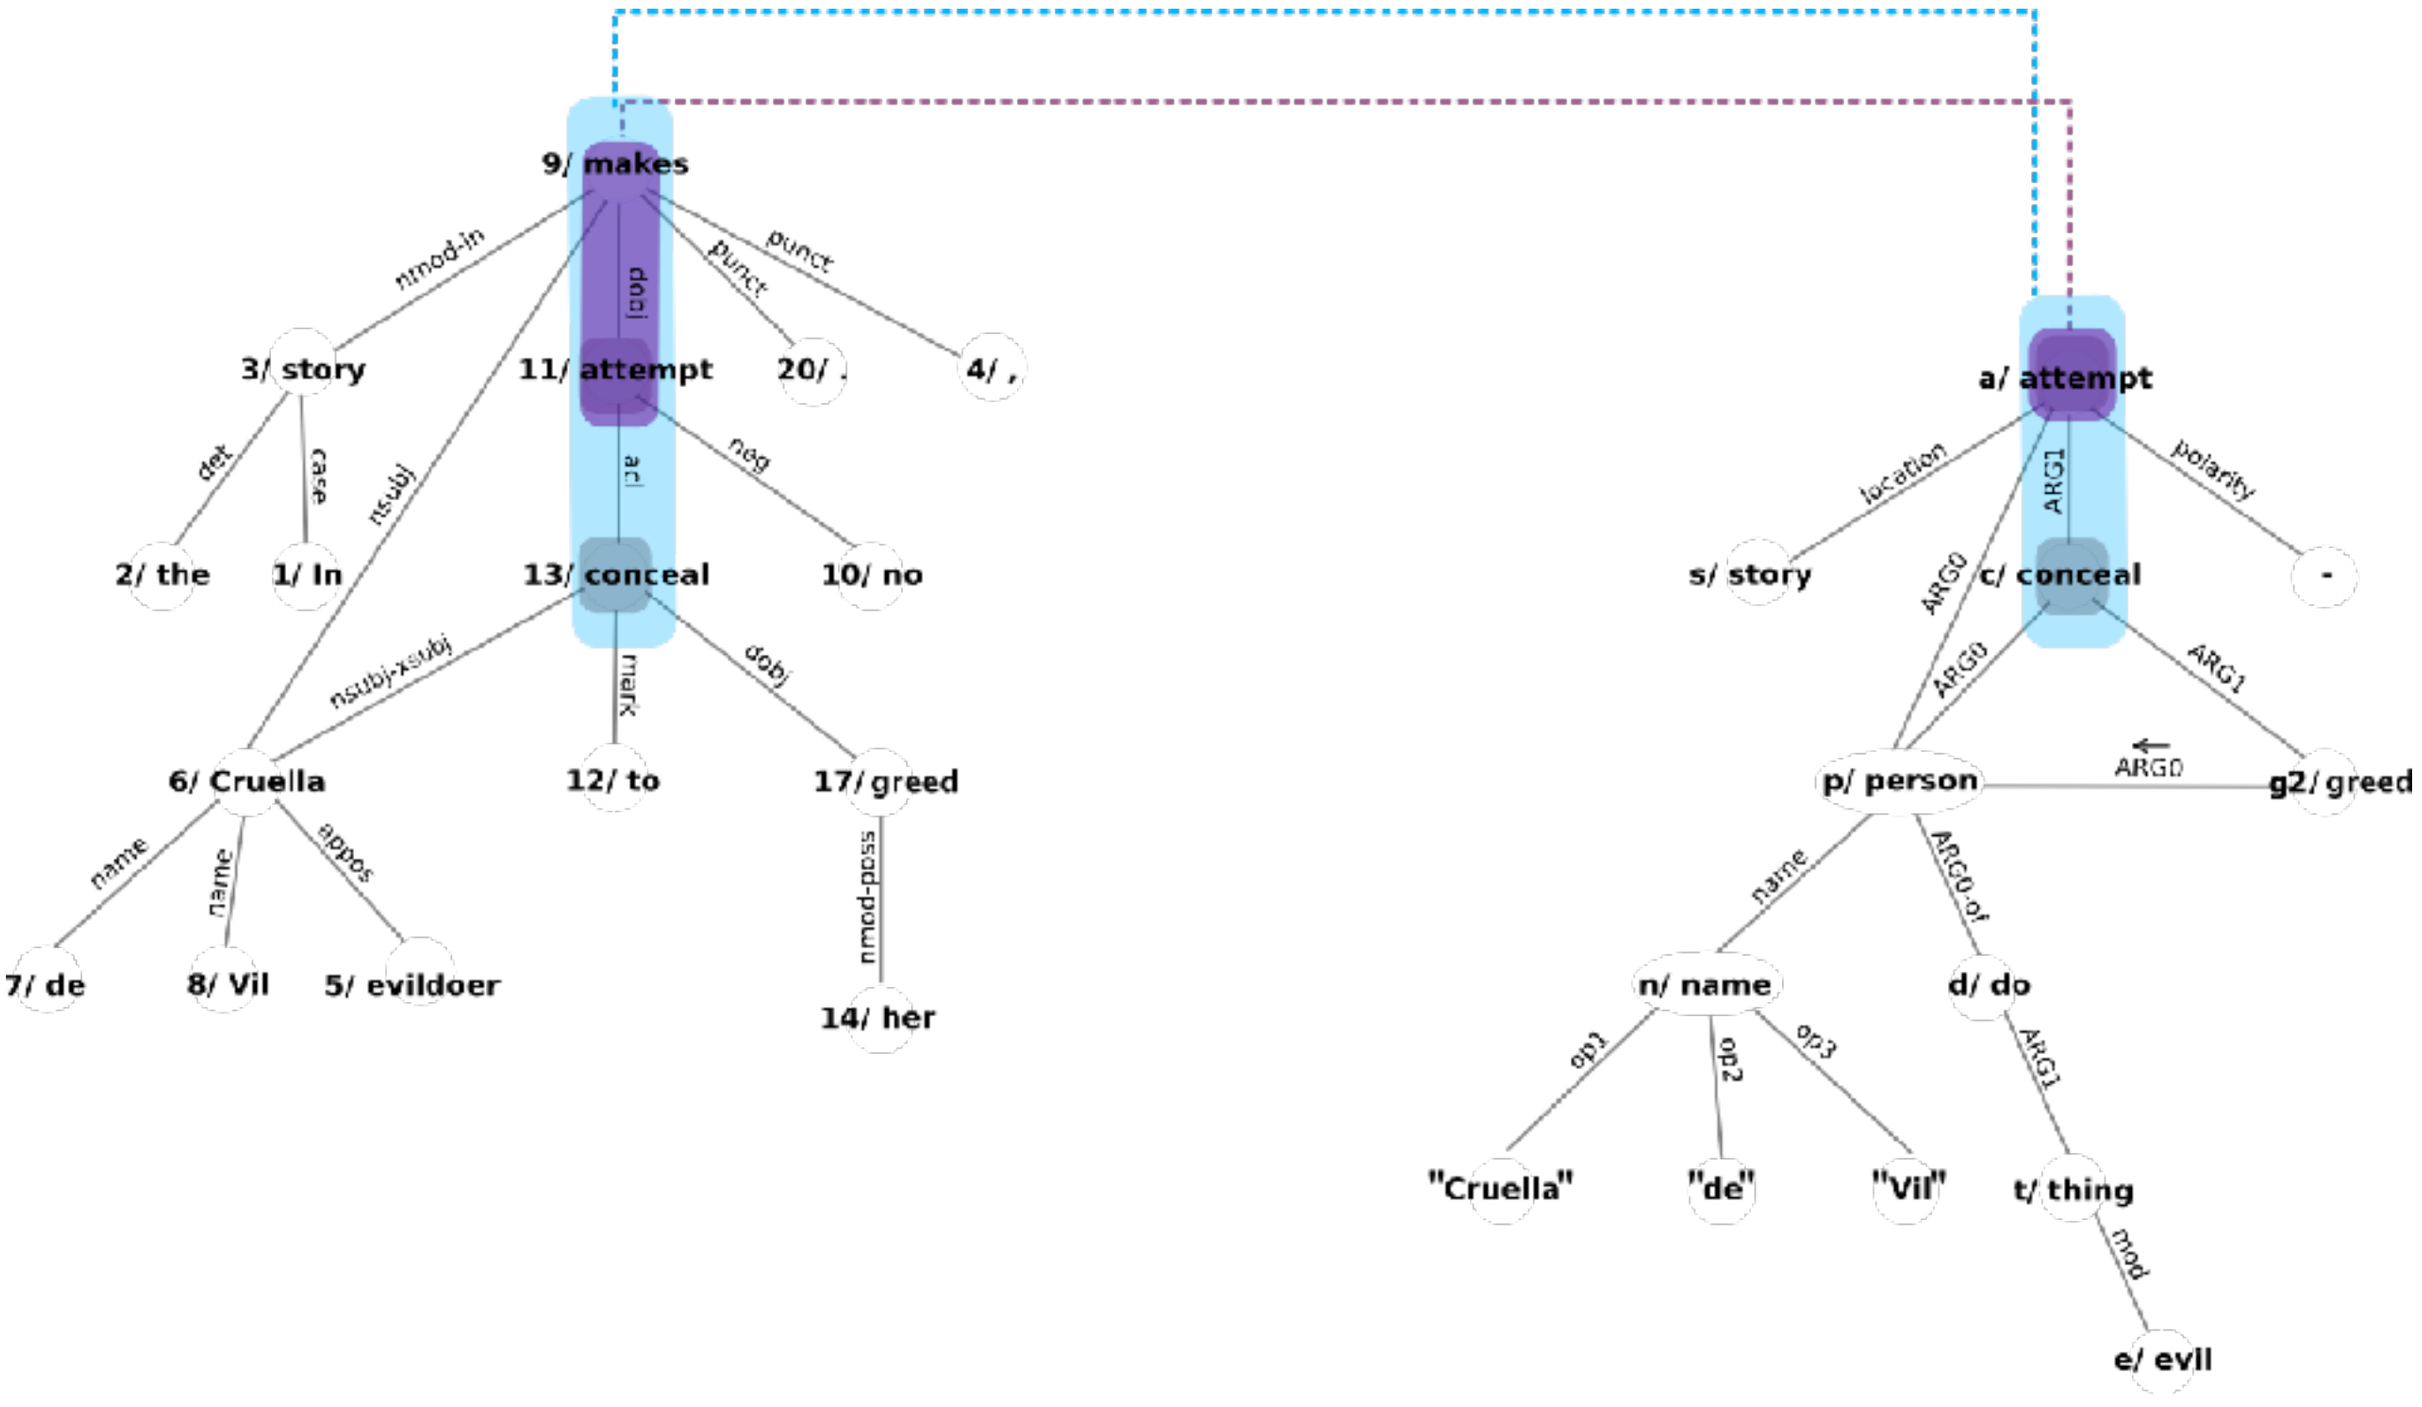
9/ makes
nmod-in
punct
punct
dooj
3/ story
11/ attempt
20/ .
4/ ,
det
case
nsubj
acl
neg
2/ the
1/ In
13/ conceal
10/ no
nsubj-xsubj
mark
dobj
6/ Cruella
12/ to
17/ greed
name
name
appos
nmod-poss
7/ de
8/ Vil
5/ evildoer
14/ her
a/ attempt
location
ARG1
polarity
s/ story
ARG0
c/ conceal
-
ARG0
ARG1
ARG0
p/ person
g2/ greed
name
ARG0-of
op1
op2
op3
n/ name
d/ do
ARG1
"Cruella"
"de"
"Vil"
t/ thing
mod
e/ evil

Control

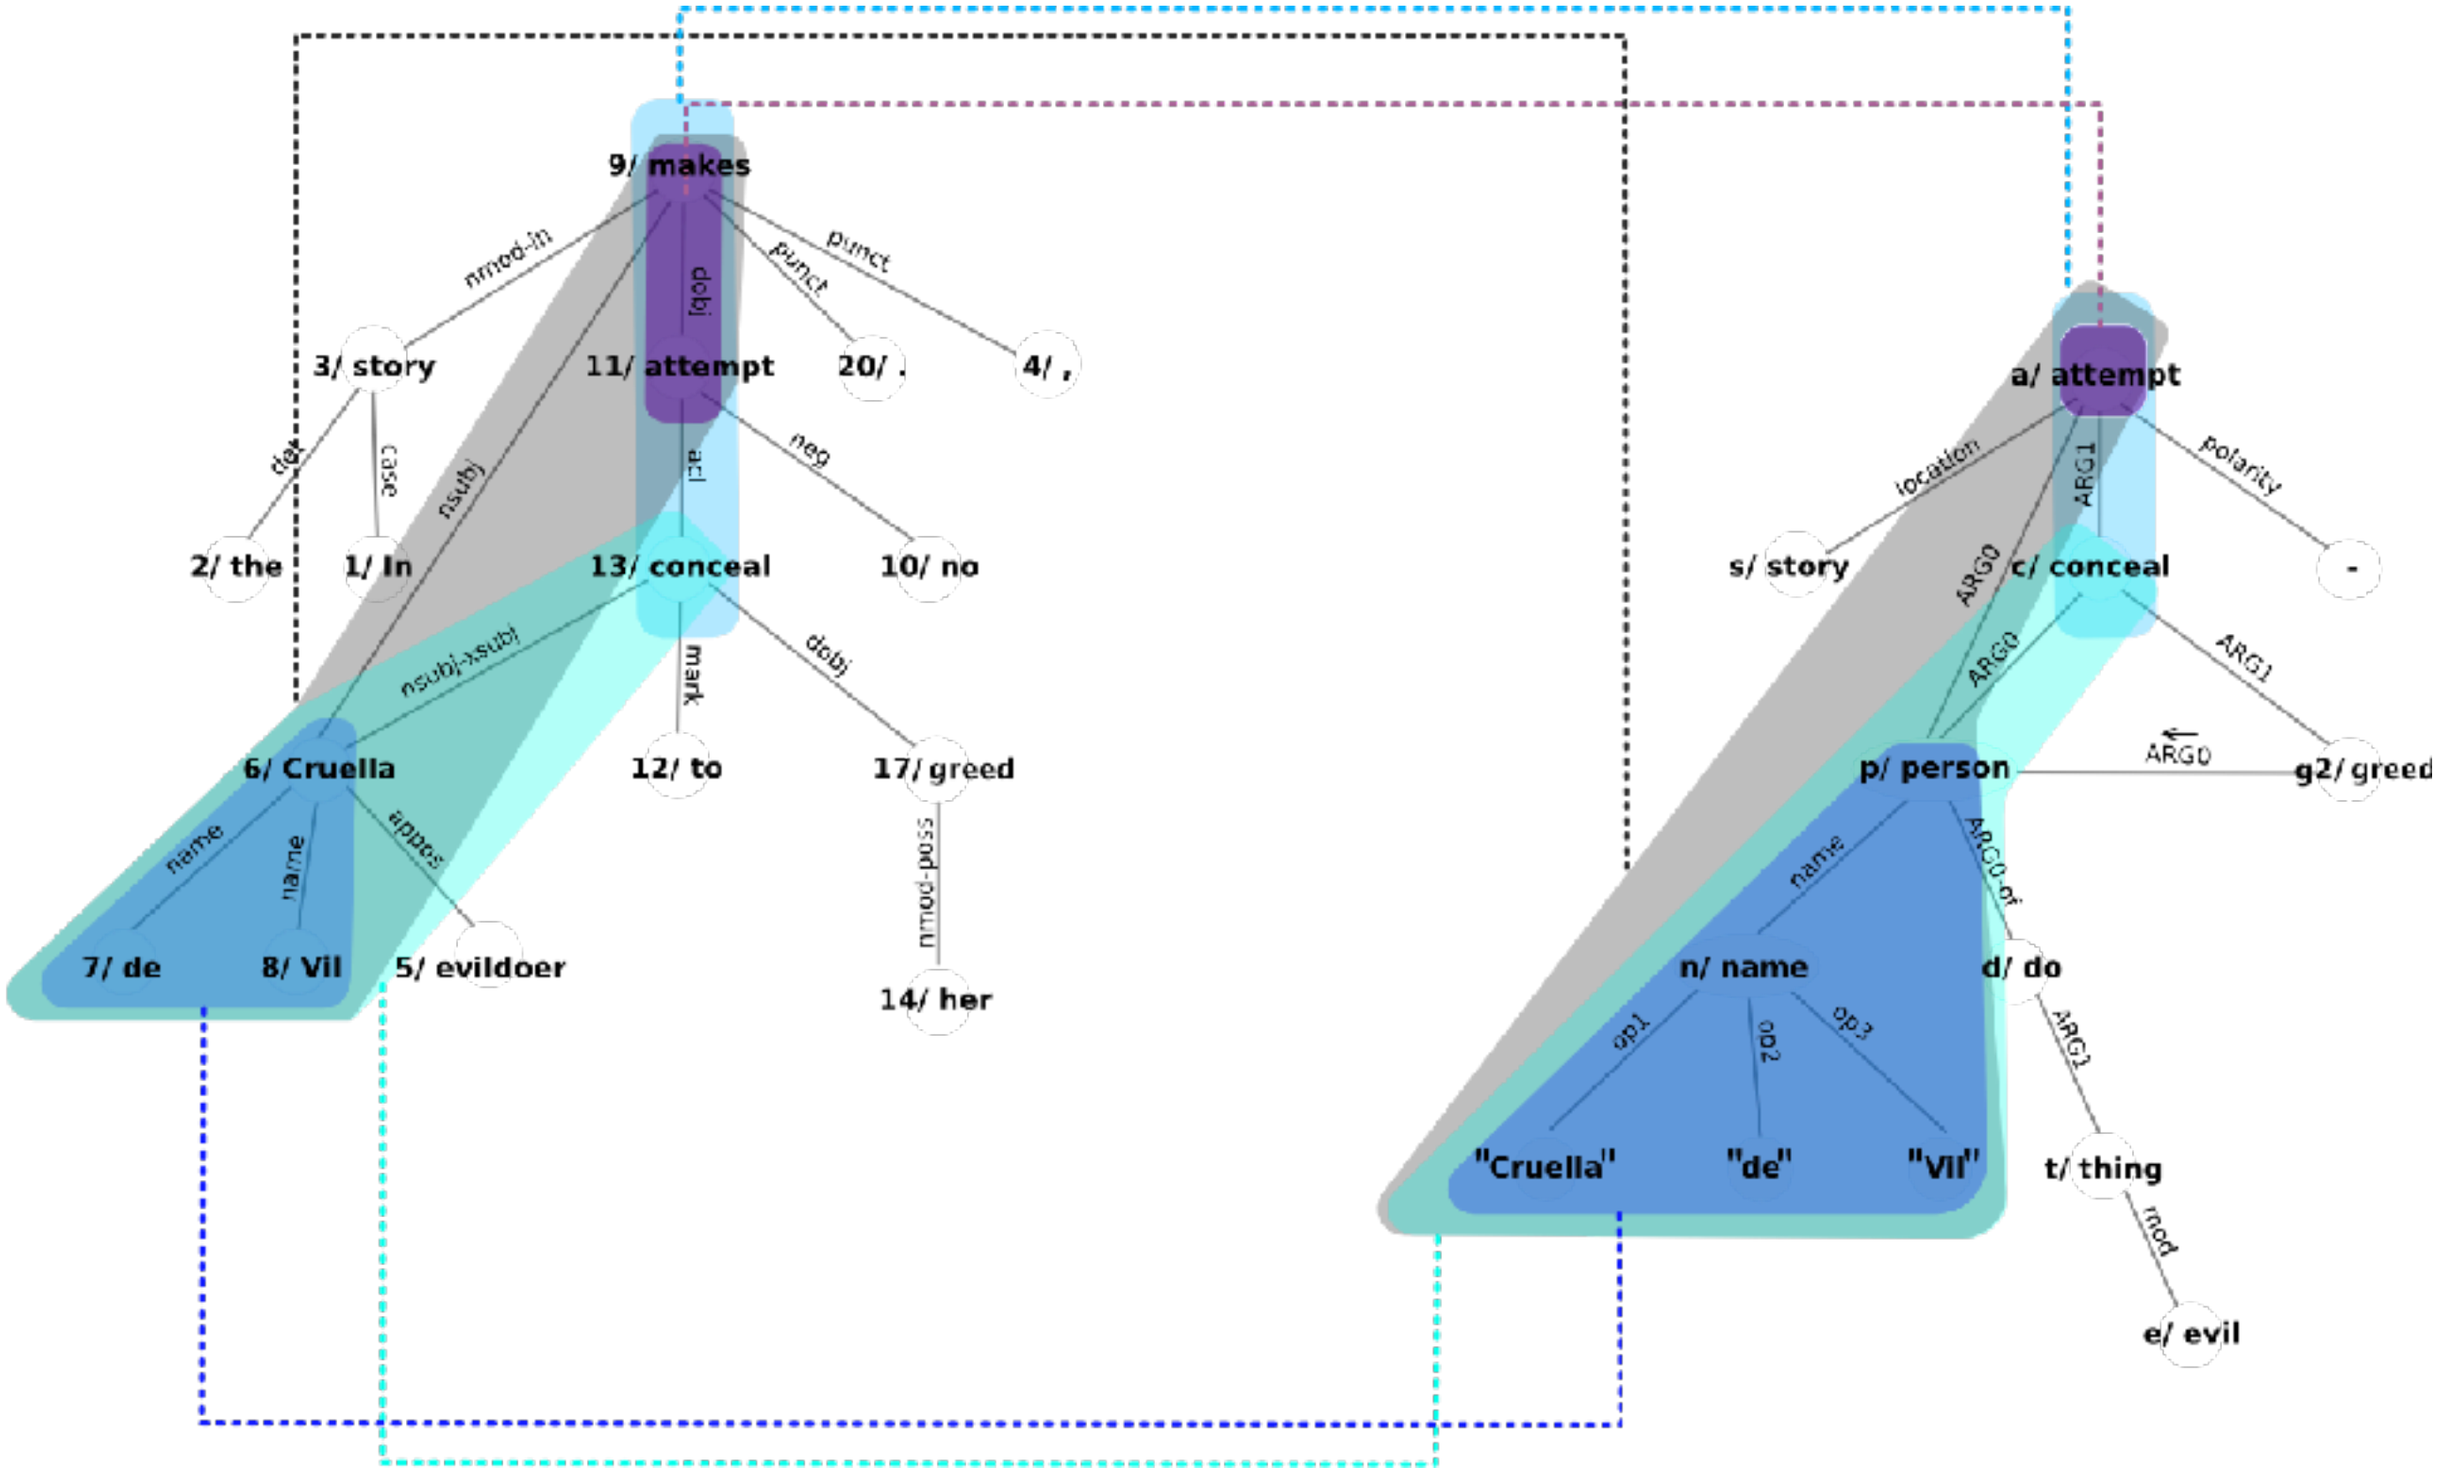
9/ makes
nmod:in
punct
punct
3/ story
11/ attempt
20/ .
4/ ,
dep
case
nsubj
xcomp
2/ the
1/ In
13/ conceal
10/ no
nsubj:xsubj
mark
dobj
6/ Cruella
12/ to
17/ greed
name
name
appos
nmod:poss
7/ de
8/ Vil
5/ evildoer
14/ her
a/ attempt
location
ARG1
polarity
s/ story
ARG0
c/ conceal
-
ARG0
ARG1
ARG0
p/ person
g2/ greed
name
ARG0-of
n/ name
d/ do
op1
op2
op3
ARG1
"Cruella"
"de"
"Vil"
t/ thing
mod
e/ evil

enhanced++ UD annotation

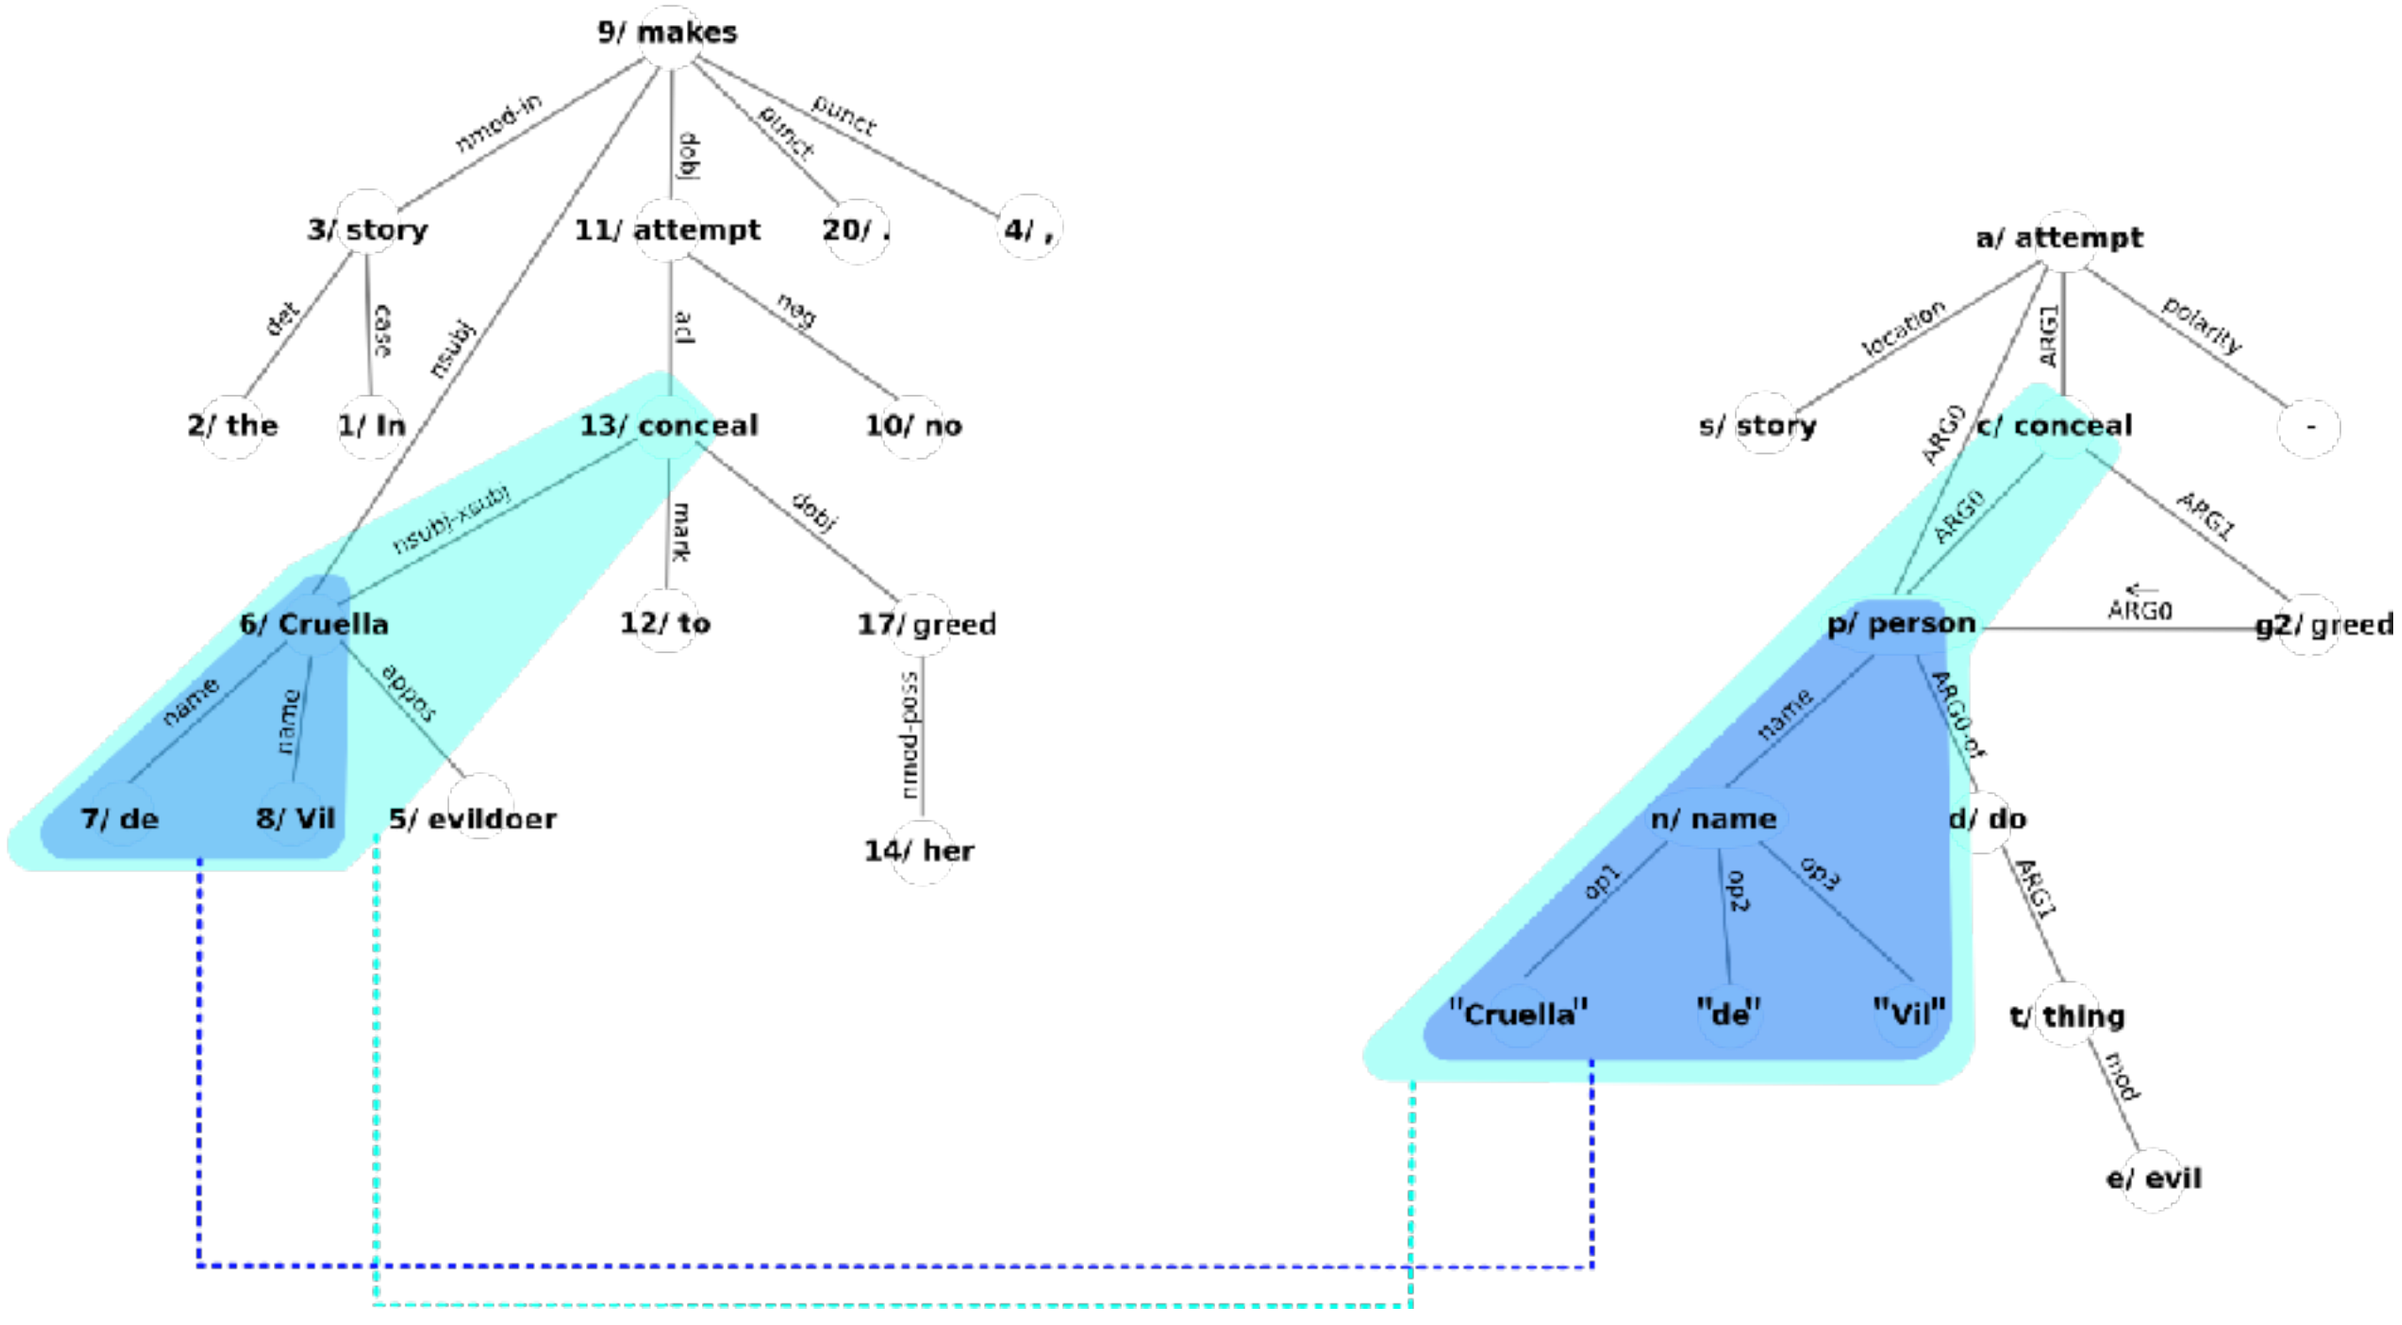

Automatic aligner

- standard label-based node alignment

	dataset								
aligner	our			ISI			JAMR		
our	89	85	**87**	88	77	82	55	81	65
ISI	71	68	70	96	85	**90**	47	67	55
JAMR	86	63	72	95	66	78	92	85	**88**

Table 3: Lexical alignment (precision, recall, F_1-score). Our *lexical* alignment algorithm does not use syntax.

* data used for experiments: our corpus, ISI corpus (Pourdamghani et al., 2014), and JAMR corpus (Flanigan et al., 2014)