

GUMBridge: a Corpus for Varieties of Bridging Anaphora

Lauren Levine, Amir Zeldes

Georgetown University, Department of Linguistics
{lel76, amir.zeldes}@georgetown.edu

Abstract

Bridging is an anaphoric phenomenon where the referent of an entity in a discourse is dependent on a previous, non-identical entity for interpretation, such as in "There is a house. **The door** is red," where the door is specifically understood to be the door of the aforementioned house. While there are several existing resources in English for bridging anaphora, most are small, provide limited coverage of the phenomenon, and/or provide limited genre coverage. In this paper, we introduce GUMBridge, a new resource for bridging, which includes 24 diverse genres of English, providing both broad coverage for the phenomenon and granular annotations for the subtype categorization of bridging varieties. We also present an evaluation of annotation quality and report on baseline performance using open and closed source contemporary LLMs on three tasks underlying our data, showing that bridging resolution and subtype classification remain difficult NLP tasks in the age of LLMs.

Keywords: bridging anaphora, corpus, annotation, English, genre, LLM evaluation

1. Introduction

When an entity is mentioned in a discourse, it is generally either already familiar to the participants or it is new to them. However, there are also instances where an entity is newly introduced to the discourse but the identity of its referent is still inferable to participants. When the referent is inferable specifically due to the presence of a previous, non-identical entity in the discourse, we consider it to be a case of "bridging anaphora". Consider the following sentence:

- (1) While dancing ballet takes dedication and requires serious training, you can learn **the basics** to prepare yourself for **further study**.¹

In (1) above, the bold entities "the basics" and "further study" are newly introduced, but it is inferable that they specifically refer to "the basics of ballet" and "further study of ballet". These entities are bridging anaphora, and "ballet" is their associative antecedent. "Bridging" occurs when a discourse participant constructs an implicature from the entity they are currently processing back to an antecedent entity (Clark, 1975). The associative relations between entities that give rise to bridging links can manifest in a variety of different ways in a discourse, including part-whole relations (a house → **the door**), relative adjectives (a dog → **a larger dog**), and prototypical associations (a library → **the books**). The correct interpretation of such relations is essential for a variety of NLP tasks, including question-answering, summarization and more, and it is an open question to what extent current

LLMs can resolve them (see Section 5).

Although bridging is relatively understudied when compared with other entity related tasks such as named entity recognition (NER) and coreference resolution, bridging resolution – the automatic identification of bridging anaphora and resolution back to their associative antecedents – has gained attention in recent years (Kobayashi et al., 2023; Hou, 2020). Additionally, efforts have been made to include bridging anaphora in recent shared task settings (Khosla et al., 2021; Yu et al., 2022). However, despite some attempts at harmonization (Levine and Zeldes, 2024), the current resource landscape for bridging remains disparate, containing a number of different annotation formalisms with varying definitions of bridging and annotation schemes for subcategorization (Kobayashi and Ng, 2020).

Since its introduction, bridging has been studied from a variety of theoretical perspectives (e.g., Hawkins, 1978; Prince, 1981; Asher and Lascarides, 1998; Baumann and Riester, 2012), and linguistic resources have been constructed for various languages, including German (e.g., GRIN (Schweitzer et al., 2018) and DIRNDL (Eckart et al., 2012)), Polish (e.g., PCC (Ogrodniczuk and Zawisławska, 2016)), and Czech (e.g., PDT (Nedoluzhko et al., 2009)). In English, there are several influential resources which have been used for the evaluation of bridging resolution: ISNotes (Markert et al., 2012), BASHI (Rösiger, 2018), and ARRAU RST (Poesio and Artstein, 2008; Uryupina et al., 2019). ARRAU annotates instances of bridging by identifying entity pairs that establish cohesion, using a set of predefined semantic relations, while the bridging annotations of ISNotes and BASHI are based on the information status of entities, regarding bridging as a type of mediated information. Such differences complicate setting

¹Bridging anaphora are marked in bold face, and their associative antecedents are underlined.

up a standard benchmark for evaluation, as it is difficult to make a meaningful comparison of results of bridging resolution systems on different resources. Additionally, these resources are mostly small, severely lacking in genre diversity, and/or provide coverage of the phenomenon that is limited in various ways.

In order to bridge this gap in the landscape of bridging resources, we present GUMBridge, a new resource for bridging, which contains **5.7k bridging instances** over **291k tokens** of English, covering **24 diverse genres**. GUMBridge is constructed on top of GUM v12, an existing multi-genre corpus of English (Zeldes, 2017), and it aims to unite aspects of current bridging formalisms for English. GUMBridge provides broad coverage for the phenomenon of bridging, using rigorously defined information status based criteria to identify instances of bridging (similar to ISNotes and BASHI), which does not place structural limitations on the manner in which bridging can manifest in a discourse. Additionally, GUMBridge provides annotations for categorizing subtypes of bridging (as in ARRAU), introducing a new schema with 10 subtypes divided into 3 main categories: COMPARISON, ENTITY, and SET (with an additional OTHER category). GUMBridge also allows for the annotation of multiple subtypes on a single instance of bridging, which has not been possible in any previous English resource. Leveraging this data, we provide a baseline evaluation of three LLMs on the task of bridging resolution, with the additional task of subtype categorization. Our results show that bridging resolution remains a difficult task, even for state-of-the-art large language models.

2. Background

As noted above, the current landscape of bridging resources is heterogeneous, with resources covering the scope of the phenomenon differently. In this section, we discuss the relevant theoretical divides and gaps in coverage in the landscape of English bridging resources. We give details on the ISNotes, BASHI, and ARRAU corpora, which are commonly used for the evaluation of bridging resolution systems (Yu et al., 2022; Kobayashi et al., 2023). We also give details on the original GUM corpus, the resource on top of which GUMBridge is constructed.

Information Status & Bridging The information status (IS) of an entity refers to the extent to which that entity is familiar to the participants of a discourse (Nissim et al., 2004). Broadly speaking, an entity can either be "New" information, where the entity is not yet known to the participant, or it can be "Given" information, where the entity is

	Tokens	Bridging Instances	Bridging per 1k Tokens
ISNotes	40k	916	22.7
BASHI	58k	459	7.9
GUM v11	268k	2.2k	8.2
ARRAU			
RST	229k	3.8k	16.5
GNOME	21k	692	32.2
PEAR	14k	333	23.7
TRAINS	84k	710	8.5
Total	348k	5.5k	15.8
GUMBridge v1			
Train	213k	4k	18.9
Dev	30k	732	24.5
Test	30k	562	18.6
Test2	18k	379	21.2
Total	291k	5.7k	19.6

Table 1: Comparison of the size of several bridging resources for English.

recognized by the participant. If an entity is recognized, this can be for various reasons: (1) the entity may have been already introduced in the discourse (coreference), (2) the entity may be inferable via generic/world knowledge or deixis (pointing to something in context), or (3) the entity may be inferable from a previous, non-identical entity in the discourse, which is the case of bridging anaphora. Although entities understood through bridging and generic/world knowledge are both considered "Accessible" (as they are recognized by discourse participants when first introduced to the discourse), only instances of bridging depend on the presence of an associative antecedent for comprehension.

Lexical Bridging vs. Referential Bridging The distinction between lexical bridging and referential bridging is a theoretical divide that stems from the decision of most corpora to classify bridging as an anaphoric phenomenon based on information status. When introducing the BASHI corpus, Rösiger (2018) coined the terms for this distinction as a means of differentiating the types of bridging annotations present in different corpora. Referential bridging refers to truly anaphoric instance of bridging where the bridging anaphor requires an antecedent to be interpretable, as in (2):

- (2) She likes the house because **the windows** are large.

Lexical bridging refers to lexical semantic relations between pairs of entities, such as part-whole or set-member relations, which may or may not be anaphoric, as in (3) where the antecedent is not strictly necessary for interpretation:

- (3) I went to the United States last month. My

first stop was **Washington, DC**.

While ARRAU includes instances of both lexical bridging and referential bridging, ISNotes and BASHI both limit themselves to including only referential bridging. From an information status based perspective, instances of lexical bridging are not considered bridging at all if they are not also anaphoric in nature. As such, resources such as ARRAU, which include various types of lexical bridging, have received criticism for being over-annotated (Roesiger et al., 2018).

Corpora The following are short descriptions of the resources which have most strongly informed the construction of GUMBridge. A comparison of the size of these resources in tokens and instances of bridging is shown in Table 1.

ISNotes is a corpus for annotations of fine-grained information status, and it is composed of 50 Wall Street Journal (WSJ) documents (40k tokens). Considering bridging to be a type of mediated information, it contains 663 instances of bridging in the `mediated/bridging` subcategory and an additional 253 instances of comparative anaphora (commonly considered to be a subvariety of bridging) in the `mediated/comparison` subcategory (e.g. cases like **a larger dog** above), for a total of 916 instances.

BASHI is a corpus that specifically targets the annotation of bridging anaphora, and it also covers a set of 50 WSJ documents (58k tokens, different documents than ISNotes). As noted above, Rösiger (2018) introduces the distinction between referential and lexical bridging, and BASHI explicitly includes only instances of referential bridging in its construction. It contains just 459 annotated bridging pairs.

ARRAU is composed of four subcorpora, covering various genres: ARRAU RST (WSJ news), ARRAU PEAR (spoken narrative), ARRAU TRAINS (dialogues), and ARRAU GNOME (medical/art history). However, the annotations for the subcorpora are not consistent, and the largest subcorpus, ARRAU RST (229k tokens) is the only one that has been used in evaluation of bridging resolution systems. It includes 3,777 of the 5,512 bridging annotations in ARRAU. ARRAU's bridging annotation links related mentions that establish entity coherence through non-identity relations. Because this definition is intentionally broad, annotation is constrained to a fixed set of semantic relations.

GUM (291k tokens, v12) is a resource that covers a variety of linguistic phenomena for 16 diverse genres of English: academic writing, biographies, courtroom transcripts, essays, fiction, how-to guides, interviews, letters, news, online forum discussions, podcasts, political speeches, sponta-

neous face to face conversations, textbooks, travel guides, and vlogs. The extended test set of the corpus, GENTLE (Aoyama et al., 2023), is included in GUM v12 as the 'test2' partition, and contains an additional 8 genres: dictionary entries, live esports commentary, legal documents, medical notes, poetry, mathematical proofs, course syllabuses, and threat letters.

Version 11 of the GUM corpus had a total of 2,192 bridging instances (both lexical and referential bridging). Although GUM v11 already contained annotations for bridging, the annotation guidelines relied upon were minimal (Zeldes, 2024), and GUM-Bridge aims is to broaden the coverage and robustness of the bridging annotations for GUM v12. In addition to the standalone data release accompanying this paper, the annotations of GUMBridge are included in the GUM v12 release,² and GUM-Bridge style annotations will continue to be included in future GUM releases.

Gaps in the Resource Landscape From the descriptions of the English bridging resources above, we can see that no single one of them can be considered a reasonable "reference" corpus for bridging in English. First, the number of bridging instances present in existing resources is relatively low, as we can see in Table 1. The resources that exclusively contain referential bridging are quite small, with ISNotes and BASHI containing only 916 and 459 instances of bridging respectively. And while ARRAU contains more instances of bridging (5,512), it is likely to be "over annotated" for studies only concerned with referential bridging, and it is not trivial to separate referential and lexical bridging in the data. Additionally, we can see in Table 1 that the frequency of bridging instances annotated in existing resources has been quite variable. ARRAU and ISNotes have approximately twice the frequency of bridging annotations per 1k tokens when compared with BASHI and GUM v11 (15.8 and 22.7 vs 7.9 and 8.2 respectively). This suggests that existing resources are likely under-annotated, and that broader coverage for bridging annotations is required.

There is also a complete lack of genre diversity amongst the resources that have commonly been used for the evaluation of bridging resolution systems; ISNotes, BASHI, and ARRAU RST are all solely composed of WSJ data from more than 30 years ago, which predates basic modern expressions such as "cell phone" and "the internet". And while ARRAU provides some genre coverage in its sub-corpora, they are unbalanced, both in their size and in the density of their bridging annotations (see Table 1). As a result of their internal inconsistency,

²<https://github.com/amir-zeldes/gum>

they are not commonly used for evaluation. Additionally, ARRAU is the only resource that provides subtype annotations, meaning that more granular linguistic analysis is not possible for the resources focused on referential bridging. These gaps motivate the current paper’s introduction of GUMBridge, a new resource which provides a substantial increase in both genre diversity and annotation coverage, as well as the addition of multi-subtype annotations for every instance of bridging.

3. Resource Description

GUMBridge is a new resource for varieties of bridging anaphora, constructed on top of the existing GUM corpus. The bridging annotations for GUMBridge v1 are released with this paper under the Creative Commons Attribution (CC BY-NC-SA) version 4.0 license via our GitHub repository.³ GUMBridge contains 5,698 instances of bridging over 291k tokens, corresponding to 19.6 instances per 1k tokens. As we can see in Table 1, after ISNotes, this is the highest overall coverage of the resources discussed. GUMBridge follows the train/dev/test/test2 splits established by GUM v12. Although we do not leverage our splits for model training in this paper, we provide official, genre balanced splits in order to promote consistency in future work.⁴ The size of the splits in terms of tokens and bridging instances is shown in Table 1.

As previously noted, GUMBridge aims to unite desirable aspects of existing bridging resources: like ISNotes and BASHI, we take an information status informed approach to identifying instances of bridging, and like ARRAU, we provide subtype annotations for each instance of bridging, which allows for the analysis of different varieties of bridging anaphora. In the following subsections, we give details on GUMBridge’s information status based annotation procedures for train/dev/test/test2 and its subcategorization schema for varieties of bridging.

3.1. Annotation Procedure

The dev and test sets of GUMBridge were manually annotated by the authors of this paper and graduate students trained on the annotation guidelines below.⁵ The annotation guidelines for GUMBridge were developed after careful review of the bridging definitions, annotation guidelines, and data

³<https://github.com/lauren-lizzy-levine/gumbridge>

⁴Kobayashi and Ng (2020) note that when bridging resources lack official splits (like ISNotes and BASHI), evaluation is done using k-fold cross-validation, which hinders direct comparison of results.

⁵Full annotation guidelines included in supplementary materials and GitHub repository.

of prominent English bridging resources (ISNotes, BASHI, ARRAU, and GUM), aiming to unite desirable aspects of each. An initial pilot was conducted on the test set, which achieved low inter-annotator agreement on the recognition of bridging anaphora (Levine and Zeldes, 2025). Following a careful analysis of the poor agreement results, which focused on the subjectivity of annotator judgments, the scheme for bridging subtype classification was revised and the annotation guidelines were significantly extended. Following the finalization of the scheme/guidelines, the authors of this paper revised the annotations on the test set. An agreement study on the annotation of the dev set using the finalized annotation scheme and guidelines is included in Section 4.

The train set and test2 of GUMBridge were manually annotated by the authors of this paper. While the bridging annotations in dev and test were done from scratch, the train and test2 documents were corrected and expanded from the original bridging annotations in GENTLE and GUM v11.⁶ The original GUM corpus also already contains entity and coreference annotations following the same mention definitions used in this paper,⁷ as well as information status annotations, i.e., "New", "Given", and "Accessible" (not including accessibility from instances of bridging), which were provided to the annotators of GUMBridge. Additionally, common structural triggers of bridging anaphora (e.g., 2 token definite entities, relative adjectives) were highlighted based on a review of the annotated dev and test data in order to help focus annotators and improve recall. The annotation project was completed using an adapted version of the GitDox annotation interface (Zhang and Zeldes, 2017).

Identifying Bridging Anaphora GUMBridge adopts an information status informed perspective to identifying instances of bridging anaphora. As explained in Section 2, the information status of an entity refers to the level of its “Accessibility” to the participants of a discourse. An entity is considered to be “Accessible” if it is newly mentioned yet is still comprehensible to the participants of a discourse. In the case of bridging, the “Accessible” status of the entity in question is specifically due to an inference being made from a previous, non-identical antecedent. This is in contrast with entities that are

⁶The train documents added for GUM v12 were first annotated by students in a classroom setting using the GUMBridge guidelines (as part of GUM’s annual expansion procedure), and then corrected by authors of this paper.

⁷See Zeldes (2022) for discussion of coreference guidelines in GUM, and an alternative version following OntoNotes guidelines (Weischedel et al., 2012) in Zhu et al. (2021)

considered “Accessible” for other reasons, such as generics/world knowledge or situational context (e.g., unanchored deictic expressions). Bridging anaphora are not accessible by themselves, but dependent on previous entities for interpretation. Note this does not preclude named entities, accessible through world knowledge, from being bridging anaphora if they instantiate a bridging relation, though this is rare: using the existing Wikification annotations in GUM (Lin and Zeldes, 2021) we find only 98 cases of bridging with linked named entities, as in (4).

- (4) (he) remained a lifelong Roman Catholic, at a time when **the Church** was being persecuted

Annotators receive a detailed explanation of this definition of bridging and accessibility and are instructed to apply the following three criteria when determining whether an entity constitutes a bridging anaphor:

1. Is this entity Accessible (known/inferable at first mention) in the discourse?
2. Is this entity Accessible because of a relationship to a previous, non-identical entity in the discourse? If so, identify the antecedent entity.
3. Rule out the following scenarios which are **not** considered to be bridging:

coreference: If an entity has a previous mention (and thus an IS of "Given"), it cannot be a bridging anaphor.

bridging contained: The antecedent must be outside of the nominal phrase containing the anaphor, e.g., “**the focus of the story**”, should not be annotated as bridging.

generic/world knowledge: Entities that are accessible due to general world knowledge or situational context are not considered instances of bridging, e.g. “the world” or “here” should not be annotated as bridging.

explicit possessives: If the anaphor contains an explicit possessive which corefers with the antecedent, no bridging relation is necessary, e.g., “[Mark]... **[his] house**” should not be annotated as bridging.

If an entity pair passes all 3 of the criteria above, then it is an instance of bridging, and the annotator proceeds to the subcategorization task.

3.2. Bridging Subtype Classification

With the introduction of GUMBridge, we present a new subcategorization scheme for varieties of

Subtypes	
COMPARISON	
RELATIVE	1,200
TIME	436
SENSE	521
Subtotal	2,157
ENTITY	
ASSOCIATIVE	2,880
MERONOMY	460
PROPERTY	148
RESULTATIVE	102
Subtotal	3,590
SET	
MEMBER	391
SUBSET	178
SPAN-INTERVAL	169
Subtotal	738
OTHER	67
Total Subtype Count	6,552
Multi-subtype Occurrences	
1 subtype	4,887
2 subtypes	770
3 subtypes	39
4 subtypes	2
Total Instances	5,698

Table 2: Bridging subtypes in GUMBridge v1.

bridging anaphora. This system for subtype classification includes 3 main categories: COMPARISON relations, ENTITY relations, and SET relations. Within these 3 categories, there are 10 subcategories, and there is an additional OTHER category, for a total of 11 sub-varieties of bridging anaphora captured in GUMBridge. Figure 1 shows the full subcategorization scheme for GUMBridge v1.

As the criteria for identifying instances of bridging is anaphoric, there is no theoretical need to limit the number of subtypes that can apply to an instance of bridging. The subtype labels are primarily a means of understanding how the phenomenon manifests in a discourse. As such, annotators were instructed to select all subtypes that apply to a bridging pair. To our knowledge, GUMBridge is the first English bridging resource to allow for such multi-subtype annotation. Table 2 shows the counts of the bridging subtypes in the final version of GUMBridge v1. Below is a brief description and example of each bridging subtype (when multiple subtypes apply, it has been included in the footnotes):

COMPARISON-RELATIVE The anaphor is preceded by a comparative marker (other, another, same, more, etc.), ordinal (second, third, etc.), comparative adjective (larger, smaller, etc.), or superlative (best, worst, etc.) which implies a comparison to

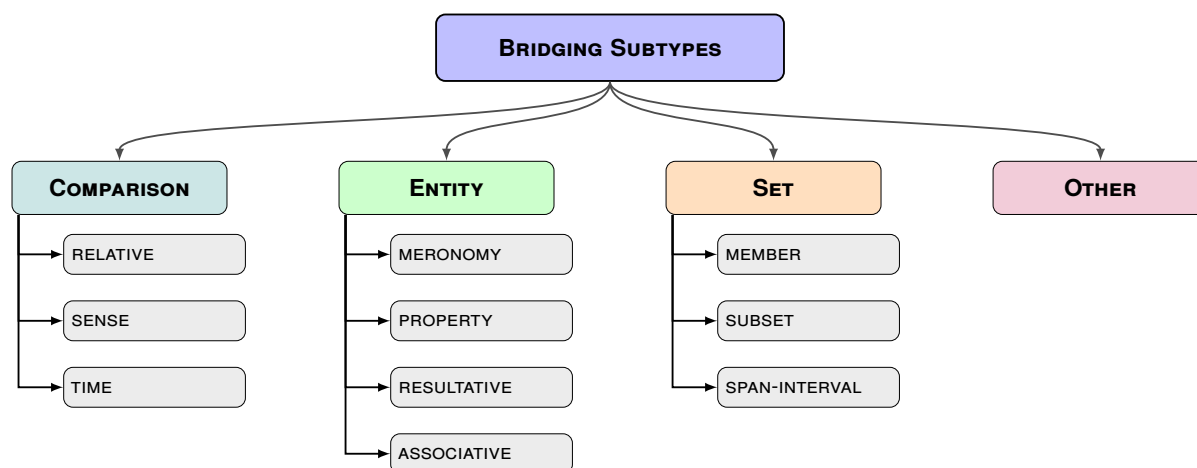


Figure 1: Bridging Subtype Classification in GUMBridge v1

the antecedent (or vice versa).

- (5) Several women walked into the room. **Other women** soon followed.

COMPARISON-SENSE The type of the anaphor is omitted but inferable via comparison to the antecedent (or vice versa).

- (6) I've been to the Chinese restaurant. I want to go to **the Italian one**.

COMPARISON-TIME The anaphor refers to a specific time/time frame which is understandable with reference to the time/time frame expressed by the antecedent (or vice versa).

- (7) I went shopping Wednesday, March 3rd. I will go again **the following Wednesday**.⁸

ENTITY-MERONOMY The anaphor is a subunit of the antecedent (or vice versa), i.e., there is some part-whole relation between the anaphor and the antecedent, including physical subparts, substance-portion, and regions/subsections.

- (8) I saw a large house by the lake. **The door** was red.

ENTITY-PROPERTY The anaphor is a physical or intangible property of the antecedent (or vice versa). For example: smell, length, style, etc.

- (9) I picked up a bouquet of roses. **The scent** was lovely.

ENTITY-RESULTATIVE The anaphor is logically inferable from the antecedent (or vice versa). This is

often the result of a transformative/product producing process, like cooking/baking.⁹

- (10) Though my flour was a strange texture, **the bread** came out perfectly.

ENTITY-ASSOCIATIVE The anaphor is an attribute or closely associated entity of the antecedent (or vice versa). This frequently manifests as implicit arguments of a predicate (11), relational nouns (12), and prototypical associations (13):

- (11) There was a murder last night. **The victim** has yet to be identified.

- (12) There is a child in the park. **The parent** must be nearby.

- (13) I went to a wedding last week. **The reception** was really fun.

SET-MEMBER The anaphor is an element of the antecedent set (or vice versa). This includes group-member and class-instance relations.

- (14) I got several books for my birthday. **The mystery novel** was my favorite.

SET-SUBSET The anaphor is a subset of the antecedent set (or vice versa).

- (15) A group of students entered the hall. **The boys** wore neckties with their uniforms.

SET-SPAN-INTERVAL The anaphor is a sub-span of the spatial or temporal antecedent interval (or vice versa).

⁸also COMPARISON-RELATIVE

⁹This covers the TRANSFORMED type proposed by Fang et al. (2022) specifically for recipe outcomes.

- (16) If you want to meet up on Sunday, I will be free in **the morning**.¹⁰

OTHER The OTHER category is for instances which fit the information status based definition of a bridging pair but do not fall into any of the bridging subtype categories outlined above. For example, [Ogrodniczuk and Zawisławska \(2016\)](#) discuss metareference, which allows for reference back to a name or label:

- (17) I went to Sensational Cakes yesterday, but I didn't think **the cakes** were very good.

While such instances are not numerous enough to warrant additional subtype categories, they are still unique and worth capturing for future investigation. When we consider the distribution of subtypes shown in Table 2, we see that only 67 of the 5,698 instances of bridging in GUMBridge are labeled with OTHER. This indicates that the subtype scheme provides good coverage for the varieties of bridging anaphora present in the data.

However, it is notable that some subtypes are far more frequent than others. For example, 51% (2,880 instances) of the bridging pairs are labeled with the ENTITY-ASSOCIATIVE subtype, while only 9% (521 instances) are labeled with COMPARISON-SENSE. While differences in the distribution of subtype categories are linguistically interesting and worth investigating, it is also worth noting that ENTITY-ASSOCIATIVE is the most broadly construed subtype, and can manifest in a variety of linguistic ways (prototypical associations, implicit possessives, implicit arguments, etc.). While we choose to group certain varieties together into one subtype category because they are difficult to reliably distinguish, we acknowledge that in the future it may be worthwhile to further delineate larger subtype categories, like ENTITY-ASSOCIATIVE.

4. Inter-Annotator Agreement Study

In order to validate the robustness of our annotation procedure and proposed schema for subtype categorization, we conducted an Inter-Annotator Agreement study on the dev set of GUMBridge. Each of the 32 documents in the dev set (~30k tokens) was double annotated by 4 different pairs of annotators, with each pair of annotators completing 8 documents. We note that the annotators in pair B are the authors of this paper, and therefore the most expert annotators in the pool. The other pairs of annotators are linguistics graduate students who completed several hours of training prior to the pilot. The results of this study are shown in Table 3. We

provide agreement numbers for each stage of annotation in our procedure: (1) recognition of bridging anaphora, (2) resolution back to the associative antecedent, and (3) subcategorization of the bridging pair. We provide the aggregate scores for the annotators together, as well as the scores of the individual pairs.

For the recognition of the bridging anaphor, we give mutual P/R/F of Annotator 1 relative to Annotator 2. The micro avg. F1 for bridging anaphor recognition is only 0.62, indicating that anaphor recognition is a difficult subtask. As the selection of a correct antecedent is limited by the recognition of the anaphor, we give the accuracy of Annotator 2 selecting the antecedent entity when both annotators agree on the bridging anaphor, which is 84% of a total of 336 cases. For the bridging subtype scores, accuracy shows the proportion of exact matches, while Cohen's κ is reported on the individual subtype judgments (e.g., a bridging pair with 2 subtype annotations is counted as 2 annotator judgments). For the 283 instances where both annotators agreed on the anaphor and antecedent of a bridging pair, the Cohen's κ for the bridging subtype annotation is 0.76, which indicates substantial agreement. From these scores, we see that the detection of bridging anaphora is the limiting factor in producing reliable annotations for instances of bridging.

Looking at the individual annotator pair scores, we see very clear divides in performance. While pairs A and B achieve high agreement on all tasks, pairs C and D perform substantially worse, particularly on the task of anaphor recognition. This inconsistency is in line with recent work which has shown annotation of bridging anaphora to be a highly subjective task, dependent on an annotator's personal judgment of how entities are related in a discourse ([Levine and Zeldes, 2025](#)). We note that our most expert annotators (pair B) were able to achieve good agreement. Amongst the 3 other pairs, who all received the same amount of training, the agreement ranged from high (pair A), to moderate (pair C), to low (pair D). From these results, we see that the various subtasks of bridging annotation are possible to learn and agree on, particularly for expert annotators. However, it is still a very difficult set of tasks, and it cannot be taken for granted that all annotators will perform reliably.

In Figure 2, we show a confusion matrix of the bridging subtype labels assigned by annotators to overlapping bridging pairs (bridging instances agreed on by both annotators in an annotator pair). We see that while there is overlap for a number of different subtypes, most confusions only occur one or twice. The subtypes with the greatest overlap are ENTITY-MERONOMY and ENTITY-ASSOCIATIVE, with a total of 12 disagreements between the two cate-

¹⁰also COMPARISON-TIME

Annotator Pairs	Anaphor Recognition			Antecedent Resolution	Bridging Subtype	
	Precision	Recall	F1 Score	Accuracy	Accuracy	Cohen's κ
A	0.84	0.81	0.83	0.94	0.90	0.86
B	0.93	0.68	0.79	0.90	0.82	0.80
C	0.66	0.36	0.47	0.70	0.50	0.58
D	0.22	0.27	0.24	0.45	0.60	0.58
Micro Avg.	0.69	0.56	0.62	0.84	0.79	0.76

Table 3: GUMBridge dev pilot inter-annotator agreement.




Figure 2: Confusion matrix of bridging subtypes for bridging instances with matching anaphor and antecedent annotations.

gories. With a Cohen's κ of 0.76 for the aggregate of annotator performance, we see the schema for bridging type classification presented here is robust, but it is clear from the scattering of disagreements that some clarification on confusable edge cases would be beneficial.

5. LLM Evaluation

In this section, we provide an LLM baseline for the task of bridging resolution evaluated on the test set of GUMBridge (32 documents, ~30k tokens, 562 instances of bridging). We provide scores for two open models `meta-llama/llama-3.3-70b-instruct` and `qwen/qwen3-235b-a22b-instruct-2507`, and one closed SoTA model, GPT-5. Following the set up of the evaluation of our human annotators, we test the models on 3 subtasks: (1) anaphora recognition, (2) antecedent selection, and (3) subtype categorization. For each subtask, the models are provided with a separate prompt which gives a detailed explanation of the task with a series of few-shot examples.¹¹ For

¹¹Templates of the exact prompts are included in Appendix A

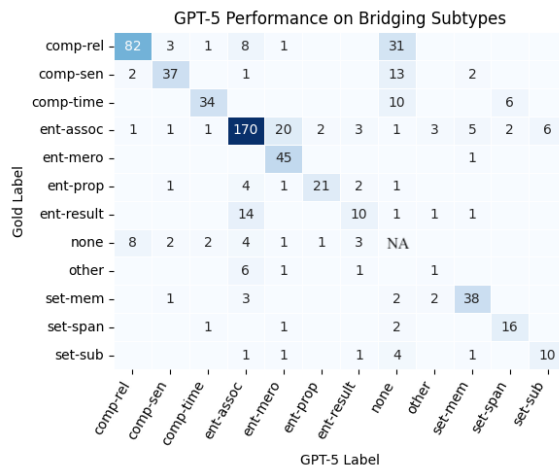


Figure 3: Confusion matrix of bridging subtypes for GPT-5 predicted labels.

the anaphor recognition subtask, the models are queried sentence by sentence through the document, with 50 tokens of additional context on each side of the sentence. For the antecedent selection subtask, the models are queried once for each bridging anaphor in the gold annotations, getting the sentence containing the anaphor with 150 tokens of preceding context. For the subtype categorization subtask, the models are queried once for each bridging pair in the gold annotations, getting the sentences containing the anaphor and the antecedent, each with 50 tokens of additional context on each side. The results of this evaluation are shown in Table 4. The reported scores are from single runs on each model.

While GPT-5 performs substantially worse than the most expertly trained annotators, we see that it scores substantially better than under-performing annotators. Comparing the results of GPT-5 to the aggregate of the human annotators, we see that the scores are lower overall. However, considering that the SoTA scores on bridging resolution (for anaphor recognition and antecedent selection together) are under 0.5 P/R/F for evaluation on ISNotes, BASHI, and ARRAU (Kobayashi et al., 2023; Kobayashi and Ng, 2020), we consider this GPT-5 baseline

Model	Anaphor Recognition			Antecedent Resolution	Bridging Subtype	
	Precision	Recall	F1 Score	Accuracy	Accuracy	Cohen's κ
GPT-5	0.43	0.38	0.40	0.49	0.69	0.64
Llama-3.3	0.15	0.30	0.20	0.28	0.61	0.27
Qwen3	0.14	0.21	0.17	0.17	0.53	0.34

Table 4: Results of LLM baselines on GUMBridge test.

to be reasonably impressive. On the other hand, `llama-3.3-70b-instruct` and `qwen3-235b-a22b-instruct-2507` perform quite poorly at the tasks of anaphor recognition (F1 0.20 and 0.17 respectively) and antecedent selection (Acc 0.28 and 0.17 respectively), and perform adequately at subtype categorization in terms of accuracy (0.61 and 0.53 respectively). This suggests that models of this size/capacity may not have the ability to recognize bridging pairs in context. While bridging resolution and subtype categorization remain difficult NLP tasks, even while leveraging LLMs, the results of the GPT-5 baseline suggest that LLMs have the capacity be a useful tool in improving the performance of bridging resolution systems.

In Figure 3 we provide a confusion matrix of the bridging subtype labels assigned by GPT-5 and the gold labels from the test set of GUMBridge. Compared with the results of the human annotators in Figure 2, we see that there are considerably more categories being confused, but most of the disagreements still only occur a handful of times. The subtypes which were most commonly confused, with 20 instances, are ENTITY-MERONOMY and ENTITY-ASSOCIATIVE, which was also the case for the human annotation. Additionally, we see that GPT-5 misses a number of instances for COMPARISON relations in cases when they co-occur with other subtype labels.




Figure 4: Performance of GPT-5 on spoken genres vs. written genres.

In Figure 4 we compare the performance of GPT-5 on spoken vs. written genres for our evaluations metrics. To our knowledge, this is the first cross-genre comparison of performance on the task of bridging resolution for English (while the

sub-corpora of ARRAU cover multiple genres, evaluation has primarily been conducted on ARRAU RST’s WSJ news data alone). Looking across the different metrics, we consistently see that spoken genres are more difficult for the model than written genres. We also see that the relative difficulty of the subtasks remains consistent between written and spoken data (anaphor recognition is the hardest, followed by antecedent selection, and subtype categorization is the easiest). While brief, we include this analysis to illustrate the important types of linguistic research which will now be possible with the additional genre diversity provided by GUMBridge.

6. Conclusion

In this paper, we introduce GUMBridge, a new English resource for varieties of bridging anaphora. Taking an information status informed perspective, we provide broad coverage, having one of the highest densities of bridging annotation in available English resources (19.6 instances per 1k tokens). We also provide subtype annotations for sub-varieties of bridging anaphora, which will allow for fine-grained linguistic analysis of how bridging instances manifest in a discourse. We note that this is the first English bridging resource to formally recognize the theoretical compatibility of multiple subtypes within a single instance and to allow for multi-subtype annotation. This dataset is the largest English resource available for instances of bridging (containing more bridging instances than the combined sub-corpora of ARRAU), particularly for referential bridging (containing more than 4 times the bridging instances of ISNotes and BASHI combined). GUMBridge also has by far the greatest range of genre diversity, covering 24 different genres, including both written and spoken data.

We also provide an LLM baseline for bridging resolution and subtype categorization evaluated on the test set of GUMBridge. Results from this evaluation indicate that bridging related tasks remain difficult even with the use of SoTA LLMs, achieving only an F1 Score of 0.40 for the task of anaphor recognition. However, given the low performance of current SoTA bridging resolution systems, the baseline also indicates the potential of LLMs to help improve bridging resolution systems in the future.

Limitations

This work is limited by the fact that GUMBridge contains only English language data, and more work will need to be done to address bridging in a multilingual capacity. While the information status based approach to identify bridging anaphora is largely language agnostic, the subtype categories of GUMBridge were developed with English data in mind and may not straightforwardly extend to other languages. Additionally, the LLM baselines included in this paper are minimal, only testing three models in a single configuration, reporting scores from a single run. This baseline is meant to illustrate that bridging resolution is still an open, challenging task, rather than to present a full system for bridging resolution.

Bibliographical References

- Nicholas Asher and Alex Lascarides. 1998. Bridging. *Journal of Semantics*, 15(1):83–113.
- Stefan Baumann and Arndt Riester. 2012. Referential and lexical givenness: Semantic, prosodic and cognitive aspects. *Prosody and meaning*, 25:119–162.
- Herbert H. Clark. 1975. Bridging. In *Theoretical Issues in Natural Language Processing*.
- Biaoyan Fang, Timothy Baldwin, and Karin Verspoor. 2022. What does it take to bake a cake? the RecipeRef corpus and anaphora resolution in procedural text. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3481–3495, Dublin, Ireland. Association for Computational Linguistics.
- John A. Hawkins. 1978. Definiteness and indefiniteness: A study in reference and grammaticality prediction. *Journal of Linguistics*, 27:405–442.
- Yufang Hou. 2020. Bridging anaphora resolution as question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1428–1438, Online. Association for Computational Linguistics.
- Sopan Khosla, Juntao Yu, Ramesh Manuvinakurike, Vincent Ng, Massimo Poesio, Michael Strube, and Carolyn Rosé. 2021. The CODI-CRAC 2021 shared task on anaphora, bridging, and discourse deixis in dialogue. In *Proceedings of the CODI-CRAC 2021 Shared Task on Anaphora, Bridging, and Discourse Deixis in Dialogue*, pages 1–15, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hideo Kobayashi, Yufang Hou, and Vincent Ng. 2023. PairSpanBERT: An enhanced language model for bridging resolution. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6931–6946, Toronto, Canada. Association for Computational Linguistics.
- Hideo Kobayashi and Vincent Ng. 2020. Bridging resolution: A survey of the state of the art. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3708–3721, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Lauren Levine and Amir Zeldes. 2024. Unifying the scope of bridging anaphora types in English: Bridging annotations in ARRAU and GUM. In *Proceedings of the Seventh Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 41–51, Miami. Association for Computational Linguistics.
- Lauren Levine and Amir Zeldes. 2025. Subjectivity in the annotation of bridging anaphora. In *Proceedings of the 19th Linguistic Annotation Workshop (LAW-XIX-2025)*, pages 48–59, Vienna, Austria. Association for Computational Linguistics.
- Jessica Lin and Amir Zeldes. 2021. WikiGUM: Exhaustive entity linking for wikification in 12 genres. In *Proceedings of the Joint 15th Linguistic Annotation Workshop (LAW) and 3rd Designing Meaning Representations (DMR) Workshop*, pages 170–175, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Malvina Nissim, Shipra Dingare, Jean Carletta, and Mark Steedman. 2004. An annotation scheme for information status in dialogue. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal. European Language Resources Association (ELRA).
- Maciej Ogrodniczuk and Magdalena Zawistawska. 2016. Bridging relations in Polish: Adaptation of existing typologies. In *Proceedings of the Workshop on Coreference Resolution Beyond OntoNotes (CORBON 2016)*, pages 16–22, San Diego, California. Association for Computational Linguistics.
- Ellen F. Prince. 1981. Toward a taxonomy of given-new information. *Radical pragmatics*, pages 223–255.
- Ina Roesiger, Arndt Riester, and Jonas Kuhn. 2018. Bridging resolution: Task definition, corpus resources and rule-based experiments. In *Proceedings of the 27th International Conference*

- on *Computational Linguistics*, pages 3516–3528, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Ina Rösiger. 2018. [BASHI: A corpus of Wall Street Journal articles annotated with bridging links](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Juntao Yu, Sopan Khosla, Ramesh Manuvinakurike, Lori Levin, Vincent Ng, Massimo Poesio, Michael Strube, and Carolyn Rosé. 2022. [The CODI-CRAC 2022 shared task on anaphora, bridging, and discourse deixis in dialogue](#). In *Proceedings of the CODI-CRAC 2022 Shared Task on Anaphora, Bridging, and Discourse Deixis in Dialogue*, pages 1–14, Gyeongju, Republic of Korea. Association for Computational Linguistics.
- Amir Zeldes. 2022. [Can we fix the scope for coreference? problems and solutions for benchmarks beyond OntoNotes](#). *Dialogue & Discourse*, 13(1):41–62.
- Amir Zeldes. 2024. [Entity and Information Status Annotation](#). GUM Annotation Wiki.
- Shuo Zhang and Amir Zeldes. 2017. [GitDOX: A linked version controlled online XML editor for manuscript transcription](#). In *Proceedings of the Thirtieth International Florida Artificial Intelligence Research Society Conference (FLAIRS 2017)*, pages 619–623.
- Yilun Zhu, Sameer Pradhan, and Amir Zeldes. 2021. [OntoGUM: Evaluating contextualized SOTA coreference resolution on 12 more genres](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 461–467, Online. Association for Computational Linguistics.
- Eckart, Kerstin and Riestler, Arndt and Schweitzer, Katrin. 2012. [A Discourse Information Radio News Database for Linguistic Analysis](#). Springer Berlin Heidelberg.
- Markert, Katja and Hou, Yufang and Strube, Michael. 2012. [Collective Classification for Fine-grained Information Status](#). Association for Computational Linguistics.
- Nedoluzhko, Anna and Mírovský, Jiří and Pajas, Petr. 2009. [The Coding Scheme for Annotating Extended Nominal Coreference and Bridging Anaphora in the Prague Dependency Treebank](#). Association for Computational Linguistics, ISLRN 942-053-729-014-3.
- Ogrodniczuk, Maciej and Zawisławska, Magdalena. 2016. [Bridging Relations in Polish: Adaptation of Existing Typologies](#). Association for Computational Linguistics.
- Poesio, Massimo and Artstein, Ron. 2008. [Anaphoric Annotation in the ARRAU Corpus](#). European Language Resources Association (ELRA), ISLRN 462-157-606-044-8.
- Rösiger, Ina. 2018. [BASHI: A Corpus of Wall Street Journal Articles Annotated with Bridging Links](#). European Language Resources Association (ELRA).
- Schweitzer, Katrin and Eckart, Kerstin and Gärtner, Markus and Falenska, Agnieszka and Riestler, Arndt and Rösiger, Ina and Schweitzer, Antje and Stehwien, Sabrina and Kuhn, Jonas. 2018. [German Radio Interviews: The GRAIN Release of the SFB732 Silver Standard Collection](#). European Language Resources Association (ELRA).
- Olga Uryupina and Ron Artstein and Antonella Bristot and Federica Cavicchio and Francesca Delogu and Kepa Joseba Rodríguez and Massimo Poesio. 2019. [Annotating a broad range of anaphoric phenomena, in a variety of genres: the ARRAU Corpus](#).
- Ralph Weischedel and Sameer Pradhan and Lance Ramshaw and Jeff Kaufman and Michelle Franchini and Mohammed El-Bachouti and Nianwen Xue and Martha Palmer and Jena D. Hwang and Claire Bonial and Jinho Choi and Aous Mansouri and Maha Foster and Abdel-aati Hawwary and Mitchell Marcus and Ann Taylor and Craig Greenberg and Eduard Hovy and Robert Belvin and Ann Houston. 2012. [OntoNotes Release 5.0](#).
- Amir Zeldes. 2017. [The GUM Corpus: Creating Multilayer Resources in the Classroom](#). ISLRN 421-566-418-865-2.

Language Resource References

- Aoyama, Tatsuya and Behzad, Shabnam and Gessler, Luke and Levine, Lauren and Lin, Jessica and Liu, Yang Janet and Peng, Siyao and Zhu, Yilun and Zeldes, Amir. 2023. [GENTLE: A Genre-Diverse Multilayer Challenge Set for English NLP and Linguistic Evaluation](#). Association for Computational Linguistics.

A. LLM Baseline Prompts

A.1. Anaphor Recognition

You are a linguistic analyst whose job is to find cases of bridging anaphora: mentions of newly introduced entities (noun phrases) in a text, for which a reader would need to refer back to a previously mentioned, non-identical entity to resolve their meaning. There are several classes of bridging anaphors, any of which should be identified in the text being analyzed. In the following examples, the bridging anaphor is surrounded by *asterisks*.

comparison-relative: The anaphor is preceded by a comparative marker (other, another, same, more, ordinal modifiers, comparative adjectives, superlatives, etc.) which implies a comparison to the antecedent. For example: "The children... **another child**" (=another with comparison to the aforementioned children); similar cases may be **similar children**, **older children** (compared to the aforementioned children), etc.

comparison-sense: the semantic type of a phrase requires a previous mention to identify it, for example "the Italian restaurant... **a Chinese one**" (we can't know "a Chinese one" is a restaurant without referring back to the Italian restaurant), or **"another one"*, **"the others"* etc.

comparison-time: the anaphor refers to a specific time/timeframe which is understandable with reference to the antecedent, for example: "Tuesday, February 2nd ... **the following week**"

entity-meronymy: the anaphor is a subunit of the antecedent (part-whole), including physical subunits, portion-substance relations, and regions/subsections. For example: "the house ... **the door**" (=of the house).

entity-associative: the anaphor is an attribute or closely associated entity of the antecedent, including both prototypical and inducible associations: "a wedding ... **the bride**" (=the bride at that wedding), implicit arguments of a predicate or a verbal nominalization: "a play... **the performance**" (=of the play), relational nouns: "a murder ... **the victim**"

entity-property: the anaphor is a physical or intangible property of the antecedent (e.g., smell, length, size, style, etc.): "the tea... **the sweet aroma**"

entity-resultative: the anaphor is logically inferable from the antecedent (e.g., result, transformation/transmutation, cause): "the dough ... **the bread**" (=the dough becomes bread after baking)

set-member: the anaphor is an element of the antecedent set, including groups-member relations and classes-instances: "the cars ... **the Mazda**", additionally indefinite members to definite sets: "a candle on each cupcake... **the candles**"

set-subset: the anaphor is a subset of the antecedent set: "the cars ... **the Mazdas**" (not all Mazdas, just the subset among the aforementioned cars)

set-span-interval: the anaphor is a sub-span of a spatial or temporal interval defined by the antecedent: "last week... **Wednesday**" (=Wednesday of last week), "Sunday... **the morning**" (=the morning portion of that Sunday)"

other: the anaphor requires a previous entity for interpretation, but it doesn't fit into any of the above categories. This is a rare class.

There are also some exceptions which should NOT be identified as bridging anaphora:

Coreference: If an entity has a previous mention, it cannot be an instance of bridging. For instance, in "Catherine and Henry had their wedding last week. The bride was very beautiful", even though there is an associative relationship between the wedding and the bride, since "the bride" corefers with "Catherine", which has already been introduced to the discourse, "the bride" is not eligible to be an instance of bridging.

Bridging-contained: If the entity one would need to refer back to in order to understand the bridging anaphor is a direct modifier in the noun phrase of the potential bridging anaphor, e.g. "the focus of the story" or "two of them", it should not be annotated as bridging. In other words, the previous antecedent entity must be outside of the nominal phrase containing the anaphor. An entity that is followed by a prepositional phrase or a relative clause is sufficiently qualified and is thus NOT an instance of bridging.

Generics/Situational bridging: Entities that are accessible due to general world knowledge or situational context are not considered instances of bridging, i.e., if it doesn't have a previous associated antecedent entity to be bridging from, it cannot be bridging.

Possession with an explicit possessive: If the potential bridging anaphor contains an explicit possessive which corefers with the associative antecedent entity, no bridging relation is necessary. Explicit coreference between the associative antecedent and the possessive is sufficient (e.g., [Mark]. . . [his] house → no bridging, coreference between "Mark" and "his"). Contrast this with [the family] . . . **the house** → bridging, since we cannot interpret which house it is (the house of the family) without referring to "the family", which is outside of the anaphor phrase.

Here are 2 an examples of the task:

Please return a list all of the bridging anaphors in the following text in the order in which they appear.

Output the anaphor mention phrase exactly as it appears in the text. If there are no bridging anaphors, return an empty list.

Text:

... with their friends to a picnic. The picnic was supposed to take place in a grove, but the shade wasn't enough, so they had to find a different place. Conny started to say ...

Answer(s):

["the shade", "a different place"]

Please return a list all of the bridging anaphors in the following text in the order in which they appear. Output the anaphor mention phrase exactly as it appears in the text. If there are no bridging anaphors, return an empty list.

Text:

... making this technique the basis of training for all types of dance . While dancing ballet takes dedication and requires serious training , you can learn the basics to prepare yourself for further study . Learn to get ready for practicing...

Answer(s):

["the basics", "further study"]

Please return a list all of the bridging anaphors in the following text in the order in which they appear. Output the anaphor mention phrase exactly as it appears in the text. If there are no bridging anaphors, return an empty list.

Text:

{text}

Answer(s):

A.2. Antecedent Selection

You are a linguistic analyst whose job is to select the associative antecedent for of a bridging anaphor: mentions of newly introduced entities (noun phrases) in a text, for which a reader would need to refer back to a previously mentioned, non-identical entity (the antecedent) to resolve their meaning. There are several classes of bridging instances, defined by the associative relationship between the bridging anaphor and its associative antecedent. In the following examples, the bridging antecedent is surrounded by *asterisks* and the bridging anaphor is surrounded by {{double curly brackets}}.

comparison-relative: The anaphor is preceded by a comparative marker (other, another, same, more, ordinal modifiers, comparative adjectives, superlatives, etc.) which implies a comparison to the antecedent. For example: "*The children* ... {{another child}}" (=another with comparison to the aforementioned children); similar cases may be

{{similar children}}, {{older children}}(compared to the aforementioned children), etc.

comparison-sense: the semantic type of a phrase requires a previous mention to identify it, for example "*the Italian restaurant* ... {{a Chinese one}}" (we can't know "a Chinese one" is a restaurant without referring back to the Italian restaurant), or "{{another one}}", "{{the others}}" etc.

comparison-time: the anaphor refers to a specific time/timeframe which is understandable with reference to the antecedent, for example: "*Tuesday, February 2nd* ... {{the following week}}"

entity-meronymy: the anaphor is a subunit of the antecedent (part-whole), including physical subunits, portion-substance relations, and regions/subsections. For example: "*the house* ... {{the door}}" (=of the house).

entity-associative: the anaphor is an attribute or closely associated entity of the antecedent, including both prototypical and inducible associations: "*a wedding* ... {{the bride}}" (=the bride at that wedding), implicit arguments of a predicate or a verbal nominalization: "*a play* ... {{the performance}}" (=of the play), relational nouns: "*a murder* ... {{the victim}}"

entity-property: the anaphor is a physical or intangible property of the antecedent (e.g., smell, length, size, style, etc.): "*the tea* ... {{the sweet aroma}}"

entity-resultative: the anaphor is logically inferable from the antecedent (e.g., result, transformation/transmutation, cause): "*the dough* ... {{the bread}}" (=the dough becomes bread after baking)

set-member: the anaphor is an element of the antecedent set, including groups-member relations and classes-instances: "*the cars* ... {{the Mazda}}", additionally indefinite members to definite sets: "*a candle* on each cupcake ... {{the candles}}"

set-subset: the anaphor is a subset of the antecedent set: "*the cars* ... {{the Mazdas}}" (not all Mazdas, just the subset among the aforementioned cars)

set-span-interval: the anaphor is a sub-span of a spatial or temporal interval defined by the antecedent: "*last week* ... {{Wednesday}}" (=Wednesday of last week), "*Sunday* ... {{the morning}}" (=the morning portion of that Sunday)

other: the anaphor requires a previous entity for interpretation, but it doesn't fit into any of the above categories. This is a rare class.

Here are 2 an examples of the task:

Please return a single string for associative antecedent of the bridging anaphor surrounded by {{double curly brackets}}. Output the antecedent mention phrase exactly as it appears in the text.

If there is no associative antecedent, return "no antecedent". The antecedent you are returning CANNOT be the same as the bracketed anaphor.

Text:

... with their friends to a picnic. The picnic was supposed to take place in a grove, but `{{the shade}}` wasn't enough, so they had to find a different place. Conny started to say ...

Answer:

a grove

Please return a single string for associative antecedent of the bridging anaphor surrounded by `{{double curly brackets}}`. Output the antecedent mention phrase exactly as it appears in the text. If there is no associative antecedent, return "no antecedent".

Text:

... making this technique the basis of training for all types of dance . While dancing ballet takes dedication and requires serious training , you can learn the basics to prepare yourself for `{{further study}}` . Learn to get ready for practicing ...

Answer:

ballet

Please return a single string for associative antecedent of the bridging anaphor surrounded by `{{double curly brackets}}`. Output the antecedent mention phrase exactly as it appears in the text. If there is no associative antecedent, return "no antecedent".

Text:

{text}

Answer:

A.3. Subtype Categorization

You are a linguistic analyst whose job is to select the subtype classification for a bridging anaphor - antecedent pair: mentions of newly introduced entities (the anaphor) in a text, for which a reader would need to refer back to a previously mentioned, non-identical entity (the antecedent) to resolve their meaning. There are several classes of bridging instances, defined by the associative relationship between the bridging anaphor and its associative antecedent. In the following subtype examples, the bridging antecedent is surrounded by *asterisks* and the bridging anaphor is surrounded by `{{double curly brackets}}`.

comparison-relative: The anaphor is preceded by a comparative marker (other, another, same, more, ordinal modifiers, comparative adjectives, superlatives, etc.) which implies a comparison to the antecedent. For example: `"*The children* ... {{another child}}"` (=another with comparison to the

mentioned children); similar cases may be `{{similar children}}`, `{{older children}}` (compared to the mentioned children), etc.

comparison-sense: the semantic type of a phrase requires a previous mention to identify it, for example `"*the Italian restaurant* ... {{a Chinese one}}"` (we can't know "a Chinese one" is a restaurant without referring back to the Italian restaurant), or `"{{another one}}"`, `"{{the others}}"` etc.

comparison-time: the anaphor refers to a specific time/timeframe which is understandable with reference to the antecedent, for example: `"*Tuesday, February 2nd* ... {{the following week}}"`

entity-meronymy: the anaphor is a subunit of the antecedent (part-whole), including physical subunits, portion-substance relations, and regions/subsections. For example: `"*the house* ... {{the door}}"` (=of the house).

entity-associative: the anaphor is an attribute or closely associated entity of the antecedent, including both prototypical and inducible associations: `"*a wedding* ... {{the bride}}"` (=the bride at that wedding), implicit arguments of a predicate or a verbal nominalization: `"*a play*... {{the performance}}"` (=of the play), relational nouns: `"*a murder* ... {{the victim}}"`

entity-property: the anaphor is a physical or intangible property of the antecedent (e.g., smell, length, size, style, etc.): `"*the tea* ... {{the sweet aroma}}"`

entity-resultative: the anaphor is logically inferable from the antecedent (e.g., result, transformation/transmutation, cause): `"*the dough* ... {{the bread}}"` (=the dough becomes bread after baking)

set-member: the anaphor is an element of the antecedent set, including groups-member relations and classes-instances: `"*the cars* ... {{the Mazda}}"`, additionally indefinite members to definite sets: `"*a candle* on each cupcake ... {{the candles}}"`

set-subset: the anaphor is a subset of the antecedent set: `"*the cars* ... {{the Mazdas}}"` (not all Mazdas, just the subset among the mentioned cars)

set-span-interval: the anaphor is a sub-span of a spatial or temporal interval defined by the antecedent: `"*last week* ... {{Wednesday}}"` (=Wednesday of last week), `"*Sunday* ... {{the morning}}"` (=the morning portion of that Sunday)

other: the anaphor requires a previous entity for interpretation, but it doesn't fit into any of the above categories. This is a rare class.

Here are 2 an examples of the task:

In the following text, a bridging anaphora is marked with `{{double curly brackets}}` and the corresponding antecedent is surrounded by *asterisks*. Read

the following text and for the bridging anaphor-antecedent pair, classify the variety of bridging subtype relation (defined above) that holds between the two entities. Multiple subtypes may apply to a single pair. Output a string of all applicable subtypes, connected by semicolons (no spaces).

The possible subtype labels are as follows:

comparison-relative
comparison-sense
comparison-time
entity-associative
entity-meronymy
entity-property
entity-resultative
set-member
set-subset
set-span-interval
other

Antecedent Text:

... with their friends to a picnic. The picnic was supposed to take place in *a grove*, but the shade wasn't enough, so they had to find a different place. Conny started to say ...

Anaphor Text:

... to a picnic. The picnic was supposed to take place in a grove, but {{the shade}} wasn't enough, so they had to find a different place. Conny started to say ...

Answer:

entity-associative

In the following text, a bridging anaphora is marked with {{double curly brackets}} and the corresponding antecedent is surrounded by *asterisks*. Read the following text and for the bridging anaphor-antecedent pair, classify the variety of bridging subtype relation (defined above) that holds between the two entities. Multiple subtypes may apply to a single pair. Output a string of all applicable subtypes, connected by semicolons (no spaces).

The possible subtype labels are as follows:

comparison-relative
comparison-sense
comparison-time
entity-associative
entity-meronymy
entity-property
entity-resultative
set-member
set-subset

set-span-interval

other

Antecedent Text:

... this technique the basis of training for all types of dance . While dancing *ballet* takes dedication and requires serious training , you can learn the basics to prepare yourself for further study . Learn to get ready for practicing ...

Anaphor Text:

... making this technique the basis of training for all types of dance . While dancing ballet takes dedication and requires serious training , you can learn the basics to prepare yourself for {{further study}} . Learn to get ready for ...

Answer:

comparison-relative

In the following text, a bridging anaphora is marked with {{double curly brackets}} and the corresponding antecedent is surrounded by *asterisks*. Read the following text and for the bridging anaphor-antecedent pair, classify the variety of bridging subtype relation (defined above) that holds between the two entities. Multiple subtypes may apply to a single pair. Output a string of all applicable subtypes, connected by semicolons (no spaces).

The possible subtype labels are as follows:

comparison-relative
comparison-sense
comparison-time
entity-associative
entity-meronymy
entity-property
entity-resultative
set-member
set-subset
set-span-interval
other

Antecedent Text:

... {antecedent_text} ...

Anaphor Text:

... {anaphor_text} ...

Answer: