

Datasets for Verb Alternations across Languages: BLM Templates and Data Augmentation Strategies

Giuseppe Samo¹, Paola Merlo^{1,2}

¹IDIAP Research Institute, ²University of Geneva
{giuseppe.samo, paola.merlo}@idiap.ch

Abstract

Large language models (LLMs) have shown remarkable performance across various sentence-based linguistic phenomena, yet their ability to capture cross-sentence paradigmatic patterns, such as verb alternations, remains underexplored. In this work, we present curated paradigm-based datasets for four languages, designed to probe systematic cross-sentence knowledge of verb alternations (change-of-state and object-drop constructions in English, German and Italian, and Hebrew binyanim). The datasets comprise thousands of the Blackbird Language Matrices (BLMs) problems. The BLM task – an RPM/ARC-like task devised specifically for language – is a controlled linguistic puzzle where models must select the sentence that completes a pattern according to syntactic and semantic rules. We introduce three types of templates varying in complexity and apply linguistically-informed data augmentation strategies across synthetic and natural data. We provide simple baseline performance results across English, Italian, German, and Hebrew, that demonstrate the diagnostic usefulness of the datasets.

Keywords: Verb Alternation, Synthetic curated datasets, Corpus Data, LLMs, Syntax/Semantics

1. Introduction

Large language models (LLMs) have demonstrated successful performance in a vast range of practical language-based applications and the nature and extent of their linguistic competence have been investigated across a wide range of linguistic phenomena, with impressive results in areas such as negation, subject-verb agreement or other dependency-related phenomena (Warstadt et al., 2019; Ravichander et al., 2021; Belinkov, 2022; Acs et al., 2024; Linzen and Baroni, 2021; Morante and Blanco, 2021; Gautam et al., 2024).

A further dimension of complexity arises with paradigmatic phenomena such as verb alternations or inflectional patterns, since uncovering the underlying rules requires considering relations across multiple sentences rather than within a single sentence. These cross-sentence dependencies reveal broader systematic structures that go beyond local grammatical constraints (Bobaljik, 2015; Klein and Tsarfaty, 2020). In this work, we focus on causative alternations, a phenomenon where certain verbs can appear in different forms and structures, one of them conveying causative meaning. Consider the English verbs *paint* and *break*, which can be transitive or intransitive. In the transitive use, an agent acts on a patient, causing a change of state (e.g., *The artist painted/broke the vase*). In the intransitive use, *paint* retains its agentive subject (*The artist painted*), but *break* shows the patient surfacing as the subject of the clause (*The vase broke*). This alternation in *break*, a change-of-state verb (COS), involves an argument shift, that does not occur for verbs like *paint* (Object Drop class,

OD) (Levin, 1993; Merlo and Stevenson, 2001).

Languages differ considerably in how they encode the causative alternation morphologically (Haspelmath et al., 2014; Samardžić and Merlo, 2018). Figure 1 illustrates this variation in causative and object drop alternations across English, Italian and German. In change-of-state verbs, English exhibits no overt marking, Italian employs a reflexive-like marker (*si*) in the intransitive form (Jezek, 2024), and German signals the shift through case morphology, where the patient surfaces with nominative rather than accusative case, a reflex of its syntactic subject position. Modern Hebrew, and other Semitic languages, show the complex non-concatenative morphology typical of this family, where verb formation relies on consonantal roots combined with vocalic templates (McCarthy, 1979; Arad, 2005; Tsarfaty, 2004; Kastner, 2019).

In this paper, we present curated datasets across languages for a particular task, the Blackbird Language Matrices (BLMs) (Merlo, 2023b), inspired by the visual analytical reasoning test of Raven Progressive Matrices (RPM, Raven 1938) and analogous to the visual ARC corpus (Chollet, 2019). This task matches well the underlying linguistic rules of alternations. BLMs are linguistic puzzles that implicitly describe a paradigmatic linguistic system. The BLM task requires selecting the sentence that satisfies an underlying linguistic rule in an incomplete template. Specifically, the task consists in a context set of sentences that implicitly provides all necessary information for continuation of a given linguistic pattern, and an answer set of minimally differing contrastive sentences, where only one – the missing element in the pattern described by the

Language	Structure	Change of state verbs	Object Drop verbs
English	Transitive Intransitive	The artist broke the vase The vase broke	The artist paints the vase The artist paints
Italian	Transitive Intransitive	L' artista ruppe il vaso Il vaso si ruppe	L' artista dipingeva il vaso L' artista dipingeva
German	Transitive Intransitive	Der Künstler zerbrach den Becher Der Becher zerbrach	Der Künstler malte den Becher Der Künstler malte

Figure 1: Example of Change of state verbs and Object drop verbs alternation and morphological marking across languages. The verb is highlighted in bold while colours code semantic roles (agents, patients), morphological marking on the verb (si in Italian) and morphological case marking on the argument nominative and accusative case.

context – is correct. Such template is then instantiated with curated datasets.

In this work, we explore three typologies of templates and three types of methodology for data creation and augmentation. We also provide baseline experiments to demonstrate the diagnostic usefulness of the data and offer a point of comparison of model abilities. These experiments are not intended as a comprehensive evaluation.

2. Related work

Verb alternations have long been a central topic in linguistic theory (Levin, 1993; Rappaport Hovav and Levin, 2024) and have recently attracted increasing attention in computational studies. LLMs have demonstrated strong performance on a range of alternation phenomena (Kann et al., 2019; Warstadt et al., 2019; Wilson et al., 2023). For instance, Yi et al. (2022) show that LLMs with contextual embeddings capture information about verb alternation classes at both the word and sentence levels in English. Similarly, the semantic properties of verb arguments, such as agents and patients, have been analyzed using transformer-based models (Proietti et al., 2022). Beyond English, recent studies have examined verb subcategorization in Italian, assessing how well models identify and classify syntactic frames (Simonetti et al., 2024; Grasso et al., 2024).

Crafting paradigm-based datasets is crucial to evaluate models and testing their ability to generalize across paradigms, for example, morphological patterns (Batsuren et al., 2022; Nicolai et al., 2024). Evaluating the linguistic competence of pretrained LLMs typically relies on benchmark suites combining synthetic and natural datasets. Synthetic datasets are designed to probe specific grammatical phenomena in a controlled way, often using minimal pairs or constructed paradigms (Warstadt et al., 2019, 2020). Recent work has explored the use of LLMs for data generation, producing synthetic corpora or targeted examples and assess their distance from natural corpus data (Schepens et al., 2025; Civico, 2025). Although automatic gen-

eration enables large-scale testing, it also raises concerns about distributional biases (Zhang and Pavlick, 2025; Nadas et al., 2025; Griffiths et al., 2024).

In contrast, natural datasets such as the Universal Dependencies treebanks (UD) provide large-scale, annotated corpora derived from authentic textual sources (De Marneffe et al., 2021). These treebanks have recently been used to study and create large-scale data targeting complex linguistic patterns (Jumelet et al., 2025). UD data have also been used to evaluate LLMs performance on tasks such as morphological analysis and syntactic annotation (Akkurt et al., 2024; Acs et al., 2024).

3. The BLM templates

CONTEXT	
1	The artist broke the vase two years ago
2	The artist broke the vase
3	The vase broke two years ago
4	???
ANSWER SET	
1	The vase broke the artist
2	The vase broke
3	The artist broke
4	The vase broke in the museum

Figure 2: Simplified example of a BLM template.

Blackbird Language Matrices (BLMs) are linguistic puzzles designed to implicitly capture linguistic paradigms (Merlo, 2023b,a; An et al., 2023; Samo et al., 2023; Nastase et al., 2024a,b; Jiang et al., 2024; Jiang and Merlo, 2026). In a BLM task, given an informative but incomplete context, the goal is to select the sentence that satisfies an underlying linguistic rule and completes the context. The sentence is selected from a contrastive answer set. Contexts are designed based on templates and are then instantiated using natural language instances. A simplified BLM is shown in Figure 2. BLMs create a controlled environment to study how internal representations encode knowledge of linguistic paradigms, such as verb alternation. To

solve the puzzle, a model must observe all variants and represent the templatic structure underlying the system (Bobaljik, 2015).

Templates Let us first consider the context. In Figure 2, verb alternants are relatively easy to locate in a BLM template: one (or more) of the alternants is to be given in the context, and the correct answer shows the missing alternant. The composition of the sentences in the context should highlight the underlying alternation rules. Formally, the relations operations R transform a list of sentences into a predictable sequence (Merlo, 2023b).

These operators can be broadly divided into three categories: operators peripheral to the phenomenon under investigation (type A), operators showing relations among components of the phenomenon through linguistic diagnostics (type B), and a combination of the two (type C).

The BLM template in Figure 2 can be seen as a case of the first type (type A). The core components of the verb alternation are the verb and its arguments. The temporal constituent *two years ago* is external to the linguistic phenomenon, but it helps create a pattern to be solved. In Figure 2, the extra-phenomenon rule is simply the presence or absence of a temporal prepositional phrase (*two years ago*) between the two alternants. The correct answer is therefore #2 (in bold).

A second type of template (type B) is designed so that models can implicitly learn internal rules by assigning properties to the relevant arguments using linguistic diagnostics drawn from grammatical descriptions (Samo et al., 2023; Jiang and Merlo, 2026). This template allows us to test whether the model encodes the abstract linguistic properties of the phenomenon. A variation of Figure 2’s context with this type of template is shown in Figure 3. The figure shows that the two arguments of the verb are plausible subjects in a subject relative clause with an active verb, therefore both constituents can be subjects in a standard canonical clause.

CONTEXT
1 The artist broke the vase
2 The artist that broke the vase is there
3 The vase that broke is there
4 ???

Figure 3: Example of template of type B

The third type of template is a combination of the two. In type C (see example in Figure 4), the context provides different kinds of cues, including linguistic diagnostics (such as relative clauses) and factors that are external to alternation (presence of prepositional phrases). The correct answer *The vase broke* appears within a more complex paradigm, since the other verbs in the context dis-

play different inflectional forms (in this case, tense) and therefore require the model to rely on different relations operations.

CONTEXT
1 The artist will break the vase on stage
2 The vase will break during the performance
3 We met the artist that broke the vase on stage
4 The vase that broke on stage now values millions
5 The artist breaks the vase
6 ???

Figure 4: Example of template of type C

The organization of a given template determines also the answer set. Answer sets contain two types of errors, that apply to all the templates described here. **SEQUENCE** errors, where the candidate answer does not follow the expected order of constituents in the target answer (e.g. answers 1 and 4 in Figure 2). **GRAMMAR** errors, where the quality or the constituents (their semantic or syntactic nature) of the sequence does not match the target answer (e.g. answer 3 in Figure 2).

Data Augmentation Templates allow us to define a specific linguistic pattern and maintain full control over its structural properties. However, to ensure that the model’s results are reliable and not overly sensitive to minimal lexical instantiation, a large amount of instantiated templates is required. The instantiation process typically begins with one or two carefully crafted examples. The examples are then augmented in a linguistically-informed way. In section 4, we illustrate three strategies: the use of masked modelling for controlled data generation, prompt-based augmentation, and the use of natural data extracted from large-scale corpora. For each BLM matrix, two different levels of instantiations are output: minimal lexical variation across constituents (henceforth, **MINLEX**), where the template is built with the same lexical material (see Figure 2 and 3), and maximal lexical variation (**MAXLEX**), in which sentences do not share the same lexical constituents, but they still encode the same syntactic and semantic structure, as in Figure 5.

In what follows, we present the datasets on verb alternations across languages.

4. Data

For each dataset, we introduce the linguistic phenomenon, the template, and the data augmentation procedure. Table 1 provides an overview of these datasets¹.

¹The datasets are available at <https://www.idiap.ch/en/scientific-research/data>, where they can be found by searching with the string ‘BLM’.

CONTEXT	
1	The queen chipped the scepter this morning
2	The artist broke the vase
3	The door opened at midnight
4	???
ANSWER SET	
1	The inspector closed the case
2	The situation improved
3	The artist broke
4	My roommate fried last night

Figure 5: A simplified example of MAXLEX of the template of Figure 2. The sentences follow the same syntax, but the words vary in each sentence.

Name	Language	Template	Augmentation Strategies
OD	En, It	A	Masking
COS	En, It	A	Masking
COS+	De, It, En	B	Prompting
Caus	He	C	Natural data

Table 1: Overview of the datasets.

4.1. BLM-COS and BLM-OD

The verbs adopted as examples in Figure 1 belong to two distinct, well-studied verb classes (Levin, 1993): Change-of-State (COS) verbs (e.g., *break*), and Object-Drop (OD) verbs (e.g., *paint*) (Levin, 1993; Merlo and Stevenson, 2001). As mentioned in the introduction, these two semantic classes of verbs differ in their idiosyncratic lexical meaning, but only minimally in their argument syntactic structure properties. We restrict our investigation to English and Italian, two languages that offer useful contrasts for examining these verb classes, as Italian overtly marks the intransitive form of COS verbs but not that of OD verbs (see Figure 1).

Template The progression relies on superficial cues. Passive constructions are adopted as forms of intermediate stages for intransitive verbs, where the patient is the subject of the clause – but also an initial constituent – but the verb is inflected with a passive voice. For COS verbs, this progression is natural, while for OD verbs, the passive acts as a confounding factor in the context (cf. Meo et al. 2007). The semantic role of the subject in the final contextual sentence is the key differentiator between the two verb classes. The context template also includes an operator peripheral to the phenomenon under investigation, the alternation between a p-NP and a by-NP phrase preceded by a by-phrase (*by* in English, *da* in Italian).

The answer set is constructed to reflect both the linguistic rules under study and the extra-linguistic operator. All sentences in the answer set follow a common XP–V–XP structure to favour compari-

son. We create eight candidates that correspond to selected combinations of four factors: agents (Ag) and patients (Pat), active (Akt) and passive (Pass), the presence or absence of a *by* preposition and the use of arguments as elements following *by* so that the correct answer requires the model to recognize the appropriate intransitive form and templatic continuation, while incorrect answers systematically deviate in terms of the semantic roles of the constituents and the voice of the verb. Errors are categorized based on their nature: if the alternation with a *by*-phrase is correct but the semantics is wrong, it is classified as a GRAMMAR error; if an NP appears in a transitive or passive sentence where it should not, it is classified as a SEQUENCE error. Among SEQUENCE errors, we emphasize that standard canonical transitive and passive clauses can appear as grammatical answers, thereby misleading the model. Figure 6 shows the template for English.

For Italian, we created two additional answers due to the use of the reflexive-like element *si*, characteristic of causative constructions (Vietri, 2019), e.g. *il burro si scioglie* ‘the butter SI melts’. Its absence in COS sentences renders them ungrammatical (e.g. **il burro scioglie*), while its presence in OD contexts produces a grammatical but transitive clause with a reflexive (*l’artista si dipinge*, lit. ‘the artist paints herself’).

COS CONTEXT					COS ANSWERS		
1	Ag	Akt	Pat	p-NP	1	Pat Akt by-NP	CORRECT
2	Ag	Akt	Pat	by-NP	2	Ag Akt by-NP	GRAMMAR
3	Pat	Pass	by-Ag	p-NP	3	Pat Pass by-Ag	SEQUENCE
4	Pat	Pass	by-Ag	by-NP	4	Ag Pass by-Pat	SEQUENCE
5	Pat	Pass		p-NP	5	Pat Akt Ag	SEQUENCE
6	Pat	Pass		by-NP	6	Ag Akt Pat	SEQUENCE
7	Pat	Akt		p-NP	7	Pat Akt by-Ag	GRAMMAR
8	???				8	Ag Akt by-Pat	GRAMMAR

OD CONTEXT					OD ANSWERS		
1	Ag	Akt	Pat	p-NP	1	Pat Akt by-NP	GRAMMAR
2	Ag	Akt	Pat	by-NP	2	Ag Akt by-NP	CORRECT
3	Pat	Pass	by-Ag	p-NP	3	Pat Pass by-Ag	SEQUENCE
4	Pat	Pass	by-Ag	by-NP	4	Ag Pass by-Pat	SEQUENCE
5	Pat	Pass		p-NP	5	Pat Akt Ag	SEQUENCE
6	Pat	Pass		by-NP	6	Ag Akt Pat	SEQUENCE
7	Ag	Akt		p-NP	7	Pat Akt by-Ag	GRAMMAR
8	???				8	Ag Akt by-Pat	GRAMMAR

Figure 6: BLM context sets (left) for the change-of-state group (COS) and the object drop (OD) class and answers sets (right). Italian COS and OD has two additional answers in which the presence/absence of the reflexive-like element *si* is absent (both GRAMMAR errors).

Data instantiation We selected 30 verbs in English from the classification in Levin (1993, 240-247) and their equivalent, in terms of meaning and syntactic behaviour, in Italian. The semi-automatic generation of the synthetic lexical data (arguments and prepositional phrases) began with the English

dataset: lexical items were manually selected from the top 25 candidates produced automatically with masked language modeling. A different language model than the one used for the sentence embeddings of our benchmark was used to minimize potential biases (*bert-base-uncased*, (Devlin et al., 2018)). Specifically, verbs were masked while only functional elements (e.g. pronominal entities) surrounded them, retrieving appropriate agents and patients. For instance, to generate arguments for the verb *break*, the patient was masked in a sentence with a pronominal subject (e.g., *she broke (the/a/some/...) <MASK>*), and the subject was masked in a configuration with a pronominal patient to retrieve plausible agents (e.g., (the/a/some/...) <MASK> broke it). Figure 7 shows an example of instantiation of the BLM-COS for English (MINLEX).

The Italian lexical elements were translations of the English data, with manual adjustments where necessary to ensure sentence acceptability, semantic plausibility, and a balanced distribution of nouns across gender and number. An example of a MAXLEX version of an OD instantiation for Italian is shown in Figure 8.

CONTEXT
The chef melted the butter on the stove
The chef melted the butter by mistake
The butter was melted by the chef on the stove
The butter was melted by the chef by mistake
The butter was melted on the stove
The butter was melted by mistake
The butter melted on the stove
?
ANSWER SET
The butter melted by mistake
The chef melted by mistake
The butter was melted by the chef
The chef was melted by the butter
The butter melted the chef
The chef melted the butter
The butter melted by the chef
The chef melted by the butter

Figure 7: An example for COS for English with minimal lexical variation (MINLEX).

4.2. BLM-COS+: Change-of-State Food dataset

As shown in Figure 1, COS verb alternation can vary in its morphosyntactic strategies across languages. We built a dataset specifically for the COS verb class, but we limited our lexical domain to food-related verbs and constituents. Food-related contexts were chosen to clearly distinguish agent-like elements from patients. In choosing food-related terms, we aimed to make this dataset dynamic, easily translatable into many languages. At this stage,

CONTEXT
<i>Gli artisti dipinsero questi paesaggi in una giornata</i> 'The artists painted these landscapes in one day'
<i>Le attrici cantavano le stesse canzoni da qualche anno</i> 'The actresses had been singing the same songs for a few years'
<i>Dei loghi saranno disegnati dalla ditta in meno di un'ora</i> 'Some logos will be drawn by the company in less than one hour'
<i>Il grano deve essere seminato dal contadino da stamattina</i> 'The wheat must be sown by the farmer since (lit. by) this morning'
<i>I messaggi saranno recitati con la giusta interpretazione</i> 'The messages will be recited with the right interpretation'
<i>Le case furono intagliate da manuale</i> 'The houses were carved to perfection (lit. by the handbook)'
<i>Il falegname scolpiva nella cava di marmo</i> 'The carpenter was carving in the marble quarry'
?
ANSWER SET
<i>Una gonna si cuciva da qualche giorno</i> 'A skirt SI sew for (lit. by) a few days'
<i>Le artiste si scolpiscono da qualche mese</i> 'The artists SI sculpt for (lit. by) a few months'
<i>I capelli possono essere lavati dall'assistente</i> 'The hair can be washed by the assistant'
<i>I suoi genitori devono essere cucinati dal pranzo</i> 'His parents must be cooked by the lunch'
<i>Le uova cucinavano le mie amiche</i> 'The eggs were cooking my friends'
<i>Gli studenti pescano salmoni</i> 'The students fish salmon'
<i>Un articolo lesse dalle autrici</i> 'An article read by the authors'
<i>Il sottocuoco deve impastare dalla pizza</i> 'The sous-chef must knead by the pizza'
<i>Gli articoli scrivono da manuale</i> 'The articles write to perfection (lit. by the handbook)'
<i>Alcuni colleghi studiano da anni</i> 'Some colleagues have been studying for (lit. by) years'

Figure 8: An example for OD for Italian (with English glosses in grey) with maximal lexical variation (MAXLEX).

we have created data for English, German, and Italian.

Template The template uses relative clauses as a linguistic diagnostic. Relative clauses reveal the syntactic roles of the constituents, showing that the object of a transitive COS verb can appear either as the object in an object relative clause or as the syntactic subject in a subject relative clause in a COS intransitive configuration. The matrices can vary in the directionality of the alternants, they can start with the transitive verb in the context and probe the intransitive (T2I) or the inverse (I2T).

The answer set is minimal, since the answer only shows the combination of the arguments, their number and their position. Specifically, it includes one correct answer (the relevant alternant) and two types of errors: one involving a reverse-role transitive (object-verb-subject), and one involving an intransitive construction in which the agent appears as the subject. The reverse-role transitive is classified as a GRAMMAR error when the target sentence is the transitive alternant, and as a SEQUENCE error when the target is the intransitive alternant. The two templates (T2I and I2T) are given in Figure 9.

Data instantiation We selected 25 verbs for each language that show the very same structure. For

T2I CONTEXT		I2T CONTEXT	
SVO	Ag Verb Pa	SVO	Pa Verb
S _{ag} R	Ag rel Verb Pa	S _{pa} R	Pa rel Verb
OR	Pa rel Ag Verb	OR	Pa rel Ag Verb
S _{pa} R	Pa rel Verb	S _{ag} R	Ag rel Verb Pa
Answer Set		Answer Set	
GRAMMAR	Ag Verb	SEQUENCE	Ag Verb
CORRECT	Pa Verb	CORRECT	Ag Verb Pa
SEQUENCE	Pa Verb Ag	GRAMMAR	Pa Verb Ag

Figure 9: BLM templates for the T2I and I2T groups. Context sets and corresponding Answer sets are shown with colour-coded elements: agent (Ag), patient (Pa), verb (V), and relative markers (rel).

example, for Italian we excluded the verb *bollire* 'boil' since the intransitive form does not require the *si* (Vietri, 2019) and for German, we did not include separable verbs (e.g. *aufwärmen* 'reheat') to have comparable results across the morphological patterns.

We augmented the data using the DeepSeek-V3 model (Guo et al., 2025), adopting a hybrid approach that combines LLM generation with explicit linguistic constraints. This ensured the generated sentences remained grammatically well-formed and semantically coherent within each alternation paradigm. For each target language, the prompt requested the generation of 80 possible patients for each verb (25 interactions) and a total of 300 possible agents across verbs (1 interaction). To assess the quality of the augmented data, native speakers evaluated the combinations of the seed constituents.

As illustrated in Figure 1, German exhibits case marking that differentiates patients in object and subject positions. This marking is overtly distinct only for masculine nouns, but not for feminine or neuter constituents. Therefore, we created two datasets for German: one containing only masculine, case-marked constituents (Case or C) and one with masculine, feminine and neuter arguments (Mix).

4.3. BLM-CausH: Hebrew binyanim system

Hebrew verbs exhibit distinct morphological patterns that encode voice and causativity called *binyanim*. In our investigation we select to work on four forms (Samo and Merlo, 2026), presented in Figure 11, that encode different types of marking: PAAL represents the basic, labile transitive form; NIFAL typically marks the passive voice (Coffin and Bolozky, 2005, 71); HIFIL and HUFAL correspond to causative active and passive forms.

Template The context set contains three complete pairs of sentences and one incomplete pair. In the first three pairs, verbs are inflected according

CONTEXT
<i>Der Chef schmilzt den Käse</i> 'The.NOM chef melts the.ACC cheese'
<i>Der Chef der den Käse schmilzt</i> 'The.NOM chef that.NOM the.ACC cheese melts'
<i>Den Käse den der Chef schmilzt</i> 'The.ACC cheese that.ACC the.NOM chef melts'
<i>Der Käse der schmilzt</i> 'The.NOM that.NOM cheese melts'
ANSWER SET
<i>Der Chef schmilzt</i> 'The.NOM chef melts'
<i>Der Käse schmilzt</i> 'The.NOM cheese melts'
<i>Den Käse schmilzt der Chef</i> 'The.NOM cheese is melted by the.ACC chef'

Figure 10: MinLex for COS+T2IC (Case) for German. English glosses in grey. NOM = Nominative case, ACC = Accusative case.

Binyan	Form	Example
PAAL	כתב	<i>katav</i> 'wrote'
NIFAL	נכתב	<i>nixtav</i> 'was written'
HIFIL	הכתב	<i>hextiv</i> 'dictated'
HUFAL	הכתב	<i>huxtav</i> 'was dictated'

Figure 11: Binyanim under investigations and relative examples for the root KTB (related to the act of writing). Adapted from Kastner 2019, 574-575. The colour codes the different overt characters representing morphemes marking the alternation.

to a specific binyan. In the fourth pair, one sentence illustrates the remaining binyan, and the task is to infer the missing sentence. The answer set comprises four sentences, each exemplifying a different binyan. This is to be considered a type C template, since the model will rely on the forms but also on the number of arguments, their semantic nature and their context.

CONTEXT			
Paal	Paal	Paal	Nifal
Paal	Paal	Paal	Nifal
Nifal	Nifal	Hifil	Hifil
Nifal	...	Nifal	...
Hifil	Hufal	Hufal	Hufal
Hifil	Hufal	Hufal	Hufal
Hufal	Hifil	Nifal	Paal
?			
ANSWER SET			
Paal			
Nifal			
Hifil			
Hufal			

Figure 12: BLM template for the CausH.

Data instantiation We create two BLM datasets. One dataset, NATURAL, was created with natural occurring sentences extracted from UD treebanks. The second dataset, SYNTHETIC, only contains the verb corresponding to each sentence in the SENTENCE dataset, with no additional lexical content. These sentences can be grammatically correct since Hebrew allows both subject and object drop (Vainikka and Levy, 1999; Erteschik-Shir et al., 2013). The data were retrieved using *match.grew.fr* (Guillaume, 2021), by querying a simple variable X with the relevant annotation in two treebanks of Hebrew containing respectively news (HBT, Tsarfaty 2013; McDonald et al. 2013; 114,648 tokens, 6,143 trees) and encyclopaedic entries (IAHLWiki, henceforth IW; Zeldes et al. 2022; 103,395 tokens; 5,039 trees). Both treebanks are version 2.15. The query collects sentences where the main verb is annotated with relevant binyan for the morphosyntactic property HEBBINYAN².

Both datasets only instantiate the BLM templates with the most complex form of lexical variation (MAXLEX). The lack of MINLEX is because not all roots are compatible with all templatic structures (Kastner, 2019). All the wrong candidates are errors of GRAMMAR.

5. Experiments

Building on prior studies (Nastase and Merlo, 2024), we primarily apply a Feed-Forward Neural Network (FFNN) to establish a benchmark across the system and various models. Specifically, we use a FFNN architecture as a baseline system, using sentence embeddings derived from the curated datasets as described in (Samo et al., 2023). For each sentence in the BLM, we compute the averaged token embeddings (the token embeddings are calculated with the multilingual model Electra (*google/electra-base-discriminator*)). The FFNN receives the stacked embeddings that represent the context, is trained with a max-margin loss, and selects the answer whose embedding has the highest cosine similarity to the network’s output. For Hebrew, we also compare the results of a monolingual model (AlephBERT, *onlplab/alephbert-base*). Each dataset is then spit into training and testing. To avoid data leakage, we use disjoint sets of training and testing instances. Statistics on the dataset are given in Table 2.

Performance in terms of F1 scores is visualized in Figure 14. English consistently achieves the highest scores across most tasks, with higher marks in Template B (En-COS+T2I: 0.983, En-COS+I2T: 0.971). Italian demonstrates performance scores generally lower than English. For template B, Italian shows higher performance in MAXLEX when

Data	Lang	Lex	Train	Test
COS	En, It	MIN/MAX	2700	300
OD	En, It	MIN/MAX	2700	300
COS+T2I	De _{Case} De _{mix} En, It	MIN/MAX	1600	400
COS+I2T	De _{Case} De _{Mix} En, It	MIN/MAX	1600	400
Caus	He _{Nat} He _{synth}	MAX	7200	800

Table 2: Datasets, languages, training and testing instances for each language and lexical variation.

the target answer is the marked intransitive (It-COS+T2I: 0.733). German shows higher performance in the dataset when the patient in COS verbs is case-marked differently in the two alternants (DeEN-COS+IT2C: 0.950). In Hebrew the monolingual (He_MO) model outperform the multilingual (He_ML) in the natural settings, but multilingual and monolingual are comparable in the synthetic setting. Finally, we can observe that MINLEX is easier than MAXLEX across languages, except for OD English.

The error analysis is presented in Figure 15. Overall, the FFNN learns the SEQUENCE pattern of the BLM, but, with many grammatical errors.

COS and OD primarily exhibit grammatical errors (see Table 3 in the Appendix). The most prominent error in both languages involves syntax-semantic mapping mismatches, where the subject of an intransitive verb is interpreted as an agent rather than as a patient (e.g., *The artist broke* instead of *The vase melts*). In English OD MaxLex, the same probability for this error (4%) is assigned to the candidate answer in which a nominal phrase introduced by the preposition *by* is incorrectly interpreted as agentive (e.g., *The vase breaks by the artist*). In English OD MinLex, a common SEQUENCE error involves a transitive sentence with an agentive subject (e.g., *The artist paints the vase*). In Italian COS, we also observe that the absence of the reflexive-like element *si* (*il burro stava sciogliendo da qualche minuto*) or its presence in OD (*Le artiste si scolpiscono da qualche mese*) is not a prominent error.

With respect to COS+, we observe that, for German, BLMs comprising only case-marked constructions perform better than mixed ones, especially in I2T conditions. English shows very few errors, while Italian does not show differences between the T2I and I2T configurations. Also in this case the errors are mainly error of syntax-semantic interface, (e.g., *The chef melted* instead of *The butter melts*).

Figure 16 shows a confusion matrix in terms of error probability of Hebrew. All the errors in the Hebrew data are grammatical. In the monolingual both datasets, we do not observe a preferred an-

²E.g. PAAL: *pattern* X [HebBinyan = "PAAL"]

NATURAL	SYNTHETIC
יציאתו של ליר גררה בתוך דקה את צימצום התוצאה ל-88 91. 'Lif's exit <u>led</u> , within a minute, to narrowing the score to 91-88.' (HTB, 5728)	גררה lead.PAAL
אתי כיום עובדת כמוקדנית במוקד שירות לקוחות. 'Etti currently <u>works</u> as a call center representative in a customer service center.' (iahltwiki_hameila-babank-lemishar-76)	עובדת work.PAAL
על כך נענשו שני הצדדים. 'For this, both sides <u>were punished</u> .' (HTB, 5908)	נענשו punish.NIFAL
הגוש המזרחי ודרום אפריקה נמנעו בהצבעה. 'The Eastern Bloc and South Africa <u>abstained</u> in the vote.' (iahltwiki_unrwa-14)	נמנעו abstain.NIFAL
אפשר כבר להכין את הכרטיסים. 'You can already <u>prepare</u> the tickets.' (HTB, 5744)	להכין prepare.HIFIL
בנוסף להשפעתו על המוח, משפיע האמפטמין גם על הגוף. 'In addition to its effect on the brain, amphetamine also <u>affects</u> the body.' (iahltwiki_amphetamine-38)	משפיע affect.HIFIL
שני ניצחונות בית צפויים הושגו בכפר בלום ובחולון. 'Two expected home victories were <u>achieved</u> in Kfar Blum and in Holon.' (HTB, 5758)	הושגו achieve.HUFAL
פריטוניטיס- דלקת בצפק, מוגדרת כעליה בתאי דם לבנים בנוזל הפרטונאלי. 'Peritonitis is <u>defined</u> as an increase in white blood cells in the peritoneal fluid.' (iahltwiki_dialysis-39)	מוגדרת define.HUFAL
זיהומים הקשורים לקתטר 'Infections that <u>associate</u> with the catheter' (iahltwiki_dialysis-44)	הקשורים associate.PAAL
על קיר זה פסיפס שנוצר בשנת 1990. 'On this wall, there is a mosaic <u>created</u> in 1990.' (iahltwiki_holy-sepulchre-256)	שנוצר create.NIFAL
ב-1985 הקליט איינשטיין עם שם טוב לוי את האלבום "תוצרת הארץ". 'In 1985, Einstein <u>recorded</u> the album Totzeret Haaretz with Shem-Tov Levi' (iahltwiki_arik-einstein-274)	הקליט record.HIFIL

Figure 13: An example of instantiation of sentences (NATURAL) extracted from UD and the derived synthetic data made by the verbs (SYNTHETIC). Glosses are provided and the verb of the sentence is highlighted in gray.




Figure 14: F1 scores across datasets.

swer. On the other hand, the multilingual natural data shows an error pattern in line with a chance level performance and in the multilingual synthetic data there is clear confusion between HIFIL and HUFAL, two verbal voices that share common morphological marking.

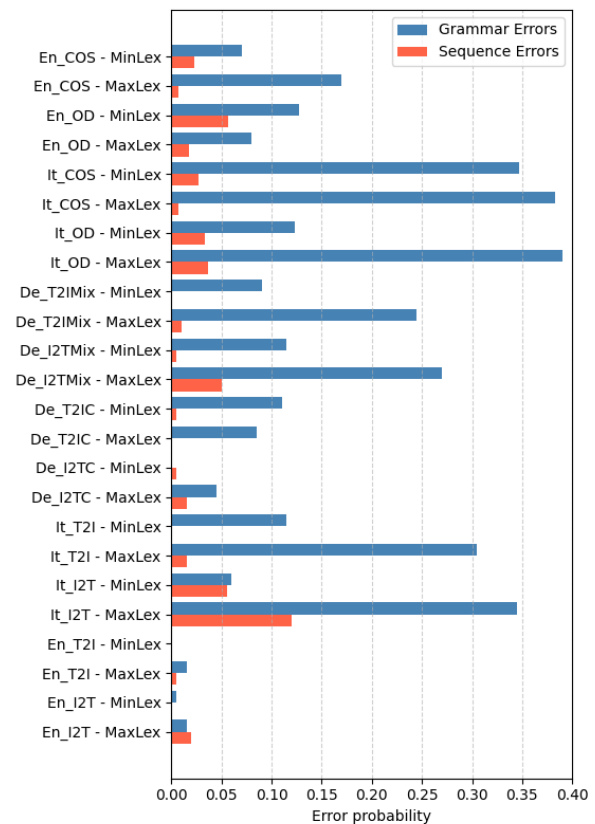


Figure 15: Error analysis across datasets.

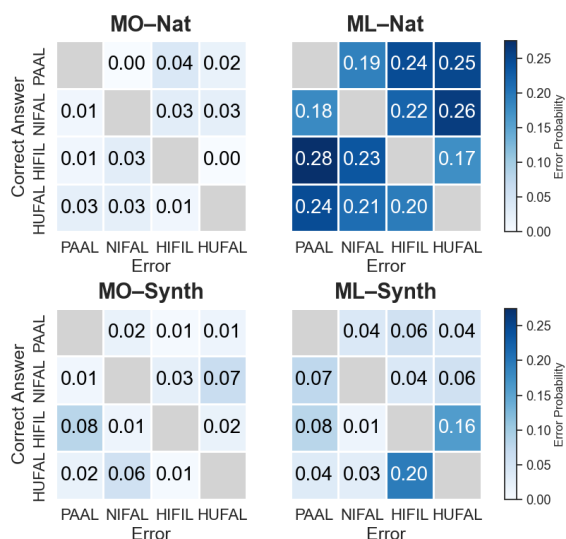


Figure 16: Confusion matrices for Hebrew across datasets and models, for each verbal voice, reported as proportions.

6. Discussion

BLM curated datasets allow us to probe systematic grammatical knowledge in models by providing controlled, contrastive templates that isolate the phenomena under investigation. Templates play a crucial role in our approach: they are designed both to isolate the linguistic phenomenon under investigation and to instruct the model in how to generalize from controlled input.

The templates presented here show different types of complexity. The BLM template used for the dataset BLM-COS and BLM-OD (type A) requires the model to pay attention to an extra-linguistic phenomenon to find the sequence, and in so doing it helps identify the linguistic problem at stake. This level can be seen as *analogical* in a weak sense: the model learns to associate surface patterns without necessarily accessing the internal structure of the target phenomenon.

The BLM template for the dataset BLM-COS+ (Type B) forces analogical representations in a stronger sense: the model must map relational structures rather than surface similarities. This step introduces a form of *abductive* inference, since the model must hypothesize the most plausible underlying relation that accounts for the observed patterns.

The BLM template adopted for the dataset BLM-CausH focuses on the relation across elements within a paradigm. This third type targets *systematicity* — the capacity to represent a grammatical paradigm. Systematicity implies that once a model acquires a relation within a system, it should be able to extend that relation to other members of the same paradigm in a consistent manner.

Finally, results from our baseline experiment suggest that multilingual models achieve comparable cross-linguistic performance with synthetic data. Comparing monolingual, as well as other generative models, will provide additional avenues for investigation in the understanding of linguistic knowledge in models.

7. Conclusions

In this work, we presented datasets in several languages within the framework of the Blackbird Language Matrices (BLMs), designed to probe linguistic knowledge in LLMs for verb alternations. By combining carefully curated templates with linguistically informed data augmentation strategies, we provide datasets for testing cross-sentence syntactic and semantic generalizations.

Overall, our work demonstrates the importance of linguistically-principled dataset design and provides resources for cross-linguistic evaluation. It also establishes a foundation for future research investigating the limits of model generalization across paradigms and languages.

Limitations

Despite the controlled design and cross-linguistic scope of the proposed datasets, several limitations should be acknowledged. The datasets focus on a restricted set of verb alternation phenomena and languages. While English, Italian, German, and Hebrew provide typologically diverse case studies, they do not fully represent the range of morphosyntactic strategies found cross-linguistically.

The experimental evaluation is intentionally limited to baseline architectures and embedding-based representations. The reported experiments are designed to demonstrate the diagnostic usefulness of the datasets and the task rather than to provide a comprehensive comparison across state-of-the-art LLMs. Consequently, the results should not be interpreted as definitive statements about current model capabilities.

Ethics

A portion of the data in the datasets used in this work are derived from publicly available corpora, including sources such as news articles and other open-access materials. These datasets may contain content that is sensitive or potentially distressing. We recognize that some individuals may find certain topics upsetting.

Acknowledgments

We gratefully acknowledge the support of this work by the Swiss National Science Foundation, through grant SNF Advanced grant TMAG-1_209426 to PM.

8. Bibliographical References

- Judit Acs, Endre Hamerlik, Roy Schwartz, Noah A. Smith, and Andras Kornai. 2024. [Morphosyntactic probing of multilingual BERT models](#). *Natural Language Engineering*, 30(4):753–792.
- Furkan Akkurt, Onur Gungor, Büşra Marşan, Tunga Gungor, Balkiz Ozturk Basaran, Arzucan Özgür, and Susan Uskudarli. 2024. [Evaluating the Quality of a Corpus Annotation Scheme Using Pre-trained Language Models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 6504–6514, Torino, Italia. ELRA and ICCL.
- Aixiu An, Chunyang Jiang, Maria A. Rodriguez, Vivi Nastase, and Paola Merlo. 2023. [BLM-AgrF: A new French benchmark to investigate generalization of agreement in neural networks](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1363–1374, Dubrovnik, Croatia.
- Maya Arad. 2005. *Roots and Patterns: Hebrew Morpho-Syntax*. Springer, New York.
- Khuyagbaatar Batsuren, Omer Goldman, Salam Khalifa, Nizar Habash, Witold Kieraś, Gábor Bella, Brian Leonard, Garrett Nicolai, Kyle Gorman, Yustinus Ghanggo Ate, Maria Ryskina, Sabrina Mielke, Elena Budianskaya, Charbel El-Khaissi, Tiago Pimentel, Michael Gasser, William Abbott Lane, Mohit Raj, Matt Coler, Jaime Rafael Montoya Samame, Delio Siconatzi Camaiteri, Esaú Zumaeta Rojas, Didier López Francis, Arturo Oncevay, Juan López Bautista, Gema Celeste Silva Villegas, Lucas Torroba Hennigen, Adam Ek, David Guriel, Peter Dirix, Jean-Philippe Bernardy, Andrey Scherbakov, Aziyana Bayyr-ool, Antonios Anatasopoulos, Roberto Zariquiey, Karina Sheifer, Sofya Ganieva, Hilaria Cruz, Ritván Karahóga, Stella Markantonatou, George Pavlidis, Matvey Plugaryov, Elena Klyachko, Ali Salehi, Candy Angulo, Jatayu Baxi, Andrew Krizhanovsky, Natalia Krizhanovskaya, Elizabeth Salesky, Clara Vania, Sardana Ivanova, Jennifer White, Rowan Hall Maudslay, Josef Valvoda, Ran Zmigrod, Paula Czarnowska, Irene Nikkarinen, Aelita Salchak, Brijesh Bhatt, Christopher Straughn, Zoey Liu, Jonathan North Washington, Yuval Pinter, Duygu Ataman, Marcin Wolinski, Totok Suhardijanto, Anna Yablonskaya, Niklas Stoehr, Hossep Dolatian, Zahroh Nuriah, Shyam Ratan, Francis M. Tyers, Edoardo M. Ponti, Grant Aiton, Aryaman Arora, Richard J. Hatcher, Ritesh Kumar, Jeremiah Young, Daria Rodionova, Anastasia Yemelina, Taras Andrushko, Igor Marchenko, Polina Mashkovtseva, Alexandra Serova, Emily Prud’hommeaux, Maria Nepomniashchaya, Fausto Giunchiglia, Eleanor Chodroff, Mans Hulden, Miikka Silfverberg, Arya D. McCarthy, David Yarowsky, Ryan Cotterell, Reut Tsarfaty, and Ekaterina Vylomova. 2022. [UniMorph 4.0: Universal Morphology](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 840–855, Marseille, France. European Language Resources Association.
- Yonatan Belinkov. 2022. [Probing classifiers: Promises, shortcomings, and advances](#). *Computational Linguistics*, 48(1):207–219.
- Jonathan David Bobaljik. 2015. Suppletion: Some theoretical implications. *Annual Review of Linguistics*, 1(1):1–18.
- François Chollet. 2019. [On the Measure of Intelligence](#). *CoRR*, abs/1911.01547.
- Mario Civico. 2025. [Linguistic statistical Universals: comparing Computer- and Human-generated texts](#). *International Journal of Digital Humanities*, 7:1–37.
- Edna Amir Coffin and Shmuel Bolozky. 2005. *A Reference Grammar of Modern Hebrew*. Cambridge University Press.
- Marie-Catherine De Marneffe, Christopher D Manning, Joakim Nivre, and Daniel Zeman. 2021. Universal dependencies. *Computational linguistics*, 47(2):255–308.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Nomi Erteschik-Shir, Lena Ibnbari, and Sharon Taube. 2013. Missing objects as topic drop. *Lingua*, 136:145–169.
- Prakhar Gautam, Jitendra Singh Thakur, and Ashish Mishra. 2024. Subject–verb agreement error handling using rnn architectures. In *International Conference on Innovations in Computational Intelligence and Computer Vision*, pages 215–224. Springer.

- Francesca Grasso, Vladimiro Lovera Rulfi, and Luigi Di Caro. 2024. [Verbalignet: Unlocking multilingual exploration of verbal arguments](#). In *Selected Challenges in Grammar-Based Text Generation from the Semantic Web*, pages 3–17. Springer Nature Switzerland.
- Thomas L Griffiths, Jian-Qiao Zhu, Erin Grant, and R Thomas McCoy. 2024. Bayes in the age of intelligent machines. *Current Directions in Psychological Science*, 33(5):283–291.
- Bruno Guillaume. 2021. [Graph matching and graph rewriting: GREW tools for corpus exploration, maintenance and conversion](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 168–175, Online. Association for Computational Linguistics.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, et al. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *arXiv preprint*, arXiv:2501.12948.
- Martin Haspelmath, Andreea Calude, Michael Spagnol, Heiko Narrog, and Elif Bamyacı. 2014. Coding causal–noncausal verb alternations: A form–frequency correspondence explanation. *Journal of linguistics*, 50(3):587–625.
- Elisabetta Jezek. 2024. Semantics first! Rethinking the role of syntax in verb classification. *Lingue e linguaggio*, 23(1):9–30.
- Chunyang Jiang and Paola Merlo. 2026. Analogical structure, minimal contextual cues and contrastive distractors: Input design for sample-efficient linguistic rule induction. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2026)*.
- Chunyang Jiang, Giuseppe Samo, Vivi Nastase, and Paola Merlo. 2024. [BLM-It - Blackbird Language Matrices for Italian: A CALAMITA Challenge](#). In *Proceedings of the 10th Italian Conference on Computational Linguistics (CLiC-it 2024)*, pages 1135–1143, Pisa, Italy. CEUR Workshop Proceedings.
- Jaap Jumelet, Leonie Weissweiler, Joakim Nivre, and Arianna Bisazza. 2025. MultiBLIMP 1.0: A Massively Multilingual Benchmark of Linguistic Minimal Pairs. *arXiv preprint arXiv:2504.02768*.
- Katharina Kann, Alex Warstadt, Adina Williams, and Samuel R. Bowman. 2019. [Verb argument structure alternations in word and sentence embeddings](#). In *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*, pages 287–297.
- Itamar Kastner. 2019. Templatic morphology as an emergent property: Roots and functional heads in Hebrew. *Natural Language & Linguistic Theory*, 37:571–619.
- Stav Klein and Reut Tsarfaty. 2020. [Getting the ##life out of living: How adequate are word-pieces for modelling complex morphology?](#) In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 204–209, Online. Association for Computational Linguistics.
- Beth Levin. 1993. *English Verb Classes and Alternations A Preliminary Investigation*. University of Chicago Press, Chicago and London.
- Tal Linzen and Marco Baroni. 2021. [Syntactic structure from deep learning](#). *Annual Review of Linguistics*, 7(1):195–212.
- John J. McCarthy. 1979. *Formal Problems in Semitic Phonology and Morphology*. Ph.D. thesis, Massachusetts Institute of Technology.
- Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar Täckström, Claudia Bedini, Núria Bertomeu Castelló, and Jungmee Lee. 2013. [Universal Dependency Annotation for Multilingual Parsing](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 92–97, Sofia, Bulgaria. Association for Computational Linguistics.
- Maria Meo, Maxwell J. Roberts, and Francesco S. Marucci. 2007. [Element salience as a predictor of item difficulty for raven’s progressive matrices](#). *Intelligence*, 35(4):359–368.
- Paola Merlo. 2023a. Blackbird language matrices (BLM), a new task for rule-like generalization in neural networks: Can Large Language Models pass the test? In *Findings of the 2023 Conference on Empirical Methods in Natural Language Processing*.
- Paola Merlo. 2023b. [Blackbird language matrices \(blm\), a new task for rule-like generalization in neural networks: Motivations and formal specifications](#). *arXiv*, <https://arxiv.org/abs/2306.11444>.
- Paola Merlo and Suzanne Stevenson. 2001. Automatic verb classification based on statistical distributions of argument structure. *Computational Linguistics*, 27(3):373–408.
- Roser Morante and Eduardo Blanco. 2021. Recent advances in processing negation. *Natural Language Engineering*, 27(2):121–130.

- Mihai Nadas, Laura Diosan, and Andreea Tomescu. 2025. Synthetic Data Generation Using Large Language Models: Advances in Text and Code. *arXiv preprint arXiv:2503.14023*.
- Vivi Nastase and Paola Merlo. 2024. [Tracking linguistic information in transformer-based sentence embeddings through targeted sparsification](#). In *Proceedings of the 9th Workshop on Representation Learning for NLP (RepL4NLP-2024)*, pages 203–214, Bangkok, Thailand. Association for Computational Linguistics.
- Vivi Nastase, Giuseppe Samo, Chunyang Jiang, and Paola Merlo. 2024a. Exploring Italian Sentence Embeddings Properties through Multi-tasking. In *Proceedings of the Tenth Italian Conference on Computational Linguistics (Clic-It 2024)*, pages 1–10.
- Vivi Nastase, Giuseppe Samo, Chunyang Jiang, and Paola Merlo. 2024b. Exploring syntactic information in sentence embeddings through multilingual subject-verb agreement. In *Proceedings of the Tenth Italian Conference on Computational Linguistics (Clic-It 2024)*, pages 1–13.
- Garrett Nicolai, Eleanor Chodroff, Frederic Mailhot, and Çağrı Çöltekin, editors. 2024. *Proceedings of the 21st SIGMORPHON workshop on Computational Research in Phonetics, Phonology, and Morphology*. Association for Computational Linguistics, Mexico City, Mexico.
- Mattia Proietti, Gianluca Leboni, and Alessandro Lenci. 2022. [Does BERT recognize an agent? modeling Dowty’s proto-roles with contextual embeddings](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4101–4112, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Malka Rappaport Hovav and Beth Levin. 2024. [Variable agentivity: Polysemy or underspecification](#). *Glossa: a journal of general linguistics*, 9(1).
- John C. Raven. 1938. Standardization of progressive matrices. *British Journal of Medical Psychology*, 19:137–150.
- Abhilasha Ravichander, Yonatan Belinkov, and Eduard Hovy. 2021. [Probing the probing paradigm: Does probing accuracy entail task relevance?](#) In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3363–3377, Online. Association for Computational Linguistics.
- Tanja Samardžić and Paola Merlo. 2018. The probability of external causation: An empirical account of crosslinguistic variation in lexical causatives. *Linguistics*, 56(5):895–938.
- Giuseppe Samo and Paola Merlo. 2026. [Modelling the Morphology of Verbal Paradigms: A Case Study in the Tokenization of Turkish and Hebrew](#). *arXiv*, <https://arxiv.org/abs/2602.05648>.
- Giuseppe Samo, Vivi Nastase, Chunyang Jiang, and Paola Merlo. 2023. [BLM-s/IE: A structured dataset of English spray-load verb alternations for testing generalization in LLMs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12276–12287, Singapore. Association for Computational Linguistics.
- Tijs Schepens, Marten van der Klis, and Willem Zuidema. 2025. Large language models as corpus generators: Opportunities and risks. *arXiv preprint arXiv:2503.11245*.
- Luca Simonetti, Elisabetta Ježek, and Guido Vetere. 2024. [Subcategorization of Italian Verbs with LLMs and T-PAS](#). In *Proceedings of the 10th Italian Conference on Computational Linguistics (CLiC-it 2024)*. CEUR Workshop Proceedings.
- Reut Tsarfaty. 2004. ‘binyanim ba’avir’: An investigation of Aspect Semantics in Modern Hebrew. Master’s thesis, Universiteit van Amsterdam, Institute for Logic Language and Computation.
- Reut Tsarfaty. 2013. [A unified morpho-syntactic scheme of Stanford dependencies](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 578–584, Sofia, Bulgaria. Association for Computational Linguistics.
- Anne Vainikka and Yonata Levy. 1999. Empty subjects in Finnish and Hebrew. *Natural Language & Linguistic Theory*, 17(3):613–671.
- Simonetta Vietri. 2019. [I verbi ad alternanza causativa. classificazione lessico-grammaticale](#).
- Alex Warstadt, Yu Cao, Ioana Grosu, Wei Peng, Hagen Blix, Yining Nie, Anna Alsop, Shikha Bordia, Haokun Liu, Alicia Parrish, Sheng-Fu Wang, Jason Phang, Anhad Mohanane, Phu Mon Htut, Paloma Jeretic, and Samuel R. Bowman. 2019. [Investigating BERT’s knowledge of language: Five analysis methods with NPIs](#). 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 2877–2887, Hong Kong, China. Association for Computational Linguistics.

- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. BLiMP: The Benchmark of Linguistic Minimal Pairs for English. In *Transactions of the Association for Computational Linguistics*, volume 8, pages 391–406.
- Michael Wilson, Jackson Petty, and Robert Frank. 2023. How abstract is linguistic generalization in large language models? experiments with argument structure. *Transactions of the Association for Computational Linguistics*, 11:1377–1395.
- David Yi, James Bruno, Jiayu Han, Peter Zukerman, and Shane Steinert-Threlkeld. 2022. [Probing for understanding of English verb classes and alternations in large pre-trained language models](#). In *Proceedings of the Fifth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 142–152, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Amir Zeldes, Nick Howell, Noam Ordan, and Yifat Ben Moshe. 2022. [A second wave of UD Hebrew treebanking and cross-domain parsing](#). In *Proceedings of EMNLP 2022*, pages 4331–4344, Abu Dhabi, UAE.
- Lingze Zhang and Ellie Pavlick. 2025. [Does training on synthetic data make models less robust?](#) *arXiv*, <https://arxiv.org/abs/2502.07164>.

Appendix

ENGLISH					
Wrong Answer	Error Type	COSMINLEX	COSMAXLEX	ODMINLEX	ODMAXLEX
1 Pat Akt by-NP	GRAMMAR	-	-	35	12
2 Ag Akt by-NP	GRAMMAR	21	49	-	-
3 Pat Pass by-Ag	SEQUENCE	1	1	0	0
4 Ag Pass by-Pat	SEQUENCE	4	0	1	2
5 Pat Akt Ag	SEQUENCE	2	1	3	0
6 Ag Akt Pat	SEQUENCE	0	0	13	3
7 Pat Akt by-Pat	GRAMMAR	0	2	0	0
8 Ag Akt by-Pat	GRAMMAR	0	0	3	12

ITALIAN					
Wrong Answer	Error Type	COSMINLEX	COSMAXLEX	ODMINLEX	ODMAXLEX
1 Pat SI Akt by-NP	GRAMMAR	-	-	0	7
2 Ag SI Akt by-NP	GRAMMAR	103	109	5	15
3 Pat Pass by-Ag	SEQUENCE	1	7	0	0
4 Ag Pass by-Pat	SEQUENCE	7	2	1	0
5 Pat Akt Ag	SEQUENCE	0	0	4	9
6 Ag Akt Pat	SEQUENCE	0	0	5	2
7 Pat Akt by-Ag	GRAMMAR	0	4	5	9
8 Ag Akt by- Pat	GRAMMAR	0	1	5	15
9 Pat Akt by-NP	GRAMMAR	26	76	1	0
10 Ag Akt by-NP	GRAMMAR	21	49	-	-

Table 3: Raw counts of errors for BLM-COS and OD in English and Italian across lexical variation conditions (test set $n = 300$).