A Crosslinguistic Database for Combinatorial and Semantic Properties of Attitude Predicates

Deniz Özyıldız
Universität Konstanz
deniz.ozyildiz@uni-konstanz.de

Ciyang Qing
University of Edinburgh
cqing@ed.ac.uk

Floris Roelofsen
Universiteit van Amsterdam
f.roelofsen@uva.nl

Maribel Romero
Universität Konstanz
maribel.romero@uni-konstanz.de

Wataru Uegaki
University of Edinburgh
w.uegaki@ed.ac.uk

Abstract

We introduce a crosslinguistic database for attitude predicates, which references their combinatorial (syntactic) and semantic properties. Our data allows assessment of crosslinguistic generalizations about attitude predicates as well as discovery of new typological/crosslinguistic patterns. This paper highlights empirical and theoretical issues that our database will help to address, motivates the predicate sample and the properties that it references, as well as our methodological choices. Two case studies illustrate how the database can be used to assess the validity of crosslinguistic generalizations.

1 Introduction

Attitude predicates are natural language expressions characterized by the fact that they combine with sentential complements and that they ascribe to their subject an attitude. They are used to talk about what people believe, wonder, hope, or say. These predicates exhibit a variety of combinatorial restrictions in terms of the types of clauses they can combine with. For example, they can be distinguished into three classes based on whether they are compatible with declarative or question complements: *Antirogatives* like *believe* combine only with declaratives, in (1a). *Rogatives* like *wonder* combine only with interrogatives, in (1b). And *responsives* like *know* combine with either, in (1c).

(1) a. Al believes that/*whether Jo is Dutch.
b. Al wonders *that/whether Jo is Dutch.
c. Al knows that/whether Jo is Dutch.

Other instances of combinatorial restrictions include responsive predicates that are compatible with constituent questions (*who*, *what*, *which*, etc.) while being incompatible with *whether* questions, e.g., *be amazed* or *be surprised*, and predicates that differ in terms of whether they are compatible with indicative or subjunctive complements in languages that make the distinction.

In a tradition tracing back at least to Frege (1898 [1948]), attitude ascriptions have been studied extensively in the philosophical and the linguistic literature. One recent strand of research argues that differences in the combinatorial properties of attitude predicates, rather than being accidental and idiosyncratic facts, can be explained generally on the basis of their semantic properties (Zuber, 1982; Egré, 2008; Mayr, 2019; Theiler et al., 2019; Uegaki and Sudo, 2019). We elaborate on some of these semantic properties and how they might relate to attitude verbs' combinatorial properties in Section 2. A second, intimately connected strand of research aims to uncover semantic properties that classes of attitude predicates have in common (in addition to places of variation), within a given language's lexicon and across languages, i.e., crosslinguistic universals in the attitude domain (White and Rawlins, 2016; Roelofsen and Uegaki, 2020; Steinert-Threlkeld, 2019; Maldonado et al., 2022).

In this paper, we present a database that will allow researchers to address these questions and explore other linguistic properties of attitude predicates in a crosslinguistic way. The database references a sample of semantic and combinatorial properties of approximately 50 attitude predicates from 15 languages. The values of these properties are based on introspective judgments of native speakers of each language, and are collected by means of a questionnaire. They are summarized in tables in CSV format, one per language and speaker, which are accompanied by text documents that contain the linguistic examples that motivate the speaker's responses and reference additional facts about the data (e.g., the variety of the language spoken by the native speaker consultant, particular clause type distinctions available in the language, etc.).

This resource adds to a set of existing databases about the properties of attitude predicates: The *Mega* databases MegaAcceptability (White and Rawlins, 2016), MegaVeridicality (White and

Proceedings of the 5th Workshop on Research in Computational Linguistic Typology and Multilingual NLP (SIGTYP 2023), pages 65–75

Rawlins, 2018), MegaNegRaising (An and White, 2020), MegaIntensionality (Kane et al., 2021) and MegaOrientation (Moon and White, 2020), as well as the ZAS Database of Clause-embedding Predicates (Stiebels et al., 2018). The contribution of our database is novel in at least three respects. First, it enables a *crosslinguistic* exploration of the properties of attitude predicates. This is important because generalizations that concern these predicates are often formulated on the basis of a single language and yet, given their nature, are expected to hold crosslinguistially. Second, it is the same speakers that provide the introspective judgments that underlie the semantic and combinatorial properties that are tested. To the extent that we can assume that these judgments come from the same source grammar, within speaker and within language comparisons can be made consistently. It has been shown that speakers may differ from one another in terms of how strongly a linguistic expression displays some property, and that correlations between syntactic or semantic properties may ultimately depend on this gradient perception (Chemla et al., 2011; Tonhauser et al., 2018). Third, the quantitative component of the database (the summary tables in CSV format) is supported by a qualitative component (the text documents with examples and other considerations supporting/qualifying the consultant's judgments). This makes it possible not only to draw broad generalizations, but also to examine the properties of specific predicates in more depth. We would finally like to highlight that the dataset may be used for a broad range of applications in NLP, including but not limited to improving and evaluating the performance of natural language understanding and machine translation systems. This, we believe, is particularly valuable in that our dataset references several 'low resource' languages, for which such systems might perform poorly.

Outline Section 2 of this paper presents the semantic properties of attitude predicates included in our database and how these have been argued to relate to these predicates' combinatorial properties. Section 3 references the predicates that we have included, as well as the response categories that were used to elicit these predicates' semantic and combinatorial properties. Section 4 contains practical information about how the database is formatted, can be accessed, and further contributed to. Section 5 presents two case studies illustrating how the database can be used to test generalizations concerning attitude predicates. Section 6 concludes. (We draw attention to Limitations in the Appendix.)

2 Semantic Properties

This section introduces the semantic properties of attitude predicates included in our database and relevant generalizations about them in the literature.

A predicate V is *veridical* iff *x Vs that S* entails S. For instance, *know* is veridical, but *be certain* is not: (2) entails that it is raining but (3) does not.

(2) Alice knows that it is raining.

(3) Alice is certain that it is raining.

Veridicality is argued to correlate with the ability to take interrogative complements (e.g., Egré, 2008).

A predicate is *projective under negation* (or *projective* for short) if one can infer the complement when the predicate is negated. For instance, *be happy* and *be surprised* are projective (4).

(4) Alice isn't happy/surprised that it is raining
$\rightsquigarrow$ It is raining

A predicate V is *neg-raising* if *not V S* is interpreted as *V not S*. For instance, *think* and *believe* are neg-raising (5), whereas *know* and *be sure* are not (6).

(5) Alice does not think/believe it is raining
$\approx$ Alice thinks/believes it is not raining.

(6) Alice doesn't know/isn't sure it is raining
$\not\approx$ Alice knows/is sure it is not raining.

It has been suggested that neg-raising predicates are generally anti-rogative, and several theoretical explanations for this have been proposed (Zuber, 1982; Mayr, 2019; Theiler et al., 2019).

Many predicates, such as *be happy* and *hope*, have meanings that intuitively involve a notion of preference. Several formal semantic accounts characterize *preferentiality* in terms of *focus sensitivity* and *gradability* (Villalta, 2008; Romero, 2015; Uegaki and Sudo, 2019). A predicate V is *focus sensitive* if its truth conditions can be influenced by the placement of focus in the embedded clause. For instance, *be happy* and *hope* are focus sensitive because the two sentences in (7) need not be true at the same time: Mary might be the best among syntax teachers, but syntax might not be the best among subjects Mary can teach.

(7) a. Alice is happy/hopes that

MARY will teach syntax.
b. Alice is happy/hopes that Mary will teach SYNTAX.

In contrast, *know* and *think* are not focus sensitive. If one sentence in (8) is true of Alice's epistemic/doxastic state, the other must be true as well.

(8) a. Alice knows/thinks that MARY will teach syntax.
b. Alice knows/thinks that Mary will teach SYNTAX.

A predicate is *gradable* if it can participate in degree constructions, e.g., intensification (9) or comparison (10).

(9) Alice is very happy that Mary is here.

(10) Alice hopes that it is raining more than Bob does.

Karttunen (1977) observes that a certain class of preferential predicates, which he calls *emotive factives*, can take *wh*-questions but not *whether* questions (11) (see Section 5.1 for further discussion, and Saebø (2007) and Abenina-Adar (2019) for challenges). Uegaki and Sudo (2019) suggest that non-veridical preferential predicates such as *hope* cannot take embedded questions altogether (12).

(11) It is amazing what they serve for breakfast / *whether they serve breakfast.

(12) *Alice hopes whether Bob left / who left.

There is no consensus on exactly how to characterize emotive factives (see, e.g., Egré, 2008, for discussion), but it is uncontroversial that when they take a declarative complement, the attitude holder must believe that the complement is true (13).

(13) Alice is happy/surprised that it is raining
⇒ Alice believes that it is raining

There is a complication, however. It is unclear what level of credence *believe* corresponds to, since this attitude predicate can often be used when the subject is not fully certain that the complement is true (e.g., Hawthorne et al., 2016). Therefore, in our database we instead directly test the compatibility between a predicate and various levels of credence. For instance, a predicate V *always implies likelihood* if *x Vs that S* entails that *x* considers *S* more likely than *not S*.

For question-embedding predicates, one important semantic property is what can be inferred about the relation between the subject's belief and possible answers to the embedded question. Some predicates, such as *know*, entail that there is a possible answer to the embedded question that the subject believes (14). Such predicates are *belief-implying*. Some predicates, such as *wonder*, entail that there is no possible answer that the subject believes (15). Such predicates are *ignorance-implying*. Other predicates, such as *care*, are *neutral wrt belief and ignorance*. *Alice cares (about) who won* can be true with or without Alice having a belief as to who won (Elliott et al., 2017).

(14) Alice knows whether Bob left.
⇒ Alice believes that Bob left or she believes that Bob didn't leave.

(15) Alice wonders whether Bob left.
⇒ Alice neither believes that Bob left nor does she believe that Bob didn't leave.

Ciardelli and Roelofsen (2015) use the fact that predicates such as *wonder* entail ignorance to explain their rogativity.

For a responsive predicate V, an important question is how the meanings of their declarative-embedding use *x Vs that S* and their interrogative-embedding use *x Vs Q* are related. V is *Q-to-P veridical* if *x Vs Q* entails *x Vs that p*, where p is the true answer to Q. For instance, if *Alice knows which player won* and in fact *Bob won*, then it follows that *Alice knows that Bob won*.

V is *Q-to-P distributive* if *x Vs Q* entails *x Vs that p* for some p that is a potential answer to Q. For instance, if *Alice is certain (about) which player won*, then there must be some player *y* such that *Alice is certain that y won*. Note that Q-to-P veridical predicates must be Q-to-P distributive but not vice versa. For instance, *be certain* is Q-to-P distributive but not Q-to-P veridical.

Finally, V is *P-to-Q distributive* if *x Vs that p*, where p is a possible answer to a question Q, entails *x Vs Q*. For instance, *Alice is certain that Bob won* entails *Alice is certain (about) which player won*.

Spector and Egré (2015) propose that responsive predicates are all Q-to-P distributive, whereas Roelofsen and Uegaki (2020) propose, instead, that they are all P-to-Q distributive (see Section 5.2 for further discussion).

Before concluding this section, we note that the semantic properties described here can in principle be applied to predicates in any language. Similarly,

Class	Verbs
Communication	*accept*, *announce*, *argue*, *assert*, *claim*, *complain*, *deny*, *explain*, *inform*, *tell*, *whisper*, *write*
Doxastic	*agree*, *assume*, *believe*, *(be) certain*, *(be) convinced*, *doubt*, *expect*, *forget*, *know*, *learn*, *prove*, *(be) right*, *suspect*, *think*, *(be) unaware*, *(be) wrong*
Perception	*see*
Directive	*decide*, *demand*, *order*, *propose*
Emotive	*fear*, *(be) happy*, *hope*, *pray*, *prefer*, *regret*, *(be) surprised*, *want*, *(be) worried*
Inquisitive	*ask*, *(be) curious*, *inquire*, *investigate*, *wonder*
Relevance	*care*

Table 1: Verb classes and verbs included in the database

the empirical generalizations proposed in the literature make crosslinguistic predictions, even though they were typically motivated by data from English or a few well-studied languages. Testing such predictions in a wider range of languages is crucial to assess the validity of existing proposals.

3 Design of the Crosslinguistic Database

Our database is designed to assess empirically the kinds of crosslinguistic generalizations described in Section 2. Furthermore, it will possibly enable discovery of previously unnoticed correlations, in particular ones involving interactions between multiple properties. In this section, we introduce the general design of the database. We will also briefly discuss practical aspects of data collection.

3.1 The properties and sample predicates

The database contains information about ∼50 clause-embedding predicates in each language. Each predicate is annotated with respect to ∼15 semantic properties and ∼12 combinatorial properties. The numbers are approximate because in some languages there are multiple attitude predicates corresponding to just one predicate in another language, and certain languages make more clause type distinctions than others. In the English database there are 48 predicates, listed in Table 1. The semantic and combinatorial properties considered are listed in Table 2.

Semantic properties The semantic properties are annotated based on inferential diagnostics and acceptability judgments. For example, the property of Veridicality is annotated based on the following inferential test:
Veridicality test Consider:

(16) Ann Vs that it is raining.

Does this sentence always imply that it is raining? If not, does it always imply that it is not raining?

Marking instructions

- If you answered *yes* to the first question, please mark V as **always veridical**.
- If you answered *yes* to the second question, please mark V as **always anti-veridical**.
- If you answered *no* to both questions, but you feel that the sentence typically implies that it is raining, please mark V as **typically veridical**.
- Similarly, if you answered *no* to both questions, but you feel that the sentence typically implies that it is not raining, please mark V as **typically anti-veridical**.
- Otherwise, please mark V as **neither**.

An example of a semantic property annotated based on acceptability judgments rather than an inferential test is Gradability. Specifically, this property is annotated based on the acceptability of sentences like (9) and (10) above. For some predicates, the judgments can be unclear, in which case the option *undecided* is used.

Combinatorial properties Combinatorial properties are annotated based on whether the predicate can take specific clause types. The relevant clause types for English are listed in the last row of Table 2, and those for other languages contain corresponding information with respect to syntactic equivalents of these clause types. Some languages involve further clause-type distinctions. For example, the data for Catalan, French, Italian, and Spanish involve an indicative/subjunctive mood distinction and the data for Greek, Hungarian, Japanese, and Turkish involve complementizer and other clause-type distinctions.

Predicate sample The sample of 48 English predicates in Table 1 has been selected from various classes of predicates investigated in the theoretical literature and cover a wide range of combinations of semantic and combinatorial properties. For languages other than English, we initially ask consul-

Semantic properties	**Response options**
Veridicality†	veridical, anti-veridical, neither
Conjunction with negation of the complement	contradictory, redundant, neither
Conjunction with the complement	contradictory, redundant, neither
Complement projection/reversal through negation†	projective, reversive, neither
Neg-raising†	neg-raising, non-neg-raising
Subject's {likelihood / unlikelihood / equal likelihood} estimation towards complement	always implies, typically implies, compatible, incompatible
Subject's {certainty / counter-certainty / uncertainty} towards complement	always implies, typically implies, compatible, incompatible
Subject's {preference / opposition / indifference} towards complement	always implies, typically implies, compatible, incompatible
Focus sensitivity	focus-sensitive, non-focus-sensitive
Grammatical gradability with declaratives	gradable, non-gradable, undecided
Belief/ignorance implications w.r.t. interrogatives†	belief-, ignorance-implying, neutral
Grammatical gradability w.r.t. interrogatives	gradable, non-gradable, undecided
Q-to-P veridicality†	veridical, anti-veridical, neither
Q-to-P distributivity†	distributive, non-distributive
P-to-Q distributivity†	distributive, non-distributive
Combinatorial properties	**Response options**
Finite & non-finite declaratives; Finite & non-finite interrogatives (polar, alternative, *which*, *who/what*); Concealed questions; Intransitive use	acceptable, unacceptable, degraded (from ? to ???), [preposition/particle/etc.] required, undecided

Table 2: All of the properties included in the questionnaire, where † indicates properties for which a graded response was elicited, e.g., *typically* or *always* veridical.

tants to provide direct translations of the English predicates, to the extent that such translations exist. If a direct translation does not exist, consultants are encouraged to consider predicates that are similar in meaning to the original English predicate and comment on the extent to which they are comparable in the text document. We further discuss this translation-based method of sampling predicates across languages in the Limitations section.

3.2 Annotation

The annotation instructions are collated in a questionnaire format, with accompanying predicate-specific notes that discuss certain confounding factors that need to be controlled for on a predicate-specific basis. Both documents are accessible at https://osf.io/vd8mg/. Data were annotated by native speakers with a background in linguistics (at least an undergraduate degree). Each consultant spent 60 to 100 hours (distributed over 3 to 4 months) on completing their dataset, and consulted regularly with at least one of the authors during this process in order to clarify difficult judgments or resolve possible complications. Annotation was performed across all properties by a single consultant for each language. This design allows a within-subject testing of possible correlations between different properties. At the same time, since the format of our database tracks consultant IDs for each data point, our design of the database does not preclude addition of data based on annotation from other speakers in the future.

4 Practical Details about the Database

4.1 Format

The database is located at https://wuegaki.ppls.ed.ac.uk/mecore/mecore-databases/. Each language has its own folder containing the following documents: (i) a README file containing basic information about the language, the list of language-specific semantic and combinatorial properties, and the data collection process, (ii) a table (a CSV file) in wide format, where each row corresponds to a predicate and each column to a combinatorial or semantic property (see Table 3), (iii) the corresponding text document containing the linguistic examples used in determining the properties and relevant discussions.

The tables are in wide format so that it is easy to visually inspect them, which is useful when one is casually exploring the database. However, as discussed in the previous section, different languages have different sets of properties. For instance, Mandarin Chinese has two negation markers which can lead to different interpretations. As a result, each negation-related property corresponds to two columns in the Mandarin table but only one in other languages. Therefore it is impossible to directly aggregate tables in wide format from different languages. They need to be converted to long format tables first to be appended to one another. In this case, the long format includes an additional column called NegationMarker. For a negation-related property, the value of this column is the negation marker under consideration. If the property does not involve negation, the value is NONE.

Other language-specific distinctions are, e.g., *mood* (for Romance languages) and *complementizer* (for Japanese, Greek, Turkish and Hungarian). Information about such distinctions is stated in the README file for the relevant languages.

4.2 Snapshot

At the time of writing, the database contains 15 languages: Catalan, Dutch, English, French, German, Greek, Hebrew, Hindi, Italian, Japanese, Kîîtharaka, Mandarin, Spanish, Swedish and Turkish. For two of the languages, German and Polish, a detailed report on the process of creating a first version of the dataset (superseded by the version that we currently release) is available as Master's theses (Naehrlich, 2022; Klochowicz, 2022).

4.3 Contributing to the database

Researchers are welcome to contribute to the database. The simplest way is to use our questionnaire and predicate-specific notes to collect data on (possibly a subset of) the translations of the 48 English predicates in the current database (as part, for example, of a student's research project or internship). While the questionnaire and predicate-specific notes are designed for trained linguists as consultants, they can be adapted to a fieldwork setup for consultants with no training in linguistics.

We emphasize that the additional data need not be about a new language. Due to intra-language variation, it is also valuable to have judgments from multiple speakers of the same language.

One can also apply the questionnaire to predicates beyond the ones in the current database. In this case, contributors are encouraged to provide predicate-specific notes on the additional predicates to faciliate future crosslinguistic investigations.

5 Two case studies

We discuss two case studies using our database. Although strong conclusions cannot be drawn from the limited sample we currently have, as a proof of concept, they show how our database informs debates about crosslinguistic generalizations.

5.1 Emotive factives and whether questions

The first case study concerns the relation between combinatorial and semantic properties. Recall that Karttunen (1977) observes that emotive factives cannot take *whether* questions (11). We aim to evaluate this generalization crosslinguistically.

In line with how this class of predicates is generally thought of in the literature, we adopt the following criteria. A predicate is emotive factive if it is (i) typically or always veridical, (ii) typically or always projective, (iii) focus sensitive, (iv) gradable, and (v) it entails that the subject believes the complement—which we operationalize as implying that, according to the subject, the complement is more likely than its negation (e.g., Egré, 2008; Villalta, 2008; Romero, 2015). In our database for English, 4 predicates satisfy these criteria: *be happy*, *be surprised*, *regret*, and *care*. The first three are indeed canonical examples of emotive factives, and intuitively *care* is an emotive predicate and shares the semantic properties listed above, e.g., it is typically veridical and projective (17).

Predicate	**English translation**	**Veridicality/ Anti-veridicality**	...	**Finite declaratives**	**Finite *which* interrogatives**	...
vergeten	*forget*	always veridical	...	acceptable	acceptable	...
ongelijk hebben	*be wrong*	always anti-veridical	...	acceptable	acceptable	...
geloven	*believe*	neither	...	acceptable	unacceptable	...
zich afvragen	*wonder*	NA	...	unacceptable	acceptable	...
...	...	...	...	...	...	

Table 3: Part of the Dutch predicate table in wide format

(17) Alice cares/does not care that Bob won
$\rightsquigarrow$ Bob won

But, while the first three predicates indeed cannot take *whether* questions, *care* can (18), which makes it a potential counterexample to the generalization.

(18) Ann cares whether Bob or Charles won.

However, note that the meaning of (18) is different from what one would expect when an emotive factive predicate takes a question complement. For instance, (19) entails that there is an answer p to the embedded question such that the subject is surprised that p. That is, canonical emotive factives are Q-to-P distributive. In contrast, (18) does not have such an entailment (20). This is because *Alice cares that x won* entails that Alice believes that x won, but (18) can be true even if Ann does not have a belief about who won at all (Elliott et al., 2017).

(19) Alice is surprised (about) who won.
$\Rightarrow \exists x$. Alice is surprised that x won.

(20) Alice cares whether Bob or Charles won.
$\not\Rightarrow \exists x$. Alice cares that x won.

This observation allows us to refine the original generalization by Karttunen. A predicate cannot take *whether* questions if it is an emotive factive (as operationalized above) *and Q-to-P distributive*.

This refined generalization is highly robust crosslinguistically. When the counterparts of canonical emotive factives *be happy*, *be surprised* and *regret* take *whether* questions, the results are consistently judged unacceptable or highly marked. The counterparts of *care* consistently lack Q-to-P distributivity and can take *whether* questions.

It is worth looking into Kîîtharaka *rigara*, offered by our consultant as the translation of English *be surprised*, in some more detail. This predicate has two senses. When it takes a declarative complement, it is translated as *be surprised*. When it takes a *wh*-question, it can mean that there is an answer p to the question such that the subject is surprised that p. In this respect *rigara* is an emotive factive predicate just like *be surprised*. However, when *rigara* takes a question, it can also be translated as *wonder*. Crucially, although *rigara* can take *whether* questions, it can only be translated as *wonder* in such cases. In particular, *Bill rigara whether Mary left* means that Bill wonders whether Mary left, and crucially, it does not entail that either *Bill rigara that Mary left* or *Bill rigara that Mary did not leave* must be true. Thus, when *rigara* takes *whether*-complements, it is not Q-to-P distributive.

There are further cases of predicates that satisfy the criteria of emotive factives while lacking Q-to-P distributivity. For instance, Swedish *vara orolig över*, unlike its English counterpart *be worried*, is always veridical (therefore a more accurate translation would be *it worries x that*). It is not Q-to-P distributive and can take *whether* questions.

This case study lends support for a modified version of Karttunen's generalization: if a predicate is an emotive factive and Q-to-P distributive, it is incompatible with *whether* questions. It also highlights the utility of our database in the investigation of crosslinguistic correlations between semantic and combinatorial properties of attitude predicates. Without the type of data available in the current database, it would be difficult to empirically assess the relevance of Q-to-P distributivity to Karttunen's original observation in a crosslinguistic context.

5.2 P-to-Q distributivity

The second case study concerns the crosslinguistic validity of the generalization that all responsive attitude predicates satisfy P(roposition)-to-Q(uestion) distributivity (Roelofsen and Uegaki, 2020). To illustrate, from (21a), we may infer (21b) and (21c), where the embedded declarative in (21a) ("P") is one of the possible answers to the embedded interrogatives in (21b) and (21c) ("Q").

(21) a. Al knows/cares that Jo is Dutch.

b. Al knows/cares whether Jo is Dutch.
c. Al knows/cares where Jo is from.

Roelofsen and Uegaki identify three classes of potential counter-examples to P-to-Q distributivity, without drawing definite conclusions. First, some predicates are non-veridical with declarative complements, but veridical with interrogative complements (Q-to-P veridical). A prototypical example is *tell* (Karttunen, 1977). Examples like (22) do not entail the embedded clause, suggesting non-veridicality with declaratives, but the conjunction of (23a) and (23b) is often judged to entail (23c), suggesting that *tell* might be Q-to-P veridical. (Note, however, that the predicate is not considered Q-to-P veridical by everyone—see Tsohatzidis 1993; Holton 1997; Spector and Egré 2015, a.o.)

(22) Al told Jo that Sue won. $\not\rightarrow$ Sue won.

(23) a. Al told Jo which runner won.
b. Zoe won.
c. $\therefore$ Al told Jo that Zoe won.

If this is correct, *tell* cannot be P-to-Q distributive as (23a) does not follow from (22) in situations where Sue did not win.

Second, there are predicates similar to Kîîtharaka *rigara*, which alternate between *surprise*- and *wonder*-like interpretations. Third, predicates like English *explain* alternate between 'explanans' ('that which explains') and 'explanandum' ('that which is explained') interpretations (Pietroski, 2000; Elliott, 2017; Bondarenko, 2021). What unifies these predicates is that they have qualitatively different meanings across declarative and interrogative embedding.

Our sample corroborates that there is a general tendency for responsive predicates to be P-to-Q distributive, but also that the identified classes of counter-examples are crosslinguistically attested: Speakers of some languages in our sample judged that every predicate obeys the property (Dutch, English, Greek, Kîîtharaka and Mandarin); for others there was a variable, but small number of exceptions (Catalan, Italian, Hebrew, Hindi, Japanese, Polish, Spanish, Swedish and Turkish). Among these exceptions, we first find communicative and doxastic predicates that are non-veridical in declarative, but veridical in interrogative embedding. Some examples include Turkish *bildir*- 'inform' and Polish *wyjaśniać* 'explain' (see also Özyıldız 2019, Bondarenko 2020, Jeong 2020).

Second, we find predicates like Swedish *tänka på*, which roughly translates sentences of the form '*think about* the fact that x won' with declaratives, and ones like '*think about* which runner won' with questions. Importantly, the former is reported to entail the belief that x won, and the latter, ignorance about which runner won. As belief is incompatible with ignorance in this situation, P-to-Q distributivity fails. One way of identifying this kind of predicate involves comparing their values for likelihood and certainty implications with the one for belief/ignorance implications w.r.t. interrogatives. Mismatching values here will point towards a shift in meaning across declarative and interrogative embedding. Among this class of predicates, we also find the counterparts of 'think' in Catalan, Spanish and Turkish, *surprise*/*wonder*-type predicates in Japanese, Spanish and Swedish, and a third set of predicates instantiated by Turkish communicatives *de*-, *yaz*- and *fısılda*- ('say,' 'write,' and 'whisper'). With declaratives, these Turkish predicates imply that their subject linguistically produced the declarative (e.g., *Al said: "Jo won."*), but with interrogatives, that the subject produced the interrogative (e.g., *Al said: "Which runner won?"*). Hence, P-to-Q distributivity fails for them as well.

This case study confirms a general tendency for predicates to be P-to-Q distributive, but also reveals variation, both within and across languages. Its results are consistent with debates in the literature, e.g., regarding the properties of *tell* and *explain*. Some exceptions to the general tendency are better understood (e.g., veridicality alternating predicates) than others (e.g., the class of *surprise*/*wonder* predicates). This, in turn, paves the way for new empirical and theoretical research.

6 Conclusions

We have presented our crosslinguistic database for combinatorial and semantic properties of attitude predicates. As our case studies show, the database enables assessment of two types of crosslinguistic generalisations: one concerning correlations between semantic and combinatorial properties of attitude predicates and the other concerning general semantic constraints on attitude predicates. The database complements existing resources due to three features: (i) crosslinguistic data; (ii) enabling within-subject comparison across properties, and (iii) accompanying text documents that allow fine-grained qualitative assessment of data.

Limitations

The data collection process was time-intensive. Each language required a total of 60 to 100 hours of work by a native speaker with a background in linguistics, typically over the course of 3 to 4 months with regular consultation sessions with one of the authors of the present paper. Because of this, the current database for the most part only features introspective judgments coming from a single speaker per language (although occasionally informants would consult other native speakers and/or corpora when they were uncertain). While this is a good place to start, the database is not yet equipped to address issues pertaining to within and across speaker variability. For the same reason, we have had to limit the number of attitude predicates that we tested to a manageable number. While we believe that our sample covers much of the logical space of possibilities for the meaning of attitude predicates, the number of predicates remains small (especially in comparison with the *Mega* datasets). The fact that our initial survey is translation based makes it also possible that certain predicates of interest in the target languages were missed.

The languages that were included in the database are typologically diverse, but they do not cover all known language families and are currently restricted to the spoken modality. There is nothing, however, that prevents the inclusion of other languages, including sign languages, and we are hopeful that our database will expand in these directions.

Regarding the tests that we have used to elicit semantic and combinatorial properties, while some are relatively easy to transpose into other languages (e.g., conjunction with the (negation of the) complement), others are harder, and their results might be less reliable. For example, the question about neg-raising is currently eliciting an inference which might be driven by factors other than the predicate actually being neg-raising. An alternative, arguably more reliable test would make use of strict Negative Polarity Items (NPIs), but identifying NPIs in a given language requires detailed knowledge of the language and may only be possible for languages the researcher is familiar with or has conducted extensive fieldwork on.

Regarding the consistency of the data, there are some values that some of the properties cannot jointly take. For example, a predicate cannot at the same time be less than always veridical, always Q-to-P veridical and always P-to-Q distributive. However, this particular combination of values has been observed for certain predicates in our sample. We have attempted to minimize such inconsistencies by conducting follow-up interviews with our speakers, and making sure that they assessed the predicates in all relevant contexts of use. Rather than being a problem, however, this can be seen as a feature of our method, as it allows us to identify strong tendencies in how speakers interpret attitude ascriptions.

Finally, we note that, while the tables that are included in the database are machine readable, the supporting text documents are currently not. They have to be processed directly by the interested researcher. We are working towards making the text documents machine readable as well.

Ethics Statement

The data collection process, described in Section 3.2, and the projected or otherwise possible applications of our data have been approved by the ethics committees of the institutions funding and hosting this research, and they conform to the ACL Ethics Policy. The language consultants who have provided their introspective judgments have been compensated in accordance with the laws in place in the UK and in Germany. The database only contains anonymized consultant IDs, and our consultants have been offered the option of remaining anonymous, or of being authors on or being acknowledged by name in relevant publications—the latter two options being relevant for the academic recognition of some of our consultants, who are also professional linguists.

Acknowledgements

We thank our consultants: Aayush Bagchi, Sjaak de Wit, Rebecka Elm, Clara Giralt, Nori Hayashi, Patrick Kanampiu, Tomasz Klochowicz, Sarah Molina Raith, Flavia Naehrlich, Aviv Schoenfeld, Anastasis Stefas, Yingyu Su, Ilaria Venagli, Caitlin Wilson, and one anonymous consultant. We thank two additional consultants, Nana Kwame and Eszter Ótott-Kóvacs, who we are still working with on collecting data for two additional languages. The results will be added to the database once the data collection is complete.

We also thank Kajsa Djärv, Jenny Doetjes, Despina Oikonomou, Jakub Szymanik and Malte Zimmermann for helping with data collection and discussion of methodological issues at various

stages of the project.

This paper is a part of the project 'MECORE: A cross-linguistic investigation of meaning-driven combinatorial restrictions in clausal embedding', supported by the AHRC-DFG Collaborative Grant in Humanities (AHRC reference: AH/V002716/1; DFG reference: RO 4247/5-1).

References

Maayan Abenina-Adar. 2019. Surprising. In *Proceedings of the 36th West Coast Conference on Formal Linguistics*, pages 41–47. Somerville, MA: Cascadilla Proceedings Project.

Hannah Youngeun An and Aaron Steven White. 2020. The Lexical and Grammatical Sources of Neg-Raising Inferences. In *Proceedings of the Society for Computation in Linguistics 3*, pages 220–233.

Tatiana Bondarenko. 2020. Factivity from pre-existence: Evidence from Barguzin Buryat. *Glossa: a journal of general linguistics*, 5:1–35.

Tatiana Bondarenko. 2021. Two paths to explain: clausal embedding with verbs of speech. Manuscript, MIT.

Emmanuel Chemla, Vincent Homer, and Daniel Rothschild. 2011. Modularity and intuitions in formal semantics: the case of polarity items. *Linguistics and Philosophy*, 34:537–570.

Ivano Ciardelli and Floris Roelofsen. 2015. Inquisitive dynamic epistemic logic. *Synthese*, 192(6):1643–1687.

Paul Egré. 2008. Question-embedding and factivity. In F. Lihoreau, editor, *Grazer Philosophische Studien 77*, pages 85–125.

Patrick Elliott. 2017. *Elements of clausal embedding*. Ph.D. thesis, UCL.

Patrick D. Elliott, Nathan Klinedinst, Yasutada Sudo, and Wataru Uegaki. 2017. Predicates of relevance and theories of question embedding. *Journal of Semantics*, 34(3):547–554.

Gottlob Frege. 1898 [1948]. Sense and reference. *The philosophical review*, 57(3):209–230.

John Hawthorne, Daniel Rothschild, and Levi Spectre. 2016. Belief is weak. *Philosophical Studies*, 173(5):1393–1404.

Richard Holton. 1997. Some telling examples: A reply to tsohatzidis. *Journal of pragmatics*, 28(5):625–628.

Sunwoo Jeong. 2020. Prosodically-conditioned Factive Inferences in Korean: An Experimental Study. In *Semantics and Linguistic Theory 30 (SALT 30)*.

Benjamin Kane, William Gantt, and Aaron Steven White. 2021. Intensional gaps: Relating doxasticity, bouleticity, veridicality, factivity, and neg-raising. In *Proceedings of Semantics and Linguistic Theory 31*.

Lauri Karttunen. 1977. Syntax and semantics of questions. *Linguistics and Philosophy*, 1:3–44.

Tomasz Klochowicz. 2022. Investigation semantic and selectional properties of clause-embedding predicates in Polish. MSc in Logic thesis, Universiteit van Amsterdam.

Mora Maldonado, Jennifer Culbertson, and Wataru Uegaki. 2022. Learnability and constraints on the semantics of clause-embedding predicates. In *Proceedings of the Annual Meeting of the Cognitive Science Society*.

Clemens Mayr. 2019. Triviality and interrogative embedding: Context sensitivity, factivity and neg-raising. *Natural Language Semantics*, 27.

Ellise Moon and Aaron Steven White. 2020. The Source of Nonfinite Temporal Interpretation. In *Proceedings of the 50th Annual Meeting of the North East Linguistic Society*, pages 11–24. Amherst, MA: GLSA Publications.

Flavia Naehrlich. 2022. Semantic and selectional properties of clause-embedding predicates in German. MSc in Logic thesis, Universiteit van Amsterdam.

Deniz Özyıldız. 2019. Potential answer readings expected, missing. In *Proceedings of Tu+ 4*.

Paul M. Pietroski. 2000. On explaining that. *The Journal of philosophy*, 97(12):655–662.

Floris Roelofsen and Wataru Uegaki. 2020. Searching for a universal constraint on the denotations of clause-embedding predicates. In *Proceedings of Semantics and Linguistic Therory 30*.

Maribel Romero. 2015. Surprise-predicates, strong exhaustivity and alternative questions. In *Semantics and Linguistic Theory*, volume 25, pages 225–245.

Kjell Johan Saebø. 2007. A whether forecast. In B.D. ten Cate and H.W. Zeevat, editors, *TbiLLC 2005*, pages 189–199. Springer-Verlag Berlin Heidelberg.

Benjamin Spector and Paul Egré. 2015. A uniform semantics for embedded interrogatives: an answer, not necessarily the answer. *Synthèse*, 192(6):1729–1784.

Shane Steinert-Threlkeld. 2019. An Explanation of the Veridical Uniformity Universal. *Journal of Semantics*, 37(1):129–144.

Barbara Stiebels, Thomas McFadden, Kerstin Schwabe, Torgrim Solstad, Elisa Kellner, Livia Sommer, and Katarzyna Stoltmann. 2018. Zas database of clause-embedding predicates, release 1.0. OWIDplus, hg. v. Institut für Deutsche Sprache, Mannheim, http://www.owid.de/plus/zasembed.

Nadine Theiler, Floris Roelofsen, and Maria Aloni. 2019. Picky predicates: Why believe doesn't like interrogative complements, and other puzzles. *Natural Language Semantics*, 27(2):95–134.

Judith Tonhauser, David I Beaver, and Judith Degen. 2018. How Projective is Projective Content? Gradience in Projectivity and At-issueness. *Journal of Semantics*, 35(3):495–542.

Savas L Tsohatzidis. 1993. Speaking of truth-telling: The view from wh-complements. *Journal of pragmatics*, 19(3):271–279.

Wataru Uegaki and Yasutada Sudo. 2019. The *hope-wh puzzle. *Natural Language Semantics*, 27:323–356.

Elisabeth Villalta. 2008. Mood and gradability: An investigation of the subjunctive mood in Spanish. *Linguistics and Philosophy*, 31(4):467–522.

Aaron Steven White and Kyle Rawlins. 2016. A computational model of S-selection. In *Proceedings of Semantics and Linguistic Theory 26*, pages 641–663.

Aaron Steven White and Kyle Rawlins. 2018. The role of veridicality and factivity in clause selection. In *Proceedings of the 48th Annual Meeting of the North East Linguistic Society*, volume 3, pages 221–234. Amherst, MA: GLSA Publications.

Richard Zuber. 1982. Semantic restrictions on certain complementizers. In *Proceedings of the 12th International Congress of Linguists, Tokyo*, pages 434–436.