

©Copyright 2022

Nathaniel Imel

Modals in Natural Language Optimize the Simplicity/Informativeness Trade-Off

Nathaniel Imel

A thesis

submitted in partial fulfillment of the
requirements for the degree of

Master of Science

University of Washington

2022

Committee:

Shane Steinert-Threlkeld, Chair

Jakub Szymanik

Program Authorized to Offer Degree:

Computational Linguistics

University of Washington

Abstract

Modals in Natural Language Optimize the Simplicity/Informativeness Trade-Off

Nathaniel Imel

Chair of the Supervisory Committee:
Assistant Professor Shane Steinert-Threlkeld
Department of Linguistics

The meanings expressed by the world’s languages have been argued to support efficient communication. Evidence for this hypothesis has drawn on cross-linguistic analyses of vocabulary in semantic domains of both content words (e.g. kinship terms as shown in Kemp and Regier [2012]; color terms as in Zaslavsky et al. [2018]) and function words (e.g. quantifiers as in Steinert-Threlkeld [2021]; indefinite pronouns in Denic et al. [2020]) approaching the hypothesis concretely in terms of a trade-off between simplicity and informativeness. I apply the analysis to *modals* (e.g. *can*, *ought*, *might*). Two proposed universals in this domain from Nauze [2008] and Vander Klok [2013] are used for generating many artificial languages with varying degrees of quasi-naturalness as a proxy for natural data. A computational experiment shows that the Vander Klok universal picks out the optimal solutions to the simplicity/informativeness trade-off, suggesting that efficient communication is a leading explanation for constraints on modal semantic variation.

TABLE OF CONTENTS

	Page
List of Figures	iii
Chapter 1: Introduction	1
1.1 Research questions	2
1.2 Structure of the thesis	4
Chapter 2: Background and Literature Review	5
2.1 The simplicity/informativeness trade-off	5
2.2 Modal semantics	6
Chapter 3: Methods	11
3.1 Measuring simplicity and informativeness	11
Chapter 4: Results	16
4.1 Computational experiment	16
Chapter 5: Discussion	20
5.1 Typology	20
5.2 Measures of complexity and informativeness	21
5.3 Alternative explanations	22
Chapter 6: Conclusion	24
Bibliography	25
Appendix A: Additional parameter results	30
A.1 Additional semantic spaces	30
A.2 Utility functions: indicator vs. half-credit	31

A.3 Language of Thought: negation	33
A.4 Discussion	33
Appendix B: Degree SAV vs. Pareto optimality	34
Appendix C: Explicit typologies	36
Appendix D: Boolean algebra (language of thought for modals)	37
Appendix E: Shortest formula heuristic search algorithm	40

LIST OF FIGURES

Figure Number	Page
4.1 The complexity/communicative cost trade-off for modal languages. The black line is the Pareto frontier (optimal solutions) and robustly consists of the Domain-Level Single-Axis of Ambiguity typology, suggesting that the known semantic variation for modals is explained by general pressures for efficient communication.	18
A.1 The complexity/communicative cost trade-off for modal languages with 4 and 5 flavors.	30
A.2 The complexity/communicative cost trade-off for a graded (main results, left) versus an all-or-nothing utility function.	31
A.3 The complexity/communicative cost trade-off with (main results, left) and without negation as a primitive in the Language of Thought.	33
B.1 Degree of Single Axis of Variability Universal satisfaction vs. Pareto optimality	34

ACKNOWLEDGMENTS

I wish to thank Shane Steinert-Threlkeld for his advising and mentoring, and the wider UW Linguistics department and Jakub Szymanik for valuable feedback. I am also grateful to my family and friends for their robust support throughout my time in the CLMS program.

DEDICATION

To my best friend, Mika Aubin, and my inspiration, David Imel

Chapter 1

INTRODUCTION

Humans speak only a subset of the mathematically possible languages. One goal of linguistics consists in accurately describing this subset —identifying the actual human languages. In the domain of semantics, we can ask: which meanings do we happen to use, and why?

While there is great typological variation in language meanings across the world, there are also constraints on this variation —*semantic universals*. There is a hypothesis going back to Zipf [1949] that the meanings observed across languages are shaped by a pressure for efficient communication. This thesis follows a line of recent work approaching this hypothesis concretely in terms of a trade-off between cognitive simplicity and communicative informativeness [Kemp and Regier, 2012], in the domain of *modal semantics*.

A modal is typically considered to be a semantic operator that qualifies the truth of an expression. In English, this can be expressed by auxiliaries including *might, may, must, could* and adverbs *probably, necessarily* among a variety of other constructions. Cross-linguistically, these meanings are expressed by diverse lexical categories and strategies.¹ Mood and modality, together with tense and aspect, exemplify the property Charles Hockett [1960] named displacement: the phenomenon of talking about beyond the actual, here and now. This work takes a computational cognitive approach to the study of modality: though many modal semantic systems can be constructed, which ones optimize efficient communication? And are

¹Not every language has a small set of lexical items used to express every modal category: Urdu/Hindi expresses modality via a specific set of morpho-syntactic constructions [Bhatt et al., 2011]. Additionally, Tlingit has few grammatical strategies to express modality, and instead uses various pragmatic strategies to communicate about its modal categories [Cable, 2017].

these systems similar to those found in natural languages? To answer these questions, I explore two proposed modal semantic universals from Nauze [2008] and Vander Klok [2013] in a computational experiment measuring the optimal solutions to a trade-off between cognitive simplicity and communicative informativeness.

Measuring the simplicity of a language will reduce to abstractly modeling a language (fragment) and choosing a meaning space for its terms. The model of the meaning space is motivated by the cross-linguistic formal semantics literature, in which the variability of modals is described in terms of points lying on a map (table) with two axes: force and flavor. I intend the simplicity of a language to track the simplicity with which its terms carve up this meaning space. The informativeness of a language aims to measure of the likelihood of successful communication between two literal agents using the language’s possibly underspecified terms.

The results of this work can bear on the generality of the hypothesis that efficient communication shapes the world’s semantic systems. Meanwhile, if modals do not appear to optimize the simplicity/informativeness trade-off, more weight ought to be given to alternative causal explanations: perhaps modals cross-linguistically have the meanings they do because they are easiest to learn [Steinert-Threlkeld and Szymanik, 2019, 2020]. Meanwhile, this work does not provide a positive argument for the truth of the proposed semantic universals that it accepts as functional starting points for research. Should the universals be verified, however, the results in efficient communication here together with additional analysis might suggest that they are not fundamental properties of modals, but rather follow from general cognitive and cultural pressures.²

1.1 Research questions

The question of whether efficient communication shapes the modal typology can be broken down into several main subquestions:

²See Steinert-Threlkeld [2021] for similar discussion about quantifier universals.

1.1.1 *How should a modal system be modeled?*

Analyses of modal semantics typically involve displacement, quantification and distinctly human-related concerns including rules, goals, knowledge [Kratzer, 1981, 2012, Portner, 2009, Matthewson, 2016]. However, this work models the constraints on *underspecification* across these features without inquiring into their internal structure. The question that I am concerned with is: why do modals in each of the world’s languages show the specific semantic variability that they do? And absent a database comprising knowledge of what the actual modal inventories are, how can we identify the relative naturalness of modal systems?

1.1.2 *What makes a modal meaning more or less cognitively demanding?*

Cognitive simplicity of meaning can be modeled abstractly in different ways, including as a measure of the resources necessary to decode, produce, or represent the meaning. Is there a natural correspondence between modal meanings as proto-linguistic (e.g. Language of Thought) expressions and meanings as specific geometries in the meaning space? How might certain carvings of the meaning space count as more simple?

1.1.3 *What makes a modal meaning more or less informative?*

What kinds of assumptions are appropriate for measuring communicative success? This work does not incorporate evidence about the frequency of meaning points used by people using modals in conversation, nor does it incorporate the rich pragmatic reasoning that people use to reliably recover a speaker’s intended meaning. On the one hand, given that each feature is atomic and unanalyzed for additional structural complexity, a realistic prior over meaning points may or may not affect results of this trade-off. Notably, context sensitivity has central tool in explaining modal semantics since Kratzer [1977], so pragmatic modeling of information transfer might *prima facie* contribute significantly to a robust analysis.

1.1.4 *How can a sufficiently general set of languages be explored for optimality?*

My modal languages are generated as random sampled, fixed-sized lists from the set of possible modals. Combinatorial explosion presents a serious challenge: for a semantic space with 2 forces and 3 flavors and a fixed vocabulary size of 5 there are $\binom{64}{5} = 7624512$ languages. But vocabulary size should vary, and differently sized semantic spaces should be explored. To overcome the combinatorial explosion, I use a genetic algorithm (from Steinert-Threlkeld [2021]) to explore a sufficiently rich subset of possible modal languages. Still, the shape of the languages explored differs substantially from the those in related work. Depending on the adequacy of the metrics of simplicity and informativeness, this uniqueness of the trade-off space might indicate that there is something special about modals.

1.2 *Structure of the thesis*

Theoretical background and relevant literature is covered in Chapters 2, including an introduction to the simplicity/informativeness trade-off and a brief overview of modal semantics and its cross-linguistic variation. Methods are described in Chapter 3, including the measures of simplicity and informativeness. The main computational experiment and results are covered in Chapter 4. I discuss some modeling decisions and areas for future work in Chapter 5 and conclude in Chapter 6. Important details in the Appendices include the pseudocode for the heuristic algorithm that searches for the simplest description of a modal term and some results for additional experimental parameters.

Chapter 2

BACKGROUND AND LITERATURE REVIEW

2.1 *The simplicity/informativeness trade-off*

Consider a language in which there is only one expression. This one expression is used to communicate every possible meaning, no semantic distinctions are made, and so in one important sense the language minimizes cognitive complexity. This language is uninformative in the sense that using the expression does not help a listener narrow down what was meant.

On the other hand, a language with a distinct expression for every meaning is likely to support perfect recovery of linguistic signals in communicative exchanges. But because each expression must distinguish itself uniquely from every other expression, language simplicity is sacrificed. The efficient communication hypothesis predicts that these two pressures—simplicity and informativeness—trade off; the natural languages are the solutions to this multi-objective optimization problem.

The efficient communication hypothesis has been studied in a variety of semantic domains including kinship terms, color terms, number terms, container terms, quantifiers, boolean connectives and indefinite pronouns [Kemp and Regier, 2012, Regier, Kemp, and Kay, Xu and Regier, 2014, Xu, Regier, and Malt, 2016, Steinert-Threlkeld, 2021, Uegaki, 2020, Denic, Steinert-Threlkeld, and Szymanik, 2020]. I follow the methodology of this recent work in using a computational experiment to simulate the simplicity/informativeness trade-off. In particular, I generate a large sample of mathematically possible languages and identify those optimizing the trade-off.

Generally, the analysis rests on the following argument: there is evidence that efficient communication is a pressure shaping the meanings of a language if optimal languages contain the natural languages. The results can also bear causal explanatory burden with respect to

the existence of semantic universals (without precluding additional explanations of universals by other general pressures, such as learnability). In this work, I do not measure the natural languages because a comprehensive corpus describing the inventories of a sufficiently general sample of languages does not exist. Instead, I follow the methodology pursued in Steinert-Threlkeld [2019] and use two proposed semantic universals to generate languages with varying degrees of quasi-naturalness. Consequently, the results of this experiment illuminate how very general pressures for efficient communication can explain the variation in modal meanings observed cross-linguistically.

2.2 *Modal semantics*

Here I briefly review the semantics for modals, which has been extensively studied in linguistics and philosophy. A modal expresses some notion of ability, permission or other kind of possibility. A common starting point in the literature is Kratzer [1977, 1981]. This semantics and the family of analyses that follows it argue that a modal is roughly a quantifier over a certain contextually restricted *flavor* of possible worlds. Nothing in this paper, however, hinges on a specifically Kratzer-based semantics for modality; in particular, my analysis requires only that a modal be represented as a pair of two sets of binary features: the **force(s)** and **flavor(s)**. Throughout this thesis I describe these features in terms of quantification and contextual restriction for convenience.

Below are examples of the *deontic* and *epistemic* flavor readings for English *must*. On a Kratzer-based analysis, (1) is true just in case all the possible worlds where the relevant code of conduct hasn't been violated are worlds where all employee wash their hands. Likewise, the truth-conditions for (2) specify that you are tired in all the possibilities not ruled out by the speaker's evidence. The idea is that neither homonymy nor polysemy are responsible for the two readings: there is just one item *must*, and its lexical entry parametrizes modal flavor to an argument resolved in context.

- (1) Employees must wash their hands.

- (2) You must be tired.

Notably, modals need not be underspecified in general, and the underspecification might not be limited to the flavor axis. For example, English *might* is fully specified for existential quantification over epistemic possibility (and no other meanings). Meanwhile, modals in St’át’imcets (Lillooet Salish) have been argued to exhibit flexibility along quantificational force (and not modal flavor). Such *variable-force* modals lie in the opposite configuration to English modals in that they lexically encode exactly the axis of meaning that English modals underspecify. For example, Rullmann, Matthewson, and Davis [2008] argue that the deontic modal *ka* can express permission and obligation (flexibility between existential and universal readings). Their examples are reproduced in (3) and (4), and the discussed semantic variation across English and St’át’imcets is summarized by Table 1.

- (3) *cúy’-lhkacw* ***ká*** *t’u7 nas áts’x-en ta* *kwámts-sw-a*
 going.to-2SG.SUBJ DEON just go see-DIR DET husband-2SG.POSS-EXIS
 “You must go see your husband.”
- (4) *lán-lhkacw* ***ka*** *áts’x-en ti* *kwámts-sw-a,* *t’u7 áoz-as*
 already-2SG.SUBJ DEON see-DIR DET husband-2SG.POSS-EXIS but NEG-3CONJ
k-wá-su *xát’-min’* *k-wá-su* *nás-al’men,* *t’u7 áma*
 DET-IMP-2SG.POSS want-RED DET-IMP-2SG.POSS go-want just good
 “You may go see your husband, but if you don’t want to go, that’s okay.”

St’át’imcets <i>ka</i>			English <i>must</i>			
	epistemic	deontic	...	epistemic	deontic	...
weak		✓		weak		
strong		✓		strong	✓	✓

Table 2.1: Two kinds of modal semantic underspecification: *variable-force* and *variable-flavor*.

2.2.1 Modal semantic universals

A significant typological survey of modality is Nauze [2008]. Using a modified version of the semantic map from van der Auwera and Plungian [1998], Nauze presents modals as usually expressing up to two forces (necessity or possibility) and combinations of four flavors (epistemic, deontic, goal-oriented, or circumstantial). Based on a small but geographically varied sample of languages, Nauze offers the following tentative semantic universal, which I will refer to as the Single Axis of Variability (SAV) Universal.

Modal elements can only have more than one meaning along a unique axis of the semantic space: they either vary on the horizontal axis and thus are polyfunctional in the original sense of expressing different types of modality or they vary on the vertical axis and can express possibility and necessity, but they cannot vary on both axes. [Nauze, 2008, p. 222]

Vander Klok [2013] proposes an even more restrictive generalization based on work on Gitksan and Javanese which I refer to as the Domain-Level Single-Axis of Variability (DL-SAV).¹ Vander Klok predicts that a modal *system* will allow elements that are ambiguous along the force axis, or else ambiguous across the flavor axis—but there cannot exist both kinds of ambiguity present in the system as a whole. The DL-SAV proposal further predicts that the ambiguity is confined within a single modal domain, defined as either the epistemic or non-epistemic (root) flavors. The modal system in Table 2.2 wouldn't exist under the DL-SAV typology because item m_2 has referential ambiguity across the force axis while m_3 is ambiguous across the flavor axis, but both m_1, m_2 are in the root domain.

¹As reported in Matthewson [2016] .

	epistemic	deontic	circumstantial	goal-oriented
weak	m_1	m_2	m_3	m_3
strong	m_1	m_2	m_4	m_5

Table 2.2: A modal system predicted to exist under the SAV typology and not the DL-SAV [Matthewson, 2016].

A language like Washo may present a counterexample to SAV’s predictions (and so also to the DL-SAV universal). Bochnak [2015] observes that the modal verb *-eʔ*, appropriate in both epistemic and root contexts, is also compatible with necessity and possibility readings. As *-eʔ* thus exhibits underspecification along both axes of modal meaning, it is excluded by the two typologies (though by my measure of communicative efficiency, Washo might still be a language optimizing the trade-off, viz., by being more simple to represent). Moreover, Močnik and Abramovitz [2019] argue that the attitude verb *ivək* in Koryak is underspecified for force and flavor (doxastic and assertive) as a counterexample to the SAV.

2.2.2 Modals as sets of force, flavor pairs they can express

I assume that force and flavor are fundamentally properties of contexts of use.² This reflects current practice in semantic fieldwork as applied to modality [Matthewson, 2004, Bochnak and Matthewson, 2020, Vander Klok, 2021].³ For example, the modal questionnaire of Vander Klok 2021 consists exactly of discourse contexts designed to isolate a single force-flavor pair. For example, English *must* can express the pair (universal, deontic) because there is a reading for that pair under the context for (1). Here I identify a modal as the set of (force, flavor) pairs that it can express. A language is (generously) identified as a list of modals.

²See Database paper, universal paper for discussion.

³In addition to the particular studies already mentioned, see ?? for more examples of the application of these methods.

I adopt this level of generality because it avoids commitment on the exact formal semantics of these expressions, which is often still being debated. For example, we can say that a *variable force modal* is one that can express more than one pair with the same force. This is useful because there are two broad approaches to the semantics of such variable force modals: they actually encode existential quantification but lack a universal scalemate [Deal, 2011] or they encode universal quantification but rely on some mechanism of domain restriction [Rullmann et al., 2008, Bochnak, 2015, Močnik and Abramovitz, 2019]. On such analyses, the underlying semantics contains one specific quantifier; in the present setting, they will still be considered variable force since bare positive sentences are used in contexts with multiple forces. This allows us to take the two-dimensional feature map as an appropriate model of modal semantic variation. I now explain the measures of simplicity and informativeness in detail.

Chapter 3

METHODS

3.1 *Measuring simplicity and informativeness*

3.1.1 *Complexity in a Language of Thought*

I model the complexity of a modal meaning as the fewest number of atoms it takes to express its meanings in a Language of Thought (LOT) [Denic et al., 2020, Piantadosi et al., 2016, Fodor, 1975]. This representation language is a boolean algebra defined on the two-dimensional semantic map.¹ Atoms in the language correspond to columns in the case of flavors and rows in the case of forces. The rest of the primitives include the operators \wedge , \vee , and \neg for conjunction, disjunction and negation of features.²

I extend heuristics described in Feldman [2000] to find the shortest boolean formula for a modal.³ In particular, I take advantage of the distributive property to capture the difference between two modals with the same number of positive features that are nonetheless found in positions that are intuitively different in terms of in how difficult it is to locate them. Consider English *may* and the hypothetical modals *mought* and *notcirc* in Table 3.1:

¹Details in **Appendix D** describe some special axioms to encode the atoms as indices of the 2-D array representation of the feature map.

²See **Appendix A** for results without negation.

³Since the DNF is not in general the shortest formula available for a meaning, this makes the well-known Quine-McCluskey algorithm [Quine, 1952] inappropriate for this approximation.

Modal	Meaning representation				Shortest Formula in LOT	Complexity (# of atoms)
		<i>e</i>	<i>d</i>	<i>c</i>	<i>t</i>	
<i>may</i>	\exists \forall	✓	✓			$\exists \wedge (e \vee d)$ 3
<i>mought</i>	\exists \forall	✓				$(\exists \wedge e) \vee (\forall \wedge d)$ 4
<i>notcirc</i>	\exists \forall	✓	✓		✓	\bar{c} 1

Table 3.1: Measuring complexity for English *may* and two hypothetical modals. First column: meaning representation. Second column: shortest LOT formula. Third column: complexity measure.

One could write the meaning of *may* as $(\exists \vee e) \cup (\exists \wedge d)$ for its complexity, which uses as many atoms as *mought*. But its shorter form can be derived by factoring out shared axes. The distributive property of the boolean algebra validates this factoring out of terms as an inference. Including this property in the LOT is a way of formally encoding the intuitions about naturalness in the two proposed modal semantic universals: modals with meanings grouped along an axis are easier to represent than modals lacking such uniformity. In addition, treating the semantic map as a finite space allows intuitively simple meanings (such as *notcirc*) to be lower complexity.

3.1.2 Informativeness

The informativeness of a language L with a set of meaning points M is modeled after the idea of successful communication of signals between abstract senders and receivers [Skyrms, 2010, Kemp, Xu, and Regier, 2018, Steinert-Threlkeld, 2019]:

$$I(L) := \sum_{m \in M} \mathbb{P}(m) \sum_{i \in L} \mathbb{P}(i|m) \sum_{m' \in i} \mathbb{P}(m'|i) \cdot u(m', m)$$

$\mathbb{P}(i|m)$ is the probability a speaker selects a specific expression i to communicate a meaning m (a single point in the semantic space). $\mathbb{P}(m|i)$ is the probability that a listener correctly guesses the literal, intended meaning, given the expression heard. A prior over meaning points $\mathbb{P}(m)$ models how often agents need to communicate about specific meanings. The results now are for a uniform prior over meanings; I leave estimating a more realistic prior (perhaps using an annotated dataset of modalized expressions) to future work.

Related work measuring efficient communication of color terms and vague terms argues for quantifying the success of listener’s guesses on a scale, as opposed to an “all-or-nothing” affair [Jager, 2007, O’Connor, 2014]. With the structure available for the modal meaning space, it is possible to model some kind of graded value for listener guesses. In particular, I define a utility scoring function $u(m', m)$ which measures how close a listener’s guess $m' \in i$ is to the speaker’s intended meaning m . The guess m' is scaled (multiplied) by 0 if it shares no axes in common with the sender’s meaning m , by 0.5 if m' shares one axis with m , and by 1 if both axes are recovered (when $m' = m$).⁴

$$u(m', m) = 0.5 \cdot \mathbf{1}\{m_{\text{force}} = m'_{\text{force}}\} + 0.5 \cdot \mathbf{1}\{m_{\text{flavor}} = m'_{\text{flavor}}\}$$

Lastly, just as *complexity* is modeled as a decreasing function of the simplicity of a language, I define the *communicative cost* of a language as the complement of its informativeness: $Cost(L) = 1 - I(L)$. I use these terms interchangeably throughout this work.

⁴See **Appendix A** for results using an indicator utility function, which only rewards perfect recovery of meanings.

3.1.3 *Additional parameters*

An important question concerns what kinds of distinctions there might be cross-linguistically in modal semantic spaces. A separate question concerns how many modal expressions there are in each language. Both parameters can have a significant effect on the efficient communication of each language. Intuitively, the complexity of a language grows along with the possible number of semantic distinctions. Indeed, in order to keep vocabulary size modest, it is best for communicative purposes to have highly underspecified modals which jointly lower the risk of omitting points in the meaning space.

The level of refinement cross-linguistically in semantic spaces can be modeled in this setup by varying the size of the feature map (there is a sense, however, in which as more semantic spaces become more elaborate, a few forces and flavors are principal and the rest can be thought of as refinements of the initial divisions). In the main experiment, I also let the number of expressions per language vary to model different vocabulary sizes. Combinatorial explosion prevents exhaustive enumeration of the possible modal languages, even for modest values for these parameters.

3.1.4 *Exploring the optimization space with an evolutionary algorithm*

To combat the difficulty of finding the optimal modal languages, I follow Steinert-Threlkeld [2021] in using an evolutionary algorithm to estimate the *Pareto* frontier —the set of approximately optimal solutions to the simplicity/informativeness tradeoff—and measure these languages together with an additional sample.

This algorithm works to explore the space of possible languages by simulating their evolution. Roughly, it generates a ‘seed population’ of languages; applies a random choice of mutations to a subset of the population, and lets the most ‘fit’ languages survive after each generation. My implementation differs only by the use of one novel mutation, `remove_bit`, which chooses a modal uniformly at random from the language and replaces it with a modal that has one less expressible meaning. The intention of this mutation is

to help explore languages with higher informativeness. Full details of the rest of the algorithm, including pseudocode, can be found in the appendices of Steinert-Threlkeld [2021]. The code and full instructions for reproduction of my experimental results are available at <https://github.com/nathimel/modals>.

Chapter 4

RESULTS

4.1 Computational experiment

In order to evaluate the simplicity/informativeness trade-off for modals, we measure the distance of modal languages to the optimal solutions along the Pareto frontier. The main computational experiment generates the Pareto frontier and a large sample of modal languages.

Specifically, the experiment involves the following steps: (1) fix a semantic feature map from which to generate meanings; (2) find the shortest expression for each meaning; (3) estimate the Pareto frontier; (4) generate a random sample of languages with varying levels of satisfaction with the modal universals; and (5) measure each language's distance to the frontier.

In the beginning of the evolutionary algorithm, a seed population of 2000 artificial modal languages is randomly generated. At the end of each iteration a random choice of between 1 and 5 mutations is applied to the dominant languages. These dominant languages represent the subset of their generation best optimizing the simplicity/informativeness trade-off. After 200 generations, the dominant languages are the estimated Pareto frontier. Then, another sample of 2000 languages is generated in the same way as the genetic algorithm's seed population: for each vocabulary size from 1 to 10, we generate 200 languages with varying degrees of satisfaction with the SAV universal by drawing random bags of modal expressions from the set of possible expressions.

Finally, each language is measured for complexity, communicative cost and its minimum Euclidean distance to a point on the Pareto frontier. The SAV typology can be measured as a continuous value of languages (fraction of the vocabulary conforming to the criterion),

while DL-SAV must be measured as a categorical variable. The main experiment centers a single semantic space with 2 possible forces and 3 possible flavors (additional configurations are in **Appendix A**).

4.1.1 Main results

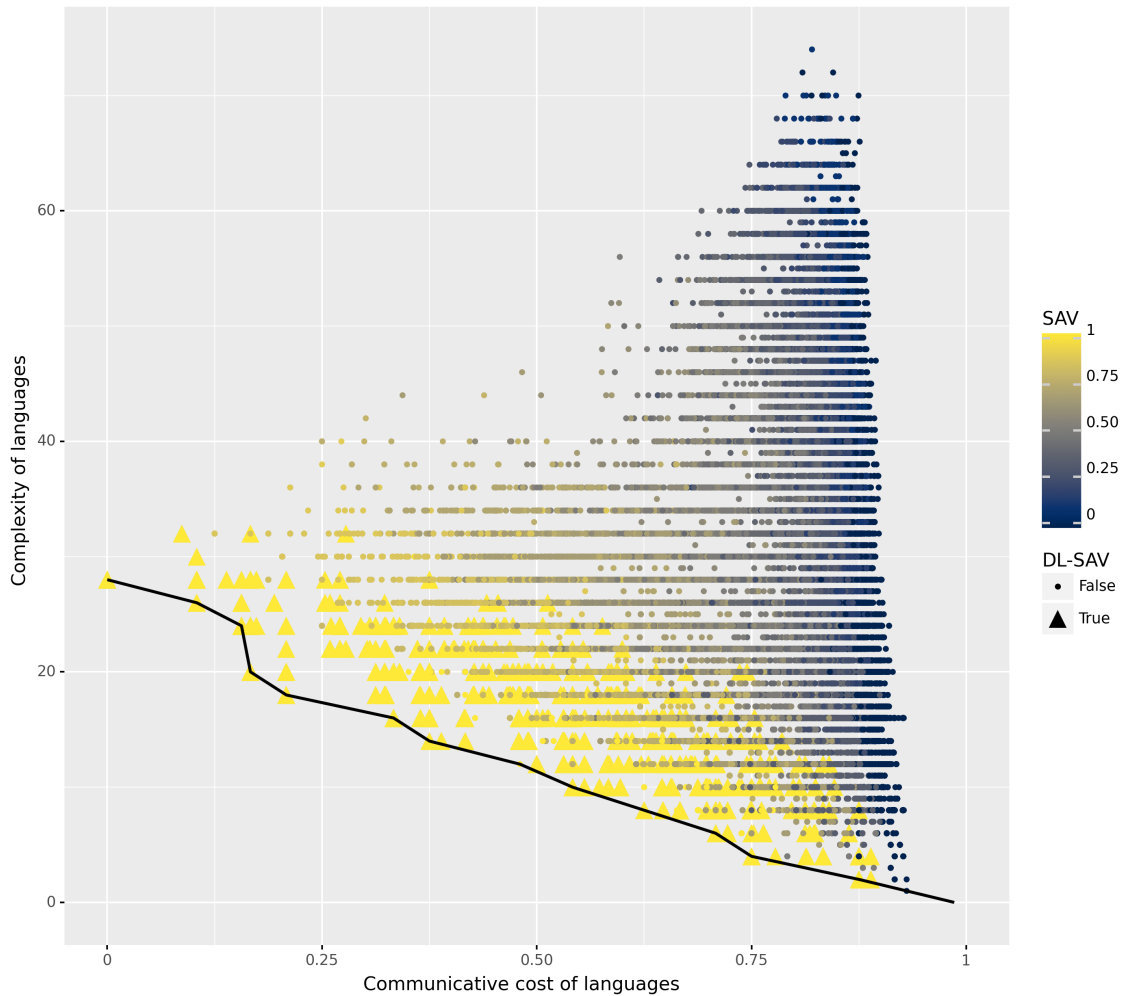


Figure 4.1: The complexity/communicative cost trade-off for modal languages. The black line is the Pareto frontier (optimal solutions) and robustly consists of the Domain-Level Single-Axis of Ambiguity typology, suggesting that the known semantic variation for modals is explained by general pressures for efficient communication.

For the main experiment of $N = 67278$ total languages, measuring Pearson coefficients showed that degree of SAV-satisfaction is indeed correlated with Pareto optimality ($r(N) =$

.59). This suggests that languages which have more modals satisfying Nauze’s SAV tend to do better at optimizing the simplicity/informativeness trade-off. In fact, degree-SAV has fairly low correlation with simplicity ($r(N) = .44$) and strong correlation with informativeness ($r(N) = .79$). This indicates that the communicative efficiency that the SAV universal enforces is mostly in virtue of a pressure for modal terms to be informative.¹ A comparison of mean simplicity, informativeness and Pareto optimality for languages satisfying the DL-SAV proposal is given in **Table 4.1**.²

language	N	simplicity	informativeness	optimality
mean DL-SAV	2255	0.8017	0.4714	0.9560
mean non-DL-SAV	65023	0.6398	0.2504	0.7974

Table 4.1: Mean simplicity, informativeness, and Pareto optimality for the DL-SAV typology.

Figure B.1 and **Table 4.1** show that the optimally simple and informative languages are overwhelmingly contained by the DL-SAV universal. In addition, the fact that the DL-SAV typology spans almost the entire range of informativeness suggests that modal semantic variation may be largely a result of the natural languages finding diverse solutions to this trade-off. A noteworthy region of the plot is the bottom right, where there is one optimal language with 0 modals compatible with the SAV proposal. Strikingly, this language contains one maximally simple modal that can express every meaning—an expected result of the trade-off. Overall, these results suggest that the DL-SAV modal semantic universal optimizes the simplicity/informativeness trade-off and that the known range of cross-linguistic modal meanings is plausibly explained by efficient communication.

¹The linear relationship between degree of SAV-satisfaction and Pareto optimality is plotted in **Appendix B**.

²Given the large sample sizes intrinsic to this setup, it is not clear how to interpret the t -values of a single samples t-test, or similarly the p -values from pearson correlations.

Chapter 5

DISCUSSION

The results of the computational experiment suggest that the two modal typologies presented by Nauze [2008] and Vander Klok [2013] are optimized for the simplicity/informativeness trade-off. I now follow Denic, Steinert-Threlkeld, and Szymanik [2020] in describing some of the choice points necessary for modeling in this context. For the present case of modals there are two related, main assumptions: (1) we have an appropriate way to measure the difference between natural languages and non-natural languages and (2) the measures of complexity and informativeness are aligned with reasonable intuitions about language naturalness. I also discuss (3) how this experiment ties in with related work investigating alternative explanations of semantic universals. Below I address these issues and areas of potential future work.

5.1 *Typology*

Two pressing questions are: what are the actual modal systems of the world, and what is the range of modal flavors that can be expressed across languages? There is a wealth of data available contained in descriptive grammars, in addition to the cross-linguistic survey offered by Nauze [2008] that is focused on the formal semantics of modality. While collecting the entire modal inventory for each language may present difficulties (since accurately measuring a language requires knowing the full modal vocabulary), I look forward to implementing in future work a systematic extraction of relevant data from such primary descriptive resources. There a technique will need to be developed which takes the measures of simplicity and informativeness across languages with possibly different force and flavor distinctions, and normalizes these values to enable a comparison of their relative optimality.

I can also ask: where does the linguistic category of *modal* begin and end? Where does this work draw the line, for example, between mood and modality? Does this analysis extend to intensional predicates in general (including adjectives such as *fragile*)? A simple answer to both questions is that the efficient communication results in this paper hold for any semantic phenomena with a meaning space always encoding underspecification across two kinds of features. This might include expressions falling under both mood and modality, in addition to other categories unrelated to displacement. Modals and attitude verbs¹ are merely natural targets for this analysis.

The Domain-Level Single Axis of Variability universal is a strikingly perspicuous description of the languages optimizing efficient communication. This raises a question: is it the most precise predicate at the level of describing specific constraints on the meaning space for modals—and independent of the axes defining the Pareto frontier—containing exactly the optimal languages? Clearly, it cannot be: because the DL-SAV captures a large superset of the Pareto frontier languages. We might ask: what *are* the specific properties shared by optimal (in possibly many senses) languages, and can such analyses be used to generate default hypotheses for cross-linguistic semantics investigations? A detailed exploration of candidate universals constructed from the results of this style of experiment represents an interesting research program, executable for modals and other semantic phenomena.

5.2 Measures of complexity and informativeness

The trade-off results here are a direct consequence of the measures of complexity and informativeness.² In this section I address some other possible models. There are other primitives

¹See Bogal-Allbritten [2016] for a variable flavor analysis of an attitude verb in Navajo and Močnik and Abramovitz [2019] for an variable-force and variable-flavor analysis of an attitude verb in Koryak.

²In other experiments on efficient communication, and for indefinite pronouns in particular the simplicity/informativeness variation looks different from my plotted sample. Normally, complexity and ambiguity are very strongly negatively correlated, which can result in a trade-off space that looks like an obtuse L .

But for the present case of modals and the deployed LOT, the high-complexity and low communicative cost languages—which could result in a more familiar trade-off shape—are absent. Conjecturing, it may be impossible to construct a highly informative and highly complex because adding more atoms in the LOT never decreases ambiguity.

for a Language of Thought than the ones selected for this paper, which could bear on the complexity of modal meanings. Rather than the boolean algebra given (propositional logic), I could have used a richer language that included first-order (or higher order) quantification. But because the feature space is more usefully interpreted as simple semantic map and not a model of objects in the world, my boolean algebra seems the simpler, preferred model. In addition, the operators in the LOT need not only be conjunction, disjunction, and negation. I could have used many combinations of connectives and operators, such as implication, XOR and NAND. Notably, there is empirical support from Piantadosi, Tenenbaum, and Goodman [2016] on the learnability of LOTs that people use a set of boolean connectives richer than the one in this experiment, but probably excluding XOR and NAND.

Consider, too, the role that certain axioms of the boolean algebra play in the modal complexity measure. The distributive property, for example, models a semantic equivalence between LOT expressions with possibly different (in particular, differing in complexity) syntactic representations. This decision was motivated by the intuition that modals with meanings grouped along an axis are easier to represent than modals lacking such uniformity, and ought to be measured as more simple. A similar motivation can be given for modeling an axis as the disjunction of all the points contained in that axis: a modal that can be used to express universal force and any modal flavor (e.g., *universal*) seems simpler than a fully unambiguous modal that must specify the exact force and flavor pair (e.g., *universal-deontic*) it can express.

5.3 *Alternative explanations*

Efficient communication appears to explain important variation in the modal typology, but it might not explain all of it. In particular, the explanation here is entirely consistent with a potential, competing explanation of the modal typology from a pressure for ease of learnability. The existence of this pressure has been demonstrated for several different semantic universals [Steinert-Threlkeld and Szymanik, 2020]. Such a paradigm is exemplified in the case of natural language quantifiers, which both optimize the simplicity/informativeness trade-off

and are among the easiest quantifiers to learn [Carcassi, Steinert-Threlkeld, and Szymanik, 2021].

The most significant challenge to constructing an alternative explanation of the modal typology in terms of ease of learnability will likely be defining a sufficiently rich feature space. In this experiment on efficient communication, I modeled *underspecification*, encoding no additional structure for the force and flavor features beyond presence/absence. In other words, there is nothing characteristically quantificational in the case of forces, and nothing particularly modal about the flavors. I look forward to developing models with additional structure capturing these aspects of meaning, which are plausibly relevant to evaluating whether natural language modals emerge as easiest to learn by, e.g., artificial agents in an iterated learning environment.

Chapter 6

CONCLUSION

This work argued that the known variation in modal semantic typology can be explained by a pressure for efficient communication. Specifically, we found that the proposed semantic universals by Nauze [2008] and Vander Klok [2013] restricting the underspecification of modal force and flavor possible in a language appear to approximate optimal solutions to the problem of trading off cognitive simplicity and informativeness. This result ties in with the recent work exploring the efficient communication hypothesis for diverse semantic phenomena. Finally, we discussed potential areas for future work, including a systematic extraction of typological data from descriptive resources, enriching the Language of Thought for modals and formulating a procedure for constructing default semantic universals from computational experiments.

BIBLIOGRAPHY

- Rajesh Bhatt, Tina Bogel, Miriam Butt, Annette Hautli, and Sebastian Sulger. Urdu/hindi modals. In Miriam Butt and Tracy Holloway King, editors, *Proceedings of LFG '11 Conference*, pages 47–67, 2011. URL <http://web.stanford.edu/group/cslipublications/cslipublications/LFG/16/papers/lfg11hattetal.pdf>.
- M. Ryan Bochnak. Variable force modality in washo. In Thuy Bui and Deniz Özyıldız, editors, *Proceedings of North-East Linguistic Society (NELS) 45*, 2015.
- M. Ryan Bochnak and Lisa Matthewson. Techniques in Complex Semantic Fieldwork. *Annual Review of Linguistics*, 6(1):261–283, 2020. ISSN 2333-9683, 2333-9691. doi: 10.1146/annurev-linguistics-011619-030452.
- Elizabeth Bogal-Allbritten. *Building Meaning in Navajo*. PhD thesis, UMass Amherst, 2016.
- Seth Cable. The expression of modality in tlingit: A paucity of grammatical devices. *International Journal of American Linguistics*, 83(4):619–678, 2017. doi: 10.1086/692975. URL <https://doi.org/10.1086/692975>.
- Fausto Carcassi, Shane Steinert-Threlkeld, and Jakub Szymanik. Monotone quantifiers emerge via iterated learning. *Cognitive Science*, 45(8):e13027, 2021. doi: <https://doi.org/10.1111/cogs.13027>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/cogs.13027>.
- Amy Rose Deal. Modals without scales. *Language (Baltimore)*, 87(3):559–585, 2011. ISSN 0097-8507.
- Milica Denić, Denić, Shane Steinert-Threlkeld, and Jakub Szymanik. Complexity/informa-

- tiveness trade-off in the domain of indefinite pronouns. In *Proceedings of Semantics and Linguistic Theory (SALT 30)*, 2020.
- Jacob Feldman. Minimization of boolean complexity in human concept learning. *Nature*, 407:630–3, 11 2000. doi: 10.1038/35036586.
- Jerry Fodor. *The Language of Thought*. Crowell, 1975.
- Charles F. Hockett. The origin of speech. *Scientific American*, 203:88–96, 1960. URL <http://www.isrl.uiuc.edu/~amag/langev/paper/charles60theOrigin.html>.
- Gerhard Jäger. The evolution of convex categories. *Linguistics and philosophy*, 30(5):551–564, 2007. ISSN 0165-0157.
- Charles Kemp and Terry Regier. Kinship categories across languages reflect general communicative principles. *Science*, 336:1049 – 1054, 2012. doi: 10.1126/science.1218811.
- Charles Kemp, Yang Xu, and Terry Regier. Semantic typology and efficient communication. *Annual review of linguistics*, 4(1):109–128, 2018. ISSN 2333-9683.
- A. Kratzer. What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy*, 1: 337–335, 1977. doi: <https://doi.org/10.1007/BF00353453>.
- Angelika Kratzer. The notional category of modality. In H.J Eikmeyer and H Reiser, editors, *Words, Worlds, and Contexts*, pages 38 – 74. de Gruyter, 1981.
- Angelika Kratzer. The Notional Category of Modality. In *Modals and Conditionals*. Oxford University Press, Oxford, 2012. ISBN 978-0-19-923468-4. doi: 10.1093/acprof:oso/9780199234684.003.0002. URL <https://oxford.universitypressscholarship.com/10.1093/acprof:oso/9780199234684.001.0001/acprof-9780199234684-chapter-2>.
- Lisa Matthewson. On the Methodology of Semantic Fieldwork. *International Journal of American Linguistics*, 70(4):369–415, 2004. ISSN 0020-7071. doi: 10.1086/429207.

Lisa Matthewson. *Modality*, page 525–559. Cambridge Handbooks in Language and Linguistics. Cambridge University Press, 2016. doi: 10.1017/CBO9781139236157.019.

Lisa Matthewson, Henry Davis, and Hotze Rullmann. Evidentials as epistemic modals: Evidence from st’át’imcets. *Linguistic Variation Yearbook*, 7(1):201 – 254, 2007. ISSN 15681483. URL <http://offcampus.lib.washington.edu/login?url=https://search.ebscohost.com/login.aspx?direct=true&db=cms&AN=32192603&site=ehost-live>.

Maša Močnik and Rafael Abramovitz. A Variable-Force Variable-Flavor Attitude Verb in Koryak. In Julian J Schlöder, Dean McHugh, and Floris Roelofsen, editors, *Proceedings of the 22nd Amsterdam Colloquium*, pages 494–503, 2019.

Fabrice Dominique Nauze. *Modality in Typological Perspective*. PhD thesis, Universiteit van Amsterdam, 2008.

Cailin O’Connor. The evolution of vagueness. *Erkenntnis*, 79(S4):707–727, 2014. ISSN 0165-0106.

Steven T Piantadosi, Joshua B Tenenbaum, and Noah D Goodman. The logical primitives of thought: Empirical foundations for compositional cognitive models. *Psychological review*, 123(4):392–424, 2016. ISSN 0033-295X.

Paul Portner. *Modality*. Oxford University Press, 2009. ISBN: 9780199292431 Place: Oxford ; New York Series: Oxford linguistics.

W. V. Quine. The problem of simplifying truth functions. *The American Mathematical Monthly*, 59(8):521–531, 1952. ISSN 00029890, 19300972. URL <http://www.jstor.org/stable/2308219>.

Terry Regier, Charles Kemp, and Paul Kay. Word meanings across languages support efficient communication. In B. MacWhinney and W. O’Grady, editors, *The handbook of language emergence*.

- Hotze Rullmann, Lisa Matthewson, and Henry Davis. Modals as distributive indefinites. *Natural language semantics*, 16(4):317–357, 2008. ISSN 0925-854X.
- Brian Skyrms. *Signals : evolution, learning, information*, 2010.
- Shane Steinert-Threlkeld. Quantifiers in natural language optimize the simplicity/informativeness trade-off. In *Proceedings of the 22nd Amsterdam Colloquium*, 2019.
- Shane Steinert-Threlkeld. Quantifiers in natural language: Efficient communication and degrees of semantic universals. *Entropy*, 23(10), 2021. ISSN 1099-4300. doi: 10.3390/e23101335. URL <https://www.mdpi.com/1099-4300/23/10/1335>.
- Shane Steinert-Threlkeld and Jakub Szymanik. Learnability and semantic universals. *Semantics and Pragmatics*, 12:1, 11 2019. doi: 10.3765/sp.12.4.
- Shane Steinert-Threlkeld and Jakub Szymanik. Ease of learning explains semantic universals. *Cognition*, 195:104076, 2020. ISSN 0010-0277. doi: <https://doi.org/10.1016/j.cognition.2019.104076>. URL <https://www.sciencedirect.com/science/article/pii/S0010027719302495>.
- Wataru Uegaki. Nand and the communicative efficiency model. Unpublished manuscript, University of Edinburgh, 2020.
- Johan van der Auwera and Vladimir A. Plungian. Modality’s semantic map. *Linguistic typology*, 2(1):79–124, 1998. ISSN 1430-0532.
- Jozina Vander Klok. Restrictions on semantic variation: a case study on modal system types. In *Workshop on Semantic Variation*, October 2013.
- Jozina Vander Klok. Revised Modal Questionnaire for Cross-Linguistic Use. unpublished, 2021. URL <https://jozinav.wordpress.com/linguistics/>.

Yang Xu and Terry Regier. Numeral systems across languages support efficient communication: From approximate numerosity to recursion. In *36th Annual Meeting of the Cognitive Science Society*, pages 1802 – 1807, 2014.

Yang Xu, Terry Regier, and Barbara C. Malt. Historical semantic chaining and efficient communication: The case of container names. *Cognitive Science*, 40(8):2081–2094, 2016. doi: 10.1111/cogs.12312.

Noga Zaslavsky, Charles Kemp, Terry Regier, and Naftali Tishby. Efficient compression in color naming and its evolution. *Proceedings of the National Academy of Sciences - PNAS*, 115(31):7937–7942, 2018. ISSN 0027-8424.

George K. Zipf. Human behavior and the principle of least effort. 1949.

Appendix A

ADDITIONAL PARAMETER RESULTS

In this appendix, I catalogue the efficient communication results for several experimental parameters: number of flavor distinctions in a language, two different utility functions for measuring informativeness, and the presence of negation in the modal LOT. **A.3** briefly explains how these results relate to the main experimental findings.

A.1 *Additional semantic spaces*




Figure A.1: The complexity/cost trade-off for modal languages with 4 and 5 flavors.

As Figure A.1 illustrates, the average informativeness of modal languages decreases as a function of the granularity (size) of the semantic space. DL-SAV languages are consistently

optimal for other configurations of semantic space size and language size. Note that although the relative complexity of languages appears to remain relatively unchanged, this is a result of normalization: languages with more forces and flavors have higher complexity in general.

A.2 Utility functions: indicator vs. half-credit




Figure A.2: The complexity/communicative cost trade-off for a graded (main results, left) versus an all-or-nothing utility function.

utility function	language	N	simplicity	informativeness	optimality
half-credit	DL-SAV	2255	0.8017	0.4714	0.9560
	non-DL-SAV	65023	0.6398	0.2504	0.7974
all-or-nothing	DL-SAV	885	0.8389	0.3085	0.9622
	non-DL-SAV	46132	0.5974	0.1425	0.7619

Table A.1: Results of DL-SAV optimality for two different communicative utility functions.

utility function	simplicity	informativeness	optimality
half-credit	0.44	0.79	0.59
all-or-nothing	0.35	0.69	0.43

Table A.2: Pearson coefficients for the correlation between degree-SAV satisfaction and simplicity, informativeness, and Pareto optimality.

Figure A.2 and Tables A.1, A.2 show that the choice of communicative utility function has an effect on the overall informativeness possible for languages (but not the qualitative results of trade-off, since both the SAV and DL-SAV proposals remain optimal). These results differ from the main experimental results only in that the listener model does not receive ‘half-credit’ for recovering merely the correct force or flavor; only if the listener correctly guesses the exact (force, flavor) pair is it rewarded. As expected, the overall values of informativeness are measured to be lower when the communicative utility of some modals are weighted by 0 instead of 0.5.

A.3 Language of Thought: negation




Figure A.3: The complexity/communicative cost trade-off with (main results, left) and without negation as a primitive in the Language of Thought.

A.4 Discussion

Overall, the result that the SAV and DL-SAV proposals appear to optimize the simplicity/informativeness trade-off appears robust to initial variations in several parameters. In particular, increasing the number of force and flavor distinctions that a language can make appears to lower the overall informativeness of any of the languages explored. The overall informativeness is also decreased by using the communicative utility function that only rewards perfect recovery of meanings (as opposed to half credit, in the main results). The presence of negation as a primitive in the Language of Thought appears insignificant.

Appendix B

DEGREE SAV VS. PARETO OPTIMALITY




Figure B.1: Degree of Single Axis of Variability Universal satisfaction vs. Pareto optimality

Figure 3 plots the optimality of modal languages against their degree of quasi-naturalness (fraction of vocabulary consistent with the SAV Universal). Notably, as languages become

more quasi-natural and optimal, the variation in optimality decreases. This suggests that SAV and DLSAV proposals are relatively accurate constraints on the best languages for efficient communication. The positive slope of the orange line of fit reflects the positive correlation reported in the main results.

Appendix C

EXPLICIT TYPOLOGIES

This section includes formal descriptions of the two proposed modal universals: Single Axis of Variability (SAV) proposed in Nauze (2008) and the Domain-Level Single Axis of Variability (DL-SAV) in Vander Klok (2013).

THE SINGLE AXIS OF VARIABILITY (SAV) UNIVERSAL: All modals in natural language satisfy the single axis of variability property: if a modal can express more than one flavor, it can only express one force (and *mutatis mutandis* for force and flavor). That is to say: a modal may exhibit variable force or variable flavor, but not both.

[Alternative formulation: $|(m)| = 1$ or $|(m)| = 1$, where $|\cdot|$ is the set cardinality function.]

THE DOMAIN LEVEL SINGLE AXIS OF VARIABILITY (DL-SAV) UNIVERSAL: All modal systems in natural language satisfy the domain-level single axis of variability property: *a modal system as a whole* may only exhibit variability on a single axis in each of the root and epistemic domains. That is: if one root modal exhibits variability on the flavor axis, no other root modal exhibits variability on the force axis (though an epistemic modal may do so) and *mutatis mutandis* for epistemic modals and also for the force axis.

For the DL-SAV, I assume that the flavor *epistemic* is not divisible, no matter the size of the semantic space. Consequently, when measuring of artificial modal languages whether they satisfy the DL-SAV, I arbitrarily but consistently identify it with the first column in the table representation of the modal meaning space.

Appendix D

BOOLEAN ALGEBRA (LANGUAGE OF THOUGHT FOR MODALS)

Here we specify a few salient properties exploited in the heuristic for finding shortest formulae. Sets and their operations are defined in the usual way as a convenient meta-level language.

A modal semantic space \mathbb{M} of size $n \times m$ is partitioned into a pair (Q, F) such that $|Q| = n$ and $|F| = m$ and $\mathbb{M} = Q \cup F$. We define a boolean algebra on the semantic space with two binary operations \vee and \wedge and a unary operation \neg .

Disjunction

Commutativity

$$\mathbf{A.1} \quad X \vee Y = Y \vee X$$

Array cover

$$\mathbf{A.2.a} \quad \bigvee_i^n Q_i = \top$$

$$\mathbf{A.2.b} \quad \bigvee_i^n f_i = \top$$

Identity

$$\mathbf{A.3} \quad \perp \vee X = X$$

Idempotence

$$\mathbf{A.4} \quad X \vee X = X$$

Conjunction

Commutativity

$$\mathbf{M.1} \quad X \wedge Y = Y \wedge X$$

Array annihilation

$$\mathbf{M.2.a} \quad \forall i \neq j : Q_i \wedge Q_j = \perp$$

$$\mathbf{M.2.b} \quad \forall i \neq j : f_i \wedge f_j = \perp$$

Idempotence

$$\mathbf{M.3.} \quad X \wedge X = X$$

Distributivity of conjunction over disjunction

$$\mathbf{D.1} \quad (X \wedge Y) \vee (X \wedge Z) = X \wedge (Y \vee Z)$$

Complementation

$$\mathbf{C.1.a} \quad \forall P \subseteq Q : Q - P = \neg P$$

$$\mathbf{C.1.b} \quad \forall G \subseteq F : F - G = \neg G$$

DeMorgan

$$\neg(X \wedge Y) = \neg X \vee \neg Y$$

Example

The array representation of a modal is mapped to a disjunctive normal form (DNF) in the boolean algebra. This DNF can then be reduced to a simpler, logically equivalent expression:

$$\begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \mapsto (Q_1 \wedge f_1) \vee (Q_1 \wedge f_2) \vee (Q_2 \wedge f_1) \vee (Q_2 \wedge f_2) \vee (Q_2 \wedge f_3) = \neg(Q_1 \wedge f_3)$$

Proof.

$$\begin{aligned} & (Q_1 \wedge f_1) \vee (Q_1 \wedge f_2) \vee (Q_2 \wedge f_1) \vee (Q_2 \wedge f_2) \vee (Q_2 \wedge f_3) && \text{(given DNF)} \\ & = Q_1 \wedge (f_1 \vee f_2) \vee Q_2 \wedge (f_1 \vee f_2 \vee f_3) && \text{(Distributivity)} \\ & = Q_1 \wedge (f_1 \vee f_2) \vee Q_2 \wedge \top && \text{(Array cover)} \\ & = Q_1 \wedge (f_1 \vee f_2) \vee Q_2 && \text{(Conjunctive Identity)} \\ & = Q_1 \wedge \neg f_3 \vee Q_2 && \text{(Complementation)} \\ & = \neg(\neg Q_1 \vee f_3) \vee Q_2 && \text{(DeMorgan)} \\ & = \neg(Q_2 \vee f_3) \vee Q_2 && \text{(Complementation)} \\ & = \neg((Q_2 \vee f_3) \wedge \neg Q_2) && \text{(DeMorgan)} \\ & = \neg((Q_2 \vee f_3) \wedge Q_1) && \text{(Complementation)} \\ & = \neg(Q_1 \wedge Q_2 \vee Q_1 \wedge f_3) && \text{(Distributivity)} \\ & = \neg(\perp \vee Q_1 \wedge f_3) && \text{(Array annihilation)} \\ & = \neg(Q_1 \wedge f_3) && \text{(Disjunctive Identity)} \end{aligned}$$

Appendix E

SHORTEST FORMULA HEURISTIC SEARCH ALGORITHM

This procedure is applied to every possible modal generated by a meaning space. Note that each operation can be seen as one of the logical inferences described in Appendix D. These operations were implemented as functions defined over the (syntactic) tree representations of each formula. Importantly, the procedure only searches the additional semantically coextensive formulae that distinct from the current formula. This is a more relaxed condition than requiring candidate formulae to be shorter than the current one, allowing this procedure to have access to solutions contained in search paths for which expression complexity is not generally decreasing.

Algorithm 1 Approximating the shortest formula for a modal.

```

function SHORTEN( $a$ , operations, max_iterations)    ▷ a modal  $a$  is a set of meaning
points
     $e \leftarrow \text{array\_to\_dnf}(a)$ 
    to_visit  $\leftarrow \{e\}$ 
     $i \leftarrow 0$ 
    shortestst  $\leftarrow e$ 
    while to_visit  $\neq \emptyset$  do                                ▷ breadth-first search
        if  $i = \text{max\_iterations}$  then
            break
        end if
        next  $\leftarrow \text{to\_visit.dequeue}()$ 
        children =  $\emptyset$ 
        for operation in operations do
            children.enqueue(operation(next))
        end for
        for child in children do
            if child  $\neq$  next then
                to_visit.enqueue(child)
            end if
        end for
         $i \leftarrow i + 1$ 
        if complexity(next) < complexity(shortest) then
            shortestst  $\leftarrow$  next
        end if
    end while
    return shortestst
end function

```
