

Finding Concept-specific Biases in Form–Meaning Associations

Tiago Pimentel^δ Brian Roark^α Søren Wichmann[†] Ryan Cotterell^{δ,‡} Damián Blasi^{¶,α,⊐}

^δUniversity of Cambridge ^αGoogle [†]Kazan Federal University [‡]ETH Zürich

[¶]Harvard University ^αMPI for Evolutionary Anthropology [⊐]HSE University

tp472@cam.ac.uk, roark@google.com, wichmannsoeren@gmail.com

ryan.cotterell@inf.ethz.ch, dblasi@fas.harvard.edu

Abstract

This work presents an information-theoretic operationalisation of cross-linguistic non-arbitrariness. It is not a new idea that there are small, cross-linguistic associations between the forms and meanings of words. For instance, it has been claimed (Blasi et al., 2016) that the word for TONGUE is more likely than chance to contain the phone [l]. By controlling for the influence of language family and geographic proximity within a very large concept-aligned, cross-lingual lexicon, we extend methods previously used to detect within language non-arbitrariness (Pimentel et al., 2019) to measure cross-linguistic associations. We find that there is a significant effect of non-arbitrariness, but it is unsurprisingly small (less than 0.5% on average according to our information-theoretic estimate). We also provide a concept-level analysis which shows that a quarter of the concepts considered in our work exhibit a significant level of cross-linguistic non-arbitrariness. In sum, the paper provides new methods to detect cross-linguistic associations at scale, and confirms their effects are minor.

1 Introduction

The arbitrariness of the sign, i.e. the principle that a word’s form is unrelated to what it denotes, was one of the cornerstones in the structuralist revolution in linguistics (Saussure, 1916). While languages do seem to adhere to the principle to a large extent, researchers have repeatedly uncovered evidence that there are preferences in form–meaning matches (Perniss et al., 2010). Indeed, the notion that these small, but systematic, form–meaning relations hold *across* the world’s languages has become a mainstream topic of research in the last couple of decades.¹

¹See §2 below for a brief literature review and Dingemanse et al. (2015) for a more comprehensive one.




Figure 1: We used a sample of 5189 languages (9148 doculects) to study cross-linguistic systematicity. In this map, the colours represent the four macroareas to which languages were assigned.

Determining effective metrics to capture meaningful form–meaning associations is far from trivial, though, and researchers have explored a substantial number of statistical and heuristic approaches (Bergen, 2004; Wichmann et al., 2010; Johansson and Zlatev, 2013; Haynie et al., 2014; Gutierrez et al., 2016; Blasi et al., 2016; Joo, 2019). Previous studies differ from each other along (at least) three axes: (i) which unit is used to measure wordform similarity (e.g., phonemes, sub-phonemic features or arbitrary sequences); (ii) how they deploy a baseline for statistical comparison (e.g. permute forms with meanings, or propose a generative model that yields wordforms uninformed by their meaning) and (iii) whether they study non-arbitrariness within or across languages.

Pimentel et al. (2019) provide the first holistic measure of non-arbitrariness (in a large vocabulary sample of a single language) using tools from information theory, and apply their measure to discover phonesthemes.² Our work extends their approach to the problem of discovering and estimating the strength of frequent *cross-linguistic* form–meaning associations (e.g. iconicity and systematicity) in individual concepts. We do this by adapting Pimentel et al.’s (2019) approach, modelling

²Phonesthemes are sub-morphemic units which are associated in a language with some small semantic domain.

form–meaning associations in a large collection of basic vocabulary wordlists covering close to $\frac{3}{4}$ of the world’s languages (see Fig. 1 and [Wichmann et al., 2020](#)). By taking the words in these lists to be random variables and asking how much information within wordforms is explained by the meaning they refer to, we obtain a quantitative estimate of cross-linguistic non-arbitrariness.

Specifically, we propose to model a universal (language-independent) form distribution (using neural language models), and then we estimate concept-specific distributions. With these in hand, we are able to determine how much the meaning of a concept predicts its form cross-linguistically—by measuring the mutual information between them; see §4 for details. This method further allows us to identify which concepts exhibit stronger non-arbitrary form–meaning association and which form patterns are more likely to occur in them.

In order to maximise the reliability of the observed associations, we implement stringent controls for genealogical and areal effects, as well as for the size of each language family. See §4.5 for details on these controls. After introducing these controls, we find that wordlists display an average of around 0.01 bits of form–meaning mutual information explained by cross-linguistic non-arbitrariness ($\approx 0.3\%$ of the wordform uncertainty) with substantial variation among concepts and languages.³ Of the 100 basic concepts in our data, we find a statistically identifiable pattern in 26 of them ($p < 0.01$). Inspection of the results show that our method recovers previously proposed associations, e.g. the association of [l] with the concept TONGUE and [p] with FULL ([Blasi et al., 2016](#)).

2 Non-Arbitrary Form–Meaning Associations

Several studies have looked at non-arbitrary patterns in languages, be it systematicity ([Shillcock et al., 2001](#); [Gutierrez et al., 2016](#); [Dautriche et al., 2017](#); [Pimentel et al., 2019](#)) or iconicity ([Dingemanse, 2012, 2018](#)). With respect to cross-linguistic non-arbitrariness specifically, the hypothesised sources of form–meaning associations range from the fact that humans are endowed with the same neurocognitive architecture ([Bankieris and Simner, 2015](#)) to their encountering similar experiences within the world ([Parise et al., 2014](#)).

³See §3 for a discussion on potential biases in the dataset that likely influence the actual numerical value derived.

While global non-arbitrary form–meaning associations have been hypothesised to exist at different levels of linguistic description ([Haiman, 1980](#)), by far the component of language that has received the most attention in this respect is the lexicon. A few circumstances facilitate this type of research in contrast to other domains of grammar. For instance, the space of possible words that could be used in a given language to refer to an arbitrary referent is large, whereas the relative canonical order of a verb with respect to its object complement is substantially smaller (which renders cross-linguistic similarities less informative than in the first case). Additionally, the sheer amount of data available in the form of wordlists exceeds other types of linguistic data for the languages of the world.

As a consequence, some of the largest evaluations of non-arbitrary form–meaning associations involve systematic wordlists with comparable referents across languages ([Wichmann et al., 2010](#); [Johansson and Zlatev, 2013](#); [Haynie et al., 2014](#); [Blasi et al., 2016](#); [Joo, 2019](#)). Most of these studies were focused on the regular association between phonemic or phonetic units with meaning, occasionally controlling for other potential sources of form–meaning association such as phonotactics or word length ([Blasi et al., 2016](#)). While useful, the estimates emerging from this type of study can be regarded as lower bounds to the total amount of non-arbitrary associations found in the vocabulary.

Recent efforts have resulted in datasets with thousands of languages ([Wichmann et al., 2020](#)), with which linguists can look for universal statistical patterns ([Wichmann et al., 2010](#); [Blasi et al., 2016](#)). These studies, though, only looked at the presence (or not) of individual phones in words, not accounting for their connections. Our methods rely on neural phonotactic models, similar to those used by [Pimentel et al. \(2020\)](#), thus capturing a broader range of potential correspondences.

3 Data

An exceptional resource with substantial cross-linguistic representation is provided in the Automated Similarity Judgment Program, better known by its acronym ASJP ([Wichmann et al., 2020](#)). ASJP is a collection of basic vocabulary wordlists, i.e. lists of words with referents that are expected to be widely attested across human societies. It involves body parts, some colour terms, lower numerals, general properties (such as big or round), and flora and fauna that are usually found in places

where humans live (e.g. trees and dogs). The individual words in ASJP are transcribed by field linguists in a specific phonetic annotation scheme that involves 41 symbols, chosen in order to maximise cross-linguistic utility by merging rare phones with similar phonetic features within the same category. These wordlists are assembled with the purpose of studying the history of languages—following the tradition established by Swadesh (1955)—under the principles of the comparative method.

ASJP has gathered, in its latest iterations, data for close to $3/4$ of the world’s languages, which makes it an unparalleled resource for evaluating form–meaning associations across spoken languages. Furthermore, the vocabulary in its wordlists was chosen as so to be resistant to borrowings—making it especially interesting for our purposes of finding universal form–meaning biases. We leave out pidgin and creole data,⁴ as defined by the World Atlas of Language Structures (Dryer and Haspelmath, 2013), since they are ambiguous with relation to their genealogical affiliation. We also omit constructed and fake languages (e.g. Esperanto and Taensa). This leaves us 9148 doculects (or wordlists) from 5189 languages.⁵

Form–meaning associations have been studied in earlier versions of this dataset. Firstly, Wichmann et al. (2010) studied the average form across different concepts in ASJP, and found a number of tentative patterns pointing to non-arbitrariness. Yet the lack of historical and statistical controls compromised the nature of such patterns: form–meaning associations could be due to widespread linguistic contact (e.g. the word for DOG, Pache et al. 2016) or to its fortuitous presence in large families. Blasi et al. (2016), however, provide a conservative evaluation of individual form–meaning associations by imposing a restrictive set of conditions. They looked for associations that were present in a minimum number of continents and language families. This resulted in a sizable number of non-arbitrary

associations, many of which had been highlighted as interesting based on behavioural and linguistic experiments in a handful of languages.

Data Disclaimer. As mentioned above, ASJP gathers lists of wordforms that are expected to be present across most human societies and their corresponding language(s). While this guarantees a fair coverage in our study, it limits the scope of our conclusions to those concepts present herein.

4 Methods

4.1 Notation

We describe each word as comprised by form and meaning, which we represent as a pair $(\mathbf{w}^{(n)}, \mathbf{v}^{(n)})$. The form $\mathbf{w}^{(n)} \in \Sigma^*$ is represented as a phone string where Σ is a phonetic alphabet. In this work, we take Σ to be the set of 41 phonetic symbols in ASJP plus the end-of-string symbol. We write W to denote a Σ^* -valued random variable. The meaning $\mathbf{v}^{(n)} \in \{0, 1\}^K$ is represented by a one-hot vector, where K is the number of analysed concepts.⁶ We write V to denote a $\{0, 1\}^K$ -valued random variable.

4.2 Non-Arbitrariness as Mutual Information

The goal of this work is to measure cross-linguistic form–meaning associations, operationalised as the **mutual information** (MI) between a form-valued random variable W and a meaning-valued random variable V . Symbolically, we are interested in computing (Cover and Thomas, 2012):

$$I(W; V) = H(W) - H(W | V) \quad (1)$$

Intuitively, this quantity captures the uncertainty we have over the form, the entropy $H(W)$, minus how much uncertainty we have over the form *given* the meaning, the conditional entropy $H(W | V)$. Thus, if eq. (1) is zero, its minimum, we have the result that meaning tells us absolutely nothing about the wordform. On the other hand, if eq. (1) is $\min\{H(W), H(V)\}$, its maximum, we have that the form is a deterministic function of the meaning (or the opposite; the meaning being deterministically determined given the form).

⁶We note Pimentel et al. (2019) used high-dimensional distributional semantic vectors to represent meaning, while we use a one-hot vector. However, their work relied on a specific language’s WORD2VEC—a choice which could potentially bias our results with that language’s properties. We did, however, run an extra experiment with English WORD2VEC; this led to similar conclusions to the ones presented here.

⁴Pidgins are believed to rely particularly on iconicity due to their smaller degree of lexicalisation; this reliance then diminishes as it morphs into a creole (Romaine, 1988). Future work could expand the methods here to study this phenomenon.

⁵When there is more than one wordlist for one language (as defined by their ISO-codes) one can sometimes refer to them as different dialects, but these are often just alternative versions of the same language as recorded by different linguists. There can be as much variation in such different recordings as among different dialects recorded by one and the same linguist. For those reasons, it is practical to use the term *doculect*, which we adopt here. This is a neutral term that refers to some dialect *as recorded in some specific source*.

That the mutual information may take values in $[0, \min\{H(W), H(V)\}]$ —together with the fact that, for our specific study, $H(W)$ is smaller than $H(V)$ —suggests a more interpretable metric called the **uncertainty coefficient**:

$$U(W | V) = \frac{I(W; V)}{H(W)} \quad (2)$$

This quantity is the proportion of uncertainty in the form reduced by knowing the meaning. Both mutual information and uncertainty coefficients are general measures of non-arbitrariness. One might also inquire about how non-arbitrary a single form–meaning pair is. To measure this, we propose **point-wise mutual information** (PMI):

$$\text{PMI}(\mathbf{w}; \mathbf{v}) = \log \frac{p(\mathbf{w} | \mathbf{v})}{p(\mathbf{w})} \quad (3)$$

4.3 Approximating Mutual Information

As noted above, we want to estimate the entropy of language agnostic wordforms, i.e.

$$H(W) = \sum_{\mathbf{w} \in \Sigma^*} p(\mathbf{w}) \log \frac{1}{p(\mathbf{w})} \quad (4)$$

Unfortunately, we do not know the exact distribution of $p(\mathbf{w})$ and, even if we did, we would need to sum over the infinite set of possible strings Σ^* to compute this entropy, which is intractable. If we have another probability distribution $p_\theta(\mathbf{w})$, though, we can calculate the cross-entropy between them as an approximation, i.e.

$$H(W) \leq H_\theta(W) \approx \frac{1}{N} \sum_{n=1}^N \log \frac{1}{p_\theta(\tilde{\mathbf{w}}^{(n)})} \quad (5)$$

where $\{\tilde{\mathbf{w}}^{(n)}\}_{n=1}^N$ are samples from the true distribution p . Throughout the paper, the tilde marks held-out data, i.e., data not used during model training. We note that the approximation becomes exact as $N \rightarrow \infty$ by the weak law of large numbers. This cross-entropy estimate gives us an upper bound on the actual entropy. This bound is tighter the closer the distributions $p(\mathbf{w})$ and $p_\theta(\mathbf{w})$ are.

4.4 Estimating the Approximator p_θ

How should we train a model to estimate this universal phonotactic distribution $p_\theta(\mathbf{w})$, though? We train a phone-level language model to predict the next phone given previous ones in a word, i.e.

$$p_\theta(\mathbf{w}) = \prod_{t=1}^{|\mathbf{w}|} p_\theta(w_t | \mathbf{w}_{<t}) \quad (6)$$

In this work, we use an LSTM as our language model (Hochreiter and Schmidhuber, 1997). Each phone w_t is represented using a lookup embedding $\mathbf{z}_t \in \mathbb{R}^d$. These are fed into the LSTM, outputting temporal representations of the sequence:

$$\mathbf{h}_t = \text{LSTM}(\mathbf{z}_{t-1}, \mathbf{h}_{t-1}) \quad (7)$$

where \mathbf{h}_0 is the zero vector. These representations are linearly transformed and used in a softmax to approximate the probability distribution:

$$p_\theta(w_t | \mathbf{w}_{<t}) = \text{softmax}(\mathbf{W} \mathbf{h}_t + \mathbf{b}) \quad (8)$$

All parameters are learned via gradient descent, minimising the cross-entropy in the training set.

4.5 Cross-linguistic Controls

As mentioned before, salient regularities between form and meaning across languages might result from large groups of genealogically or spatially related languages. In particular it is practical to consider two independent problems in this respect:

- (i) Eq. (5)’s inequality only holds if $H_\theta(W)$ is estimated on a set of datapoints sampled independently from the set of points on which the model p_θ was trained. As such, the test set should only include languages that are not genealogically or areally related to those in the training set;
- (ii) Within our dataset, the different size of areal and genealogical groups should be accounted for so that our results are not biased towards particularly large areas or language families.

Train–test split. To mitigate the problem referred to in the first item, we cross-validate our models by appealing to the notion of macroareas, large-scale regions of the world that simultaneously maximise internal historical dependency while minimising external ones. Striking a balance between historical independence and data availability, we consider the following four macroareas: the Americas, Eurasia, Africa, and the Pacific (which in this instantiation includes Papua New Guinea and Australia—see Fig. 1). We will use these macroareas as our folds. Two macroareas will be used at each time for training, while one other is used for validation and the last for testing. Some language families, though, might be present in more than one macroarea (e.g. many European languages are spoken natively in the Americas and Africa). These

families will be assigned to the one macroarea which contains most of its family members, since we believe reducing genealogical impact should be preferred over areal impact for our data and purposes, in cases for which such a choice is required.⁷

Family size bias. The second problem is tackled by weighting each example’s contribution to our loss function by the inverse of its family size $l^{(n)}$:

$$\mathcal{L}(\theta) = \frac{1}{L} \sum_{n=1}^N \frac{1}{l^{(n)}} \log \frac{1}{p_{\theta}(\mathbf{w}^{(n)})} \quad (9)$$

where $L = \sum_{n=1}^N \frac{1}{l^{(n)}}$ re-normalises the cross-entropy using the family sizes. This weighted cross-entropy loss function makes per instance contributions of large language families smaller, reducing their impact on the trained model.

To mitigate the same bias effect on the evaluation of validation and test sets, we first get cross-entropies per word. We subsequently average them per language, per family, and per macroarea. This way, each family will have the same effect per macroarea and each macroarea will have the same effect on the overall cross-entropy.

4.6 Concept-Specific Form Distributions

We want to compare per-concept phonotactic models with general ones to analyse sound–meaning associations. With that in mind, we condition phone-level language models on meaning:

$$p_{\theta}(\mathbf{w} | \mathbf{v}) = \prod_{t=1}^{|\mathbf{w}|} p_{\theta}(w_t | \mathbf{w}_{<t}, \mathbf{v}) \quad (10)$$

These models are trained following the same procedures explained above, but conditioning the LSTMs on concept specific representations. Specifically, the one-hot representation is linearly transformed and fed into the LSTM as its initial state

$$\mathbf{h}_0 = \mathbf{W}_0 \mathbf{v} \quad (11)$$

where the linear transformation $\mathbf{W}_0 \in \mathbb{R}^{d \times K}$ is randomly initialised and learned with the rest of the model. We then use this distribution to estimate the conditional entropy, analogously to eq. (5), as in

$$H(W | V) \lesssim \frac{1}{N} \sum_{n=1}^N \log \frac{1}{p_{\theta}(\tilde{\mathbf{w}}^{(n)} | \tilde{\mathbf{v}}^{(n)})} \quad (12)$$

⁷As mentioned in §3, the list of concepts in ASJP was chosen to minimise borrowings across languages. We further note here that loan words are annotated in this dataset and we drop those words for the purpose of our analysis.

where $\{\tilde{\mathbf{w}}^{(n)}, \tilde{\mathbf{v}}^{(n)}\}_{n=1}^N$ are held-out from–meaning pairs, sampled from the true distribution.⁸

4.7 Non-Arbitrariness as Information

The mutual information between wordforms and meaning can be decomposed into the difference of two entropy measures. Unfortunately, we have no way of directly measuring these entropy values without their probability distributions ($p(\mathbf{w})$ and $p(\mathbf{w} | \mathbf{v})$). We use the estimated cross-entropies as an approximation to this mutual information:

$$I(W; V) = H(W) - H(W | V) \quad (13)$$

$$\approx H_{\theta}(W) - H_{\theta}(W | V) \quad (14)$$

We note that eq. (14) is approximate because it is the difference of two upper bounds. Furthermore, while there are many ways to estimate mutual information, computing it as the difference between two cross-entropies seems to produce consistent results (McAllester and Stratos, 2020).

4.8 Bounds and Optimisation

As mentioned in §4.3, our entropy upper bounds will be tighter if our models p_{θ} better capture p . With this in mind, we optimise the hyper-parameters of our models using Bayesian optimisation with a Gaussian process prior (Snoek et al., 2012)—hyper-parameter ranges are presented in App. A. We train 25 models for each configuration and choose the best one according to the validation set, optimising our weighted cross-entropy loss using AdamW (Loshchilov and Hutter, 2019).

5 Experiments and Analysis⁹

5.1 Analysis #1: Overall Mutual Information

We are interested in estimating the cross-linguistic mutual information between meaning and wordforms. With this in mind, we follow the steps described in §4.4, but instead of only 1 model, we train 25 models using different seeds for each fold (totalling 100 models). Average results—overall and per macroarea—are shown in Tab. 1.

⁸This meaning conditioned model may potentially be better than the raw LSTMs (without conditioning on meaning; due to the extra parameters). To control for this fact, we ran an extra experiment where we estimated $H(W)$ using the meaning dependent model with shuffled concept IDs (so there is no form–meaning association). The results from this shuffled IDs model were very similar to the raw LSTM ones.

⁹Our code is available at <https://github.com/rycolab/form-meaning-associations>.




Figure 2: Uncertainty Coefficients for all 100 concepts in ASJP distributed according to their average word length.

Macroarea			Systematicity		
Train	Validation	Test	$H(W)$	$I(W; V)$	$U(W V)$
P, A _m	Eurasia	Africa	3.773	0.011*	0.279%
E, A _f	Pacific	Americas	3.901	0.007	0.173%
A _f , P	Americas	Eurasia	3.999	0.015 [‡]	0.376%
A _m , E	Africa	Pacific	3.755	0.016 [‡]	0.422%
Average			3.857	0.012 [‡]	0.312%

[‡] $p < 0.01$ * $p < 0.1$

Table 1: Mutual information (in bits per phone) between meaning and wordforms. In the train column, P: Pacific, E: Eurasia, A_f: Africa, A_m: Americas.

Across macroareas, results indicate a small average contribution of meaning into form (in all cases smaller than 1%).¹⁰ A simple permutation test (explained later in this section) indicates that, under standard levels of significance ($\alpha = 0.01$) and after controlling for multiple comparisons,¹¹ this average quantity is significant in 2 out of 4 of the macroareas. Nevertheless, this should not be overinterpreted, as unaccounted factors might be responsible for these effects; for instance, the impact of shared history across families in regions smaller than macroareas (almost all human languages have been in contact, directly or indirectly). Hence it is reasonable to conclude that there is no definitive evidence for an overall average association at this level of description of the data. We consider specific concept form–meaning associations next.¹²

¹⁰For comparison, Pimentel et al. (2019) estimate intra-language systematicity only accounts for roughly 3 ~ 5% of the entropy in wordforms in English, German and Dutch (given a characteristic sample of the vocabulary).

¹¹All our experiments rely on Benjamini and Hochberg (1995) corrections

¹²We ran an experiment changing the macroarea combinations in the train-validation-test sets and the results were stable, leading only to minor numerical changes to Tab. 1.

Paired Permutation Tests. For the permutation test, we first get the average MI over the 25 random seed results for a macroarea. We then permute the signs on these 25 results to create 10^5 new average MIs. By comparing the original result with these permutation ones we get the probability that our MI estimate is significantly larger than zero. A relevant detail is that these tests are performed in estimates—as opposed to real MI. The mutual information is always non-negative, but our estimate is not. If the MI is zero, we expect our estimates to be negative half the time, since both upper bounds should be roughly equivalent $H_{\theta}(W) \approx H_{\theta}(W | V)$.

A note on the LSTMs quality. Our results strongly rely on the quality of approximations. Our language independent $H(W)$ estimate is 3.85 bits per phone. Meanwhile, the per-language phonotactic cross-entropy found by Pimentel et al. (2020) is, on average, roughly 3 bits per phone—generally speaking, these results seem consistent.¹³ Furthermore, our model’s cross-entropy on the training set is 3.73—while it may have overfit slightly, this is not an aberration.

5.2 Analysis #2: Per Concept

In this section we focus on concept-specific form–meaning associations. With this in mind we group all words for a specific concept $c \in C$ into a set:

$$S_c = \left\{ (\tilde{\mathbf{w}}^{(n)}, \tilde{\mathbf{v}}^{(n)}) \mid \tilde{\mathbf{v}}^{(n)} = \mathbf{c} \in C \right\} \quad (15)$$

For each such set, we run a permutation test on their approximated pointwise mutual information

¹³These results are not directly comparable, though, since words are encoded with different phonetic alphabets in ASJP and NorthEuraLex (Dellert et al., 2020).

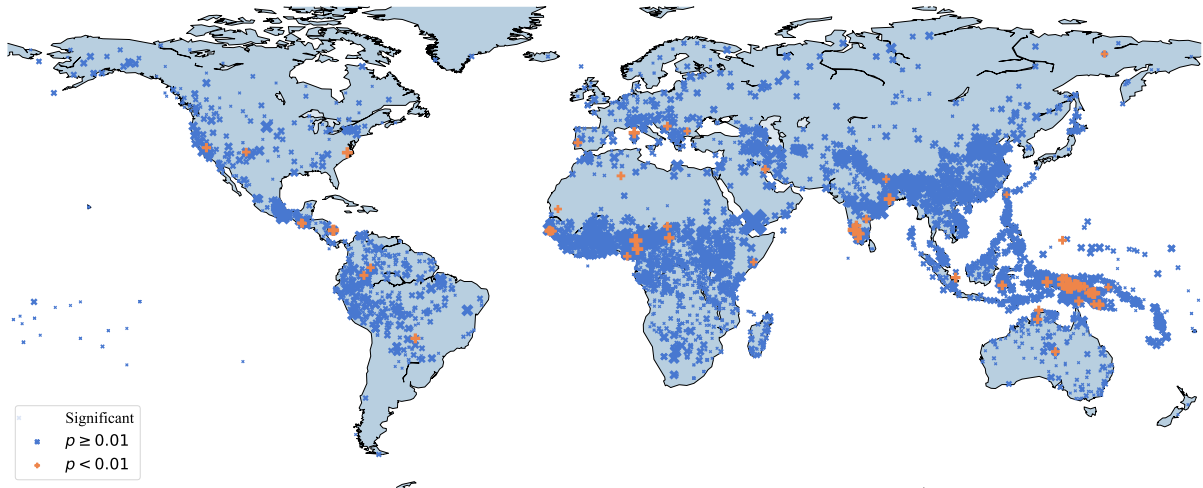


Figure 3: Uncertainty Coefficients for the 5189 languages (9148 doculects) in ASJP. Each language is represented by a point and larger points imply a larger coefficient. Significance was assessed through permutation tests.

values $\text{PMI}(\mathbf{w}^{(n)}; \mathbf{v}^{(n)})$, assessing if a concept has a statistically significant sound–meaning association.¹⁴ Of the 100 concepts in our dataset, 26 of them have positive mutual information ($p < 0.01$). This means that, at least in the set of concepts represented in our dataset, non-arbitrary form–meaning associations are not exceptions.

We present the average uncertainty coefficient per concept compared to average wordform length in Fig. 2. We do not find any correlation between these measurements. Analysing these results more closely, we see the pronouns *I* and *you* present the highest coefficient values. Most colours in our dataset (*white, red, green, yellow*) show statistically positive MI. Furthermore, some concepts related to body parts (*tongue, skin, knee, heart, claw*) and several concepts related to the environment (*water, sand, star, cloud, dry, cold*) have statistically positive results.

Wichmann et al. (2010) also looked at how concepts differ in their degree of form–meaning associations, presenting them in an ordered list together with a measure of how much they deviate from a global average phone usage. They only look at isolated phone’s frequencies, though, and do not control for word length—our mutual information metric controls for both factors. When we compare our results to Wichmann et al.’s (2010) top 10 list of concepts, we see both contain several body parts (*tongue, skin, knee*) and pronouns (*I, you*).

¹⁴This permutation test is similar to the one in §5.1, but uses the family size corrections discussed in §4.5 when averaging results—i.e., for each permutation (and the original one) we average words, languages, families, and macroareas, in this sequence, to get the MI estimate.

5.3 Analysis #3: Per Language

In their position paper, Perniss et al. (2010) argue that non-arbitrariness is a general property of language, although sometimes believed to be an exception. They further state that:

“if we look at the lexicon of English (or that of other Indo-European languages), we might be forgiven for thinking that there could be anything but a conventionally determined, arbitrary connection between a given word and its referent. For the vast majority of English words there is an arbitrary relationship between form and meaning.”

In fact, in our results we do not find positive MI values, on average, for English. In this section, we analyse results per language, trying to find signs of cross-linguistic non-arbitrary associations in them.

Analogously to what we did with concepts, we run permutations tests using the PMIs for the set of words in each language (i.e. sets S_l analogous to S_c in eq. (15)). Fig. 3 presents the per-language uncertainty coefficient values in a world map. There are 5189 languages in ASJP, out of those we find that only 85 have significantly positive mutual information ($p < 0.01$). Each language, though, has at most 100 values (the number of concepts), making this a hard statistical test after correcting for the multiple tests. If we relax our hypothesis testing thresholds to $p < 0.05$ (an admittedly much weaker test), then 242 languages present statistically positive MI—this suggests that, although maybe not common, form–meaning patterns are not a rare exception restricted to a small number of languages.

Concept	Tokens	Concept	Tokens	Concept	Tokens	Concept	Tokens	Concept	Tokens
blood	s	eye	i	liver	k l r t	path	d t	tree	#
bone	s u	fire	# t	louse	m n	say	#	two	r
breast	m u	fish	a s	mountain	b d g l o r	see	# e	water	#
come	# e	full	l o p t	name	# i	skin	k l p r t	we	# e i n
die	t	give	#	neck	o	star	k l o r s t u w	who	#
dog	k	horn	k r	new	a	stone	k t	you	# a i n
drink	# u	I	# a n	night	N d i l m p r t u	sun	e		
ear	e l t	knee	N b g k m o r t u	nose	N i u	tongue	d e l n r		
eat	#	leaf	a l p t	one	k t	tooth	e i		

Table 2: Concept–Token pairs with statistically significant ($p < 0.01$ after Benjamini–Hochberg corrections) mutual information in all 4 macro areas. # is the end-of-string token.

5.4 Analysis #4: Per Concept–Token Pair

We now turn to the relationship between concepts and the phones which appear in them, trying to assess specific concept–phone pairs which present positive MI. Such a positive value would indicate that concept informs on the presence of that specific phone, suggesting a non-arbitrary association between them. Similarly to before, we create sets of concept–token pairs:

$$S_{c,s} = \left\{ (\tilde{w}_t^{(n)}, \tilde{v}^{(n)}) \mid \tilde{v}^{(n)} = c \in C, \tilde{w}_t^{(n)} = s \in \Sigma \right\} \quad (16)$$

where (c, s) is the analysed concept–token pair and $\tilde{w}_t^{(n)}$ is the t^{th} token of word $\tilde{w}^{(n)}$. During this analysis, though, we focus on concept–phone pairs which had statistically significant PMIs in *all four* macroareas, following the controls introduced in Blasi et al. (2016), as a way of maximising the chances of finding true history-independent associations (under the risk of increasing the rate of false negatives). With that in mind, we split sets $S_{c,t}$ per macroarea and got the PMI values for each of them, similarly to §5.2. We threw away pairs which did not occur at least 1000 times together and ran a permutation test with 10^5 permutations for each concept–token–macroarea tuple. We note a concept–token association does not make a pair probable; the token is simply more likely to appear with the concept than would be without it.

Tab. 2 presents pairs which were significant in all macroareas ($p < 0.01$ after corrections). After analysis, we find a few interesting results. As mentioned in §1, we see an association between [l] and the concept TONGUE and between [p] and FULL, similarly to Blasi et al. (2016). We also see an association between pronouns—e.g. I, WE, YOU—and

the end-of-string [#].¹⁵ This was expected; pronouns are very frequent words in most languages, and such words are usually shorter (Zipf, 1949).

As previously found by Blasi et al. (2016), the concept BREAST has a significant association with both [m] and [u]. As they point out, these might be due to the mouth configuration of suckling babies or the sounds they produce when feeding (Jakobson, 1960; Traunmüller, 1994). We further find several other pairs which are supported by their findings: HORN–[k,r]; KNEE–[o,u,k]; LEAF–[l,p]; WE–[n]. Furthermore, a nice sanity check is that none of the negative concept–pair associations they found are present in our results.

5.5 Analysis #5: Macroareas vs Family

As a final experiment, we analyse the importance of splitting train–test sets according to macroareas (as discussed in §4.5) in order to minimise areal effects—versus simply splitting languages based on their families. Even though the list of concepts in ASJP was designed to be resistant to borrowings (and we further remove loan words from our analysis), language contacts beyond loan words could still impact results. One such example is the (potential) impact of Basque in Spanish phonology, which lost word initial /f/ in many words, e.g. *hablar*, during the late Middle Ages (see pg. 91 of Penny, 2002, for a longer discussion).

We create 4 folds, splitting them based on glotto-code language families, and use 4-fold cross-validation to get family-split results—in opposition to the macroarea-split results. Using family-splits we get an $I(W; V) = 0.020$ bits, with an uncertainty coefficient of 0.53% (averaged over the 4-

¹⁵The association of a concept with the [#] symbol means the model can more easily predict the end-of-word when conditioned on this concept. This means the length of that concept is not distributed as the average, being more predictable.

folds)—this is almost twice the overall MI found on the macroarea-splits. A Welch’s *t*-test between both runs shows family-splits have a larger MI than the macroarea results ($p < 0.01$), suggesting it is important to control for areal effects when evaluating sound–meaning associations.

6 Conclusion

In this paper we have provided a holistic assessment of form–meaning associations involving words found in the basic vocabulary in a large number of languages. In agreement with previous findings, we find that on average the meaning does not contribute substantially to the form of the words, but instead the most consistent associations were restricted to a specific subset of all of the words analysed. We find a list of 26 concepts (out of the 100 analysed) with statistically significant form–meaning associations—suggesting that cross-linguistic non-arbitrariness is not a rare exception. Finally, we also find a set of concept–phone pairs with a consistently positive relationship across the four analysed macroareas.

Ethical Considerations

This paper concerns itself with investigating cross-linguistic form–meaning associations. We see no direct ethical concerns relating to this work, as it only involves computational experiments on previously collected data.

Acknowledgments

Søren Wichmann’s research was partly funded by a subsidy from the Russian government to support the Programme of Competitive Development of Kazan Federal University, Russia. Damián E. Blasi acknowledges funding from the Branco Weiss Fellowship, administered by the ETH Zürich. Damián E. Blasi’s research was also executed within the framework of the HSE University Basic Research Program and funded by the Russian Academic Excellence Project ‘5-100’.

Appendix

A Hyper-parameter Optimisation Range

As mentioned in §4.8, we used Bayesian optimisation to tune the model’s hyper-parameters. We consider a log-uniform prior over the embedding size (from 4 to 1024), and over the size of the hidden state (32 to 1024). We also considered a uniform prior over the number of layers (1 to 4) and dropout (0 to 0.5).

References

- Kaitlyn Bankieris and Julia Simner. 2015. [What is the link between synaesthesia and sound symbolism?](#) *Cognition*, 136:186–195.
- Yoav Benjamini and Yosef Hochberg. 1995. [Controlling the false discovery rate: A practical and powerful approach to multiple testing.](#) *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):289–300.
- Benjamin K. Bergen. 2004. [The psychological reality of phonaesthemes.](#) *Language*, 80(2):290–311.
- Damián E. Blasi, Søren Wichmann, Harald Hammarström, Peter F. Stadler, and Morten H. Christiansen. 2016. [Sound–meaning association biases evidenced across thousands of languages.](#) *Proceedings of the National Academy of Sciences of the U.S.A.*, 113(39):10818–10823.
- Thomas M. Cover and Joy A. Thomas. 2012. *Elements of Information Theory*. John Wiley & Sons, New York.
- Isabelle Dautriche, Kyle Mahowald, Edward Gibson, and Steven T. Piantadosi. 2017. [Wordform similarity increases with semantic similarity: An analysis of 100 languages.](#) *Cognitive Science*, 41(8):2149–2169.
- Johannes Dellert, Thora Daneyko, Alla Münch, Alina Ladygina, Armin Buch, Natalie Clarius, Ilja Grigorjew, Mohamed Balabel, Hizniye Isabella Boga, Zalina Baysarova, Roland Mühlenbernd, Johannes Wahle, and Gerhard Jäger. 2020. [NorthEuraLex: A wide-coverage lexical database of Northern Eurasia.](#) *Language Resources and Evaluation*, 54:273–301.
- Mark Dingemanse. 2012. [Advances in the crosslinguistic study of ideophones.](#) *Language and Linguistics Compass*, 6(10):654–672.
- Mark Dingemanse. 2018. [Redrawing the margins of language: Lessons from research on ideophones.](#) *Glossa: a journal of general linguistics*, 3(1).
- Mark Dingemanse, Damián E. Blasi, Gary Lupyan, Morten H. Christiansen, and Padraic Monaghan. 2015. [Arbitrariness, iconicity, and systematicity in language.](#) *Trends in Cognitive Sciences*, 19(10):603–615.
- Matthew S. Dryer and Martin Haspelmath, editors. 2013. *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig.
- E. Dario Gutierrez, Roger Levy, and Benjamin Bergen. 2016. [Finding non-arbitrary form-meaning systematicity using string-metric learning for kernel regression.](#) In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2379–2388. Association for Computational Linguistics.

- John Haiman. 1980. [The iconicity of grammar: Iso-morphism and motivation](#). *Language*, 56(3):515–540.
- Hannah Haynie, Claire Bower, and Hannah LaPalombara. 2014. [Sound symbolism in the languages of Australia](#). *PLoS ONE*, 9(4):e92852.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Roman Jakobson. 1960. Why ‘mama’ and ‘papa’? In B. Kaplan and S. Wapner, editors, *Perspectives in Psychological Theory: Essays in Honor of Heinz Werner*, pages 538–549. International Universities Press, New York.
- Niklas Johansson and Jordan Zlatev. 2013. [Motivations for sound symbolism in spatial deixis: A typological study of 101 languages](#). *Public Journal of Semiotics*, 5(1):3–20.
- Ian Joo. 2019. [Phonosemantic biases found in Leipzig-Jakarta lists of 66 languages](#). *Linguistic Typology*, 24(1):1–12.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- David McAllester and Karl Stratos. 2020. [Formal limitations on the measurement of mutual information](#). In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 875–884. PMLR.
- Matthias Pache, Søren Wichmann, and Mikhail Zhivlov. 2016. [Words for ‘dog’ as a diagnostic of language contact in the Americas](#). In Andrea L. Berez-Kroeker, Diane M. Hintz, and Carmen Jany, editors, *Language Contact and Change in the Americas: Studies in honour of Marianne Mithun*, pages 385–410. John Benjamins, Amsterdam/Philadelphia.
- Cesare V. Parise, Katharina Knorre, and Marc O. Ernst. 2014. [Natural auditory scene statistics shapes human spatial hearing](#). *Proceedings of the National Academy of Sciences of the U.S.A.*, 111(16):6104–6108.
- Ralph John Penny. 2002. [A History of the Spanish Language](#), 2nd ed. edition. Cambridge University Press, Cambridge.
- Pamela Perniss, Robin Thompson, and Gabriella Vigliocco. 2010. [Iconicity as a general property of language: Evidence from spoken and signed languages](#). *Frontiers in Psychology*, 1:227.
- Tiago Pimentel, Arya D. McCarthy, Damian Blasi, Brian Roark, and Ryan Cotterell. 2019. [Meaning to form: Measuring systematicity as information](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1751–1764, Florence, Italy. Association for Computational Linguistics.
- Tiago Pimentel, Brian Roark, and Ryan Cotterell. 2020. [Phonotactic complexity and its trade-offs](#). *Transactions of the Association for Computational Linguistics*, 8:1–18.
- Suzanne Romaine. 1988. *Pidgin and Creole Languages*. Longman, London/New York.
- Ferdinand de Saussure. 1916. *Course in General Linguistics*. Columbia University Press, New York. English edition of June 2011, based on the 1959 translation by Wade Baskin.
- Richard Shillcock, Simon Kirby, Scott McDonald, and Chris Brew. 2001. [Filled pauses and their status in the mental lexicon](#). In *ISCA Tutorial and Research Workshop (ITRW) on Disfluency in Spontaneous Speech*.
- Jasper Snoek, Hugo Larochelle, and Ryan Prescott Adams. 2012. [Practical Bayesian optimization of machine learning algorithms](#). In *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc.
- Morris Swadesh. 1955. [Towards greater accuracy in lexicostatistic dating](#). *International Journal of American Linguistics*, 21(2):121–137.
- Hartmut Trautmüller. 1994. Sound symbolism in deictic words. In Hans Auli and Peter af Trampe, editors, *Tongues and Texts Unlimited. Studies in Honour of Tore Jansson on the Occasion of his Sixtieth Anniversary*, pages 213–234. Dept. of Classical Languages, Stockholm University, Stockholm.
- Søren Wichmann, Eric W. Holman, and Cecil H. Brown. 2010. [Sound symbolism in basic vocabulary](#). *Entropy*, 12(4):844–858.
- Søren Wichmann, Eric W. Holman, and Cecil H. Brown. 2020. [The ASJP database \(version 19\)](#). Accessed on April 2, 2020.
- George Kingsley Zipf. 1949. *Human Behavior and the Principle of Least Effort*. Addison-Wesley Press, Cambridge, Mass.