



Efficient compression in color naming and its evolution

Noga Zaslavsky^{a,b,1}, Charles Kemp^{c,2}, Terry Regier^{b,d}, and Naftali Tishby^{a,e}

^aEdmond and Lily Safra Center for Brain Sciences, The Hebrew University, Jerusalem 9190401, Israel; ^bDepartment of Linguistics, University of California, Berkeley, CA 94720; ^cDepartment of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213; ^dCognitive Science Program, University of California, Berkeley, CA 94720; and ^eThe Benin School of Computer Science and Engineering, The Hebrew University, Jerusalem 9190401, Israel

Edited by James L. McClelland, Stanford University, Stanford, CA, and approved June 18, 2018 (received for review January 11, 2018)

We derive a principled information-theoretic account of cross-language semantic variation. Specifically, we argue that languages efficiently compress ideas into words by optimizing the information bottleneck (IB) trade-off between the complexity and accuracy of the lexicon. We test this proposal in the domain of color naming and show that (i) color-naming systems across languages achieve near-optimal compression; (ii) small changes in a single trade-off parameter account to a large extent for observed cross-language variation; (iii) efficient IB color-naming systems exhibit soft rather than hard category boundaries and often leave large regions of color space inconsistently named, both of which phenomena are found empirically; and (iv) these IB systems evolve through a sequence of structural phase transitions, in a single process that captures key ideas associated with different accounts of color category evolution. These results suggest that a drive for information-theoretic efficiency may shape color-naming systems across languages. This principle is not specific to color, and so it may also apply to cross-language variation in other semantic domains.

information theory | semantic typology | color naming | categories | language evolution

Languages package ideas into words in different ways. For example, English has separate terms for “hand” and “arm,” “wood” and “tree,” and “air” and “wind,” but other languages have single terms for each pair. At the same time, there are universal tendencies in word meanings, such that similar or identical meanings often appear in unrelated languages. A major question is how to account for such semantic universals and variation of the lexicon in a principled and unified way.

One approach to this question proposes that word meanings may reflect adaptation to pressure for efficient communication—that is, communication that is precise yet requires only minimal cognitive resources. On this view, cross-language variation in semantic categories may reflect different solutions to this problem, while semantic commonalities across unrelated languages may reflect independent routes to the same highly efficient solution. This proposal, focused on linguistic meaning, echoes the invocation of efficient communication to also explain other aspects of language (e.g., refs. 1–4).

Color is a semantic domain that has been approached in this spirit. Recent work has relied on the notion of the “informativeness” of word meaning, has often cast that notion in terms borrowed from information theory, and has accounted for several aspects of color naming across languages on that basis (5–10). Of particular relevance to our present focus, Regier, Kemp, and Kay (ref. 8, henceforth RKK) found that theoretically efficient categorical partitions of color space broadly matched major patterns of color naming seen across languages—suggesting that pressure for efficiency may indeed help to explain why languages categorize color as they do.

However, a fundamental issue has been left largely undressed: how a drive for efficiency may relate to accounts of color category evolution. Berlin and Kay (11) proposed an evolutionary sequence by which new terms refine existing partitions of color space in a discrete order: first dark vs. light, then red, then green and yellow, then blue, followed by other basic color

categories. RKK’s efficient theoretical color-naming systems correspond roughly to the early stages of the Berlin and Kay sequence, but they leave the transitions between stages unexamined and are based on the false (9, 12, 13) simplifying assumption that color-naming systems are hard partitions of color space. In actuality, color categories are a canonical instance of soft categories with graded membership, and it has been argued (12, 13) that such categories may emerge gradually in parts of color space that were previously inconsistently named. Such soft category boundaries introduce uncertainty and therefore might be expected to impede efficient communication (9). Thus, it remains an open question whether a hypothesized drive for efficiency can explain not just discrete stages of color category evolution, but also how systems evolve continuously from one stage to the next, and why inconsistent naming patterns are sometimes observed.

Here, we argue that a drive for information-theoretic efficiency provides a unified formal explanation of these phenomena. Specifically, we argue that languages efficiently compress ideas into words by optimizing the trade-off between the complexity and accuracy of the lexicon according to the information bottleneck (IB) principle (14), an independently motivated formal principle with broad scope (15–17), which is closely related (ref. 18 and *SI Appendix*, section 1.3) to rate distortion theory (19). We support this claim by showing that cross-language variation in color naming can be explained in IB terms. Our findings suggest that languages may evolve through a trajectory of efficient solutions in a single process that synthesizes, in formal terms, key ideas from Berlin and Kay’s (11) theory and from more continuous accounts (12, 13) of color category evolution. We also show that soft categories and inconsistent naming can be information-theoretically efficient.

Significance

Semantic typology documents and explains how languages vary in their structuring of meaning. Information theory provides a formal model of communication that includes a precise definition of efficient compression. We show that color-naming systems across languages achieve near-optimal compression and that this principle explains much of the variation across languages. These findings suggest a possible process for color category evolution that synthesizes continuous and discrete aspects of previous accounts. The generality of this principle suggests that it may also apply to other semantic domains.

Author contributions: N.Z., C.K., T.R., and N.T. designed research; N.Z. performed research; N.Z. and N.T. contributed new reagents/analytic tools; N.Z. analyzed data; and N.Z. and T.R. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Published under the [PNAS license](#).

¹To whom correspondence should be addressed. Email: noga.zaslavsky@mail.huji.ac.il.

²Present address: School of Psychological Sciences, The University of Melbourne, Parkville, Victoria 3010, Australia.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1800521115/-DCSupplemental.

Published online July 18, 2018.

Our work focuses on data compression, in contrast with work that views language in information-theoretic terms but focuses instead on channel capacity (2–4, 7, 20), including work on language evolution (21). Our work also further (e.g., refs. 7 and 22) links information theory to the study of meaning, a connection that has been contested since Shannon's (23) foundational work. IB has previously been used to find semantically meaningful clusters of words (ref. 15; see also ref. 22), but has not previously been used to account for word meanings as we do here.

Communication Model

To define our hypothesis precisely, we first formulate a basic communication scenario involving a speaker and a listener. This formulation is based on Shannon's classical communication model (23), but specifically concerns messages that are represented as distributions over the environment (Fig. 1). We represent the environment, or universe, as a set of objects \mathcal{U} . The state of the environment can be any object $u \in \mathcal{U}$, and we let U be a random variable that represents a possible state. We define a meaning to be a distribution $m(u)$ over \mathcal{U} and assume the existence of a cognitive source that generates intended meanings for the speaker. This source is defined by a distribution $p(m)$ over a set of meanings, \mathcal{M} , that the speaker can represent. Each meaning reflects a subjective belief about the state of the environment. If the speaker's intention is $m \in \mathcal{M}$, this indicates that she wishes to communicate her belief that $U \sim m(u)$. We consider a color communication model in which \mathcal{U} is restricted to colors and each $m \in \mathcal{M}$ is a distribution over colors.

The speaker communicates m by producing a word w , taken from a shared lexicon of size K . The speaker selects words according to a naming policy $q(w|m)$. This distribution is a stochastic encoder that compresses meanings into words. Because we focus on the uncertainty involved in compressing meanings into words, rather than the uncertainty involved in transmission, we assume an idealized noiseless channel that conveys its input unaltered as its output. This channel may have a limited capacity, which imposes a constraint on the available lexicon size. In this case, the listener receives w and interprets it as meaning \hat{m} based on her interpretation policy $q(\hat{m}|w)$, which is a decoder. We focus on the efficiency of the encoder and therefore assume an optimal Bayesian listener with respect to the speaker (see *SI Appendix, section 1.2* for derivation), who interprets every word w deterministically as meaning




Fig. 1. (A) Shannon's (23) communication model. In our instantiation of this model, the source message M and its reconstruction \hat{M} are distributions over objects in the universe \mathcal{U} . We refer to these messages as meanings. M is compressed into a code, or word, W . We assume that W is transmitted over an idealized noiseless channel and that the reconstruction \hat{M} of the source message is based on W . The accuracy of communication is determined by comparing M and \hat{M} , and the complexity of the lexicon is determined by the mapping from M to W . (B) Color communication example, where \mathcal{U} is a set of colors, shown for simplicity along a single dimension. A specific meaning m is drawn from $p(m)$. The speaker communicates m by uttering the word "blue," and the listener interprets blue as meaning \hat{m} .

$$\hat{m}_w(u) = \sum_{m \in \mathcal{M}} q(m|w)m(u), \quad [1]$$

where $q(m|w)$ is obtained by applying Bayes' rule with respect to $q(w|m)$ and $p(m)$.

In this model, different color-naming systems correspond to different encoders, and our goal is to test the hypothesis that encoders corresponding to color-naming systems found in the world's languages are information-theoretically efficient. We next describe the elements of this model in further detail.

Encoders. Our primary data source for empirically estimating encoders was the World Color Survey (WCS), which contains color-naming data from 110 languages of nonindustrialized societies (24). Native speakers of each language provided names for the 330 color chips shown in Fig. 2, *Upper*. We also analyzed color-naming data from English, collected relative to the same stimulus array (25). We assumed that each color chip c is associated with a unique meaning m_c and therefore estimated an encoder $q_l(w|m_c)$ for each language l from the empirical distribution of word w given chip c (see data rows in Fig. 4 for examples). Each such encoder corresponds to a representative speaker for language l , obtained by averaging naming responses over speakers.

Meaning Space. In our formulation, colors are mentally represented as distributions. Following previous work (6, 8), we ground these distributions in an established model of human color perception by representing colors in 3D CIELAB space (Fig. 2, *Lower*) in which Euclidean distance between nearby colors is correlated with perceptual difference. We define the meaning associated with chip c to be an isotropic Gaussian centered at c , namely $m_c(u) \propto \exp(-\frac{1}{2\sigma^2}\|u - c\|^2)$. m_c reflects the speaker's subjective belief over colors that is invoked by chip c , and the scale of these Gaussians reflects her level of perceptual uncertainty. We take $\sigma^2 = 64$, which corresponds to a distance over which two colors can be comfortably distinguished (*SI Appendix, section 6.3*).

Cognitive Source. The cognitive source $p(m)$ specifies how often different meanings m must be communicated by a speaker. In principle, different cultures may have different communicative needs (8); we leave such language-specific analysis for future work and instead consider a universal source for all languages. Previous studies have used the uniform distribution for this purpose (8, 10); however, it seems unlikely that all colors are in fact equally frequent in natural communication. We therefore consider an alternative approach, while retaining the uniform distribution as a baseline. Specifically, we focus on a source that is derived from the notion of least informative (LI) priors (*Materials and Methods*), a data-driven approach that requires minimal assumptions. This approach also accounts for the data better than another approach based on image statistics (*SI Appendix, section 7.2*).

Bounds on Semantic Efficiency

From an information-theoretic perspective, an optimal encoder minimizes complexity by compressing the intended message M as much as possible, while maximizing the accuracy of its interpretation \hat{M} (Fig. 1A). In general, this principle is formalized by rate distortion theory (RDT) (19). In the special case in which messages are distributions, the IB principle (14) provides a natural formalization. In IB, as in RDT (*SI Appendix, section 1.3*), the complexity of a lexicon is measured by the number of bits of information that are required for representing the intended meaning. In our formulation the speaker represents her intended




Fig. 2. (Upper) The WCS stimulus palette. Columns correspond to equally spaced Munsell hues. Rows correspond to equally spaced lightness values. Each stimulus is at the maximum available saturation for that hue/lightness combination. (Lower) These colors are irregularly distributed in 3D CIELAB color space.

meaning M by W , using an encoder $q(w|m)$, and thus the complexity is given by the information rate

$$I_q(M; W) = \sum_{m,w} p(m) q(w|m) \log \frac{q(w|m)}{q(w)}, \quad [2]$$

where $q(w) = \sum_{m \in \mathcal{M}} p(m) q(w|m)$. Minimal complexity, i.e., $I_q(M; W) = 0$, can be achieved if the speaker uses a single word to describe all her intended meanings. However, in this case the listener will not have any information about the speaker's intended meaning. To enable useful communication, W must contain some information about M ; i.e., the complexity $I_q(M; W)$ must be greater than zero.

The accuracy of a lexicon is inversely related to the cost of a misinterpreted or distorted meaning. While RDT allows an arbitrary distortion measure, IB considers specifically the Kullback–Leibler (KL) divergence,

$$D[m||\hat{m}] = \sum_{u \in \mathcal{U}} m(u) \log \frac{m(u)}{\hat{m}(u)}, \quad [3]$$

which is a natural distortion measure between distributions. [For a general justification of the KL divergence see ref. 26, and in the context of IB see ref. 18.] Note that this quantity is 0 if and only if the listener's interpretation is accurate; namely, $\hat{m} \equiv m$. The distortion between the speaker and the ideal listener is the expected KL divergence,

$$\mathbb{E}_q[D[M||\hat{M}]] = \sum_{m,w} p(m) q(w|m) D[m||\hat{m}_w]. \quad [4]$$

In this case, the accuracy of the lexicon is directly related to Shannon's mutual information,

$$\mathbb{E}_q[D[M||\hat{M}]] = I(M; U) - I_q(W; U). \quad [5]$$

Since $I(M; U)$ is independent of $q(w|m)$, minimizing distortion is equivalent to maximizing the informativeness, or accuracy, of the lexicon, quantified by $I_q(W; U)$. This means that mutual information appears in our setting as a natural measure both for complexity and for semantic informativeness.

If the speaker and the listener are unwilling to tolerate any information loss, the speaker must assign a unique word to each meaning, which requires maximal complexity. However, between the two extremes of minimal complexity and maximal accuracy, an optimal trade-off between these two competing needs can be obtained by minimizing the IB objective function,

$$\mathcal{F}_\beta[q(w|m)] = I_q(M; W) - \beta I_q(W; U), \quad [6]$$

where $\beta \geq 1$ is the trade-off parameter. Every language l , defined by an encoder $q_l(w|m)$, attains a certain level of complexity and a certain level of accuracy. These two quantities can be plotted against each other. Fig. 3 shows this information plane for the present color communication model. The maximal accuracy that a language l can achieve, given its complexity, is bounded from above. Similarly, the minimal complexity that l can achieve given its accuracy is bounded from below. These bounds are given by the complexity and accuracy of the set of hypothetical IB languages that attain the minimum of Eq. 6 for different values of β . The IB curve is the theoretical limit defined by these optimal languages, and all trade-offs above this curve are unachievable.

Predictions

Near-Optimal Trade-offs. Our hypothesis is that languages evolve under pressure for efficient compression, as defined by IB, which implies that they are pressured to minimize \mathcal{F}_β for some value of β . If our hypothesis is true, then for each language l there should be at least one value, β_l , for which that language is close to the optimal $\mathcal{F}_{\beta_l}^*$. If we are able to find a good candidate β_l for every language, this would support our hypothesis, because such an outcome would be unlikely given systems that evolved independently of \mathcal{F}_β . A natural choice for fitting β_l is the value of β that minimizes $\Delta\mathcal{F}_\beta = \mathcal{F}_\beta[q_l] - \mathcal{F}_\beta^*$. We measure the efficiency loss, or deviation from optimality, of language l by $\varepsilon_l = \frac{1}{\beta_l} \Delta\mathcal{F}_{\beta_l}$.

Structure of Semantic Categories. Previous work (e.g., ref. 8) has sometimes summarized color-naming responses across multiple speakers of the same language by recording the modal naming response for each chip, resulting in a hard categorical partition of the stimulus array, called a mode map (e.g., Fig. 44). However, IB predicts that if some information loss is allowed, i.e., $\beta < \infty$, then an efficient encoder would induce soft rather than hard categories. This follows from the structure of the IB optima (14), given by

$$q_\beta(w|m) \propto q_\beta(w) \exp(-\beta D[m||\hat{m}_w]), \quad [7]$$

which is satisfied self-consistently with Eq. 1 and with the marginal $q_\beta(w)$. We therefore evaluate how well our model accounts for mode maps, but more importantly we also evaluate how well it accounts for the full color-naming distribution across




Fig. 3. Color-naming systems across languages (blue circles) achieve near-optimal compression. The theoretical limit is defined by the IB curve (black). A total of 93% of the languages achieve better trade-offs than any of their hypothetical variants (gray circles). Small light-blue Xs mark the languages in Fig. 4, which are ordered by complexity.

each language were estimated from at least five responses per chip (*SI Appendix, section 4.1*).

LI Source. A source distribution can be defined from a prior over colors by setting $p(m_c) = p(c)$. For each language l , we constructed a LI source $p_l(c)$ by maximizing the entropy of c while also minimizing the expected surprisal of c given a color term w in that language (see *SI Appendix, section 2* for more details). We obtained a single LI source by averaging the language-specific priors.

IB Curve. For each value of β the IB solution is evaluated using the IB method (14). IB is a nonconvex problem, and therefore only convergence to local optima is guaranteed. To mitigate this problem we fix $K = 330$ and use the method of reverse deterministic annealing to evaluate the IB curve (*SI Appendix, section 1.4*).

Dissimilarity Between Naming Distributions. Assume two speakers that independently describe m by $W_1 \sim q_1(w_1|m)$ and $W_2 \sim q_2(w_2|m)$. We define the dissimilarity between q_1 and q_2 by

$$\text{gNID}(W_1, W_2) = 1 - \frac{I(W_1; W_2)}{\max\{I(W_1; W'_1), I(W_2; W'_2)\}}, \quad [8]$$

where W'_i corresponds to another independent speaker that uses q_i . If q_1 and q_2 are deterministic, i.e., they induce hard partitions, then gNID reduces to NID (*SI Appendix, section 3* for more details).

Relative Accuracy. We define the informativeness of a word w by

$$I_q(w) = D[\hat{m}_w \| m_0], \quad [9]$$

where $m_0(u) = \sum_m p(m)m(u)$ is the prior over u before knowing w . Note that the accuracy of a language can be written as $I_q(W; U) = \sum_w q(w)I_q(w)$, and therefore we define the relative accuracy of w (y axis in Fig. 5) by $I_q(w) - I_q(W; U)$.

ACKNOWLEDGMENTS. We thank Daniel Reichman for facilitating the initial stages of our collaboration, Delwin Lindsey and Angela Brown for kindly sharing their English color-naming data with us, Bevil Conway and Ted Gibson for kindly sharing their color-salience data with us, and Paul Kay for useful discussions. This study was supported by the Gatsby Charitable Foundation (N.T.), IBM PhD Fellowship Award (to N.Z.), and Defense Threat Reduction Agency (DTRA) Award HDTRA11710042 (to T.R.). Part of this work was done while N.Z. and N.T. were visiting the Simons Institute for the Theory of Computing at University of California, Berkeley.

- Ferrer i Cancho R, Solé RV (2003) Least effort and the origins of scaling in human language. *Proc Natl Acad Sci USA* 100:788–791.
- Levy RP, Jaeger TF (2007) Speakers optimize information density through syntactic reduction. *Advances in Neural Information Processing Systems*, eds Schölkopf B, Platt JC, Hoffman T (MIT Press, Cambridge, MA), Vol 19, pp 849–856.
- Piantadosi ST, Tily H, Gibson E (2011) Word lengths are optimized for efficient communication. *Proc Natl Acad Sci USA* 108:3526–3529.
- Gibson E, et al. (2013) A noisy-channel account of crosslinguistic word-order variation. *Psychol Sci* 24:1079–1088.
- Jameson K, D'Andrade RG (1997) It's not really red, green, yellow, blue: An inquiry into perceptual color space. *Color Categories in Thought and Language*, eds Hardin CL, Maffi L (Cambridge Univ Press, Cambridge, UK), pp 295–319.
- Regier T, Kay P, Khetarpal N (2007) Color naming reflects optimal partitions of color space. *Proc Natl Acad Sci USA* 104:1436–1441.
- Baddeley R, Attewell D (2009) The relationship between language and the environment: Information theory shows why we have only three lightness terms. *Psychol Sci* 20:1100–1107.
- Regier T, Kemp C, Kay P (2015) Word meanings across languages support efficient communication. *The Handbook of Language Emergence*, eds MacWhinney B, O'Grady W (Wiley-Blackwell, Hoboken, NJ), pp 237–263.
- Lindsey DT, Brown AM, Brainard DH, Apicella CL (2015) Hunter-gatherer color naming provides new insight into the evolution of color terms. *Curr Biol* 25:2441–2446.
- Gibson E, et al. (2017) Color naming across languages reflects color use. *Proc Natl Acad Sci USA* 114:10785–10790.
- Berlin B, Kay P (1969) *Basic Color Terms: Their Universality and Evolution* (Univ of California Press, Berkeley).
- MacLaury RE (1997) *Color and Cognition in Mesoamerica: Constructing Categories as Vantages* (Univ of Texas Press, Austin, TX).
- Levinson SC (2000) Yéll Dnye and the theory of basic color terms. *J Linguistic Anthropol* 10:3–55.
- Tishby N, Pereira FC, Bialek W (1999) The information bottleneck method. *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing*, eds Hajek B, Sreenivas RS (Univ of Illinois, Urbana, IL), pp 368–377.
- Slonim N (2002) The information bottleneck: Theory and applications. PhD thesis (Hebrew Univ of Jerusalem, Jerusalem).
- Shamir O, Sabato S, Tishby N (2010) Learning and generalization with the information bottleneck. *Theor Comput Sci* 411:2696–2711.
- Palmer SE, Marre O, Berry MJ, Bialek W (2015) Predictive information in a sensory population. *Proc Natl Acad Sci USA* 112:6908–6913.
- Harremoës P, Tishby N (2007) The information bottleneck revisited or how to choose a good distortion measure. *IEEE International Symposium on Information Theory*. Available at <https://ieeexplore.ieee.org/document/4557285/>. Accessed July 10, 2018.
- Shannon CE (1959) Coding theorems for a discrete source with a fidelity criterion. *IRE Natl Conv Rec* 4:142–163.
- Jaeger TF (2010) Redundancy and reduction: Speakers manage syntactic information density. *Cogn Psychol* 61:23–62.
- Plotkin JB, Nowak MA (2000) Language evolution and information theory. *J Theor Biol* 205:147–159.
- Pereira F, Tishby N, Lee L (1993) Distributional clustering of English words. *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, ed Schubert LK (Association for Computational Linguistics, Stroudsburg, PA), pp 183–190.
- Shannon C (1948) A mathematical theory of communication. *Bell Syst Tech J* 27:623–656.
- Cook RS, Kay P, Regier T (2005) The World Color Survey database: History and use. *Handbook of Categorization in Cognitive Science*, eds Cohen H, Lefebvre C (Elsevier, Amsterdam), pp 223–242.
- Lindsey DT, Brown AM (2014) The color lexicon of American English. *J Vis* 14:17.
- Csiszár I, Shields P (2004) Information theory and statistics: A tutorial. *Found Trends Commun Inf Theor* 1:417–528.
- Vinh NX, Epps J, Bailey J (2010) Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *JMLR* 11:2837–2854.
- Rose K (1998) Deterministic annealing for clustering, compression, classification, regression, and related optimization problems. *Proceedings of the IEEE* 86:2210–2239.
- Lindsey DT, Brown AM (2004) Color naming and color consensus: “Blue” is special. *J Vis* 4:55.