Noise resistance in communication: Quantifying uniformity and optimality

Christine Cuskley^{*1}, Rachael Bailes¹, and Joel Wallenberg¹

¹*Centre for Behaviour and Evolution, Linguistics Section, Newcastle University, Newcastle upon Tyne, UK*

^{*}*Corresponding author: christine.cuskley@ncl.ac.uk*

Abstract

Over the past few decades, several lines of research have investigated how information is distributed in language. In particular, research has suggested that speakers distribute the information in linguistic utterances as evenly as possible, in order to make the utterance more robust against noise for the hearer. Several studies have shown evidence for this hypothesis in limited linguistic domains, showing that speakers manipulate acoustic and syntactic features by using redundancy to avoid drastic spikes or troughs in information content. However, this previous work doesn't consider information density across entire utterances, and only rarely has the proposed function of noise resistance been directly explored in detail. Here, we introduce a new descriptive statistic (*DORM*) that quantifies the uniformity of information across an entire utterance. We present this alongside an algorithm (*UIDO*) that allows us to quantify the contribution made to an utterance's uniformity by the order of its elements, independent of its level of redundancy. Using a simple simulation with data from an English corpus, we show that utterances whose order results in more uniform distributions of information are, in fact, more robust against noise.

Word Count: 6433

Keywords: Information Theory, language, noise, redundancy.

This is a preprint of the publication listed below. Please use this citation when referencing the work:

Cuskley, C., Bailes, R., & Wallenberg, J. (2021). Noise resistance in communication: Quantifying uniformity and optimality. *Cognition*, 214, 104754. <https://doi.org/10.1016/j.cognition.2021.104754>.

1 Introduction

For any event or situation with a discrete number of possible outcomes, Shannon (1948) introduced *entropy* (H) to quantify the average amount of information an observer could expect to receive about an event through some signalling process. Shannon went on to demonstrate that communication only takes place when some signal removes uncertainty in the receiver.

The receiver, or **observer**, crucially begins with a prior set of possible outcomes in mind, and then hopes to observe some **signal** about the event that narrows the set of possible outcomes down to one specific outcome (or a few; any narrowing of the set is a reduction in her or his uncertainty about the event). For an event with n discrete outcomes, which have probabilities $p_1 \dots p_n$, the equation gives the average amount of information about the event in “bits” (which is also the number of binary circuits it would take to encode the outcome).

$$H = \sum_1^n p_i \log_2 \frac{1}{p_i} \quad (1)$$

A core concept in information theory is Shannon’s (1948) noisy channel coding theorem: (a) the speed with which a signal can communicate information depends on the quantity of noise in the channel which must transmit the signal, and (b) for any amount of noise, information can theoretically be transmitted with no errors by a message in some code, though possibly a very slow one (i.e. a signal with a large amount redundancy). In contexts with more interference, signals must be longer in order to transmit the same amount of information. For example, FM radio carries more information per time unit than AM radio, but AM radio is subject to less fade over long distances for exactly this reason (Walker, 2017, p.50).

The implications of information theory for language have been noted for over half a century (Jakobson, 1961), and ongoing efforts to apply information theoretic approaches to natural language continue to concentrate at the nexus of computational and cognitive science, mathematics, and linguistics (Debowski & Bentz, 2020). In terms of noise specifically, a family of hypotheses has suggested that the distribution of information across an utterance provides resistance against disruption caused by noise (henceforth referred to as *information density accounts*; Fenk & Fenk, 1980; Aylett & Turk, 2004; Turk, 2010; Jaeger, 2010). This was first articulated by Fenk & Fenk (1980, p.402) (English translation by C. Bentz, as appears in Ferrer-i-Cancho (2017, p.46)):

“A communication system, which is supposed to deliver messages without loss, should not only be required to have a certain average level of redundancy (not

34 *exceeding the short term memory capacity), but also, that the information is*
35 *distributed as uniformly as possible across small time spans.”*

36 The crucial components of this idea are that (i) linguistic utterances can be structured
37 in a way that achieves resistance to noise, by means of (ii) signal redundancy, such as
38 synonymous lexical items or reduplication of phonetic segments (both of which lengthen
39 the signal), *and additionally* (iii) modulation of the information distribution without
40 adding redundancy, such as by arranging the order of elements so that information is
41 portioned out as evenly as possible across the whole utterance (i.e., over time).

42 The ultimate aim of deploying these strategies is to achieve some constant flow of
43 linguistic information across the utterance (Fenk-Oczlon, 2001). Following on from the
44 work of Fenk & Fenk and Fenk-Oczlon, a family of hypotheses has since been developed
45 that is unified by the idea of a speaker bias toward a constant flow of linguistic infor-
46 mation (Genzel & Charniak, 2002; Aylett & Turk, 2004; Levy & Jaeger, 2007, see also
47 Ferrer-i-Cancho, 2017 for further discussion). The common principle is that the linguistic
48 encoding of information is manipulated in order to achieve uniformity. While redundancy
49 and ordering are conceptually distinct means of achieving this, they can easily be con-
50 flated, partly because they are intricately related. Reduplicating or ‘stretching’ the same
51 information across multiple linguistic elements (signal redundancy) can be applied where
52 peaks of information occur, as a means of flattening the attendant part of the information
53 distribution. This does not, however, address the role played by the *order* of a given set
54 of items in an utterance’s distribution of information. Achieving noise resistance through
55 reduplication (i.e. redundancy) follows directly from Shannon’s (1948) noisy channel
56 coding theorem. The modulation of an information distribution through the ordering
57 of constituent items, however, is not directly outlined in the theorem. The *position* of
58 information peaks and troughs across the entire distribution of an utterance may con-
59 stitute a distinct, but important, strategy for noise resistance. If information is evenly
60 distributed across an utterance, any noise event which disrupts the signal is unlikely to
61 cause a complete communication failure. On the other hand, if crucial information is
62 clustered at a particular point in the utterance, noise is more likely to completely disrupt
63 communication.

64 Given this, we should expect that sentences where low probability words and phrases
65 are evenly distributed throughout are preferred to those where low probability items clus-
66 ter together. A description of this ordering at the level of the whole utterance has so far
67 remained largely unexplored and unaddressed. Rather, existing studies on information
68 distribution uniformity or ‘smoothing’ have been principally concerned with the strate-
69 gic deployment of redundancy on limited portions of utterances (Aylett & Turk, 2004;
70 Pluymaekers, Ernestus, & Baayen, 2005; Levy & Jaeger, 2007; Frank & Jaeger, 2008;
71 Jaeger, 2010; Bergen & Goodman, 2015; Zhan & Levy, 2018), and for the most part do

not distinguish this from modulating the order of peaks and troughs across the whole distribution.

A simple measure of uniformity for whole utterances has many potential applications in the study of information distributions in language. But it also affords an opportunity for greater theoretical clarity on two counts. First, a measure that is particular to an ordering strategy for noise resistance allows us to disentangle redundancy and ordering as strategies for noise resistance. Second, some characteristics of smoothing through redundancy are consistent with other accounts of linguistic planning (such as availability-based production models), and a uniformity measure particular to distribution ordering may help to differentiate between these and information uniformity accounts. The next section demonstrates these points in more detail.

1.1 Disentangling redundancy and order

Over the past decade and a half, several lines of research have demonstrated the use of redundancy in language as a means of flattening peaks of information. Information content in bits¹ is defined as the $\log_2 \frac{1}{p_i}$ portion of equation 1 above, where p_i could be the probability of any linguistic unit (e.g., syllable, word, phrase), usually operationalized as its normalized frequency in a dataset. Logging the inverse probability means low probability units have high information content, while high probability units have low information content (e.g., the lower frequency word “ocelot” is more informative than the higher frequency word “cat”).

One such study by Aylett and Turk (2004) demonstrates this as evidence for an information density account, which they term the *smooth signalling redundancy hypothesis* (see also Turk, 2010). By analysing naturalistic speech in the MapTask corpus (Anderson et al., 1991), Aylett and Turk (2004) showed that speakers use prosody and duration to lengthen particular segments, thereby achieving more uniform distributions of information. Bergen and Goodman (2015) also show that redundancy is strategically deployed in real-time: when noise is deliberately introduced into a communication channel, speakers prosodically emphasize and lengthen words with higher information content. Similarly, Levy and Jaeger (2007) show that contractions (e.g., *you are* vs *you’re*) and deletions (e.g., *I know that she is*, vs *I know she is*) - which pack the information of two words into the time of one word - are avoided if they create large informational peaks in a sentence (see also Frank & Jaeger, 2008).

A corollary of stretching units with high information content to improve noise resistance is the compression of units with low information content: less informative units

¹As a reviewer points out, the measure would have the same properties for our purposes regardless of the base of the logarithm. We choose bits as our unit of information, using $\log_2$, simply because that unit has been very commonly used since Shannon’s work, but units such as e.g. the hartley, using $\log_{10}$ (Hartley, 1928), are equally valid.

are more likely to be reduced. The phenomenon of compressed length for high frequency (low information content) words has long been known in the form of Zipf’s Law of Abbriviation (Zipf, 1935), and may be a particular hallmark of the referential functionality of language (Ferrer-i-Cancho & Solé, 2003). This is underpinned by Zipf’s accompanying premise, the Principle of Least Effort, which recognises that cognitive constraints on speakers and hearers introduce a pressure for reducing effort. For our purposes, it is notable that high-frequency, low-information words are where language users strategically deploy reduction while also maintaining communicative function (Kanwal, Smith, Culbertson, & Kirby, 2017). This also applies to purely phonetic forms of reduction in language; more frequent words (i.e., words with lower information content) are more likely to be reduced in duration during speech (Barinova, 1971; Fenk-Oczlon, 1989, 2001; Pluymaekers et al., 2005).

While the strategic use of redundancy is consistent with accounts which suggest its function is to smooth information distributions for noise resistance, it’s also consistent with other accounts of linguistic planning. Availability-based production (ABP) accounts hold that utterances are constructed according to the linguistic items that are most cognitively accessible to the speaker as production unfolds (Bock, 1982, 1987; Levelt & Maasen, 1981; Ferreira, 1996; Ferreira & Dell, 2000), rather than being constrained by factors affecting comprehension (e.g., noise). To investigate whether utterances are planned according to speaker constraints or to reduce ambiguity for the hearer, Ferreira and Dell (2000) compared instances of complementizer deletion before the embedded subjects *you* vs *I* (e.g. *You know (that) I missed practice* vs *You know (that) you missed practice*) in a series of experimental tasks. In the example *You know you missed practice*, the embedded subject *you* creates a temporary syntactic ambiguity for the hearer as to whether it begins a relative clause or is a direct object, whereas *I* in the same position does not create this temporary ambiguity (because *I* cannot be an object). For cases like this, speakers more often omitted *that* before the repeated subject *you* than the less ambiguous subject *I* in the embedded clause. They argue this shows that the cognitive pressure to use primed subjects as soon as possible induces speakers to forego a complementizer despite the ambiguity this may cause for the hearer. However, it is notable that *you* is more ambiguous precisely because it has lower information content than *I*. In other words, an information density account also predicts that speakers will be more likely to moderate the informational peak that accompanies *I* (compared with *you*) by deploying *that* in order to add redundancy and stretch the signal.

In another test of an information density account, Jaeger (2010) analyzed complementizer deletion in a corpus of spontaneous speech, examining conditions which give rise to *that* deletion (e.g., *I know he’s hungry* vs *I know that he’s hungry*). This work found that high-information embedded subjects and markers of disfluency in preceding speech both predicted *that* insertion. As with earlier results, this is consistent with both

ABP and information density accounts. While Jaeger (2010) was able to use a model to estimate the relative contribution of different constraints on linguistic planning (e.g., ambiguity avoidance vs information density based on redundancy), this approach cannot distinguish between the two different strategies for manipulating information density originally identified by Fenk and Fenk (1980).

More recent work has attempted to explicitly adjudicate between ABP and information density accounts of linguistic planning. Zhan and Levy (2018) analysed the use of “general” and “specific” numeral classifiers in a corpus of Mandarin online news articles. In Mandarin Chinese, a numeral classifier must be produced after its associated numeral and before the noun phrase it applies to. The predictions of ABP followed from three premises particular to the case of numeral classifiers: (i) noun lemmas must be accessed first in order to select the appropriate specific classifier for that noun, (ii) unpredictable (less frequent, higher information content) noun lemmas are harder to access, and (iii) the general classifier (which can be used with most nouns) is always cognitively available.

ABP accounts predict that general classifiers will be used more often with less predictable nouns, as this buys the speaker more time to access the noun lemma. On the other hand, Zhan and Levy (2018) argue that information density accounts predict specific classifiers will be used more often with less predictable (higher information content) nouns. However, this prediction arises, at least in part, from representing distribution uniformity only in terms of redundancy. On this view, a specific classifier reduces the conditional probability of the following noun over the time-course of production, therefore ‘stretching’ a given peak of information across two units. We also note that implicit in their account is that the relevant information content is conditional entropy given preceding word (left-to-right bigram), rather than e.g. unconditional entropy or entropy conditioned on bidirectional context (cf. a similar approach in Xu & Reitter, 2016, 2018 who use preceding two words, as do Levy & Jaeger, 2007, Maurits, Navarro, & Perfors, 2010, Jaeger, 2010; Genzel & Charniak, 2002 use three preceding words). This is certainly reasonable, but not the only possible perspective (cf. Temperley & Gildea, 2015 for an example of entropy conditioned left-to-right but at a greater distance).

On a view that is principally concerned with the modulating the distribution of peaks throughout an entire utterance, we may expect that preceding an unpredictable (i.e., information-dense) item with a *less* informative item such as a general classifier would serve to ‘space out’ the peaks of information across the utterance (particularly if we also consider unconditional word entropies). Note here that the envisioned optimum in the latter case is not the flattest possible (highest redundancy) distribution, but one where neither peaks nor troughs of information are clustered together across the utterance. In this way, the lack of evidence supporting the information density account reported in Zhan and Levy (2018) may result in part from only considering redundancy-based strategies for modulating informational distributions. Conceiving of the choice between

classifiers as a choice between more or less redundancy cannot speak to the question of whether, given the choice between two strings of items with the same information content, speakers favour a distribution of those items that portions out its peaks across the whole utterance.

One line of evidence that may support an ordering strategy of noise resistance for language - which cannot be explained by redundancy - comes from cross-linguistic tendencies in basic word order. Word order in languages is often described in terms of subject (S), object (O) and verb (V). English is an SVO language, e.g., *The dog* (S) *bites* (V) *the man* (O). Maurits et al. (2010) suggests that the typological distribution of basic word orders (where object initial orders are particularly rare) can be explained in terms of broad preferences for uniform information density. By simulating the distribution of information in different word orders using data from child directed speech in English and Japanese, Maurits et al. (2010) showed that subject and verb initial word orders had smoother distributions than object initial orders, though they make the wrong typological prediction for verb initial vs subject initial orders. (However, see Ferrer-i-Cancho, 2015, 2017, for alternative accounts of the distribution of word order typology). Gibson et al. (2013) compare SVO and SOV orders in a gestural experiment, and found a surprisingly high use of SOV even among their English speakers (when the resulting string is not ambiguous, in which case SVO is preferred). The overall bias in favor of SOV could be due to the fact that SOV produces more informationally uniform strings on average, in the sense of information peak dispersal, but they do not explicitly make this point.

Overall, this previous research shows that speakers often deploy redundancy, potentially as a strategy for information distribution uniformity. However, this is difficult to disentangle from availability based production accounts. Existing studies do not address a strategy of ordering the peaks and troughs of information across the whole utterance, and so the noise resistance function of manipulating the distribution of information content remains under-explored.

Here, we provide a measure which can quantify the ordering-based uniformity of information distributions across entire utterances, and use a simulation to test the effectiveness of more even distributions (in the ordering sense alone) in the presence of noise. First, we introduce a novel descriptive statistic for assessing the uniformity of information content across whole utterances,² and propose an algorithm which compares the actual statistic to a distribution of information that is optimized for uniformity. Second, using a simple

²Maurits et al. (2010) actually stop just short of doing this: they do have a way of comparing information uniformity amongst the set of 3-word sentences and 3-word experimental stimuli they consider, and they describe an idea that could be fleshed out as a more general method: estimating the degree of linearity in a model of a series of left-to-right conditional trigram word entropies. We note, however, that this is more difficult to operationalize than the method we propose. First, it involves linear model estimation and comparing derived goodness of fit statistics. Secondly, such a method can only consider information distributions from left-to-right but not right-to-left. Thirdly, this method will be difficult to use to compare highly non-linear distributions to each other.

simulation, we show that a) utterances that are more optimized for uniformity are significantly more resistant to noise, and b) actual utterances are closer to the optimum than would be expected by chance. Finally, we discuss the implications of these results and suggest how this might contribute to future work.

2 Measuring information density across an utterance

We have developed a simple descriptive statistic (deviation of the rolling mean, or DORM) for assessing information uniformity across a whole utterance (i.e. in a string of words, morphemes, or phrases), and an algorithm (uniform information density optimization, or UIDO) that allows for comparison between utterances that contain different units. The combination allows us to gauge the uniformity or “clumpiness” of an utterance based on its order of information-bearing units (in comparison with other possible orders of those units, but not with signals that contain more redundancy). We can also readily compare the uniformity of one whole utterance to another, regardless of how much information is being transmitted by the utterance overall, by using UIDO to calibrate the DORM of a given utterance. This calibration allows us to assess the distance between the uniformity in a real sentence and its optimum, allowing us to compare DORMS across utterances. Again, this allows us to compare uniformity based on ordering across utterances while controlling for redundancy.

Below we apply the DORM measure for informational uniformity to sentences from the *The Penn-York Computer-annotated Corpus of a Large amount of English* (PYCCLE; Ecay, 2015³), and simulate the effect of noise on those sentences alongside theoretical reorderings of their information content. Specifically, we take all 628,083 sentences from the corpus that are exactly 10 words in length, measure their uniformity as described below (Section 2.1), and then use a simple simulation to compare how they fare under conditions of random single-word (categorical) noise events and simulated adjacent word (“clustered”) noise.

We consider the original sentences from the corpus alongside three other orders of the same words (1) a sorted, and thus, maximally non-uniform version, (2) a random permutation of their words, and (3) an approximate theoretical optimum of uniformity produced by UIDO. We have restricted our sample to sentences of the same length so that we can directly test the effect of uniformity on noise disturbance while experimentally controlling for sentence length.

The study will show that an artificial, optimally uniform version of each sentence (UIDO) results in the fewest catastrophic losses of information, but that nevertheless, the information distributions found in natural sentences are more uniform and result in fewer

³This corpus is an annotated version of the *Eighteenth Century Collections Online* and *Early English Books Online* corpora, which together are 1.05 billion words.

catastrophic communication failures than either the artificial asymmetric or randomized orders. First, however, we introduce the DORM measure and the UIDO optimization algorithm.

2.1 DORM: A descriptive statistic for information uniformity

The descriptive statistic, the “Deviation of the Rolling Mean” (DORM), provides a concise and easily interpretable summary of how uniform or clumpy a particular utterance is. Given a sentence and a list of information content measures for its unit probabilities (which could be based on words, bigrams, phrases, etc., and conditional or unconditional probabilities of the above – the method is fully general), we take the arithmetic mean of every adjacent, overlapping pair of information content values (i.e. a rolling mean). We then compute the sample variance of the resulting list of means (an unbiased estimator, using denominator $n - 1$) and arrive at a single number, the DORM.⁴

The DORM concisely measures the overall uniformity of the sentence’s information distribution. A high DORM means that overlapping sequences of adjacent words are often very different to each other in information content. In other words, a high value indicates clumps of information, creating peaks and troughs that mean the sentence is less uniform in terms of ordering. On the other hand, a low DORM indicates an utterance with a more uniform distribution overall.

This measure quantifies the uniformity of an utterance, while also allowing for some comparison of uniformity between utterances – even where these have different values of information content. However, it does not address the *optimality* of actual distributions of information. In other words, given a certain array of information content values, what would be the optimal way to order these to achieve the smoothest possible distribution? Below, we introduce an algorithm that calculates the distribution of a given array of information content values that is optimized for uniformity.

2.2 UIDO: Calculating the uniformity optimum for an utterance

The DORM’s companion algorithm is “Uniform Information Density Optimization” (UIDO). For any array of information content values, UIDO systematically reorders them, seeking to minimize the DORM for the whole utterance. We use UIDO experimentally in the simulation below, but it can also be used to find a minimum DORM for an utterance with which to calibrate the observed DORM of an utterance so that we can compare DORMs between utterances of different lengths. The by-sentence calibration is not necessary in the simulation below, since we control the length of the strings experimentally,

⁴We thank Ramon Ferrer-i-Cancho for the suggestion to use sample variance rather than sample standard deviation, which was in our first version of the function.

but future empirical work can use the $DORM(utterance) - DORM(UIDO(utterance))$ measure, calibrating DORM for an utterance by the approximate theoretical optimum for that length of utterance and set of linguistic elements. This per-utterance distance-from-UIDO measure can straightforwardly be used to compare different utterances to each other in terms of the unconditional entropies of linguistic elements.⁵

The UIDO reordering proceeds in the following way. First, the array of information content values is ordered greatest to least. Secondly, starting with the second and penultimate value in the array and moving inward, every other number is swapped, mixing up the large and small values. The algorithm calculates the DORM for the resulting array, and if this is lower than the original DORM, the array is kept; otherwise, the algorithm proceeds with the original array.

Whichever version of the array is kept, the next stage of the algorithm moves sequentially left-to-right through the array, one step at a time, and checks whether permuting the current value and the next yields a lower DORM than the previous. If not, the algorithm moves one step rightward and continues without changing the array. If the swap of adjacent numbers does yield a lower DORM, the version of the array with the swap is kept, and the algorithm goes back to the leftmost numbers and begins again. UIDO halts once it has reached the rightmost number and no additional pairwise permutation yields a lower DORM. This does not always yield the absolute lowest DORM, which would require enumerating all possible permutations and comparing the DORM, but it is much more computationally tractable on large datasets and arrays with more than just a few numbers. Even without finding the absolute lowest possible DORM for all permutations of a given array, UIDO provides a useful and consistent baseline that better allows us to interpret the DORM for a given sentence (and across many sentences), by comparing the observed DORM to the DORM of that sentence’s UIDO order. Code for calculating DORM and UIDO, validation simulations, and the PYCCLE simulation are all available at <https://github.com/CCuskley/NoiseResistance>.

We also validated the UIDO method, showing that it produces arrays that are consistently more uniform than random sequences (i.e., have consistently lower DORM values), even as the arrays become very long. To demonstrate this we drew arrays of probabilities of varying lengths (10 units to 100) from both uniform and Zipfian distributions.⁶

⁵We note that UIDO, as currently defined, cannot perform the same calibration of DORM if DORM is applied to conditional entropies of linguistic items, since it relies on reordering the items, and thus would change the estimated conditional entropies. We intend to develop a calibration mechanism for conditional entropies in future work. One straightforward possibility is to permute the string as UIDO currently does, but re-estimate conditional entropies for all the items after each permutation, given the permuted string, and then compute the DORM on the result. Again, as UIDO currently does, the permutation procedure would then continue until a minimum DORM is reached.

⁶Probabilities from the Zipfian distribution are word frequencies from the Brown Corpus (Francis & Kucera, 1979), drawn using the `zipfR` package (Evert, Baroni, & Evert, 2006). Probabilities from a uniform distribution were generated in R with: `-log2(runif(n=stringLength,min=0,max=1))`. Full code available at: <https://github.com/CCuskley/NoiseResistance>

Figure 1: Comparison of DORM values between random sequences of probabilities drawn from a uniform or Zipfian distribution, alongside the DORM for the UIDO optimized sequence. We examined generated arrays of lengths 10 through 100 units, and 1000 iterations per length (91,000 arrays considered in total). DORMs are calculated using information content based on probabilities drawn from either a uniform or Zipfian distribution. GAM smoothing curves are fit to the DORMs.

Random sequences drawn from either uniform or Zipfian distributions (Figure 1) have consistently higher DORM values (less informational uniformity) than their optimized versions created using the UIDO method. Even as random sequences of probabilities become very long, their DORMs do not reach zero in either case (at least up to sequences of 100 units long). It is worth noting, however, that the DORM of the UIDO optimized arrays converges to a very small number as the arrays get longer, perhaps more so in the Zipfian case.⁷

The UIDO-optimized values of DORM for random sequences remain lower than the DORM values for random sequences at every length, and their average seems to stabilize at values just above 1 for arrays longer than 25 elements in the uniform case (Figure 1, solid orange line). In the Zipfian case, DORMs continue to increase as the arrays get longer (dashed blue line, Figure 1). Interestingly, in neither case does the DORM of the purely random sequences decrease as the arrays get longer, which suggests that peak uniformity as measured by DORM (i.e. the lower bound of DORM) is not achieved by random sequences of arbitrarily long length; the DORM of the UIDO-permuted sequences (solid lines in 1) shows there is almost always a way to do better than the random sequence by hyper-dispersing high information units relative to low ones (and vice-versa) for each

⁷We thank Ramon Ferrer-i-Cancho, p.c., for raising the question of what happens to DORM as the arrays get longer, and asking whether DORM converges to 0 for very long arrays.

length.

This appears, at first glance, to be contrary to a mathematically very well-motivated expectation in Ferrer-i-Cancho (2017): that information uniformity can be achieved if the symbols in a sequence are “independent identically distributed (i.i.d.) variables”, and the sequence is long enough. That is to say, the amount of information per symbol will average out to maximal uniformity as a sequence of completely random symbols gets longer. However, this result is permitted by a “strong” definition of uniform information density that is formulated in terms of redundancy, in which (1) only conditional entropy of a variable is considered, and (2) the goal is to minimize differences in conditional entropy from one variable to the next (i.e. such that, in the ideal case, the conditional entropy of every linguistic element is the same). Here, we consider unconditional entropies, and perhaps crucially, DORM and UIDO together represent a different conceptualization to other information density accounts which tend to focus on redundancy. DORM allows peaks and troughs of entropy that are one unit long (e.g. from one word to the next), but penalizes peaks and troughs in any sub-string longer than 2, and UIDO minimizes DORM. Therefore, DORM measures the dispersal of peaks and troughs, and UIDO hyper-disperses peaks and troughs across the sequence. This allows us to consider unconditional word entropies while taking their order into account in the DORM statistic (i.e. a window of length 2 in computing the rolling mean represents the order of elements in the string), and it also sets the optimum for uniformity at something more achievable in natural language morphosyntax, where there is necessarily some alternation between high frequency, low information closed-class (“function”) and lower frequency, higher information open-class (“content”) morphemes.

3 Is uniformity adaptive in noisy conditions?

Below, we present a simulation which uses DORM and UIDO to show that optimally uniform utterances are, in fact, more resistant to noise than other sequences.

3.1 Method

3.2 Simulating noise

For each 10-word sentence from the PYCCLE corpus, we simulate a noise disturbance for each actual sentence and for three theoretical versions of the actual sentence made by permuting its word order: (1) a maximally asymmetric version, with word information content values ordered greatest-to-least, (2) a randomized version, and (3) an optimally smooth, uniform version made using UIDO on its word information content values. These information content values were derived by $\log_2 \frac{1}{p_i}$ with word probabilities estimated from

PYCCLE itself. Our method means we can directly compare different information distributions using the same information content values in each sentence, while controlling for length of sentence.

To simulate continuous noise, we randomly selected a sequence of three numbers from each version of a sentence to be “knocked out”, as if some continuous noise event had occurred during the sentence and obscured three words in a row. To simulate single-word noise events, we also included a condition in which we randomly selected any three words from each sentence.^{8 9}

For each type of noise, the same three units were knocked out of each version of the sentence – for example, for categorical single-word noise, if the 1st, 4th, and 5th items were knocked out of the actual sentence, the 1st, 4th, and 5th items would be knocked out of the optimized, randomized, and asymmetric versions. For clustered noise, the three-item sequence began at the same position in the string in each of the four versions of each of the ~600K sentences extracted from the corpus. In other words, for a given trial, if we suppose noise knocked out the first three words, it knocks out the first three words in the actual sentence, the first three in the randomized version, the first three in the asymmetric version, and the first three in the UIDO version. We consider clustered noise to be more realistic as a model for many linguistic situations than the single-word noise (though this may also exist, and so we include it in our simulation): for example, interfering acoustic noise, momentary lapses in attention, and some types of short-term memory failures can all be thought of as types of continuous noise that often span over multiple linguistic units.

We have, of course, made simplifying assumptions about the nature of the noise events for the purposes of our simulation, and there are certainly other types of noise which interfere with utterances in various ways we have not considered here. Regarding the choice of noise “window” that is 3 words long, our goal was to use a standard, controlled amount of noise for every trial which was longer than 1 word (so that the information *distribution* across the utterance could reasonably be affected), but not 2 words long, since 2 words is exactly the window that the UIDO algorithm has access to (and so the results in that case would have been built-in to the experiment). Therefore, we settled on a 3-word-long noise segment as a test for the robustness of 10-word sentences in our model. Any longer a noise window would have approached destroying more than 50% of the information in most trials regardless of condition, and so would not allow

⁸We note here that to the extent that the randomly selected words happen to occur sequentially in the string, these two conditions become equivalent. The probability of any k categorical noise events all being adjacent in a string is $\frac{n-k+1}{\binom{n}{k}}$, or the number of k -length slots available in n positions, out of the total number of possible combinations of k positions from n positions. For the case of 3 noise events in a 10 word string, this probability is $\frac{1}{15} \approx 0.067$.

⁹In each case, one can think of this noise as either “deletion” noise or “erasure” noise (Futrell & Levy, 2017); the results below should apply to either type (though the effect of information uniformity on erasure vs deletion is left as an interesting topic for further research.)

Figure 2: Raincloud plots of (Allen et al., 2019) distributions of the square root of DORM values (the standard deviation instead of the variance) for the 628,083 sentences in each different version. Note we display only one plot here, representing the DORM values for all sentences prior to the noise simulation.

us to compare the different word orders on the 50%-information-retained criterion (i.e. catastrophic failure of communication).

3.3 Results

Figure 2 shows the distribution of DORM values for each kind of sentence version; note that although single-word noises and clustered noise were implemented in separate simulations, the DORM distributions for the altered sentences don’t differ between the two simulations.¹⁰ In both cases, the UIDO algorithm is successful in optimizing for the lowest DORM value: the “uniform” sentences show the lowest DORM values with the highest peak, significantly lower than actual ($t = -993.61$, $df = 623082$, $p < 2.2 \times 10^{-16}$, Cohen’s $D = 1.26$), random ($t = -1079.8$, $df = 623082$, $p < 2.2 \times 10^{-16}$, Cohen’s $D = 1.36$), or asymmetric ($t = -1321.9$, $df = 623082$, $p < 2.2 \times 10^{-16}$, Cohen’s $D = 1.67$) versions. The asymmetric versions show the broadest range of DORM values, which skew generally higher. The actual and randomized sentences fall somewhere in between, with actual sentences having lower DORM values on average than random ($t = -221.31$, $df = 623082$, $p < 2.2 \times 10^{-16}$, Cohen’s $D = 0.28$) or asymmetric sentences ($t = -1009.4$, $df = 623082$, $p < 2.2 \times 10^{-16}$, Cohen’s $D = 1.28$), though the difference between actual and random sentences is small (Cohen, 2013). The first important result of the study is that the naturally occurring sentences do indeed show smoother information distributions on average than chance orderings (randomized) or asymmetric orderings, based on the DORM measure of uniformity.

¹⁰The only difference between the two simulations, other than the noise condition, is that the “Random” order for each sentence may not be identical between the simulations; otherwise, all sentences and all orders of sentences were identical.

Figure 3 shows the distributions of the proportion of information lost in each sentence version due to the noise event. First, note that the proportion of information lost does not vary across sentence types for 3 single-word categorical noise events. In other words, natural language sentences are no more robust against categorical noise than optimized, random, or asymmetric orderings of the same information content values.

In the case of clustered noise, higher DORMs significantly predict more bits lost ($p < 2 \times 10^{-16}$, $\beta = 0.157$, in a simple linear regression¹¹ with dependent variable = bits lost due to noise from each string, and independent variable = DORM for the given string).¹² To compare the distributions information lost under clustered noise across actual sentences, uniform, random, and asymmetric versions in more detail, we use a Kolmogorov-Smirnov test, a general comparison of cumulative distribution function (CDF) of two distributions. We use a one-tailed version of the test given our specific predictions: we predict that the distribution of information lost in uniform sentences will be lower than the actual sentences (i.e., the CDF will for the uniform sentences will fall below the CDF for the actual sentences, whether by lower mean or skew or both), while the random and asymmetric will be higher (i.e., the CDF for the actual will fall below the CDF for the random and asymmetric). This test confirms that the distribution of information lost with clustered noise skews higher for the original sentences than for the optimized ones ($D = 0.10, p < 2.2 \times 10^{-16}$, two sampled, one tailed), while it skews significantly lower than the random or asymmetric versions (random: $D = 0.023, p < 2.2 \times 10^{-16}$; asymmetric: $D = 0.18, p < 2.2 \times 10^{-16}$; both two sampled, one tailed). The mean proportion of information content lost varies little across the sentence versions (0.3 for all versions, unsurprisingly, given that 3 values were removed regardless). This combination of tests confirms that more information is lost in the asymmetric and randomized conditions under clustered noise than in the original, actual sentences, and that the UIDO-optimized versions lose the least.

While the variance also exhibits little difference between conditions for single-word noise events (all ± 0.00377), there are interesting patterns for clustered noise. Under clustered noise, the information lost in the original sentences has a smaller variance around the shared mean (± 0.003) than the randomized (± 0.004) or asymmetric (± 0.009) versions, and the UIDO-optimized versions have the lowest variance of all (± 0.002). In other words, the versions with smoother information distributions (lower DORM values) tend to lose less variable amounts of information to noise across trials (as shown in Figure 3).

We also considered the sentence versions in light of what we might call catastrophic

¹¹We used simple regressions because mixed-effects models with random intercepts for trial would not converge on a dataset with this 628,083 trials. However, since each trial consists only of the same sentence in the 4 versions we engineered, the concern over random bias in multiple observations per trial should not apply to this case.

¹²Call: `lm(formula = absNoise ~ dorms, data = propnoise.df)`

Figure 3: Raincloud plots (Allen et al., 2019) of distributions of the proportion of information lost in 628,083 trials in each condition under 3 single-word noises (A) and clustered noise (B).

Figure 4: The number of sentences with “catastrophic” information failure in each condition (a noise event knocking out $\geq 50\%$ of the total information content) under 3 single-word noises (a) and clustered noise (b).

communication failures, i.e. situations under which communication of an utterance may fail completely. Beyond simply minimizing information lost, such situations may be characterized by information loss reaching some critical threshold, after which an utterance becomes communicatively useless. For the purposes of this simulation, we suppose that a loss of more than 50% of the total information content in a string is “catastrophic”.

As Figure 4 shows, there is no difference in the rate of catastrophic communication failures across conditions for single-word noise.

Under clustered noise, on the other hand, the optimally uniform information distributions produced by UIDO result in zero catastrophic communication failures in all 628,083 sentences. The original, actual order of each sentence results in fewer catastrophic failures (207; 0.03%) than either the randomized (649; 0.1%) or asymmetric (6824; 1.1%) orders. In a logistic regression with DORM as the predictor and catastrophic loss as a binary outcome variable, lower DORM values are significantly less likely to lead to catastrophic

failure ($p < 2 \times 10^{-16}$, $\beta = 0.124$) under clustered noise.¹³ Together with the fact that the DORM significantly predicts the overall bits lost, this strongly suggests that smoother distributions more effectively prevent communication failures.

4 Discussion & Conclusions

Since Aylett and Turk (2004) proposed the smooth signalling redundancy hypothesis, important corpus-based and experimental work has shown that speakers prefer to avoid large spikes in information content. However, the proposed mechanism behind this - that this arises from an adaptation for noise resistance - has not been extensively tested. Bergen and Goodman (2015) showed that speakers strategically varied prosody to reduce peaks under artificially noisy conditions, but more extensive experimental tests that directly link linguistic information-packaging strategies (e.g., syntactic reduction/expansion/ordering, lexical choice) to the noisy channel have yet to be explored.

In order to make progress in this growing area, we have introduced a novel but simple measure, the DORM, which quantifies information density over an entire utterance, while previous work has only focused on measuring local spikes or troughs. We also introduced an algorithm, UIDO, which can take any array of information content values - calculated using anything from word probabilities to phrase probabilities - and provide the order of these values which is optimized for uniformity. Importantly, UIDO focuses on an ordering strategy of noise resistance, optimizing sequences by hyper-dispersing peaks and troughs throughout a sequence. Redundancy optimization can be achieved by ‘smoothing’ the amount of information from unit-to-unit; this makes them more similar, such that the most uniform string using this strategy could be an infinite string of random symbols (cf. Ferrer-i-Cancho (2017)). Ordering optimization, on the other hand, allows for large differences between local symbols to be maintained in a sequence, reflecting an optimum more commensurate with the observed patterns of content and function words in natural language morphosyntax.

Finally, we used these measures in a simulation, making the first effort to quantify exactly how advantageous uniformity is in actual language, independent of length or other aspects of a string, for communication under noisy conditions. Using word probabilities from the PYCCLE corpus in actual, asymmetric, random, and optimized orders, we showed that optimal orders are significantly less likely to result in catastrophic communication failures particularly for clustered noise. In other words, this simulation provides tangible evidence that more uniform information distributions do, in fact, function to ensure more effective communication under noisy conditions.

While the current simulation used word frequencies (in this case from PYCCLE), we

¹³Call: `glm(formula = bigLoss ~ dorms, family = binomial, data = propnoise.df)`

stress that the method is much more general. In particular, the descriptive measure we introduce, the DORM, can work to quantify information distributions of any sequence of units for which it is possible to calculate reliable probabilities. For example, information theory has led to important advances in fields as diverse as genetics and neuroscience (Debowski & Bentz, 2020). The DORM measure could be used by any researchers using information content measures on temporal sequences. In terms of specific linguistic applications, the DORM measure could also be readily applied to more complex lexical measures such as n-gram probabilities, conditional probabilities, or probabilities of larger syntactic units available from parsed corpora (e.g., for a recent application of DORM to examine changes in word order in English and Icelandic, see Wallenberg, Bailes, Cuskley, & Ingason, 2021).

The measure should also be applied to cross-linguistic research – an intriguing possibility is that different languages use different strategies to accomplish optimized uniformity given other structural constraints on the process. For example, a language with more head-final structures may not alternate low and high frequency words or morphemes as often as e.g. English does, purely as a result of structural constraints. And for many languages, the difference between computing uniformity on words vs morphemes will lead to non-trivially different results, which our measure enables researchers to examine systematically.

Future experimental and corpus-based work will continue to shed light on how speakers manipulate the information distributions of linguistic utterances. The current contribution has not only provided a new descriptive measure to quantify the uniformity of information across utterances, but is the first study to capitalize on the control afforded by a simulation to specifically show the utility of uniformity independent of utterance length.

Contributions

On the whole, CC and JW contributed equally to the work; RB provided theoretical contributions and substantial sections of writing. **CC**: Conceptualization, Validation, Formal analysis, writing (original draft), visualization, **JW**: Conceptualization, Methodology, Software, Formal Analysis, writing; **RB**: Theoretical framing, conceptualization, writing.

Acknowledgements

This work was funded in part by ESRC-SDAI grant ES/T005955/1, *Constraints on the Adaptiveness of Information in Language*, awarded to CC, JW, and RB. Thanks go to Anthony Kroch and Tom Smulders for discussion of various aspects of uniformity and information theory, and to reviewers.

Conflict of interest

The authors declare no conflict of interest

References

- Allen, M., Poggiali, D., Whitaker, K., Marshall, T. R., & Kievit, R. A. (2019). Raincloud plots: a multi-platform tool for robust data visualization [version 2; peer review: 2 approved]. *Wellcome Open Research*, 4.
- Anderson, A. H., Bader, M., Bard, E. G., Boyle, E., Doherty, G., Garrod, S., ... others (1991). The HCRC map task corpus. *Language and speech*, 34(4), 351–366.
- Aylett, M., & Turk, A. (2004). The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech. *Language and speech*, 47(1), 31–56.
- Barinova, G. (1971). Redukcija vypadenie intervokal'nych soglasnyh v razgovornoj reci. *Fonologic eskie podsistemy*, 117–27.
- Bergen, L., & Goodman, N. D. (2015). The strategic use of noise in pragmatic reasoning. *Topics in cognitive science*, 7(2), 336–350.
- Bock, K. (1982). Toward a cognitive psychology of syntax: Information processing contributions to sentence formulation. *Psychological Review*, 89(1), 1.
- Bock, K. (1987). An effect of the accessibility of word forms on sentence structures. *Journal of memory and language*, 26(2), 119–137.
- Cohen, J. (2013). *Statistical power analysis for the behavioral sciences*. Academic press.
- Debowski, L., & Bentz, C. (2020). Information theory and language. *Entropy*, 22(4), 435–441. doi: 10.3390/e22040435
- Eccy, A. (2015). *The Penn-York Computer-annotated Corpus of a Large amount of English based on the TCP (PYCCLE-TCP)*. Public release 1. <https://github.com/uoy-linguistics/pyccle>.
- Evert, S., Baroni, M., & Evert, M. S. (2006). The zipfr package. Available on-line at URL: <http://cran.r-project.org/doc/packages/zipfr.pdf>.
- Fenk, A., & Fenk, G. (1980). Konstanz im kurzzeitgedächtnis-konstanz im sprachlichen informationsfluß. *Zeitschrift für experimentelle und angewandte Psychologie*, 27, 402.
- Fenk-Oczlon, G. (1989). Geläufigkeit als determinante von phonologischen backgrounding-prozessen. *Papiere zur Linguistik*, 40(1), 91–103.
- Fenk-Oczlon, G. (2001). Familiarity, information flow, and linguistic form. *Typological Studies in Language*, 45, 431–448.
- Ferreira, V. S. (1996). Is it better to give than to donate? syntactic flexibility in language production. *Journal of memory and language*, 35(5), 724–755.
- Ferreira, V. S., & Dell, G. S. (2000). Effect of ambiguity and lexical availability on syntactic and lexical production. *Cognitive psychology*, 40(4), 296–340.
- Ferrer-i-Cancho, R. (2015). The placement of the head that minimizes online memory: a complex systems approach. *Language Dynamics and Change*, 5(1), 114–137.
- Ferrer-i-Cancho, R. (2017). The placement of the head that maximizes predictability. an information theoretic approach. *arXiv preprint arXiv:1705.09932*.
- Ferrer-i-Cancho, R., & Solé, R. V. (2003). Least effort and the origins of scaling in human language. *Proceedings of the National Academy of Sciences*, 100(3), 788–791.

- Francis, W. N., & Kucera, H. (1979). Brown corpus manual. *Letters to the Editor*, 5(2), 7.
- Frank, A. F., & Jaeger, T. F. (2008). Speaking rationally: Uniform information density as an optimal strategy for language production. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 30).
- Futrell, R., & Levy, R. (2017). Noisy-context surprisal as a human sentence processing cost model. In *Proceedings of the 15th conference of the european chapter of the association for computational linguistics: Volume 1, long papers* (pp. 688–698).
- Genzel, D., & Charniak, E. (2002). Entropy rate constancy in text. In *Proceedings of the 40th annual meeting of the association for computational linguistics* (pp. 199–206).
- Gibson, E., Piantadosi, S. T., Brink, K., Bergen, L., Lim, E., & Saxe, R. (2013). A noisy-channel account of crosslinguistic word-order variation. *Psychological science*, 24(7), 1079–1088.
- Hartley, R. V. (1928). Transmission of information 1. *Bell System technical journal*, 7(3), 535–563.
- Jaeger, T. F. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive psychology*, 61(1), 23–62.
- Jakobson, R. (1961). Linguistics and communication theory. In *Proceedings of symposia in applied mathematics* (Vol. 12, pp. 245–52).
- Kanwal, J., Smith, K., Culbertson, J., & Kirby, S. (2017). Zipf’s law of abbreviation and the principle of least effort: Language users optimise a miniature lexicon for efficient communication. *Cognition*, 165, 45–52.
- Levelt, W., & Maasen, B. (1981). Lexical search and order of mention in sentence production. In *Crossing the boundaries in linguistics* (pp. 221–252). Springer.
- Levy, R. P., & Jaeger, T. F. (2007). Speakers optimize information density through syntactic reduction. In *Advances in neural information processing systems* (pp. 849–856).
- Maurits, L., Navarro, D., & Perfors, A. (2010). Why are some word orders more common than others? a uniform information density account. In *Advances in neural information processing systems* (pp. 1585–1593).
- Pluymaekers, M., Ernestus, M., & Baayen, R. (2005). Articulatory planning is continuous and sensitive to informational redundancy. *Phonetica*, 62(2-4), 146–159.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell system technical journal*, 27(3), 379–423.
- Temperley, D., & Gildea, D. (2015). Information density and syntactic repetition. *Cognitive science*, 39(8), 1802–1823.
- Turk, A. (2010). Does prosodic constituency signal relative predictability? a smooth signal redundancy hypothesis. *Laboratory phonology*, 1(2), 227–262.
- Walker, M. (2017). *Why we sleep: Unlocking the power of sleep and dreams*.
- Wallenberg, J. C., Bailes, R., Cuskley, C., & Ingason, A. K. (2021). Smooth signals and syntactic change. *Languages*, 6(2), 60.
- Xu, Y., & Reitter, D. (2016). Entropy converges between dialogue participants: explanations from an information-theoretic perspective. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 537–546).
- Xu, Y., & Reitter, D. (2018). Information density converges in dialogue: Towards an information-theoretic model. *Cognition*, 170, 147–163.

- Zhan, M., & Levy, R. P. (2018). Comparing theories of speaker choice using a model of classifier production in mandarin chinese..
- Zipf, G. K. (1935). 1965. *the psychobiology of language*. Cambridge, MA: The MIT Press.