


The time scale of redundancy between prosody and linguistic context

Tamar I. Regev,¹ Chiebuka Ohams,¹ Shaylee Xie,² Lukas Wolf,³
Evelina Fedorenko,¹ Alex Warstadt,⁴ Ethan G. Wilcox,⁵ Tiago Pimentel³

¹MIT, ²Wellesley College, ³ETH Zürich, ⁴UCSD, ⁵Georgetown
{tamarr,cohams,evelina9}@mit.edu, sx102@wellesley.edu,
{lukas.wolf,tiago.pimentel}@inf.ethz.ch,
awarstadt@ucsd.edu, ethan.wilcox@georgetown.edu

Abstract

In spoken communication, information is transmitted not only via words, but also through a rich array of non-verbal signals, including prosody—the non-segmental auditory features of speech. Do these different communication channels carry distinct information? Prior work has shown that the information carried by prosodic features is substantially redundant with that carried by the surrounding words. Here, we systematically examine the time scale of this relationship, studying how it varies with the length of past and future contexts. We find that a word’s prosodic features require an extended past context (3-8 words across different features) to be reliably predicted. Given that long-scale contextual information decays in memory, prosody may facilitate communication by adding information that is locally unique. We also find that a word’s prosodic features show some redundancy with future words, but only with a short scale of 1-2 words, consistent with reports of incremental short-term planning in language production. Thus, prosody may facilitate communication by helping listeners predict upcoming material. In tandem, our results highlight potentially distinct roles that prosody plays in facilitating integration of words into past contexts and in helping predict upcoming words.

 <https://github.com/Chief-Buka/contextual-redundancy>

1 Introduction

Auditory features of speech such as pitch, loudness, and tempo—collectively termed prosody—play a crucial role in conveying meaning. Prosody influences sentence-level interpretation, encoding both linguistic and para-linguistic cues relevant to the communicative context (Cole, 2015; Wagner and Watson, 2010; Breen et al., 2010). For example, prosody can signal phrase boundaries, emphasize key elements, transform statements into questions,

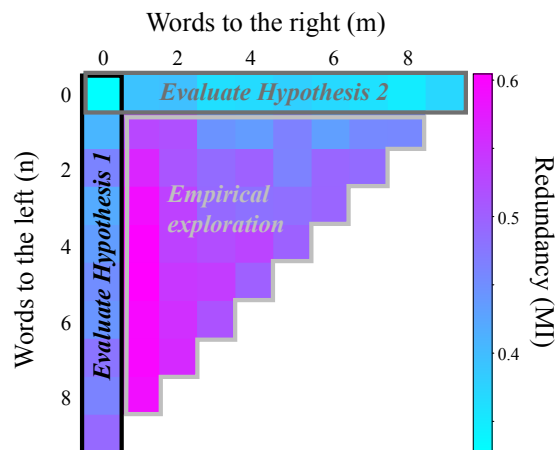


Figure 1: Redundancy (MI) between prosody (P_t) and linguistic context ($W_{n,m}$) as a function of the number of words contained in the past (n) and/or future (m) context. The values are averaged across 6 prosodic features (see Fig. 3 for each of these features separately). The first column is used to evaluate Hypothesis 1; the first row, Hypothesis 2 (see section §1). Time scales are defined as the time after which the MI does not increase.

and express sarcasm, excitement, or doubt. However, much of the information conveyed by prosody is redundant with the information encoded in the words themselves (Wolf et al., 2023)—that is, it is possible to predict a word’s prosodic features from its **linguistic context**, i.e., its surrounding words.¹

Past work, however, only quantified the redundancy between prosody and the *entire* linguistic context. The question of the *dynamics* of this redundancy is thus left open. Here, we investigate the time scale of redundancy between prosodic features of a given word and its linguistic context by systematically varying the length of past or future context from 0 to 9 words in each direction (0 includes only the word whose prosody is being tested).

¹In both Wolf et al. (2023) and our study, text is used as a proxy to measure information present in the words themselves, which we refer to as “linguistic information” (also termed “segmental information” in the phonology literature).

A priori, what might one expect regarding the time scale of redundancy between prosodic features and **past linguistic context**? During language comprehension, context is often critical to interpreting the current word and integrating it into the linguistic representations built thus far (Hale, 2001; Levy, 2008). However, human memory is limited: contextual information decays over time, and connecting an incoming word to one far back is costly (Gibson, 1998; Lewis et al., 2006; Futrell et al., 2020). Given these memory limitations, prosody may be critically helpful in conveying information that is redundant with past context but is *locally unique*, such that if a listener’s representation of the past context is imperfect, they would benefit from the prosodic information contained by the word itself. We therefore hypothesize that prosody carries information that is recoverable from *long-term past context*, but not from the short-term past (see Hypothesis 1 below).

With respect to **future linguistic context**, lexical and syntactic planning in language production is known to be relatively restricted (Brown-Schmidt and Konopka, 2008; Bock et al., 1994): we may have a general idea of what we want to say next, but the particular words and constructions are selected incrementally, often with revisions. As a result, how we pronounce a current word and what prosodic features we associate with it should not be affected by long-term future linguistic context. We therefore hypothesize that prosody carries information that is not recoverable from long-term future context (see Hypothesis 2 below).

The above logic is laid out in terms of “short” vs. “long” contexts, but what constitutes a short or long time scale during language processing? Following recent work in neuroscience (Jain and Huth, 2018; Regev, Casto, et al., 2024; Shain, Kean, et al., 2024) as well as work investigating intonational units in prosody itself (Inbar et al., 2020), we set the cutoff between “short” and “long” at *roughly* 1-2 words, although see section §6.1 for more discussion.

We thus formulate two main hypotheses:

Hypothesis 1. Long-scale past redundancy. *The redundancy between prosody and past linguistic context unfolds across a long time scale (longer than 1-2 words).*

Hypothesis 2. Short-scale future redundancy. *The redundancy between prosody and future linguistic context unfolds across a short time scale (of 1-2 words).*

To test these hypotheses, we quantify the redundancy between prosody and linguistic context as their mutual information, which we estimate using pre-trained language models, following Wolf et al. (2023). Notably, we extend their approach to investigate the time scale of this redundancy by jointly varying context lengths parametrically, from 0 to 9 words for both past and future contexts.² We then analyze how redundancy changes across time for several commonly discussed prosodic features: pitch, loudness, duration, pause, and prominence.

Our main results, averaged across all prosodic features considered are shown in Fig. 1 (see Fig. 3 for feature-specific results). In line with our hypotheses, these results demonstrate that prosody is redundant with relatively long-term past linguistic context (up to 8 words), but is only redundant with short-term future context (up to 2 words). This finding advances our understanding of the dynamics of the redundancy between linguistic context and prosody, with potential implications for the role of prosody in natural spoken communication.

2 Prosodic Features

Prosodic information is conveyed through multiple acoustic features of speech. Here, we examine the redundancy between linguistic context and: pitch, loudness, duration, pause and two versions of prominence (as detailed in section §4.1). We chose these prosodic features as they have been commonly investigated in prior prosody research (e.g., Breen et al., 2010; Gibson, 1998; Cole, 2015).

Pitch Pitch is the perceptual dimension over which listeners order sounds on a scale from low to high. The acoustic correlate giving rise to this perception is the periodicity of sound signals; pitch is thus often measured as the fundamental frequency (f_0) of the sound. Pitch is (arguably) the hallmark feature of prosody, having been extensively studied and characterized (Pierrehumbert, 1980; Silverman et al., 1992; Jun, 2006). In stress-accent languages like English, pitch contours carry contextual information that can signal a wealth of information, including emphasizing specific words, signaling boundaries, speech act type, and the speaker’s intent (e.g., interrogation, sarcasm, or affective state). Some typical pitch curves are the rise of pitch to

²In addition to examining the effects of the length of the past and future contexts separately, we also considered different combinations of past and future context lengths in order to identify the most predictive combination.

wards the last word of a question in yes/no questions in American English, the rise and then fall of pitch on a specific word to emphasize it, and a fall toward the end of phrases. Note that we do not study here the relationship between pitch and lexical identity, as in tonal-languages. However, see [Wilcox et al. \(2025\)](#), who use similar methods to study typological pitch variation across languages.

Loudness Loudness is the perceptual dimension over which listeners can order sounds on a scale from quiet to loud. The acoustic correlate giving rise to this perception is sound pressure, being measured as the intensity of acoustic energy. The loudness of speech can be used to transmit information, such as to emphasize important words or convey emotion. The correlation between pitch and loudness is partly explained by vocal production constraints; producing speech with higher energy helps raise and better control the fundamental frequency. However, loudness variations may also convey independent information from pitch.

Duration A word’s duration is the difference between its offset (end) and onset (start) times. The relationship between word duration and linguistic information has long been studied as a signature of efficiency in communication, such that more predictable words are reduced to a shorter duration ([Jurafsky et al., 1998](#); [Bell et al., 2009](#); [Seyfarth, 2014](#); [Coupé, Oh, et al., 2019](#); [Pimentel et al., 2021](#)). Further, elongating a word is a common way to emphasize it, or signal prosodic boundary. Duration is thus also highly correlated with pitch and loudness in natural speech, but it can also be used independently to convey meaning, or to compress words of low information content.

Pause A pause following a word is the time difference between the word’s offset (end) time and the next word’s onset (start), being another way to emphasize an important word in context, or to signal phrase boundaries ([Hawkins, 1971](#)). In contrast to phrase boundaries, within the phrase speech tends to be ‘connected’ such that there are usually no pauses between words; i.e., most pauses are of zero seconds (see section §4.1).

Prominence Prosodic prominence is a term that describes how salient a linguistic entity—in our case, a single word—is perceived relative to the words surrounding it in an utterance ([Terken and Hermes, 2000](#)). Unlike the previously described

prosodic features, prominence is a higher-level percept in the sense that it is not elicited by a single acoustic dimension. The perception of prominence is affected by multiple acoustic features ([Cole et al., 2010](#))—elongating a word, increasing the speech energy or modulating the f_0 contour of the word can all make this word be perceived as more prominent in context. Although other acoustic features, like timbre, can also affect prominence, and factors like word frequency influence it as well ([Cole et al., 2010](#)), a combination of duration, loudness and pitch has been proposed as an effective acoustic measure to quantify prosodic prominence ([Talman et al., 2019](#)), which we use here.

Notably, the prosodic features above present a high degree of correlation in natural speech ([Ladd, 1996](#)). For instance, as noted above, producing a higher pitch generally requires greater vocal effort, resulting in increased intensity; similarly, it may take longer to reach a higher pitch target, creating dependencies between pitch, loudness, and duration. Furthermore, prominence is strongly correlated to pitch, loudness and duration, as it is directly computed using a combination of these acoustic measures. We did not model these correlations explicitly; instead, we quantified the redundancy of each feature with linguistic context separately. This design choice does not affect our central research question, which concerns the time scale of these redundancies. Furthermore, although these prosodic features are correlated, the differences observed across them in our results may indicate that they also carry distinct information.

3 Redundancy between Prosody and Linguistic Context

This paper concerns the time scale of the redundancy between prosody and linguistic context, where ‘linguistic context’ here refers to the segmental information of an utterance, represented in our experiments as text. In this section, we first explain how this redundancy can be formalized as a mutual information, following [Wolf et al. \(2023\)](#). We then expand on this framework by discussing how context-length manipulations allow us to investigate the time scale of this redundancy.

3.1 Redundancy as Mutual Information

Let P_t be a prosody-valued random variable, which takes values $p_t \in \mathbb{R}$. Further, let \mathbf{W} be a words-valued random variable, which takes values $\mathbf{w} \in \mathcal{W}^*$,

where \mathcal{W} is a language’s lexicon. We follow Wolf et al. (2023) in formalizing the redundancy between prosody and linguistic context as the mutual information: $\text{MI}(\mathbf{P}_t; \mathbf{W})$. Under a few technical assumptions (e.g., the good mixed-pair assumption, see Wolf et al. for details), we can write this as:

$$\text{MI}(\mathbf{P}_t; \mathbf{W}) = H(\mathbf{P}_t) - H(\mathbf{P}_t | \mathbf{W}) \quad (1)$$

In this equation, **unconditional entropy** $H(\mathbf{P}_t)$ serves as a baseline, representing how much uncertainty there is about \mathbf{P}_t , in general. In turn, **conditional entropy** $H(\mathbf{P}_t | \mathbf{W})$ represents how much uncertainty remains about \mathbf{P}_t once we know \mathbf{W} . Their difference then represents how much information \mathbf{W} contains about \mathbf{P}_t (and *vice versa*).

We are now left with the problem of estimating these entropies. While these values are unknown, we only require two things to estimate them: a corpus of prosodic values coupled to linguistic contexts, sampled from the ground-truth distribution,

$$\mathcal{D}_{\text{tst}} = \{\mathbf{p}'_t, \mathbf{w}'\}_{n=1}^N, \mathbf{p}'_t, \mathbf{w}' \sim p(\mathbf{p}_t, \mathbf{w})$$

and models p_θ of distributions $p(\mathbf{p}_t)$ and $p(\mathbf{p}_t | \mathbf{w})$. We can then use a cross-entropy upper-bound (Pimentel et al., 2019) to estimate these entropies:

$$\begin{aligned} H(\mathbf{P}_t | \mathbf{W}) &\leq H_\theta(\mathbf{P}_t | \mathbf{W}) \\ &\approx \frac{1}{|\mathcal{D}_{\text{tst}}|} \sum_{\mathbf{p}'_t, \mathbf{w}' \in \mathcal{D}_{\text{tst}}} \log \frac{1}{p_\theta(\mathbf{p}'_t | \mathbf{w}')} \end{aligned} \quad (2)$$

where $p_\theta(\mathbf{p}_t | \mathbf{w})$ is replaced with $p_\theta(\mathbf{p}_t)$ when estimating $H(\mathbf{P}_t)$. We describe our dataset \mathcal{D}_{tst} and how to estimate p_θ in section §4.

3.2 Manipulating Context Length

To analyze the time scale of the redundancy between prosody and linguistic context, we will estimate $\text{MI}(\mathbf{P}_t; \mathbf{W})$ while systematically manipulating the amount of context, i.e. the number of words, in \mathbf{W} . We thus quantify ‘time’ in units of words, as opposed to seconds, acknowledging the discrepancy between these concepts due to varying duration of words and speaking rates. To this end, we define $\mathbf{W}_{\overset{\leftarrow}{n}, \vec{m}}$ as the linguistic context comprising n words before and m words after word \mathbf{W}_t , including the word itself:

$$\mathbf{W}_{\overset{\leftarrow}{n}, \vec{m}} = \langle \mathbf{w}_{t-n}, \dots, \mathbf{w}_t, \dots, \mathbf{w}_{t+m} \rangle \quad (3)$$

Thus, for instance, $\mathbf{W}_{\overset{\leftarrow}{0}, \overset{\rightarrow}{0}}$ corresponds to the word \mathbf{W}_t by itself, and $\mathbf{W}_{\overset{\leftarrow}{5}, \overset{\rightarrow}{3}}$ corresponds to the word




Figure 2: Estimation procedure for $H(\mathbf{P}_t | \mathbf{W}_{\overset{\leftarrow}{n}, \vec{m}})$. A span of words $\mathbf{W}_{\overset{\leftarrow}{n}, \vec{m}}$ which includes word \mathbf{W}_t is used as input to a model which predicts that word’s prosody \mathbf{P}_t . The loss function that the model minimizes estimates the conditional entropy.

\mathbf{W}_t with 5 words in its past context and 3 words in its future context (see Fig. 2).

Given this definition, we can then explore the time scale we are interested in by estimating the mutual information $\text{MI}(\mathbf{P}_t, \mathbf{W}_{\overset{\leftarrow}{n}, \vec{m}})$ while varying n and m . This amounts to estimating $H(\mathbf{P}_t)$ and $H(\mathbf{P}_t | \mathbf{W}_{\overset{\leftarrow}{n}, \vec{m}})$. The unconditional entropy does not depend on context; thus, to estimate it, we only need to compute a distribution over the domain of each prosodic feature. On the other hand, estimating the conditional entropy requires computing a family of conditional distributions for each prosodic feature, with one distribution for each n, m combination. In other words, we need a model $p_\theta(\mathbf{p}_t | \mathbf{w}_{\overset{\leftarrow}{n}, \vec{m}})$ that works for any n, m pair. We elaborate on this model in section §4.2.

Importantly, the $\text{MI}(\mathbf{P}_t; \mathbf{W}_{\overset{\leftarrow}{n}, \vec{m}})$ is a monotonically increasing function of both m and n ; larger contexts must contain at least as much (but maybe more) information about prosody than smaller contexts.³ However, at some value of m and some value of n , the MI might reach a plateau. We consider the **time scale** of the redundancy between prosody and linguistic context to be the value from which increasing context length does not significantly increase the MI.

4 Methods

4.1 Dataset

Our data-extraction process uses Wolf et al.’s (2023) proposed and publicly available pipeline (see their paper for more details).

Speech Data We use the LibriTTS spoken language corpus (Zen et al., 2019),⁴ which contains

³We note that—while the MI is monotonically increasing in theory—it is not necessarily the case that this underlying monotonicity will be reflected by our estimation methods.

⁴This dataset is derived from LibriSpeech audiobooks corpus (Panayotov et al., 2015), which is itself derived from

public domain audiobook materials (audio and text) recorded by volunteer narrators. This dataset contains 585 hours of English speech data at a 24kHz sampling rate, and includes recordings from 2,456 speakers reading aloud books which are paired with the corresponding transcripts. We filtered out texts from LibriTTS that contained less than three words (such as book and chapter titles) and we eliminated punctuation marks, since these can be very informative regarding prosody, and are not explicitly present in spoken communication.

Prosody Feature Extraction For extracting prosodic features we began by aligning the audio and text using the Montreal Forced Aligner (MFA; McAuliffe et al., 2017). With the alignment in place, both the duration and pause of each word can be easily computed from the words’ offset and onset times. Notably, 89.4% of the words in our dataset have a pause of 0 seconds. Duration was normalized by the number of syllables (from CELEX; Baayen et al., 1996) to obtain duration per syllable, reducing variability across word identities.

Pitch and loudness were extracted using the algorithms from Suni et al. (2017). These yield pitch—measured as fundamental frequency, f_0 —curves that we z -scored per speaker to normalize for interspeaker pitch range and avoid bimodal pitch distributions due to sex. For each word, we analyzed pitch in a window of up to 250ms (or the word’s full duration, if shorter) centered on the primary (stressed) syllable, as identified using CELEX. This choice reflects the fact that English pitch accents typically align with stressed syllables (Pierrehumbert, 1980) and reduces phonological variability due to word length. However, this approach focuses on lexical stress and may underrepresent boundary tones. We then averaged the pitch values within the selected window to yield a single mean pitch value per word. Loudness, measured as energy, was similarly extracted, with energy curves being averaged per word.

Prominence was computed using a composite acoustic measure from Suni et al. (2017), which integrates time-frequency variation across duration, energy, and f_0 . This measure has been validated in prior work as capturing prosodic prominence (Suni et al., 2017; Talman et al., 2019; Terken and Hermes, 2000). In addition to the absolute prominence values, we also computed relative prominence by subtracting the mean prominence

LibriVox (Kearns, 2014).

of the preceding three words from the current word’s value, emphasizing local contrast. We used the word-level prominence scores released with this dataset (Talman et al., 2019).

Splitting the Data into Train, Validation and Test

The dataset was divided into train, validation and test sets, following standard practice in machine learning. We used splits of the dataset provided by Talman et al. (2019). For training, we used a data split (termed train-360) containing 904 speakers, 11,262 sentences and 2,076,289 words. For validation, we used a data split (termed dev) containing 40 speakers, 5,726 sentences and 99,200 words. For testing, we used two data splits (both non-overlapping with the Train and Validation sets): one (termed train-100) containing 247 speakers, 33,041 sentences and 570,592 words, and the other (termed test) containing 39 speakers, 4,821 sentences and 90,063 words. We used the train-100 split as the main test set, except when analyzing word prominence; for those analyses, we combined the two test sets in order to obtain more stable results.

4.2 Estimating the Cross-Entropies

We now explain how we model the probability distributions $p_\theta(\mathbf{p}_t)$ and $p_\theta(\mathbf{p}_t \mid \mathbf{w}_{\leftarrow n, m}^{\leftarrow})$, which serve to estimate the unconditional and conditional cross-entropies, respectively.

Modeling $p_\theta(\mathbf{p}_t)$ To model this unconditional distribution over prosodic values, we simply follow Wolf et al. (2023) in using a Gaussian kernel density estimator (KDE). Given a training set \mathcal{D}_{trn} , sampled from $p(\mathbf{p}_t, \mathbf{w})$, this model is defined as:

$$p_\theta(\mathbf{p}_t) = \frac{1}{|\mathcal{D}_{\text{trn}}|} \sum_{\mathbf{p}'_t \in \mathcal{D}_{\text{trn}}} \mathcal{N}(\mathbf{p}_t; \mu = \mathbf{p}'_t, \sigma = \hat{\sigma}) \quad (4)$$

where \mathcal{N} are Gaussian distributions, each centered at a prosodic value $\mu = \mathbf{p}'_t$ and all having the same variance $\sigma = \hat{\sigma}$. We choose $\hat{\sigma}$ that achieves the highest likelihood on our validation set.

Modeling $p_\theta(\mathbf{p}_t \mid \mathbf{w}_{\leftarrow n, m}^{\leftarrow})$ To model this conditional distribution, we again follow Wolf et al. in finetuning a language model (LM), with an added linear layer on top, to predict the parameters of a conditional distribution over \mathbf{p}_t . Using language models to represent text offers greater expressivity than traditional regression-based approaches. Unlike Wolf et al., however, we limit our model’s input to include only $\mathbf{w}_{\leftarrow n, m}^{\leftarrow}$ instead of the entire \mathbf{w}

(Wolf et al. used the entire available context in the segments from the LibriTTS corpus, which varied in length between 4 and 60 words).

We assume the conditional distribution over prosody follows a parametric distribution \mathcal{Z} , and use a LM to predict this distribution’s parameters $\hat{\phi}$:⁵

$$\hat{\phi} = \text{LM}(\mathbf{w}_{\overleftarrow{n},\overrightarrow{m}}) \quad (5)$$

$$p_{\theta}(\mathbf{p}_t \mid \mathbf{w}_{\overleftarrow{n},\overrightarrow{m}}) = \mathcal{Z}(\mathbf{p}_t; \phi = \hat{\phi}) \quad (6)$$

We finetune this model by minimizing its cross-entropy on a training set \mathcal{D}_{trn} ; which amounts to minimizing the right-hand side of Eq. (2). As the cross-entropy is an upper bound on the entropy, the lower its value (and consequently the better our model) the tighter the estimate we get for the entropy. Due to compute time considerations, we estimate all n, m combinations using a single finetuned LM: During training, we sampled inputs of varying lengths, spanning 1 to 10 words. These inputs were obtained by first sampling an item from the dataset, and then randomly cutting it into the desired length (between 1-10 words). For each sample, the model then predicts the prosody of each of the words in this span in parallel; in a 7-word span, thus, the first word’s prosody is predicted as $p_{\theta}(\mathbf{p}_t \mid \mathbf{w}_{\overleftarrow{0},\overrightarrow{6}})$ and the 5th word’s prosody is predicted as $p_{\theta}(\mathbf{p}_t \mid \mathbf{w}_{\overleftarrow{4},\overrightarrow{2}})$. Importantly, we made sure that the models saw an equal number of samples for each n, m combination. For each prosodic feature, we tested several models, namely BERT, BERT-large and RoBERTa-large, and selected the model that gave the best results. In most cases, this was BERT-large, except for pauses and pitch where it was BERT. An early stopping criterion was applied, such that if the loss did not decrease for 3 epochs the model stopped training.

5 Results

We present our main results in Figs. 1 and 3. Namely, Fig. 1 displays the average mutual information (MI) across the 6 prosodic features tested here (pitch, loudness, duration, pause, absolute prominence, relative prominence), while Fig. 3 displays those results for each feature separately. See section App. A for the values of the unconditional entropies comprising those MIs.

⁵We evaluate models with Gaussian, Gamma and Laplace distributions, choosing the distribution that leads to the lowest cross-entropy on a validation set. Parameters ϕ are, e.g., the mean and standard deviation for a Gaussian.

The overall trend is consistent with the hypotheses outlined above. In line with the long-scale past redundancy hypothesis (Hypothesis 1), the mutual information between prosody (averaging across all prosodic features used here) and past context (first column in Fig. 1) increases as a function of context length until up to about 5-8 words, at which point it plateaus. In line with the short-scale future redundancy hypothesis (Hypothesis 2), the MI between prosody (averaging across all features) and future context (first row in Fig. 1) increases only until up to one or two words and then plateaus. Furthermore, at context lengths above one word, the MI with the past context tends to be higher than the MI with the future context, with this trend becoming more pronounced at longer contexts.

We also empirically explore other combination of past and future contexts (i.e., $n > 0$ and $m > 0$; the n -th column and m -th row in Fig. 1). This analysis revealed that prosodic features at the current word are best predicted by a combination of 5-8 words in the past and about one word in the future. Notably, this n, m combination led to higher mutual information than even other combinations with larger context (i.e., with $n', m', n' \geq n$, and $m' \geq m$). While this is theoretically impossible (adding more context can never decrease mutual information), this is likely due to our models’ training procedure not being able to ignore unhelpful contributions of long-scale contexts.

Interestingly, individual prosodic features show somewhat distinct patterns (see Fig. 3). The long-scale past hypothesis (Hypothesis 1) is supported for most prosodic features individually, but not for duration and pauses. For pitch, loudness, and prominence, the MI with past context (red curves in the lower plot for each feature in Fig. 3) increases until up to at least 3 words (although the curve for absolute prominence is a little noisy). For duration, however, the MI with past context never rises significantly above the 0, 0 point, indicating the past context does not add information beyond the word identity itself; and for pauses, the past MI saturates after two words.

The short-scale future hypothesis (Hypothesis 2) is also supported for most individual prosodic features, except for duration and pause. For pitch, loudness, and prominence, the MI with future context is shorter than with past context, saturating between 1 and 3 words (blue curves in the lower plot in each feature in Fig. 3). For duration, the MI with future context saturates after 1 word, which is

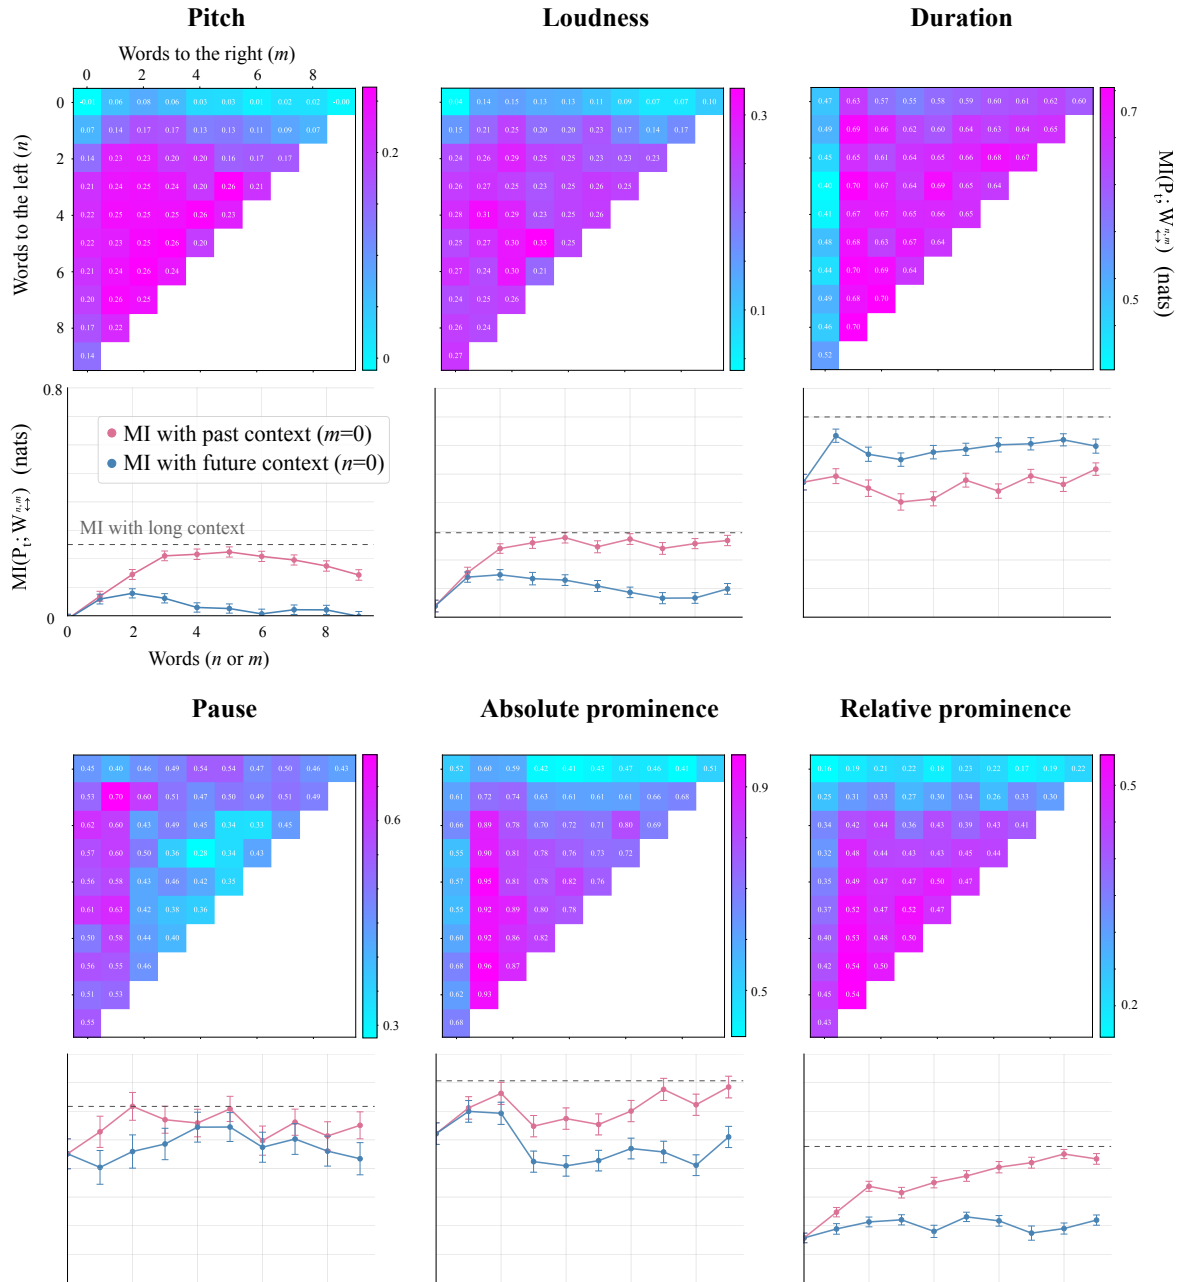


Figure 3: For each of the 6 tested prosodic features, two plots are presented. The upper plots are similar to Fig. 1, and display the redundancy, quantified as mutual information between a prosodic feature at a given word p_t and the linguistic context $w_{n,m}$, which includes the word itself and a varying length context, n words before, and m words after it. The lower plots display either: (a) in red, just the first column from the upper plots (corresponding to past linguistic context, which includes the word itself and a gradually increasing number of previous words, red curve) and (b), in blue, the first row from the upper plots (corresponding to future linguistic context, which includes the word itself and a gradually increasing number of future words, blue curve). Details displayed only for Pitch or Duration (labels, numbers, legend) are the same for all features. The MI values correspond to the mean across all test data. Error bars correspond to standard errors of the mean. Dashed horizontal gray lines corresponds to the MI with the longest available context, taken from Wolf et al. (2023) (between 6 and 60 words into the past or future). See section App. B for the distributions of these features in our dataset.

longer than with past context (the past curve shows no increasing trend). For pauses, the future MI saturates after about 4 words, which is longer than with past context—although these MI estimates are noisy for this feature, likely due to data sparsity. Notably however, the MI with past context is higher than with future context for all prosodic features except for duration. We thus conclude that the results from these individual prosodic features generally support the short-scale future hypothesis.

6 Discussion

This work aimed to characterize the time scale of the redundancy between prosodic and linguistic information. We built on the foundation of past work, which established that linguistic information is partially redundant with prosodic features (Wolf et al., 2023), to systematically manipulate the length of past and future context, from 1 to 10 words. Our extension of this pipeline is publicly available for researchers wishing to extend our analysis, potentially exploring the time scale of redundancy between other communication channels of interest.

Based on what is known about human memory limitations and linguistic planning, we hypothesized that redundancy between linguistic information and prosodic features extends over a relatively long time scale into the past, and a relatively short time scale into the future (see section §1). We find support for these hypotheses: For most prosodic features we tested (apart from duration and pause, see section §6.4), redundancy with past linguistic context unfolds across a long time scale (of 3-8 words), whereas redundancy with future linguistic context is shorter-scale (of 1-2 words). We next discuss the implications of these results with respect to the previous literature.

6.1 The Time Scales of Linguistic and Prosodic Information

Sentence comprehension is constrained by cognitive demands such as attention and working memory, leading to lossy representations of past linguistic context (Gibson, 1998; Vasishth et al., 2010; Futrell et al., 2020; Kuribayashi et al., 2022). It is, however, challenging to estimate the precise number of words maintained in working memory during language processing, as this likely depends on many factors. Neural evidence suggests that language-selective neural populations⁶ integrate

⁶Language-selective neural populations are clusters of neurons in the human brain that respond selectively to linguistic

information across time scales varying between 1-6 words (Jain and Huth, 2018; Shain, Kean, et al., 2024; Regev, Casto, et al., 2024). Moreover, work in prosody has put forward the notion of "intonation units", which follow a rhythmic structure of about 1 Hz (Inbar et al., 2020). Given the average speech rate of 200 words per minute (Yuan et al., 2006), each intonation unit contains about 3-4 words, which aligns with the relatively short context of linguistic processing.⁷

Our findings are generally in accord with these previous studies, and suggest that redundancy with past context spans around 3-8 words—a scale comparable to both the time scale of linguistic integration and prosodic segmentation.

6.2 Prosody as an Audience-Design Tool in Communication

A longstanding debate in linguistics concerns the extent to which language production is shaped by audience design considerations, with speakers tailoring their utterances to facilitate comprehension. Some evidence suggests that syntactic choices are not strongly adapted for listener needs but rather reflect the speaker's own constraints (Ferreira, 2008; Morgan and Ferreira, 2022). This apparent lack of audience design in syntax may stem from the rigid structural constraints imposed by the rules of the language (at least in English, which is used in much past work). In contrast, prosodic choices may offer greater flexibility for accommodating audience-centered considerations, allowing speakers to dynamically modulate pitch, loudness, and rhythm in real time. This flexibility may afford prosody a larger role in audience design (Clark et al., 2025), serving as a communicative channel that facilitates comprehension. For example, prior work shows a trade-off between a word's duration and its information content (Jurafsky et al., 1998; Bell et al., 2009; Coupé, Oh, et al., 2019; Pimentel et al., 2021), which is typically interpreted as arising to facilitate comprehension by more evenly

tic experimental manipulations, but not to manipulations in other perceptual or cognitive domains such as math and music. These populations are thought to constitute a brain network specialized for language processing (Fedorenko et al., 2024).

⁷To be precise about the generality of the claims above, we note that: the claims in (Gibson, 1998; Vasishth et al., 2010; Futrell et al., 2020; Kuribayashi et al., 2022) are intended as universal and were shown for English, German, and Japanese; the experiments in (Jain and Huth, 2018; Shain, Kean, et al., 2024; Regev, Casto, et al., 2024) are in English; the experiments in (Inbar et al., 2020) cover 6 typologically diverse languages; the experiments in (Yuan et al., 2006) are in English and Chinese.

distributing information across time (known as the uniform information density hypothesis; [Fenk and Fenk, 1980](#); [Genzel and Charniak, 2002](#); [Levy and Jaeger, 2007](#); [Meister et al., 2021](#)). Our findings may be taken as further evidence for the audience-design view of prosody: the fact that its redundancy with past context extends over long time scales suggests that prosody may help clarify the relationship between the current word and the broader preceding context, thus facilitating the extraction of the intended meaning. Given that prosody appears to carry locally unique information, prosodic features may aid comprehension when representations of longer-scale past are lossy.

6.3 The Relationship Between Prosody and Future Words

Our findings indicate that the redundancy between prosody and future words is overall lower compared to past words. However, prosody still exhibits a strong relationship with the upcoming word or two. This short-range relationship may arise from motor constraints on prosodic articulation, from local linguistic dependencies such as fixed expressions, or from short-scale prosodic planning—consistent with theories of language production that emphasize incremental, short-term planning ([Brown-Schmidt and Konopka, 2008](#); [Bock et al., 1994](#)). In addition, prosody may help listeners form expectations about upcoming words through cues such as duration, pauses, or pitch modulations. This interpretation aligns with recent findings that humans are better at predicting upcoming words given spoken, compared to written, context ([Botch and Finn, 2025](#)). Relatedly, prior research has shown that a word’s duration correlates with its predictability given future context ([Bell et al., 2009](#)). Along similar lines, reading times also seem to correlate with features of upcoming words (such as frequency, predictability and entropy; [Roark et al., 2009](#); [Angele et al., 2015](#); [van Schijndel and Schuler, 2017](#)). Future work should investigate these potential mechanisms to clarify the role of prosody in forward-looking processing.

6.4 Duration and Pause Show Weak Past, Strong Future Redundancy

Compared to the other prosodic features, duration and pause stand out in their relatively short time scale of redundancy with both past and future words, as well as in their relatively high MI with future words. This pattern may suggest that length-

ening a word and pausing after it primarily serves to prepare for the upcoming word, facilitating its processing by slowing the rhythm of speech. Alternatively, this pattern may stem from sentence boundaries: As discussed above, pauses are most common after sentence-final words, which are also often elongated ([Seifart et al., 2021](#); [Paschen et al., 2022](#)). The high predictability of duration and pause from future context may therefore partly reflect our models’ ability to detect a sentence-final vs. sentence-medial word, an important structural cue which could, in turn, help predict the values of these prosodic features. While we removed punctuation from the input, therefore preventing sentence-boundary information from being conveyed explicitly through punctuation marks, we did not lowercase all the data and therefore sentence-initial words are clearly detectable by their capitalization. Furthermore, sentence boundary cues could be inferred from the lexical content itself, since certain words tend to occur more frequently at the beginning or end of sentences.

7 Conclusion

Our findings reveal a fundamental asymmetry in the time scale of redundancy between prosody and linguistic information: while prosody exhibits redundancy with both past and future words, this relationship extends across a longer span for past words (3–8 words) than for future words (1–2 words). This suggests that prosody’s relationship with future words primarily reflects short-term effects such as next-word prediction, local word dependencies, or other production factors—future work should try to distinguish those explanations. In contrast, its relationship with past words operates over a broader scale, potentially serving to reinforce or highlight information that may be cognitively demanding for listeners to process in real-time communication. These results provide new insights into the role of prosody in spoken language.

Limitations

Our study has several limitations that should be considered when interpreting the results.

Data-related Limitations. The first set of limitations relates to the dataset used. Our dataset consists of audiobooks, which do not necessarily reflect natural prosody in real time communication, potentially affecting the generalizability of

our findings. Redundancy may be higher in audiobooks than in spontaneous speech, because the text is written with the assumption that it must convey all necessary information without relying on prosody. We address this concern to some extent by removing punctuation marks, which serve as a substitute for prosody in written text. Another dataset-related limitation is the sample size. Furthermore, the syntactic structures used in audiobooks may differ from those in spontaneous speech. Because pause and duration often signal syntactic boundaries, they may be particularly sensitive to such structural differences. This may help explain why pause and duration did not follow our hypotheses. Larger datasets may be required for more stable estimations, especially given that we compute 55 different values (for $w_{n,m}$, n and m from 0 to 9). Furthermore, our dataset is limited to the English language. Therefore, our findings are English-specific and may not generalize across languages, given known cross-linguistic variation in prosody. Future work should generalize this work to other typologically distinct languages (see Wilcox et al., 2025, for work using similar methods to study diversity in prosodic typology).

Estimation-related Limitations. The second set of limitations has to do with the estimation procedure. The mutual information we compute tries to approximate the true value, and is constrained by the quality of our models $p_{\theta}(p_t | w_{n,m})$. One of our modeling assumptions is the functional form of the conditional distribution of prosody given a context (namely, Gaussian, Gamma or Laplace distributions depending on the prosodic feature). However, this parametric assumption may limit the model’s performance and future work should explore alternative conditional distributions which may improve results (as done by, e.g., Wilcox et al., 2025). This assumption is particularly violated for features that manifest different distributions; pause, for instance, takes a 0 value in 89.4% of the data and may therefore be better modeled by a zero-inflated distribution. Indeed, our results for pause seem particularly noisy. Additionally, to estimate $w_{n,m}$, we provided the models with short segments of 1 to 10 words. However, the language models used here (e.g., BERT) were not pretrained on such short segments, but rather on longer spans of text; this might have impacted their efficiency in extracting the information from short segments. While finetuning likely mitigates this issue, it remains a

potential limitation. Furthermore, we train a single model for all combinations of n, m , which does not guarantee that the model is optimal for each combination separately. Finally, we observed cases where the mutual information decayed for longer contexts, which contradicts expectations from information theory, as additional context can never reduce information. This phenomenon likely stems from issues in training the models, which could be biased toward under-utilizing the available context for longer spans. Future work should address these limitations to refine our understanding of redundancy between prosody and linguistic information.

Acknowledgments

The authors would like to acknowledge Cui Ding and Giovanni Acampa for technical assistance. TIR was supported by the Poitras Center for Psychiatric Disorders Research, McGovern Institute for Brain Research, MIT.

Ethics Statement

We foresee no potential ethical concerns or risks associated with this study, although we acknowledge the inherent risks in using any AI system.

References

- Bernhard Angele, Elizabeth R. Schotter, Timothy J. Slattery, Tara L. Tenenbaum, Klinton Bicknell, and Keith Rayner. 2015. [Do successor effects in reading reflect lexical parafoveal processing? Evidence from corpus-based and experimental eye movement data.](#) *Journal of Memory and Language*, 79-80:76–96.
- R Harald Baayen, Richard Piepenbrock, and Leon Gulikers. 1996. The celex lexical database (cd-rom). *Linguistic Data Consortium*.
- Alan Bell, Jason M Brenier, Michelle Gregory, Cynthia Girand, and Dan Jurafsky. 2009. [Predictability effects on durations of content and function words in conversational english.](#) *Journal of Memory and Language*, 60(1):92–111.
- Kathryn Bock, Willem Levelt, and Morton A Gernsbacher. 1994. *Handbook of Psycholinguistics*, chapter Language production: Grammatical encoding. Academic Press, San Diego, CA.
- Thomas Lasman Botch and Emily S Finn. 2025. [Sensory context as a universal principle of language in humans and llms.](#) *PsyArXiv Preprints*.
- Mara Breen, Evelina Fedorenko, Michael Wagner, and Edward Gibson. 2010. [Acoustic correlates of information structure.](#) *Language and cognitive processes*, 25(7-9):1044–1098.

- Sarah Brown-Schmidt and Agnieszka E Konopka. 2008. [Little houses and casas pequeñas: Message formulation and syntactic form in unscripted speech with speakers of english and spanish.](#) *Cognition*, 109(2):274–280.
- Thomas H Clark, Moshe Poliak, Tamar I Regev, AJ Haskins, Edward Gibson, and Caroline Robertson. 2025. [The relationship between surprisal, prosody, and backchannels in conversation reflects intelligibility-oriented pressures.](#) *PsyArXiv preprints*.
- Jennifer Cole. 2015. [Prosody in context: A review.](#) *Language, Cognition and Neuroscience*, 30(1-2):1–31.
- Jennifer Cole, Yoonsook Mo, and Mark Hasegawa-Johnson. 2010. [Signal-based and expectation-based factors in the perception of prosodic prominence.](#) *Laboratory Phonology*, 1(2):425–452.
- Christophe Coupé*, Yoon Mi Oh*, Dan Dediu, and François Pellegrino. 2019. [Different languages, similar encoding efficiency: Comparable information rates across the human communicative niche.](#) *Science Advances*, 5(9):eaaw2594.
- Evelina Fedorenko, Anna A Ivanova, and Tamar I Regev. 2024. [The language network as a natural kind within the broader landscape of the human brain.](#) *Nature Reviews Neuroscience*, 25(5):289–312.
- August Fenk and Gertraud Fenk. 1980. [Konstanz im Kurzzeitgedächtnis - Konstanz im sprachlichen Informationsfluß?](#) *Zeitschrift für Experimentelle und Angewandte Psychologie*, 27(3):400–414.
- Victor S Ferreira. 2008. [Ambiguity, accessibility, and a division of labor for communicative success.](#) *Psychology of Learning and motivation*, 49:209–246.
- Richard Futrell, Edward Gibson, and Roger P Levy. 2020. [Lossy-context surprisal: An information-theoretic model of memory effects in sentence processing.](#) *Cognitive science*, 44(3):e12814.
- Dmitriy Genzel and Eugene Charniak. 2002. [Entropy rate constancy in text.](#) In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 199–206, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Edward Gibson. 1998. [Linguistic complexity: Locality of syntactic dependencies.](#) *Cognition*, 68(1):1–76.
- John Hale. 2001. [A probabilistic Earley parser as a psycholinguistic model.](#) In *Second Meeting of the North American Chapter of the Association for Computational Linguistics*, pages 1–8.
- Peter R Hawkins. 1971. [The syntactic location of hesitation pauses.](#) *Language and Speech*, 14(3):277–288.
- Maya Inbar, Eitan Grossman, and Ayelet N Landau. 2020. [Sequences of intonation units form a ~ 1 Hz rhythm.](#) *Scientific reports*, 10(1):15846.
- Shailee Jain and Alexander Huth. 2018. [Incorporating context into language encoding models for fmri.](#) *Advances in neural information processing systems*, 31.
- Sun-Ah Jun. 2006. *Prosodic typology: The phonology of intonation and phrasing*, volume 1. Oxford University Press.
- Daniel Jurafsky, Alan Bell, Eric Fosler-Lussier, Cynthia Girand, and William Raymond. 1998. [Reduction of english function words in switchboard.](#) In *5th International Conference on Spoken Language Processing (ICSLP 1998)*, page paper 0669.
- Jodi Kearns. 2014. [Librivox: Free public domain audiobooks.](#) *Reference Reviews*, 28(1):7–8.
- Tatsuki Kuribayashi, Yohei Oseki, Ana Brassard, and Kentaro Inui. 2022. [Context limitations make neural language models more human-like.](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10421–10436, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- D. Robert Ladd. 1996. *Intonational Phonology*. Cambridge University Press.
- Roger Levy. 2008. [Expectation-based syntactic comprehension.](#) *Cognition*, 106(3):1126–1177.
- Roger P. Levy and Tim Florian Jaeger. 2007. [Speakers optimize information density through syntactic reduction.](#) In *Advances in Neural Information Processing Systems*, pages 849–856.
- Richard L Lewis, Shravan Vasishth, and Julie A Van Dyke. 2006. [Computational principles of working memory in sentence comprehension.](#) *Trends in cognitive sciences*, 10(10):447–454.
- Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. 2017. [Montreal forced aligner: Trainable text-speech alignment using kald.](#) In *Interspeech*, volume 2017, pages 498–502.
- Clara Meister, Tiago Pimentel, Patrick Haller, Lena Jäger, Ryan Cotterell, and Roger Levy. 2021. [Revisiting the Uniform Information Density hypothesis.](#) In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 963–980, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Adam M Morgan and Victor S Ferreira. 2022. [Still no evidence for audience design in syntax: Resumptive pronouns are not the exception.](#) *Journal of Memory and Language*, 127:104368.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. [Librispeech: an asr corpus based on public domain audio books.](#) In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5206–5210. IEEE.

- Ludger Paschen, Susanne Fuchs, and Frank Seifart. 2022. [Final lengthening and vowel length in 25 languages](#). *Journal of Phonetics*, 94:101179.
- Janet Breckenridge Pierrehumbert. 1980. *The phonology and phonetics of English intonation*. Ph.D. thesis, Massachusetts Institute of Technology.
- 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, Clara Meister, Elizabeth Salesky, Simone Teufel, Damián Blasi, and Ryan Cotterell. 2021. [A surprisal–duration trade-off across and within the world’s languages](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 949–962, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tamar I Regev*, Colton Casto*, Eghbal A Hosseini, Markus Adamek, Anthony L Ritaccio, Jon T Willie, Peter Brunner, and Evelina Fedorenko. 2024. [Neural populations in the language network differ in the size of their temporal receptive windows](#). *Nature Human Behaviour*, 8(10):1924–1942. Equal first authors: Regev and Casto.
- Brian Roark, Asaf Bachrach, Carlos Cardenas, and Christophe Pallier. 2009. [Deriving lexical and syntactic expectation-based measures for psycholinguistic modeling via incremental top-down parsing](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 324–333, Singapore. Association for Computational Linguistics.
- Frank Seifart, Jan Strunk, Swintha Danielsen, Iren Hartmann, Brigitte Pakendorf, Søren Wichmann, Alena Witzlack-Makarevich, Nikolaus P Himmelmann, and Balthasar Bickel. 2021. [The extent and degree of utterance-final word lengthening in spontaneous speech from 10 languages](#). *Linguistics Vanguard*, 7(1).
- Scott Seyfarth. 2014. [Word informativity influences acoustic duration: Effects of contextual predictability on lexical representation](#). *Cognition*, 133(1):140–155.
- Cory Shain*, Hope Kean*, Colton Casto, Benjamin Lipkin, Josef Affourtit, Matthew Siegelman, Francis Mollica, and Evelina Fedorenko. 2024. [Distributed sensitivity to syntax and semantics throughout the language network](#). *Journal of Cognitive Neuroscience*, 36(7):1427–1471.
- Kim Silverman, Mary Beckman, John Pitrelli, Mori Ostendorf, Colin Wightman, Patti Price, Janet Pierrehumbert, and Julia Hirschberg. 1992. [ToBI: A standard for labeling English prosody](#). In *Second international conference on spoken language processing*.
- Antti Suni, Juraj Šimko, Daniel Aalto, and Martti Vainio. 2017. [Hierarchical representation and estimation of prosody using continuous wavelet transform](#). *Computer Speech & Language*, 45:123–136.
- Aarne Talman, Antti Suni, Hande Celikkanat, Sofoklis Kakouros, Jörg Tiedemann, and Martti Vainio. 2019. [Predicting prosodic prominence from text with pre-trained contextualized word representations](#). In *Proceedings of the 22nd Nordic Conference on Computational Linguistics*, pages 281–290, Turku, Finland. Linköping University Electronic Press.
- Jacques Terken and Dik Hermes. 2000. *The Perception of Prosodic Prominence*, pages 89–127. Springer Netherlands, Dordrecht.
- Marten van Schijndel and William Schuler. 2017. [Approximations of predictive entropy correlate with reading times](#). In *Proceedings of the Cognitive Science Society*, pages 1260–1265.
- Shravan Vasishth, Katja Suckow, Richard L Lewis, and Sabine Kern. 2010. [Short-term forgetting in sentence comprehension: Crosslinguistic evidence from verb-final structures](#). *Language and Cognitive Processes*, 25(4):533–567.
- Michael Wagner and Duane G Watson. 2010. [Experimental and theoretical advances in prosody: A review](#). *Language and cognitive processes*, 25(7-9):905–945.
- Ethan Gotlieb Wilcox, Cui Ding, Giovanni Acampa, Tiago Pimentel, Alex Warstadt, and Tamar I Regev. 2025. [Using information theory to characterize prosodic typology: The case of tone, pitch-accent and stress-accent](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.
- Lukas Wolf, Tiago Pimentel, Evelina Fedorenko, Ryan Cotterell, Alex Warstadt, Ethan Wilcox, and Tamar Regev. 2023. [Quantifying the redundancy between prosody and text](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9765–9784, Singapore. Association for Computational Linguistics.
- Jiahong Yuan, Mark Liberman, and Christopher Cieri. 2006. [Towards an integrated understanding of speaking rate in conversation](#). In *Ninth International Conference on Spoken Language Processing*.
- Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu. 2019. [Libritts: A corpus derived from librispeech for text-to-speech](#). *arXiv preprint arXiv:1904.02882*.

A Unconditional Entropies

Prosodic Feature	Unconditional Entropy
Absolute Prominence	0.536
Relative Prominence	1.355
Energy	0.815
Duration	-0.920
Pause	-5.193
f_0	3.469

Table 1: Unconditional entropies of each prosodic feature.

B Prosodic Features' Histograms

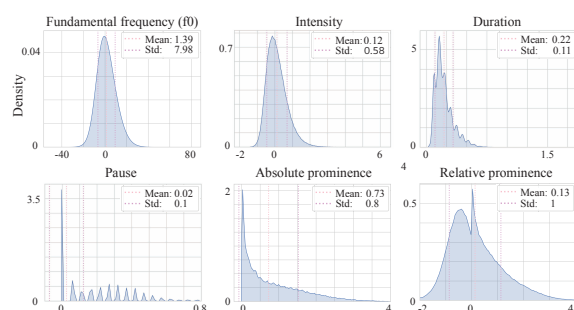


Figure 4: Histogram of prosodic features. All y -axes correspond to probability densities. (top-left) Pitch (fundamental frequency), units are arbitrary (z-scored); (top-center) Loudness, units are arbitrary (z-scored); (top-right) Duration, units are seconds per syllable; (bottom-left) Pause, units are seconds; (bottom-center) Absolute prominence, units are arbitrary; (bottom-right) Relative prominence, units are arbitrary.