

Quantitative measurement of prosodic strength in Mandarin

Greg Kochanski, Chilin Shih, and Hongyan Jing
Bell Laboratories, Lucent Technologies
600 Mountain Ave
Murray Hill, NJ 07974

March 16, 2003

Abstract

We describe models of Mandarin prosody that allow us to make quantitative measurements of prosodic strengths. These models use Stem-ML, which is a phenomenological model of the muscle dynamics and planning process that controls the tension of the vocal folds, and therefore the pitch of speech. Because Stem-ML describes the interactions between nearby tones, we were able to capture surface tonal variations using a highly constrained model with only one template for each lexical tone category, and a single prosodic strength per word. The model accurately reproduces the intonation of the speaker, capturing 87% of the variance of f_0 with these strength parameters. The result reveals alternating metrical patterns in words, and shows that the speaker marks a hierarchy of boundaries by controlling the prosodic strength of words. The strengths we obtain are also correlated with syllable duration, mutual information and part-of-speech.

Corresponding author: Greg Kochanski, gpk@alum.mit.edu
324 Third St.
Dunellen, NJ 08812-1140
Submitted to Speech Communication, 2001.

35 pages (including this one).

1 table.

15 figures.

Keywords: Intonation, tone, tonal variation, prosodic structure, metrical pattern, prosodic strength, prosody modeling, muscle dynamics, Text-to-Speech

1 Introduction

Intonation production is generally considered a two-step process: an accent or tone class is predicted from available information, and then the tone class is used to generate f_0 as a function of time. Historically, most attention has been paid to the first, high level, step of the process. We here show that by focusing on f_0 generation, one can build a model that starts with acoustic data and reaches far enough up to connect directly to linguistic factors such as part-of-speech, word length and position in the text.

Specifically, we present a model of Mandarin Chinese intonation that makes quantitative f_0 predictions in terms of the lexical tones and the prosodic strength of each word. The model is able to generate tonal variations from a few tone templates that correspond to lexical tones, and accurately reproduce f_0 in continuous Mandarin speech with a 13 Hz RMS error. The result is comparable to machine learning systems that may use more than one hundred tone templates to account for Mandarin tonal variations.

We find that some parameters of the model can be interpreted as the prosodic strength of a tone. We determine the prosodic strengths (and the values of the other global parameters) by executing a least-squares fit of the model to the time-series of f_0 from a corpus of speech data. The resulting best-fit strengths, tone shapes, and metrical patterns of words can be associated with linguistic properties. We show that strengths computed from the model exhibit strong and weak alternation as in metrical phonology [31], and the values are correlated with the part-of-speech of words, with mutual information, and with the hierarchy of the prosodic structure [49, 46, 26] such as the beginning and ending of sentences, clauses, phrases, and words.

We will also show that values of parameters from a fit to one half of the corpus match equivalent parameters fit to the other half of the corpus. Further, we can change the details of the model, and show that the values of many parameters are essentially unaffected by the change. This consistency is important because if we hope to interpret these parameters (and thus the models that contain them) as statements about the language as a whole, they must at least be consistent across the corpus and between similar models.

The model we use is described in Section 3. It is written in Soft Template Mark-up Language (Stem-ML)[24, 22], and depends upon its underlying mathematical model of prosody control. We write a Stem-ML model in terms of a set of tags (parameters) then find the parameter values that best reproduce f_0 in a training corpus. Fitting the model to the data can be done automatically.

Stem-ML calculates an intonational contour from a set of tags. Some of the tags set global parameters that correspond to speaker characteristics, such as pitch range, while others represent intonational events such as lexical tone categories and accent types. The tags can contain adjustable parameters that can explain surface variations.

Stem-ML does not impose restriction on how one define tags. In our view, a meaningful way is to use the tags to represent linguistic hypotheses such as Mandarin lexical tones, or English accent types. We call tags that define tones

or accents *templates* because they define the ideal shapes of f_0 in their vicinity. In this paper, our usage of tone tags (tone templates) corresponds directly to Mandarin lexical tone categories, and we interpret the Stem-ML *strength* parameters as the prosodic strengths of these tone templates. The actual realization of f_0 depends on the templates, their neighbors, and the prosodic strengths. We show in the paper that this treatment successfully generates continuous tonal variations from lexical tones.

Described another way, a Stem-ML model is a function that produces a curve of f_0 *vs.* time. The resulting curve depends on a set of adjustable (free) parameters which describe things like the shape of tones, how tones interact, and the prosodic strength of each syllable. When Stem-ML is generating a f_0 curve, one can set these parameters to any values, and each setting will get you a different curve. In the reverse, one can find the best values for the parameters via data fitting procedures.

We use a least-squares fitting algorithm to find the values for the parameters that best describe the data. The algorithm operates iteratively by adjusting the parameter values, and accepting steps that reduce the sum of the squared differences between the model and the data. The values of the parameters that make the summed squared difference as small as possible, for a given model, are called the best-fit (or fitted) parameters.

2 Chinese Tones

Tonal languages, such as Chinese, use variations in pitch to distinguish otherwise identical syllables. Mandarin Chinese has four lexical tones with distinctive shapes: high level (tone 1), rising (tone 2), low (tone 3), and high falling (tone 4). The syllable *ma* with a high level tone means *mother*, but it means *horse* with a low tone. Thus, in a text-to-speech (TTS) system, good pitch prediction is important not just for natural sounding speech but also for good intelligibility. There is a fifth tonal category, traditionally named *neutral tone*, which refers to special syllables with no lexical tone assignment. The pitch values of such syllables depend primarily on the tone shape of the preceding syllable.

Superficially, modeling Chinese tones seems straightforward. One might concatenate lexical tones to generate continuous speech. The challenge is that tone shapes vary in natural speech to the extent that the realized f_0 contour sometimes bears little obvious relationship to the concatenation of the tones. Figure 1 shows a Mandarin phrase *fan3 ying4 su4 du4* (“reaction time”), along with the tones from which it is constructed [57, 54]. The last three syllables are all recognized as tone 4 by native speakers, but have drastically different f_0 contours. The second syllable *ying4* has an inverted tone shape while the last syllable *du4* is lower than expected.

In previous Chinese intonation generation models, variations of a lexical tone are either ignored, or are treated as discrete classes. These discrete classes may be linked to the lexical tone by rules [51, 28], or by a machine learning method such as a neural network [3, 4]. It is not uncommon for these systems to use up




Figure 1: *Tones vs. realization.* The upper panels show shapes of tones 3 and 4 taken in a neutral environment and the lower panel shows the realization of an actual sentence containing those tones. The grey curves (or green in color display) show the templates, and the black curve (or red) shows the f_0 vs. time data.

to a hundred discrete classes to represent tonal variations. Both rule-based and machine learning methods link the lexical tone and their surface forms in an *ad hoc* manner, using factors such as lexical tones, tonal contexts, and positions in the sentence, yet neither method offers an explanation of the relations between lexical tone and their variations, or the relationship among discrete classes.

We explain the phenomenon displayed in Figure 1 as a natural consequence of tone shapes interacting via articulatory constraints. These severely distorted tone shapes occur when the shape of a weak tone is contradictory to the trajectory defined by strong neighbors. In those cases the weak tone accommodates the shapes of neighboring strong tones to maintain smooth surface f_0 contours.

Our model of Chinese intonation starts with a linguistically reasonable assumption: that all tonal variations of a lexical tone are generated from the lexically determined tonal templates. From these, we calculate f_0 at each time point as a function of the nearby templates and their prosodic strengths. We will show that this conceptually simple representation is capable of capturing the drastic tonal variations such as shown in Figure 1.

Given surface f_0 curves, and assuming that the lexical tone is known, learning the Chinese prosody description reduces to learning the lexical tone templates and the prosodic strengths of the templates.

3 Modeling Intonation

We build our model for Mandarin on top of Stem-ML[22] because it captures several desirable properties. A positive feature of Stem-ML is that the representation is understandable, adjustable, and can be transported from one situation to another.

Unlike most engineering approaches, this model cleanly separates into local (word-dependent) and global (speaker-dependent) parameters. For instance, one can generate acceptable speech by using the templates of one speaker with prosodic strengths from another[54], where a female speaker’s tone templates were used as part of a model to predict a male speaker’s f_0 contours. Unlike some descriptive models, we predict numerical f_0 values, and so our model is subject to quantitative test. Few other approaches to intonation have all these properties.

3.1 Concepts Behind the Model

Stem-ML brings together several ideas into intonation modeling:

- we assume that people plan their utterances several syllables in advance,
- we assume that people produce speech that is optimized to meet their needs,
- we apply a physically reasonable model for the dynamics of the muscles that control pitch and skilled movements [19, 40], and

- we introduce the concept of prosodic strength, a continuous parameter associated with linguistic units such as syllable, tone, and word, to control variations.

Pre-planning in speech was first shown in terms of the control of inhaled air volume [65, 66, 64, 37]: people will inhale more deeply when confronted with longer phrases, hence we see a positive correlation of longer phrase and higher initial f_0 [52]. This fact implies that at least a rough plan for the utterance has been constructed about 500 ms before speech begins. As another example, Figure 8 in Bellegarda *et al.* [1] shows that in an upwards pitch motion, the rate of the motion is reduced as the motion becomes longer, presumably to avoid running above the speaker’s comfortable pitch range. We take this as evidence for pre-planning of f_0 over a 1.5 second range, at least in practiced, laboratory speech.

Next, we assume that speech is optimized for the speaker’s purposes. The idea of representing muscle motions as the solution to an optimization problem has been developed in the biomechanics literature [73, 18, 48], and there have been comparisons of these models to actual movements [12] and to electromyogram signals [7]. Nelson [40] modeled jaw movement and arm movement during violin bowing and showed that skilled movements are influenced by minimum-cost solutions which balance performance trade-offs between competing objectives.

Speech is a skilled movement, and native speakers of Mandarin are skilled practitioners of tonal production. A speaker of Mandarin has the opportunity to practice and optimize all the common 3-tone or perhaps 4-tone sequences, even if one assumes that each tone needs to be practiced at several distinct strength levels. For instance, if we count tone N-grams in the ROCLING Chinese corpus [6], we find that the most common 64 of the 179 tone 3-grams cover 90% of the corpus (we count phrase boundaries in the N-grams). Likewise, the most common 358 of the 881 4-grams cover 90% of the corpus. A speaker could practice the common tonal combinations in an hour of speech.

A more realistic model, such as the one we propose in this paper, would add a strength parameter to each tone, but one could then still expect to practice the common tonal combinations with several levels of strength in a short time.

The question then arises, “optimal in what sense?” It has been proposed that optimality be defined by a balance between the ability to communicate accurately and the effort required to communicate [41, 24, 22], and such models have been applied by ourselves [24, 55, 23] and others [45, 44].

Our works extend the concept of optimizing communication needs and the ease of articulatory efforts to account for tonal variations in continuous speech [24, 22]. The optimal pitch curve is the one that minimizes the sum of effort plus a scaled error term. Certainly, when we speak, we wish to be understood, so the speaker must consider the error rate on the speech channel to the listener. Likewise, much of what we do physically is done smoothly, with minimum muscular energy expenditure, so minimizing effort in speech is also a plausible goal. Different from most previous works, our view is that the

trade-off relations between different objectives change dynamically during continuous speech. We introduce a scale factor (the prosodic strength) to describe the shifting dynamics of how the speaker optimize communication needs and articulatory efforts in continuous speech.

3.2 Mathematical Definition of Model

The assumption that pitch is produced to optimize the sum of effort plus error can be converted into a quantitative mathematical model. We will describe the equations below, and the variables involved will be defined in Table 1.

The effort expended in speech, G (Equation 1), is based upon the literature on muscle dynamics and energetics [61, 72, 12, 71, 68]. Qualitatively, our effort term behaves like the physiological effort: it is zero if muscles are stationary in a neutral position, and increases as motions become faster and stronger. Minimizing G tends to make the pitch curve smooth and continuous, because it minimizes the magnitude of the first and second derivatives of the pitch.

Note that we do not depend on the assumption that the effort term is an actual measurement of the energy expenditure in the muscle. The effort term could very well be a measure of competition for resources in the central nervous system, could be due to neural feedback loops local to the muscle (similar to the Equilibrium Point Hypothesis [25, 11]) or could be entirely phenomenological. It does seem, however, that the effort term is not just a way to express the nonzero response time of a muscle fiber: measurements of single-fiber twitches (*i.e.* the force-*vs*-time curve of a single muscle fiber triggered by a single nerve impulse) show a contraction time of ≈ 19 ms [34], which is too short to account for inverted tone shapes and other phenomena that can last for 100 ms or more.

The error term, R (Equations 2 and 3), behaves like a communications error rate: it is zero if the prosody exactly matches an ideal tone template, and it increases as the prosody deviates from the template. The choice of template encodes the lexical information carried by the tones. The speaker tries to minimize the deviation, because if it becomes too large, the speaker will expect the listener to mis-classify the tone and possibly misinterpret the utterance.

Stem-ML makes one physically motivated assumption. It assumes that f_0 is closely related to muscle tensions [38]. There must then be smooth and predictable connections between neighboring values of f_0 because muscles cannot discontinuously change position. Most muscles cannot respond faster than 150 ms, a time which is comparable to the duration of a syllable, so we expect the intonation of neighboring syllables to affect each other. Because our model derives a smooth f_0 contour from muscle dynamics, our model is an extension of those of Ohman[42], Fujisaki[15], and Lindblom[39, 33], and is similar to that of Xu and Sun[70].

In Stem-ML, a “tag” is a tone template, along with a few parameters that describe the scope of the template and how the template interacts with its environment. It corresponds to the mathematical description of an intonation event (*e.g.*, a tone or an accent). Tags have a parameter, *type*, which controls whether

errors in the shape or average value of the pitch curve are most important. In this work, the targets, y , consist of a tone component riding on top of the phrase curve, p .

In order to efficiently solve the optimization problem, and calculate the surface realization of prosody, we write simple approximations to G and R so that the model can be solved efficiently as a set of linear equations:

$$G = \sum_t \dot{e}_t^2 + (\pi \cdot \text{smooth}/2)^2 \ddot{e}_t^2 + \text{adroop}^2 \cdot e_t^2 \quad (1)$$

$$R = \sum_{k \in \text{tags}} s_k^2 r_k \quad (2)$$

$$r_k = \sum_{t \in \text{tag } k} \cos(\text{type} \cdot \pi/2)(e_t - y_{k,t})^2 + \sin(\text{type} \cdot \pi/2)(\bar{e}_k - \bar{y}_k)^2, \quad (3)$$

where

$$\bar{e}_k = \sum_{t \in \text{tag } k} e_t / \sum_{t \in \text{tag } k} 1 \quad (4)$$

and

$$\bar{y}_k = \sum_{t \in \text{tag } k} y_t / \sum_{t \in \text{tag } k} 1. \quad (5)$$

and

Finally, f_0 is e , scaled to the speaker’s pitch range:

$$\hat{f}_0 = g(e, \text{add}) \cdot \text{range} + \text{base} \quad (6)$$

the scaling allows p and e to be dimensionless quantities, typically between 0 and 1. The function $g()$ handles linear ($\text{add} = 1$) or log ($\text{add} = 0$) scaling, and has the properties that $g(e, 1) = e$ for any e , and that $g(0, \text{add}) = 0$ and $g(1, \text{add}) = 1$ for any add .

Figure 2 shows how the G (effort) term depends on the shape of e . The curves we show all go through the same set of pitch targets (dashed circles). The G values increase with the RMS curvature and slope of e . In this case, optimal pitch curve has the smallest value of G , G_1 .

Note that there are two distinct optimizations in this paper, and they should not be confused. First (Section 3.2), we represent the Stem-ML model as an optimization problem, minimizing *effort+error* to find f_0 as a function of the model parameters. This first minimization is actually done analytically, to convert the Stem-ML model into a set of linear equations that are solved by matrix techniques.

Second (Section 4.2), we adjust the parameters to minimize the difference between the model and the data. This gives us best-fit values for the parameters that best describe the data. This second minimization treats the evaluation of the Stem-ML model as a black box, calculating many models to find the best fit.




Figure 2: Schematic diagram showing the dependence of G on the shape of the pitch curve. The large, left axis shows values of G (speech effort) for each of the displayed curves ($G_1 \dots G_5$). Each small axes show sample curves of pitch as a function of time. The resulting Stem-ML pitch curve is the one with the optimal (smallest) value of $G + R$. Because we have chosen $R = 0$ in this example, the solution here is G_1 , the one with the smallest G .

Symbol	Location	Meaning
add†	Eq. 6	Controls the mapping between e and f_0 . See $g()$.
adroop†	Eq. 1	Rate at which e droops toward the phrase curve in the absence of a tag.
base†	Eq. 6	The speaker’s relaxed f_0 .
smooth†	Eq. 1	Response time of muscles.
type†	Eq. 3	Is tone defined by its shape (0) or f_0 value (1).
$M_{L,i}$	Eq. 8	Metrical pattern of the i^{th} syllable in a L syllable word.
s_k †	Eqs. 2 , 7, 8	Strength of syllable k .
S_w	Eqs. 8	Strength of word w .
atype	Eq. 7	Controls how the size of the template depends on the strength of a syllable.
ctrshift	§4.3	Position of center of template relative to center of syllable.
wscale	§4.3	Width of a tone template, relative to a syllable.
P, D, d	Eq. 9	Parameters defining the phrase curve.
f_0	many places	Measured pitch.
\hat{f}_0	Eq. 6	Modeled pitch.
p †	Eq. 9	Phrase curve.
e †, e_t	§3.2	Emphasis, <i>i.e.</i> , \hat{f}_0 relative to the speaker’s range.
\bar{e} †	Eqs. 3, 4	Mean emphasis over the scope of a tag.
y †, y_t	§3.2	Tone template.
\bar{y} †	Eqs. 3, 5	Mean value of a tone template.
G †	Eq. 1	Effort expended in realizing the pitch contour.
r_i	Eq. 3	The summed error for syllable i between the template and the realized pitch.
R †	Eq. 2	The summed error for an utterance between the ideal templates and the realized pitch contour.
$g()$ †	Eq. 6	Function to map between subjective emphasis (e) and objective f_0 .

Table 1: *Definitions of parameters and variables used in this paper. Daggers (†) denote parameters defined more fully in [22].*

As an additional complication, we then take some of the best-fit parameter values (specifically the values of the prosodic strength parameters) and fit *them* with an additive linear model (Section 5.5). This final fit helps us understand which linguistic features have the most influence on the strength of a syllable.

3.3 Prosodic Strength

Effort is ultimately measured in physical units, while the communication error probability is dimensionless. Since one can only sensibly add numbers with the same units (*e.g.*, 1 kilogram + 1 meter = ?), a scale factor is needed to convert one into the units of the other. This scale factor, s_k (in Equation 2), can vary from a tone, a syllable or a word to the next, and we identify it with the prosodic strength.

If a syllable’s strength is large, the Stem-ML pitch contour will closely approximate the tone’s template and the communication error probability will be small. In other words, a large strength indicates that the speaker is willing to expend enough effort to produce precise intonation on a syllable. On the other hand, if the syllable is de-accented and its strength is small, the produced pitch will be controlled by other factors: neighboring syllables and ease of production. For prosodically weak syllables, minimizing the effort term will have the most effect: when s_k is small, smoothness becomes more important than accuracy. The listener then may not be able to reliably identify the correct tone on that syllable. Presumably, the listener can infer the tone from the surrounding context.

The concept that strength is related to how carefully speech is articulated was discussed by Browman and Goldstein [2], in the context of phoneme changes in casual speech. Flemming [14, 13] discusses optimization models with continuous parameters (into which class this model falls), and their relationship with Optimality Theory [47].

Traditionally, prosodic strength is expressed as abstract categories S (strong) and W (weak) in metrical phonology [31], where one of the goals is to capture the rhythmic alternation in natural sentences even though words typically do not come in iambic or trochaic pairs. One can build a prosodic structure with strong and weak nodes to describe sentence prosody in relative terms.

Our model is related to conventional views of accents and intonation, except that we consider strength to be a continuous parameter associated with a word or a syllable. We suggest that listeners *might* treat strong tones as categorically different from weak tones, so these strength measurements might be equivalent to the presence or absence of accents (strong implies present). The strength numbers are associated with a particular rendition of the sentence. They vary somewhat even among utterances that were spoken with the same intent, but they seem to vary more between utterances where the sentence focus, the intonation type, or other prosodic features differ.

4 Experiment

4.1 Data Collection

The corpus was obtained from a male native Mandarin speaker reading paragraphs from newspaper articles, selected for broad coverage of factors in the text that are associated with prosodic effects, including tonal patterns in the beginning, medial, and final positions of utterances, phrases, and words. To select sentences from a corpus, we used the greedy algorithm described in [63]. Pause and emphasis were transcribed manually after text selection and recording. A complete description of the factors, procedures, and evaluation of the algorithm were described in [53].

We fit two subsets (10 sentences each, 347 and 390 syllables), that were randomly chosen from the corpus. The speaking rate was 4 ± 1.4 syllables per second, with a phrase duration of 1.2 ± 0.7 s. We define a phrase as speech materials separated by a perceptual pause. We measured these pauses acoustically, and found that the speech power dropped by at least 10 dB relative to a 50 ms window on either side in 94% of the pauses, and the median duration of pauses was 240 ms.

Tones were identified by automatic text analysis, including the tone sandhi rule in (Shih [50]), then checked by two native speakers. Neutral tones were manually identified prior to fitting, because they cannot be reliably identified from a dictionary. Phone, syllable, and phrase boundaries were hand-segmented, based on acoustic data.

We computed f_0 with an automatic pitch tracker [62], then cleaned the data by hand, primarily to repair regions where the track was an octave off. If uncorrected, the octave errors would have doubled the ultimate error of the fit, and systematically distorted tone shapes.

Because word boundaries are not marked in Chinese text, different native speakers can assign word boundaries differently. Even so, the concept of a word is present, and is reflected in the prosody. We obtained word boundaries independently from three native Mandarin speakers: A, J, and S (J and S are authors). All three had a generally consistent segmentation of the text into words. Pairwise comparison indicates that J and S have the highest level of agreement: J identified 395 word boundaries, S identified 370 boundaries, 99% of which were also identified by J. A identified 359 word boundaries, of which 98% were also marked by J and 92% were also marked by S.

Most disagreements were related to the granularity of segmentation: whether longer units were treated as single words or multiple words, and whether neutral tone syllables were attached to the preceding words. The labelers exhibited strong and consistent personal preferences on words that could be segmented more than one way. Labeler A had the longest words, 2.04 syllables on average. J and S divided words at a finer granularity: S's words averaged 1.98 syllables, and J's words averaged 1.86 syllables per word. Labeler A consistently cliticized neutral tone syllables to the preceding word, while the other two labelers rarely did so.

We also created a random word segmentation (called “R”). The random segmentation provides a check that the metrical patterns (Section 5.4) we found are indeed significant.

4.2 Fitting

The Stem-ML model is built by placing tags on tone templates, with adjustable parameters defining the tag shapes and positions (details below). We built several different models, focusing on models with one parameter (prosodic strength) for each word, plus a set of 36, 39, or 42 shared parameters. The models discussed here have between 210 and 246 free parameters, or an average of 0.6 parameters per syllable. The parameters that define the strength of words are correlated only with a few neighbors, but the shared parameters are correlated with everything.

The algorithm obtains the parameters’s values by minimizing the RMS frequency difference between the data and the model. Unvoiced regions were excluded. We fit the two subsets separately, to allow comparisons.

We used a Levenberg-Marquardt algorithm [29, 35] with numerical differentiation to find the parameter values that give the best fit. The algorithm requires about 30 iterations before the RMS error and parameter values stabilize.

Levenberg-Marquardt, like many data fitting algorithms, can become trapped in a local minimum of χ^2 , and may miss the global best-fit. If we start the fit with parameter values randomly chosen from “reasonable” ranges, it will converge to what we believe to be the global minimum in about 1 in 4 tries. Consequently, we believe there are only a small number of minima. The global minimum seems to be characterized by values of *adroop* < 1 (see Table 3.2), and its χ^2 is often 10% smaller than the next best minimum. Convergence to the global minimum seems fairly reliable if a fit is started with values of the shared parameters taken from a previous successful fit, even if the model or data subset differs, and even if the strengths are initialized randomly.

4.3 Mandarin-specific Model

Our model for Mandarin is a more predictive, stronger model than bare Stem-ML, and is stronger than our previous works on Mandarin tone modeling [21] where an independent strength parameter is fitted for every syllable.

The current model, which is an extension of [23], starts with a Stem-ML *stress* tag specifying the lexical tone templates associated with the syllable. The syllabic strength is tied to the strength of the word via metrical patterns. This model fits less parameters but still achieve comparable results.

We assume that each of the five lexical tone classes is described by one template. A template is defined by 5 (2 for neutral tones) pitch values, spaced across its scope. It is merely stretched (in time) and scaled (changing its pitch range) to describe all syllables which have that tone. Each tone class has a Stem-ML *type* parameter. Tone classes also have an *atype* parameter, which

controls how the template scaling depends on each syllable’s strength. The pitch excursions of the template on syllable k are scaled by a factor

$$F_k = atype \cdot s_k^{|atype|} \quad (7)$$

before the Stem-ML tag is generated. Thus, if $|atype| > 1$, the pitch range of the generated Stem-ML tag will change a lot for a small change in strength, while if $|atype| < 1$, the pitch range of the tag will be relatively independent of strength.

In the general Stem-ML model, each tone template has a strength value, which controls how it interacts with its environment. In a pitch generation process this gives us enough parameters to describe a pitch contour [54, 22]. In the reverse process of fitting the strength values from data [21], we found that the data cannot support the estimation of one parameter per syllable and that the fitting process was often trapped in a local minima. Increasing the size of the database would not help to disambiguate syllable strengths, since the number of strength parameters to be estimated increases with the number of syllables in the database.

However, we noticed that syllable strength within a word is not independent of each other, and that they tend to exhibit alternating metrical patterns. If there are consistent strength patterns within a word, then we should be able to describe the observed prosody with word-level strength and a few metrical patterns. In the current model, we allow words of different length to have their own metrical patterns. This turns out to be a viable method. Compared to the syllable model, the word model reduces the number of parameters by 40% while maintaining the same level of performance.

In the model, each word has its own adjustable *strength* parameter, S_w , and we derive strengths for each syllable via

$$s_{w,i} = S_w \cdot M_{L(w),i}, \quad (8)$$

where $s_{w,i}$ is the strength of the i^{th} syllable of word w , $M_{L,i}$ is the metrical strength of the i^{th} position in a word of L syllables, and $L(w)$ is the length of word w . That means we allow the strengths of words to vary independently¹ while restricting the strength relationship of syllables within the word. Each word is associated with a word strength and the strengths of the component syllables are derived from the word strength and the metrical pattern. This metrical pattern is assumed to be the same for all words that have the same number of syllables. The word strengths, S_w , are the only place in our model where linguistic information can influence the f_0 contour beyond selection of the lexical tone. In Section 4.2, the word strengths will be adjusted to fit the model to the data.

¹One alternative to the assumption that each word has its own strength parameter would be to assume that (for example) all sentence-initial words have the same strength. Instead, we chose to let each word have its own strength so that we could search for relationships among the strengths we obtain by fitting the model to a corpus of data.

There are several parameters that are shared by all syllables. Two parameters describe the scope of templates: *ctrshift* is the offset of the template’s center from the syllable’s center, and *wscale* sets the length of the template relative to the syllable. Phrases are described by a straight-line phrase curve:

$$p(t) = P \cdot L - (D \cdot L^d) \cdot t, \quad (9)$$

where t is time, $p(t)$ is the phrase curve, and L is the length of the phrase (in seconds). All phrase curves share three parameters: D , the declination rate; d , the dependence of the declination on the sentence length; and P , which tells how the initial height of the phrase curve depends on sentence length. To complete the model, We used Stem-ML *step_to* tags to implement the phrase curve, and *phrase* tags were placed on phrase boundaries. Four other Stem-ML parameters control overall properties: *adroop*, *add*, *smooth*, and *base*.

We created and fit a set of different models to the data, using a factorial design. We used two subsets of the corpus times the four different word segmentations (A, J, S, R) times three different parameterizations. We refer to the three parameterizations as ‘w’, ‘wA’, and ‘wAT’. These form a nested set of models with a decreasing number of parameters. In the ‘w’ parameterization, each tone class has its own *atype* and *type* parameters: we allow tone templates to scale differently as the strength increases, and we allow some tones to be defined by their shape while others are defined by their position relative to the phrase curve. In the ‘wA’ parameterization, we force all tone classes to share one *atype* parameter, so that all tone templates scale with the same function of strength. Finally, in the ‘wAT’ parameterization, we force all tones to share the *type* parameter, so all tone classes exercise the same trade-off between control of shape and control of average pitch.

Of these 24 models, 15 converged to comparably small χ^2 values, and we believe those sets of parameters to be globally optimal for their model. Of the remainder, several were not attempted, due to limits on the available CPU time, and the rest seemed to land on local minima, with χ^2 values more than 30% larger than the global minimum.

5 Analysis of Best-Fit Parameters

5.1 Results of Fit

Overall, our word-based models fit the data with a 13 Hz RMS error, approximately 1.5 semitones. In Figure 3, we show the beginning of an utterance from the best-fit model (subset1-J-wA). In Figure 4, we show the phrase with median error from that model, and in Figure 5, the phrase containing the worst-fit pair of syllables in the worst model (subset2-S-wAT). Generally, the worst-fitting syllables tend to be the ones with the largest and fastest pitch excursions. These are conditions where Stem-ML’s approximation to muscle dynamics may break down, or where the simple approximation that we use to estimate the error




Figure 3: *The beginning of an utterance. Fit (solid) vs. data (dots). Syllable centers are marked with vertical dashed lines. The tones are marked above (in an open face font) and fitted prosodic strength, s_i , is marked as a subscript. (Syllable strength is calculated from word strength and metrical patterns.) The text is marked below. Stem-ML phrases, as defined by pauses, are marked with “┆”.*

between templates and the realized pitch curve may be furthest from the actual perceptual metric.

These models explain 87% of the variance of the data, and much of the rest may be explainable by phoneme-dependent segmental effects [27, 59]. Thus, essentially all the prosodic information in the f_0 contour is captured by the parameter values we obtain from the fits. Of the parameters, only the word strengths have localized effects so that only they can capture localized prosodic features like emphasis, focus, and marking of sentence structure. We expect, then, that the word strengths resulting from the Stem-ML analysis are nearly a complete description of Mandarin prosody². The rest of the paper will attempt to show that they are simple, useful descriptions of prosody in addition to being nearly complete descriptions.

We can show that the strength values that we obtain are robust against small changes in the assumptions that define the model. For example, Figure 6 shows a plot of syllable strengths obtained for the first subset with the S-wA model, plotted against strengths obtained from the J-wAT model. Despite the different word segmentations and the different sets of shared parameters the strength values are quite consistent. Comparisons between different models using the same segmentation are even closer. Nearly all of the values fall on a narrow band about a smooth curve that maps the strength from one fit to the other. This mapping is the result of differences of shared parameters (most importantly *atype*) between the two fits. The strength values that are least reproducible are single syllable words, especially single syllable neutral tones.

For Stem-ML to be a model of a language, instead of just a scheme for efficiently coding f_0 contours, we should be able to correlate the results of the fit with linguistic features. In the following sections, we will discuss the results of the fit and see how they correlate with linguistic expectations.

²Prosody as it affects intonation, not necessarily duration or articulation.




Figure 4: *Typical fit (solid) vs. data (dots), for model subset1-J-wA. Displayed as above.*




Figure 5: *Phrase containing the worst-fit pair of syllables in the worst model. Displayed as above.*




Figure 6: Comparisons of strength values of syllables between the S-wA and J-wAT models. The strength of most syllables is measured nearly independently of the details of the model.

5.2 Analysis of Phrase Curve

Our phrase curve is Equation 9: simple linear declination. We included a phrase curve in the model and fit it, because phrase curves are a common feature in many qualitative descriptions of intonation. However, the data shows no evidence that the phrase curve is necessary, and we see no systematic declination. Neither $P = -4 \pm 3 \text{ Hz}\cdot\text{s}^{-1}$ nor $D = 0 \pm 4 \text{ Hz}\cdot\text{s}^{-1}$ is very large, and neither is substantially different from zero (the error bars are derived from the standard deviation of the values of equivalent parameters among the models).

In our model of Mandarin, a positive D would correspond to a systematic decrease in f_0 during a phrase. This is distinguishable from a systematic decrease in strength, which causes the magnitude of f_0 swings to become smaller as the phrase progresses. Our phrase curve roughly corresponds to the reference line of Liberman and Pierrehumbert [30], and our strength is similar to the difference between their base line and their top line.

5.3 Analysis of Tone Shapes

First, the fitted scope of the templates is well matched to a syllable. The best fit templates are $68 \pm 4\%$ of the length of their syllable, and the centers of the tone templates are just $18 \pm 8\%$ of the length of the syllable after the center. This matches well with the intuition that tones are associated with syllables (but see




Figure 7: *Modeled shapes of isolated tones. The shapes match standard descriptions, and interact to reproduce continuous speech. The two dashed vertical bars mark the syllable boundaries, and dots mark the boundaries of the tone’s template in each of the models (random segmentations were excluded). Each tone was calculated with its strength set to the median of all the strengths in the corpus.*

Xu and Wang [69]).

Figure 7 shows the shapes of the four main Mandarin tone templates, calculated for each of our models. The tone shapes are consistent among different models, and across subsets. Overall, the shapes match standard descriptions of Mandarin tones. The symmetry between tones 1 and 3 and tones 2 and 4 is striking, and was in no way imposed by the analysis procedure. The four tones appear to have evolved to be nearly as different as possible, under the constraint that the pitch changes can be accomplished by human muscles within one syllable [70].

5.4 Analysis of Metrical Patterns

The RMS error from these word-based models, 13 Hz, is nearly the same as the 12 Hz RMS error we obtain from similar models [21] that do not impose a metrical pattern, but instead allow the strength of each syllable to vary independently. Clearly, the metrical patterns in the words are successful at capturing much of the strength variation from syllable to syllable within a word. The models in this paper have approximately half as many free parameters (and thus are more predictive) than our earlier models [21], and yet still provide an accurate representation of the actual speech.




Figure 8: *Metrical patterns for the J and S segmentations of 4, 3, and 2 syllable words. The words (ω) are plotted as trees, and syllables (σ) are represented by the black arrowheads at the end of the lines. The vertical position of the i^{th} arrowhead is proportional to the metrical strength of the i^{th} syllable: $\log(M_{L,i}) \cdot atype^{1/2}$. Differences of $\log(M)$ among leaves and nodes are shown numerically, with the parenthesized number showing the uncertainty in the last digit, as determined from the scatter among different models. The patterns for four syllable words have larger errors, as they are rare: they are drawn with double arrows to display the range of fitted solutions.*

Figure 8 shows a tree diagram of the metrical patterns we observe. A direct comparison of the metrical patterns from different models is not useful, because *atype* differs from model to model. The metrical patterns are really measures of relative syllable strength, and *atype* controls how the strength is related to the amplitude of the template. Looking back at Equation 7, we see that tags with a small value of *atype* will need a broad range of strengths to get a relatively small change in the pitch excursion, and *vice versa*. This happens because the pitch excursion is proportional to F_k (Equation 7), thus it increases at least as fast³ as the strength raised to the power *atype*. Since the pitch excursions are fit to the data, we expect that models with a small *atype* will have the largest range of strengths. This correlation between *atype* and $\text{variance}(\log(s_k))$ is indeed strong. In order to make comparisons clearer, we scale the metrical patterns, $\log(M_{L,i})$, by $atype^{1/2}$ to make the strengths of different models comparable. Recall that *atype* is a global parameter, so that this scaling does not change the shapes or the metrical patterns, nor the relationship between metrical patterns for different words.

All the real segmentations (A, J, S) show a clear strong-weak pattern for two syllable words. This means that the initial syllable’s tone is realized more precisely, and the f_0 swings will tend to be larger. Although the details vary by model, and depend on the neighboring words, our results indicate that RMS swings on the first syllable should be about 30% larger than the second syllable. While it has been generally expected that Mandarin words would show a consis-

³It will actually increase faster, because as the strength increases from zero, the f_0 curve will tend to follow the templates more and more closely. Note that this argument applies to typical pitch excursions, and is not necessarily true for each syllable: the excursion in a particular syllable depends on its tone class and the strengths and tone classes of its neighbors.




Figure 9: *Metrical patterns for the A-segmentation, plotted as above.*




Figure 10: *Metrical patterns for random word segmentation, plotted as above. As expected, the residual patterns are weak and inconsistent.*

tent metrical pattern, previous expectations [32] tended more to a weak-strong pattern, based primarily on evidence from duration and perceptual judgments.

In the A, J, and S segmentations, three-syllable words are predominantly left-branching. Because of this, we applied the same metrical pattern to all three-syllable words, and did not attempt to see if words with different internal structure had different metrical patterns. Again, we see strong-weak patterns at both levels of the metrical hierarchy, though the patterns are weaker than the two-syllable case.

All of the four-syllable words in the data could be broken up into pairs of two-syllable words. We know this both from comparison of the J and S segmentations, where the primary difference was just such a splitting, but also from plausibility judgments of the labelers. Consequently, we adopted the metrical tree shown in Figure 8. Expressed on that tree, we again get strong-weak patterns at both levels.

In Figure 9, we show the metrical trees from the A-segmentation. While the patterns differ in detail because of A’s tendency to attach particles to words, the overall picture is similar to the J and S segmentations.

Figure 10 shows the corresponding pattern for a random word segmentation (R). As expected, the R-segmentation does not yield a strong metrical pattern, because there is no consistent relationship between the spoken words and the random model. Further, the R-segmentation does not give as good of a fit to the data: the χ^2 values are 11% to 21% above the corresponding models with real (A, J, or S) segmentations. This change in χ^2 is substantial: it is an order of magnitude larger than necessary for statistical significance at the 1% level,

even if one makes allowance for correlations among the f_0 measurements.

Our results are consistent with the prediction of metrical phonology [31]. We find an alternating strong/weak relation within bisyllabic words. This pattern repeats in a four syllable word with a higher order hierarchical relation that also shows strong/weak alternation.

5.5 Analysis of Word Strengths

The strengths that result from the above fitting process can be correlated with linguistic factors. We considered three features: the number of syllables in the word, the position of the word in the utterance, and the part-of-speech of the word. We did not include any semantic features, and syntax was only included through part-of-speech and (to some extent) through our definition of pauses. Also, there was no feature in the model equivalent to the concept of “the focus of a sentence.” We limited ourselves to features that could be derived from the text alone (with the exception of phrasal pauses). Phrasal pauses seem to be clear enough to a listener, and their perception seems relatively independent of the pitch, so we tolerated the slight circularity introduced by their use as features.

We then fit the strengths with a trimmed linear regression [36] to separate out the effects of the different factors. The model for the observed word strength, S_w , is

$$\hat{S}_w = c_0 + \sum_i c_i \cdot f_{i,w}, \quad (10)$$

where \hat{S}_w is the modeled strength. In the sum, i ranges over the features described below, $f_{i,w}$ is 0 or 1, depending on whether the i^{th} feature is present on word w , and c_i is the regression coefficient for the i^{th} feature. Coefficient c_0 shows the strength of words without any features. In this trimmed linear regression, we find the regression coefficients that minimize $\sum'_w (S_w - \hat{S}_w)^2$, where the primed sum excludes the 5 largest errors. Excluding a handful of wild points prevents the regression from being dominated by words whose strength cannot be accurately measured (*i.e.*, monosyllabic words that have a neutral tone), and leads to a much more reliable result. Such outliers comprise about 2% of the strength measurements, and can be clearly seen in Figure 6. We calculated this regression separately for each of our models. In Figure 11, we plot the distribution of the regression coefficients across models for each factor.

Overall, predicting strength via this linear model reduces the median absolute deviation by 17%: these factors do not provide more than a partial prediction of the strengths or f_0 . Again, we use a robust estimator like median absolute deviation instead of variance to reduce the effect of the outliers. If the strength distribution were Gaussian, this regression would have Pearson’s $r = 0.31$.

We found that:




Figure 11: *Relation between strength and word positions. Each box shows the range of the data (the shaded region extends from the 25th and 75th percentiles), the median (white stripe in the box), and outlying points (brackets on the border). All boxes are referenced to words that are not at any kind of boundary, which are shown as the zero line.*

5.5.1 Words at the beginning of a sentence, clause, or phrase have greater strengths than words at the final positions.

Figure 11 shows the regression coefficients at different positions. We define a sentence as a grammatical utterance that is marked with a period at the end, a clause as a subset of a sentence that is marked by a comma, and a phrase as a group of words that are separated by pause.

The hierarchy of linguistic units is displayed with strengths that increase with the size of the unit. Note that the regression coefficient of words not at a boundary is defined to be zero, and that zero (horizontal line) neatly divides the initial words of units (sent-i, clau-i, phr-i) from the final words of the units (phr-f, clau-f, sent-f). These results are consistent with previous findings that speakers use high pitch to mark discourse initial segments [17].

5.5.2 Nouns and adverbs typically have more strength than words of other part of speech, and particles have the lowest strengths.

Figure 12 shows the regression coefficients for different parts-of-speech (Equation 10). As we can see, adverbs on average have a greater strength than words of other parts-of-speech. The strengths for nouns, verbs, and conjunctions are slightly weaker than that for adverbs and their strengths are close to each other. In contrast, the strength for particles (*e.g.*, neutral tones) are much weaker than that for other parts-of-speech. This may be related to the low information content of function words. These results are consistent with previous results which




Figure 12: *Relation between part-of-speech and strength. Regression coefficients for Equation 10 are shown.*

were obtained using human-annotated accents (e.g. Hirschberg [16]).

5.5.3 Words with more syllables have greater strength than words with smaller number of syllables.

Figure 13 shows the regression coefficients (Equation 10) for strengths for words of different lengths. The regression coefficient for 3-syllable words is defined as zero, which is shown as the horizontal line in the figure. The plot shows three populations of monosyllabic words, bisyllabic words, and longer words, where word strength increases as a function of word length. The weak status of a monosyllabic word is consistent with previous linguistic observations, where such phenomenon prompted the postulation of the *Monosyllabic De-stressing Rule* (Selkirk [49]).

The correlations between strength in our Stem-ML models and the above linguistic features suggest that the strengths indeed represent the prosodic relations of syllables and words. This has two consequences: First, this knowledge allows us to use features such as position, part-of-speech, and number of syllables in word to predict the strength of a word, and thus improve prediction of f_0 in a Mandarin speech synthesizer. Second, it may be possible to apply it to speech recognition systems, so that the recognizer can detect word boundaries and to deduce whether a word is being emphasized (see Shih *et al.* [56] for discussion).

5.6 The Correlation of Strength and Duration

We can also calculate the correlation between the fitted strength values with acoustic measurements such as duration. Many duration studies reported a




Figure 13: *Relation between strength and the number of syllables in a word. The boxes are plotted relative to three-syllable words, which are shown as the zero line.*

lengthening effect of stressed vowels [20, 8]. It is generally expected that, everything else being equal, strong words would have longer duration than weak words.

We calculated the correlation scores between strength and duration in our models, excluding the models using random word segmentation. Outliers are trimmed by excluding the 5% of the population that is farthest from the regression line that defines the correlation, again using a trimmed linear regression. The mean correlation scores of these models are 0.40 in the sentence final position, and 0.27 in the non-final positions.

Figure 14 show the strength/duration correlation from one of the models. The left panel shows the population in the sentence final position, and the right panel shows the population in the non-final position. All sample points are used in these plots, where the correlation scores are 0.45 in the sentence final syllables and 0.34 in the non-final syllables.

Phrase final syllables are subject to final lengthening effect [9] and this trend is reflected in the discrepancies between the strength values of final and non-final populations. The phrase final population is characterized by lower strength values and longer duration.

5.7 Mutual Information and Observed Metrical Structure

Why should we observe word initial syllables with higher strength than other syllables in the word? We investigate the hypothesis that the speaker is willing to spend more effort to articulate a speech sound clearly when the material is less predictable, but will accept sloppy pronunciation when the material is




Figure 14: *Correlation of strength values and duration. The left panel shows the sentence final syllables, and the right panel shows the non-final syllables.*

predictable. In this section, we use the point-wise mutual information between adjacent syllables to estimate how well a syllable can be predicted from the preceding one, and show that there is a correlation between mutual information scores and prosodic strength.

Point-wise mutual information [10, 5] is a measure of how strongly two events are associated, and is defined as

$$I(a; b) = \log_2(P(a, b)/P(a)P(b)) \quad (11)$$

where $P(a)$ is the probability of the event a , $P(b)$ is the probability of the event b , and $P(a, b)$ is the probability of a and b occurring together.

If a and b are independent events, then the probability of them occurring together is the product of the probabilities of a and b : $P(a, b) = P(a)P(b)$ and the mutual information is zero. Applying this measure to text, we can estimate mutual information of two words by using frequency information obtained from a database.

If two words tend to occur together, their mutual information score is positive. Negative mutual information scores suggest some level of mutually exclusiveness so that the two syllables occur together less often than chance.

In the speech channel, orthographic information is not represented. Therefore, instead of using units like words or Chinese characters [60] that apply to written text, we use the syllable, a sound-based unit, to compute mutual




Figure 15: *Mutual information scores of syllables, based on the preceding syllable. The mutual information is lower for initial syllables (i.e. a prediction across a word boundary), thus they are less predictable from the preceding syllable than word internal syllables are.*

information. Syllables with different tones are considered different events.

We used a database with 15000 sentences (half a million characters). We converted written text into syllable transcriptions using the text analysis component of a text-to-speech system [58]. The system uses a dictionary together with a homograph disambiguation component to allow context sensitive character-to-sound mapping. We then compute the frequency count of each syllable and each syllable pair from the database, and estimate their probability by dividing the frequency with the total syllable count of the database.

Figure 15 compares the mutual information scores of the 737 pairs of adjacent syllables in the speech corpus. The figure compares syllable pairs where the second member is word initial (the syllable pair straddles a word boundary) *vs.* pairs where both syllables are within the same word. The mutual information is high within a word: if you hear the beginning of a word, you have more information about the next syllable. On the other hand, knowing the syllable at the end of one word is not as helpful for predicting the beginning of the next word. We suggest that where the speech sound is less predictable, speakers spend more effort in pronunciation to make the speech clearer. This may be part of the explanation of the higher prosodic strength we obtained in the word initial positions. Figure 15 uses word boundaries from the *J* segmentation, but plots from the other two labelers are nearly identical to the one shown. These results are consistent with those obtained by Pan and Hirschberg [43], using human identification of accent locations.

5.8 Correlation of Strength and Mutual Information

We compute the correlation between mutual information and the prosodic strength of the word initial syllables from three ‘wAT’ models, one from each

word boundary labeler. We do not include word internal syllables in this computation, because the prosodic strength of the word internal syllables is distributed by the metrical structure. The correlation scores of the three models for labelers A, J, and S are -0.20, -0.17, and -0.16, all significant at the 95% confidence level. As expected, there is a negative correlation between mutual information scores and fitted prosodic strengths. Again we see that the less predictable syllable is spoken with higher prosodic strength.

We note that the available database is barely sufficient for calculating mutual information scores across word boundaries: the median syllable occurs only 135 times, thus most possible pairs of syllables simply are not sampled. Consequently, we view these correlations as suggestive, rather than conclusive. However, the observed correlations in Section 5.5.1 are consistent with this hypothesis that strength is at least partially controlled by mutual information. We expect words at the beginning of sentences, clauses, and phrases to be less predictable than words in the middle, because these boundaries can introduce new topics.

As a comparison, we calculated the correlation between mutual information and the high f_0 region in each word. It has been generally expected that a speaker will raise pitch to signal less predictable information. We calculated the f_0 mean of three consecutive voiced samples and took the highest value in each word. The correlation scores of the three segmentations are -0.14, -0.12, and -0.11, smaller than the correlation obtained from fitted prosodic strength, and only the first is significant at the 95% level.

There are several reasons why the fitted strength performs better than surface f_0 values. The raw f_0 values are not corrected for tone class or the effects of the neighboring tones, while the Stem-ML strengths include those basic normalizations. For example, high f_0 may be the result of a preceding rising tone, especially if that tone is emphasized. Not all high f_0 correspond to local intentional emphasis [51, 56]. Furthermore, speakers may use tone-dependent strategies to convey the same prosodic meaning. For example, to express emphasis, people may raise pitch for a high tone but lower pitch for a low tone.

5.9 The Scope of Prosodic Strength

What is the scope of prosodic strength in Mandarin? We cannot directly answer this question because we assume that Equation 8 relates the word strengths to the syllable strengths. All of our models assume that one is exactly proportional to the other, therefore the models do not distinguish between the two.

However, we can compare our results here to previous work by Kochanski and Shih [21] where we built models with a separate strength value for each syllable (thus syllable-scope strengths) to fit the same corpus. Since the RMS errors are only marginally worse when we tie the syllable strengths together to make a word strength (13 Hz in this work, *vs.* 12 Hz in [21]), we can see that associating strength with words works just as well as associating it with syllables, but leads to a much simpler, more compact model with fewer parameters. Occam's razor thus leads us to associate strengths with words.

However, a comparison of RMS errors has its limitations. It averages over the entire data set, and so cannot exclude the possibility that while most words are spoken in the default word-scope manner, the speaker exercises more detailed syllable-scope control over a few words.

6 Conclusion

We have used Stem-ML to build a model of continuous Mandarin speech that connects the acoustic level up to the results of text analysis (part-of-speech information, and word, phrase, clause, and sentence boundaries). When fit to a corpus, the model shows that prosody is used in a consistent way to mark divisions in the text: sentences, clauses, phrases, and words start strong and end weak. Our prosodic measurements also show a useful correlation with word length, and the part-of-speech of words. We also show that the strength values correlate in expected ways with other acoustic observables such as duration. There is also a correlation between the strength values and mutual information, which suggests that speakers apply a higher prosodic strength to less predictable materials.

The results point to the conclusion that the mathematical models behind Stem-ML provide a quantitative method for measuring prosodic strength. The simplicity and compactness with which one can describe Mandarin using this representation suggests that it captures some important aspects of human behavior during speech. For more information, see <http://prosodies.org>.

References

- [1] BELLEGARDA, J., SILVERMAN, K., LENZO, K., AND ANDERSON, V. Statistical prosodic modeling: from corpus design to parameter estimation. *IEEE Transactions on Speech and Audio Processing* 9, 1 (2001), 52–66.
- [2] BROWMAN, C. P., AND GOLDSTEIN, L. Tiers in articulatory phonology, with some implications for casual speech. In *Papers in Laboratory Phonology I: Between the Grammar and Physics of Speech*, J. Kingston and M. Beckman, Eds. Cambridge University Press, 1990, pp. 341–376.
- [3] CHEN, S.-H., HWANG, S. H., AND TSAI, C.-Y. A first study of neural net based generation of prosodic and spectral information for Mandarin text-to-speech. In *Proceedings of IEEE ICASSP (1992)*, vol. 2, pp. 45–48.
- [4] CHEN, Y., GAO, W., ZHU, T., AND MA, J. Multi-strategy data mining on Mandarin prosodic patterns. In *Proceedings of the Sixth International Conference on Spoken Language Processing (ICSLP)* (Beijing, China, October 16-20 2000).

- [5] CHURCH, K. W., AND GALE, W. A comparison of the enhanced Good-Turing and deleted estimation methods for estimating probabilities of English bigrams. *Computer Speech and Language* 5, 1 (1991), 19–54.
- [6] COMPUTATIONAL LINGUISTIC SOCIETY OF THE REPUBLIC OF CHINA. *RO-CLING Chinese Corpus*. Institute of Information Science, Academia Sinica, Nankang, Taipei, Taiwan 11529, R.O.C., 1993. Newspaper texts collected in 1979 and 1980 in Taiwan. Announced Linguist List 4.191.
- [7] CROWNINSHIELD, R. D., AND BRAND, R. A. Physiologically based criterion of muscle force prediction in locomotion. *J. Biomechanics* 14, 11 (1981), 793–801.
- [8] CRYSTAL, T. H., AND HOUSE, A. S. Segmental durations in connected speech signals: syllabic stress. *J. Acoustical Soc. Am.* 83 (1988), 1574–1585.
- [9] EDWARDS, J., BECKMAN, M., AND FLETCHER, J. The articulatory kinematics of final lengthening. *J. Acoustical Soc. Am.* 89 (1991), 369–382.
- [10] FANO, R. *Transmission of Information*. MIT Press, 1961.
- [11] FELDMAN, A. G., ADAMOVICH, S. V., OSTRY, D. J., AND FLANAGAN, J. R. The origin of electromyograms - explanations based on the equilibrium point hypothesis. In Winters and Woo [67], pp. 195–213. and references therein.
- [12] FLASH, T., AND HOGAN, N. The coordination of arm movements: An experimentally confirmed mathematical model. *The Journal of Neuroscience* 5, 7 (1985), 1688–1703.
- [13] FLEMMING, E. Phonetic optimization: Compromise in speech production. In *University of Maryland Working Papers in Linguistics* (1997), vol. 5, pp. 72–91. See <http://www.stanford.edu/flemming/>.
- [14] FLEMMING, E. Scalar and categorical phenomena in a unified model of phonetics and phonology. *Phonology* 18 (2001), 7–44.
- [15] FUJISAKI, H. Dynamic characteristics of voice fundamental frequency in speech and singing. In *The Production of Speech*, P. F. MacNeilage, Ed. Springer-Verlag, 1983, pp. 39–55.
- [16] HIRSCHBERG, J. Pitch accent in context: Predicting intonational prominence from text. *Artificial Intelligence* 63 (1993), 305–340.
- [17] HIRSCHBERG, J., AND PIERREHUMBERT, J. The intonational structuring of discourse. In *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics* (1986), vol. 24, pp. 136–144.
- [18] HOGAN, N., AND WINTERS, J. M. Principles underlying movement organization: Upper limb. In Winters and Woo [67], pp. 182–194. and references therein.

- [19] HOLLIEN, H. In search of vocal frequency control mechanisms. In *Vocal Fold Physiology: Contemporary Research and Clinical Issues*, D. M. Bless and J. H. Abbs, Eds. College-Hill Press, San Diego, CA, 1981, pp. 361–367.
- [20] KLATT, D. H. Interaction between two factors that influence vowel duration. *J. Acoustical Soc. Am.* 54 (1973), 1102–1104.
- [21] KOCHANSKI, G., AND SHIH, C. Automated modelling of Chinese intonation in continuous speech. In *Proceedings of Eurospeech 2001* (October 2001), International Speech Communication Association. Aalborg, Denmark.
- [22] KOCHANSKI, G., AND SHIH, C. Prosody modeling with soft templates. *Speech Communication* 39, 3-4 (February 2003), 311–352.
- [23] KOCHANSKI, G., SHIH, C., AND JING, H. Hierarchical structure and word strength prediction of Mandarin prosody. *International Journal of Speech Technology* 6, 1 (January 2003), 33–43.
- [24] KOCHANSKI, G. P., AND SHIH, C. Stem-ML: Language independent prosody description. In *Proceedings of the International Conference on Spoken Language Processing* (October 2000), vol. 3, pp. 239–242. Beijing, China.
- [25] LABOISSIÈRE, R., OSTRY, D. J., AND FELDMAN, A. G. The control of multi-muscle systems: human jaw and hyoid movements. *Biological Cybernetics* 74 (1996), 373–384.
- [26] LADD, D. R. *Intonational Phonology*. Cambridge University Press, 1996.
- [27] LEA, W. Segmental and suprasegmental influences on fundamental frequency contours. In *Consonant Types and Tones*, L. Hyman, Ed. University of Southern California, Los Angeles, 1973, pp. 15–70.
- [28] LEE, L.-S., TSENG, C.-Y., AND HSIEH, C.-J. Improved tone concatenation rules in a formant-based Chinese text-to-speech system. *IEEE Transactions on Speech and Audio Processing* 1, 3 (1993), 287–294.
- [29] LEVENBERG, K. A method for the solution of certain problems in least squares. *Quart. Applied Math.* 2 (1944), 164–168.
- [30] LIBERMAN, M. Y., AND PIERREHUMBERT, J. B. Intonational invariance under changes in pitch range and length. In *Language Sound Structure*, M. Aronoff and R. Oehrle, Eds. M.I.T. Press, Cambridge, Massachusetts, 1984, pp. 157–233.
- [31] LIBERMAN, M. Y., AND PRINCE, A. On stress and linguistic rhythm. *Linguistic Inquiry* 8 (1977), 249–336.

- [32] LIN, M.-C., AND YAN, J. The stress pattern and its acoustic correlates in Beijing Mandarin. In *Proceedings of the 10th International Congress of Phonetic Sciences* (1983), pp. 504–514.
- [33] LINDBLOM, B. Spectrographic study of vowel reduction. *J. Acoustical Soc. Am.* 35, 11 (November 1963), 1773–1781.
- [34] MACNEILAGE, P. F., SUSSMAN, H. M., WESTBURY, J. R., AND POWERS, R. K. Mechanical properties of single motor units in speech musculature. *J. Acoustical Soc. Am.* 65, 4 (April 1979), 1047–1052.
- [35] MARQUARDT, D. An algorithm for least-squares estimation of nonlinear parameters. *SIAM J. Applied Math* 11 (1963), 431–441.
- [36] MATHSOFT, INC.I. *Plus Online Documentation*, 3.3 ed., 1995. Subroutine *ltsreg()*, set to exclude the 5 most extreme data points from the objective function.
- [37] MCFARLAND, D. H., AND SMITH, A. Effects of vocal task and respiratory phase on prephonatory chest-wall movements. *J. Speech and Hearing Research* 35, 5 (1992), 971–982.
- [38] MONSEN, R. B., ENGBRETSON, A. M., AND VEMULA, N. R. Indirect assesment of the contribution of subglottal air pressure and vocal fold tension to changes in the fundamental frequency in English. *J. Acoust. Soc. Am.* 64, 1 (1978), 65–80.
- [39] MOON, S.-J., AND LINDBLOM, B. Interaction between duration, context, and speaking style in English stressed vowels. *J. Acoustical Soc. Am.* 96, 1 (July 1994), 40–55.
- [40] NELSON, W. L. Physical principles for economies of skilled movements. *Biological Cybernetics* 46 (1983), 135–147.
- [41] OHALA, J. J. The segment: Primitive or derived? In *Papers in Laboratory Phonology II: Gesture, Segment, Prosody*, G. J. Docherty and D. R. Ladd, Eds. Cambridge University Press, 1992, pp. 166–183. ISBN 0-521-40127-5.
- [42] ÖHMAN, S. Word and sentence intonation, a quantitative model. Tech. rep., Department of Speech Communication, Royal Institute of Technology (KTH), 1967.
- [43] PAN, S., AND HIRSCHBERG, J. Modeling local context for pitch accent prediction. In *Proc. 38th Ann. Mtg. Association for Computational Linguistics (ACL2000)*, L. Hyman, Ed. Association for Computational Linguistics, Hong Kong, October 2000, pp. 15–70.
- [44] PERKELL, J. S., AND ZANDIPOUR, M. Economy of effort in different speaking conditions. II. Kinematic performance spaces for cyclical and speech movements. *J. Acoustical Soc. Am.* 112, 4 (2002), 1642–1651.

- [45] PERKELL, J. S., ZANDIPOUR, M., MATTHIES, M. L., AND LANE, H. Economy of effort in different speaking conditions. I. a preliminary study of intersubject differences and modeling issues. *J. Acoustical Soc. Am.* 112, 4 (2002), 1627–1641.
- [46] PIERREHUMBERT, J. B., AND BECKMAN, M. E. *Japanese Tone Structure*. The MIT Press, 1988.
- [47] PRINCE, A., AND SMOLENSKY, P. *Optimality Theory: Constraint Interaction in Generative Grammar*. MIT Press, to be published. Also available as Technical Report 2 from the Center for Cognitive Science (RuCCS), Rutgers University, Busch Campus, New Brunswick, NJ 08903.
- [48] SEIF-NARAGHI, A. H., AND WINTERS, J. M. Optimized strategies for scaling goal-directed dynamic limb movements. In Winters and Woo [67], pp. 312–334. and references therein.
- [49] SELKIRK, E. O. *Phonology and Syntax: The Relation between Sound and Structure*. The MIT Press, Cambridge, MA, 1984.
- [50] SHIH, C. *The Prosodic Domain of Tone Sandhi in Chinese*. PhD thesis, University of California, San Diego, 1986.
- [51] SHIH, C. Tone and intonation in Mandarin. In *Working Papers of the Cornell Phonetics Laboratory, Number 3: Stress, Tone and Intonation*. Cornell University, 1988, pp. 83–109.
- [52] SHIH, C. A declination model of Mandarin Chinese. In *Intonation: Analysis, Modelling and Technology*, A. Botinis, Ed. Kluwer Academic Publishers, 2000, pp. 243–268.
- [53] SHIH, C., AND AO, B. Duration study for the Bell Laboratories Mandarin text-to-speech system. In *Progress in Speech Synthesis*, J. van Santen, R. Sproat, J. Olive, and J. Hirschberg, Eds. Springer-Verlag, New York, 1997, pp. 383–399.
- [54] SHIH, C., AND KOCHANSKI, G. P. Chinese tone modeling with StemML. In *Proceedings of the International Conference on Spoken Language Processing* (2000), vol. 2, pp. 67–70. Beijing, China.
- [55] SHIH, C., AND KOCHANSKI, G. P. Prosody control for speaking and singing styles. In *Proceedings of Eurospeech 2001* (October 2001), International Speech Communication Association, pp. 669–672. Aalborg, Denmark.
- [56] SHIH, C., KOCHANSKI, G. P., FOSLER-LUSSIER, E., CHAN, M., AND YUAN, J.-H. Implications of prosody modeling for prosody recognition. In *Proceedings of the ISCA Tutorial and Research Workshop on Prosody in Speech Recognition and Understanding* (October 2001), M. Bacchiani, J. Hirschberg, D. Litman, and M. Ostendorf, Eds., International Speech Communication Association, pp. 133–138. Red Bank, NJ.

- [57] SHIH, C., AND SPROAT, R. W. Variations of the Mandarin rising tone. In *Proceedings of the IRCS Workshop on Prosody in Natural Speech* (1992), University of Pennsylvania, pp. 193–200.
- [58] SHIH, C., AND SPROAT, R. W. Issues in text-to-speech conversion for Mandarin. *Computational Linguistics and Chinese Language Processing* 1, 1 (1996), 37–86.
- [59] SILVERMAN, K. E. *The Structure and Processing of Fundamental Frequency Contours*. PhD thesis, University of Cambridge, UK, 1987.
- [60] SPROAT, R. W., AND SHIH, C. A statistical method for finding word boundaries in Chinese text. *Computer Processing of Chinese and Oriental Languages* 4, 4 (1990), 336–351.
- [61] STEVENS, K. N. *Acoustic Phonetics*. The MIT Press, 1998.
- [62] TALKIN, D., AND LIN, D. *ESPS/waves online documentation*, 5.31 ed., 1996. Program get_f0. ESPS was purchased by Microsoft in 2000. Algorithm is based on: Talkin, D. (1995). A Robust Algorithm for Pitch Tracking (RAPT). In Kleijn, W. B. and Paliwal, K. K. (Eds.), *Speech Coding and Synthesis*. New York: Elsevier.
- [63] VAN SANTEN, J. P. H., AND BUCHSBAUM, A. L. Methods for optimal text selection. In *EuroSpeech 97* (Rhodes, Greece, 1997), vol. 2, pp. 553–556.
- [64] WHALEN, D., AND KINSELLA-SHAW, J. M. Exploring the relationship of inspiration duration to utterance duration. *Phonetica* 54 (1997), 138–152.
- [65] WILDER, C. N. Chest wall preparation for phonation in female speakers. In *Vocal Fold Physiology: Contemporary Research and Clinical Issues*, D. M. Bless and J. H. Abbs, Eds. College-Hill Press, San Diego, CA, 1981, pp. 109–123. ISBN 0-933014-87-2.
- [66] WINKWORTH, A. L., DAVIS, P. J., ADAMS, R. D., AND ELLIS, E. Breathing patterns during spontaneous speech. *Journal of Speech and Hearing Research* 38, 1 (1995), 124–144.
- [67] WINTERS, J., AND WOO, S., Eds. *Multiple Muscle Systems: Biomechanics and Movement Organization*. Springer-Verlag, New York, 1990.
- [68] WINTERS, J. M. Hill-based muscle models: A systems engineering perspective. In Winters and Woo [67], pp. 69–93. and references therein.
- [69] XU, Y. Pitch targets and their realization: Evidence from Mandarin Chinese. *Speech Communication* 33 (2001), 319–337.
- [70] XU, Y., AND SUN, X. J. How fast can we really change pitch? maximum speed of pitch change revisited. In *Proceedings of the Sixth International Conference on Spoken Language Processing (ICSLP)* (Beijing, China, October 16-20 2000).

- [71] ZAHALAK, G. I. Modeling muscle mechanics (and energetics). In Winters and Woo [67], pp. 1–23. and references therein.
- [72] ZAJAC, F. E. Muscle and tendon: Properties, models, scaling, and application to biomechanics and motor control. *Critical Reviews in Biomedical Engineering* 17, 4 (1989), 359–411.
- [73] ZAJAC, F. E., AND WINTERS, J. M. Modeling musculoskeletal movement systems: Joint and body segmental dynamics, musculoskeletal actuation, and neuromuscular control. In Winters and Woo [67], pp. 139–146. and references therein.