Decomposition of Pitch Curves in the General Superpositional Intonation Model

Taniya Mishra, Jan van Santen and Esther Klabbers

Center for Spoken Language Understanding
OGI School of Science & Engineering
Oregon Health & Science University
20000 NW Walker Road, Beaverton, OR 97006, USA
{mishra, vansanten, klabbers}@cslu.ogi.edu

Abstract

This paper describes and applies a new algorithm for decomposing pitch curves into component curves, in accordance with the General Superpositional Model of Intonation. According to this model, which is a generalization of the Fujisaki model [3], a pitch contour can be described as the sum of component curves that are each associated with different phonological levels, including the phrase, foot, and phoneme. The algorithm assumes that the phrase curve is locally linear during intervals spanned by a foot. The algorithm was evaluated using synthetically generated curves, and was found to accurately recover the synthetic component curves. The algorithm was also evaluated in a perceptual experiment, where speech generated by concatenation of accent curves was shown to produce better speech quality than speech based on direct concatenation of "raw" pitch curve fragments.

1. Introduction

Recently, we proposed an approach to speech synthesis in which a multi-level search process is used to find sequences of acoustic units covering the target phoneme sequences, (left-headed) foot sequences, and phrase sequences [13]. This approach, which we shall refer to as *Multi-Level Unit Sequence Synthesis*, is based on the General Superpositional Model of Intonation, according to which a pitch contour can be described as the sum of component curves that are each associated with different phonological levels, specifically the phoneme, foot, and phrase [10, 12]. During synthesis, segmental perturbation curves, accent curves and phrase curves are extracted from the acoustic unit, foot, and phrase sequences, respectively, and are combined into *target pitch curves*; these target curves are then imposed on the acoustic unit sequences using standard pitch modification methods. (Segmental perturbation curves represent short-lived spikes around nasal-vowel boundaries, sharp elevations in frequency during initial portions of vowels following obstruents, and effects of intrinsic pitch)

This approach represents an attempt to combine the strengths of the two different approaches to speech synthesis that are currently dominant: *Unit selection synthesis*, which preserves all details of natural speech but struggles with coverage of the very large combinatorial space of phoneme sequences and prosodic contexts; and *diphone synthesis*, which addresses coverage by generating rule-based synthetic target prosody and then imposing it on acoustic units using signal modification methods. In the Multi-Level Unit Sequence Synthesis approach, we use quasi-natural target contours but impose these on acoustic unit sequences, thereby making it unnecessary to cover all *combinations* of phoneme sequences and prosodic contexts. We avoid the discontinuities that would occur if we created target pitch contours by raw concatenation of pitch contour fragments by instead concatenating extracted accent curves and adding these to smooth, continuous phrase curves.

One of the unsolved problems in this approach is the automatic decomposition of pitch curves into component curves. This problem is far from trivial because (i) successive accent curves may overlap in time and (ii) we want to impose few if any constraints on the shapes of accent curves and phrase curves. Two other superpositional approaches have been used for F_0 decomposition, one based on the Fujisaki model [3] and the other based on the SFC model [21]. Each approach differs from ours. The Fujisaki based approach imposes strong constraints on the curve shapes, while we impose few constraints on the shapes of the accent curves. The SFC based approach estimates the component curves using neural-nets based global optimization, while we perform optimization on a per utterance basis. The presented paper proposes and evaluates our algorithm for decomposing F_0 curves.

2. Pitch Decomposition Algorithm

2.1. Basic Assumptions

The pitch decomposition algorithm is based on the following four assumptions:

1. **Additive Decomposition of the Pitch Curve**: The General Superpositional Model of Intonation [10, 12] states that,

$$F_0(t) = \bigoplus_{c \in C} \bigoplus_{k \in c} f_{c,k}(t) \tag{1}$$

where C is a set of curve classes (e.g., {phrase, accent, perturbation}) and k is a particular curve from a particular class. The operator $\bigoplus$ represents *generalized addition* – an operator that includes addition and multiplication as special cases. When we assume that $C = \{P, A_i\}$, where P = phrase curve, A_i = accent curve, and operator $\bigoplus$ represents addition, then:

$$F_0(t) = P(t) + \Sigma_i^n A_i(t) \tag{2}$$

For our present study, we are not considering perturbations to be a separate curve class - high frequency perturbations are

This research was conducted with support from NSF grant 0205731, "Prosody Generation in Child-Oriented Speech".

smoothed out, and the remaining are grouped together with accent curves.

2. **Different Temporal Scopes of the Component Curves**: As in the General Superpositional Model, it is assumed that each of the component curves is tied to a distinct phonological entity and follows a distinct time course. We assume here that the phrase curve is tied to the phonological phrase and spans the entire phrase length, whereas each accent curve is tied to a distinct *left-headed foot* and is left-aligned with the start of the foot. The *left-headed foot* is defined as an accented syllable followed by one or more unaccented syllables [5, 6]. The sequence of syllables preceding the first accented syllable in a phrase is called the *anacrusis*.

3. **Piece-wise Linear Phrase Curve**: It is assumed that a phrase curve can be approximated by n line segments, p_i, where n is the number of feet in the phrase (Eq. 3). Each p_i spans the length of a foot. The points at which the phrase curve change direction are called *inflection points*.

$$p_i(t) = \beta_i(t) + \gamma_i, P = concatenate(p_i) \tag{3}$$

4. **Accent Curves from Common Templates**: Accent curve shapes are unspecified in the General Superpositional Model, except for the abstract property that accent curves for a given class (e.g., single-peaked, yes/no rise) are all generated from a class-specific template, E_i, using parameterized time warps [11]. Thus, for every accent curve A_i, there exists a template curve E_i (shared with similarly-shaped accent curves), such that A_i equals a scaled (by a height factor, h_i) and time warped replica of E_i (Eq. 4).

$$A_i(t) = h_i \times time\_warped(E_i(t)) \tag{4}$$

Based on previous research [11], we assume that a single-peaked accent curve, which ranges from 0 to a peak value and again descends to 0, starts at the beginning of the associated left-headed foot and finishes at *or beyond* the end of the foot, but, conservatively, no later than the point where the next accent curve reaches its peak.

2.2. Algorithm

Given these assumptions, the pitch decomposition algorithm involves the following steps, formulated here for the special case of single-peaked accent curves in declarative phrases:

1. **Peak Detection**: Given the raw pitch curve associated with a given phrase, the number of feet in the phrase, and the locations of all left-headed feet, one peak per foot is detected using a robust peak detection algorithm based on isotonic regression [23].

2. **Parametrization**: A template is characterized by 9 pitch values that approximate a Gaussian sampled at equal time intervals, with the first and last values rounded down to 0. A given accent curve is characterized by a time warp, consisting of 9 time points associated with these 9 points in the template, and a height parameter, h_i. Because it is assumed that accent curves start at the beginning of the foot, the first time warp parameter is set equal to the foot start time. Thus, there are 9 parameters (8 remaining time warp parameters and one height parameter) to be estimated for each of the n accent curves.

A phrase curve is characterized by n interconnected line segments. Each segment begins at the starting point of a distinct foot and spans the length of the foot. Since we are only considering all-sonorant declarative phrases for the present study, we can assume that the phrase curve values at the start and end of the phrase equal the values of the raw pitch curve at those time points. Because of continuity of the phrase curve, and because the times of the inflection points are known, $n-1$ parameters are needed to describe the interconnected line segments making up the phrase curve. For the present study, the warping function has no constraint other than monotonicity. However, when we consider more complex types of utterances, more constraints may need to be applied to warping function.

3. **Optimization**: Parameters are estimated using the Nelder-Mead simplex method [7] as implemented in the routine fminsearch in MATLAB [8], using as error criterion the weighted root mean squared deviation between predicted and observed pitch contours; weights are given by the product of the get_f0 [17] voicing flag and energy.

3. Experiment with Synthetic Curves

To test the pitch decomposition algorithm, we synthesized a set of 75 declarative F_0 contours with single-peaked pitch accents, and used the algorithm to extract phrase and the accent curves.

3.1. Materials

The pitch contours were generated using a simplified version of the Bell Labs' Linear Alignment model [10], called SLAM (or the Simplified Linear Alignment Model). In this model, the pitch curve is a summation of a phrase curve and n accent curves. The phrase curve is created by interpolation between three points: start of the phrase, start of the last foot, end of the phrase. The accent curves are created by cosine interpolation between the start, peak location, and end. Peak location is a function of foot duration and the number of syllables in the foot.

We created 32 contours with two accent curves and 43 with three accent curves. For every contour, the accent curves were asymmetric in shape, and non-phrase-final curves overlapped with the next curve. The following curves were generated. (i) *Two-accent case*. Accent curve heights were set at 50 Hz and 75 Hz, slopes for the pre-nuclear segment of the phrase curves were set at 50 Hz/s and 70 Hz/s, the initial boundary tone was always 110 Hz, and successive accent curve overlap was set at 10% and 20%. An example is shown in Figure 1. (ii) *Three-accent case*. Similar to the two-accent case, except that accent curve heights were set at 50 Hz, 75 Hz and 95 Hz and that only a random subset of all combinations (43 out of 432) was used. An example is shown in Figure 2.

Overlapping accent curves illustrate a key strength of the proposed pitch decomposition algorithm over other decomposition approaches. In a *filtering method* [14, 9], the assumption must be made that the phrase curve is completely smooth with no inflection points. A weakness of the *wavelet-based approach* [12] is that it only performs a partial decomposition: it returns the phrase curve and a summation of the accent and segmental perturbation curves.

3.2. Results

The accuracy with which the algorithm estimates component curves was measured by the Root Mean Squared Error (RMSE) between the estimated and the original component curves for the 75 instances. The results (Table 1) show an encouraging goodness-of-fit.

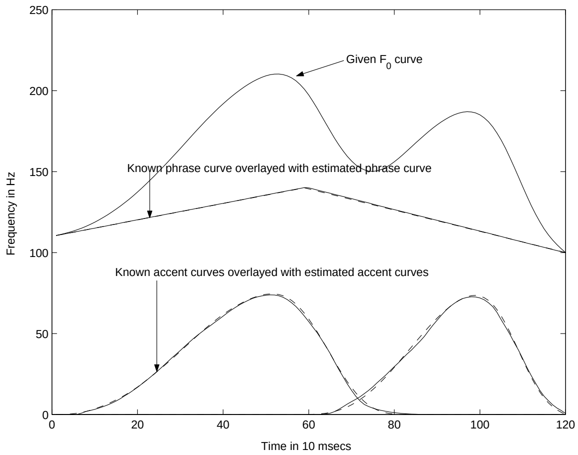

Figure 1: Example of a two-accent synthetic F_0 curve used to test the decomposition algorithm.

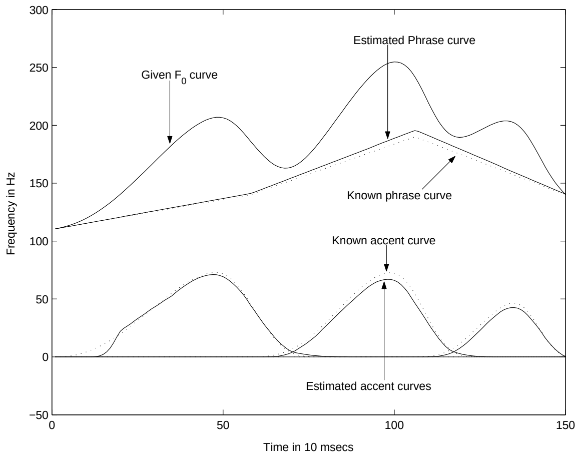

Figure 2: Example of a three-accent synthetic F_0 curve used to test the decomposition algorithm.

4. Perceptual Study

A perceptual study was conducted to test the claim of the Multi-Level Unit Sequence Synthesis method that it should produce better quality intonation than that produced by synthetic target pitch contours (e.g., diphone synthesis) or synthesis based on raw concatenation of pitch contour fragments (e.g., [20]).

4.1. Textual Materials

The recordings used for the pilot study were elicited from one male speaker. Each utterance set had the following format, where x denotes utterance set number. $x = 1, 2, ..., 13$:

1. a_x: **Leena**$_a$ Roy$_a$.
2. b_x: **Leena**$_b$ Weller$_b$ and **Ann**$_b$ Roy$_b$.
3. c_x: Leena$_c$ **Weller**$_c$ and Ann$_c$ **Roy**$_c$.

The speaker was instructed to put a relatively higher emphasis on the highlighted word, and to pronounce the b_x and c_x items in a 'list-like' manner, in order to elicit variability in pitch range more typical of natural speech used in unit selection synthesis than of the stilted speech in diphone synthesis. A total of 39 utterances were recorded.

	2-accent contours	3-accent contours
Phrase	4.16	9.11
Accent-1	3.99	4.58
Accent-2	2.21	6.76
Accent-3	N/A	3.20

Table 1: Performance of the decomposition algorithm (quantities in Hz). Accent-i is the i-th accent curve in the phrase.

4.2. Stimulus Generation

Every utterance was manually segmented into left-headed feet using Wavesurfer [15]. Pitch curves were extracted (at 10 ms interval) using ESPS get_f0 utility [17]; based on [19], we removed high-frequency noise using the Savitzky-Golay [16] smoothing filter of order 3 and length 5.

As an accent curve template, we use a Gaussian curve. We note that the exact shape of the template is not critical as long as it is single-peaked and has initial and final values of 0, because the unconstrained time warping procedure renders different templates with these features equivalent. Given these inputs, the algorithm decomposes every natural pitch curve into the estimated phrase curve and the estimated accent curves.An example of the natural pitch curve decomposition is given in Figure 3. The average warping function obtained by the decomposition of the utterance set b_x is shown in Figure 4. Note that for the first three accent curves, the average warping function goes beyond 100% of the associated foot duration, indicating that each of the first three accent curves overlaps with the following accent curve. In each of the 13 utterance sets (shown in Section 4.1),

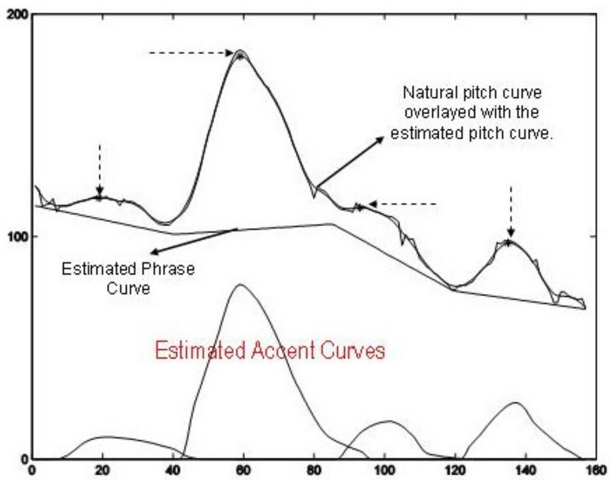

Figure 3: Decomposition of a natural pitch curve: The dashed arrows indicate the peaks detected by the algorithm. The x-axis is time in 10 ms intervals, and the y-axis is in frequency in Hz.

utterance a_x is considered the target utterance. So, for every set, utterance a_x is resynthesized five times using STRAIGHT [4]; each time with a different pitch contour. Duration mapping between the generated F_0 and the original F_0 was done at the foot level. The five different pitch contours are as follows:
1. Original F_0 (ORIG): The first target pitch curve is simply the pitch curve extracted using ESPS.
2. Raw F_0 concatenated (CONCAT): For each utterance set, raw F_0 will be extracted from the first and the last units of utterance b_x (because $Leena_a$ and $Leena_b$ are in the same prosodic

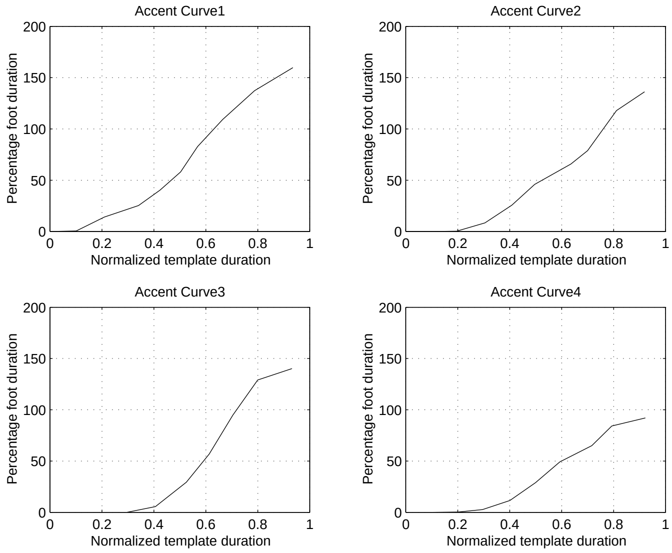

Figure 4: Average warping function related to the decomposition of utterance set b_x

context, as are Roy_a and Roy_b) and concatenated to generate the second target pitch contour. We tried to match the peaks of the concatenated F_0 to the peaks of the natural F_0 of the target utterance by multiplying each of the units by a height factor. However, doing so sometimes resulted in more drastic jumps at the unit boundaries, as shown in Figure 5; we therefore decided not to perform the peak-matching.

3. SLAM F_0 (SYNTHETIC): We use the SLAM model to gen-

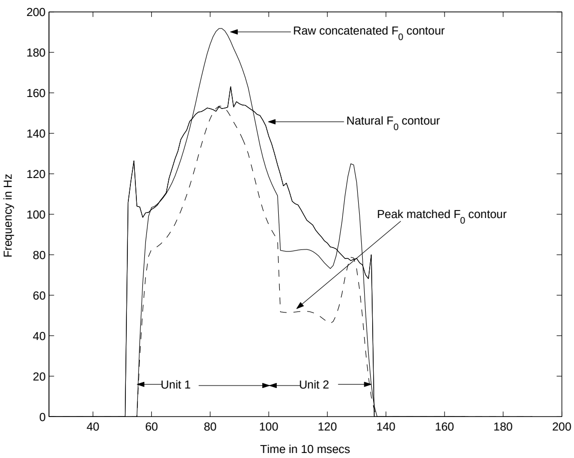

Figure 5: Peak matching concatenated pitch curves can make the pitch mismatches at unit boundaries more drastic

erate the third target pitch contour by means of the OGI version of the Festival Speech Synthesis system [2]. The phrase curve height parameters, as well as the accent height parameters, are set to mimic the actual speaker's average phrase start, phrase end, and peak heights closely.

4. 'Semi-natural' F_0 with accent curves from the right prosodic context (DECOM$_1$): For every utterance set, the accent curves associated with the first and the last units of b_x are obtained. The obtained accent curves are scaled by a suitable height factor, and added to a synthetic phrase curve to generate the fourth target F_0 curve. The phrase curve was generated to approximately mimic the average phrase curve obtained by decomposition of a_x.

5. 'Semi-natural' F_0 with accent curves from the wrong prosodic context (DECOM$_2$): In the previous method of F_0 generation, we took care to find accent curves from the same prosodic context as the target accent curves. However, we may not be able to find accent curves in the corpus that match the needed target accent curves completely in terms of prosodic context. To examine the effects of using such accent curves, for every utterance set, we use the accent curves associated with the first and last units of c_x to produce the target contour. Note that the relative emphasis on $Leena_c$ and Roy_c are the reverse of the relative emphasis on $Leena_a$ and Roy_a, respectively. However, the position in phrase and the number of syllables for the matched pairs, ($Leena_a$, $Leena_c$) and (Roy_a, Roy_c), are identical. We conjecture that given the matching position in phrase and the identical number of syllables, the effect of emphasis translates to only a difference in the height factor of the accent curve. So, the accent curves were appropriately scaled and added to a synthetic phrase curve to generate the fifth target contour.

4.3. Listening Protocol

We created 13 utterance sets comprising 5 target pitch contours applied to the same utterance. The ordering of the 13 sets was randomized, and within each set, the order of presentation of the 5 target utterances was randomized as well. The listening test was presented to six listeners using a CGI-based script (WWStim [22]). The five target utterances of every randomized set were presented on the same page; each set on a different page. The test was performed on one computer with a M-Audio Duo USB audio interface and a high quality AKG headset. The listeners were asked to listen to each of the 65 target utterances, one at a time, and rate the naturalness of the pitch on a five-point scale. Of the six listeners, four were students from our university and two were staff members. All listeners are fluent in spoken English, and four are involved in speech research.

4.4. Results

The data form a 65 stimuli $\times$ 6 listeners score matrix S. We first performed planned t-tests for three key predictions, according to which we expect DECOM$_1$ to be better than CONCAT, SYNTHETIC, and DECOM$_2$. All t-tests were significant at 0.025, with t values of 4.62, 4.70, and 3.58. Since all six listeners having results in the predicted direction, also sign tests were significant (p= 0.016).

Second, to obtain a picture of the combined scores corrected for some subjects using a different range of ratings or not being in line with the majority of subjects, we performed a principal components analysis (PCA) on S after its columns were transformed into z-scores. This analysis produces a weighted combination of the ratings, assigning larger weights to mainstream listeners, and eliminating any differences in individual usage of the rating scales due to the z-transformation. The resulting averages are shown in Figure 7.

From Figure 7, we conclude that method $DECOM_1$ predicts more natural sounding F_0 compared to methods SYNTHETIC, CONCAT and $DECOM_2$. We hypothesize that the superior performance of $DECOM_1$ over the other methods can be attributed to the following key differences:(i) The key difference between $DECOM_1$ and SYNTHETIC is that $DECOM_1$ generates pitch curves containing details of natural pitch, whereas SYNTHETIC does not. An example illustrating this difference

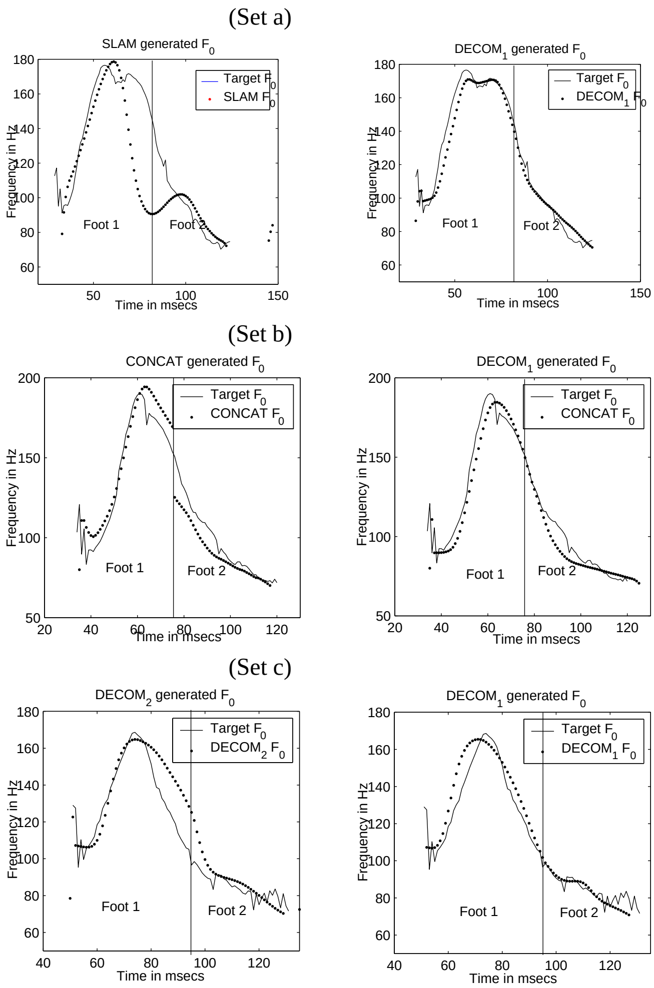

Figure 6: Main differences in F_0 prediction between $DECOM_1$, and each of SYNTHETIC, CONCAT and $DECOM_2$

is shown in Figure 6 (set a). This difference arises from the fact that $DECOM_1$ uses component accent curves extracted from natural pitch curves to generate target F_0 curves, whereas SYNTHETIC generates target F_0 curves using statistically-based rules. (ii) The main difference between $DECOM_1$ and CONCAT is that $DECOM_1$ *always* generates continuous F_0 curves whereas CONCAT sometimes generates F_0 curves that have audible pitch discontinuities, as shown in Figure 6 (set b). In the illustrated example, the discontinuity between the two parts of the pitch predicted by CONCAT occurred in spite of each of the pieces being chosen from appropriate prosodic contexts. Though chosen from appropriate contexts, the two parts were chosen from a list-type utterance. List-type utterances are a type of expressive speech, and hence have extreme variations in pitch. Thus, even two pitch pieces chosen from the same list utterance and concatenated together result in a large pitch jump at the boundary between the units. One might argue that this discontinuity can be smoothed, however most smoothing techniques will cause the natural details inherent in each of the pitch pieces to be lost in the smoothed pitch curve. Since our goal is to preserve the details of natural pitch, we have refrained from performing such smoothing. Another intuitive argument may be that the pitch pieces should be extracted from less expressive speech. Doing so will perhaps result in no (or small) pitch discontinuities, however, the resultant pitch curve will also be less expressive. (iii) Finally, the main difference between $DECOM_1$ and

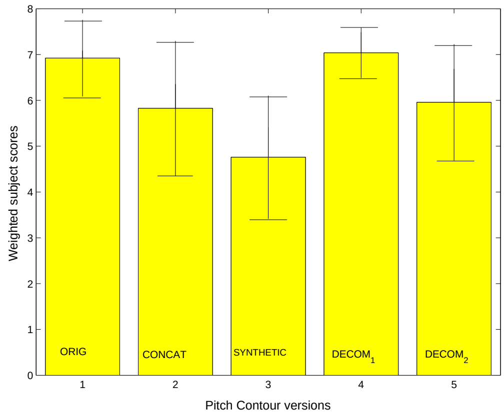

Figure 7: Means and standard deviations of different methods of F_0 generation: The bar graph shows the means and the line above each bar indicates the standard deviation.

$DECOM_2$ is that since in method $DECOM_2$, the component accent curves are extracted from sub-optimal prosodic contexts, the resultant pitch curve might not match the target pitch curve as well as the pitch curve generated by method $DECOM_1$, which is composed of accent curves extracted from the optimal prosodic context. This key difference is illustrated in Figure 6 (set c).

5. Conclusion and Future Work

The core goal of this paper was to describe an algorithm for decomposition of pitch contours into accent curves and phrase curves while making minimal assumptions about the shapes of the underlying curves. Two assumptions were made. One is that the accent curves are single-peaked, and that their start and end values are zero. This assumption in effect says that the model can only be applied to declaratives with high pitch accents. The second assumption is that the phrase curve is piece-wise linear, with the inflection points occurring at foot boundaries. Clearly, this assumption is at most approximately valid.

The evaluation on synthetic data showed that the algorithm can recover the underlying component curves with good accuracy. Evidence for robustness is provided by the fact that the synthetic accent curves (asymmetric curves cobbled together via cosine interpolation) were different in shape from the Gaussian templates.

The perceptual evaluation showed that the Multi-Level Unit Sequence Synthesis method is a promising alternative to synthesis approaches that use synthetic pitch target curves, and to synthesis approaches that generate target pitch curves by concatenation of raw pitch curves. Comparing the superpositional method with raw pitch concatenation, we can conclude that the continuity of the target pitch contour provided by the new method clearly outweighed any pitch modification distortion.

Future work on the decomposition algorithm includes, first of all, the further decomposition of the accent curves as estimated by the proposed method into "true" accent curves and segmental perturbation curves – these two are combined in the proposed estimation procedure. Second, the algorithm has to be extended to process speech with non-sonorants (in the current paper, all-sonorant text was used). Third, many other pitch accent shapes need to be explored. Fourth, better approximation of the phrase curve shapes.

6. References

[1] A.W. Black and P. Taylor. "Automatically clustering similar units for unit selection in speech synthesis." In *Proceedings of Eurospeech 1997*, Rhodes, Greece, 601-604, 1997.

[2] A.W. Black and P. Taylor. "Festival speech synthesis system: System documentation (1.1.1)." In *Human Communication Research Centre Technical Report HCRC/TR-83*, Edinburgh, U.K., 1997.

[3] H. Fujisaki. "Dynamic characteristics of voice fundamental frequency in speech and singing." In Peter MacNeilage, editor, *The Production of Speech*, p39-55, Springer, New York, NY, 1983.

[4] H. Kawahara, J. Estill and O. Fujimura. "Aperiodicity extraction and control using mixed mode excitation and group delay manipulation for a high quality speech analysis, modification and synthesis system STRAIGHT." In *Proceedings of MAVEBA 2001*, Firenze, Italy, 2001.

[5] E. Klabbers, J. van Santen and J. Wouters. "Prosodic factors for predicting local pitch shape." In *Proceedings of the 2002 IEEE Workshop on Speech Synthesis*, Santa Monica, CA, 2002.

[6] E. Klabbers and J. van Santen. "Control and prediction of the impact of pitch modification on synthetic speech quality." In *Proceedings of EUROSPEECH'03*, Geneva, Switzerland, p317-320, 2003.

[7] J.C. Lagarias, J.A. Reeds, M.H. Wright and P.E. Wright. "Convergence properties of the Nelder-Mead simplex method in low dimensions." In *SIAM Journal on Optimization*, 9(1):112-147, 1998.

[8] MathWorks, Inc. FMINSEARCH. MATLAB function. From the MATLAB command 'type fminsearch' (2002).

[9] H. Mixdorff. "A novel approach to the fully automatic extraction of Fujisaki model parameters." In *Proceedings of ICASSP'02*, 3:1281-1284, 2002.

[10] J. van Santen and B. Möbius. "A quantitative model of F_0 generation and alignment." In A. Botinis, editor, *Intonation: Analysis, Modeling and Technology*, p269-288, Kluwer Academic Publishers, Netherlands, 1999.

[11] J. van Santen. "Quantitative modelling of pitch accent alignment." In *Proceedings of Speech Prosody'02*, Aix-en-Provence, France, 2002.

[12] J. van Santen, T. Mishra, and E. Klabbers. "Estimating phrase curves in the general superpositional intonation model." In *Proceedings of the ISCA Speech Synthesis Workshop'04*, Pittsburgh, PA, 2004.

[13] J. van Santen, A. Kain, E. Klabbers, and T. Mishra. "Synthesis of prosody using multi-level unit sequences." In *Speech Communication*, Volume 46, 365-375, 2005.

[14] A. Sakurai and K. Hirose. "Detection of phrase boundaries in Japanese by low-pass filtering of fundamental frequency contours." In *Proceedings of ICSLP'00*, Philadephia, PA, p817-820, 1996.

[15] K. Sjölander and J. Beskow. "Wavesurfer - An open source speech tool." In *Proceedings of ICSLP'00*, Beijing, China, 2000.

[16] A. Savitzky and M.J.E. Golay. "Smoothing and differentiation of data by simplified least squares procedures." In *Analytical Chemistry*, 36:1627-1639, 1964.

[17] D. Talkin. *ESPS.* Entropic Research Lab Inc., 1993.

[18] N. Thorsen. "Intonation contours and stress group patterns in declarative sentences of varying length in ASC Danish." In *ARIPUC 14*, p1-24, 1980.

[19] S.J. Orfanidis. *Introduction to Signal Processing.* Prentis-Hall, Englewood Cliffs, NJ, 1996.

[20] A. Raux and A.W. Black. "A unit selection approach to $F0$ modeling and its application to emphasis." In *Proceedings of ASRU 2003*, St Thomas, US Virgin Islands, 2003.

[21] S. Raidt, G. Bailly, B. Holm, and H. Mixdorff, "Automatic generation of prosody: comparing two superpositional systems." In *Proceedings of Speech Prosody 2004*, Nara, Japan, 2004.

[22] T. Veenker. *WWStim: A CGI script for presenting web-based questionnaires and experiments.* Retrieved September 19, 2005, from Utrecht University, Department of Linguistics: http://www.let.uu.nl/~Theo.Veenker/personal/projects/wwstim/doc/en/

[23] E. Wegman. "Two approaches to nonparametric regression: splines and isotonic inference." In K. Matusita, editor, *Recent Developments in Statistical Inference and Data Analysis*, p323-334, North-Holland Publishing Company, Amsterdam, The Netherlands, 1980.