# Imperfect pitch: Gabor's uncertainty principle and the pitch of extremely brief sounds 

I-Hui Hsieh ${ }^{1}$ • Kourosh Saberi ${ }^{2}$

Published online: 29 May 2015
(C) Psychonomic Society, Inc. 2015


#### Abstract

How brief must a sound be before its pitch is no longer perceived? The uncertainty tradeoff between temporal and spectral resolution (Gabor's principle) limits the minimum duration required for accurate pitch identification or discrimination. Prior studies have reported that pitch can be extracted from sinusoidal pulses as brief as half a cycle. This finding has been used in a number of classic papers to develop models of pitch encoding. We have found that phase randomization, which eliminates timbre confounds, degrades this ability to chance, raising serious concerns over the foundation on which classic pitch models have been built. The current study investigated whether subthreshold pitch cues may still exist in partial-cycle pulses revealed through statistical integration in a time series containing multiple pulses. To this end, we measured frequencydiscrimination thresholds in a two-interval forced-choice task for trains of partial-cycle random-phase tone pulses. We found that residual pitch cues exist in these pulses but discriminating them requires an order of magnitude (ten times) larger frequency difference than that reported previously, necessitating a reevaluation of pitch models built on earlier findings. We also found that as pulse duration is decreased to less than two cycles its pitch becomes biased toward higher frequencies, consistent with predictions of an auto-correlation model of pitch extraction.


Keywords Discrimination • Computational modeling • Pitch

[^0]
## Introduction

In 1946, Dennis Gabor published his seminal work on communication theory based on Heisenberg's uncertainty principle in quantum physics. He showed that one cannot simultaneously specify a sound's exact frequency and time of occurrence. Encapsulated in the mathematical identity $\Delta \mathrm{f} \Delta \mathrm{t} \geq 0.5$, the theory states that there is a tradeoff between temporal and spectral resolution. In colloquial terms, the briefer the sound, the broader is its observed spectrum. Transient sounds such as clicks have broad bandwidths. Pure tones of long durations have narrow bandwidths. The question then arises as to the efficiency with which the auditory system can perceptually encode the pitch of very brief sounds given the limitations imposed on physical stimuli by Gabor's uncertainty principle.

Several studies have investigated the minimum number of pure-tone periods required for reliable identification or discrimination of pitch (Freyman and Nelson, 1986; Henning, 1970; Hsieh and Saberi, 2007; Kietz, 1963; Konig, 1957; Moore, 1973; Patterson et al., 1983; Robinson and Patterson, 1995; Ronken, 1971; Savart, 1830; Sekey, 1963; Turnbull, 1944; von Békésy, 1972). The question has been of interest not only for what it can reveal about how pitch salience declines as a function of duration, but also for what it may contribute to models of pitch encoding (Freyman and Nelson, 1986; Hsieh and Saberi, 2007; Moore, 1973; Patterson et al., 1983; Robinson and Patterson, 1995; Zwicker, 1970). To our knowledge, two studies have attempted to evaluate pitch extraction from partial- or single-cycle tones. Sipovsky et al. (1972) reported a $2 \%$ frequency discrimination threshold for a 0.5 -cycle pure tone $(\Delta \mathrm{f}=30 \mathrm{~Hz}$ at 1500 Hz$)$ and Mark and Rattay (1990) reported thresholds as low as $5 \%$ for singlecycle tones. One difficulty with interpreting the results of
these studies is that discrimination thresholds may not have represented pitch extracted from waveform fine structure as intended, but on confounds associated with pulse duration and phase. Given a fixed number of cycles, changing stimulus frequency results in a change in duration and a detectable change in timbre associated with burst duration in a two-alternative forced-choice (2IFC) task. This is especially problematic for very brief tone pulses. Decreasing pulse duration results in an upward shift in the cutoff frequency of the pulse spectrum and hence an increase in high-frequency energy that may be used in a frequency discrimination task. Using zero-phase pulses also introduces a timbre confound in a 2IFC frequency discrimination task.

The current study was designed to investigate whether pitch cues may be extracted from the fine structure of partial-cycle pure tones under conditions that appropriately control for confounds. This has not been previously demonstrated. When confounds are accounted for, pitch discrimination performance is at chance for a 0.5 -cycle pulse. However, this does not mean that fine-structure pitch cues are inaccessible to the system. Subthreshold pitch cues may be detected (and hence quantified) if vectorially summed in a time series containing multiple pulses. In the current study, we measured pitch-discrimination thresholds at two frequency regions for pulse trains comprising partial-cycle tone pulses with random phases, amplitudes, and interpulse intervals (IPIs). We found that listeners can reliably discriminate the pitch of partial-cycle pulse trains, suggesting that pitch cues may in fact be extracted from a partial-cycle fine structure, but that this requires at least an order of magnitude higher thresholds than those reported in prior studies.

## Experiment 1: Discriminating the pitch of partial-cycle pure-tone pulses

## Methods

## Subjects

Five normal-hearing subjects, including one of the authors (IH), participated in Experiment 1. Three subjects participated in each part of this experiment, with one subject participating in both parts. All subjects were experienced as subjects in psychoacoustic experiments, and were additionally practiced on the various conditions of the experiment prior to data collection. The experimental protocol was approved by the Institutional Review Board of the University of California, Irvine,
and all subjects signed informed-consent forms before participating in the experiments.

## Part 1: Frequency discrimination for a single partial-cycle pulse

Part 1 was a brief experiment designed to measure frequencydiscrimination ability for a random-phase partial cycle pulse, to be contrasted with the results of Sipovsky et al. (1972) who used zero-phase pulses. Pure tones were generated at a sampling rate of 44.1 kHz using Matlab software (Mathworks) and presented monaurally (right ear) through Sennheiser headphones (HD 380 Pro) in a double-walled steel acoustically isolated chamber (Industrial Acoustics Company). Stimuli were 0.5 -cycle of a pure tone at 1500 Hz , the same frequency used by Sipovsky. The starting phase of the pulse was randomized on each presentation. Frequency discrimination performance was measured in a 2IFC task for $\Delta \mathrm{f}=30$, which is equal to the threshold value reported by Sipovsky, and 60 Hz twice the reported threshold value. Each subject completed three runs of 100 trials each with the two $\Delta \mathrm{fs}$ presented with equal prior probabilities within the same run.

## Part 2: Frequency discrimination for partial-cycle pulse

 trainsThe top panel of Fig. 1 shows a brief segment of an example stimulus used in this experiment. Stimuli were trains of 25 partial-period pulses with interpulse intervals (IPIs) randomly selected from a uniform distribution with a $20-50 \mathrm{~ms}$ range (i.e., successive pulse-pairs within a train had different IPIs). Within a train, the pulse-tone frequency and the number of cycles per pulse were held constant. Within a run, the frequency difference between the standard and comparison stimuli was also held constant. Three stimulus parameters were investigated: (1) base frequency ( 100 or 1000 Hz ), (2) number of cycles within a pulse ( $0.5,0.75$, and one), and (3) the frequency difference between the base and comparison tones (onethird, one-half, and one octave). ${ }^{1}$ The ability to discriminate the pitch of the base tone from that of the comparison was examined for trains of fixed-frequency pulses in a 2IFC block design. The interstimulus interval (ISI) between the first and second trains was 250 ms . Subjects had to determine which of two intervals contained the higher pitch sound by pressing a number key ( 1 or 2 ) to record their response.

To eliminate discrimination based on energy, loudness, or timbre, the level of each pulse of the train was independently perturbed by 12 dB , its phase selected from a uniform $(0-2 \pi)$

[^1]0.5 cycles at 1000 Hz (random phase)


Fig. 1 Top panel: Stimuli used in the current experiment comprised partial cycle tone pulses with randomized phase, amplitude, and interpulse intervals (see Methods for details). Bottom panel: Sipovsky et al. (1972) reported a frequency-discrimination $(\Delta f)$ threshold of 30 Hz for a $1500-\mathrm{Hz}$ zero-phase half-cycle tone (blue line). We found that phase randomization reduces performance to chance (0.5) even for a $\Delta \mathrm{f}$ twice that used by Sipovsky et al. Each open symbol represents data from one subject, with red lines representing mean performance
distribution, and the overall level of the train in each interval of the trial was randomized by 6 dB about a nominal value of 73 dB SPL. Levels were calibrated using a 6-cc coupler, $0.5-$ in. microphone (Brüel\&Kjær, Model 4189), and a Precision Sound Analyzer (Brüel\&Kjær, Model 2260). To eliminate pitch cues associated with pulse duration, the duration of each pulse was set to a constant mean value for both intervals of the 2IFC and additionally perturbed about this mean by $20 \%$. The mean value was selected based on the number of cycles of the higher frequency train in the 2IFC task. For example, for a $100-\mathrm{Hz}$ base tone, a period fraction of 0.5 per pulse, and a frequency difference of one-third octave between base and comparison trains, the comparison tone has a frequency of 126 Hz and a half-period of 3.97 ms . This value was used as the mean duration (perturbed by $20 \%$ ) of each pulse of both the higher and lower frequency trains. Therefore, each pulse of the lower frequency train ( $100-\mathrm{Hz}$ base) had a cycle fraction of $\sim 0.4$ ( $20 \%$ randomized) instead of the nominal 0.5 cycles. We selected the higher frequency pulse to set pulse durations for both trains to ensure that no pulse cycle fraction exceeded the target cycle fraction (i.e., 0.5 in this example).

Each subject completed four runs of 50 trials per each experimental condition in a random-block design. Prior to each
run, one experimental condition was randomly selected for that run. This procedure was continued until every condition was run at least once before a second set of runs began and until a total of four runs per condition per subject were completed. The order of runs was different for each subject. Response feedback was provided after each trial.

## Results

The bottom panel of Fig. 1 shows results for part 1 of Experiment 1 (single-pulse condition). Each circle represents the averaged data from one subject. Red horizontal lines near the 0.5 proportion correct show averaged performance across the three subjects for each of the two $\Delta \mathrm{f}$ conditions. The blue horizontal line at the 0.75 proportion correct shows averaged threshold reported by Sipovsky et al. (1972) for a zero-phase 0.5 -cycle pulse at the same base frequency used in the current experiment $(1500 \mathrm{~Hz})$. Note that for partial cycle pulses with a random phase, performance is at chance level, even for a $\Delta \mathrm{f}$ twice that reported by Sipovsky ( $60 \mathrm{~Hz}: \mathrm{t}(2)=0.994$, n.s.).

Figures 2 and 3 show results of part 2 of this experiment for the $100-$ and $1000-\mathrm{Hz}$ base frequencies, respectively, with mean performance shown in the lower-right panel of each figure. On average, subjects performed above chance in nearly all conditions, with performance monotonically increasing as the frequency difference between the base and comparison trains increased. The lower base frequency $(100 \mathrm{~Hz})$ produced slightly lower performance levels for the 0.75 and one-cycle conditions, and more variable performance relative to the higher base frequency $(1000 \mathrm{~Hz})$. Nonetheless, subjects reliably discriminate the pitch of two trains of partial-cycle tones even for 0.5 -cycle pulse trains. A two-way repeated measures analysis of variance on the data of Fig. 2 showed no significant effect of the number of cycles $(F(2,4)=4.05, p=0.1)$, a significant effect of frequency difference $(\mathrm{F}(2,4)=25.92$, $\mathrm{p}<0.005$ ), and a significant interaction between number of cycles and frequency difference $(\mathrm{F}(4,8)=7.00, \mathrm{p}<0.05)$. For the $1000-\mathrm{Hz}$ base frequency, subjects were able to discriminate trains of 0.5 -cycle pulses that were one-third octave $(\sim 260 \mathrm{~Hz})$ apart at an average accuracy rate of $70 \%$. This is nearly ten times higher than that reported by Sipovsky et al. (1972), who measured a $30-\mathrm{Hz}$ frequency discrimination threshold for a $0.5-c y c l e ~ p u l s e ~ o f ~ a ~ 1500-H z ~ p u r e ~ t o n e ~$ ( 0.028 octaves). For 0.75 - and one-cycle pulses, performance was near ceiling levels when the frequency difference was at least half an octave (open symbols). A two-way repeated measures analysis of variance on the data of Fig. 3 showed a significant effect of the number of cycles $(\mathrm{F}(2,4)=34.1$, $\mathrm{p}<0.005$ ), a significant effect of frequency difference ( $\mathrm{F}(3$, $6)=177.24, \mathrm{p}<0.001$, and a significant interaction between number of cycles and frequency difference $(\mathrm{F}(6,12)=9.63$,


Fig. 2 Results of Part 2 of Experiment 1 for for a base frequency of 100 Hz . Each panel shows data from one subject, except for the bottom-right panel which shows mean data across the three subjects. The parameter is number of periods in each pulse. Error bars are one standard deviation
$\mathrm{p}<0.005)$. The interaction effect is largely due to ceiling effects as performance converges to near perfect levels for the 0.75 - and one-cycle conditions (lower right panel of Fig. 3).

## Experiment 2: Effects of phase, duration, and level perturbation

In experiment 2 we investigated frequency-discrimination ability for: (1) fixed- versus random-phase pulses, (2) fixed versus random duration pulses, and (3) very large level perturbations ( 40 dB ) across the two intervals of a trial. In Condition 1 , we compared frequency-discrimination ability for a 0.5 -cycle tone pulse that either had a zero or random starting phase, but equal durations across the two intervals of a trial. The purpose was to determine if pitch discrimination in the random-phase condition is based on statistical "reconstruction" of a full cycle by sampling different segments of the sinusoidal curvature across pulses. Note that because pulse durations in both intervals of a trial were equal, no spectraledge artifacts were present contrary to prior studies (Mark and Rattay, 1990; Sipovsky et al., 1972). If the random-phase condition produces significantly lower thresholds relative to the zero-phase condition, it may suggest that a statistical
reconstruction of the full waveform could be a viable explanation for the observed performance in Experiment 1. However, if thresholds for the zero-phase condition are significantly lower than those for the random-phase condition, then one may conclude that either phase perturbation partially disrupts pitch integration or timbre cues facilitate discrimination performance in the zero-phase case. If they are equal, it would suggest that spectral and timbre distortions resulting from phase perturbation do not significantly interfere with pitch integration across brief tone pulses.

In Condition 2, we compared frequency-discrimination ability between two conditions: (1) fixed-duration, where pulses in both intervals of the 2IFC had the same durations, and (2) fixed number of cycles, in which each pulse was exactly 0.5 cycles, and hence had different durations in the two intervals of a trial because pulses in the higher frequency interval had correspondingly shorter durations. The goal was to determine the extent to which the pitch associated with pulse duration (i.e., $1 /$ duration) contributes to frequency discrimination. All other stimulus parameters were randomized as described in the methods section of Experiment 1 (i.e., IPI, level of each pulse, level of each train, phase of each pulse, etc.). Our expectation was that phase randomization would distort the duration-difference cue, yielding near equal


Fig. 3 Results of Part 2 of Experiment 1 for a base frequency of 1000 Hz . Parameters and subjects are the same as those described for Fig. 2
performance for the two conditions (fixed vs. random duration). In Condition 3, we examined the effects of very large level perturbations ( 40 dB ) across the two intervals of the 2IFC. In Experiment 1, the overall level of each train was perturbed by 6 dB . Although we felt that 6 dB is sufficient to eliminate energy-based cues, we decided to use extreme level perturbations (Green, 1988) partly because for frequency-discrimination tasks, changes in tone frequency are often correlated with a concomitant change in loudness (Henning, 1966), and partly to evaluate the resiliency with which subjects can extract pitch cues from highly variable stimulus contexts.

## Methods

The same subjects who participated in part 2 of Experiment 1 also participated in this experiment. All procedures and stimulus parameters were the same as those described for Experiment 1 except for the following: The base frequency was 1 kHz for all conditions. In addition, the following changes were made specific to each condition.

Condition 1: Pitch-discrimination performance for random-phase pulses was compared to zero-phase pulses. The frequency difference between base and comparison
pulses was 0.5 octaves ( 1 vs .1 .41 kHz ). The duration of each pulse was 0.35 ms , the half period of the higher-frequency tone (i.e., 0.5 cycles at $1.41 \mathrm{kHz}, 0.35$ cycles at 1 kHz ). This duration was further randomized by $20 \%$ for each pulse of each train.

Condition 2: Pitch-discrimination performance for equal durations of base and comparison pulses (and hence varying number of cycles) was compared to that for an equal number of cycles for the base and comparison pulses (and hence varying durations). The frequency difference between base and comparison trains was half an octave ( 1 vs .1 .41 kHz ). Each pulse either had a mean duration of 0.5 ms in both intervals of the equal-duration condition, or had 0.5 cycles in the equalcycle condition. Starting phases were randomized for both conditions.

Condition 3: The level of each pulse train was randomized by 40 dB across the two intervals of the 2IFC task. Performance in this condition was compared to no-level randomization across intervals of a trial. As before, the level of each individual pulse within a train was randomized by 12 dB in both conditions. Based on pilot listening, we selected a frequency difference of one octave between the base and comparison tone frequencies ( 1 and 2 kHz ). Pulses within the two trains had equal durations, anchored to the duration of the base frequency pulse at a mean
value of 0.75 ms (or 0.75 cycles). This duration was independently selected for each pulse of each train from a uniform distribution with a range of $0.65-0.85 \mathrm{~ms}$. The purpose of randomizing durations, phases, levels of pulses within a train (12 dB), and across trains ( 40 dB ) was to isolate discrimination to pitch differences.

## Results

Figure 4 shows mean results from three subjects for the three experimental conditions. Error bars are one standard deviation. Left pair of bars show that randomizing the starting phase of partial-cycle pulses does not affect pitch-discrimination ability $(\mathrm{t}(2)=0.56$, n.s.). Averaged performance is near $90 \%$ and nearly equal in the two conditions. It is important to note that the durations of all pulses, in both intervals of the 2IFC task, were randomized by $20 \%$ about a fixed value of 0.35 ms to eliminate timbre and spectral-edge cues. Thus, it does not appear that the high levels of pitch-discrimination performance observed in Experiment 1 are the results of statistical reconstruction of the full curvature of tone cycle from sampling different starting phases. One should, however, be cautious in this interpretation since performance for the equalcycle condition is near ceiling levels.

The middle pair of bars in Fig. 4 show that the equalduration pulses produce only slightly lower performance than the equal-cycle pulses ( $97 \%$ vs. $90 \%$ ). This difference is not statistically significant $(\mathrm{t}(2)=1.21$, n.s.), and suggests that the pitch associated with pulse duration (i.e., $1 /$ duration) contributes little to pitch discrimination above that extracted from the fine structure of each pulse. As was the case for Condition 1, one must be cautious in interpretation of these results since performances in these conditions are near ceiling levels.


Fig. 4 Results of Experiment 2. The left pair of bars show the difference in frequency discrimination performance between random-phase and zero-phase pulse trains. The base and comparison frequencies were half an octave apart ( 1 and 1.41 kHz ). Each pulse was 0.5 cycles, randomized by $20 \%$. The middle pair of bars show the effects of using pulses of equal durations versus pulses that had an equal number of cycles (and hence different durations). The right pair of bars show the effects of perturbing the level of each pulse by 40 dB . Error bars are one standard deviation

Furthermore, note that the starting phases were randomized for both cases, and hence, spectral cues from waveform fine structures were severely distorted in both cases.

The right pair of bars in Fig. 4 show that perturbing the overall level of each pulse train by 40 dB still allows for pitchdiscrimination performance above $80 \%$ for 0.75 -cycle pulse trains that are one octave apart in frequency. Performance, however, is near perfect ( $100 \%$ correct) when the level of each train is not randomized, even though the level of each pulse within the train was perturbed by 12 dB . There is a statistically significant difference between pitchdiscrimination performance for these two conditions $(\mathrm{t}(2)=$ 8.44, $\mathrm{p}=0.014$ ), suggesting that extreme level randomization reduces, but does not eliminate, the ability to discriminate the pitch of partial-cycle pulses.

## Discussion

The current findings suggest that there likely is sufficient finestructure pitch information in less than one cycle of a pure tone for reliable discrimination of low frequency tones, provided that this information is accumulated across multiple pulses. The two base frequencies employed in our study are in the region of the spectrum typically associated with temporal models of pitch extraction, and, in particular, autocorrelation. To determine the accuracy with which a model of the auditory periphery can predict pitch-discrimination performance for the stimuli used in our study, we examined the output of an autocorrelation model with several frontend preprocessing stages. The model consisted of a bank of 30 fourthorder GammaTone bandpass filters spaced logarithmically from 50 to 3000 Hz (Holdsworth et al., 1988; Hsieh and Saberi, 2007, 2009; Hsieh et al., 2010, 2011; Saberi and Petrosyan, 2005). Filter bandwidths were based on human auditory filter estimates measured in notched-noise (Glasberg and Moore, 1990). The filterbank was followed by half-wave rectification and square-law nonlinearity (Saberi et al., 2004; Shear, 1987). The signal output within each channel was weighted by a frequency-dependent function representing outer- and middle-ear attenuation. This weighting function was derived from a logistic fit to the data shown in Fig. 2 of Meddis and Hewitt (1991). The filter outputs were followed by autocorrelation within frequency channels, frequency integration, and a decision device.

The left panels of Fig. 5 show the output of this model prior to frequency integration for a $1-\mathrm{kHz}$ tone pulse containing either 0.5 , one, or two cycles (top to bottom panels, respectively). Note that as cycle number increases the bandwidth decreases and energy becomes more focused at 1 kHz . The middle panels show these outputs after integration across frequency channels. The vertical dashed red line shows the expected position of the first autocorrelation peak at a positive


Fig. 5 Output of an autocorrelation model of the auditory periphery. The three rows of panels show analysis for 0.5 , one, and two cycles of a $1-\mathrm{kHz}$ tone, respectively. The first column shows the model output prior to integration across frequency channels. The second column shows this
output after frequency integration. The red dashed line is the expected peak for a $1-\mathrm{kHz}$ tone (i.e., 1 ms ). Right panels show model output for the type of pulse trains used in the current study (see text for details)
lag (i.e., 1 ms ). No peak is observed at 1 ms for the 0.5 -cycle pulse. A small peak is observed for the one-cycle pulse (middle panel), but at a lag that is slightly smaller than the predicted peak, yielding an overestimation of the expected pitch. For a two-cycle pulse, the autocorrelation peak occurs very near but still below 1 ms . The right panels of Fig. 5 show the model output for the 25 -pulse train used in the current study. Note that peaks are evident for non-zero lags even for the 0.5 -cycle condition (top right), and are even more pronounced for oneand two-cycle conditions relative to single-pulse conditions (middle panels). The 0.5- and one-cycle pulse trains produce peaks that underestimate the stimulus fine structure period, with a larger undershoot associated with the 0.5 -cycle pulse train. The two-cycle pulse train again produces autocorrelation peaks very near, but slightly lower than the expected period. In addition, as the carrier frequency is increased (not shown), the expected autocorrelation peaks move closer to zero (i.e., 0.5 ms for a $2-\mathrm{kHz}$ tone). This suggests that sufficient information is available in the autocorrelation function of partial-cycle pulse trains for pitch discrimination, i.e., a judgment based on the relative positions of peaks within the autocorrelation function. We observed nearly identical patterns for a $100-\mathrm{Hz}$ base tone.

To quantitatively determine if the predictions described above follow patterns similar to those observed in our data
(e.g., Fig. 3), Monte Carlo simulations were run for the $1000-\mathrm{Hz}$ condition from Experiment 1. Because the pulsetrain contains random perturbations (e.g., phase, amplitude, and interpulse interval), the autocorrelation output will also contain peaks that are perturbed in position on each trial. For each of the 12 combinations of frequency separation (zero, one-third, one-half, and one octave) and number of cycles ( $0.5,0.75$, and one), we ran 5000 runs of 100 trials each, with


Fig. 6 Predictions of the autocorrelation model for the stimulus conditions used in the current experiment $(1-\mathrm{kHz}$ carrier). Each point is based on 5000 runs of 100 trials each in a Monte Carlo simulation
the single free parameter of the model being the magnitude of internal noise added independently on each trial to each delay-by-frequency channel. The decision device generated a binary choice on which of two intervals of the 2IFC trial contained a higher pitch estimate based on the relative positions of autocorrelation peaks in the two intervals. Figure 6 shows predictions of this simulation (cf. Fig. 3). The model captures two main trends in the data. First, performance monotonically improves with increased separation between the base and comparison tone frequencies. Second, performance improves monotonically with an increase in the number of cycles per pulse. The model also predicts a larger difference in performance between the 0.5 - and 0.75 -cycle conditions relative to that between the 0.75 - and one-cycle conditions, a trend which is also observed in the data of Fig. 3. The model, however, predicts better performance for the 0.5 -cycle condition relative to what we have observed in our data. This prediction may be improved, but at the cost of an additional free parameter which we wanted to avoid.

Although pitch salience is relatively weak for a single pulse containing fewer than two cycles of a pure tone, a stronger sense of pitch does arise when listening to trains of such pulses. Savart (1830) and von Békésy (1972) have suggested that two cycles of a sinusoid are sufficient to give rise to a sense of pitch. Although a sense of pitch may arise with two cycles, the pitch generated by these stimuli are possibly biased toward higher frequencies, as suggested by autocorrelation analysis (Fig. 5). Our own experiments with absolute pitch listeners show that a minimum of four cycles are required for accurate identification of the pitch of a sinusoid in isolation (Hsieh and Saberi, 2007).

In summary, the current study aimed to determine whether there are extractable pitch cues in the finestructure of very brief tone pulses. Our results are different from those reported in prior work in that when fine-structure pitch is isolated for discrimination, performance is at chance and only exceeds chance when this information is integrated over many pulses. Furthermore, these pitch cues are likely to be biased toward higher frequencies and significantly weaker than those reported in previous work. We suggest that pitchencoding models which have employed findings from these earlier reports be re-evaluated to account for how these higher than reported thresholds affect their predictions.

Acknowledgments This work was supported by grants from the National Science Council, Taiwan (NSC 98-2410-H-008-081) and the National Institutes of Health (R01DC009659).

## References

Freyman, R. L., \& Nelson, D. A. (1986). Frequency discrimination as a function of tonal duration and excitation-pattern slopes in normal and hearing-impaired listeners. Journal of the Acoustical Society of America, 79, 1034-1044.
Gabor, D. (1946). Theory of communication. Journal of the Institute of Electrical Engineering, 93, 429-457.
Glasberg, B. R., \& Moore, B. C. J. (1990). Derivation of auditory filter shapes from notched-noise data. Hearing Research, 47, 103-138.
Green, D. M. (1988). Profile analysis: Auditory intensity discrimination. New York: Oxford University Press.
Henning, G. B. (1966). Frequency discrimination of random amplitude tones. Journal of the Acoustical Society of America, 39, 336-339.
Henning, G. B. (1970). Effects of duration on frequency and amplitude discrimination. In R. Plomp \& G. F. Smoorenburg (Eds.), Frequency analysis and periodicity detection in hearing. Leiden: A.W. Sijthoff.
Holdsworth, J., Nimmo-Smith, I., Patterson, R., \& Rice, P. (1988). Implementing a Gammatone filterbank [SVOS final report: Annex C. Part A, the Auditory Filter Bank]. Cambridge, United Kingdom: Medical Research Council Applied Psychology Unit.
Hsieh, I., Petrosyan, A., Goncalves, O., Hickok, G., \& Saberi, K. (2010). Cross- modulation interference with lateralization of mixedmodulated waveforms. Journal of Speech, Language, and Hearing Research, 53, 1417-1428.
Hsieh, I., Petrosyan, A., Goncalves, O., Hickok, G., \& Saberi, K. (2011). Observer weighting of interaural cues in positive and negative envelope slopes of amplitude modulated waveforms. Hearing Research, 277, 143-151.
Hsieh, I., \& Saberi, K. (2007). Temporal integration in absolute identification of musical pitch. Hearing Research, 233, 108-116.
Hsieh, I., \& Saberi, K. (2009). Detection of spatial cues in linear and logarithmic frequency-modulated sweeps. Attention, Perception, \& Psychophysics, 71, 1876-1889.
Kietz, H. (1963). Frequency-difference limens with extremely short tonal duration. International Audiology, 2, 48-51.
Konig, E. (1957). Effect of time on pitch discrimination thresholds under several psychophysical procedures - comparison with intensity discrimination thresholds. Journal of the Acoustical Society of America, 29, 606-612.
Mark, H. E., \& Rattay, F. (1990). Frequency discrimination of singlecycle, double-cycle, and triple-cycle sinusoidal acoustic signals. Journal of the Acoustical Society of America, 88, 560-563.
Meddis, R., \& Hewitt, M. J. (1991). Virtual pitch and phase sensitivity of a computer model of the auditory periphery. I: Pitch identification. Journal of the Acoustical Society of America, 89, 2866-2882.
Moore, B. C. J. (1973). Frequency difference limens for short-duration tones. Journal of the Acoustical Society of America, 54, 610-619.
Patterson, R. D., Peters, R. W., \& Milroy, R. (1983). Threshold duration for melodic pitch. In R. Klinke \& R. Hartmann (Eds.), Hearing: Physiological bases and psychophysics. Berlin: Springer Verlag.
Robinson, K., \& Patterson, R. D. (1995). The stimulus duration required to identify vowels, their octave, and their pitch chroma. Journal of the Acoustical Society of America, 98, 1858-1865.
Ronken, D. A. (1971). Some effects of bandwidth-duration constraints on frequency discrimination. Journal of the Acoustical Society of America, 49, 1232-1242.
Saberi, K., Antonio, J. V., \& Petrosyan, A. (2004). A population study of the precedence effect. Hearing Research, 191, 1-13.
Saberi, K., \& Petrosyan, A. (2005). Neural cross-correlation and signal decorrelation: Insights into coding of auditory space. Journal of Theoretical Biology, 235, 45-56.

Savart, F. (1830). Notes sur la sensibilité de l'organe de l'ouïe. Annales de Chimie et de Physique, [Notes on the sensitivity of the organ of hearing. Annals of Chemistry and Physics], 44, 337-352.
Sekey, A. (1963). Short-term auditory frequency discrimination. Journal of the Acoustical Society of America, 35, 682-690.
Shear, G.D. (1987). Modeling the dependence of auditory lateralization on frequency and bandwidth. Master's Thesis, Carnegie Mellon University, Pittsburgh, PA.
Sipovsky, A. V., Gershuni, G. V., Gorelik, B. M., Korotkin, I. I., \& Lubinsky, I. A. (1972). The determination of differential frequency thresholds for short tone signals. Biofizika, 17, 495-502.

Turnbull, W. W. (1944). Pitch discrimination as a function of tonal duration. Journal of Experimental Psychology, 34, 302-316.
von Békésy, G. (1972). The missing fundamental and periodicity detection in hearing. Journal of the Acoustical Society of America, 51, 631-637.
Zwicker, E. (1970). Masking and psychological excitation as consequences of the ear's frequency analysis. In R. Plomp \& G. F. Smoorenburg (Eds.), Frequency analysis and periodicity detection in hearing (pp. 376-394). Leiden, Netherland: Sijthoff.


[^0]:    I-Hui Hsieh
    ihsieh@ncu.edu.tw
    Kourosh Saberi
    saberi@uci.edu
    1 Institute of Cognitive Neuroscience, National Central University, Taoyuan City, Taiwan
    2 Department of Cognitive Sciences, University of California, Irvine, CA 92697-5100, USA

[^1]:    ${ }^{1}$ These parameter values were selected based on initial pilot runs. With one exception (one-third octave, 0.5 cycle, $100-\mathrm{Hz}$ base), all combinations of frequency difference by cycle numbers were used for both base frequencies.

