Perceptual timing precision with vibrotactile,
auditory, and multisensory stimuli

Mercedes B. Villalonga

Volen Center for Complex Systems and Department of Psychology

Brandeis University, Waltham MA USA

Rachel F. Sussman

Department of Psychology and Program in Neuroscience

Brandeis University, Waltham MA USA

Robert Sekuler

Volen Center for Complex Systems and Department of Psychology

Brandeis University, Waltham MA USA

Author Note

Corresponding author: Robert Sekuler, Volen Center for Complex Systems,
Brandeis University, Waltham, MA, USA. Email: vision@brandeis.edu

Abstract

1 The growing use of vibrotactile signaling devices makes it important to understand the
2 perceptual limits on vibrotactile information-processing. To promote that understanding,
3 we carried out a pair of experiments on vibrotactile, auditory, and bimodal (synchronous
4 vibrotactile and auditory) temporal acuity. On each trial, subjects experienced a set of
5 standard intervals (400 ms each), followed by one interval of variable duration (400
6 $\pm$ 1-80 ms). Intervals were demarcated by short vibrotactile, auditory, or bimodal pulses.
7 Subjects categorized the timing of the last interval by describing the final pulse as either
8 “early” or “late” relative to its predecessors. In Experiment One, each trial contained
9 three isochronous standard intervals, followed by an interval of variable length. In
10 Experiment Two, the number of isochronous standard intervals per trial varied, from one
11 to four. Psychometric modeling revealed that vibrotactile stimulation produced poorer
12 temporal discrimination than either auditory or bimodal stimulation. Moreover, auditory
13 signals dominated bimodal sensitivity, and inter-individual differences in temporal
14 discriminability were reduced with bimodal stimulation. Additionally, varying the number
15 of isochronous intervals in a trial failed to improve temporal sensitivity in either modality,
16 suggesting that memory played a key role in judgments of interval duration.

17 *Keywords:* Temporal sensitivity, auditory perception, vibrotactile perception,
18 multisensory perception, timing, memory, individual differences

Perceptual timing precision with vibrotactile, auditory, and multisensory stimuli

Introduction

Our sensory systems provide us with timely and reliable warnings about consequential changes in our surroundings. Often, these changes originate from technologies that emit sensory alerts to notify their users. One way technologies can create salient signals is by generating rhythmic patterns of light, sound, or vibrations. Our growing reliance on auditory and vibrotactile signaling devices like cell phones makes it important to understand the capabilities and limitations of these sensory systems as communication channels for temporal information (Meng & Spence, 2015).

Surprisingly little is understood about the potential of the skin, our largest sensory organ, for communicating temporal information. An influential essay by von Békésy (1959) provides a valuable introduction to the skin's advantages and disadvantages in this domain. Comparing the temporal responsiveness of audition and somatosensation, von Békésy (1959) wrote that "all the phenomena in which time patterns are involved can be discriminated by the skin only if the changes are slow relative to changes that the ear can recognize" (p. 6). While von Békésy's generalization has enjoyed considerable support (Azari, Mioni, Rousseau, & Grondin, 2020; Grondin & Rousseau, 1991; Jones, Poliakoff, & Wells, 2009; Lauzon, Russo, & Harris, 2016; Mayer, Di Luca, & Ernst, 2014), there has been at least one dissenting view, namely an assertion that the two modalities possess similar temporal acuity (Nagarajan, Blake, Wright, Byl, & Merzenich, 1998).

As a way to clarify the matter, we decided that studying bimodal auditory and somatosensory stimulation could afford additional insight into each modality's individual capacity for communicating temporal information. Our approach builds on the modality appropriateness hypothesis (Welch & Duttonhurt, 1986; Welch & Warren, 1980), which states that the spatial and temporal characteristics of different modalities are optimized for different types of processing, where vision is purportedly specialized for processing

spatial information and audition is instead specialized for processing temporal information. Studies of audio-visual temporal processing support this hypothesis, having shown that auditory cues dominate bimodal judgments (Burr, Banks, & Morrone, 2009; Repp & Penel, 2002; Walker & Scott, 1981). Adapting this approach, we studied duration discrimination for vibrotactile and auditory stimuli, and also for bimodal stimuli in which the two were combined. Following the modality appropriateness hypothesis, if auditory temporal precision were actually superior to vibrotactile temporal precision, auditory stimuli should dominate when the two are paired.

Here, we report two experiments comparing temporal acuity for auditory and vibrotactile stimuli. In both experiments, subjects were presented with a series of empty intervals demarcated by auditory, vibrotactile, or bimodal pulses. Isochronous standard intervals were followed by a comparison interval of variable duration. Subjects compared the duration of the final interval relative to that of the isochronous standard intervals. Our first experiment evaluated auditory, vibrotactile, and bimodal temporal sensitivity by comparing performance across those three conditions, keeping the number of standard intervals presented constant. There is some evidence, however, that varying the number of repetitions of a standard stimulus affects temporal sensitivity (Grondin, 2001; Grondin & McAuley, 2009; Miller & McAuley, 2005; ?, ?) It is unclear whether this effect arises from influences in early sensory processing (*e.g.*, von Békésy, 1959), from influence by higher-order cognitive processes, or both (VanRullen & Thorpe, 2001; Agam & Sekuler, 2007; Cavanagh, 2011; Waskom, Okazawa, & Kiani, 2019). Our second experiment tested whether the number of standard intervals on each trial affected temporal sensitivity in each unimodal condition.

Experiment One

Our first experiment adapted a task introduced by Lumaca, Trusbak Haumann, Brattico, Grube, and Vuust (2019). In our adaptation, shown schematically in Fig. 1, subjects received sequences of four empty intervals, each demarcated by brief auditory,

73 vibrotactile, or bimodal pulses, indicated by the solid vertical lines in the figure. The first
 74 three intervals were isochronous, each 400 ms long. The duration of the fourth interval,
 75 however, was randomized so that its terminating pulse could arrive either earlier or later
 76 than 400 ms after the previous pulse. Subjects categorized the timing of the final pulse
 77 as either “early” or “late” relative to the preceding interval(s). We used these judgments
 78 as an index of subjects’ perception of the duration of the final interval.

Figure 1. Sample “early” (A) and “late” (B) trials. Each 50-ms stimulus pulse is represented by a solid vertical line. Three standard intervals, each 400 ms, preceded a final comparison interval of variable duration. The sign and magnitude of the variable x represents the trial’s *rhythmic deviancy*: on an “early” trial (A), the comparison interval was shorter than 400 ms (*i.e.*, $400-x$), and on a “late” trial (B), the comparison interval was longer than 400 ms (*i.e.*, $400+x$). Dashed vertical lines represent when the final stimulus pulse would have occurred with no rhythmic deviancy.

Methods

Subjects. Twenty-one subjects (13 female; mean age = 20.8 years, SD = 2.1 years) participated in this experiment. All subjects denied having difficulty in sensing vibrations on the skin or in hearing. We performed an *a priori* power analysis to assess whether our sample size was large enough to detect a previously reported effect size in a similar task (Jones et al., 2009) with 85% power. The *Superpower* package in R (Caldwell & Lakens, 2019) showed that a sample of 10 subjects would be needed to observe a similar effect, confirming the sufficiency of our sample size. All experimental procedures had been approved by Brandeis University's Institutional Review Board, and work was conducted in accordance with the Declaration of Helsinki. All subjects provided written informed consent prior to their participation. During the consent process, subjects were informed that they would receive at least \$12.00 (USD) for their time, as well as an additional \$0.25 for each block on which their accuracy exceeded 80%.

Apparatus and stimuli. Stimulus presentation and response collection were controlled by a Macintosh computer running MATLAB (version 9.3.0; 2017) and extensions from the Psychophysics Toolbox (version 3.0.14; Brainard, 1997). The computer worked in tandem with an Arduino UNO microprocessor that directly controlled the stimulus generation devices (Fig. 2).

Tactile (hereafter, **T**) stimulation was produced by a linear resonant actuator vibrotactor (Engineering Acoustics, Inc., Casselberry, FL). The active element of the vibrotactor was 7.90 mm in diameter, and produced vibrations at 250 Hz with < 2 ms rise time and 0.5 mm peak vertical displacement. Note that the frequency of our **T** stimuli, 250 Hz, is near the peak sensitivity for Pacinian corpuscles, the mechanoreceptor type most sensitive to vibrations (Makous, Friedman, & Vierck, 1995; Mountcastle, LaMotte, & Carli, 1972). For Auditory (hereafter, **A**) stimuli we chose a higher frequency because detection of auditory stimuli is appreciably slower for low-frequencies like 250 Hz, compared to detection of higher-frequency tones (Woods, Alain, Covarrubias, & Zaidel,

1993). Our **A** stimuli were ~1000-Hz pulses at ~70 dB generated by a piezo active buzzer. Bimodal auditory-vibrotactile (hereafter, **AT**) stimuli were created by activating the vibrotactor and the piezo buzzer simultaneously, with a negligible ($<10\mu\text{s}$) delay between sending the signals to activate the vibrotactor and buzzer. Throughout the experiment, the intensities of the vibrotactile and auditory stimuli were held constant at levels well above threshold.

Vibrations were delivered to the ventral surface of the subject's left index finger, which rested on the vibrotactor's active element. To stabilize the finger's contact with the vibrotactor, the subject's hand was positioned within a close-fitting, custom 3D-printed container. This container also included insulation around the vibrotactor to reduce unexpected auditory emissions resulting from vibrations against the plastic of the 3D-printed container. The piezo buzzer used to generate the auditory stimuli was located next to the container housing the vibrotactor (Fig. 2).

Brownian noise calibration. Despite the insulation in its 3D-printed container, activation of the vibrotactor produced an audible, low-frequency sound. This sound was masked by continuous Brownian noise whose level was customized for each subject using a standard calibration procedure as described below to prevent this sound from contaminating the conditions in which the vibrotactor was used. Subjects heard the continuous Brownian noise through supraural headphones, and, at subjects' option, also wore in-ear, sound-absorbing earplugs.

The calibration procedure used a modified method of limits to find a level of Brownian noise that completely masked the sound of the vibrotactor and still allowed the piezo buzzer to be heard. At each level of masking noise, five **A** stimuli were presented with the buzzer; the subject then reported how many tones they heard. Brownian noise amplitude was decreased following incorrect responses. This procedure allowed us to find the maximal level of noise through which subjects could still hear the **A** stimuli. With this level of masking noise, we repeated the procedure with **T** stimuli, increasing noise

Figure 2. A test subject whose left hand is cradled within a 3D-printed support that positioned their index finger on the vibrotactor's active element. A subject used their right hand to communicate their binary judgments on the computer's numerical keypad. In the photo, the computer is displaying one of the experiment's instruction screens. The Arduino microprocessor and associated circuitry are enclosed in the 3D-printed box. Photo used with permission from this subject.

amplitude if subjects verbally noted they could hear the sound of the vibrotactor. This iterative process produced a level of noise that allowed the subject to hear **A** stimuli but not the sound generated by the vibrotactor.

Experimental task. On each trial, subjects received five 50-ms pulses of **A**, **T**, or **AT** stimulation. These five pulses demarcated four empty intervals, each defined by the temporal separation between successive pulses. The first three intervals were isochronous, at 400 ms each. However, the fourth interval (before the fifth and final pulse) was equally likely to be less than or greater than 400 ms. This made the final pulse occur either “earlier” or “later”, respectively, than a pulse following a standard,

400-ms interval. The difference between the duration of the final (comparison) interval and the duration of each standard interval defined what we will call the trial's *rhythmic deviancy*. The subject's task was to characterize the arrival of the fifth pulse as either "early" or "late" compared to the timing of the previous pulses. If the fourth interval had a *negative rhythmic deviancy* (duration < 400 ms; Fig. 1A), the last pulse arrived "earlier" than a pulse terminating a standard interval would have, and if the fourth interval had a *positive rhythmic deviancy* (duration > 400 ms; Fig. 1B), the last pulse arrived "later" than a pulse terminating a standard interval would have.

In each trial, a white fixation cross appeared on the center on a black computer screen for 500 ms. Then, the stimulus sequence began. After the fifth pulse, subjects signaled their responses with their right hands, using the "1" and "2" keys on a numeric computer keypad (Fig. 2). The mapping of the keys onto the response categories ("early" and "late") was counterbalanced across subjects. Following each response, immediate feedback ("Correct!" or "Incorrect.") was displayed on the screen for 1000 ms, along with the subject's average accuracy so far during that block of trials. Subjects had up to 1000 ms to enter their responses; if they failed to respond within that time, the trial was categorized as incorrect.

Preliminary testing confirmed that a range of rhythmic deviancies of ± 1 -80 ms produced response accuracies which ranged from 50% to nearly perfect (100%). Rhythmic deviancies for the experiment were therefore drawn from two ranges: 320-399 ms ("early" trials) and 401-480 ms ("late" trials). To facilitate our planned data analysis, each range was subdivided into four time bins, each spanning 20 ms (e.g., 320-339 ms). Within each time 20 ms bin, values were drawn randomly from uniform distributions. This arrangement of bins (i) allowed each psychometric function to be based on eight values of the independent variable, and (ii) facilitated serial communication between MATLAB and the Arduino. In total, for each modality, a subject received 20 trials with equal numbers of rhythmic deviancies drawn randomly from the eight different time bins (four

“early” and four “late”). This yielded a total of 480 trials per subject (3 [modality conditions] $\times 8$ [rhythmic deviancy bins] $\times 20$ [trials per modality condition and rhythmic deviancy bin]).

Subjects first completed the calibration procedure described above, and then completed an familiarization phase. During the familiarization phase, subjects received 2 example trials (1 “early” and 1 “late”), and then completed 6 practice trials (3 “early” and 3 “late”) for each modality (3 rounds total). The rhythmic deviancy on all pre-experimental trials was ± 60 -80 ms. All trials in the familiarization phase followed the general trial structure described above. Following the familiarization phase, each subject completed 12 blocks of 40 trials per block. The experiment followed a block-randomized design, with all experimental trials in a single block devoted exclusively to one modality. Subjects were permitted *ad lib* breaks between blocks, and received a bonus of \$0.25 for each block on which their accuracy exceeded 80%. During breaks, an on-screen text display informed the subject which modality would be tested in the next block of trials. Subjects completed the entire experiment, including instructions and practice, in ~ 45 -60 minutes.

Data curation. Subjects failed to respond on 0.8% of all trials. One-way ANOVAs and *t*-tests showed that the number of missed trials was unrelated to stimulus modality ($p = 0.170$), response key mapping ($p = 0.884$), or response category (“early” or “late”, $p = 0.358$). We subsequently excluded these trials from our analyses.

Psychometric modeling. Psychometric functions (PFs) supplied the main analytical parameters used to characterize subjects’ temporal sensitivity. Each PF modeled the likelihood that a subject would categorize a trial’s final pulse as “late” given the trial’s rhythmic deviancy value and its modality. To visualize the main trends and overall differences among modality conditions, we first fit a group-level model by aggregating data from all subjects. We then fit a subject-level model to generate distributions of estimated parameters for our analyses. We fit logistic, hierarchical Bayesian PFs using the simulation engine Just Another Gibbs Sampler (JAGS, Plummer,

2003), as implemented in the Palamedes toolbox for MATLAB (Prins & Kingdom, 2018). We used uninformative, flat priors for all model parameters, with one chain, 5000 (group-level)/1000 (subject-level) samples per chain, and 2000 discarded samples per parameter. Each PF was defined by four parameters: location along the rhythmic deviancy axis (α), slope (β), lapse rate (γ), and guess rate (λ).

We assessed goodness-of-fit of the subject-level model by the width of the 95% high-density interval (HDI) for the location parameter. Briefly, the 95% HDI bounds the inner 95% of the estimated parameter's posterior distribution, with smaller HDIs indicating more precise parameter estimation (Calin-Jageman & Cumming, 2019; Kruschke, 2013).

Treatment of model parameters. We used the estimated model parameters to derive two separate dependent measures: the just noticeable difference (JND) and the point of subjective equality (PSE). Each of these measures captures different aspects of subjects' timing judgments.

In duration discrimination studies, the JND captures sensitivity to changes in interval duration (Bausenhardt, Di Luca, & Ulrich, 2018). In our model, the JND indicates how much rhythmic deviance (in ms) is needed to produce a reliable change in "early" and "late" responses, with smaller values of JND denoting higher sensitivity. We calculated the subject-level JND for each modality by averaging the values of the inverse PF at $P(\text{"late"}) = 0.25$ and $P(\text{"late"}) = 0.75$.

The location parameter (α) provided a different perspective on temporal resolution, namely, in terms of the PSE. This parameter provides the interval duration subjects were equally likely to categorize as either "early" or "late", or the duration deemed perceptually equivalent to that of the standard intervals (400 ms). If comparison intervals were perceptually represented with perfect fidelity to the standard, that is, with no constant error, the PSE would be 400 ms. Subject-level PSEs for each modality were generated directly by our model-fitting procedures, where $\alpha = \text{PSE}$.

Quantitative evaluation of PSE and JND estimates. We used one-way, within-subject ANOVAs to guide interpretations of modality effects on subject-level distributions of JND and PSE values. Mauchly's test for sphericity assessed potential violations of the assumption of sphericity; if violated, we corrected degrees of freedom and p -values using the Greenhouse-Geisser epsilon. Because we expected a significant main effect of modality, we performed planned pairwise comparisons using estimated marginal means following the ANOVAs. In addition to p -values and effect sizes for each comparison, we report Bayes factors (BF_{10}) calculated using the *BayesFactor* package in R (Morey & Rouder, 2011). Reporting the BF_{10} allows us to present evidence in support of not only our alternative hypotheses (*i.e.*, observing a meaningful difference between conditions, $BF_{10} \geq 3$) but also our null hypotheses ($BF_{10} < 3$; Schönbrodt & Wagenmakers, 2018; Jeffreys, 1961).

Cue combination. For a quantitative account of how each type of unimodal cue influenced bimodal sensitivity, we adopted a model-based approach introduced by To, Baddeley, Troscianko, and Tolhurst (2011). As the slope β of a PF directly reflects sensitivity to temporal change, we used models of cue combination to estimate m , an index of the way that information from the two modalities was combined:

$$\beta_{AT} = (\beta_A^m + \beta_T^m)^{1/m} \quad (1)$$

Specifically, we tested whether the combination of unimodal cues followed linear summation (both cues weighted equally, $m = 1$) or a MAX metric (complete dominance by one of the two inputs, $m = \infty$). We also tested whether audio-tactile cue combination in our task mirrored the mode of audio-visual cue combination reported by To et al. ($m = 2.56$; 2011). Summed squared error (SSE) values were used to estimate model fit.

Post hoc analyses. We used *post hoc* tests to explore trends in the results of our primary analyses. To compare PSE estimates to 400 ms, we calculated the constant error (CE), defined as $CE = 400 \text{ ms} - \text{PSE}$. To test whether $CE = 0$, we ran Bayesian t -tests using R's *BEST* package (Kruschke, 2013). We also used Bayesian t -tests to

compare the standard deviations of the JND estimates across modality conditions. We considered standard deviations to be reliably different if the 95% HDI of the difference distribution did not include zero (Kruschke, 2013). We report adjusted R^2 values for each regression model.

Results

Table 1

Mean JNDs and PSEs for each modality, in milliseconds. Values of standard deviation are shown in parentheses.

Modality	Mean JND (SD)	Mean PSE (SD)	BF ₁₀ [95% HDI]
Auditory (A)	19.32 (4.90)	401.52 (5.40)	0.47 [-4.11, 1.04]
Tactile (T)	27.45 (5.80)	394.82 (4.43)	797.39 [3.06, 7.20]
Bimodal (AT)	18.03 (1.05)	397.95 (3.48)	3.88 [0.57, 3.69]

Note: BF₁₀ = Bayes factor quantifying support for H₁: CE = 400 ms - PSE ≠ 0.

Figure 3 shows the results of our model fits. Group-level PFs are shown in Fig. 3A, and two representative subject-level PFs are shown in Fig. 3B (best fit) and Fig. 3C (worst fit). Subject-level 95% HDI widths (Fig. 3D) suggest that **T** measures were estimated with less precision than either **A** or **AT** measures.

Subsequent analyses performed on subject-level JNDs (Table 1, Fig. 3E) revealed that temporal discriminability significantly differed among the three modalities ($F(2, 40) = 42.13$, $p < 0.001$, $\eta_p^2 = 0.68$, BF₁₀ > 100). Contrasts showed that this effect was driven by higher **T** JNDs compared to **A** and **AT** JNDs (**T** vs. **A**: $F(1, 40) = 53.37$, $p < 0.001$, $\eta_p^2 = 0.57$, BF₁₀ > 100; **T** vs. **AT**: $F(1, 40) = 71.67$, $p < 0.001$, $\eta_p^2 = 0.64$, BF₁₀ > 100). **A** and **AT** discriminability did not significantly differ ($F(1, 40) = 1.35$, $p = 0.253$, $\eta_p^2 = 0.03$, BF₁₀ = 0.24).

The analysis of subject-level PSEs showed a significant effect of modality, $F(2, 40)$

Figure 3. Experiment One: Psychometric model fits. (A) Group-level psychometric functions fit via Bayesian estimation for the three modalities show that subjects were least sensitive to rhythmic deviancy in the **T** condition (red). (B, C) Representative subject-level PFs showing good and poor model fit, respectively. (D) Boxplots of the 95% HDI widths for the subject-level PFs reveal that the worst goodness-of-fits were from the **T** condition, compared to the **A** and **AT** conditions. (E) Boxplots of subject-level JND estimates for the three modalities show that **T** JNDs were larger than **AT** JNDs ($p < 0.001$, $BF_{10} > 100$) and **A** JNDs ($p < 0.001$, $BF_{10} > 100$). **AT** JNDs were not different from **A** JNDs ($p = 0.253$, $BF_{10} = 0.24$). In each boxplot, the vertical line inside each box signifies that distribution's median.

266 = 17.55, $p < 0.001$, $\eta_p^2 = 0.47$, $BF_{10} > 100$ (Table 1). All pairwise comparisons were
 267 statistically significant (**A** vs. **T**: $F(1, 40) = 35.04$, $p < 0.001$, $\eta_p^2 = 0.47$, $BF_{10} > 100$; **A** vs.
 268 **AT**: $F(1, 40) = 9.98$, $p = 0.003$, $\eta_p^2 = 0.20$, $BF_{10} = 14.45$; **T** vs. **AT**: $F(1, 40) = 7.62$, $p =$
 269 0.009 , $\eta_p^2 = 0.16$, $BF_{10} = 5.76$). We also tested how veridical each PSE was with respect

to the 400 ms standard interval duration (Table 1). The mean PSE with **A** stimuli showed the highest fidelity, followed by closely by the mean PSE for **AT** stimuli. The least veridical PSEs were found in the **T** condition.

As explained above, we used Equation 1 to examine the separate, unimodal contributions to bimodal temporal sensitivity. Results showed that **AT** slopes were best predicted from individual **A** and **T** slopes using a MAX metric, mean SSE < 0.001. The MAX rule outperformed both linear summation, $m=1$ (SSE = 0.915), and a previously reported metric of audio-visual cue combination, $m = 2.56$ (SSE = 0.473). This suggests that one of the two unimodal inputs strongly dominated bimodal sensitivity, and the similarity between the **A** and **AT** JNDs points to **A** signals as the dominant input to the bimodal condition.

Exploratory analysis

An additional analysis of Experiment One's data uncovered a result that, despite its exploratory nature, deserves to be reported here. This exploratory analysis began with an observation about Experiment One results with bimodal stimulation. Although the mean JND did not differ between **AT** and **A** conditions, the *variability* in subject-level JND estimates was reduced five-fold in the bimodal condition compared to the unimodal auditory condition (Fig. 3E, Table 1). A Bayesian *t*-test showed that these standard deviations were reliably different, 95% HDI = [2.17, 5.88], suggesting that inter-individual differences in sensory discrimination were minimized in the bimodal condition. Notably, Gamache and Grondin (2010) found a similar result: reduced variance in temporal reproductions of intervals demarcated by audio-visual stimuli compared to intervals demarcated by visual cues only; Murray, Thelen, Ionta, and Wallace (2019) observed an analogous effect in response times to audio-visual stimuli.

We asked what might account for this finding. Møller et al. (2018) suggest that, in accordance with the inverse effectiveness principle, individual differences in unimodal sensitivity mediate the differential effects of bimodality across subjects. This allows for

reduced bimodal inter-individual differences despite large unimodal inter-individual differences. In particular, Møller et al. propose that subjects with lower unimodal sensitivity benefit more from bimodality than subjects with higher unimodal sensitivity. We tested this hypothesis by correlating each subject's unimodal discriminability (JND_A or JND_T) with the improved discriminability they showed with the *other* unimodal component added. We represent the bimodal discrimination benefit of adding **T** stimuli to unimodal **A** stimuli ($JND_A - JND_{AT}$) in Fig. 4A, and we represent the bimodal discriminability benefit of adding **A** stimuli to unimodal **T** cues ($JND_T - JND_{AT}$) in Fig. 4B. In Fig. 4, data points in the green region represent subjects who showed *improvements* in bimodal sensitivity relative to their discriminability in one unimodal condition (shown along the abscissa); data points in the red region represent subjects whose bimodal sensitivity was *impaired* relative to their discriminability of that unimodal condition.

Figure 4 shows two important results. First, the amount of multisensory gain increases linearly with unimodal JND in both the **A** and **T** conditions, following from the principle of inverse effectiveness. Poorer unimodal discriminability, characterized by larger **A** or **T** JNDs, is associated with higher discriminability gain from combining both forms of stimulation (supplementing **A** cues with **T** signals: $F(1, 19) = 605.0$, $p < 0.001$, $R^2 = 0.968$, Fig. 4A; supplementing **T** cues with **A** signals: $F(1, 19) = 703.7$, $p < 0.001$, $R^2 = 0.972$, Fig. 4B). Second, it is worthwhile to note that *all* subjects benefited from adding **A** to **T** stimuli, but only subjects with poorer **A** discriminability benefited when **T** stimuli were added to **A** stimuli. The performance of subjects with exceptional **A** discriminability actually lost ground from the addition of **T** stimulus in the bimodal condition (red region of Fig. 4A).

Note that although **T** signals clearly affected performance by reducing variance in estimates of bimodal discriminability and reducing bimodal PSEs (Table 1), the **T** modality's variable effects on participants and overall weaker temporal acuity allowed **A**

Figure 4. Individual subjects' benefit in discriminability with bimodal stimulation as a function of their unimodal discriminability. We define "bimodal discrimination benefit" as the difference between a subject's unimodal and bimodal JNDs, where this difference measure reflects the effect of information added from the *other* modality. (Panel A) Individual subjects' unimodal auditory sensitivity (x-axis) predicts their bimodal discrimination benefit from the addition of **T** stimuli (y-axis), $p < 0.001$, $R^2 = 0.968$. Only subjects with poor **A** discriminability (higher JNDs) benefited from the addition of **T** stimuli (green region), whereas subjects with better **A** discriminability (lower JNDs) detrimented from the addition of **T** stimuli (red region). (Panel B) Individual subjects' unimodal **T** discriminability (x-axis) predicts that subject's bimodal discrimination benefit from the addition of **A** stimuli as well (y-axis), $p < 0.001$, $R^2 = 0.972$. All subjects benefited from the addition of **A** stimuli to **T** signals (green region).

inputs to disproportionately influence estimates of bimodal sensitivity. Indeed, the bimodal discrimination benefit across all subjects conferred by **T** signals, as shown along the ordinate in Fig. 4A, was not significantly different from zero, $t(20) = 1.41$, $p = 0.175$, $d = 0.31$, $BF_{10} = 0.54$. In comparison, adding **A** stimuli to **T** signals *consistently* aided discrimination (green region of Fig. 4B), further confirming auditory dominance, $t(20) = 8.21$, $p < 0.001$, $d = 1.79$, $BF_{10} > 100$.

Discussion

Our analysis of JNDs showed better discrimination of duration when intervals were demarcated by **A** pulses than **T** pulses. Additionally, analysis of PSEs showed that comparison intervals between **A** pulses were perceived with greater fidelity to the standard than those between **T** pulses. Both results are consistent with previous research comparing judgments with unimodal auditory and tactile durations (Azari et al., 2020; Grondin & Rousseau, 1991; Jones et al., 2009; Lauzon et al., 2016; Mayer et al., 2014). Importantly, the bimodal (**AT**) test condition assessed directly how the result of concurrent stimulation with both modalities compared to the result with either modality alone. Bimodal discrimination (measured with JND values) closely approximated auditory-only discrimination (Fig. 3D), but PSEs in the **AT** condition were significantly shorter than those with **A** stimuli. This divergence between the two psychophysical measures suggest that the addition of a tactile pulse to an auditory pulse affected perceived duration but not sensitivity to changes in duration (Fig. 4A).

Our finding of strong but somewhat incomplete auditory dominance when **A** and **T** stimuli were combined reflects the complexities of multisensory combination. Often, discriminability with bimodal stimuli surpasses what would have been achieved with either modality alone (e.g. Diederich & Colonius, 2004; Hershenon, 1962; Meredith & Stein, 1986; Murray et al., 2019). However, if one unimodal input is less reliable than the other, their bimodal combination is dominated by the more reliable of the unimodal inputs (Ernst & Banks, 2002; Burr et al., 2009; Repp & Penel, 2002; Walker & Scott, 1981). Our

results are consistent with this latter outcome: as unimodal auditory processing was considerably more precise than tactile processing, the dominant input in determining bimodal discriminability was the auditory one. Finally, to complete the picture, it should be noted that under some conditions, individual subjects' responses to bimodal combinations suggest differences in the weights they place upon the separate components in multisensory combination (Zhou, Wong, & Sekuler, 2007), a fact further supported by the exploratory analyses whose results are shown in Fig. 4.

We believe that our exploratory analysis offers a useful lesson for developers of various human-interface devices. Incorporating bimodal stimuli can reduce individual differences in performance by compensating for subject-level differences in processing unimodal stimuli: bimodal stimulus presentation reduced inter-individual variability in JND (Figure 3E), supporting the results from Gamache and Grondin (2010) and Murray et al. (2019). For technology developers using time-locked stimuli, this means that supplementing unimodal signals with information from another modality could help homogenize different users' ability to exploit the information from the signals. However, the measurements we made are merely snapshots of individual differences in bimodal combination, and are mute as to whether those striking differences are stable over time, whether they extend to other modality combinations or tasks (Murray et al., 2019), and how they might relate to other, well-established predictors of bimodal combination (Stevenson, Zemtsov, & Wallace, 2012). We believe that follow-up research could profitably explore individual differences in multisensory processing.

Experiment Two

In Experiment One, subjects judged the duration of a trial's final interval relative to its isochronous predecessors. To make these comparisons, subjects must have formed some internal representation of the duration of the standard interval (*i.e.*, 400 ms). We cannot tell, though, what information subjects actually used to construct that representation. For example, in every trial, did subjects incorporate each and every one

of the preceding intervals into an internal representation against which to compare the trial's final interval?

The “multiple-look model” proposed by Drake and Botte (1993) hypothesizes that repeated exposure to a stimulus promotes formation of a memory trace, or perceptual memory, of that stimulus. With additional “looks”, the memory trace strengthens and sharpens, presumably up to the precision limit of the system. While there is some evidence that increasing the number of repetitions of a standard stimulus aids temporal sensitivity (Grondin, 2001; Grondin & McAuley, 2009; Miller & McAuley, 2005; Ten Hoopen, Van Den Berg, Memelink, Bocanegra, & Boon, 2011), several studies have failed to show a “multiple-look effect” (Grondin & McAuley, 2009; Miller & McAuley, 2005; Ten Hoopen et al., 2011). This failure could be explained by subjects’ ability to draw on some robust memory trace that developed early in the experiment. Such a memory trace would render the standard intervals on subsequent trials superfluous (*e.g.*, Morgan, Watamaniuk, & McKee, 2000; Nachmias, 2006; Dyjas, Bausenhardt, & Ulrich, 2012; Zimmermann & Cicchini, 2020).

These findings prompted us to conduct a second experiment in which we varied the number of standard intervals that preceded the final comparison interval. If subjects based their judgments on each and every one of the standard intervals presented on a trial, performance should improve with added standard interval presentations. Alternatively, if subjects developed and based their judgments on some robust, persistent memory trace of the isochronous interval, then the number of standard intervals presented before the final comparison interval should be inconsequential.

Methods

Except where explicitly noted, the methods were the same as in Experiment One.

Subjects. Fifteen subjects (8 female, mean age = 19.3 years, SD = 1.3 years) participated in this experiment; none had participated in Experiment One. No subject reported issues related to processing auditory or vibrotactile signals. However, data from

two subjects were excluded from analysis: one subject withdrew before completing the experiment, and data from another subject were lost because of a computer error. To confirm the sufficiency of our sample size, we conducted an *a priori* power analysis for independent variable following the procedures used in Experiment One. We found that, with 85% power, we would need $n=12$ subjects to observe an effect of modality similar to that reported by Jones et al. (2009) and $n=6$ subjects to observe a similar effect of standard intervals to Grondin and McAuley (2009), so our sample size of $n=13$ was adequate. All experimental procedures were approved by Brandeis University's Institutional Review Board and were conducted in accordance with the Declaration of Helsinki. All subjects provided written informed consent prior to participation. Subjects were awarded credits towards a university course requirement for their time.

Apparatus and stimuli. With the exception of some additional sound-absorbing insulation inserted beneath the vibrotactor, the apparatus was unchanged from Experiment One. Note that this extra sound-absorbing material did not affect the intensity of **T** stimulation. Additionally, the strong similarity of **AT** and **A** performance in Experiment One led us to cull the bimodal condition, testing performance with just the two unimodal stimuli.

Brownian masking noise calibration. Although we had no reason to suspect that subjects could hear the **T** stimuli in Experiment One (since our analyses revealed significant differences between the **A** and **T** JNDs and PSEs), we made our calibration procedures more rigorous for Experiment Two. We used a two-alternative forced choice task to find an optimal level of Brownian noise needed to mask the auditory sounds of the vibrotactor for each subject. Subjects were told that the vibrotactor would emit an auditory sound when activated, and were then instructed to listen for any such sound on the calibration trials. Subjects removed their hand from the vibrotactor before the calibration started so they could not feel any of the vibrations. Each trial consisted of two components presented in random order. During one component, the vibrotactor was

actively vibrating, while in the other component, it was inactive. Subjects responded by indicating which trial component, the first or the second, contained the sound. Accuracy $\sim 50\%$ constituted chance performance, and was taken as an indication that the subject could not hear the sounds produced by the vibrotactor. Each block contained 24 trials. Until subject accuracy fell in the middle 90% of a binomial distribution, the amplitude of the masking noise was increased and another block of calibration trials was presented.

Task. We modified Experiment One's task by varying the number of standard intervals, hereafter "nSTD", delivered on each trial. Across different sets of trials, nSTD was block randomized from 1STD (one standard interval preceding the comparison interval) to 4STD (four standard intervals preceding the comparison interval). All 48 trials in a block presented stimuli from the same modality and had the same nSTD, but rhythmic deviancy was randomized from trial to trial. Before each each block of experimental trials, subjects completed six practice trials with that block's modality and nSTD.

The complete set of 16 blocks of trials took ~ 1.5 hours to complete. In total, for each modality and nSTD, each subject received 12 trials with rhythmic deviancies drawn randomly and equally from each of the eight bins (four "early" and four "late"). Each subject completed 768 total trials (2 [modalities] $\times 4$ [nSTD] $\times 8$ [rhythmic deviancy bins] $\times 12$ [trials per condition]). As in Experiment One, subjects were allowed short, self-timed breaks between blocks. Subjects could earn additional course credits with accuracy scores $\geq 80\%$ on 8 or more of the 16 blocks.

Data curation. Except where noted, we used the same analytic procedures as in Experiment One. Subjects failed to respond before the 1000 ms maximum response time on just 0.9% of all trials; these trials were excluded from analysis. One-way ANOVAs and t -tests showed that trials with missing responses were not disproportionately represented for modality ($p = 0.293$), nSTD ($p = 0.769$), response key mapping ($p = 0.520$), or response category ("early" vs. "late", $p > 0.999$) conditions.

Psychometric modeling. The group- and subject-level models each contained eight conditions ($2 [\text{modalities}] \times 4 [\text{nSTD}]$). From the psychometric functions, we derived values of JND and PSE as in Experiment One. We used two-way, within-subjects ANOVAs to test for additive and interactive effects of modality and nSTD on subject-level JND and PSE estimates.

Results

Table 2

Mean JND for each modality and nSTD value. Standard deviations are shown in parentheses.

#Std	JND Mean (SD)		PSE Mean (SD)	
	Auditory	Vibrotactile	Auditory	Vibrotactile
1	20.29 (5.32)	29.47 (16.12)	399.46 (11.25)	406.01 (8.28)
2	20.91 (7.57)	34.75 (15.23)	402.59 (2.35)	404.29 (11.13)
3	20.04 (9.86)	44.96 (28.19)	400.83 (8.20)	398.40 (5.17)
4	21.71 (9.11)	28.75 (8.45)	403.38 (5.03)	396.27 (12.86)

Table 3

Results of two-way [modality $\times$ nSTD] within-subjects ANOVAs on JNDs and PSEs.

Effect	JND					PSE				
	df	F	p	η_p^2	BF ₁₀	df	F	p	η_p^2	BF ₁₀
Modality	1, 12	26.66	<0.001	0.69	>100	1, 12	0.06	0.808	0.005	0.22
#Std	1.78, 21.38 [†]	2.90	0.082 [†]	0.19	1.05	3, 36	1.65	0.194	0.121	0.35
Modality $\times$ #Std	1.84, 22.07 [†]	5.61	0.012 [†]	0.32	4.62	3, 36	4.15	0.013	0.260	8.92

Note: [†]Corrected df and p. BF₁₀ = Bayes factor; df = degrees of freedom.

Figure 5 shows the model-fitting results. Table 2 lists the mean JND and PSE for each experimental condition. Table 3 reports full omnibus ANOVA results for both JND

Figure 5. Experiment Two: Psychometric model fits. The psychometric functions (PFs), high-density intervals (HDIs), and just noticeable difference values (JNDs) comparing **A** and **T** sensitivity (panels A-C) were estimated by averaging across all nSTD, and *vice versa* (panels D-F). (A) Group-level PFs fit via Bayesian estimation show that, like Experiment One, **T** temporal sensitivity was poorer than **A** temporal sensitivity. (B) Boxplots of 95% HDI widths for subject-level PSEs show that model fits were poorer for the **T** condition than for the **A** condition. (C) Boxplots of JNDs derived from subject-level models show that **A** JNDs were smaller than **T** JNDs. (D, E) Temporal sensitivity across all of the nSTD conditions were similar, when averaging **A** and **T** performance. (F) No significant differences existed among the JNDs from the four nSTD conditions.

and PSE; only p -values for these effects are reported in the following sections.

Analysis of JNDs. Replicating a main finding from Experiment One, an omnibus ANOVA showed that JNDs were larger for **T** stimuli than for **A** stimuli, $p < 0.001$ (Fig. 5C). Averaging over both modalities, nSTD had no significant effect on JND, $p = 0.082$

(Fig. 5F); however, the interaction between Modality and nSTD was significant, corrected $p = 0.012$. To dissect the ANOVA's significant interaction, we separated JNDs by modality and performed two separate one-way ANOVAs with polynomial contrasts, testing for linear and quadratic trends across the nSTD conditions. Neither linear nor quadratic components were statistically significant in the **A** ANOVA, confirming that JNDs with **A** stimuli were invariant as nSTD increased ($F(1, 36) = 7.4$, $p = 0.547$, $\eta_p^2 = 0.01$; and $F(1, 36) = 3.6$, $p = 0.674$, $\eta_p^2 = 0.005$, respectively). For **T** stimuli, the picture was a bit more complex. First, a non-significant linear component suggested that there was no monotonic change in **T** JNDs with nSTD, $F(1, 36) = 0.26$, $p = 0.616$, $\eta_p^2 = 0.007$. However, there was a significant quadratic component, reflecting the non-monotonicity that is shown in Table 2, $F(1, 36) = 9.14$, $p = 0.005$, $\eta_p^2 = 0.20$.

Analysis of PSEs. Averaging across levels of nSTD, an omnibus ANOVA showed that PSEs did not significantly differ between modalities, $p = 0.808$. Also, collapsing across modalities, variation in nSTD did not affect PSEs, $p = 0.194$. However, there was significant interaction between nSTD and modality, $p = 0.013$. We again examined this significant interaction by separating PSEs by modality and performing two separate one-way ANOVAs with linear and quadratic contrasts. PSEs for **A** stimuli were constant over values of nSTD, with non-significant linear and quadratic components ($F(1, 36) = 1.65$, $p = 0.207$, $\eta_p^2 = 0.04$; and $F(1, 36) = 0.03$, $p = 0.869$, $\eta_p^2 < 0.01$, respectively). However, the decrease in **T** PSEs as nSTD increased was significantly linear, $F(1, 36) = 10.71$, $p = 0.002$, $\eta_p^2 = 0.23$. The quadratic trend was not significant, $F(1, 36) = 0.007$, $p = 0.932$, $\eta_p^2 < 0.001$.

We performed additional post hoc tests to compare the PSE in each condition (see Table 2) to the value of the standard interval, 400 ms. In three of the eight conditions, subjects' mean PSE significantly differed from the standard: **A**-2STD, **A**-4STD, and **T**-1STD ($\text{BF}_{10} = 23.99$, 95% HDI = [1.15, 3.95]; $\text{BF}_{10} = 2.23$, 95% HDI = [0.09, 6.66]; and $\text{BF}_{10} = 3.01$, 95% HDI = [1.03, 11.64], respectively).

Discussion

Overall, Experiment Two showed that increases in nSTD brought little or no systematic benefit to temporal sensitivity as measured by JNDs. Specifically, JND values for **A** stimuli remained stable despite variation in nSTD, and JNDs for **T** stimuli varied some, but did so non-monotonically (Table 2). So, neither modality exhibited what might be expected had there been a multiple-look effect (Drake & Botte, 1993). At first glance, this null result seemed surprising; after all, comparable studies showed performance improvements with additional presentations of a standard stimulus (*e.g.*, Drake & Botte, 1993; Grondin, 2001; Grondin & McAuley, 2009; Miller & McAuley, 2005; Ten Hoopen et al., 2011). The explanation for the divergence between our and others' results may lie in experimental design. Specifically, multiple-look effects are more likely when the number of comparison intervals varies, rather than the number of standard intervals (Grondin & McAuley, 2009; Miller & McAuley, 2005; Ten Hoopen et al., 2011). The number of standard intervals becomes a more critical factor when the duration of the standard interval varies (Miller & McAuley, 2005), probably because a fixed standard interval duration promotes the formation of robust memory traces that renders additional standard presentations on each trial superfluous (Morgan et al., 2000; Nachmias, 2006; Dyjas et al., 2012; Zimmermann & Cicchini, 2020). Such speculation about the role of memory raises but does not answer a number of questions, such as the rate at which such a memory forms, its robustness (Tong, Dubé, & Sekuler, 2019), and whether, once formed, memories for multiple stimuli can co-exist as demonstrated in psychophysical experiments with standard stimuli very different from those used here (Morgan et al., 2000).

Although the analysis of JNDs in Experiment Two echoed Experiment One's modality difference, the analysis of PSEs did not: there was no main effect of modality on PSE. For a more apples-to-apples comparison with Experiment One, we evaluated differences between modalities in Experiment Two within the 3STD condition specifically.

The analysis showed no significant difference between the modalities, $t(12) = 1.16$, $p = 0.269$, $d = 0.32$, $BF_{10} = 0.49$. While PSE values did not differ *between* modalities, we found differences in how PSE varied *within* each modality as nSTD varied. The linear reduction in **T** PSE with increasing nSTD suggests that the comparison interval seemed longer when preceded by more standards. No such effect was present in the **A** condition.

Both JND and PSE analyses echo the results of other studies conducted using auditory stimuli (*e.g.*, Grondin & McAuley, 2009; Miller & McAuley, 2005; Ten Hoopen et al., 2011). Importantly, however, our data also extend to the vibrotactile domain, showing that increasing standard intervals in this modality did not consistently affect temporal sensitivity but did have an impact on perceived duration.

Summary and conclusions

The first experiment compared auditory, vibrotactile, and bimodal temporal discriminability using a task in which subjects compared the duration of a variable interval with the duration of three preceding standard intervals. JND and PSE estimates for each modality were derived and compared using Hierarchical Bayesian psychometric fits, and the unimodal contributions to bimodal processing were explored by modeling how stimuli from different modalities were combined, and by examining the degree to which individual subjects benefited from bimodal stimulation. Our second experiment also compared auditory and vibrotactile temporal discriminability, as well as the effect of ranging the number of standard intervals from one to four. These data were also interpreted using psychometric fits, JND estimates, and PSE estimates.

Overall, temporal sensitivity measured with JND was lower for vibrotactile than for auditory signals (Experiments One and Two), and bimodal audio-tactile processing was dominated by auditory cues (Experiment One). Following from the principle of inverse effectiveness (Meredith & Stein, 1983; Ross, Saint-Amour, Leavitt, Javitt, & Foxe, 2007), inter-individual differences in discriminability were significantly reduced in the bimodal condition, suggesting that developers of human-interface technology would be

well-advised to utilize bimodal stimuli for a more uniform user experience (Experiment One). Moreover, the number of standard intervals presented per trial had no consistent effect on temporal discriminability in either modality, pointing to the likely contribution of memory for the non-roving standard duration (Experiment Two).

Combined with results from another study from our lab (Villalonga, Sussman, & Sekuler, 2020), our findings suggest that different modalities occupy different levels in a hierarchy of temporal acuities, with audition at the top, followed by vibrotactile sensation and then vision. This general framework is consistent with a comparable ordering of Weber fractions reported by Azari et al. (2020): smallest for auditory stimuli (0.121) vibrotactile stimuli at an intermediate level (0.170), and visual stimuli slightly higher (0.186). However, attempts to draw parallels between multiple sets of results need to be tempered by a recognition of differences among tasks and among the particular exemplar stimuli used to represent various modalities.

Open Practices Statement

Datasets generated during and analyzed during the current study are available in the Open Science Framework (OSF) repository, found at https://osf.io/q8y6j/?view_only=76df6923d10a498b84ec44883a75b525. Neither of the experiments was preregistered.

Acknowledgments

MBV was supported by the NIGMS training grant T32 GM132498. RFS was supported by the NIH grant R90 DA033463 for Undergraduate and Graduate Training in Computational Neuroscience. Adam Fleishaker and Jacob Smith provided coding support. The content is solely the responsibility of the authors and does not represent the official views of any funding agency. The authors declare they have no conflicts of interest.

References

- Agam, Y., & Sekuler, R. (2007). Interactions between working memory and visual perception: An ERP/EEG study. *NeuroImage*, 36(3), 933-942.
- Azari, L., Mioni, G., Rousseau, R., & Grondin, S. (2020). An analysis of the processing of intramodal and intermodal time intervals. *Attention, Perception, & Psychophysics*, 82, 1473-1487.
- Bausenhardt, K. M., Di Luca, M., & Ulrich, R. (2018). Assessing duration discrimination: Psychophysical methods and psychometric function analysis. In *Timing and Time Perception: Procedures, Measures, & Applications* (pp. 52–78). Brill.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433-436.
- Burr, D., Banks, M. S., & Morrone, M. C. (2009). Auditory dominance over vision in the perception of interval duration. *Experimental Brain Research*, 198, 49-57.
- Caldwell, A., & Lakens, D. (2019). *Power Analysis with Superpower*. Retrieved from <https://arcaldwell149.github.io/SuperpowerBook> (Self Published on GitHub)
- Calin-Jageman, R. J., & Cumming, G. (2019). Estimation for better inference in neuroscience. *eNeuro*, 6(4).
- Cavanagh, P. (2011). Visual cognition. *Vision Research*, 51(13), 1538-1551. doi: 10.1016/j.visres.2011.01.015
- Diederich, A., & Colonius, H. (2004). Bimodal and trimodal multisensory enhancement: Effects of stimulus onset and intensity on reaction time. *Perception & Psychophysics*, 66(8), 1388–1404.
- Drake, C., & Botte, M. C. (1993). Tempo sensitivity in auditory sequences: Evidence for a multiple-look model. *Perception & Psychophysics*, 54(3), 277-286.
- Dyjas, O., Bausenhardt, K. M., & Ulrich, R. (2012). Trial-by-trial updating of an internal reference in discrimination tasks: Evidence from effects of stimulus order and trial sequence. *Attention, Perception, & Psychophysics*, 74, 1819-1841.
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a

- statistically optimal fashion. *Nature*, 415(6870), 429-433. doi: 10.1038/415429a
- Gamache, P., & Grondin, S. (2010). Sensory-specific clock components and memory mechanisms: investigation with parallel timing. *European Journal of Neuroscience*, 31(10), 1908-1914.
- Grondin, S. (2001). Discriminating time intervals presented in sequences marked by visual signals. *Perception & Psychophysics*, 63(7), 1214–1228.
- Grondin, S., & McAuley, J. (2009). Duration discrimination in crossmodal sequences. *Perception*, 38, 1542-1559.
- Grondin, S., & Rousseau, R. (1991). Judging the relative duration of multimodal short empty time intervals. *Perception & Psychophysics*, 49, 245-256.
- Hershenson, M. (1962). Reaction time as a measure of intersensory facilitation. *Journal of Experimental Psychology*, 63(3), 289.
- Jeffreys, H. (1961). *Theory of Probability*. Oxford, UK: Oxford University Press.
- Jones, L., Poliakoff, E., & Wells, J. (2009). Good vibrations: Human interval timing in the vibrotactile modality. *The Quarterly Journal of Experimental Psychology*, 62, 2171-2186.
- Kruschke, J. K. (2013). Bayesian estimation supersedes the t test. *Journal of Experimental Psychology: General*, 142(2), 573.
- Lauzon, A. P., Russo, F. A., & Harris, L. R. (2016). There's something about that groove: Rhythm improves detection of audio but not vibrotactile asynchronies. *Canadian Acoustics*, 44(3).
- Lumaca, M., Trusbak Haumann, N., Brattico, E., Grube, M., & Vuust, P. (2019). Weighting of neural prediction error by rhythmic complexity: A predictive coding account using mismatch negativity. *The European Journal of Neuroscience*, 49(12), 1597-1609.
- Makous, J. C., Friedman, R. M., & Vierck, C. J. (1995). A critical band filter in touch. *Journal of Neuroscience*, 15(4), 2808–2818.

MATLAB. (2017). *version 9.3.0 (r2017b)*. Natick, Massachusetts: The MathWorks Inc.

Mayer, K. M., Di Luca, M., & Ernst, M. O. (2014). Duration perception in crossmodally-defined intervals. *Acta Psychologica*, 147, 2–9.

Meng, F., & Spence, C. (2015). Tactile warning signals for in-vehicle systems. *Accident Analysis and Prevention*, 75, 333-346.

Meredith, M. A., & Stein, B. E. (1983). Interactions among converging sensory inputs in the superior colliculus. *Science*, 221(4608), 389-391. doi: 10.1126/science.6867718

Meredith, M. A., & Stein, B. E. (1986). Spatial factors determine the activity of multisensory neurons in cat superior colliculus. *Brain Research*, 365(2), 350–354.

Miller, N. S., & McAuley, J. (2005). Tempo sensitivity in isochronous tone sequences: The multiple-look model revisited. *Perception & Psychophysics*, 67, 1150-1160.

Møller, C., Højlund, A., Bærentsen, K. B., Hansen, N. C., Skewes, J. C., & Vuust, P. (2018). Visually induced gains in pitch discrimination: Linking audio-visual processing with auditory abilities. *Attention, Perception, & Psychophysics*, 80(4), 999–1010.

Morey, R. D., & Rouder, J. N. (2011). Bayes factor approaches for testing interval null hypotheses. *Psychological Methods*, 16(4), 406–419.

Morgan, M. J., Watamaniuk, S. N. J., & McKee, S. P. (2000). The use of an implicit standard for measuring discrimination thresholds. *Vision Research*, 40, 2341-2349.

Mountcastle, V. B., LaMotte, R. H., & Carli, G. (1972). Detection thresholds for stimuli in humans and monkeys: comparison with threshold events in mechanoreceptive afferent nerve fibers innervating the monkey hand. *Journal of Neurophysiology*, 35(1), 122–136.

Murray, M. M., Thelen, A., Ionta, S., & Wallace, M. T. (2019). Contributions of intraindividual and interindividual differences to multisensory processes. *Journal of*

Cognitive Neuroscience, 31(3), 360–376.

Nachmias, J. (2006). The role of virtual standards in visual discrimination. *Vision Research*, 46(15), 2456-2464.

Nagarajan, S. S., Blake, D. T., Wright, B. A., Byl, N., & Merzenich, M. M. (1998).

Practice-related improvements in somatosensory interval discrimination are temporally specific but generalize across skin location, hemisphere, and modality. *Journal of Neuroscience*, 18(4), 1559–1570.

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing* (Vol. 124, pp. 1–10).

Prins, N., & Kingdom, F. A. A. (2018). Applying the model-comparison approach to test specific research hypotheses in psychophysical research using the Palamedes toolbox. *Frontiers in Psychology*, 9:1250. doi: 10.3389/fpsyg.2018.01250

Repp, B. H., & Penel, A. (2002). Auditory dominance in temporal processing: new evidence from synchronization with simultaneous visual and auditory sequences. *Journal of Experimental Psychology: Human Perception and Performance*, 28, 1085-1099.

Ross, L. A., Saint-Amour, D., Leavitt, V. M., Javitt, D. C., & Foxe, J. J. (2007). Do you see what I am saying? Exploring visual enhancement of speech comprehension in noisy environments. *Cerebral Cortex*, 17(5), 1147-1153. doi: 10.1093/cercor/bhl024

Schönbrodt, F. D., & Wagenmakers, E.-J. (2018). Bayes factor design analysis: Planning for compelling evidence. *Psychonomic Bulletin & Review*, 25(1), 128–142. Retrieved from <https://doi.org/10.3758/s13423-017-1230-y> doi: 10.3758/s13423-017-1230-y

Stevenson, R. A., Zemtsov, R. K., & Wallace, M. T. (2012). Individual differences in the multisensory temporal binding window predict susceptibility to audiovisual illusions.

685 *Journal of Experimental Psychology: Human Perception and Performance*, 38(6),
686 1517.

687 Ten Hoopen, G., Van Den Berg, S., Memelink, J., Bocanegra, B., & Boon, R. (2011).
688 Multiple-look effects on temporal discrimination within sound sequences. *Attention*,
689 *Perception, & Psychophysics*, 73(7), 2249–2269.

690 To, M. P. S., Baddeley, R. J., Troscianko, T., & Tolhurst, D. J. (2011). A general rule for
691 sensory cue summation: evidence from photographic, musical, phonetic and
692 cross-modal stimuli. *Proceedings of the Royal Society B: Biological Sciences*,
693 278(1710), 1365-1372.

694 Tong, K., Dubé, C., & Sekuler, R. (2019). What makes a prototype a prototype?
695 Averaging visual features in a sequence. *Attention, Perception & Psychophysics*,
696 81(6), 1962-1978. doi: 10.3758/s13414-019-01697-5

697 VanRullen, R., & Thorpe, S. J. (2001). The time course of visual processing: from early
698 perception to decision-making. *Journal of Cognitive Neuroscience*, 13(4), 454-461.
699 doi: 10.1162/08989290152001880

700 Villalonga, M. B., Sussman, R. F., & Sekuler, R. (2020). Feeling the beat (and seeing it,
701 too): Vibrotactile, visual, and bimodal rate discrimination. *Multisensory Research*,
702 33(1), 31–59. doi: 10.1163/22134808-20191413

703 von Békésy, G. (1959). Similarities between hearing and skin sensations. *Psychological*
704 *Review*, 66, 1-22.

705 Walker, J. T., & Scott, K. J. (1981). Auditory–visual conflicts in the perceived duration of
706 lights, tones, and gaps. *Journal of Experimental Psychology: Human Perception*
707 *and Performance*, 7(6), 1327.

708 Waskom, M. L., Okazawa, G., & Kiani, R. (2019). Designing and interpreting
709 psychophysical investigations of cognition. *Neuron*, 104(1), 100-112. doi:
710 10.1016/j.neuron.2019.09.016

711 Welch, R. B., & Duttonhurt, L. (1986). Contributions of audition and vision to temporal

rate perception. *Perception & Psychophysics*, 39, 294-300.

Welch, R. B., & Warren, D. H. (1980). Immediate perceptual response to intersensory discrepancy. *Psychological Bulletin*, 88(3), 638-667.

Woods, D. L., Alain, C., Covarrubias, D., & Zaidel, O. (1993). Frequency-related differences in the speed of human auditory processing. *Hearing Research*, 66(1), 46–52.

Zhou, F., Wong, V., & Sekuler, R. (2007). Multi-sensory integration of spatio-temporal segmentation cues: one plus one does not always equal two. *Experimental Brain Research*, 180(4), 641-654. doi: 10.1007/s00221-007-0897-0

Zimmermann, E., & Cicchini, G. M. (2020). Temporal context affects interval timing at the perceptual level. *Scientific Reports*, 10(1), 8767. doi: 10.1038/s41598-020-65609-6