Threshold adaptation and its time course: An investigation of gradable adjectives

Ming Xiang, Chris Kennedy, Allison Kramer (University of Chicago)

It is well established that the interpretation of gradable adjectives is heavily context dependent, since the threshold of determining whether an object qualifies as having a particular adjectival property "X" varies across contexts. One source of variability is the speaker: different speakers may have different thresholds for the same adjective. For successful communication to happen, speakers need to dynamically align thresholds. This study is a first step towards experimentally investigating whether and how hearers adjust their adjective thresholds based on speaker input.

The procedure was adopted and modified from the speech adaptation studies in Vroomen et al. (2007) and Kleinschmidt and Jager (2015). A one-to-five scalar continuum was created for three adjective properties: a relative adjective "tall", a minimum adjective "bent", and a maximum adjective "plain" (see Figure 1A), which represent three different kinds of lexical scales that lead to distinct thresholds (Kennedy 2007). The experiment consisted of a pre-calibration, an exposure/testing and a post-calibration session. In the **pre-calibration session**, each participant was presented with all the images from each adjective scale multiple times, and made a binary judgment to the question "Is this tall/bent/plain?". The presentation of images followed a Gaussian distribution, with the image from scale position 3 presented most frequently for each adjective. For each participant, the most ambiguous scale point image for each adjective was individually determined, and was used for the next session. In the **exposure/testing session**, for each adjective, participants were exposed to a sequence of 24 repetitions of an auditory statement that describes a simultaneously visually presented image. At six different intermediate exposure trial positions (the 2^{th}, 4^{th}, 8^{th}, 13^{th}, 20^{th} and 24^{th}), participants received three test trials. The test trials always consist of the "most ambiguous" image individually chosen for each participant from his/her pre-calibration, and also an additional image from the scale position right above the most ambiguous image or right below it. Participants made a binary judgment as to whether each test image was "tall/bent/plain". Since the goal of the testing trials was to examine whether and how participants adapt and revise their own thresholds under the influence of the exposure trials, in a between-subject design, we manipulated four different kinds of exposure sequences, defined by the combination of the auditory statements and the images they describe (Figure 1B). In the Ambiguous.Positive and the Ambiguous.Negative conditions, the exposure image was the most ambiguous image, and the auditory description was "This is X" or "This is not X". In the Prototypical.Positive condition, the exposure image was from the highest #5 scale position, and the auditory description was "This is X". In the Prototypical.Negative condition, the exposure image was from the lowest #1 scale position, and the auditory description was "This is Not X". For each of the four exposure sequences, before the first auditory statement, the speaker set up a discourse scenario that specified a relevant conversational goal (e.g. "For a party, I need a tall candle" or "For a party, I need a candle that is not tall"). After completing all the exposure/testing trials, participants carried out a **post-calibration session** for each adjective, following the same procedure as the pre-calibration.

Results and discussion Thirty native speakers were recruited for each exposure condition (120 total MTurkers). Figure 2 shows that participants' thresholds were significantly influenced by the exposure sequences, measured by the change from pre-calibration to the post-calibration acceptance judgments. There are two main findings. First, the Ambiguous.Negative exposure made participants less likely to accept an image as having the property X, whereas the Ambiguous.Positive exposure made participants more likely to do so. However, the effect of the Negative and Positive statements with the Prototypical images is the opposite from the Ambiguous image exposure conditions. To explain this, we propose that participants initially maintain probabilistic distributions over the possible threshold values for each gradable adjective, and upon hearing "This is X" with an exposure image, they recalibrate the mean of the initial threshold distribution by shifting the mean towards the direction of the scale position indicated by the exposure trial image. Importantly, in addition to maintaining and adjusting the threshold distribution for a category X, they also simultaneously maintain and adjust the threshold distribution for a category "Not X". Upon observing a negative exposure trial "This is not X", they shift the mean of the "Not X" threshold distribution towards the exposure image as well. Since the two categories "X" and "Not X" are not completely independent from each other, shifting the mean of one distribution simultaneously leads to the shift of the other mean as well. Figure 3 demonstrates the process of threshold adaptation with three examples. Adaptation that involves shifting the mean of an original category boundary distribution has been found for categorical perception in speech (Kleinschmidt and Jager 2015) and quantifier interpretation (e.g. *some* and *many*, Yildirim et al., 2016). The current study extends this general adaptation strategy to adjective processing. In addition, our results also showed that even exposures to statements that the hearer would have no dispute about can still trigger readjustment behavior for the hearer (e.g. the two Prototypical exposure conditions).

The second finding in Figure 2 is that for absolute adjectives *plain* and *bent*, although the direction of the adaption is similar to the relative adjective *tall*, the adaptation effect is only present under some of the exposure conditions. We suggest that absolute and relative adjectives are sensitive to the same adaption strategy, but the end-point oriented lexical semantics of absolute adjectives leads to a very different initial threshold distribution for the participants, which in turn shapes the output of the adaptation strategy.

Finally, the testing trials during the exposure sequence were designed to examine the time course of the adaptation behavior. Analysis on these trials revealed that the adaptation effect emerged at the earliest point we tested (i.e. after the 2nd trial during the exposure sequence), and the effect size did not change at the later testing trials, suggesting that adaptation happens very early, and is insensitive to the frequency of the exposure trial (c.f. Kleinschmidt and Jager 2015 for categorical speech perception).

Conclusion Listeners actively and quickly revise their thresholds of gradable adjectives by re-estimating the mean of the original threshold distribution based on speaker input. The lexical semantics of different adjectives also modulates the output of the adaptation strategy.

Figure 1: **A**. (**Left)**: Image Stimuli on a 1-5 scale. B. (**Right)**: The four exposure sequences.

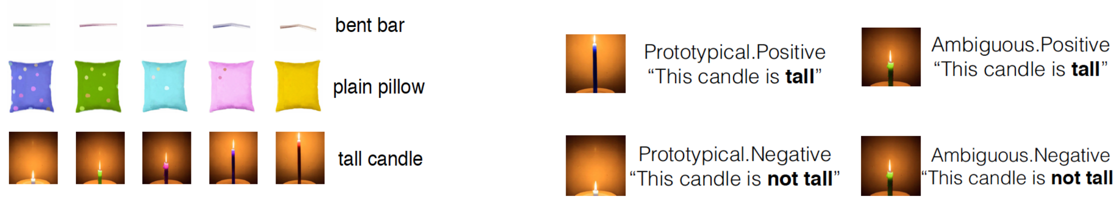

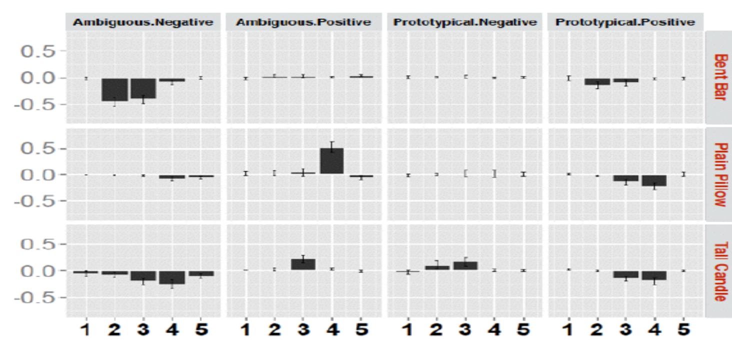

Figure 2: Post-calibration minus Pre-calibration difference scores for each adjective and exposure condition, and for each scale position. A positive score indicating increased acceptance of a token as having a property X in the post-calibration phase; and a negative score indicating decreased acceptance. X-axis plots the 1-5 scale positions**.**

Figure 3: Schematic adaption behavior for "tall" under three different exposure conditions. The x-axis of each plot represents degrees on a height scale. Solid line distributions represent hearers' original threshold distributions, and dashed line distributions represent the new distributions after adaptation.