Adaptation Aftereffects as a Result of Bayesian Categorization

Marina Dubova (marina.dubova.97@gmail.com)
Saint Petersburg State University, Department of Psychology, 6 Makarova Embankment
Saint Petersburg, 199034 Russia

Arseny Moskvichev (amoskvic@uci.edu)
University of California Irvine, 2277 Social
Behavioral Sciences Gateway Building
Irvine, CA 92697 USA

Abstract

We propose a unified explanation of contrastive and assimilative adaptation aftereffects from the perspective of higher-level cognitive processes: rational category learning and categorical perception. We replicate (twice) previously reported assimilative and contrastive effects (Uznadze illusion in visual modality), propose a rational computational model of the process, and evaluate our model performance against the Bayesian logistic regression baseline. We conclude by discussing theoretical implications of our study and directions for further research.

Keywords: adaptation aftereffects, perceptual biases, set illusion, Uznadze illusion, computational modeling, categorical perception

Introduction and Background

In many experimental settings, repeated exposure to stimuli affects the perception of subsequent ones. These phenomena are often referred to as the aftereffects of adaptation (Gibson & Radner, 1937). For example, if a participant is repeatedly presented with two circles, one bigger than another, she might perceive equal circles as being different during the test trial (Figure 1). Similar effects are manifest across a wide range of experimental conditions, in different modalities, and on different levels of abstraction. Behavioral studies demonstrate adaptation aftereffects in situations that run the gamut from simple shape and motion perception under brief presentation (Suzuki & Cavanagh, 1998; Chalk et al., 2010) to perception and recognition of faces, facial expressions, gender, and race (Webster & MacLeod, 2011; Leopold et al., 2001).

Contrastive and assimilative effects

It is possible to split all known adaptation aftereffects into two broad categories: **contrastive** and **assimilative** (Howard & Rogers, 1995). Contrastive aftereffects take place when the test stimulus seems more **different** from those seen during the adaptation phase (adaptors) than it would be perceived under normal conditions. Assimilative aftereffects, in turn, produce a reversed effect: the test stimulus is perceived as being more **similar** to adaptors. There is evidence that these two types of effects could occur in very similar and even identical experimental settings (Uznadze, 1958; Fritsche et al., 2017; Chopin & Mamassian, 2012). This raises a question: **what determines whether a contrastive or assimilative aftereffect will be present in a given trial**?

A broadly accepted view is that the probability of contrastive aftereffects occurrence grows with increasing difference between the test stimuli and the adaptors, increased length of adaptor presentation, as well as with the increase of overall stimuli salience and contrast (Howard & Rogers, 1995; Palumbo et al., 2017; Fritsche et al., 2017; Chopin & Mamassian, 2012).

Finding a mechanism that would explain the onset of both types of adaptation aftereffects turned out to be challenging. Previously dominant framework of adaptation as neural fatigue proved unsuccessful in accounting for the wide range of observed phenomena (Thompson & Burr, 2009). Recent studies predominantly focused on uncovering the mechanisms of a particular type of aftereffect: either contrastive (Webster & MacLeod, 2011; Rhodes & Jeffery, 2006; Grill-Spector et al., 2006; Stocker & Simoncelli, 2006; Chopin & Mamassian, 2012) or assimilative (Chalk et al., 2010; Palumbo et al., 2017).

There are, however, models that propose potential mechanisms of both contrastive and assimilative effects in visual (Wei & Stocker, 2015) and audial (Kleinschmidt & Jaeger, 2011) modalities. Wei and Stocker (2015) explained the opposite perceptual biases as a result of efficient coding constraints in a rational observer framework. Unfortunately, this model falls short in accounting for the influence of the difference between the test stimulus and the adaptors on illusion type (it predicts that this factor has no impact). At the same time, similar aftereffects in phonetic adaptation were modeled as Bayesian belief updating over two competing phonetic categories (Kleinschmidt & Jaeger, 2011). The limitation of this model is that it is designed for the task of forced choice between two categories that are given in advance. In most real-world and experimental adaptation scenarios, however, the alternative categories are implicit.

Overall, none of the existing models provide a complete account of the existing phenomena, which warrants further research in this direction.

Adaptation aftereffects and categorization

We propose a high-level interpretation of adaptation biases from a categorization standpoint. We argue that during the adaptation phase a person forms the categories of "typical" and "other" (atypical) stimuli. Learning is formalized using the ideal observer approach. The structure of the "typi-

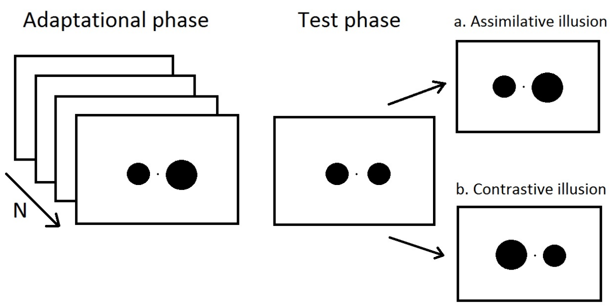

Figure 1: Uznadze visual illusion experimental procedure During the adaptation phase, a subject is repeatedly exposed to two circles, one being bigger than another. On the test trial, two equal circles are presented, and the subject responds whether they appear **the same (no illusion)** or, if not, which one appears bigger (**contrastive** or **assimilative** illusion).

cal" category is estimated from observed adaptors, while the "other" category is determined through its relationship to the already learned one. The main assumption is that an observer expects different visual categories to lie relatively far from each other in the feature space. On the test phase, the observer reconstructs the most likely true stimulus, given the learned category structure and the noisy sensory observation.

There is some evidence that provides conceptual support for our approach. First, in the domain of category learning, there is a notion of **categorical perception** which refers to phenomena whereby the same stimuli seem more different or similar, depending on whether they belong to the same or different categories in the learned conceptual structure (Goldstone, 1994, 1995; Goldstone & Hendrickson, 2010; Kuhl & Iverson, 1995). Second, in the domains of color and speech perception, perceptual bias toward the category prototype was formalized as an optimal statistical inference of real stimulus in high uncertainty conditions (Feldman et al., 2009; Cibelli et al., 2016). These perceptual shifts resemble the assimilative aftereffects. Third, a similar idea was successfully applied earlier in the domain of face perception: it was shown that the contrastive aftereffects are directed precisely toward the anti-prototype of the seen examples (Leopold et al., 2001, 2005; Rhodes & Jeffery, 2006). Assimilative aftereffects were not, however, considered in these studies.

Overall, there is evidence that category attribution plays an important role in perception. Although visual adaptation is most commonly viewed as a low-level process, current low-level models may not be able to fully capture the broad spectrum of visual adaptation aftereffects and their dynamics (Leopold et al., 2005) and are hardly compatible with interocular transfer of adaptation biases (Raymond, 1993). We believe that the difficulties encountered by low-level explanations, together with the successes of categorical perception models, warrant considering alternative, high-level explanations of perceptual aftereffects.

Our model builds upon the previous results and provides a simple and unified interpretation of both assimilative and contrastive aftereffects from a categorical perception standpoint.

To test our interpretation, we use a visual version of the Uznadze illusion (Figure 1). We replicate previously reported results on the association between the probabilities of opposite illusions with the length of the adaptation phase and the difference between the adaptation stimuli (Uznadze, 1958, 1966). After that, we evaluate the performance of our model on these data.

Experiment 1

This experiment replicated the findings reported in Uznadze (1958, 1966).

Hypothesis: Difference between the adaptors and the test stimulus, together with the number of adaptation trials, determine the probability of assimilative vs contrastive aftereffect occurrence. In particular, the assimilative aftereffect is associated with lower differences between stimuli sizes and smaller numbers of adaptation trials, while the contrastive aftereffect onset probabilities follow a reversed pattern.

Procedure

Pairs of circles of different sizes were presented as adaptors. We varied the magnitude of difference between adaptation stimuli (from 1 to 3 individual differential thresholds) and the number of adaptation trials (from 1 to 8) to evaluate their effect on the probabilities of assimilative and contrastive illusions. The procedure is illustrated in 1.

1. **Estimation of individual differential thresholds.** Two circles (diameters: left 2.5cm, right 2.5 or 3.0cm) were presented to participants. They were asked to focus on the dot in the center of the screen. We estimated participants' differential thresholds by the method of adjustment (Gescheider, 1997). That is, subjects saw two different circles and altered the size of one of them until the circles appeared equal to each other. In the second condition, the circles were initially the same and subjects made them different. We repeated this procedure six times and averaged the results to obtain the differential threshold estimate.

2. **Adaptation phase.** Subjects focused at a central dot on the screen, while they were exposed to two circles (for 150 ms) several (1-8) times. The difference in size between the two circles was 1, 2 or 3 individual differential thresholds.

3. **Test phase.** Participants saw two equal circles for 150ms and reported whether they appeared the same. If there was a perceived difference, participants identified which of the two appeared larger. They were instructed to respond as fast as possible and to rely only on their sensations. The test trial was repeated until the "same" relationship was reported 3 times in a row. This ensured that the aftereffect has faded before the start of the next trial. We did not analyze

the fading dynamics and only used the first test response in further analysis.

This procedure was repeated 24 times for every participant using all the combinations of experimental conditions. The order of conditions was randomized.

4. **Post-experimental interview.** Participants shared their experience and strategy. The results of this stage were used to check whether subjects responded purely based on what they saw (as opposed to realizing that they experience an illusion and correcting their answers).

Experiment was programmed and presented using PsychoPy software package (Peirce, 2007).

Participants

The initial sample consisted of 30 adult participants. Data from 4 participants were excluded after the post-experimental interview: they figured out that test circles are always equal, and based their answers on this assumption, not on their actual perception. This results in a final sample of 26 participants (11 male, 15 female) aged from 18 to 47 years (mean age: 22.27, sd: 5.65). All had normal or fully corrected vision.

Experiment 2

The second experiment investigated how robust are the observed regularities. In particular, whether it is necessary to account for individual differential thresholds.

Procedure

Experiment 2 replicates Experiment 1 with one qualitative change: the difference between adaptation circles varies in **absolute units**, not in individual differential thresholds. Therefore, there is no stage of differential threshold estimation. The left circle again has the diameter of 2.5cm, and the diameter of the right circle is 0.1, 0.2, 0.3, 0.4, or 0.5cm bigger. The number of adaptation trials varies from 1 to 6. The conditions are randomized.

Participants

Initial sample consisted of 55 adults. 5 adults were excluded from subsequent analyses, because they figured out that test circles are always equal and based their responses on this assumption. This results in a final sample of 50 participants (22 male, 28 female) aged from 18 to 34 years (mean age: 22.91, sd: 3.47). All of them had normal or fully corrected vision. The sample was divided into two groups based on the results of post-experimental interview:

1. **Naive (35 adults).** These participants did not realize that test circles are always equal.
2. **Non-naive (15 adults).** These participants realized that test circles are always equal, but followed the instruction and tried to base their responses only on their sensations.

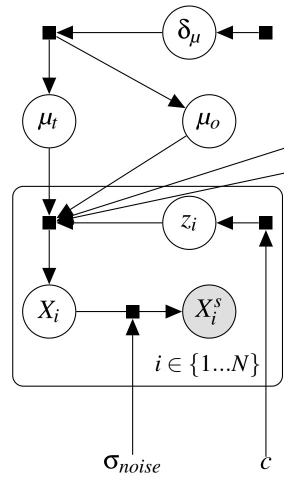

$$\delta_\mu \sim Gamma(10,1)$$
$$\mu_t \sim Normal(0,6)$$
$$\mu_o = \mu_t - \delta_\mu$$
$$\sigma_t, \sigma_o \sim Cauchy(0,5)$$
$$X_i \sim Normal(\mu_i, \sigma_i)$$
$$X_i^s \sim Normal(X_i, \sigma_{noise})$$

Figure 2: Graphical model

Variables: X_i - real stimulus; X_i^s - perceived stimulus (after adding perceptual noise); z_i - indicator variable for the class from which a real stumulus was generated (distributed according to the Chinese Restaurant Process); μ_t and σ_t - μ and σ of the typical class; μ_o and σ_o - μ and σ of another (unobserved) class; δ_μ - the expected difference between two classes; c - coupling probability for CRP; σ_{noise} - perceptual noise.

Computational Model

We formalize the process of adaptation as rational acquisition of the "typical" and "other" stimuli categories. Perception is modeled as an optimal probabilistic inference over the true stimulus parameters given the learned category structure and the noisy sensory input. Graphical model is presented on Figure 2.

1. **Category learning**. Category learning during the adaptation phase is modeled as Bayesian inference of the "typical" category structure. Stimuli in our experiment could be aligned along one relevant dimension (the magnitude of difference between two circles) so we formalize a category as a univariate normal distribution in this dimension. An observer assumes that the true stimuli come from a normal distribution with a mean difference between two circles μ and a standard deviation σ. These are the parameters she estimates to represent a category. Priors on the category's μ and σ are chosen arbitrarily and are set to be relatively diffuse (given the scale of our feature): $\mu \sim normal(0,6)$ and $\sigma \sim cauchy(0,5)$. These parameters are estimated purely from the adaptive stimuli for every particular experimental trial. The adaptive stimuli, in turn, are randomly generated from a normal distribution where μ is the difference between adaptional stimuli in a given condition, and σ_{noise} is some random perceptual noise. We used 0.2 as a noise in final model evaluation, however, we also checked that changing this number does not influence the results. At the end of this stage, the observer has estimations of μ and σ

of the category.

2. **Representation of the unknown category**. The key assumption of a rational observer in our model is that the **centers of two categories are more likely to be relatively distant from each other** than to be close (Figure 3). It is formulated by adding a new parameter: *difference between category prototypes* δ_μ with a prior $Gamma(10,1)$ (modeling the opposite symmetric tail is not necessary in this case, as its likelihood is always practically zero on the test trial). Then, the estimation of the "other" category's mean (μ_o) for each experimental condition is simply $\hat{\mu}_t - \delta_\mu$. Hence, the prior assumptions on the structure of the unknown category are shifted outward from the learned one. Thus, the prior on μ_o is completely defined by an estimated μ of the typical stimuli and the assumption on the difference between categories.

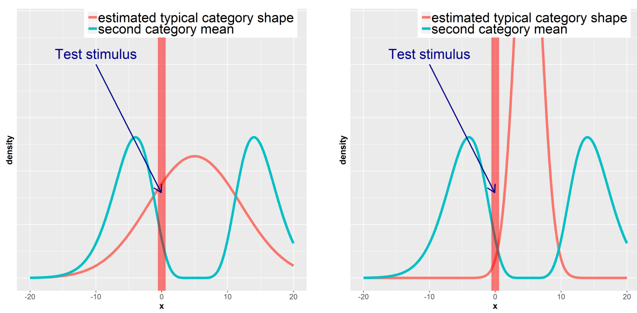

(a) After a small number of trials, there is still a lot of variation in the estimated typical distribution shape. The noisy test stimulus is attributed to the typical category with higher probability and thus the reconstructed true stimulus is shifted towards the "typical" category prototype.

(b) After more trials, the estimated typical distribution shrinks, thus making the attribution to the typical category unlikely. Thus the reconstructed source of the noisy stimulus is shifted towards the closest peak of the "other" category.

3. **Test phase**. Perception of the test stimulus is determined by the decision of what category ("typical" or "other") is more likely to have generated it. Conditional probabilities of the categories are calculated using Bayes' rule (where z_i is a variable indicating category membership):

$$P(z_i = j|X_i^s) = \frac{f(X_i^s|z_i = j)\cdot P(z_i = j)}{f(X_i^s)} = \\ = \frac{f(X_i^s|z_i = j)\cdot P(z_i = j)}{\sum_{j=1}^{\#cat} f(X_i^s|z_i = j)\cdot P(z_i = j)} \qquad (1)$$

Likelihoods of the test stimulus for both categories are taken from the corresponding estimated normal probability density functions. The priors on whether a new stimulus is coming from the known or a new category are estimated using the Chinese Restaurant Process (Anderson, 1991; Navarro & Kemp, 2017). Thus, the prior probability that a new stimulus is generated from the "typical" category is

$$P(z_{n+1} = typical) = \frac{c\cdot n}{1-c+c\cdot n} \qquad (2)$$

where c is a probability that two observation come from the same category (the coupling probability) and n is a number of adaptation trials. The prior probability that a new stimulus comes from an unknown category is

$$P(z_{n+1} = other) = \frac{1-c}{1-c+c\cdot n} \qquad (3)$$

To efficiently reconstruct a real stimulus, perception is shifted toward the probability density of its category. Due to the aforementioned inference bias, the "atypical" and "typical" category densities are shifted in opposite directions. Thus, assimilative illusion onset is formalized as *Bernoulli* random variable with $p = P(typical|test)$, and the contrastive - as *Bernoulli* r.v. with $p = P(other|test)$ respectively.

Model fits 3 parameters: c (coupling probability), δ_μ (difference between prototypes of two categories), and σ (standard deviation of the "other" category).

Bayesian modeling for the paper was implemented using Stan probabilistic language (Carpenter et al., 2017).

Results

Experiment 1

Assimilative aftereffect appeared 103 times (17%), contrastive - 153 times (25%). Notably, more than 50% of the data consisted of the reports of stimuli equality, which correspond to no illusion registered. "No illusion" instances were excluded from the analysis. We applied mixed effects logistic regression and Bayesian mixed effects logistic regression (with non-informative priors). We used the difference between adaptation circles and the number of adaptation trials as predictors, and the illusion type (contrastive (1) vs assimilative (0)) as the outcome variable. This model can be expressed using the following formula:

$$\text{illusion type} \sim \text{number of adaptation trials} + \\ \text{difference between stimuli sizes} + (1|\text{participant})$$

The ANOVA comparison with a zero model was significant ($p < .001$), as well as the tests for both individual coefficients: number of adaptation trials ($p < .001$, $est. = .133$, $sd = .036$, $BF_{10} = 2.5$) and difference between stimuli' sizes ($p < .001$, $est. = .365$, $sd = .108$, $BF_{10} = 56.3$). Both estimates are positive, in line with the the initial hypotheses.

Experiment 2

Assimilative aftereffect appeared 164 times (11%), contrastive - 402 times (28%). To analyse these data, we applied the same frequentist and Bayesian mixed effects models to the three (all, naive, and non-naive) groups separately.

1. For the whole sample, the difference between adaptation circles and the number of adaptation trials are significant predictors with $p < .001$ ($est. = .537$, $sd = .093$, $BF_{10} = 47995.7$) and $p < .05$ ($est. = .171$, $sd = .07$, $BF_{10} = 4.8$) respectively.

Table 1: Performance of Cognitive Model and Bayesian Logistic Regression. Standard deviations are indicated in parentheses.

	Measure	Bayesian LR	Cognitive Model
Experiment 1: assimilative	Recall	0.296 (0.086)	**0.577** (0.082)
	Precision	0.521 (0.12)	**0.522** (0.034)
Experiment 1: contrastive	Recall	**0.817** (0.065)	0.65 (0.056)
	Precision	0.637 (0.018)	**0.701** (0.032)
Experiment 2: assimilative	Recall	0.057 (0.0049)	**0.293** (0.057)
	Precision	0.378 (0.228)	**0.426** (0.044)
Experiment 2: contrastive	Recall	**0.97** (0.03)	0.845 (0.032)
	Precision	0.73 (0.006)	**0.754** (0.012)

2. For the group of naive participants, the difference between adaptation circles is a statistically significant predictor ($p < .01$, $est. = .489$, $sd = .163$, $BF_{10} = 215.5$). The number of adaptation trials is not significant ($p > .05$, $est. = .105$, $sd = .9$, $BF_{10} = 0.5$).

3. For the non-naive participants, both predictors are statistically significant: the number of adaptation trials ($p < .05$, $est. = 1.21$, $sd = .523$, $BF_{10} = 23.8$) and the difference between stimuli sizes ($p < .001$, $est. = .518$, $sd = .151$, $BF_{10} = 254.6$).

The subsequent ANOVA model test (frequentist) was significant ($p < .05$) for all groups. All the estimates are positive, in line with the initial hypotheses.

Model Evaluation

We compared our cognitive model against the Bayesian logistic regression baseline:

$$\textit{illusion type} \sim \textit{number of adaptation trials} + \textit{difference between stimuli sizes}$$

Both the regression and the cognitive model have 3 parameters. Bayesian logistic regression was chosen as a baseline, since it is a very successful descriptive model with the same amount of parameters. In particular, it outperforms a frequentist logistic regression for our data.

The cognitive model fits the whole dataset better than the baseline logistic regression models, but this does not guarantee that the cognitive model would demonstrate better performance on the out-of-sample data as well. Therefore, we used random subsample cross-validation in order to evaluate and compare the **generalization** performance of the models.

1. The data were randomly split into two subsets: train (50% of assimilative data, 50% of contrastive data) and test (remaining 50% of assimilative and 50% of contrastive data)

2. Parameters of the models were estimated on the training set

3. The performance measures (precision and recall) were calculated for models' predictions for the upheld test subset.

We repeated the above steps 50 times and calculated mean precision and recall measures for both assimilative and contrastive classes, along with their standard deviations. The results are shown in the Table 1.

Evaluation metrics: **Precision** and **Recall** measures allow us to compare models based on their sensitivity and accuracy for both classes. **Recall** shows the proportion of the target class occurrences that were accurately predicted. **Precision** shows the proportion of the target class occurrences among the predictions of that class.

The cognitive model repeats the main regularities found in both experiments. In particular, it predicts assimilative illusion more frequently for the smaller differences between adaptive stimuli and number of trials, while the predictions of contrastive illusion follow the reverse pattern. Importantly, the logistic regression does not yield these types of regularities when it predicts assimilative illusion.

The estimates of the "other" category center were always negative, which corresponds to the **contrastive** shifts in perception.

Discussion

Replication

We replicated the results reported in Uznadze (1958, 1966). The difference between adaptation stimuli sizes was a significant predictor of the aftereffect type in all collected datasets. The number of adaptational trials was not a significant predictor for naive participants in the second experiment, but it was significant in the remaining datasets. The signs of all the coefficients were consistent with the initial hypotheses. The effect proved robust to the scale of the differences between stimuli, and overall, Experiments 1 and 2 yielded similar results.

We view this replication as an important impact of our paper. The works of Uznadze are predominantly focused on the study of "set", or "set illusions", which denote the same group of phenomena as perceptual aftereffects. He performed extensive studies of these effects in visual, auditory and haptic modalities (Uznadze, 1966). Nevertheless, although the so-called "Uznadze illusion" (perceptual aftereffect in haptic modality) received some attention (Janzen et al., 1976;

Wohlwill, 1960), most of his contributions remain untranslated and almost entirely unknown to the scientific community outside the post-Soviet space. We find, however, that some of his findings are still relevant and could lead to a better understanding of perceptual aftereffects. We hope that our results would encourage further use of Uznadze visual illusion in the studies of perceptual adaptation aftereffects.

Modeling

The proposed cognitive model performs better than Bayesian logistic regression, which makes it a useful baseline for further research. In particular, it is sensitive to both types of aftereffect and yields more accurate predictions within these categories.

More importantly, our model provides a simple and unified interpretation of seemingly disparate phenomena of assimilative and contrastive aftereffects. This explanation is based on the principles of rational analysis and an intuitive assumption about the category structure inductive bias (different categories have non-coinciding centroids). Thus, our model shows that the apparently low-level perceptual aftereffects may be explained from the logic of higher-level cognitive processes, such as categorization. Moreover, it allows us to view the role of adaptation aftereffects in perception from a new angle: we demonstrate that they may serve as an important part of an optimal stimuli reconstruction process, as opposed to being an artifact or an epiphenomenon.

Future directions

Our model is based on the high-level logic of perception and is not bound to specific low-level mechanisms. This greatly broadens the scope of its potential applications.

Firstly, there is a number of promising extensions of our model **within the domain of adaptation aftereffects**:

- The model could be extended to account for the cases of **illusion absence** (this could be done by incorporating individual perceptual differential thresholds). Since the "no illusion" case is very common in our data, this would make our account of the perceptual aftereffect phenomenon much more complete.

- The model can be scaled to higher dimensions by using a multivariate normal distribution for category representation. This makes it a good candidate for describing perception of high-dimensional realistic objects, such as faces. In case of success, such a unified explanation of higher- and lower-level perceptual processes may contribute to the ongoing debate about the role and even mere presence of top-down effects in perception (Firestone & Scholl, 2016).

Secondly, our model may be broadly applicable **outside of the domain of visual perceptual adaptation**:

- There is a number of well-known spatial context effects in the visual modality (demonstrated by Delboeuf, Ebbinghaus, and Müller–Lyer illusions, among many others (Goto et al., 2007)). The patterns of contrastive and assimilative bias onsets in this domain are very similar to the temporal illusion we studied in this paper (Goto et al., 2007) and may be interpreted in an analogous fashion.

- Our proposed rational categorical perception model could account for **enhanced discriminability** and **perceptual tuning** effects resulting from long-term adaptation. Chinese Restaurant Process used in our model allows to optimally refine the learned category structure as a number of seen examples grows. Shifting percepts towards the true category will be more and more beneficial as the category structure is updated and refined.

Overall, the proposed model has a high promise in demonstrating the role of category learning in perception. While the potential importance of categorical perception has been studied before (e.g. Kuhl & Iverson (1995)), such studies focus on situations when the category structure is known in advance. Our results suggest that assuming that a person always tries to group stimuli into categories (even in the short-term experiments where no obvious categories are apparent) can greatly broaden the scope of this approach and provide a unified explanation to a wide range of perceptual effects.

To facilitate further research, we make all the data, analyses, and code openly available [1].

References

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological review*, *98*(3), 409.

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., et al. (2017). Stan: A probabilistic programming language. *Journal of statistical software*, *76*(1).

Chalk, M., Seitz, A. R., & Seriès, P. (2010). Rapidly learned stimulus expectations alter perception of motion. *Journal of Vision*, *10*(8), 2–2.

Chopin, A., & Mamassian, P. (2012). Predictive properties of visual adaptation. *Current biology*, *22*(7), 622–626.

Cibelli, E., Xu, Y., Austerweil, J. L., Griffiths, T. L., & Regier, T. (2016). The sapir-whorf hypothesis and probabilistic inference: Evidence from the domain of color. *PloS one*, *11*(7), e0158725.

Feldman, N. H., Griffiths, T. L., & Morgan, J. L. (2009). The influence of categories on perception: Explaining the perceptual magnet effect as optimal statistical inference. *Psychological review*, *116*(4), 752.

Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for" top-down" effects. *Behavioral and brain sciences*, *39*.

Fritsche, M., Mostert, P., & Lange, F. P. de. (2017). Opposite effects of recent history on perception and decision. *Current Biology*, *27*(4), 590–595.

[1] https://github.com/blinodelka/Illusions_of_set

Gescheider, G. (1997). Chapter 3: The classical psychophysical methods. *Psychophysics: the fundamentals. 3rd ed. Mahwah: Lawrence Erlbaum Associates*.

Gibson, J. J., & Radner, M. (1937). Adaptation, after-effect and contrast in the perception of tilted lines. i. quantitative studies. *Journal of experimental psychology*, *20*(5), 453.

Goldstone, R. L. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, *123*(2), 178.

Goldstone, R. L. (1995). Effects of categorization on color perception. *Psychological Science*, *6*(5), 298–304.

Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(1), 69–78.

Goto, T., Uchiyama, I., Imai, A., Takahashi, S., Hanari, T., Nakamura, S., et al. (2007). Assimilation and contrast in optical illusions 1. *Japanese Psychological Research*, *49*(1), 33–44.

Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: neural models of stimulus-specific effects. *Trends in cognitive sciences*, *10*(1), 14–23.

Howard, I. P., & Rogers, B. J. (1995). *Binocular vision and stereopsis*. Oxford University Press, USA.

Janzen, H. L., et al. (1976). A developmental analysis of set patterns in children: A normative study.

Kleinschmidt, D., & Jaeger, T. F. (2011). A bayesian belief updating model of phonetic recalibration and selective adaptation. In *Proceedings of the 2nd workshop on cognitive modeling and computational linguistics* (pp. 10–19).

Kuhl, P. K., & Iverson, P. (1995). Linguistic experience and the perceptual magnet effect. *Speech perception and linguistic experience: Issues in cross-language research*, 121–154.

Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature neuroscience*, *4*(1), 89.

Leopold, D. A., Rhodes, G., Müller, K.-M., & Jeffery, L. (2005). The dynamics of visual adaptation to faces. *Proceedings of the Royal Society of London B: Biological Sciences*, *272*(1566), 897–904.

Navarro, D. J., & Kemp, C. (2017). None of the above: A bayesian account of the detection of novel categories. *Psychological review*, *124*(5), 643.

Palumbo, R., DAscenzo, S., Quercia, A., & Tommasi, L. (2017). Adaptation to complex pictures: exposure to emotional valence induces assimilative aftereffects. *Frontiers in psychology*, *8*, 54.

Peirce, J. W. (2007). Psychopypsychophysics software in python. *Journal of neuroscience methods*, *162*(1-2), 8–13.

Raymond, J. (1993). Complete interocular transfer of motion adaptation effects on motion coherence thresholds. *Vision Research*, *33*(13), 1865–1870.

Rhodes, G., & Jeffery, L. (2006). Adaptive norm-based coding of facial identity. *Vision research*, *46*(18), 2977–2987.

Stocker, A. A., & Simoncelli, E. P. (2006). Sensory adaptation within a bayesian framework for perception. In *Advances in neural information processing systems* (pp. 1289–1296).

Suzuki, S., & Cavanagh, P. (1998). A shape-contrast effect for briefly presented stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(5), 1315.

Thompson, P., & Burr, D. (2009). Visual aftereffects. *Current Biology*, *19*(1), R11–R14.

Uznadze, D. N. (1958). Experimental basis of the psychology of set. *Experimental Studies on the Psychology of the Set*, 5–79.

Uznadze, D. N. (1966). The psychology of set.

Webster, M. A., & MacLeod, D. I. (2011). Visual adaptation and face perception. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1571), 1702–1725.

Wei, X.-X., & Stocker, A. A. (2015). A bayesian observer model constrained by efficient coding can explain'anti-bayesian'percepts. *Nature neuroscience*, *18*(10), 1509.

Wohlwill, J. F. (1960). Developmental studies of perception. *Psychological Bulletin*, *57*(4), 249.