Adaptation and Learning Priors in Visual Inference

Position Paper

Alex Kale[*]
University of Washington

Jessica Hullman[†]
Northwestern University

Abstract

Users' prior expectations are an important but understudied degree of freedom in visual inferences. We ask: To what extent are priors learned through visual experience? How do they impact behavior? Can we design visual analytics systems to manipulate users' priors and calibrate their sensitivity to the signal in data? We connect theoretical accounts of priors in visual inference with the empirical results from psychophysical and physiological studies of *visual adaptation*: a ubiquitous process by which the neural code calibrates to statistics of the immediate environment. The way the visual system adapts its internal representations based on experience explains implicit learning of empirical priors for the purpose of visual inference. Drawing on the visual adaptation literature, we present a framework for researching and designing for priors in visual inference.

Index Terms: Human-centered computing—Visualization—Visualization theory, concepts and paradigms

1 Visual Inference as Model Checking

Both exploratory data analysis (EDA) and statistical inference are concerned with the discrepancy between observed data and the predictions of a **reference model** [6,13,14]. While often implicit in the case of EDA, the reference model represents the analyst's prior belief about how the data should be distributed and what relationships are expected. In the case of confirmatory statistical inference, the reference model represents a null hypothesis (e.g., no difference between two groups). Thus, visual statistical inference usually involves some sort of implicit or explicit comparison between the observed data y and predicted data y^{rep} replicated from the reference model. Gelman [13,14] analogizes this formalization to model checking.

In recent work [19], Hullman argues that how we show uncertainty in visualizations informs the user's reference model for the purpose of visual inference. Specifically, if we do not show users uncertainty in a visualization, then their notion of the reference model is unconstrained. Users are left to imagine a data generating process and how it might map onto a visualization of the observed data. On the other hand, when we visualize uncertainty we expose a "seam" in the process by which the data were constructed and visualized. From this perspective, *effective* uncertainty visualization signals the author's intended reference distribution to the user and enables them to easily compare y to y^{rep}, and any uncertainty visualization reduces degrees of freedom in the user's inference of y^{rep}.

That visual analytics platforms should enable comparisons between the observed data and a reference or null model is perhaps most memorably illustrated by visualization **lineups** [6,40](Fig. 2). The lineup technique aims to apply confirmatory procedures to the typically informal process of visually assessing a plot to determine if it contains a pattern. A lineup is composed of a set of small multiples where one plot depicts observed data y and the others depict replications y^{rep} from a null model. Prior work [27] characterized the extent to which users are able to recognize real data plots among null plots in lineups at different levels of signal strength, and they found large individual differences which they attribute to visual skill. We consider the extent to which such individual differences may reflect exposure and *visual adaptation* to charts, through which users form *empirical priors* about expected visual properties of y^{rep}.

Priors describing reference models used in visual inference are an important factor in graphical perception, yet visualization researchers and practitioners may find it difficult to reason systematically about their impacts in visual analytics systems. In particular, we focus on the capacity for visual experience to influence priors [1]. We call expectations learned from experience **empirical priors**, in contrast with expectations informed by declarative knowledge. It is prudent for our community to better understand: *What is reasonable to assume about how a user's previous exposure to charts will impact their perception of patterns in a visualization?*

2 Learning Priors Through Experience

We examine whether it would be reasonable to assume that people learn priors for graphical statistical inference through an optimal Bayesian accumulation of experiences. In previous theoretical work, DeWeese and Zador [8] make predictions about the how a Bayesian would learn the underlying mean and variance of a sequence of observed quantities. They derive a recursive definition of the Bayesian estimate of these parameters θ_i after a set observations j through i,

$$P(\theta_i \mid s_{j\leq i}) = \frac{P(s_i \mid \theta_i, s_{j<i})P(\theta_i \mid s_{j<i})}{P(s_i \mid s_{j<i})}$$

where $s_{j<i}$ are observations before the current observation s_i. The running estimate of mean and variance for a set of observations depends on previous observations and the prior they inform.

This formulation of Bayesian learning makes several *noteworthy predictions* about the way people learn priors from experience [8]:

1. It takes more observations to accurately learn the mean and variance parameters underlying a stream of observations when the variance is high.
2. Learning variance accurately requires more observations than learning means.
3. When means change over time, it does not take many observations to detect the new mean.
4. When variance changes over time, there is an *asymmetry* in how long it takes to detect increments vs decrements: *Increments of variance are noticed after fewer observations than decrements of variance because they are not expected.*
5. Previous learning interferes with later learning because priors start flat but over time become more specific and stable.

These are all consequences of a theoretical statistical formulation of learning. *Does human behavior conform to these predictions?*

[*]e-mail: kalea@uw.edu
[†]e-mail: jhullman@northwestern.edu

[1] There is a body of literature on implicit learning from non-visual experiences (e.g., [30]) that falls outside the scope of our review.

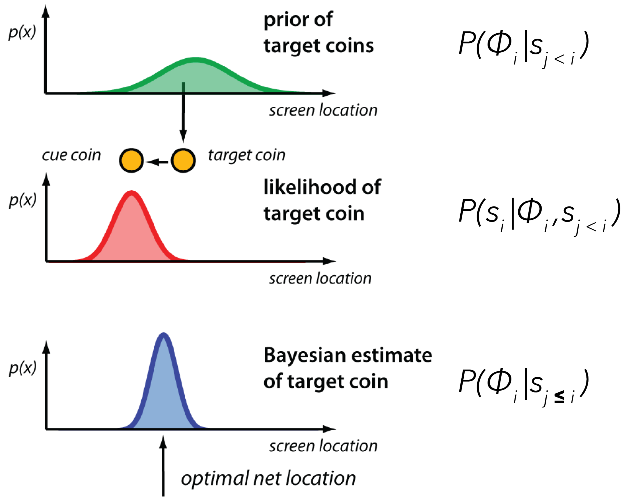

Figure 1: An illustration of how the locations of the coin and precue are sampled and used in Bayesian update (adapted from Berniker et al. [3]). Formulas match the DeWeese and Zador's [8] recursive formulation of Bayesian adaptation.

2.1 Reviewing a Study on Bayesian Visual Learning

Berniker and colleagues [3] measure visual learning and compare observed behavior to a benchmark model resembling DeWeese and Zador's [8] formulation of Bayesian update. They find close *correspondence between human behavior and the predictions of a Bayesian benchmark*, suggesting that people are optimal Bayesians when it comes to visual learning. Note that this study sets up a situation where knowledge outside the experimental task is irrelevant, so consequently their findings reflect how priors are learned in the absence of previous conceptual knowledge. Before we interrogate the generalizability of this finding, we explicate this study as a concrete example of how people learn priors through visual experience.

2.1.1 Task Structure

Participants try to catch a coin in a virtual net based on a noisy precue about the coin's location. The **precue** is a visual cue at the beginning of each trial which gives an uncertain indication of where the actual coin will fall along a horizontal axis. On each trial, the location of the coin is drawn from a Gaussian distribution, and the location of the precue is drawn from a second Gaussian distribution with a mean equal to the coin's location and standard deviation equal to a tenth the width of the display (Fig. 1). The participant must place their net where they expect the coin to fall, and they receive feedback as to whether they catch the coin. This task requires a Bayesian visual inference where participants must integrate the location of the precue (i.e., likelihood) with their prior about the parameters of the underlying Gaussian distribution from which the coin's location is drawn. The experimenters manipulate the mean and standard deviation of the coin location distribution between blocks of trials and infer the participant's prior from their net placement across many repeated trials.

2.1.2 Findings

Berniker and colleagues [3] benchmark participants' implicit estimates of the mean and variance of coin locations against an idealized Bayesian model representing the upper bound of inferential accuracy after a given sequence of trials. Human performance mirrors the model relatively closely, corroborating evidence that the human visual system is able to automatically extract summary statistical information from visual ensembles presented over time [1, 18, 26, 29]. Their findings mostly follow the predictions of DeWeese and Zador [8]:

1. Participants' sense of the mean is less accurate when the variance of the coin location distribution is high. However, since participants seem to start the task with a flat prior, their sense of variance converges more quickly when variance is high.
2. Participants learn the mean very quickly, within about ten trials, whereas it takes them about 200 trials to learn the variance.
3. When the mean coin position changes abruptly, participants notice after only a couple trials.
4. When variance changes, participants update their priors more easily when variance increases than when variance decreases.
5. Previous learning interferes with participants' ability to update their priors after changes in the variance of coin locations.

Given the correspondence between DeWeese and Zador's theory [8] and empirical evidence from Berniker [3] and others [2, 32], it is enticing to conclude that visual learning is approximately Bayesian. However, the foregoing evidence leaves us with unanswered practical questions. For example, in the context of visual analysis, *do chart users learn priors in the same way regardless of variations in marks and encodings? How does conceptual knowledge about chart types and data context integrate with Bayesian visual learning?* In order to design for visual inferences which incorporate user priors, we need to anticipate how these priors are formed and their impacts on graphical perception. In search of a framework for reasoning about these questions in relation to empirical priors, we review the visual neuroscience literature on a process supporting implicit learning from experience: adaptation.

3 Visual Adaptation to the Environment

A large body of literature in vision science points to a physiological mechanism by which empirical priors may be represented. **Adaptation** is a ubiquitous process by which the representational codes of neurons in the brain are continuously updated relative to the norm of an environment [23, 36, 37], which is learned through experience. We argue that this sense of the norm functions like an empirical prior about what statistical properties are expected in a given environment. If we view visual analytics as a kind of artificial environment, literature on adaptation suggests open questions for visualization research as well as design patterns for systems supporting visual inferences.

3.1 What is Adaptation?

Adaptation is continuous tuning of the way neurons represent information. It weakens responses to stimuli which are common in the environment and relatively strengthens responses to novel stimuli [23]. The result is a *perceptual calibration* of which stimuli appear subjectively neutral and which stimuli attract attention [37].

Some well known visual illusions such as color afterimages and motion aftereffects (e.g., the "waterfall illusion") are caused by visual adaptation [15, 23, 37]. These illusions are used in teaching and research to illustrate and study opponent processing mechanisms in the visual system. In each case prolonged exposure diminishes the response of neurons to the adapting stimulus, and consequently the viewer experiences an afterimage of the opposite percept when the adapting stimulus stops: In color afterimages, exposure to a red or blue stimulus yields a green or yellow afterimage, respectively. In motion aftereffects, exposure to a moving stimulus yields illusory motion in the opposite direction.

Accumulating evidence of adaptation to visual stimuli with varying complexity (reviewed separately by Kohn [23] and Webster [37]) suggests that adaptation is *ubiquitous*. For example, adaptation has been measured for stimuli as complex as faces, viewpoints on objects, and the perceived navigability of landscapes.

The strength and duration of adaptation effects depend on the duration of exposure to the adapting stimulus [8, 9, 15, 23, 37]. Chopin and Mamassian [7] even find *opposite effects of adaptation at different timecourses*: Perception is *repelled* (i.e., biased away) from the adapting stimulus within two or three minutes of viewing that stimulus. However, perception is *attracted* toward adapting stimuli shown in the more distant past. Extensive evidence shows both repulsive and attractive effects of adaptation [23].

We account for these opposite effects by distinguishing two components of adaptation, the *adjusted sensory signal* and the *sense of the statistical regularities of an environment*. We call the calibration of sensory signal to a sense of the norm in a specific environment an **adaptation state**. Friston [11, 12] argues that this sense of the norm represents an **empirical prior**, which encodes our predictions about the possible causes of sensory input. According to Friston, sensory signals are adjusted to represent **prediction errors**, such that signals are coded as differences from our expectations. In this framework, the short-term repulsive effects of adaptation reflect rapid adjustments to incoming sensory signal, and longer-term attractive effects reflect the formation of empirical priors.

3.2 Why and How Does the Visual System Adapt?

The brain is composed of cells called **neurons** which communicate through electrochemical pulses called **action potentials**. The patterns and frequencies of these pulses encode information about the combination of inputs to each neuron and constitute the cell's output.

The rate at which neurons can produce action potentials is limited, so they need to map a large domain of stimulation onto a smaller dynamic range of action potential frequencies. For example, light-sensitive cells in the retina have a response range of about two orders of magnitude, yet they must encode light levels that vary over ten orders of magnitude in the natural environment [2]. Adaptation enables this flexibility of representation resulting in a maximally **efficient code** for information transmission [5, 35], which reduces redundancy in and improves the metabolic efficiency of the brain [23].

In the efficient neural code, *the meaning of a specific sequence of action potentials depends on the norm of the environment* to which the visual system is adapted [9, 36, 37]. In order to disambiguate the meaning of this **normalized** code, the visual system needs to somehow represent the the statistical properties of its current environment [8]. In a landmark study, Fairhall and colleagues [9] analyze direct measurements of action potentials [2] and find evidence that *separate components of the neural code represent normalized sensory signal and immediate statistical context*, respectively. Sensory information is represented in brief patterns of action potentials, whereas the variance among stimuli presented is represented in slower changes of the overall pulse rate. They find that the component of adaptation states which represents variance is slow to learn variance to a high degree of accuracy and learns increases in variance faster than decreases, mirroring optimal Bayesian learning [8].

We argue that the confluence of findings from theoretical, physiological, and behavioral research suggest that learning environmental norms from visual experience to support adaptation is the mechanism through which we form empirical priors for visual inference.

3.3 How Do Empirical Priors Influence Perception?

As visualization researchers, the crux of our interest in adaptation is how normalization to the environment impacts perception. We argue that the statistics which the visual system learns about different environments to support adaptation function like Bayesian priors about what *expected* versus *unexpected* visual patterns look like.

[2]Fairhall and colleagues [9] take direct measurements of action potentials from visual motion-sensitive neurons in the house fly, a procedure too invasive for human subjects.

3.3.1 Overemphasizing the Mean

There is a long history of research on a phenomenon called *the central tendency of judgments* [17] in which estimates of a sequence of quantities are attracted toward the mean of the previously presented quantities. We argue that this perceptual attraction toward the mean is a consequence of adaptation states. The visual system learns the mean through ensemble processing to form a *prior about what is normal* in a sequence of quantities. This prior induces an adaptation state in which the neural code represents quantities as deviations from the norm [9, 36–38] or prediction errors [11, 12], allocating the extremes of its dynamic range to encode unexpected stimuli which are consequently more salient but perceived with lower fidelity.

Empirical priors result in a pattern of perceptual bias where values close to the mean are estimated with greater accuracy than values further from the mean [37, 38]. Zhang and colleagues [41] have proposed that this is a ubiquitous pattern of bias in judgments of frequency and probability. Similarly, others have proposed that perceptual bias toward the mean can account for seemingly irrational behaviors of risk seeking and aversion [21], which have traditionally been characterized as framing effects [34]. Note that a qualitatively similar but even stronger bias would occur if the mean itself is *heuristically substituted* [33] for individual quantities in a set. Thus, perceptual biases induced by adaptation states may interact with or be exacerbated by cognitive biases such as representativeness [33].

3.3.2 Inferring Different Environments

The visual system is organized as a **hierarchical network** such that successive stages of information processing build up representations of increasing abstraction, from low-level visual features such as orientation, color, and motion to categories of objects such as faces, places, and perhaps even charts. The visual system adapts [23, 37] and learns priors [11, 12] at each level of abstraction such that high-level hyperpriors entail predictions about what low-level visual features are expected in different environments.

The visual system's ability to adapt to the norms of different environments helps to *stablize perception across different contexts*. Because natural environments vary widely in their statistical properties (e.g., lightness [2, 31] and color [39]), we need a neural code which is relative to immediate context in order to maintain perceptual constancy [10, 37, 38]. For example, the visual system learns priors about objects' surface reflectances and illumination in different environments (e.g., times of day) to maintain constant mapping between colors and objects despite changes in raw sensory signal under different sources of illumination [4]. In this sense, the visual system infers the conditions that might produce a pattern of sensory input by learning priors for different environments and attempting to minimize prediction error [11, 12]. Thus, understanding adaptation states gives us a framework for understanding visual inferences.

4 Designing for Adaptation in Visual Analytics

In visual analytics contexts, users bring prior knowledge and experience to their interpretations of signal in charts. It also seems very likely that the visual system adapts to the statistical regularities of a set of charts viewed in close spatial and temporal proximity. We can think of modern work situations involving visual skill, such as visual analytics, as artificial environments [37] with distinct adaptation states contingent on the user's experience. This raises questions about how we should design visual analytics interfaces with adaptation states and consequent perceptual biases in mind.

4.1 Defining Reference Models and Setting Priors

Consider the visualization lineup [6, 40] where a user's task is to recognize which of an array of charts shows real data y rather than simulated data y^{rep}. *How should we define the reference model from which to sample y^{rep}?* In analogy to null hypothesis significance

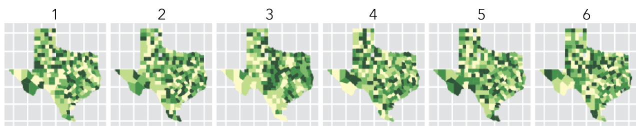

Figure 2: A lineup of cancer rates in Texas (from Wickham et al. [40]). Chart 3 shows real data while other charts show random noise.

testing, Wickham and colleagues [40] suggest using relatively uninformed reference models which represent the absence of a trend. Consider their example (Fig. 2) of a lineup showing cancer rates in each county of Texas where the suggested null model is a random spatial trend. This implies an assumption that comparing real data to a *null model* is the right design choice.

However, it might be more appropriate to use reference models which are informed by prior knowledge about one or more possible data generating processes. For example, a more appropriate reference model for cancer rates in each county of Texas would be the population in each county multiplied by a known average cancer rate plus simulated sampling error based on the number of people surveyed in each county. Such a reference model would suggest that a spatial pattern resembling population density should be expected.

Similar to the idea of a null reference model, visual analysis software designers might think they ought to enforce a flat prior if there is concern that a user might not be adapted in a way that serves the task at hand. We argue that *we probably should not try to enforce flat priors*. For highly experienced analysts, prior expectations inform reference models which are valuable for inference, and it might be harmful to try to extinguish them in favor of a flat prior.

Outside of lineups, visualization users still compare observed data to an *often implicit reference model* when making visual inferences [13, 14, 19]. We argue for showing users instantiations of reference models as an additional step in data analysis which can be used to *calibrate user priors to any reference model*[3]. For example, in a related graphical perception study [20] we showed users uncertainty visualizations of two different reference distributions representing two alternative data generating models y^{rep}, and users chose which of these two models was more likely to have produced a given chart of observed data y. Users were more sensitive to signal in y differentiating these two reference models when we showed them y^{rep} in animated hypothetical outcome plots (HOPs) rather than static uncertainty visualizations. Perhaps HOPs improve sensitivity to signal because their temporal nature is more in line with how the visual system learns priors by adapting over time. We argue that using HOPs to manipulate priors through adaptation is a promising design pattern that could be used to convey any reference model.

4.2 Adaptation States, Expertise, and Knowledge

Recall that the visual system learns priors at multiple levels of abstraction [11, 12], invoking different adaptation states when stimuli are recognized as belonging to different categories or environments [36]. This means that the visual system is probably capable of learning separate priors for different chart types, distributions, and data generating processes. In this sense, the specificity of users' adaptation states probably depends on their familiarity with chart types and graphical conventions, their previous exposure to data visualizations, and their domain knowledge about the data context. We consider two examples of users with different levels of expertise.

Imagine a *data scientist conducting exploratory data analysis* (EDA). She see a series of charts with varying geometries and distributions[4]. Because her visual system recognizes different chart types and distributions as special categories of stimuli, she interprets each chart in terms of the visual features she expects and does not expect to see based on her prior for a given chart type and distribution.

Importantly, the visual system's ability to differentiate between distributions based on their visual properties may be diminished by smart defaults in visual analytics software which scale axes in order to reduce unused white space. This artificially stabilizes the mean and variance of the positions of marks on charts regardless of important differences in the distributions they represent.

In general, users may perceive each chart using the adaptation state, among a set of candidate adaptation states, which minimizes prediction error [11, 12]. In this sense, they implicitly search the space of possibilities suggested by their experience to select a prior under which they are least surprised by what they see in a chart. Supporting the idea that adaptation states can be specific to task environments, experiments find adaptation states which are contingent on task-relevant properties of radiological images [24, 25].

Imagine a *user who has little experience with visualizations viewing a single chart in the news*. Since this user may not reliably employ separate priors for different chart types or data distributions, his adaptation state is more likely to be tuned to expectations about low-level visual features rather than categorical distinctions. For example, if he sees a multiple line chart showing concurrent trends for different countries, the ensemble of lines might inform his sense of what slope or contour is the norm and which countries look like outliers. In this sense, a user's prior may be informed by the space of possibilities implied by a single chart [19]. Especially in cases where users lack relevant experience and background knowledge, showing predictions from different reference models y^{rep} may help to calibrate users to possible interpretations of the data.

However, *manipulating users' priors through visual experience alone is unlikely to work*. At short time scales the effects of adaptation are weak [3, 9, 23, 37] and opposite of the effect of forming priors [7]. Further, when adaptation to a new stimulus disrupts an older, more entrenched adaptation state, the older adaptation state "spontaneously recovers" after a short time [28]. This suggests that attempts to manipulate priors through visual adaptation alone will result only in transient changes. In order to manipulate the adaptation states of users in ways that are likely to stick, we need to design adaptation interfaces so that users *associate samples from y^{rep} with semantic content* such as possible explanations for a phenomenon.

Future research should investigate how the visual system learns categorical priors in visual inference environments by benchmarking behavior against hierarchical Bayesian models, such as those described by Griffiths [16]. To measure user priors, visualization researchers and software designers might follow Berniker and colleagues [3] by inferring user priors from judgments, or they might employ the approach of Kim and colleagues [22] by asking users to draw their expectations.

5 Conclusion

We review vision science literature on adaptation, a ubiquitous process by which neural representation is normalized to the statistical regularities of the environment. We argue that the sense of what is normal or expected in an environment, which underlies adaptation, functions like a prior supporting visual inferences. We explicate predictions of how these empirical priors are learned through visual experience and discuss implications for design in visual analytics systems. Our hope is that this work will inspire a wave of graphical perception studies on adaptive learning of priors, and will ultimately result in new ways of designing for user priors in visualization.

[3]This echos Buja and colleagues' [6] proposal of the visualization *Rorschach*, where users are shown many charts of model predictions y^{rep} to train them to recognize the statistical signatures of random variability.

[4]Irregular distributions especially need to be studied since the predictions of DeWeese and Zador [8] only generalize to distributions with thin tails.

ACKNOWLEDGMENTS

The authors thank the member of the University of Washington Vision-Cognition Lab for contributing references and feedback.

REFERENCES

[1] G. A. Alvarez. Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Sciences*, 15(3):122–131, 2011. doi: 10.1016/j.tics.2011.01.003

[2] H. Barlow and J. D. Mollon. The Senses. In *Cambridge Texts in the Physiological Sciences*, vol. 3. Cambridge University Press, 1982.

[3] M. Berniker, M. Voss, and K. Kording. Learning Priors for Bayesian Computations in the Nervous System. *PLoS ONE*, 5(9), 2010. doi: 10.1371/journal.pone.0012686

[4] D. H. Brainard and W. T. Freeman. Bayesian color constancy. *Journal of the Optical Society of America A*, 14(7):1393–1411, 1997.

[5] N. Brenner, W. Bialek, and R. D. R. V. Steveninck. Information Transmission. *Neuron*, 26:695–702, 2000.

[6] A. Buja, D. Cook, H. Hofmann, M. Lawrence, E.-k. Lee, D. F. Swayne, and H. Wickham. Statistical Inference for Exploratory Data Analysis and Model Diagnostics. *Royal Society Philosophical Transactions*, Series A(367):4361–4383, 2009.

[7] A. Chopin and P. Mamassian. Predictive Properties of Visual Adaptation. *Current Biology*, 22:622–626, 2012. doi: 10.1016/j.cub.2012.02.021

[8] M. DeWeese and A. Zador. Asymmetric dynamics in optimal variance adaptation. *Neural Computation*, 10(5):1179–1202, 1998.

[9] A. L. Fairhall, G. D. Lewen, W. Bialek, and R. R. D. R. V. Steveninck. Efficiency and ambiguity in an adaptive neural code. *Nature*, 412(23):787–792, 2001.

[10] D. H. Foster. Color constancy. *Vision Research*, 51(7):674–700, 2011. doi: 10.1016/j.visres.2010.09.006

[11] K. Friston. The free-energy principle: A rough guide to the brain? *Trends in Cognitive Sciences*, 13(7):293–301, 2009. doi: 10.1016/j.tics.2009.04.005

[12] K. Friston and S. Kiebel. Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B*, 364(March):1211–1221, 2009. doi: 10.1098/rstb.2008.0300

[13] A. Gelman. A Bayesian Formulation of Exploratory Data Analysis and Goodness-of-Fit Testing. *International Statistical Review*, 2003.

[14] A. G. Gelman. Exploratory Data Analysis for Complex Models. *Journal of Computational and Graphical Statistics*, 13(4):755–779, 2004. doi: 10.1198/106186004X11435

[15] D. M. Glasser, J. M. G. Tsui, C. C. Pack, and D. Tadin. Perceptual and neural consequences of rapid motion adaptation. 108(45), 2011. doi: 10.1073/pnas.1101141108

[16] T. L. Griffiths, C. Kemp, and J. B. Tenenbaum. *Bayesian models of cognition*, chap. 3, pp. 59–100. The Cambridge handbook of computational psychology. Cambridge University Press, New York, NY, US, 2008.

[17] H. L. Hollingworth. The Central Tendency of Judgment. *The Journal of Philosophy, Psychology and Scientific Methods*, 7(17):461–469, 1910.

[18] B. Hubert-Wallander and G. M. Boynton. Not all summary statistics are made equal: Evidence from extracting summaries across time. *Journal of Vision*, 15(2015):1–12, 2015. doi: 10.1167/15.4.5.doi

[19] J. Hullman. Why authors don't visualize uncertainty. *IEEE transactions on visualization and computer graphics*, in press.

[20] A. Kale, F. Nguyen, M. Kay, and J. Hullman. Hypothetical outcome plots help untrained observers judge trends in ambiguous data. *IEEE transactions on visualization and computer graphics*, 2018.

[21] M. W. Khaw, Z. Li, and M. Woodford. Risk Aversion as a Perceptual Bias. Working paper, 2017.

[22] Y.-S. Kim, L. A. Walls, P. Krafft, and J. Hullman. A bayesian cognition approach to improve data visualization. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pp. 682:1–682:14. ACM, New York, NY, USA, 2019. doi: 10.1145/3290605.3300912

[23] A. Kohn. Visual Adaptation: Physiology ,Mechanisms, and Functional Benefits. *Journal of Neurophysiology*, 97:3155–3164, 2007. doi: 10.1152/jn.00086.2007.

[24] E. Kompaniez, C. K. Abbey, J. M. Boone, and M. A. Webster. Adaptation Aftereffects in the Perception of Radiological Images. *PloS One*, 8(10):1–12, 2013. doi: 10.1371/journal.pone.0076175

[25] E. Kompaniez-dunigan, C. K. Abbey, J. M. Boone, and M. A. Webster. Adaptation and visual search in mammographic images. *Attention, Perception, {&} Psychophysics*, 77(4):1081–1087, 2015. doi: 10.3758/s13414-015-0841-5.Adaptation

[26] A. Y. Leib, A. Kosovicheva, and D. Whitney. Fast ensemble representations for abstract visual impressions. *Nature Communications*, 7:1–10, 2016. doi: 10.1038/ncomms13186

[27] M. Majumder, H. Hofmann, and D. Cook. Validation of Visual Statistical Inference , Applied to Linear Models Validation of Visual Statistical Inference , Applied. *Journal of the American*, 108(503):942–956, 2013. doi: 10.1080/01621459.2013.808157

[28] J. Mesik, M. Bao, and S. A. Engel. Spontaneous recovery of motion and face aftereffects. *Vision Research*, 89:72–78, 2013. doi: 10.1016/j.visres.2013.07.004

[29] C. Oriet and K. Hozempa. Incidental statistical summary representation over time. 16:1–14, 2016. doi: 10.1167/16.3.3.doi

[30] A. S. Reber. Implicit Learning and Tacit Knowledge. *Journal of Experimental Psychology: General*, 118(3):219–235, 1989. doi: 10.1037/0096-3445.118.3.219

[31] O. Schwartz and E. P. Simoncelli. Natural signal statistics and sensory. *Nature Neuroscience*, 4(8):819–825, 2001.

[32] S. M. Smirnakis, M. J. Berry, D. K. Warland, W. Bialek, and M. Meister. Adaptation of retinal processing to image contrast and spatial scale. *Nature*, 386(6620):69–73, 1997. doi: 10.1038/386069a0

[33] A. Tversky and D. Kahneman. Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157):1124–1131, 1974.

[34] A. Tversky and D. Kahneman. The Framing Of Decisions And The Psychology Of Choice. *Science*, 211(30):453–458, 1981.

[35] M. J. Wainwright. Visual adaptation as optimal information transmission. *Vision Research*, 39:3960–3974, 1999.

[36] M. A. Webster. Adaptation and visual coding. *Journal of Vision*, 11(5), 2011. doi: 10.1167/11.5.3

[37] M. A. Webster. Visual Adaptation. *Annual Review of Vision Science*, Nov 1(1):547–567, 2015. doi: 10.1146/annurev-vision-082114-035509.Visual

[38] M. A. Webster and D. Leonard. Adaptation and perceptual norms in color vision. *Journal of the Optical Society of America. A, Optics, Image Science, and Vision*, 25(11):2817–2825, 2008.

[39] M. A. Webster and J. D. Mollon. Adaptation and the Color Statistics of Natural Images. *Vision Research*, 37(23):3283–3298, 1997.

[40] H. Wickham, D. Cook, H. Hofmann, and A. Buja. Graphical inference for infovis. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):973–979, 2010. doi: 10.1109/TVCG.2010.161

[41] H. Zhang, L. T. Maloney, and D. Lange. Ubiquitous log odds: A common representation of probability and frequency distortion in perception, action, and cognition. *Frontiers in Neuroscience*, 6(January):1–14, 2012. doi: 10.3389/fnins.2012.00001