Structure learning principles of stereotype change

Samuel J. Gershman^{1,2,*} and Mina Cikara¹

¹Department of Psychology, Harvard University, Cambridge, MA

²Center for Brains, Minds, and Machines, Cambridge, MA

*Corresponding author: gershman@fas.harvard.edu

August 25, 2021

Abstract

Why, when, and how do stereotypes change? This paper develops a computational account based on the principles of structure learning: stereotypes are governed by probabilistic beliefs about the assignment of individuals to groups. Two aspects of this account are particularly important. First, groups are flexibly constructed based on the distribution of traits across individuals. This allows the model to explain the phenomenon of subtyping, whereby deviant individuals are segregated from a group, thus protecting the group's stereotype. Second, groups are hierarchically structured, such that groups can be nested. This allows the model to explain the phenomenon of subgrouping, whereby a collection of deviant individuals is organized into a refinement of the group. The structure learning account also sheds light on several factors that determine stereotype change, including perceived group variability, individual typicality, cognitive load, and sample size.

Keywords: Bayesian modeling; intergroup cognition; social structure learning; stereotypes

Author note: We are grateful to Mahzarin Banaji for inspiring discussions and to Joel Martinez for comments on a draft of the manuscript. This work was supported by a grant from the National Science Foundation (BCS-2116543) and the Center for Brains, Minds, and Machines (CBMM), funded by NSF STC award CCF-1231216. Simulation code is publicly available at <https://github.com/sjgershm/Hierarchical-groups>.

Introduction

Stereotypes are notoriously resistant to change. Lippmann (1922) famously quipped that “there is nothing so obdurate to education or criticism as the stereotype” (p. 99). Eight decades later, this view was echoed by Banaji (2002):

Stereotypes are the vehicles of essentialist thinking about social groups. Dispositional group attributions, or the belief that groups are inherently the way they are, can lead to the assessment that attributes associated with groups are stable and unchanging. (p. 15102)

According to this cognitive view, the obduracy of stereotypes is grounded in beliefs about the structure of social groups. Stereotypes do not change because we believe groups do not change.

And yet stereotypes *do* change over longer periods of time (Bergsieker et al., 2012). For example, public opinion polls show that gender stereotypes track changes in social and occupational roles (Eagly et al., 2020; Miller et al., 2015), a pattern also reflected in measures of implicit attitudes (Charlesworth and Banaji, 2021). As we review below, stereotypes can be changed under certain circumstances, even in the short term. Our goal in this paper is to understand the principles governing these changes: why, when, and how do stereotypes change in response to experience?

The reputation for obduracy derives in part from studies showing that people often fail to update group stereotypes in response to counter-stereotypical individuals. Evidence suggests that this occurs because observers mentally segregate counter-stereotypical individuals into “subtypes” such that the stereotype is effectively protected from disconfirmation (see Hewstone, 1994, for a review).¹ Subtyping is employed when counter-stereotypical individuals are sufficiently deviant that they can be classified as outliers; by “deviant” we mean distance from the mode of the trait distribution. Under lower levels of deviance, counter-stereotypical individuals may be assimilated into the group, resulting in stereotype change. If multiple individuals share a common pattern of deviance, then they may be assimilated into a “subgroup” of the superordinate group (Maurer et al., 1995; Park and Judd, 1990; Park et al., 1992; Richards and Hewstone, 2001). In this case, the group stereotype undergoes a lesser degree of change, while still accommodating deviation within the subgroup.

While the concepts of subtyping and subgrouping have strongly shaped the study of stereotyping, a precise understanding remains elusive. Our goal in this paper is to formalize these concepts and incorporate them into a computational model of stereotype change. We will use this model to provide answers to the why, when, and how questions raised above. First, we introduce a “rational analysis” (cf. Anderson, 1990) that addresses *why* we have stereotypes. The answer, in brief, is that stereotypes enable probabilistic inferences about traits in the presence of unreliable individuating information, or even in the absence of individuating information altogether. This analysis also clarifies why stereotypes change: they must track the probability distribution of traits within a group.²

¹Allport (1954) used the term “refencing,” which he characterized as a mental device for holding onto prejudgments in the face of contradictory evidence. The first direct references to subtyping appear in Ashmore (1981), Brewer et al. (1981), and Taylor (1981).

²The concept of trait requires some clarification. We use traits to refer to *any* features of the targets under consideration, including behavior. Deviance may therefore arise both from features that are intrinsic to the target (e.g., counter-stereotypical personality traits, physical features) as well as those that are manifested in their behavior (e.g., neighborhoods they choose to live in, clothes they wear, protests they attend). This flexibility expands the modeling scope considerably.

Importantly, observing a counter-stereotypical individual does not necessarily indicate that the distribution has changed. It could alternatively indicate a new distribution (i.e., a new group), and hence generate a new stereotype while leaving the original stereotype intact. Thus, the answer to the *when* question is that stereotypes change when counter-stereotypical individuals are assimilated into an existing group. The phenomenon of subgrouping exposes a more nuanced answer to this question, whereby individuals can be *partially* assimilated into a group when they cohere with a subset of individuals in the group.

We will answer the *how* question in several steps, building from simple to more complex models of stereotyping. We begin with the setting in which group membership is observable. This allows us to formalize a principled probabilistic definition of stereotypes as beliefs about the distribution of traits conditional on group membership. Stereotype formation is then modeled as learning about the parameters governing the trait distribution for each group. We show that this model can account for several well-known aspects of stereotyping, including illusory correlation, accentuation, and outgroup homogeneity effects.

We then consider the setting in which group membership is unobservable and hence must be inferred—the major innovation of our paper. We show how this setting corresponds to a *structure learning* problem similar to the kind that has been studied in other cognitive domains (Austerweil et al., 2015). Structure learning refers to the acquisition of representations that organize domain-specific knowledge—the discovery of hypothesis spaces. For example, before we can learn about the perceptual similarities between colors or the biological relationships between animals, we need to learn that colors are organized in a circle and animals are organized in a taxonomic tree (Kemp and Tenenbaum, 2008). We apply this framework to the discovery of an important social hypothesis space: the organization of individuals into groups, what we refer to as *social structure learning* (Gershman and Cikara, 2020; Lau et al., 2020; Gershman et al., 2017b; Lau et al., 2018). We show how social structure learning can provide an account of subtyping and the various factors that moderate it. To formalize subgrouping, we extend the model to hierarchically structured groups, where subordinate groups can be nested within superordinate groups. The final result is a model that can explain how both the *structure* and *content* of stereotypes are acquired from experience.

Finally, a few words about our modeling approach are in order before proceeding. We have chosen to present a sequence of models rather than a single model that we use to simulate all the relevant empirical phenomena. This choice was guided by our goal of identifying minimal principles sufficient to explain a set of phenomena. These principles integrate coherently as we progress to more complex models, but we have chosen to present them in their simplest form in order to avoid post hoc (and possibly ambiguous) dissection of the more complex models into simpler constituents.

The function of stereotypes

The question of why we have stereotypes has exercised social psychologists since the earliest studies (see Snyder and Miene, 1994, for a review). Why should the mind equip itself with representations that contribute to biases in social perception and prejudice in intergroup attitudes? This question is not likely to be resolved by any single answer, because stereotypes play a multifaceted role in social cognition. A classic answer locates the function of stereotypes in their benefit to information processing economy: by simplifying the representation of individuals, fewer resource de-

mands are made on the observer (Allport, 1954; Tajfel, 1969; Macrae et al., 1994; Bodenhausen and Lichtenstein, 1987). We will return to this idea later in connection with our own account. Other answers locate the function of stereotypes in the maintenance and justification of social structures, such as social roles, power hierarchies, and coalitions: describing not only how certain groups are perceived but also prescribing how those groups *should* think, feel, and behave (Jost and Banaji, 1994; Koenig and Eagly, 2014; Cikara, 2021; Fiske and Neuberg, 1990). These structural functions can interact with cognitive functions; for example, more powerful individuals don't need to pay as much attention to less powerful individuals, and hence will rely more on stereotypes (Fiske, 1993). Reduced attention in turn reinforces the power hierarchy.

A third answer to the why question is also cognitive, but focuses on *statistical*, rather than information processing, constraints (Lee et al., 2013; McCauley et al., 1980). When we meet a new person, there is much that we do not know about them. However, we are not completely in the dark because unobserved traits may covary with observed traits. Thus, if we could learn the patterns of covariation, we could exploit them in the service of social inference. Stereotypes, on this view, correspond to beliefs about the covariation between a set of traits and a group label (a particular kind of trait). This view accords with several quantitative measures of stereotyping, such as the conditional probability that an individual has a particular trait given that they belong to a particular group (Brigham, 1971; Krueger, 1996), or the ratio between the conditional probability and the trait base rate (McCauley and Stitt, 1978). Empirical support for the statistical view comes from studies showing that observers rely more on covariation information when individuating information is absent, ambiguous, or uninformative (Krueger and Rothbart, 1988; Crawford et al., 2011; Nelson et al., 1990; Locksley et al., 1980). We will discuss other sources of empirical support in subsequent sections.

If stereotypes are to fulfill their statistical function, the group labels must be chosen such that traits can be effectively predicted. This implies that stereotypes should be accurate—a highly contentious proposition (Jussim et al., 2009). While studies have indicated that some stereotypes are moderately accurate (e.g., Chan et al., 2012; Diekmann et al., 2002; Rogers and Wood, 2010), they are often not, at which point they merely serve to induce systematic biases, such as distortions of perceived covariation and homogeneity. The existence of such biases does not, however, imply that stereotypes make irrational use of evidence. In fact, an ideal observer *must* be biased (Gershman, 2021), but note that we do not mean “bias” in the colloquial sense of exhibiting a *preference*. Because the data available to an observer are typically insufficient to completely disambiguate all unobserved traits, accurate trait inference requires an inductive bias that favors some inferences over others. This will inevitably produce systematic errors, despite reducing error on average. Although he did not employ the technical vocabulary for formalizing this idea, Allport (1954) recognized the necessity of inductive bias for statistical reasoning and prediction:

Open-mindedness is considered to be a virtue. But, strictly speaking, it cannot occur. A new experience must be redacted into old categories. We cannot handle each event freshly in its own right. (p. 27)

A primary goal of this paper is to place some mathematical flesh on the bones of Allport's insight.

In the next section, we introduce a simple Bayesian model that formalizes the statistical function of stereotypes. The purpose of this model is not to account for all aspects of stereotyping, but rather to establish the *prima facie* plausibility of the general approach, which we elaborate in subsequent sections to explain more complex aspects of stereotyping. The simple model will also serve a didactic purpose for those readers unfamiliar with Bayesian statistics.

Stereotype formation as Bayesian inference

An observer has access to information about N individuals, where the information about individual n is represented by a vector $\mathbf{x}_n = [x_{n1}, \dots, x_{nD}]$ consisting of D trait values. These vectors are collected into the $N \times D$ matrix $\mathbf{X}$, where each row corresponds to an individual. In addition to this matrix, the observer has access to each individual's group membership, represented by the binary vector $\mathbf{z}_n = [z_{n1}, \dots, z_{nK}]$, where K is the number of groups. For now, we will assume that each individual belongs to a single group, such that $z_{nk} = 1$ if individual n belongs to group k , and $z_{nk} = 0$ otherwise (we will relax this constraint later). The group membership vectors are collected into the $N \times K$ matrix $\mathbf{Z}$, where (like the trait matrix) each row corresponds to an individual.

In our model, a stereotype for group k corresponds to the conditional distribution over traits given membership in group k , $P(\mathbf{x}_n | z_{nk} = 1)$. This formalizes the widely cited definition given by Ashmore and Del Boca (1981): "A stereotype is a set of beliefs about the personal attributes of a social group" (p. 21).³ It is important to distinguish between the objective conditional distribution (which the observer does not access directly) and the observer's subjective *beliefs* about the conditional distribution. It is only the latter which defines a stereotype. The accuracy or inaccuracy of these beliefs lies outside the bounds of this inquiry.

From the observer's point of view, the observed data are assumed to be generated from some parametric distribution $P(\mathbf{x}_n | z_{nk} = 1, \theta)$ with unobserved parameters θ . We will refer to this subjective conditional distribution as the *observation model*, which may differ from the objective conditional distribution. For example, an observer might assume that heights (x) are drawn from a gender-specific distribution with an unknown mean (θ_k , where k indexes gender).

In order to form stereotypic beliefs, an ideal observer needs to *marginalize* over her uncertainty about the parameters:

$$P(\mathbf{x}_n | z_{nk} = 1) = \int_{\theta} P(\mathbf{x}_n | z_{nk} = 1, \theta) P(\theta | \mathbf{X}, \mathbf{Z}) d\theta. \quad (1)$$

Intuitively, this equation says that the observer weights each possible stereotype by the *posterior probability* of its parameter value. This posterior probability is the observer's belief about the stereotype parameters given her observations ($\mathbf{X}$ and $\mathbf{Z}$). Continuing the example from the previous paragraph, an observer who wants to predict the height of a man (without knowing any other individuating information) would weight each possible average height (θ_{male}) by its posterior probability conditional on the set of observed heights and genders.

Bayes' rule stipulates how to compute the posterior probability:

$$P(\theta | \mathbf{X}, \mathbf{Z}) \propto P(\mathbf{X} | \mathbf{Z}, \theta) P(\theta), \quad (2)$$

where $P(\theta)$ is the *prior probability* assigned by the observer to θ . This is the formal definition of inductive bias referred to in the previous section. Bayes' rule says that the prior should be combined multiplicatively with the *likelihood* of θ (the probability of the data conditional on θ) and renormalized to obtain the posterior probability. Thus, the observer's belief about the average height for men depends both on the likelihood of each hypothetical average height (how consistent is each average height with the observed data) and the observer's beliefs about average height before observing any data.

³See Greenwald and Banaji (1995) for a compendium of other definitions.

The Bayesian model sketched above can be understood as a formalization of the classic “book-keeping” model proposed by Rothbart (1981), according to which all evidence, both stereotype confirming and disconfirming, is assimilated into the estimated trait distribution. Hewstone (1994) described the bookkeeping model as a “feature-frequency” model in the same class as similar models used to describe non-social category learning (e.g., Fried and Holyoak, 1984). The representation of a stereotype corresponds to the set of sufficient statistics (feature frequencies in the case of discrete traits) for the trait distribution. As we will show below, this kind of model has broad explanatory power. Nonetheless, research on subtyping and subgrouping has called into question some of its basic assumptions. Later in this paper we will introduce a structure learning framework that addresses the deficiencies of the bookkeeping model.

A Gaussian model: application to illusory correlation

Here we assume that the observation model is parametrized as a factorized Gaussian, where the traits are conditionally independent given the group membership:

$$P(\mathbf{x}_n | z_{nk} = 1, \theta) = \prod_{d=1}^D \mathcal{N}(x_{nd}; \mu_{kd}, \sigma_{kd}^2), \quad (3)$$

with mean μ_{kd} and variance σ_{kd}^2 for group k and trait d . Initially, let us assume that there is a single trait ($D = 1$) and that the variances are known; therefore the unobserved parameters are the means, $\theta = [\mu_1, \dots, \mu_K]$, where we have dropped the trait index d for notational simplicity. If the prior over the means is Gaussian, $\mu_k \sim \mathcal{N}(m_k, \sigma_0^2)$, then the posterior after observing N_k individuals in group k is also Gaussian, with mean

$$\hat{\mu}_k = w_k \bar{x}_k + (1 - w_k) m_k, \quad (4)$$

where $\bar{x}_k$ is the trait value averaged across individuals in the group, and

$$w_k = \frac{1}{1 + \frac{\sigma_k^2}{N_k \sigma_0^2}} \quad (5)$$

is the weight attached to the observed data.

Eq. 4 shows that when the number of observations is small, the posterior mean is pulled toward μ_0 because the weight on the average trait value is smaller. Thus, if the mean of the prior is less than the average trait value ($m_k < \bar{x}_k$), the posterior mean $\hat{\mu}_k$ will always be less than the average trait value, but this bias diminishes with a larger sample size. This property is sufficient to reproduce several well-known aspects of stereotyping.

First, consider two groups (A and B) consisting of individuals with positive ($x_n > 0$) and negative ($x_n < 0$) traits. Both groups have an equal proportion of positive and negative traits represented across individuals (hence $\bar{x}_A = \bar{x}_B$), but group A is larger ($N_A > N_B$). According to Eq. 4, as long as $\bar{x}_A$ and $\bar{x}_B$ are greater than μ_0 , the “majority” group A will be *perceived* as more positive than the minority group B ($\hat{\mu}_A > \hat{\mu}_B$) despite having identical underlying parameters, a phenomenon known as *illusory correlation* (Hamilton and Gifford, 1976). The model also predicts that illusory correlation should be most pronounced at intermediate sample sizes, because when N_k is close to 0 the posterior mean will be dominated by the prior mean (same for both groups),

and when N_k gets very large the posterior mean will be dominated by the true mean (also the same for both groups). This prediction is consistent with the experimental results reported by Murphy et al. (2011), where absolute sample sizes were manipulated while holding the relative sample size fixed.

The model will produce illusory correlation even when the sample sizes for the two groups are equal, provided the prior means are unequal. Because the posterior mean will in general be biased towards the prior mean, identical trait averages will produce non-identical posterior means. This prediction is consistent with experimental results reported by Hamilton and Rose (1980), where prior expectations were manipulated while holding sample sizes fixed (see Spears et al., 1987, for similar results).

Inferences about individuals: application to accentuation

Accentuation refers to the exaggeration of between-group differences in the judgment of individuals (Tajfel and Wilkes, 1963; Krueger, 1992; McGarty and Penny, 1988; McGarty and Turner, 1992). For example, in the classic demonstration by Tajfel and Wilkes (1963), lines were categorized into two groups (corresponding to short and long lines). Compared to conditions in which labels were randomly assigned or no labels were assigned at all, the lengths of lines from different categories were perceived as more different.

Within the Bayesian framework laid out above, accentuation arises from the regularizing effect of the prior: trait inferences for group members are biased towards their associated group stereotype, simultaneously pulling them away from the other group’s stereotypes.⁴ It follows essentially the same logic that we used to explain illusory correlation, but now applied to inferences about individuals rather than groups.

To model inferences about individuals, we assume that memory of an individual’s trait vector is corrupted by Gaussian noise, yielding a noisy memory trace $\tilde{x}_n \sim \mathcal{N}(x, \tau^2)$, where τ is the perceptual noise standard deviation (for simplicity we continue to work with 1-dimensional trait vectors, but the theory generalizes straightforwardly to multiple dimensions). This leads to an expression for the posterior mean similar to Eq. 4:

$$\hat{x}_n = b_k \tilde{x}_n + (1 - b_k) \mu_k, \quad (6)$$

where k refers to the group membership of individual n , and

$$b_k = \frac{1}{1 + \frac{\tau^2}{\sigma_k^2}} \quad (7)$$

is the weight attached to the memory trace, a monotonically decreasing function of the ratio between the memory noise variance and within-group trait variance.⁵ One implication of this model is that memory recall of an individual’s traits should be biased towards the group mean, thereby accentuating between-group differences. The bias should weaken when memory is more unreliable. For example, accentuation is stronger in the Tajfel-Wilkes paradigm when the units of line

⁴Tajfel’s pioneering study of accentuation can be understood as a form of *categorical perception* (Harnad, 1987). Our Bayesian explanation of accentuation follows the same line of reasoning that has been applied to other categorical perception phenomena (Feldman et al., 2009).

⁵Note that we have assumed here direct access to the group mean and variance, though these could be estimated as in the previous section.

lengths are unfamiliar (Corneille et al., 2002), possibly because these units are more easily confusable in memory. A second implication of this model is that greater dispersion of traits within the group (higher σ_k) should reduce the bias towards the group mean. Indirect evidence comes from a study of racial stereotypes (Ryan et al., 1996), which found that high dispersion groups were perceived as less stereotypic (the degree to which a group is perceived to conform to the group stereotype), and stereotypicality predicted the bias towards the group stereotype in judgments of individuals.⁶

Eq. 6 assumes that the memory for each individual’s group membership is perfect, but evidence suggests that memory errors also occur for group membership, and that these errors dilute accentuation effects (Krueger and Rothbart, 1990). To capture unreliability of memory for group membership, we assume that people reconstruct membership from the noisy memory trace. Assuming a uniform prior over groups, the posterior is given by:

$$P(z_{nk} = 1 | \tilde{x}_n) \propto P(\tilde{x}_n | z_{nk} = 1) = \mathcal{N}(\tilde{x}_n; \mu_k, \sigma_k^2 + \tau^2). \quad (8)$$

The posterior mean is then obtained by marginalizing over group membership:

$$\hat{x}_n = \sum_k P(z_{nk} = 1 | \tilde{x}_n) [b_k \tilde{x}_n + (1 - b_k) \mu_k]. \quad (9)$$

One implication of Eq. 8 is that group membership errors during memory retrieval will be more likely for individuals in regions of overlap between stereotypes (i.e., near category boundaries), and therefore (by Eq. 9) accentuation effects will be weaker in these regions. Consistent with this hypothesis, Krueger and Rothbart (1990) showed that increasing the variance of the trait distribution resulted in more memory errors, and that excluding miscategorized individuals resulted in stronger accentuation effects.

Inferences about variability: application to the outgroup homogeneity effect

So far, we have assumed that the variances of the observation model were known; we now extend the model to inferences about variances.⁷ A simple form for the posterior mean can be obtained if we assume an uninformative (Jeffreys) prior over variance, according to which $P(\sigma_k^2) \propto 1/\sigma_k^2$:

$$\hat{\sigma}_k^2 = \frac{s^2}{1 + 2/N_k}, \quad (10)$$

where $s^2 = \frac{1}{N_k} \sum_n z_{nk} (x_n - \mu_k)^2$ is the sample variance. A key feature of this model is that the estimated variance is shrunk when the sample size is small. For a single sample ($N_k = 1$), the estimated variance is 1/3 of the sample variance. As N_k grows, the estimated variance becomes progressively closer to the sample variance.

⁶The results of this study are somewhat hard to interpret because perceptions of stereotypicality and dispersion were negatively correlated, raising the question of whether these are truly measuring different constructs. When they were simultaneously used as predictors of individual trait judgments, stereotypicality (but not dispersion) predicted bias towards the stereotype. However, the collinearity of the predictors means that some null effects could be type II errors.

⁷For simplicity, we will assume the means to be known, but it is straightforward to analyze the case where both means and variances are unknown.

This model offers an explanation of the *outgroup homogeneity effect*: outgroups tend to be viewed as more homogeneous than ingroups, even when their true variances are identical (Linville et al., 1989; Ostrom and Sedikides, 1992; Judd and Park, 1988; Quattrone and Jones, 1980; Park and Rothbart, 1982). Importantly, evidence suggests that this effect is not mediated by the ingroup-outgroup distinction *per se*, but rather by the differential amount of information about each group available to subjects (who necessarily must belong to either the ingroup or the outgroup). When the ingroup is smaller than the outgroup, the outgroup heterogeneity effect is reversed, with the ingroup now being perceived as more homogeneous (Simon and Brown, 1987; Simon and Pettigrew, 1990; Mullen and Hu, 1989). This is consistent with the idea that perception of variability tracks sample size. Eq. 10 formalizes this idea: the same sample variance translates to larger or smaller estimated variance depending on the group size.

As we will discuss later, inferences about variability are more complex than the picture developed above. For example, beliefs about subgroups play an important role in determining inferences about variability (Park et al., 1992; Kraus et al., 1993). A complete account will thus require us to explain the structure and origin of subgroups. We now develop such an account using principles of structure learning.

A structure learning model of stereotype change

The models developed thus far concern parameter learning: how do we build our representation of the trait distribution within a group? The parameters governing each distribution represent the content of the corresponding group stereotype. What such models do not tell us is where the groups come from in the first place. They do not solve the structure learning problem.

One answer to the structure learning problem is that groups consist of individuals who share a common group identity. This is consistent with the modeling of the previous section, where we assumed that individuals have access to group identity labels of all individuals, except in the context of memory-based judgments about individuals, where they have to reconstruct those labels from information stored in memory. There are three problems with this answer. First, it doesn't explain subtyping: why do we mentally segregate deviants even when they share a group identity label? Second, it doesn't explain subgrouping: why do we mentally construct hierarchically organized sets of groups when all the members share the same group identity label? Third, it doesn't explain how people reason about groups in the absence of explicit identity labels. For example, when police officers patrol college campuses they generate guesses about which people are students versus unauthorized visitors, in the absence of explicit labels. We argue that, irrespective of the (in)accuracy of their guesses, they accomplish this by inferring *latent groups* on the basis of observable features.

To realize this idea computationally, we extend the Bayesian model of stereotyping to incorporate uncertainty about the group membership matrix $\mathbf{Z}$. We will proceed in two steps, again building from simple to complex models. First, we will analyze the case where each individual can only belong to a single group. This will provide an account of subtyping. Then we will analyze the case where each individual can belong to multiple, hierarchically organized groups. This will provide an account of subgrouping.

Inferences about latent groups: application to subtyping

The *latent groups model* follows closely the setup of previous sections, but now the posterior is defined over both parameters (stereotype content) and group membership (stereotype structure):

$$P(\mathbf{Z}, \theta | \mathbf{X}) \propto P(\mathbf{X} | \mathbf{Z}, \theta) P(\theta) P(\mathbf{Z}), \quad (11)$$

where we have now introduced a prior over the membership matrix $\mathbf{Z}$. Recall that $z_{nk}=1$ if individual n belongs to group k , and 0 otherwise. In this section, we assume that each individual belongs to a single group, and hence each row contains a single 1. Because the number of groups is unknown *a priori*, we define a prior over membership matrices with an unbounded number of columns. A standard nonparametric prior for such matrices is the *Chinese restaurant process* (Aldous, 1985; Gershman and Blei, 2012):

$$P(z_{nk} = 1) = \begin{cases} \frac{g_{nk}}{n-1+\alpha}, & k \leq K_n \\ \frac{\alpha}{n-1+\alpha}, & k = K_n + 1, \end{cases} \quad (12)$$

where $g_{nk} = \sum_{j=1}^{n-1} z_{jk}$ is the number of individuals assigned to group k prior to n , K_n is the number of unique groups created prior to n , and $\alpha \geq 0$ is a concentration parameter that probabilistically controls the number of groups. When $\alpha = 0$, all individuals are assigned to the same group; in the limit $\alpha \rightarrow \infty$, all individuals are assigned to their own group. The expected number of groups $\mathbb{E}[K_n]$ scales according to $\alpha \log n$.

Variants of this model have been widely applied in cognitive science (see Austerweil et al., 2015, for a review). Most relevant for present purposes is the model first presented in Gershman et al. (2017b) to analyze patterns of social influence. According to that model, social influence between individuals is stronger to the extent that they believe they belong to the same group. Membership is inferred on the basis of observed choices. We subsequently applied this idea to real-world political attitudes (Lau et al., 2018) and used it to understand the neural correlates of social influence (Lau et al., 2020). In the same vein, we assume that observers use trait data (which may include an individual’s preferences) to infer latent group membership. They then use these inferences about group membership to structure their inferences about the group trait distribution (i.e., the stereotype), following the logic of the previous sections.

To obtain the posterior over group membership, we marginalize over the parameters θ :

$$\begin{aligned} P(\mathbf{Z} | \mathbf{X}) &= \int_{\theta} P(\mathbf{Z}, \theta | \mathbf{X}) d\theta \\ &\propto \mathcal{N}(\mathbf{X}; \boldsymbol{\mu}, \sigma_0^2 \mathbf{Z} \mathbf{Z}^\top + \sigma^2 \mathbf{I}) P(\mathbf{Z}), \end{aligned} \quad (13)$$

where we have assumed for simplicity that all groups share a common prior ($\boldsymbol{\mu} = \mathbf{0}$ in our applications) and that the prior and likelihood covariances are isotropic (i.e., different traits are uncorrelated with each other conditional on group membership).⁸ We have also assumed here that the variances are known, though it is possible to extend the model to inferences about variances, as we did above.

Finding the group assignment with highest posterior probabilities requires an intractable search in the space of group assignments. In practice, most of these assignments have negligible probability. We therefore enumerate a small set of plausible group assignments that have relatively high

⁸The isotropic covariance assumption is unlikely to be true in general, and is not mathematically necessary, but makes the model presentation simpler and is sufficient to account for the phenomena we address here.

Figure 1: Simulation of dispersion, variability, and sample size effects on subtyping. (A) Probability that a deviant individual is assigned to the group as a function of whether perceived variability is high or low and whether the counter-stereotypical trait are concentrated in a single individual or dispersed across multiple individuals. (B) Degree of change in the prediction of trait values for the group after observing the deviant. (C) Degree of change as a function of sample size and dispersion.

probability and score these using Eq. 13. To obtain parameter estimates, we marginalize over this plausible set (denoted by $\mathcal{Z}$):

$$P(\theta|\mathbf{X}) = \sum_{\mathbf{Z} \in \mathcal{Z}} P(\mathbf{Z}, \theta|\mathbf{X}). \quad (14)$$

Because we are interested in qualitative correspondences between the model and empirical data, we do not fit the parameters of the model. Instead, we fix the parameters across all simulations except those where they are explicitly manipulated. This allows us to demonstrate the model's breadth of explanatory power without parameter tuning. We used the following parameter values, excepted where noted otherwise: $\sigma_0^2 = 2$, $\sigma^2 = 0.1$, $\alpha = 10$.

Simulations

One of the most important and well-documented findings in the subtyping literature is the effect of dispersion: stereotype change is greater when stereotype-inconsistent information is dispersed across multiple individuals, compared to when it is concentrated in a single individual (Weber and Crocker, 1983; Johnston and Hewstone, 1992; Hewstone and Hamberger, 2000; Hantzi, 1995; Johnston et al., 1994). This finding has traditionally been interpreted as evidence that strongly deviant individuals in the concentrated condition are mentally segregated into subtypes, which allows the existing stereotype associated with the rest of the group to remain intact, whereas the weakly deviant individuals in the dispersed condition are assimilated into the group, driving change in the stereotype.

The latent groups model reproduces the dispersion effect (Figure 1). In the dispersed condition, the model places more probability on a single latent group for all individuals (including

the deviant) compared to in the concentrated condition, where it favors segregating the deviant into a separate group. Importantly, the model also explains several moderating factors of the dispersion effect. First, the dispersion effect is much stronger when variability is low (Hewstone and Hamberger, 2000); see panel A of Figure 1. We manipulated dispersion by stretching out the range of attribute values in the observed population. The interaction between dispersion and variability occurs in the model because high variability makes the stereotype more “tolerant” of the deviant, even in the concentrated condition, thereby diluting the contrast between dispersed and concentrated conditions. As predicted, we observe the most stereotype change in the low-variability/dispersed condition (Figure 3, panel B). Second, the dispersion effect is weaker when the sample size is larger (Weber and Crocker, 1983); see panel C of Figure 3. In the model, a larger sample size increases the probability that the deviants will be subtyped even in the dispersed condition, because there is greater certainty about the trait distribution and hence the stereotype is less tolerant of the deviants.

Another determinant of subtyping is deviants’ typicality; individuals who are counter-stereotypical on one dimension but are otherwise more typical of their group produce more stereotype change than more atypical individuals (Hewstone and Hamberger, 2000; Hewstone et al., 2000; Johnston and Hewstone, 1992; Hantzi, 1995; Rothbart and John, 1985; Wilder, 1984). A frequently cited example is that White, middle-class, high-earning, but introverted lawyers bring about more lawyer-related stereotype change compared to introverted black lawyers (Weber and Crocker, 1983). There is some controversy around whether typicality mediates stereotype change or is simply correlated with it (see Richards and Hewstone, 2001, for discussion). For our purposes, the important thing to note is that according to our model, higher typicality makes it more likely that a deviant will be assimilated into the group and hence drive stereotype change. In fact, even a “neutral” (i.e., unrelated) trait can contribute to perceived atypicality Kunda and Oleson (1995). In one of their studies, Kunda and Oleson presented subjects with information about an introverted lawyer. One set of subjects was given additional information (that the lawyer worked for a large or small firm) which separate pretests had established to be neutral with respect to introversion. Subjects who received no additional information updated their stereotypes about lawyers more (i.e., they reported viewing lawyers in general as being higher in introversion) compared to subjects who also received the neutral information. Our model explains this finding in terms of the same mechanism producing typicality effects in subtyping: anything that makes an individual unique relative to the group also reduces the probability of their membership, “fencing” them off, thereby suppressing stereotype change. Figure 2 presents the results of simulations that capture both findings. Probability of group membership increases with a typical trait and decreases with a neutral trait, relative to no trait. Stereotype change tracks these differences in probability.

Kunda and Oleson (1997) investigated the impact of deviance magnitude (along a single trait) on stereotype change, finding that moderate deviants produced more stereotype change than extreme deviants (see also Dannals and Miller, 2017). Our model recapitulates this finding (Figure 3). The probability of assimilation into the group decreases monotonically with deviance magnitude, but this has a non-monotonic effect on stereotype change. At very low deviance, membership probability is high but since deviance is also small there is little change in the stereotype. At very high deviance, membership probability is low, weakening the impact of the deviant. Stereotype change is maximized in between these two extremes, at moderate deviance magnitudes. This explanation mirrors structure learning accounts that have been proposed in other domains (Gershman et al., 2013, 2017a).

Figure 2: Simulation of typicality effects on subtyping. (A) Probability that a deviant individual is assigned to the group as a function of whether the deviant trait type is absent, neutral, or typical. (B) Degree of change in the prediction of trait values for the group after observing the deviant.

Figure 3: Simulation of deviance effect on subtyping. (A) Probability that a deviant individual is assigned to the group as a function of deviance magnitude. (B) Degree of change in the prediction of trait values for the group after observing the deviant.

Figure 4: Simulation of cognitive load effects on subtyping. (A) Probability that a deviant individual is assigned to the group as a function of the concentration parameter α . Larger values of alpha tend to produce more complex group structures and hence hypothetically demand more cognitive resources. Cognitive load is hypothesized to push the concentration parameter lower, thereby increasing the probability of assigning the deviant to the group. (B) Degree of change in the prediction of trait values for the group after observing the deviant. If the concentration parameter decreases under cognitive load, then stereotype change will be greater.

Finally, we explore the effects of cognitive load. Several studies have found reduced subtyping, and therefore greater stereotype change in the direction of the disconfirming individual, under cognitive load (Moreno and Bodenhausen, 1999; Yzerbyt et al., 1999). These studies were motivated by the idea that stereotypes serve an economizing function by simplifying mental computation (Allport, 1954; Tajfel, 1969; Macrae et al., 1994; Bodenhausen and Lichtenstein, 1987). The effect of load on subtyping was interpreted as evidence that constructing richer representations of social groups demands greater cognitive resources. Our model sheds light on the mechanistic nature of the economizing function of stereotype. In particular, fewer latent groups place smaller demands on mental computation because they require the observer to marginalize over a smaller number of hypotheses. We can operationalize cognitive demand in terms of the concentration parameter α , which controls the expected complexity of stereotype representations (Figure 4). Recently, Dasgupta and Griffiths (2021) have suggested a specific link between the concentration parameter and cognitive demand measured information-theoretically. The quantitative predictions of this model for stereotyping remain untested.

Inferences about hierarchically organized groups: application to subgrouping

We now describe a generalization of the latent groups model to hierarchically structured groups, which we represent as a tree structure (Figure 5). Each node in the tree picks out a collection of individuals, defining a subgroup relative to the superordinate node. Each individual is identified by a path through the tree. For example, the root node could correspond to “all people,” and below it are nodes corresponding to “men” and “women.” Below each of these are nodes corresponding

Figure 5: **Schematic of the hierarchical latent groups model.** (Top left) The observation matrix, where grayscale values denote the magnitude of a trait for a given individual. (Top right) Illustration of a grouping structure, where one subgroup is nested within a superordinate group, and one individual is fenced off as a subtype. (Bottom) Application of Bayes' rule to structure learning. The posterior distribution over structures is obtained by multiplying the prior probability of each structure with the likelihood of the observations under that structure and then renormalizing. The height of each bar denotes the probability of the corresponding structural hypothesis, shown as a set of trees. Some of these trees are "degenerate" (only a single level deep), in which case they correspond to flat (non-hierarchical) groups. The individuals assigned to each node are shown in brackets. Note that although the prior on the hierarchically organized group is low, the likelihood given the data is highest, giving rise to a posterior that favors the nested group structure.

to “old” and “young.” An old man would then correspond to the path people→men→old.

Each node indexes a trait distribution, just as in our treatment of flat clusters in the previous section. Let $\text{pa}(k)$ denote the parent of node k in the tree structure. The mean for group k is sampled from a Gaussian with mean $\mu_{\text{pa}(k)}$ and covariance $\sigma_0^2 \mathbf{I}$. Thus, subgroups will tend to share traits with their superordinate groups, and this similarity will decrease as a function of distance in the tree. We will assume that there is a root node in the tree ($k = 0$) that defines the prior for any groups that are not subgroups, with mean $\mu_0 = \mathbf{0}$ and covariance $\sigma_0^2 \mathbf{I}$.

It will be convenient for us to describe the inferences in terms of the *relative* group-level mean, $\tilde{\mu}_k = \mu_k - \mu_{\text{pa}(k)}$ (i.e., how much more or less of each trait is a subgroup associated with relative to its superordinate group). For a given tree structure, the posterior expectation of the relative mean for group k is given by:

$$\mathbb{E}[\tilde{\mu} | \mathbf{X}, \mathbf{Z}] = \left(\mathbf{Z}^\top \mathbf{Z} + \frac{\sigma^2}{\sigma_0^2} \mathbf{I} \right)^{-1} \mathbf{Z}^\top \mathbf{X}; \quad (15)$$

The group-level means μ_k can be obtained by summing the relative means along the path in the tree leading to group k .

We define a prior distribution over trees known as the *nested Chinese restaurant process* (nCRP; Blei et al., 2010), which has previously been applied to modeling hierarchical category learning (Canini and Griffiths, 2011) and perception (Gershman et al., 2016). The basic building block is the Chinese restaurant process, defined in Eq. 12. In the nCRP, this process recurses to some depth d_n , so that each individual is assigned to d_n groups, defining a path through the tree. We place a uniform distribution over depths up to 2 (in principle, there is nothing stopping us from considering deeper depths, but this was not necessary for our simulations).

Simulations

In the first study to directly compare subtyping and subgrouping, Maurer et al. (1995) showed that subjects given subtyping instructions (distinguishing stereotype-consistent individuals from stereotype-inconsistent individuals) viewed the group as more stereotypical and homogeneous compared to subjects given subgrouping instructions (sorting individuals into multiple groups based on their similarities and differences). Subtyping instructions also led subjects to perceive a greater difference in typicality between confirming and disconfirming individuals. In a subsequent study using the Weber and Crocker (1983) dispersed vs. concentrated manipulation, Hewstone and Hamberger (2000) showed that subgrouping instructions eliminated the difference between dispersed and concentrated conditions, instead showing increased stereotype change in both conditions. Figure 6 shows a simulation of this finding, which the model captures by allowing deviants to be assimilated into subgroups (thereby leading to change in the superordinate stereotype) rather than segregated into subtypes (which prevents superordinate stereotype change).

The study of subgrouping led to an important insight into the origin of the outgroup homogeneity effect. Park et al. (1992) showed that perceived variability is closely linked to the number of subgroups that subjects generate, and that more subgroups were generated for ingroups compared to outgroups (see also Kraus et al., 1993). Indeed, the outgroup homogeneity effect was eliminated entirely when the difference in number of generated subgroups was controlled for. If we start from the assumption that a key difference between ingroups and outgroups is sample

Figure 6: Simulation of dispersion effects on subgrouping. (A) Probability that a deviant individual is assigned to the group as a function of whether subgroups were invoked during the categorization of group members. (B) Degree of change in the prediction of trait values for the superordinate group after observing the deviant.

Figure 7: Probability that a deviant individual is assigned to the group as a function of sample size.

size (people are exposed to more individuals in the ingroup than in the outgroup), we can ask how sample size contributes to the formation of subgroups. Our model provides an answer to this question: the number of subgroups increases with sample size. A corollary of this property is that the probability of assigning a new individual to the superordinate group increases with sample size (Figure 7), because as subgroups proliferate it becomes increasingly easier to assimilate a new individual into one of them. The expanding diversity of the superordinate group renders it more tolerant of deviants.

General discussion

We have presented a sequence of increasingly sophisticated stereotyping models, culminating in an account of hierarchically organized latent groups. This model captures many important phenomena in the literature on stereotypes and stereotype change: illusory correlation, accentuation, outgroup homogeneity, subtyping, and subgrouping, as well as the effects of moderating factors such as sample size, variability, and cognitive load. The latter two phenomena are explained by the model as the result of structure learning mechanisms, a distinctive feature of our account. Subtyping arises from the segregation of deviant individuals from the group, preventing them from driving stereotype change. This is a formalization of “refencing” as proposed by Allport (1954) and many subsequent researchers. Subgrouping arises from the assimilation of deviant individuals into a subordinate category that shares some features with the superordinate category. This allows deviants to drive stereotype change, because the subordinate category exerts a pull on the overall stereotype representation.

We now turn to a discussion of other models and the limitations of our model, with an eye towards future directions of research.

Comparison to other models

Several other models of stereotyping have been developed which explain some of the phenomena discussed here. In this section, we compare and contrast these models.

Early models of stereotyping can be categorized into one of two classes: exemplar models and abstraction models (see Linville and Fischer, 1993, for a review). Exemplar models take a similarity function defined over individual trait vectors as their basic primitive (Linville et al., 1989; Fiedler, 1996; Smith and Zárate, 1992). Inductive reasoning and generalization operate through similarity computation. For example, predicting whether an individual has a particular trait is computed by generalizing from other similar individuals. Likewise, inferences at the group level are computed by generalizing from individuals that belong to a particular group (i.e., the group identity enters into the similarity computation). Abstraction models, by contrast, typically take trait distributions as their basic primitive, abstracting over individual exemplars (Kraus et al., 1993; Park and Judd, 1990; Nisbett and Kunda, 1985; Bordalo et al., 2016).

A second wave of stereotyping models adopted the connectionist framework, where inductive reasoning and generalization emerge from the dynamics of spreading activation between elementary processing units. In some connectionist models (Van Rooy et al., 2003; Vanhooymissen and Van Overwalle, 2010; Kunda and Thagard, 1996) the units are “localist” in the sense that they represent distinct symbolic variables (e.g., particular traits), while in other models the representation of variables is distributed across multiple units (Smith and Decoster, 1998; Queller and Smith,

2002; Kashima et al., 2000), so that the collection of traits for an individual is represented algebraically as a matrix. All of these models have in common the property that learning is driven by some form of mutual constraint satisfaction, although the constraints, learning rules, and dynamics differ between the models.

All of the connectionist models also have in common the property that memory for individual exemplars is discarded, in contrast to exemplar models in which the trait vectors for all exemplars are stored in memory (possibly corrupted by noise). So in that respect the connectionist models are similar to abstraction models, except that they do not explicitly represent trait distributions. Rather, they store some internal encoding of the observed trait vectors, for example through error-driven (Van Rooy et al., 2003) or associative (Kashima et al., 2000) learning. This means that the connectionist models cannot directly answer queries about trait distributions (e.g., perceived variability), which figure prominently in the experimental literature. Bayesian models specify a “natively probabilistic” representation that can directly answer queries about trait distributions.

Exemplar models can also answer queries about trait distributions, by computing statistics on the distribution of exemplars stored in memory. However, it has been argued that certain experimental observations are problematic for exemplar models. First, an exemplar model can only learn from information about exemplars, yet people are able to learn about groups from abstract statements (Park and Hastie, 1987). Moreover, Park and Hastie argued that subjects in their studies were reporting perceived variability based on abstract information about feature frequency rather than specific exemplars, because their reports were unaffected by making memory for some exemplars more or less accessible. Second, an outgroup homogeneity effect can sometimes be observed even when subjects are exposed to equivalent numbers of exemplars from both the ingroup and the outgroup, simply by invoking a competitive context (Judd and Park, 1988). This finding is problematic for the Bayesian model we presented as well, and speaks to a fundamental limitation of models that are based on purely statistical information (see next section).

Our work is most closely related to abstraction models that have been studied in the categorization literature, especially Anderson’s rational model of categorization (Anderson, 1991; Sanborn et al., 2010), which also uses the Chinese restaurant prior over a set of underlying “clusters.” To our knowledge, however, this paper is the first to directly apply this kind of model to the stereotyping literature. A key innovation of this kind of model relative to the other models reviewed above is that there is provision for creation of novel groups. The other models generally assume a fixed number of groups and lack a mechanism for creating new ones. We have argued that group creation occurs through Bayesian structure learning.

Limitations and future directions

Following a venerable tradition in social psychology (e.g., Taylor, 1981; Brewer et al., 1981), our treatment of group stereotypes conceptualized them as a form of category knowledge. Consequently, our modeling assumptions, explanations and predictions mirrored those of models applied to non-social domains (Anderson, 1991; Sanborn et al., 2010; Fried and Holyoak, 1984; Feldman et al., 2009). However, we should not lose sight of the fact that social groups are special, in that they serve functions that do not exist in non-social domains. For example, as recently emphasized by Cikara (2021), group identification—figuring out who belongs to “us” versus “them”—is critical to intra-group cooperation and inter-group competition. Individuals form coalitions, in part, based on their beliefs about group identity, which are in turn strengthened by the resulting coalitions. These identities may or may not map onto the categories (e.g., race, gender) people

sometimes use to organize their beliefs about social groups. Thus, flexible beliefs about group identity are both the input and the output of affiliative social behavior. This implies that our explanations of social structure learning will be incomplete as long as they are based on purely non-social categorization processes. We see the integration of distinctively social structures and processes into our theoretical framework as an important direction for future work.

Another limitation of our model is that the structure learning mechanism does not take into account explicit group labels. These explicit labels may provide constraints on the structure of stereotypes that get integrated with unsupervised structure learning, in a manner similar to models of “semi-supervised” learning (Gibson et al., 2013; Vong et al., 2016).

Finally, the computational intractability of Bayesian structure learning necessitated approximations, but we have not here committed to a particular process model of human approximate inference. This is a rich topic of active research (see Gershman and Beck, 2017), and several possibilities are viable. For example, people might try to identify a single high probability point estimate of the latent structure, and thus ignore their uncertainty. Alternatively, people might use a Monte Carlo strategy in which they sample structures in proportion to their posterior probability. The empirical literature at present does not seem to place strong constraints on which of these approximate inference strategies is most plausible.

Conclusion

Our model seeks to answer the why, when, and how questions of stereotype change as follows. Stereotypes change in order to track the distribution of traits in a group. This happens when the underlying parameters governing the trait distribution are inferred to have changed. In some cases, exposure to stereotype-inconsistent traits does not drive stereotype change because deviant individuals are assigned to new groups. The cognitive engine underlying both parameter and structure learning is Bayesian inference. Though a number of questions remain, understanding the interplay between stereotype content (parameters) and stereotype organization (structure) is an important step towards a complete theory of stereotype change.

References

- Aldous, D. J. (1985). Exchangeability and related topics. In *École d’Été de Probabilités de Saint-Flour XIII*, pages 1–198. Springer.
- Allport, G. W. (1954). *The Nature of Prejudice*. Addison-Wesley.
- Anderson, J. R. (1990). *The Adaptive Character of Thought*. Lawrence Erlbaum Associates.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98:409–429.
- Ashmore, R. D. (1981). Sex stereotypes and implicit personality theory. In Hamilton, D., editor, *Cognitive Processes in Stereotyping and Intergroup Behavior*, pages 37–81. Erlbaum.
- Ashmore, R. D. and Del Boca, F. K. (1981). Conceptual approaches to stereotypes and stereotyping. In Hamilton, D., editor, *Cognitive Processes in Stereotyping and Intergroup Behavior*, pages 1–35. Erlbaum.

- Austerweil, J. L., Gershman, S. J., Tenenbaum, J. B., and Griffiths, T. L. (2015). Structure and flexibility in Bayesian models of cognition. *Oxford Handbook of Computational and Mathematical Psychology*, pages 187–208.
- Banaji, M. R. (2002). Stereotypes, social psychology of. In Smelser, N. J. and Baltes, P. B., editors, *International Encyclopedia of the Social and Behavioral Sciences*, pages 15100–15104. Pergamon New York.
- Bergsieker, H., Leslie, L., Constantine, V., and Fiske, S. (2012). Stereotyping by omission: eliminate the negative, accentuate the positive. *Journal of Personality and Social Psychology*, 102:1214–1238.
- Blei, D. M., Griffiths, T. L., and Jordan, M. I. (2010). The nested Chinese restaurant process and Bayesian nonparametric inference of topic hierarchies. *Journal of the ACM*, 57:1–30.
- Bodenhause, G. and Lichtenstein, M. (1987). Social stereotypes and information-processing strategies: The impact of task complexity. *Journal of Personality and Social Psychology*, 52:871–880.
- Bordalo, P., Coffman, K., Gennaioli, N., and Shleifer, A. (2016). Stereotypes. *The Quarterly Journal of Economics*, 131:1753–1794.
- Brewer, M., Dull, V., and Lui, L. (1981). Perceptions of the elderly: Stereotypes as prototypes. *Journal of Personality and Social Psychology*, 41:656–670.
- Brigham, J. (1971). Ethnic stereotypes. *Psychological Bulletin*, 76:15–38.
- Canini, K. R. and Griffiths, T. L. (2011). A nonparametric Bayesian model of multi-level category learning. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*.
- Chan, W., McCrae, R. R., De Fruyt, F., Jussim, L., Löckenhoff, C. E., De Bolle, M., Costa Jr, P. T., Sutin, A. R., Realo, A., Allik, J., et al. (2012). Stereotypes of age differences in personality traits: Universal and accurate? *Journal of Personality and Social Psychology*, 103:1050–1066.
- Charlesworth, T. E. and Banaji, M. R. (2021). Patterns of implicit and explicit stereotypes III: Long-term change in gender stereotypes. *Social Psychological and Personality Science*, pages 1–13.
- Cikara, M. (2021). Causes and consequences of coalitional cognition. *Advances in Experimental Social Psychology*.
- Corneille, O., Klein, O., Lambert, S., and Judd, C. M. (2002). On the role of familiarity with units of measurement in categorical accentuation: Tajfel and Wilkes (1963) revisited and replicated. *Psychological Science*, 13:380–383.
- Crawford, J. T., Jussim, L., Madon, S., Cain, T. R., and Stevens, S. T. (2011). The use of stereotypes and individuating information in political person perception. *Personality and Social Psychology Bulletin*, 37:529–542.
- Dannals, J. E. and Miller, D. T. (2017). Social norm perception in groups with outliers. *Journal of Experimental Psychology. General*, 146:1342–1359.

- Dasgupta, I. and Griffiths, T. L. (2021). Clustering and the efficient use of cognitive resources. *PsyArXiv*.
- Diekmann, A. B., Eagly, A. H., and Kulesa, P. (2002). Accuracy and bias in stereotypes about the social and political attitudes of women and men. *Journal of Experimental Social Psychology*, 38:268–282.
- Eagly, A., Nater, C., Miller, D., Kaufmann, M., and Sczesny, S. (2020). Gender stereotypes have changed: A cross-temporal meta-analysis of us public opinion polls from 1946 to 2018. *American Psychologist*, 75:301–315.
- Feldman, N. H., Griffiths, T. L., and Morgan, J. L. (2009). The influence of categories on perception: Explaining the perceptual magnet effect as optimal statistical inference. *Psychological Review*, 116:752–782.
- Fiedler, K. (1996). Explaining and simulating judgment biases as an aggregation phenomenon in probabilistic, multiple-cue environments. *Psychological Review*, 103:193–214.
- Fiske, S. (1993). Controlling other people: The impact of power on stereotyping. *American Psychologist*, 48:621–628.
- Fiske, S. T. and Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. *Advances in Experimental Social Psychology*, 23:1–74.
- Fried, L. S. and Holyoak, K. J. (1984). Induction of category distributions: A framework for classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10:234–257.
- Gershman, S. J. (2021). *What Makes Us Smart*. Princeton University Press.
- Gershman, S. J. and Beck, J. M. (2017). Complex probabilistic inference. In Moustafa, A., editor, *Computational Models of Brain and Behavior*, pages 453–464. Wiley Hoboken, NJ.
- Gershman, S. J. and Blei, D. M. (2012). A tutorial on Bayesian nonparametric models. *Journal of Mathematical Psychology*, 56:1–12.
- Gershman, S. J. and Cikara, M. (2020). Social-structure learning. *Current Directions in Psychological Science*, 29:460–466.
- Gershman, S. J., Jones, C. E., Norman, K. A., Monfils, M.-H., and Niv, Y. (2013). Gradual extinction prevents the return of fear: implications for the discovery of state. *Frontiers in Behavioral Neuroscience*, 7:164.
- Gershman, S. J., Monfils, M.-H., Norman, K. A., and Niv, Y. (2017a). The computational nature of memory modification. *Elife*, 6:e23763.
- Gershman, S. J., Pouncy, H. T., and Gweon, H. (2017b). Learning the structure of social influence. *Cognitive Science*, 41:545–575.
- Gershman, S. J., Tenenbaum, J. B., and Jäkel, F. (2016). Discovering hierarchical motion structure. *Vision Research*, 126:232–241.

- Gibson, B. R., Rogers, T. T., and Zhu, X. (2013). Human semi-supervised learning. *Topics in Cognitive Science*, 5:132–172.
- Greenwald, A. G. and Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, 102:4–27.
- Hamilton, D. and Rose, T. (1980). Illusory correlation and the maintenance of stereotypic beliefs. *Journal of Personality and Social Psychology*, 39:832–845.
- Hamilton, D. L. and Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology*, 12:392–407.
- Hantzi, A. (1995). Change in stereotypic perceptions of familiar and unfamiliar groups: The pervasiveness of the subtyping model. *British Journal of Social Psychology*, 34:463–477.
- Harnad, S. (1987). *Categorical Perception*. Cambridge University Press.
- Hewstone, M. (1994). Revision and change of stereotypic beliefs: In search of the elusive subtyping model. *European Review of Social Psychology*, 5:69–109.
- Hewstone, M. and Hamberger, J. (2000). Perceived variability and stereotype change. *Journal of Experimental Social Psychology*, 36:103–124.
- Hewstone, M., Hassebrauck, M., Wirth, A., and Waenke, M. (2000). Pattern of disconfirming information and processing instructions as determinants of stereotype change. *British Journal of Social Psychology*, 39:399–411.
- Johnston, L. and Hewstone, M. (1992). Cognitive models of stereotype change: 3. subtyping and the perceived typicality of disconfirming group members. *Journal of Experimental Social Psychology*, 28:360–386.
- Johnston, L., Hewstone, M., Pendry, L., and Frankish, C. (1994). Cognitive models of stereotype change (4): Motivational and cognitive influences. *European Journal of Social Psychology*, 24:237–265.
- Jost, J. T. and Banaji, M. R. (1994). The role of stereotyping in system-justification and the production of false consciousness. *British Journal of Social Psychology*, 33:1–27.
- Judd, C. and Park, B. (1988). Out-group homogeneity: Judgments of variability at the individual and group levels. *Journal of Personality and Social Psychology*, 54:778–788.
- Jussim, L., Cain, T. R., Crawford, J. T., Harber, K., and Cohen, F. (2009). The unbearable accuracy of stereotypes. In Nelson, T. D., editor, *Handbook of Prejudice, Stereotyping, and Discrimination*, pages 199–227. Psychology Press.
- Kashima, Y., Woolcock, J., and Kashima, E. (2000). Group impressions as dynamic configurations: The tensor product model of group impression formation and change. *Psychological Review*, 107:914–942.
- Kemp, C. and Tenenbaum, J. B. (2008). The discovery of structural form. *Proceedings of the National Academy of Sciences*, 105:10687–10692.

- Koenig, A. M. and Eagly, A. H. (2014). Evidence for the social role theory of stereotype content: Observations of groups' roles shape stereotypes. *Journal of Personality and Social Psychology*, 107:371–392.
- Kraus, S., Ryan, C. S., Judd, C. M., Hastie, R., and Park, B. (1993). Use of mental frequency distributions to represent variability among members of social categories. *Social Cognition*, 11:22–43.
- Krueger, J. (1992). On the overestimation of between-group differences. *European Review of Social Psychology*, 3:31–56.
- Krueger, J. (1996). Probabilistic national stereotypes. *European Journal of Social Psychology*, 26:961–980.
- Krueger, J. and Rothbart, M. (1988). Use of categorical and individuating information in making inferences about personality. *Journal of Personality and Social Psychology*, 55:187–195.
- Krueger, J. and Rothbart, M. (1990). Contrast and accentuation effects in category learning. *Journal of Personality and Social Psychology*, 59:651–663.
- Kunda, Z. and Oleson, K. C. (1995). Maintaining stereotypes in the face of disconfirmation: Constructing grounds for subtyping deviants. *Journal of Personality and Social Psychology*, 68:565–579.
- Kunda, Z. and Oleson, K. C. (1997). When exceptions prove the rule: How extremity of deviance determines the impact of deviant examples on stereotypes. *Journal of Personality and Social Psychology*, 72:965–979.
- Kunda, Z. and Thagard, P. (1996). Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological Review*, 103:284–308.
- Lau, T., Gershman, S. J., and Cikara, M. (2020). Social structure learning in human anterior insula. *Elife*, 9:e53162.
- Lau, T., Pouncy, H. T., Gershman, S. J., and Cikara, M. (2018). Discovering social groups via latent structure learning. *Journal of Experimental Psychology: General*, 147:1881–1891.
- Lee, Y.-T., McCauley, C., and Jussim, L. (2013). Stereotypes as valid categories of knowledge and human perceptions of group differences. *Social and Personality Psychology Compass*, 7:470–486.
- Linville, P., Fischer, G., and Salovey, P. (1989). Perceived distributions of the characteristics of in-group and out-group members: empirical evidence and a computer simulation. *Journal of personality and social psychology*, 57:165–188.
- Linville, P. W. and Fischer, G. W. (1993). Exemplar and abstraction models of perceived group variability and stereotypicality. *Social Cognition*, 11:92–125.
- Lippmann, W. (1922). *Public Opinion*. Harcourt, Brace, and Company.
- Locksley, A., Borgida, E., Brekke, N., and Hepburn, C. (1980). Sex stereotypes and social judgment. *Journal of Personality and Social psychology*, 39:821–831.
- Macrae, C., Milne, A., and Bodenhausen, G. (1994). Stereotypes as energy-saving devices: A peek inside the cognitive toolbox. *Journal of Personality and Social Psychology*, 66:37–47.

- Maurer, K., Park, B., and Rothbart, M. (1995). Subtyping versus subgrouping processes in stereotype representation. *Journal of Personality and Social Psychology*, 69:812–824.
- McCauley, C. and Stitt, C. (1978). An individual and quantitative measure of stereotypes. *Journal of Personality and Social Psychology*, 36:929–940.
- McCauley, C., Stitt, C., and Segal, M. (1980). Stereotyping: From prejudice to prediction. *Psychological Bulletin*, 87:195–208.
- McGarty, C. and Penny, R. (1988). Categorization, accentuation and social judgement. *British Journal of Social Psychology*, 27:147–157.
- McGarty, C. and Turner, J. C. (1992). The effects of categorization on social judgement. *British Journal of Social Psychology*, 31:253–268.
- Miller, D. I., Eagly, A. H., and Linn, M. C. (2015). Women's representation in science predicts national gender-science stereotypes: Evidence from 66 nations. *Journal of Educational Psychology*, 107:631–644.
- Moreno, K. N. and Bodenhausen, G. V. (1999). Resisting stereotype change: The role of motivation and attentional capacity in defending social beliefs. *Group Processes & Intergroup Relations*, 2:5–16.
- Mullen, B. and Hu, L.-T. (1989). Perceptions of ingroup and outgroup variability: A meta-analytic integration. *Basic and Applied Social Psychology*, 10:233–252.
- Murphy, R. A., Schmeer, S., Vallee-Tourangeau, F., Mondragon, E., and Hilton, D. (2011). Making the illusory correlation effect appear and then disappear: The effects of increased learning. *Quarterly Journal of Experimental Psychology*, 64:24–40.
- Nelson, T. E., Biernat, M. R., and Manis, M. (1990). Everyday base rates (sex stereotypes): Potent and resilient. *Journal of Personality and Social Psychology*, 59:664–675.
- Nisbett, R. and Kunda, Z. (1985). Perception of social distributions. *Journal of Personality and Social Psychology*, 48:297–311.
- Ostrom, T. and Sedikides, C. (1992). Out-group homogeneity effects in natural and minimal groups. *Psychological Bulletin*, 112:536–552.
- Park, B. and Hastie, R. (1987). Perception of variability in category development: Instance-versus abstraction-based stereotypes. *Journal of Personality and Social Psychology*, 53:621–635.
- Park, B. and Judd, C. (1990). Measures and models of perceived group variability. *Journal of Personality and Social Psychology*, 59:173–191.
- Park, B. and Rothbart, M. (1982). Perception of out-group homogeneity and levels of social categorization: Memory for the subordinate attributes of in-group and out-group members. *Journal of Personality and Social Psychology*, 42:1051–1068.
- Park, B., Ryan, C., and Judd, C. (1992). Role of meaningful subgroups in explaining differences in perceived variability for in-groups and out-groups. *Journal of Personality and Social Psychology*, 63:553–567.

- Quattrone, G. and Jones, E. (1980). The perception of variability within in-groups and out-groups: Implications for the law of small numbers. *Journal of Personality and Social Psychology*, 38:141–152.
- Queller, S. and Smith, E. (2002). Subtyping versus bookkeeping in stereotype learning and change: Connectionist simulations and empirical findings. *Journal of Personality and Social Psychology*, 82:300–313.
- Richards, Z. and Hewstone, M. (2001). Subtyping and subgrouping: Processes for the prevention and promotion of stereotype change. *Personality and Social Psychology Review*, 5:52–73.
- Rogers, K. H. and Wood, D. (2010). Accuracy of united states regional personality stereotypes. *Journal of Research in Personality*, 44:704–713.
- Rothbart, M. (1981). Memory processes and social beliefs. In Hamilton, D., editor, *Cognitive Processes in Stereotyping and Intergroup Behavior*, pages 145–181. Erlbaum.
- Rothbart, M. and John, O. P. (1985). Social categorization and behavioral episodes: A cognitive analysis of the effects of intergroup contact. *Journal of Social Issues*, 41:81–104.
- Ryan, C. S., Judd, C. M., and Park, B. (1996). Effects of racial stereotypes on judgments of individuals: The moderating role of perceived group variability. *Journal of Experimental Social Psychology*, 32:71–103.
- Sanborn, A. N., Griffiths, T. L., and Navarro, D. J. (2010). Rational approximations to rational models: alternative algorithms for category learning. *Psychological Review*, 117:1144–1167.
- Simon, B. and Brown, R. (1987). Perceived intragroup homogeneity in minority–majority contexts. *Journal of Personality and Social Psychology*, 53:703–711.
- Simon, B. and Pettigrew, T. F. (1990). Social identity and perceived group homogeneity: Evidence for the ingroup homogeneity effect. *European Journal of Social Psychology*, 20:269–286.
- Smith, E. and Decoster, J. (1998). Knowledge acquisition, accessibility, and use in person perception and stereotyping: Simulation with a recurrent connectionist network. *Journal of Personality and Social Psychology*, 74:21–35.
- Smith, E. R. and Zárate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review*, 99:3–21.
- Snyder, M. and Miene, P. (1994). On the functions of stereotypes and prejudice. In Zanna M., O. J., editor, *The Psychology of Prejudice*, pages 33–42. Psychology Press.
- Spears, R., Eiser, J. R., and Van Der Pligt, J. (1987). Further evidence for expectation-based illusory correlations. *European Journal of Social Psychology*, 17:253–258.
- Tajfel, H. (1969). Cognitive aspects of prejudice. *Journal of Biosocial Science*, 1:173–191.
- Tajfel, H. and Wilkes, A. L. (1963). Classification and quantitative judgement. *British Journal of Psychology*, 54:101–114.
- Taylor, S. E. (1981). A categorization approach to stereotyping. In Hamilton, D., editor, *Cognitive Processes in Stereotyping and Intergroup Behavior*, pages 83–115. Erlbaum.

- Van Rooy, D., Van Overwalle, F., Vanhoomissen, T., Labiouse, C., and French, R. (2003). A recurrent connectionist model of group biases. *Psychological Review*, 110:536–563.
- Vanhoomissen, T. and Van Overwalle, F. (2010). Me or not me as source of ingroup favoritism and outgroup derogation: A connectionist perspective. *Social Cognition*, 28:84–109.
- Vong, W. K., Navarro, D. J., and Perfors, A. (2016). The helpfulness of category labels in semi-supervised learning depends on category structure. *Psychonomic Bulletin & Review*, 23:230–238.
- Weber, R. and Crocker, J. (1983). Cognitive processes in the revision of stereotypic beliefs. *Journal of Personality and Social Psychology*, 45:961–977.
- Wilder, D. A. (1984). Intergroup contact: The typical member and the exception to the rule. *Journal of Experimental Social Psychology*, 20:177–194.
- Yzerbyt, V. Y., Coull, A., and Rocher, S. J. (1999). Fencing off the deviant: The role of cognitive resources in the maintenance of stereotypes. *Journal of Personality and Social Psychology*, 77:449–462.