Applying Hierarchical Bayesian Modeling to Experimental Psychopathology Data: An Introduction and Tutorial

Ivy F. Tso
Stephan F. Taylor
Timothy D. Johnson

University of Michigan, Ann Arbor

Author Note

Ivy F. Tso, Department of Psychiatry, University of Michigan, Ann Arbor; Stephan F. Taylor, Department of Psychiatry, University of Michigan, Ann Arbor; Timothy D. Johnson, Department of Biostatistics, University of Michigan, Ann Arbor. This research was approved by the University of Michigan Medical School IRB and supported by the National Institute of Mental Health (K23 MH108823 and R01 MH122491 to IFT; R21 MH101676 to SFT) and University of Michigan Depression Center (Rachel Upjohn Clinical Scholar Award to IFT). The authors thank Merranda McLaughlin, Tyler Grove, Savanna Mueller, and Beier Yao for their assistance in data collection. Preliminary analyses of some of the data included in this paper have been presented at the 2018 Society for Biological Psychiatry annual meeting (*Biological Psychiatry*, 83(9): S427, 2018) and the 2019 Congress of Schizophrenia International Research Society meeting (*Schizophrenia Bulletin*, 45(Supplement_2): S114–S115, 2019). Correspondence concerning this article should be addressed to Ivy Tso, 4250 Plymouth Road, Rachel Upjohn Building, Department of Psychiatry, University of Michigan, Ann Arbor, MI 48109. E-mail: ivyts@umich.edu.

© 2021, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, via its DOI: 10.1037/abn0000707

Abstract

Over the past two decades Bayesian methods have been gaining popularity in many scientific disciplines. However, to this date, they are rarely part of formal graduate statistical training in clinical science. Although Bayesian methods can be an attractive alternative to classical methods for answering certain research questions, they involve a heavy “overhead” (e.g., advanced mathematical methods, complex computations), which pose significant barriers to researchers interested in adding Bayesian methods to their statistical toolbox. To increase the accessibility of Bayesian methods for psychopathology researchers, this paper presents a gentle introduction of the Bayesian inference framework and a tutorial on implementation. We first provide a primer on the key concepts of Bayesian inference and major implementation considerations related to Bayesian estimation. We then demonstrate how to apply hierarchical Bayesian modeling (HBM) to experimental psychopathology data. Using a real dataset collected from two clinical groups (schizophrenia and bipolar disorder) and a healthy comparison sample on a psychophysical gaze perception task, we illustrate how to model individual responses and group differences with probability functions respectful of the presumed underlying data-generating process and the hierarchical nature of the data. We provide the code with explanations and the data used to generate and visualize the results to facilitate learning. Finally, we discuss interpretation of the results in terms of posterior probabilities and compare the results with those obtained using a traditional method.

General Scientific Summary: This paper presents a gentle introduction of the hierarchical framework of Bayesian inference and a tutorial on implementing hierarchical Bayesian modeling (HBM) in experimental psychopathology research.

Keywords: Bayesian estimation, hierarchical models, schizophrenia, bipolar disorder, social cognition, visual perception

Applying Hierarchical Bayesian Modeling to Experimental Psychopathology Data: An Introduction and Tutorial

Bayesian statistics is rapidly gaining popularity in many disciplines, especially STEM (Science, Technology, Engineering, and Mathematics) (Van de Schoot et al., 2017), as it can serve as a good alternative to classical (frequentist) methods (Krypotos et al., 2017; Wagenmakers et al., 2018). On the philosophical level, the Bayesian framework aligns more closely with what a scientist generally wants to know: *the probability that their theory is true given the observed data*, $P(\text{theory}|\text{data})$ (Cohen, 1994). This is in contrast to the classical framework in which the p -value represents long-run probability of obtaining a finding as or more extreme as the one found *given that the null hypothesis (H_0) is true*, $P(\text{data}|H_0)$ (Wasserstein & Lazar, 2016). On the pragmatic level, Bayesian results (posterior probabilities, Bayes factors) are generally more straightforward and intuitive to interpret than p -values or confidence intervals and can provide graded evidence (Greenland et al., 2016; Hoekstra et al., 2014; Wasserstein & Lazar, 2016). Some research questions, such as those seeking to quantify the evidence *in favor of* the null hypothesis or the *comparative evidence* of the null vs. an alternative hypothesis, can only be addressed by Bayesian methods as the p -value provides no information about the probability of the null nor any alternative hypothesis conditional on the observed data (Dienes, 2014).

Although the interest in learning and applying Bayesian statistics among the psychology research community has been growing (Etz & Vandekerckhove, 2016), and the computational costs associated with Bayesian methods have largely been overcome with technological advances in the past two decades, classical inference remains dominant in published psychology research. A major reason is that Bayesian statistics is not included in the curriculum of mainstream graduate statistical training (Wagenmakers et al., 2018). Bayesian analysis requires an understanding of advanced mathematical concepts and implementation of complex statistical computations, which often make the self-learning curve prohibitively steep.

This paper aims to increase the accessibility of Bayesian methods to psychopathology researchers who believe that their work can be enhanced by adding this alternative framework to their statistical repertoire. Here, we provide an overview of the Bayesian inference framework, highlighting some features of the hierarchical Bayesian framework that are especially suitable for modeling experimental psychopathology studies. We discuss some of the key issues and implementation considerations in Bayesian estimation, followed by a demonstration on how to apply hierarchical Bayesian modeling (HBM) to a dataset of an experimental psychopathology study. The tutorial includes details about the assumptions, procedure, data, and code (explained line to line) used to estimate the models and visualize the results. Finally, we compare the results with those obtained using a traditional method (least-squares estimation of individual parameters and group comparisons using t - or F -tests) in previous publications and discuss the sources of the similarities and differences.

A (Super) Gentle Introduction of Bayesian Inference

Bayes' Rule: The Essence of Bayesian Inference

Bayesian inference is a probabilistic framework in which conclusions (*posterior*) are drawn by a rational process of integrating prior belief (*prior*) with new data (*likelihood*). This integration process, referred to as Bayesian updating, involves re-allocating credibility over possibilities (i.e., parameter values) in light of new data through Bayes' rule (Bayes & Price, 1763). The most basic form of Bayes' rule is as follows. Suppose D denotes the observed data and α is the parameter in the model used to describe D . Then the posterior (i.e., the probability of α given D), denoted by $p(\alpha|D)$, can be derived by knowing three things: 1) the prior probability

of α , denoted by $p(\alpha)$; 2) the likelihood of the data (i.e., the probability of the data given the parameter), denoted by $p(D|\alpha)$; and 3) the probability of the data, denoted by $p(D)$:

$$p(\alpha|D) = p(\alpha) p(D|\alpha) / p(D) \quad (1)$$

where $p(D)$, also called the “marginal likelihood” of the data or “evidence,” can be obtained by integrating the likelihood over all possible values of the parameter (denoted by α^*), weighted by the prior probability of each specific α value:

$$p(D) = \int p(D|\alpha^*) p(\alpha^*) d\alpha^* \quad (2)$$

Note that Equation (2) is an integral when the parameter is continuous but a summation when the parameter is discrete.

Let's illustrate the application of Bayes' rule with a classic example of a diagnostic test. Suppose there is a diagnostic blood test for schizophrenia. Here the parameter, α , is the true disease status, which has two possible values: schizophrenia (SZ) or not schizophrenia ($\sim$ SZ). The data, D , is the test result, which also has two possible values: positive (+) or negative (-). This test has a sensitivity of 98% and a specificity of 95%. In other words, the probability that someone with schizophrenia will test positive, $p(+|SZ)$, is 98%, while the probability the someone without schizophrenia will test negative $p(-|\sim SZ)$, is 95%. We know that the prevalence of schizophrenia is 1%, meaning that the probability of having schizophrenia, $p(SZ)$, is 1%, and the probability of not having schizophrenia, $p(\sim SZ)$, is 99%. Suppose a patient is tested positive using this test, we are able to use the Bayes rule as appears in Equation (1) to calculate the probability of this person having schizophrenia:

$$\begin{aligned} p(SZ|+) &= p(SZ) \times p(+|SZ) / p(+) \\ &= 0.01 \times 0.98 / p(+) \end{aligned} \quad (3)$$

To complete this calculation, we need to find the value of $p(+)$. The probability of $p(+)$ is essentially the probability of tested positive whether or not one has schizophrenia. This can be found by summing the likelihood of tested positive over all possible values of the parameter (SZ or $\sim$ SZ), weighted by the prior probability of each specific parameter value:

$$\begin{aligned} p(+) &= [p(+|SZ) \times p(SZ)] + [p(+|\sim SZ) \times p(\sim SZ)] \\ &= (0.98 \times 0.01) + [1 - p(-|\sim SZ)] \times 0.99 \\ &= (0.98 \times 0.01) + [(1 - 0.95) \times 0.99] \\ &= 0.0593 \end{aligned} \quad (4)$$

Substituting Equation (4) into Equation (3) allows us to determine that the value of the posterior $p(SZ|+) = 0.165261$. To sum up, even though the blood test has excellent sensitivity (98%) and specificity (95%), the probability of having schizophrenia given a positive test result is only a dismal 16.5%. This is largely due to the low base-rate of schizophrenia (1%).

Bayesian Estimation in Practice

The above example of estimating a posterior using Bayes' rule is an extremely simple case compared with almost all Bayesian estimations in real life data analyses for several reasons. First, while the parameter and data point of the blood test example have only two possible values, many analyses involve parameters and data that can have infinitely many possible values (as in continuous variables). In these cases, all the terms in Equation (1) are probability densities rather than discrete probabilities. This substantially increases the complexity of the computation. Second, the probability of the disease is known in the blood test example. However, in real life analyses, it is almost never the case that the probability of the parameter is known. Third, while the blood test example involves only one data point (subject) and one parameter, real data analyses always involve multiple data points and parameters. For instance, in a very simple analysis that aims to estimate difference in a normally distributed measure between two

diagnostic groups, at least four parameters are involved: a mean and a variance parameter for each group. The *joint* posterior distribution describing all combinations of parameter values becomes a 4-dimensional space and is very difficult (if possible at all) to derive analytically. For detailed mathematical mechanics of Bayesian estimations, please see excellent accounts provided elsewhere (Gelman et al., 2013; Gill, 2008).

Because of these issues, real life Bayesian data analyses are never as simple as the arithmetic solution presented in the blood test example. Rather, making explicit assumptions about the distribution of the parameters or data, selecting probability distributions to represent the researcher's prior beliefs about the values of the true parameters (i.e., prior selection), and implementing computationally heavy simulations are required to actualize Bayesian estimation. These steps are non-trivial and require careful considerations.

Distributional Assumptions

After identifying the variables or measures relevant in answering the research questions, the first step of Bayesian analysis is to construct a model that specifies the probability distributions of all data and parameters of interest and links them via mathematical functions. For example, in a yes-no sensory detection paradigm, while the outcome of each independent trial with the same probability of detection ("yes") can be described by a Bernoulli distribution, if the same trial is repeated multiple times (as is the case in most experimental psychopathology studies), and if the outcome variable of interest is cumulative detections, then it should be modeled by a binomial probability distribution. Further, the probability-of-success parameter of the binomial distribution of an individual can in turn be modeled a priori to come from a beta distribution that describes the variability of detection rate (continuous between 0 and 1) in the population to which this individual belongs. Different types of data and research questions require different probability distributions and models to describe them accurately, and parameters need to be linked to properly reflect their statistical properties and theoretical relationships with one another. Misspecifications of data/parameter distributions and models can potentially lead to estimation bias (Müller, 2013; Z. Yang & Zhu, 2018).

Choice of Priors

Selecting appropriate prior distributions is a key step in Bayesian inference, since prior distributions are combined with the data to yield the posterior. In general, researchers select priors for their parameters based on prior knowledge from the literature and/or the researcher's belief about the values of the parameters. Prior selection can be a controversial topic (Gelman, 2008) and is often criticized because it can be subjective and can greatly influence the results, especially for small-sample studies. There are helpful guidelines provided by leading Bayesian researchers on how to select and parameterize priors for different types of data and models (e.g., Gelman, 2019; Kass & Wasserman, 1996).

To minimize the influence of priors on the inference, one common strategy is to choose vague, noncommittal distributions. Some examples include: a uniform distribution between 0 and 1 for proportion parameters; a uniform distribution with no bounds ("flat") or a very diffuse normal distribution (i.e., one with a very large variance) for a mean parameter that can take any value between $-\infty$ and $+\infty$ and the researcher does not have a specific hypothesis about its true value; or a wide gamma, beta, or half-Cauchy distribution for parameters that are restricted to be positive. It should be noted that in some situations, using non-informative priors indiscriminately may lead to scientifically or clinically unreasonable posterior values (Gelman, 2002); using weakly informative priors may be needed to regularize and rule out unreasonable parameter values (Gelman, 2019). In general, using relatively flat priors in Bayesian models often lead to

results very similar frequentist maximum likelihood estimation, as long as the models are the same.

Another way to identify prior distributions that have minimal influence over the posterior is to use reference priors (Berger et al., 2009; Bernardo, 1979). The idea is to choose an objective prior distribution that would maximize the discrepancy between the prior and the posterior (for example, as measured with Kullback-Leibler divergence), so that the posterior estimates are maximally based on the data. However, reference priors are complex to construct and available only for few (and generally simple) models (Berger et al., 2001; Natarajan & Kass, 2000; Sun & Ye, 1995; R. Yang & Berger, 1994; Ye, 1994).

To check that the posterior is not unduly influenced by the prior choice, one can conduct a sensitivity analysis to evaluate the extent to which the inference is robust with respect to prior specifications. This can be performed formally (quantifying the effects of various perturbations to the priors through formal mathematical derivation and computation) or informally (fitting the model with a few alternative sets of prior distributions and then assessing subjectively whether the differences in the posteriors are tolerable). Because formal prior sensitivity analysis involves complex computation and is not readily implemented in popular Bayesian analysis software (Roos et al., 2015), most prior sensitivity analyses in practice are in the informal category. Some authors have provided helpful guidance on how to interpret the results of prior sensitivity analysis (van Erp et al., 2018).

Estimation

As discussed above, most of the time the joint posterior distribution for Bayesian analyses cannot be derived analytically. There are two general approaches to overcome this problem: variational methods and sampling methods. Variational methods refer to the modifications and transformations of the equations so that an analytical solution to an approximation of the joint posterior becomes possible. An alternative is sampling methods, referring to the drawing of representative samples from the joint posterior distribution to approximate numerically the posterior. Parameter estimation is typically faster with variational methods than sampling methods, but variational methods tend to underestimate variance of the posterior density (Blei et al., 2017). Deriving the sets of equations to carry out variational Bayes is also complex and involves specialized mathematical skills. Comparatively, sampling methods are much easier to execute and thus more accessible to a wide range of researchers.

Here, we will focus on Bayesian estimation via sampling methods, in particular, Markov chain Monte Carlo (MCMC) methods, because many of modern-day software for Bayesian estimation use some variant of MCMC methods. In statistics, a Monte Carlo simulation refers to the method of assessing the properties of a target distribution by generating representative random values of the distribution. A “Markov chain” is a succession of steps of a random walk (i.e., given the current value, the next value is independent of all past values). Mathematicians have developed numerous MCMC sampling algorithms, with the Metropolis-Hastings algorithm, Gibbs sampling (a special case of the Metropolis-Hastings algorithm) (Geman & Geman, 1984), and slice sampling (Neal, 2003) being the most commonly used sampling algorithms in popular Bayesian estimation software such as BUGS and JAGS (see for details Kruschke, 2015). These sampling methods allow users to use independent (2 or more) Markov chains to generate samples of parameter values to approximate the posterior. The more MCMC samples are generated (typically in the order of tens of thousands, depending on the model complexity), the more accurate and refined is the representation of the posterior distribution. Because of this, Bayesian estimations are much more computationally expensive than classical methods and were generally

infeasible until the last 20 years, when computers with sufficient computational power became widely accessible and affordable.

When the sampling algorithm works properly, a Markov chain should converge to the target distribution as long as it is run long enough. However, some considerations and strategies are needed to improve the process of MCMC and ensure that the results are representative of the target distribution. Because it takes a certain number of iterations for a Markov chain to travel to the representative region of a probability distribution, the initial samples are often discarded to minimize the influence of these non-representative values on the results. This “burn-in” period can range from tens to thousands of iterations, depending on the complexity of the model and the speed of “mixing.” Mixing refers to how well the Markov chain explores the posterior distribution. It depends heavily on the autocorrelation of the Markov chain. If the autocorrelation is large, it mixes poorly; in contrast, if the Markov chain samples are nearly independent, the chain mixes well. Determining how long the chain should run and how long the burn-in period should be is a trial-and-error process, but techniques such as the Raftery and Lewis diagnostic (Raftery & Lewis, 1992) can help determine (after the fact) if the chains and burn-in periods were sufficiently long given certain user-defined acceptable tolerance criteria. Another issue with MCMC is that sometimes a chain can get stuck in local maxima or low-density regions. If the chain terminates before having sufficiently sampled high-density regions, this can substantially distort the posterior distribution and parameter estimates. Setting up multiple independent chains with overdispersed starting values (Gelman & Rubin, 1992) and visually inspecting the traceplots (plots of value of the MCMC sample against the iteration number) and autocorrelations between the draws for each parameter can help determine if the chains are mixing well. Heavily overlapping traceplots of independent chains after the burn-in period and near-zero autocorrelation between samples at a reasonably small (e.g., 30th) lag are good signs. Additionally, formal diagnostics should be routinely run to assess convergence quantitatively. Some commonly used quantitative diagnostics include Gelman and Rubin’s convergence diagnostic (Gelman & Rubin, 1992) or the more generalized version, Gelman-Rubin-Brooks (GRB) diagnostic (Brooks & Gelman, 1998); Geweke’s diagnostic (Geweke, 1992); and Heidelberg and Welch’s diagnostics (Heidelberger & Welch, 1981, 1983). A caveat of convergence diagnostics is that they can identify non-convergence, but cannot prove convergence.

Hierarchical Modeling in the Bayesian Framework

Within the Bayesian framework, hierarchical modeling may be particularly suitable for experimental psychopathology research. Many experimental psychopathology studies use case-control designs and behavioral tasks or paradigms to measure specific cognitive or psychological functions to elucidate the etiology or disease process of the psychopathology(ies) being investigated. In these types of studies, there is a clear hierarchy in the data structure: trials of the behavioral task are nested within individuals, which are in turn nested within groups. Often times behavior data of each individual are averaged across trials to provide a point estimate, which is then taken to the group level to estimate group differences. Although averaging across trials can help reduce noise in the measurement, it removes important information about inter-trial variability that may be a defining characteristic of the psychopathology being studied. This practice of obtaining point estimates at the individual level also essentially gives equal weights to individual data of different quality (i.e., inter-trial variability), contributing to less accurate group-level estimates.

These problems are nicely addressed in hierarchical Bayesian methods by explicitly modeling the probability distributions giving rise to the data and estimating the uncertainty at each level (trial, individual, and group), allowing the estimation of uncertainty at both the individual and group levels. Additionally, setting up a hierarchical model to reflect the hierarchical structure of the data enables mathematical mechanics that use individual estimates to inform group estimates, which in turn constrain the estimates of individual parameters (Kruschke, 2015). The result of this bi-directional constraint imposed by individual and group estimates on one another is that individual estimates (especially ones with lower quality, or higher uncertainty) are pulled toward the group mean—a phenomenon called shrinkage. Shrinkage improves both individual and group estimates (Kruschke, 2015; Kruschke & Vanpaemel, 2015), and the benefit is especially appreciable when the number of observations is small (Farrell & Ludwig, 2008). Although shrinkage can also be achieved in hierarchical models using the frequentist framework (e.g., mixed-effect models, which often use maximum likelihood estimation or restricted MLE procedures), when conducted in a Bayesian framework there is an additional hierarchy of priors that further constrain parameter values within reasonable ranges. This is in contrast to MLE, where a maximum can be found outside the region of realistic parameters (e.g., negative attention lapse rate) (Schütt et al., 2016; Wichmann & Hill, 2001). Although the parameter space can also be restricted when performing MLE, sometimes a maximum may not be found within the interior of such restricted parameter space.

Workflow of Hierarchical Bayesian Modeling

Before we proceed to a demonstration, the typical workflow of conducting a hierarchical Bayesian modeling HBM analysis is summarized in Figure 1.

[Figure 1]

Demonstration: Deconstructing the Cognitive Components of Eye Gaze Perception in Schizophrenia and Bipolar Disorder Using Hierarchical Bayesian Modeling

Hierarchical Bayesian modeling can be applied straightforwardly to psychopathology studies that use experimental paradigms to psychometrically measure perceptual or cognitive differences between two or more clinical populations. Here, we use data from a psychophysical experiment to assess two key cognitive components involved in eye gaze perception (perceptual precision and self-referential bias) in schizophrenia and bipolar disorder. Eye gaze is a ubiquitous social cue that conveys attention and mental state, and the ability to accurately and efficiently discriminate others' gaze direction (especially self-directed gaze) is critical to understanding others and navigating the complex social world. Impairment in social cognition in schizophrenia (Green et al., 2015) and bipolar disorder (Brotman et al., 2008; Lahera et al., 2012; Van Rheeën & Rossell, 2014) has been documented in numerous studies and linked to social and community dysfunction (Fett et al., 2011). Therefore, understanding gaze perception—a basic building block of higher-order social cognition (Itier & Batty, 2009)—and its constituent cognitive processes could help identify the sources of deficits in these clinical populations and inform intervention. Because schizophrenia and bipolar disorder have significant genetic, clinical, and neurocognitive overlaps (Maier et al., 2006), investigating group differences in the cognitive components of gaze perception can reveal if similar mechanisms underlie social cognitive deficits observed in the two disorders.

Gaze perception entails sensory processing of visual stimuli (position of the eyes in the context of a face) and making a self-referential judgment about gaze (i.e., self-directed or not). Disruption to either of these processes could lead to altered gaze perception. For example,

disrupted early visual processing (which is well documented in schizophrenia and psychotic disorders) (Butler et al., 2008) could lead to noisy (or, imprecise) gaze perception and a tendency to perceive social information as self-directed—a common phenomenon in paranoid or grandiose delusions, highly prevalent in schizophrenia or bipolar psychosis (An et al., 2010; Combs et al., 2007)—can lead to a self-referential bias in gaze perception. Visual perception precision and self-referential tendency can be dissociated and delineated by systematically studying perception of self-directed gaze as a function of gaze direction using psychophysics methods. In psychophysics experiments, stimulus signal strength is systematically manipulated, so that the observer's response as a function of signal strength, typically sigmoidal, can be derived. This psychometric function allows the estimation of two important perceptual properties: threshold and slope. Threshold, in experiments using a yes-no detection task, is the signal strength that elicits positive response 50% of the time, indexing the signal intensity that the observer requires in order to meaningfully detect the stimulus. Slope of the psychometric function, when measured at detection rate of 50%, indexes the sensitivity of the sensory system to discriminate ambiguous stimuli. These two metrics can be readily applied to the case of gaze perception.

We devised a psychophysical gaze perception task in which participants see faces with different gaze directions, ranging from averted to direct in gradual increments. To each face, they need to make yes-no judgments of eye contact. In this context, threshold is the gaze direction that elicits perception of self-directed gaze half of the time in an individual. Because there is no objectively correct or incorrect answer to the perception of self-directed gaze, threshold essentially reflects the observer's subjective bias to perceive gaze as self-directed, with lower thresholds (corresponding to lower “eye-contact signal strengths,” meaning more averted gaze angles) indicating a stronger self-referential bias. The slope of the perception curve indicates the rate of change of the observer's perception with respect to change in gaze direction, reflecting the precision of the observer's visual perception. See Figure 2 for an illustration of the gaze perception curve and metrics derived from behavioral responses in this gaze perception task. In other words, the threshold and slope of the gaze perception curve nicely quantify the two prime candidate cognitive processes, respectively, underlying gaze perception. In this demonstration, we used hierarchical Bayesian modeling to test the hypotheses that 1) individuals with schizophrenia would show both increased self-referential bias and reduced visual perceptual sensitivity compared with healthy controls; and 2) individuals with bipolar disorder would fall intermediate between schizophrenia and healthy controls in both metrics.

[Figure 2]

Participants

The dataset ($N = 157$) consisted of 55 healthy controls (HC), 55 individuals with bipolar I or II disorder (BP), and 47 individuals with schizophrenia or schizoaffective disorder (SZ). It included new data ($n = 44$) as well as previously published data (Tso et al., 2012; Yao et al., 2018). All participants gave written informed consent, and the study was conducted in accordance with protocols approved by the University of Michigan Medical School Institutional Review Board. See previous publications (Tso et al., 2012; Yao et al., 2018) for detailed inclusion/exclusion criteria.

Psychophysical Eye Gaze Perception Task

Details of the eye gaze perception task have been described in our previous publications (Tso et al., 2012; Yao et al., 2018). Briefly, the stimuli of the task consist of faces with different gaze angles, ranging from direct (0°) to 30° averted in ten 10% increments (see Figure 3A for example stimuli). Since self-directed gaze is perceived more frequently with smaller gaze angles

(i.e., more direct gaze), gaze angle was converted to a scale of “eye-contact signal strength,” with 0.0 corresponding to gaze angle of 30° averted and 1.0 corresponding to gaze angle of 0° (direct gaze). The face stimuli also varied in head orientations (forward, 30° horizontally rotated) and emotion (neutral, fearful), but for simplicity we included only responses to forward, neutral faces in the analyses in this paper. There were 12 trials (6 actors $\times$ 2 left-right directions) for each of the 11 gaze angles, totaling 132 trials per participant. In the task, participants viewed the faces one at a time in a pseudorandomized order. For each face, they made a two-forced choice (“Looking at you?” yes or no) to indicate perception of self-directed gaze. The task was self-paced and participants were allowed to pause and take a brief break whenever they needed. Trials with reaction time shorter than 300 ms were considered invalid and excluded. Responses to the Gaze Perception Task of the HC, BP, and SZ groups are displayed in Figure 3B.

[Figure 3]

Software and Packages/Toolboxes Required

Below, we demonstrate how to model and analyze the data using hierarchical Bayesian modeling, and then compare the results with those obtained using a traditional method (least-squared fitting and F or t-tests). The software and toolboxes/packages required to run the analyses and visualize the results in this demonstration are listed in Supplementary Information Table S1. Briefly, R is used to call the WinBUGS software (Version 1.4) (Spiegelhalter et al., 2003), via the R2WinBUGS package (Sturtz et al., 2005), to run the hierarchical Bayesian model and diagnostics. WinBUGS uses Gibbs, Metropolis-Hastings, slice sampling, or adaptive rejection sampling, depending on the problem. Subsequently, R is also used to visualize the posterior density distributions of the parameters of interest. For the traditional analyses involving least-squared fitting, we used MATLAB and associated toolboxes.

Data Files and Analysis Code

The files containing data and code used in this demonstration are listed in Supplementary Information Table S2 and available as online supplementary materials for download.

Hierarchical Bayesian Modeling (HBM)

The Model

The structure of the HBM of gaze perception is schematically depicted in Figure 4. Excerpts of the model code are presented in Listing 1 and Listing 2 in the text here, the details of which we will go over in the following explanation of the model. HBM begins with identifying the probability distributions associated with the data-generating process within each subject. So, the model code begins with an i loop over $N = 157$ subjects (Listing 1, Lines 13 – 34). In the gaze perception task, a subject made a binary (yes or no) response in each trial to indicate their perception of self-directed gaze. They completed 12 trials (or fewer if some responses were invalid and excluded) for each of the 11 eye-contact signal strengths. Therefore, a j loop over 11 eye-contact signal strengths is embedded within the i loop (Listing 1, Lines 15 – 24).

[Figure 4]

[Listing 1]

Let θ_{ij} be the unknown, true probability underlying the perception of self-directed gaze. Then, the observed number of “yes” responses, Y_{ij} , is a random variable that follows a binomial probability distribution dependent on θ_{ij} and T_{ij} (the number of valid trials):

$$Y_{ij} \sim \text{Binomial}(\theta_{ij}, T_{ij})$$

where i = subject index,

j = gaze angle index corresponding to the 11 gaze angles

This is captured in Listing 1, Line 19. The code so far indicates that each subject i has 11 θ_{ij}

values (one for each gaze angle, j). However, since we are interested in θ_{ij} as a function of eye-contact signal strength, X_j , rather than the individual values of θ_{ij} per se, we need to link θ_{ij} to X_j via a mathematical function. Perception of self-directed gaze is a putatively categorical function such that the probability of perceiving self-directed gaze (θ_{ij}) would increase as eye-contact signal increases. Fittingly, our prior work demonstrated that the probability of perceiving self-directed gaze approximates a logistic function of eye-contact signal strength (Tso et al., 2012). So, for each subject, we linked θ_{ij} to X_j via a logit function using two free parameters, α_i and β_i :

$$\text{logit}(\theta_{ij}) = \alpha_i + \beta_i \cdot X_j \quad (5)$$

This step is captured in Listing 1, Line 22.

Next, the hierarchical nature of the data is modeled using the knowledge that the subjects come from three different diagnostic groups. This is done by modeling individual α_i and β_i parameters to come from two normal distributions centered around the respective mean of the subject's corresponding group k :

$$\begin{aligned} \alpha_i &\sim \text{Normal}(\mu_{\alpha k}, \tau_{\alpha k}) \\ \beta_i &\sim \text{Normal}(\mu_{\beta k}, \tau_{\beta k}) \end{aligned}$$

where $k = 1$ (HC) for $i = 1 - 55$
 $k = 2$ (BP) for $i = 56 - 110$
 $k = 3$ (SZ) for $i = 111 - 157$

and the normal distributions are parameterized by μ (mean) and τ (precision or reciprocal of variance). This step is captured in Listing 1, Lines 27 – 28. This completes the modeling of individual subjects. (The i loop does not finish until Line 34; we will explain what Lines 31 – 32 do later.)

Next, we assign priors to the group parameters (μ 's and τ 's). In Bayesian estimation, researchers select priors for their parameters based on prior knowledge from the literature and/or the researcher's belief about the values of the parameters. Here, we choose vague, noncommittal distributions to minimize the influence of priors on the results. In the code, we use a k loop over the three diagnostic groups (Listing 2, Line 36) to assign priors to the group parameters. For μ_α 's, flat priors are used because they can take on any value and we do not have specific expectations of their values (Listing 2, Line 38). For μ_β 's, we also used flat priors (Listing 2, Line 39). Although β (representing the slope of the psychometric function) should be non-negative in the context of this experiment, there are no specific restrictions in the logit equation that would constrain this value to non-negative. Therefore, mathematically μ_β 's can take on any value, making flat priors an appropriate (non-informative) choice. For precision parameters (i.e., τ 's), gamma priors are used because variances must be non-negative (Listing 2, Lines 40 – 41). Note that gamma distributions are often the standard choice of prior for non-negative, continuous variables (such as variances and reaction time) because they place all probability between 0 and ∞ , with a long tail to the right indicating decreasingly smaller probability for increasingly large values. This completes the modeling of group parameters. (The k loop does not finish until Line 33, and we will explain what Lines 44 – 45 do next.)

[Listing 2]

The formulation so far would yield individual (α_i, β_i) and group estimates ($\mu_{\alpha k}, \mu_{\beta k}, \tau_{\alpha k}, \tau_{\beta k}$). These estimates are used to compute the two gaze perception metrics of interest (self-referential bias and precision) for each individual and each group. As stated above, at the individual level, self-referential bias is represented by the threshold of the logistic function (i.e.,

X_j value at $\theta_{ij} = 0.5$). Therefore, it can be found by substituting 0.5 for θ_{ij} in Equation (5) and solving for X_j :

$$X_j|_{\theta_{ij}=0.5} = -\frac{\alpha_i}{\beta_i} \quad (6)$$

Perceptual precision is represented by the slope of the logistic function at $\theta_{ij} = 0.5$, and can be found by rewriting Equation (5) as:

$$\theta_{ij} = \frac{1}{1 + e^{-(\alpha_i + \beta_i \cdot X_j)}} \quad (7)$$

then deriving the first derivative of Equation (7):

$$\frac{d\theta_{ij}}{dX_j} = \frac{\beta_i \cdot e^{-(\alpha_i + \beta_i \cdot X_j)}}{[1 + e^{-(\alpha_i + \beta_i \cdot X_j)}]^2} \quad (8)$$

and finally evaluating Equation (8) at $\theta_{ij} = 0.5$ by substituting Equation (6):

$$\frac{d\theta_{ij}}{dX_j}|_{\theta_{ij}=0.5} = \frac{\beta_i}{4} \quad (9)$$

These mathematical operations are included in the code so that the values of these measures of interest can be tracked in the MCMC samples. Specifically, Equations (6) and (9) are included within the i loop (Listing 1, Lines 31 and 32). Similarly, group thresholds = $-\mu_{\alpha k}/\mu_{\beta k}$ and group slopes = $\mu_{\beta k}/4$, and these computations are included in the code within the k loop (Listing 2, Lines 44 and 45). Finally, group differences in threshold and slope are computed (Listing 2, Lines 51 – 57).

MCMC Estimation

Part 1 of the R script file “Bayes.R” contains code to load the data from the ‘Gaze_data.csv’ file and convert them into list structures format for the HBM analysis. Part 2 of the R script supplies the initial values for each of the three MCMC chains to be generated. Part 3 of the R script is the code to run the model (specified in the “HBM_model_one.txt” file) by calling WinBUGS. For each MCMC chain, we generated 80,000 samples and discarded the first 5,000 samples as burn-in. The values of the 75,000 samples were saved into a coda output file for each chain. Note that this part of the code also includes code for running a different model (specified in the “HBM_model_two.txt” file), which used different values for the gamma priors. This serves as a sensitivity test to check that the results are robust against choices of hyperparameter values for the priors. Readers are encouraged to run this code and compare the results. Part 4 of the R script file contains code for conducting convergence diagnostics.

HBM Results

Part 5 of the R script contains code to use the resulting $75,000 \times 3$ chains = 225,000 MCMC samples to make statistical inference on group estimates of the gaze perception measures and group differences. Since group differences in threshold and slope were predicted in the direction of HC > BP > SZ, one-tailed posterior probabilities were used to evaluate the credibility of group differences. Part 6 of the R script contains code used to generate posterior density plots of these measures of interest as presented in Figure 5.

[Figure 5]

For threshold, there was an overall pattern of HC > BP > SZ (Figure 5A). While there was a clear HC > SZ difference (posterior probability = 98.50%) and a likely HC > BP difference (posterior probability = 92.90%), the BP > SZ difference was much less clear (posterior probability = 76.68%). For slope, it also showed an overall pattern of HC > BP > SZ (Figure 5B), but this time the posterior distributions of BP and SZ overlapped substantially (with posterior probability of BP > SZ down to 66.27%) and were far apart from HC (posterior

probability of $HC > BP = 99.97\%$; $HC > SZ = 99.97\%$). Taken together, there was strong evidence for reduced perceptual precision during gaze processing in both SZ and BP, and a self-referential bias was present in SZ and probably in BP. These findings support that altered cognition at both perceptual and interpretation levels are involved in abnormal gaze perception in SZ and BP, and the differences between the two clinical groups are at best equivocal.

Traditional Method: Least-squares Fitting and F- or t-tests

The traditional approach begins with fitting a linear function to each subject's logit-transformed response to find the two free parameters, a_i and b_i :

$$\text{logit}(P_{ij}) = a_i + b_i \cdot X_j \quad (10)$$

where P = proportion of valid trials where a "yes" response was given

i = subject index,

j = gaze angle index corresponding to the 11 gaze angles

X = eye-contact signal strength (0, 0.1, ..., 1.0)

Estimation was done using the MATLAB *polyfit* function (MATLAB R2016b, The Mathworks, Natick, MA) to find a best-fit line in a least-squares sense. The MATLAB code (LeastSquares.m) used to automatically load the data from the "Gaze_data.csv" file to perform least-squares fitting and estimation for each individual.

Using the same mathematical derivations in Equations (6) and (9), threshold and slope were then computed for each subject: threshold = $-a_i/b_i$, and slope = $b_i/4$. Then, two separate one-way ANOVAs were performed to compare the three diagnostic groups for the two gaze perception metrics, followed up with post hoc Dunn-Sidak (one-sided) t -tests.

Traditional Results

For threshold, there was a significant group effect, $F(2, 154) = 5.026$, $p = .008$. Post hoc Dunn-Sidak (one-sided) t -test suggested $HC > SZ$, $p = .034$, 95% C.I. bound = 0.034, but not $HC > BP$, $p = .052$, 95% C.I. bound = -0.00046, or $BP > SZ$, $p = .325$, 95% C.I. bound = -0.039. As for slope, the overall group effect was not significant, $F(2, 154) = 2.059$, $p = .131$. Normally, no post hoc tests should be done with non-significant omnibus tests. For the sake of comparison, we followed up with post hoc Dunn-Sidak (one-sided) t -test, which also showed none of the group differences reached statistical significance: $HC > BP$, $p = .50$, 95% C.I. bound = -0.035; $HC > SZ$, $p = .46$, 95% C.I. bound = -0.027; $BP > SZ$, $p = .46$, 95% C.I. bound = -0.027.

Comparing HBM vs. Traditional Results

Even though the pairwise group comparison results for threshold were quite consistent between the two methods, the results regarding group differences in slope were dramatically different between the two methods. Examination of the individual estimates yielded by the two methods gives us some clues why. In Figure 6 (generated by running Parts 7 and 8 of the R script), we can see that although the group mean estimates were generally similar between the two methods, the distribution of individual estimates are tighter when estimated using HBM than the least-squares method.

[Figure 6]

As mentioned in the main text, the hierarchical structure of the Bayesian model allows us to constrain the individual estimates using the group estimates, thus pulling the individual estimates toward their respective group means. The resulting lower within-group variances enhanced the ability to detect group differences. This advantage is especially pronounced for measures with ex-Gaussian distribution, including not only the slope metric in this analysis, but other common dependent variables in experimental psychopathology studies such as reaction time (Matzke et al., 2013). Additionally, the HBM analysis has the advantage of estimating

uncertainty at the individual level and then taking it into account in the estimation of group-level parameters. This is in contrast to the traditional method, which only obtains point estimates for individuals, thus neglecting the associated uncertainty when making group-level inferences.

Although the HBM results that are more consistent with our hypotheses, this by itself does not necessarily mean higher accuracy. As researchers, we seek results that can represent (or are a closer approximation of) the truth. How do we know which method yielded results that are closer to the truth? In this analysis, we used real rather than simulated data and thus were unable to compare the results to the true (unknown) values of the parameters. Nevertheless, simulations of psychophysical experimental data similar to the one described in this paper have shown that Bayesian inference is more able to recover the true parameters compared with classical approaches (Fründ et al., 2011; Kuss et al., 2005). We also assessed how well the gaze perception estimates obtained with the two methods were able to explain behavior of interest. Since eye gaze perception is a basic social cognitive function, we expect the two gaze perception metrics to be able to predict higher-level social cognition (particularly constructs that are related to understanding others). Among our sample, 119 subjects had available an ability-based measure of emotional intelligence, the Mayer-Salovey-Caruso Emotional Intelligence Test (MSCEIT; Mayer, Salovey, & Caruso, 1999), which encompasses measures of one's ability to perceive, use, understand, and regulate emotions. As we can see in Figure 7 (generated by running Part 9 of the R script), both gaze perception metrics estimated using HBM explained more variance in MSCEIT compared to the traditional method. Although this is not direct evidence that the HBM estimates are more accurate, this finding provides some support that the HBM estimates may be more useful in prediction.

[Figure 7]

Conclusions

In this paper, we presented a gentle introduction to the Bayesian inference framework and discussed some common characteristics of experimental psychopathology studies that are especially suitable to take advantage of hierarchical modeling in the Bayesian framework. We used an example of a psychophysical study of eye gaze perception in psychosis to demonstrate how to apply HBM to estimate parameters and uncertainty at the individual and group levels to answer specific research questions in psychopathology. The model used was relatively straightforward. In some cases, researchers might be interested to examine additional factors in the data, such as sex differences. This can be readily performed by including additional parameters to the model (as in Listing 1, Line 22) to index the fixed effect of sex (invariant across diagnostic groups) and/or its interactions with diagnostic groups. Although we used an example in which experimental parameters were estimated with a logit function, the application of HBM is not limited to this type of function. Others have provided accounts in which HBM can be readily applied to estimate other common experimental measures (e.g., reaction time) and perform better in parameter recovery compared with traditional methods (Matzke et al., 2013). HBM can also be readily adopted in computational models of cognitive processes in clinical populations (Kruschke & Vanpaemel, 2015). Taken together, hierarchical Bayesian modeling offers a high degree of flexibility in modeling data of simple and complex hierarchical structure as well as estimating parameter values and uncertainty for individuals and groups. We hope that the gentle introduction and concrete demonstration provided in this paper would serve as a friendly entry point for experimental psychopathology researchers who are curious about venturing into Bayesian analysis.

References

- An, S. K., Kang, J. I., Park, J. Y., Kim, K. R., Lee, S. Y., & Lee, E. (2010). Attribution bias in ultra-high risk for psychosis and first-episode schizophrenia. *Schizophrenia Research*, 118(1–3), 54–61. <https://doi.org/10.1016/j.schres.2010.01.025>
- Bayes, T., & Price, R. (1763). An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, F. R. S. communicated by Mr. Price, in a letter to John Canton, A. M. F. R. S. *Philosophical Transactions of the Royal Society of London*, 53, 370–418. <https://doi.org/10.1098/rstl.1763.0053>
- Berger, J. O., Bernardo, J. M., & Sun, D. (2009). The formal definition of reference priors. *The Annals of Statistics*, 37(2), 905–938. <https://doi.org/10.1214/07-AOS587>
- Berger, J. O., De Oliveira, V., & Sansó, B. (2001). Objective Bayesian Analysis of Spatially Correlated Data. *Journal of the American Statistical Association*, 96(456), 1361–1374. <https://doi.org/10.1198/016214501753382282>
- Bernardo, J. M. (1979). Reference Posterior Distributions for Bayesian Inference. *Journal of the Royal Statistical Society. Series B (Methodological)*, 41(2), 113–147.
- Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. (2017). Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, 112(518), 859–877. <https://doi.org/10.1080/01621459.2017.1285773>
- Brooks, S. P., & Gelman, A. (1998). General methods for monitoring convergence of iterative simulations? *Journal of Computational and Graphical Statistics*, 7(4), 434–455. <https://doi.org/10.1080/10618600.1998.10474787>
- Brotman, M. A., Guyer, A. E., Lawson, E. S., Horsey, S. E., Rich, B. A., Dickstein, D. P., Pine, D. S., & Leibenluft, E. (2008). Facial emotion labeling deficits in children and adolescents at risk for bipolar disorder. *The American Journal of Psychiatry*, 165(3), 385–389. <https://doi.org/10.1176/appi.ajp.2007.06122050>
- Butler, P. D., Silverstein, S. M., & Dakin, S. C. (2008). Visual Perception and Its Impairment in Schizophrenia. *Biological Psychiatry*, 64(1), 40–47. <https://doi.org/10.1016/j.biopsych.2008.03.023>
- Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist*, 49(12), 997–1003. <https://doi.org/10.1037/0003-066X.49.12.997>
- Combs, D. R., Penn, D. L., Wicher, M., & Waldheter, E. (2007). The Ambiguous Intentions Hostility Questionnaire (AIHQ): A new measure for evaluating hostile social-cognitive biases in paranoia. *Cognitive Neuropsychiatry*, 12(2), 128–143. <https://doi.org/10.1080/13546800600787854>
- Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*, 5, 781. <https://doi.org/10.3389/fpsyg.2014.00781>
- Etz, A., & Vandekerckhove, J. (2016). A Bayesian perspective on the reproducibility project: Psychology. *PLoS ONE*, 11(2), 1–12. <https://doi.org/10.1371/journal.pone.0149794>
- Farrell, S., & Ludwig, C. J. H. (2008). Bayesian and maximum likelihood estimation of hierarchical response time models. *Psychonomic Bulletin and Review*, 15(6), 1209–1217. <https://doi.org/10.3758/PBR.15.6.1209>
- Fett, A. K. J., Viechtbauer, W., Dominguez, M. de G., Penn, D. L., van Os, J., & Krabbendam, L. (2011). The relationship between neurocognition and social cognition with functional outcomes in schizophrenia: A meta-analysis. *Neuroscience and Biobehavioral Reviews*, 35(3), 573–588. <https://doi.org/10.1016/j.neubiorev.2010.07.001>
- Fründ, I., Haenel, N. V., & Wichmann, F. A. (2011). Inference for psychometric functions in the

- presence of nonstationary behavior. *Journal of Vision*, 11(6), 1–19.
<https://doi.org/10.1167/11.6.16>
- Gelman, A. (2002). Prior distribution. In *Encyclopedia of Environmetrics* (pp. 1634 – 1637). John Wiley & Sons, Ltd.
- Gelman, A. (2008). Objections to Bayesian statistics. *Bayesian Analysis*, 3(3), 445–450.
<https://doi.org/10.1214/08-BA318>
- Gelman, A. (2019). *Prior Choice Recommendations*. <https://github.com/stan-dev/stan/wiki/Prior-Choice-Recommendations>
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). Chapman & Hall/CRC.
- Gelman, A., & Rubin, D. B. (1992). Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4), 457–472.
- Geman, S., & Geman, D. (1984). Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6), 721–741. <https://doi.org/10.1109/TPAMI.1984.4767596>
- Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In J. M. Bernardo, J. O. Berger, A. P. Dawid, & A. F. M. Smith (Eds.), *Bayesian Statistics 4*. Clarendon Press.
- Gill, J. (2008). *Bayesian Methods: A Social and Behavioral Sciences Approach* (2nd ed.). CRC Press.
- Green, M. F., Horan, W. P., & Lee, J. (2015). Social cognition in schizophrenia. *Nat Rev Neurosci*, 16(10), 620–631. <https://doi.org/10.1038/nrn4005>
- Greenland, S., Senn, S. J., Rothman, K. J., Carlin, J. B., Poole, C., Goodman, S. N., & Altman, D. G. (2016). Statistical tests, P values, confidence intervals, and power: a guide to misinterpretations. *European Journal of Epidemiology*, 31(4), 337–350.
<https://doi.org/10.1007/s10654-016-0149-3>
- Heidelberger, P., & Welch, P. D. (1981). A Spectral Method for Confidence Interval Generation and Run Length Control in Simulations. *Communications of the ACM*, 24(4), 233–245.
- Heidelberger, P., & Welch, P. D. (1983). Simulation Run Length Control in the Presence of an Initial Transient. *Operations Research*, 31(6), 1109–1144.
- Hoekstra, R., Morey, R. D., Rouder, J. N., & Wagenmakers, E. J. (2014). Robust misinterpretation of confidence intervals. *Psychonomic Bulletin and Review*, 21(5), 1157–1164. <https://doi.org/10.3758/s13423-013-0572-3>
- Itier, R. J., & Batty, M. (2009). Neural bases of eye and gaze processing: The core of social cognition. *Neuroscience & Biobehavioral Reviews*, 33(6), 843–863.
<https://doi.org/10.1016/j.neubiorev.2009.02.004>
- Kass, R. E., & Wasserman, L. (1996). The selection of prior distributions by formal rules. *Journal of the American Statistical Association*, 91(435), 1343–1370.
<https://doi.org/10.1080/01621459.1996.10477003>
- Kruschke, J. K. (2015). *Doing Bayesian data analysis: A tutorial with R and BUGS* (2nd ed.). Academic Press.
- Kruschke, J. K., & Vanpaemel, W. (2015). Bayesian estimation in hierarchical models. In J. R. Busemeyer, Z. Wang, J. T. Townsend, & A. Eidels (Eds.), *The Oxford Handbook of Computational and Mathematical Psychology* (pp. 279–299). Oxford University Press.
- Krypotos, A. M., Blanken, T. F., Arnaudova, I., Matzke, D., & Beckers, T. (2017). A Primer on Bayesian Analysis for Experimental Psychopathologists. *Journal of Experimental*

- Psychopathology*, 8(2), 140–157. <https://doi.org/10.5127/jep.057316>
- Kuss, M., Jäkel, F., & Wichmann, F. A. (2005). Bayesian inference for psychometric functions. *Journal of Vision*, 5(5), 478–492. <https://doi.org/10.1167/5.5.8>
- Lahera, G., Ruiz-Murugarren, S., Iglesias, P., Ruiz-Bennasar, C., Herreria, E., Montes, J. M., & Fernandez-Liria, A. (2012). Social cognition and global functioning in bipolar disorder. *J Nerv Ment Dis*, 200(2), 135–141. <https://doi.org/10.1097/NMD.0b013e3182438eae>
- Maier, W., Zobel, A., & Wagner, M. (2006). Schizophrenia and bipolar disorder: differences and overlaps. *Current Opinion in Psychiatry*, 19(2), 165–170. <https://doi.org/10.1097/01.yco.0000214342.52249.82>
- Matzke, D., Dolan, C. V., Logan, G. D., Brown, S. D., & Wagenmakers, E.-J. (2013). Bayesian parametric estimation of stop-signal reaction time distributions. *Journal of Experimental Psychology: General*, 142(4), 1047–1073. <https://doi.org/10.1037/a0030543>
- Mayer, J. D., Salovey, P., & Caruso, D. R. (1999). *Mayer-Salovey-Caruso Emotional Intelligence Test*. Multi-Health Systems Inc.
- Müller, U. K. (2013). Risk of Bayesian Inference in Misspecified Models, and the Sandwich Covariance Matrix. *Econometrica*, 81(5), 1805–1849. <https://doi.org/10.3982/ecta9097>
- Natarajan, R., & Kass, R. E. (2000). Reference Bayesian Methods for Generalized Linear Mixed Models. *Journal of the American Statistical Association* ISSN:, 95(449), 227–237.
- Neal, R. M. (2003). Slice sampling. *Annals of Statistics*, 31(3), 705–767.
- Raftery, A. E., & Lewis, S. M. (1992). How Many Iterations in the Gibbs Sampler? In J. M. Bernardo, J. O. Berger, A. P. Dawid, & A. F. M. Smith (Eds.), *Bayesian Statistics 4* (pp. 763–773). Oxford University Press.
- Roos, M., Martins, T. G., & Held, L. (2015). *Sensitivity Analysis for Bayesian Hierarchical*. 2, 321–349. <https://doi.org/10.1214/14-BA909>
- Schütt, H. H., Harmeling, S., Macke, J. H., & Wichmann, F. A. (2016). Painfree and accurate Bayesian estimation of psychometric functions for (potentially) overdispersed data. *Vision Research*, 122, 105–123. <https://doi.org/10.1016/j.visres.2016.02.002>
- Spiegelhalter, D. J., Thomas, A., Best, N., & Lunn, D. (2003). *WinBUGS User Manual Version 1.4*. [https://faculty.washington.edu/jmiyamot/p548/spiegelhalter winbugs user manual.pdf](https://faculty.washington.edu/jmiyamot/p548/spiegelhalter%20winbugs%20user%20manual.pdf)
- Sturtz, S., Ligges, U., & Gelman, A. (2005). R2WinBUGS : A Package for Running WinBUGS from R. *Journal of Statistical Software*, 12(3). <https://doi.org/10.18637/jss.v012.i03>
- Sun, D., & Ye, K. (1995). Reference Prior Bayesian Analysis for Normal Mean Products. *Journal of the American Statistical Association*, 90(430), 589–597. <https://doi.org/10.1080/01621459.1995.10476551>
- Tso, I. F., Mui, M. L., Taylor, S. F., & Deldin, P. J. (2012). Eye-contact perception in schizophrenia: Relationship with symptoms and socioemotional functioning. *Journal of Abnormal Psychology*, 121(3), 616–627. <https://doi.org/10.1037/a0026596>
- Van de Schoot, R., Winter, S. D., Ryan, O., Zondervan-Zwijnenburg, M., & Depaoli, S. (2017). A Systematic Review of Bayesian Articles in Psychology: The Last 25 Years. *Psychological Methods*, 22(2), 217–239. <https://doi.org/10.1037/met0000100.supp>
- van Erp, S., Mulder, J., & Oberski, D. L. (2018). Prior sensitivity analysis in default bayesian structural equation modeling. *Psychological Methods*, 23(2), 363–388. <https://doi.org/10.1037/met0000162>
- Van Rheen, T. E., & Rossell, S. L. (2014). Let's face it: facial emotion processing is impaired in bipolar disorder. *Journal of the International Neuropsychological Society : JINS*, 20(2), 200–208. <https://doi.org/10.1017/S1355617713001367>

- Wagenmakers, E. J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., Selker, R., Gronau, Q. F., Šmíra, M., Epskamp, S., Matzke, D., Rouder, J. N., & Morey, R. D. (2018). Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. *Psychonomic Bulletin and Review*, 25(1), 35–57. <https://doi.org/10.3758/s13423-017-1343-3>
- Wasserstein, R. L., & Lazar, N. A. (2016). The ASA Statement on p -Values: Context, Process, and Purpose. *The American Statistician*, 70(2), 129–133. <https://doi.org/10.1080/00031305.2016.1154108>
- Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics*, 63(8), 1293–1313. <https://doi.org/10.3758/BF03194544>
- Yang, R., & Berger, J. O. (1994). Estimation of a Covariance Matrix Using the Reference Prior. *The Annals of Statistics*, 22(3), 1195–1211.
- Yang, Z., & Zhu, T. (2018). Bayesian selection of misspecified models is overconfident and may cause spurious posterior probabilities for phylogenetic trees. *Proceedings of the National Academy of Sciences of the United States of America*, 115(8), 1854–1859. <https://doi.org/10.1073/pnas.1712673115>
- Yao, B., Mueller, S. A., Grove, T. B., McLaughlin, M., Thakkar, K., Ellingrod, V., McInnis, M. G., Taylor, S. F., Deldin, P. J., & Tso, I. F. (2018). Eye gaze perception in bipolar disorder: Self-referential bias but intact perceptual sensitivity. *Bipolar Disorders*, 20(1), 60–69. <https://doi.org/10.1111/bdi.12564>
- Ye, K. (1994). Bayesian reference prior analysis on the ratio of variances for the balanced one-way random effect model. *Journal of Statistical Planning and Inference*, 41, 267–280.

Listing 1*BUGS Code for the Hierarchical Bayesian Model of Eye Gaze Perception, Part 1.*

```

1 #####
2 ### Hierarchical Bayesian Modeling of Gaze Perception Task ###
3 #####
4 ## N = number of subjects (assign value in data file)
5 ## Y = number of "yes" responses for each gaze angle
6 ## T = number of valid trials
7 ## G = gaze angle index (11 levels)
8 ## X = eye-contact signal strength (0, 0.1, 0.2, ..., 1.0)
9
10 model {
11
12     ## 1. Model the hierarchical data-generating process
13     for (i in 1:N) {          # Loop over N subjects
14
15         for (j in 1:11) {      # Loop over 11 gaze angles
16
17             # Number of yes responses Y for each gaze angle G follows a binomial distribution
18             # dependent on the unknown underlying probability theta and the number of trials T
19             Y[subj[i],G[j]] ~ dbin(theta[subj[i],G[j]], T[subj[i],G[j]])
20
21             # Link theta to eye-contact signal strength (X) using 2 parameters (alpha, beta) via a logit function
22             logit(theta[subj[i],G[j]]) <- alpha[subj[i]] + beta[subj[i]] * X[j]
23
24         } # Finish j loop
25
26         # Individual parameters come from normal distributions around group means
27         alpha[subj[i]] ~ dnorm(mu.alpha[grp[i]],tau.alpha[grp[i]])
28         beta[subj[i]] ~ dnorm(mu.beta[grp[i]],tau.beta[grp[i]])
29
30         # Compute individual threshold and slope using alpha and beta
31         thres[i] <- -alpha[i]/beta[i]
32         slope[i] <- beta[i]/4
33
34     } # Finish i loop
35

```

Listing 2*BUGS Code for the Hierarchical Bayesian Model of Eye Gaze Perception, Part 2.*

```

36     ## 2. Assign priors for the group parameters, mu's and tau's
37     for (k in 1:3) {           # Loop over 3 groups (1=HC, 2=BP, 3=SZ)
38
39         mu.alpha[k] ~ dflat()      # Flat prior for mean parameter mu.alpha (any value)
40         mu.beta[k] ~ dflat()       # Flat prior for mean parameter mu.beta (any value)
41         tau.alpha[k] ~ dgamma(1,1) # Gamma prior for precision parameter tau.alpha (must be positive)
42         tau.beta[k] ~ dgamma(1,1)  # Gamma prior for precision parameter tau.beta (must be positive)
43
44         # Compute group threshold and slope
45         grp.thres[k] <- -mu.alpha[k]/mu.beta[k]
46         grp.slope[k] <- mu.beta[k]/4
47
48     } # Finish k loop
49
50     ## 3. Compute group differences in threshold and slope
51     grp.thresdiff[1] <- grp.thres[1] - grp.thres[2] # Threshold: HC-BP
52     grp.thresdiff[2] <- grp.thres[1] - grp.thres[3] # Threshold: HC-SZ
53     grp.thresdiff[3] <- grp.thres[2] - grp.thres[3] # Threshold: BP-SZ
54
55     grp.slopediff[1] <- grp.slope[1] - grp.slope[2] # Slope: HC-BP
56     grp.slopediff[2] <- grp.slope[1] - grp.slope[3] # Slope: HC-SZ
57     grp.slopediff[3] <- grp.slope[2] - grp.slope[3] # Slope: BP-SZ
58
59 } # Finish model
60

```

Figure 1

Typical Workflow of Hierarchical Bayesian Modeling using MCMC Estimation

Note. Grey boxes (left) represent major steps in order. White boxes (right) represent steps to troubleshoot, finesse model, and conduct informal sensitivity analysis.

Figure 2

Probability of Perception of Self-directed Gaze as a Function of Eye-contact Signal Strength

Note. Round markers represent observed data (proportion of the number of trials in which self-directed gaze was endorsed out of the total number of trials completed). The curve represents a theoretical, logistic function generative of the data. The absolute threshold (x value at $y = 0.5$) indexes self-referential tendency, and the slope of the curve at $y = 0.5$ indexes perceptual precision.

Figure 3*Sample Face Stimuli and Group Responses of the Eye-contact Perception Task*

Note. A) Sample face stimuli. From left (gaze angle = 30°) to right (gaze angle = 0°), eye-contact signal strength increases in 10% increments from 0.0 (averted) to 1.0 (direct). B) Percentage of trials where self-directed gaze was perceived according to eye-contact signal strength in healthy controls (HC), bipolar disorder (BP), and schizophrenia (SZ) participants. Vertical bars indicate standard errors of means.

Figure 4*Schematic Structure of the Hierarchical Bayesian Model of the Eye Gaze Perception Task*

Note. Lowest level depicts the observed data Y_{ij} for each individual i and each gaze angle j . Y_{ij} is the total number of “yes” responses over T_{ij} trials, each of which has a true probability of θ_{ij} ; therefore, Y_{ij} follows a binomial distribution dependent on θ_{ij} and T_{ij} . Within each subject i , θ_{ij} varies with respect to eye-contact signal strength X_{ij} ; this relationship approximates a logistic function parameterized by α_i and β_i . The values of α_i and β_i come from two normal distributions centered around the respective mean of the subjects’ corresponding diagnostic group k . Each of these group-level normal distributions is parameterized by μ (mean parameter) and τ (precision parameter). Flat priors are assigned to the μ parameters as there is no specific expectations of their values; non-informative gamma distributions (parameterized by 1, 1) are selected as the prior distributions of the τ parameters because precision must be positive.

Figure 5

Posterior Density Plots of Group Self-referential Bias (Threshold)) and Perceptual Precision (Slope) Estimates using Hierarchical Bayesian Modeling

Note. Dashed vertical lines and numbers at the top indicate median values of the group means or group differences. Numbers at the bottom indicate the posterior probabilities of one-tailed group differences (i.e., area under the curve on the right hand side of the solid line at zero).

Figure 6

Individual Estimates Obtained using Least-squares Method vs. Hierarchical Bayesian Modeling (HBM)

Note. Upper panel: threshold (self-referential bias) estimates. Lower panel: slope (perceptual precision) estimates. Thick black lines indicate group means.

Figure 7

Scatterplots of Mayer-Salovey-Caruso Emotional Intelligence Test (MSCEIT) against Gaze Perception Metrics Estimated using Least-squares Method vs. Hierarchical Bayesian Modeling

Supplementary Information

Table 1S

Software and packages/toolboxes required for running analyses and visualizing results of the demonstration

Software	Package or toolbox	Information
WinBUGS		<p><u>Usage</u>: Standalone program to analyze statistical models using Markov Chain Monte Carlo (MCMC) methods.</p> <p><u>Download</u>: https://www.mrc-bsu.cam.ac.uk/software/bugs/the-bugs-project-winbugs/</p> <p><u>Documentation</u>: https://www.mrc-bsu.cam.ac.uk/wp-content/uploads/manual14.pdf</p>
R		<p><u>Usage</u>: Free statistical platform</p> <p><u>Download</u>: https://cran.r-project.org/mirrors.html</p>
	R2WinBUGS	<p><u>Usage</u>: To run WinBUGS or OpenBUGS from R.</p> <p><u>Documentation</u>: https://cran.r-project.org/web/packages/R2WinBUGS/vignettes/R2WinBUGS.pdf</p>
	coda	<p><u>Usage</u>: Provides functions for summarizing and plotting the output from MCMC simulations, and performing convergence diagnostic tests.</p> <p><u>Documentation</u>: https://cran.r-project.org/web/packages/coda/coda.pdf</p>
	tidyverse	<p><u>Usage</u>: Includes a number of useful packages for data analyses and visualization.</p> <p><u>Documentation</u>: https://www.tidyverse.org/packages/</p>
MATLAB		<p><u>Usage</u>: Programming platform using matrix-based language to analyze data, develop algorithms, and create applications. Requires an annual license fee; free short-term trial licenses are available.</p> <p><u>Download</u>: https://www.mathworks.com/products/matlab.html</p>
	Curve Fitting Toolbox	<p><u>Usage</u>: Provides functions for fitting curves and surfaces to data.</p> <p><u>Download</u>: https://www.mathworks.com/products/curvefitting.html</p>
	Statistics & Machine Learning Toolbox	<p><u>Usage</u>: Provides functions and apps to analyze and model data.</p> <p><u>Download</u>: https://www.mathworks.com/products/statistics.html</p>

Table 2S

Data and Code Files Required for the Demonstration Analyses

File name	Usage
Gaze_data.csv	Raw data used for both HBM and least-squared fitting
Bayes.R	R code used to run HBM, visualize results, and compare HBM and least-squared fitting results
HBM_model_one.txt	Text file containing model information of the first HBM model
HBM_model_one.txt	Text file containing model information of an alternative HBM model (with different prior values) as sensitivity analysis
LeastSquares.m	MATLAB code to run least-squares fitting
GAZE_LS_output.csv	This file can be generated by running LeastSquares.m; it contains individual gaze perception metrics estimated using least-squared fitting. It is provided along with other files because it is called for in later sections of the R code (Bayes.R) when comparing HBM and least-squared results.