

Meta-Analysis of Bayesian Analyses

Paul Blomstedt^{*,†,**}, Diego Mesquita^{*,‡,**}, Omar Rivasplata^{§,**}, Jarno Lintusaari^{*},
Tuomas Sivula^{*}, Jukka Corander^{¶,||}, and Samuel Kaski^{*,§}

Abstract. Meta-analysis aims to generalize results from multiple related statistical analyses through a combined analysis. While the natural outcome of a Bayesian study is a posterior distribution, traditional Bayesian meta-analyses proceed by combining summary statistics (i.e., point-valued estimates) computed from data. In this paper, we develop a framework for combining posterior distributions from multiple related Bayesian studies into a meta-analysis. Importantly, the method is capable of reusing pre-computed posteriors from computationally costly analyses, without needing the implementation details from each study. Besides providing a consensus across studies, the method enables updating the local posteriors *post-hoc* and therefore refining them by sharing statistical strength between the studies, without rerunning the original analyses. We illustrate the wide applicability of the framework by combining results from likelihood-free Bayesian analyses, which would be difficult to carry out using standard methodology.

Keywords: meta-analysis, Bayesian analyses, likelihood-free inference, approximate Bayesian computation, machine learning.

1 Introduction

Meta-analysis comprises a collection of methods designed to combine results from related statistical analyses carried out in separate studies. Traditionally, the per study results, taken as inputs for meta-analysis, are summary statistics computed from data, such as point estimates for a treatment’s effect size. To combine point estimates, there exist well-established Bayesian and classical methodologies (see, e.g., Higgins et al., 2009, and literature referenced therein). However, while the natural outcome of a Bayesian analysis is a posterior distribution, the task of combining posteriors obtained from multiple related studies into a meta-analysis has not been explored in the literature.

In standard Bayesian *random effects meta-analysis*, a summary statistic D_j , designed to convey information about a local effect of interest θ_j , is observed for each of a number of studies $j = 1, \dots, J$. Furthermore, the studies are assumed to be related via exchangeability of the local effects (see Gelman et al., 2013), making them conditionally independent given a shared overall effect φ . Note that the presence of φ is not a property of the original studies but rather a *post-hoc* assumption made by the meta-analyst.

arXiv: [1904.04484](https://arxiv.org/abs/1904.04484)

*Department of Computer Science, Aalto University

†WithSecure

‡School of Applied Mathematics, Getulio Vargas Foundation, diego.mesquita@fgv.br

§Department of Computer Science, University of Manchester

¶Department of Biostatistics, University of Oslo

||Department of Mathematics and Statistics, University of Helsinki

**Equal contribution.

Given the above assumptions, meta-analysis takes the shape of a hierarchical model:

$$\begin{aligned}\varphi &\sim Q \\ \theta_j &\sim P_\varphi \\ D_j &\sim F_{\theta_j}.\end{aligned}$$

In the case that D_j is a direct estimate of θ_j , the distribution F_{θ_j} is usually modeled as $\mathcal{N}(\theta_j, \hat{\sigma}_j^2)$, with $\hat{\sigma}_j^2$ estimated from data. Let us denote the density functions of F_{θ_j} , P_φ and Q as $f_j(\cdot|\theta_j)$, $p(\cdot|\varphi)$ and $q(\cdot)$, respectively. One of the primary goals of this model is to estimate the overall effect φ , which has marginal posterior density

$$q(\varphi|D_1, \dots, D_J) \propto \prod_{j=1}^J \left[\int f_j(D_j|\theta_j) p(\theta_j|\varphi) d\theta_j \right] q(\varphi). \quad (1)$$

It is worth mentioning that in some cases a meta-analysis can be carried out using data on the level of individual data points instead of aggregated statistics (Riley et al., 2010). In this case D_j can be taken to represent a set of data points. However, regardless of the granularity of the data and the ensuing distributional assumptions, the standard Bayesian meta-analysis still takes the form of a hierarchical model as described above, where the inputs for the model are point-valued observations.

In this work, we consider a setting in which the meta-analysis aims to combine *posterior distributions* on the local effects θ_j instead of summary statistics. In other words, instead of being point-valued, the observable result from each study is a distribution. While summarizing studies via posterior distributions is most intuitive from a Bayesian perspective, traditional Bayesian meta-analysis using the standard machinery of hierarchical models is not equipped to handle distributional observations.

1.1 Problem description and proposed solution

This paper sets out to create a novel methodology and underlying theory to combine posterior distributions on local effects into a meta-analysis. To the best of our knowledge, there is no previous work in the literature dealing with the problem just described, and therefore our solution (outlined next) is unprecedented.

Given this setting, and continuing to assume exchangeability of the local effects, we wish to address two questions: (i) How should we utilize locally computed posteriors to update our belief about the overall effect in analogy with (1)? And (ii) does the combined model allow us to refine the local posteriors *post-hoc* through sharing of statistical strength?

Let Π_j be a posterior distribution resulting from the j th study. We may identify Π_j with its density π_j when the latter exists, which arguably is the case in many practical situations. As in conventional random-effects meta-analysis, we assume that the local effects are related via exchangeability, and therefore the effects of the different studies are conditionally independent given an overall effect φ . Unlike in the conventional random-effects setting, which uses marginalized likelihood factors to update the prior, we propose

using expected likelihood factors as functionals of the observed distributions, each factor thus being the integral of $p(\theta_j|\varphi)$ against the density π_j . Accordingly, the proposed rule to update the prior $q(\varphi)$ on the overall effect φ , given π_1, \dots, π_J , is as follows:

$$q^*(\varphi|\pi_1, \dots, \pi_J) \propto \prod_{j=1}^J \left[\int p(\theta_j|\varphi) \pi_j(\theta_j) d\theta_j \right] q(\varphi). \quad (2)$$

While this expression may look surprising at first sight, we prefer to suspend judgement at this point and resume its discussion in Section 4 where its formal definition is presented and its properties are discussed; let us only highlight at this point that (2) is our answer to question (i) formulated in the beginning of this section.

We refer to the distribution defined by (2) as a *superposterior distribution* to convey the notion of it being the result of an update on top of posterior distributions π_1, \dots, π_J . Notice that there is no explicit data level below the level of the local effects θ_j in the update defining $q^*(\cdot|\pi_1, \dots, \pi_J)$. This is a substantial difference with the marginal posterior of the hierarchical model presented in (1). We show in Section 4 below that the proposed update emerges from an extended *non-hierarchical* model with measure-valued observations Π_1, \dots, Π_J . Interestingly, it can easily be shown that a standard non-hierarchical model (i.e., one with point-valued observations) can be recovered as a special case by choosing Π_j to be a Dirac delta measure; this is discussed further in Section 4, where we also show that (2) retains some basic properties of standard Bayesian posterior updates, such as order-invariance in successive updates and concentration as $J \rightarrow \infty$.

We shall also see later on that the proposed update rule defines a joint distribution over φ and $\theta_1, \dots, \theta_J$, from which we can derive a joint density for $\theta_1, \dots, \theta_J$ and, in turn, the latter can then be used to update the local posteriors π_j by marginalizing over all quantities except the one to be updated (further on this in Section 2.2); thus answering question (ii) formulated in the beginning of the section. This is an important additional feature of our method, which enables updating the local posteriors to share statistical strength between the studies, without the need of rerunning the original analyses.

1.2 An illustrative example

To gain an intuitive understanding of the proposed update rule (i.e., the rule outlined in (2), see also (15) below), consider the following model:

$$\begin{aligned} \varphi &\sim \text{Beta}(\alpha, \beta) \\ \theta_j &\sim \text{Bernoulli}(\varphi), \quad j = 1 \dots, J. \end{aligned}$$

Suppose initially (but unrealistically) that the θ_j are observable. Then, a standard application of Bayes' rule to update the prior distribution $\text{Beta}(\alpha, \beta)$, conditional on observed values θ_j , yields a posterior distribution which is $\text{Beta}(\alpha + r, \beta + J - r)$ with $r = \sum_{j=1}^J \theta_j$, by conjugacy. For reference, if $q(\varphi)$ denotes the density of the $\text{Beta}(\alpha, \beta)$ prior, we note that the posterior density for this model takes the form

$$q(\varphi|\theta_1, \dots, \theta_J) \propto \prod_{j=1}^J \varphi^{\theta_j} (1 - \varphi)^{1 - \theta_j} q(\varphi).$$




Figure 1: Updated density functions for φ with updates as per (2), each computed under the model $\varphi \sim \text{Beta}(2, 2)$, $\theta_j \sim \text{Bernoulli}(\varphi)$, $j = 1, \dots, 10$. The distributional observations are $\pi_j(k) = \Pi_j(\theta_j = k)$, $k \in \{0, 1\}$, assumed to be identical for all j . The different densities are obtained by varying $\pi_j(1)$ from 0 to 1 by increments of 0.1. The solid curves correspond to $\pi_j(1) = 0$ and $\pi_j(1) = 1$, equivalent to posteriors computed conditional on all θ_j fully observed with values 0 and 1, respectively. The dashed curve corresponds to $\pi_j(1) = 0.5$, and equals the prior $\text{Beta}(2, 2)$.

Next, suppose that the values of θ_j cannot be observed, but instead, the observables are distributions on these variables. Distributions on unobservable quantities are in practice often taken to be continuous, but for the purpose of this example we let the distributions on θ_j be discrete. Specifically, we assume they are binary distributions.

Thus, let the binary distributions π_j be the observables, where $\pi_j(k) = \Pi_j(\theta_j = k)$, $k \in \{0, 1\}$, and $\sum_{k=0}^1 \Pi_j(\theta_j = k) = 1$. We now wish to update the prior $\text{Beta}(\alpha, \beta)$ conditional on the observed π_j . Following (2), the update takes the form

$$q^*(\varphi | \pi_1, \dots, \pi_J) \propto \prod_{j=1}^J \left[\sum_{k=0}^1 \varphi^k (1 - \varphi)^{(1-k)} \pi_j(k) \right] q(\varphi).$$

While the modified model no longer permits analytical calculations through conjugacy in general, two analytically tractable special cases can be identified. Specifically, if $\pi_j(0) = \pi_j(1) = 0.5$ for all j , then the observed distributions are uninformative about the value of θ_j , and $q^*(\cdot | \pi_1, \dots, \pi_J)$ coincides with the prior $\text{Beta}(\alpha, \beta)$. On the other hand, if for each j there is a k_j with $\Pi_j(\theta_j = k_j) = 1$, then the θ_j 's are in effect fully observed, and $q^*(\cdot | \pi_1, \dots, \pi_J)$ from (2) equals the standard posterior $\text{Beta}(\alpha + r, \beta + J - r)$. A numerical illustration is provided in Figure 1.

1.3 Further discussion of the setting

Our framework builds on the same general idea of updating the distribution of a parameter, conditional on observations, as classical Bayesian inference. In a meta-analysis

context, the distribution to be updated is the distribution of the overall effect φ . However, in our setting the observations themselves are distributions, namely, distributions on the local effects θ_j ; whereas in previous meta-analysis settings the observations were summary statistics on the local effects. Thus, our setting deals with distributional observations, and the observed distributions may be interpreted as expressing beliefs about the values of the local effects of interest. Apart from being direct representations of the uncertainty inherent in parameter estimates, posterior distributions may include prior knowledge not present in the data, possibly obtained by expert elicitation (e.g. Albert et al., 2012). This is particularly important in problems where not enough data are available to inform a model about some of its parameters (e.g. Kuikka et al., 2014). This novel setting is therefore naturally suited for a meta-analysis that aims to combine the posterior distributions produced by related Bayesian analyses.

An advantage of processing study-specific information as readily computed posterior distributions, instead of re-computing the local posteriors within a hierarchical model, is that the complexities regarding inference and computation in each study can be abstracted away from the meta-analysis. This is an asset since the meta-analyst may neither have access to the study-specific data or the necessary implementation details, nor the computational resources to rerun each analysis as part of a larger hierarchical model that combines them. For example, in Section 3.3, we illustrate our method in a case where the local data models lack a tractable likelihood function. While extending likelihood-free inference to hierarchical models is possible, it is computationally costly (Turner and Van Zandt, 2014), and consequently, building a hierarchical model from each dataset would be very challenging. It is important to note that our discussion here focuses specifically on scenarios where computationally complex posterior distributions are readily available, rather than serving as a comprehensive comparison between our framework and hierarchical models in general.

1.4 Summary of contributions

The main contributions of our work are the following:

- Developing a framework for meta-analysis that combines posterior distributions. The framework proposes a rule to combine locally computed posteriors to update beliefs on an overall effect. The framework also shows how to refine the local posteriors *post-hoc* through sharing of statistical strength.
- Demonstrating the framework via numerical examples with both synthetic and real-world data. To illustrate the generality of the approach, we consider in the numerical examples the problem of combining results from analyses conducted using likelihood-free Bayesian analyses, where there is no closed-form likelihood to relate data or summary statistics to the quantity or effect of interest, which poses a challenge for traditional formulations of meta-analysis.
- Proving formally the properties of the proposed rule. This entails the development of a theoretical framework to justify the proposed rule, and proving that it satisfies the properties of posterior update rules such as order-invariance in successive updates and concentration in the limit of infinite observations.

1.5 Organization of the paper

The paper is structured as follows: Section 2 summarizes the main equations for practical use and provides further insight into the framework. Importantly, we also discuss the use of our framework to refine study-specific posteriors *post-hoc* through sharing of statistical strength across studies, and we draw connections to standard Bayesian meta-analysis. Section 3 illustrates our method with posterior distributions resulting from likelihood-free Bayesian analyses, which are typically computationally complex and may require careful hyperparameter tuning. Section 4 develops a novel framework for Bayesian inference with distributional observations, which is the basis for our novel meta-analysis approach. Section 5 presents a discussion of related work. The paper ends with concluding remarks in Section 6. The appendices (in the supplementary materials) provide proofs, an interpretation of our approach as message-passing in probabilistic graphical models, a computational strategy for our method, which can be implemented using general-purpose software, and a further experiment on a synthetic dataset.

2 Meta-analysis of Bayesian analyses

We are motivated by the problem of conducting meta-analysis for Bayesian analyses summarized as posterior distributions, and refer to our framework as *meta-analysis of Bayesian analyses* (MBA). The central updates of the framework are given in (4) and (5) below, which update beliefs on global and local effects, respectively.

To reiterate, a meta-analysis aims to combine inferences produced by various related studies into a consensus analysis. Thus, a local effect of interest θ_j is targeted by each of a number of studies $j = 1, \dots, J$. Judging the quantities θ_j to be exchangeable, the *meta-analyst* may formulate a model

$$\prod_{j=1}^J p(\theta_j | \varphi) q(\varphi), \quad (3)$$

with an appropriate prior q on the parameter φ . Note that this model is formulated *as if* the θ_j 's were fully observable quantities; this has an important role for conceptualization but is unrealistic in practical situations.

2.1 Updating belief on the overall effect

In our setting of interest, the observables consist of a set of posterior density functions $\{\pi_1, \dots, \pi_J\}$, each informing us about beliefs on the value of a corresponding quantity of interest $\{\theta_1, \dots, \theta_J\}$. Then, to update q based on the observed posterior densities, we apply the rule proposed in (2), recalled here for convenience:

$$q^*(\varphi) \propto \prod_{j=1}^J \left[\int_{\Theta} p(\theta_j | \varphi) \pi_j(\theta_j) d\theta_j \right] q(\varphi), \quad (4)$$

where we used the notation $q^*(\cdot) := q^*(\cdot|\pi_1, \dots, \pi_J)$ for brevity. The formal definition of the posterior distribution is given in Section 4 (cf. (15)), from which the density form given in (4) is obtained.

Remark 1 (Fixed-effects MBA). *We can derive a fixed-effects version of MBA as a special case when the $\Phi = \Theta$ and the likelihood $p(\cdot|\varphi)$ approaches a point mass at φ . In this case, the MBA posterior for φ becomes proportional to the product of study-specific distributions π_1, \dots, π_J and the prior q , i.e.,*

$$q^*(\varphi) \rightarrow Z^{-1} \prod_{j=1}^J \pi_j(\varphi) q(\varphi).$$

Note that when the prior $q(\varphi)$ is uniform, $q^(\varphi)$ becomes proportional to the product of study-specific distributions. If we also assume the π_1, \dots, π_J are Bayesian posteriors computed using flat priors, then $q^*(\varphi)$ is equivalent to multiplying the study-specific likelihoods and normalizing the result.*

2.2 Post-hoc refinement of local posteriors

In a meta-analysis context, the parameter φ often has an interpretation as the central tendency of some shared property of $\theta_1, \dots, \theta_J$, such as the mean or the covariance (or both jointly). As such, inference on φ is often of primary interest in providing a ‘consensus’ over a number of studies. As a secondary goal, we may also be interested in updating the posterior of any individual quantity θ_j , subject to the information provided by the posteriors on the remaining quantities. To do so, we may use the joint density of $\theta_1, \dots, \theta_J$ as starting point:

$$p^*(\theta_1, \dots, \theta_J) \propto \int_{\Phi} \prod_{j=1}^J [p(\theta_j|\varphi) \pi_j(\theta_j)] q(\varphi) d\varphi.$$

The formal definition of the joint distribution is given in Section 4 (cf. (16)), from which the density form given above is obtained. To update a local posterior, one can then marginalize over all quantities but the one to be updated. Let $\mathcal{J} := \{1, \dots, J\}$ be a set of indices and let $j' \in \mathcal{J}$ be an arbitrary index in this set. The density function $\pi_{j'}$ is then updated as follows:

$$\pi_{j'}^*(\theta_{j'}) \propto \int_{\Phi} p(\theta_{j'}|\varphi) \pi_{j'}(\theta_{j'}) \prod_{j \in \mathcal{J} \setminus j'}^J \left[\int_{\Theta} p(\theta_j|\varphi) \pi_j(\theta_j) d\theta_j \right] q(\varphi) d\varphi. \quad (5)$$

It turns out that (4) and (5) can also be formulated as an instance of message-passing in probabilistic graphical models (e.g. Yedidia et al., 2001). We discuss this view further in the supplementary material (Appendix C, Blomstedt et al., 2024).

Remark 2 (Non-degeneracy of local effects). *According to Section 4.1, the density $q^*(\varphi)$, defined in (4), will under suitable conditions become increasingly peaked around*

some point φ_0 , as $J \rightarrow \infty$. That $\pi_{j'}^*(\theta_j)$ does not behave similarly, becomes clear by the following considerations. First, we note that (5) is equivalent to

$$\pi_{j'}^*(\theta_{j'}) = Z_{j'}^{-1} \pi_{j'}(\theta_{j'}) \int_{\Phi} p(\theta_{j'}|\varphi) q^*(\varphi|\pi_1, \dots, \pi_{j'-1}, \pi_{j'+1}, \dots, \pi_J) d\varphi,$$

where Z_j is a normalizing constant. As $q^*(\varphi|\pi_1, \dots, \pi_{j'-1}, \pi_{j'+1}, \dots, \pi_J)$ becomes increasingly peaked around φ_0 , the integral in the above equation converges to $p(\theta_{j'}|\varphi_0)$. Consequently,

$$\pi_{j'}^*(\theta_{j'}) \rightarrow Z_{j'}^{-1} \pi_{j'}(\theta_{j'}) p(\theta_{j'}|\varphi_0),$$

which can only be degenerate if either $\pi_{j'}(\theta_{j'})$ or $p(\theta_{j'}|\varphi_0)$ is degenerate by design. Instead of degeneracy, $\pi_{j'}^*(\theta_{j'})$ exhibits shrinkage with respect to φ_0 .

2.3 Bayesian meta-analysis as a special case

It is straightforward to show that standard Bayesian random-effects and fixed-effects meta-analyses can be recovered as special cases of MBA. In its traditional formulation (e.g. Normand, 1999), *random-effects* meta-analysis (REMA) assumes that for each of J studies, a summary statistic, D_j , $j = 1, \dots, J$, has been observed, drawn from a distribution with study-specific mean $\mathbb{E}(D_j) = \theta_j$ and variance $\text{Var}(D_j) = \sigma_j^2$:

$$D_j \sim \mathcal{N}(\theta_j, \sigma_j^2), \quad (6)$$

where the approximation of the distribution of D_j by a normal distribution is justified by the asymptotic normality of maximum likelihood estimates. The variances σ_j^2 are directly estimated from the data, while the means θ_j are assumed to be drawn from some common distribution, typically

$$\theta_j \sim \mathcal{N}(\mu, \sigma_0^2),$$

where the parameters μ and σ_0^2 represent the average treatment effect and inter-study variation, respectively. For this model, the (μ, σ_0^2) are the global effect parameters, for which the general notation φ is used elsewhere in the text. *Fixed-effects* meta-analysis (FEMA) is a special case of REMA, where $\sigma_0^2 \rightarrow 0$, such that $\theta_1 = \theta_2 = \dots = \theta_J$.

The marginal posterior density for the parameters (μ, σ_0^2) in REMA can be written as

$$\begin{aligned} q(\mu, \sigma_0^2 | D_1, \dots, D_J) &\propto q(\mu, \sigma_0^2) \prod_{j=1}^J \int_{\Theta} N(D_j | \theta_j, \hat{\sigma}_j^2) N(\theta_j | \mu, \sigma_0^2) d\theta_j \\ &\propto q(\mu, \sigma_0^2) \prod_{j=1}^J \int_{\Theta} l_j(\theta_j; D_j) N(\theta_j | \mu, \sigma_0^2) d\theta_j, \end{aligned}$$

where $N(\cdot | \cdot, \cdot)$ denotes a Gaussian density function, $l_j(\theta_j; D_j)$ is the likelihood function of θ_j given D_j , and $\hat{\sigma}_j^2$ is the empirical variance of D_j .

To study the connection between the above posterior density and (4), assume that instead of a summary statistic D_j , each study has been summarized using a posterior distribution with density $\pi_j(\theta_j)$ over its study-specific effect parameter θ_j . If each of the posteriors has been independently computed under a Gaussian data model (see (6)), and using improper uniform priors $\nu_j(\theta_j) \propto 1$ for each $j = 1 \dots J$, then the local posterior densities are

$$\pi_j(\theta_j) = N(\theta_j | D_j, \hat{\sigma}_j^2) \propto \exp \left\{ -\frac{(D_j - \theta_j)^2}{2\hat{\sigma}_j^2} \right\} = l_j(\theta_j; D_j) \quad \forall j = 1, \dots, J,$$

resulting in both the marginal posterior of (μ, σ_0^2) for REMA, and the corresponding superposterior for MBA, being equivalent.

3 Numerical studies

In this section, we illustrate the proposed novel framework, meta-analysis of Bayesian analyses (MBA), with numerical examples.¹ To illustrate the generality of the approach, we consider the problem of combining results from analyses conducted using likelihood-free models. We emphasize, however, that MBA is not restricted to this class of models. In likelihood-free inference, the data can typically be summarized by a number of different statistics but there is no closed-form likelihood to relate these to the quantity or effect of interest, which poses a challenge for traditional formulations of meta-analysis. Using our framework, we directly utilize the local posteriors to build a joint model, from which we sample using a simple computational strategy, described in the supplementary material (Appendix D). In addition to modeling the shared central tendency of the inferred model parameters, we demonstrate that weakly informative or poorly identifiable posteriors for original studies can be updated *post-hoc* through joint modeling.

We begin with a brief review of likelihood-free inference using approximate Bayesian computation.

3.1 Likelihood-free inference using approximate Bayesian computation

Approximate Bayesian computation (ABC) is a paradigm for Bayesian inference in models which either entirely lack an analytically tractable likelihood function, or for which it is costly to compute. The only requirement is that we are able to sample data from the model by fixing values for the parameters of interest, as is the case for simulator-based models. In the basic form of ABC, simulations are run for parameter proposals drawn from a prior distribution. The parameter proposals whose simulated data x_θ match the observed data x_0 are collected and constitute a sample from the posterior distribution. It can be shown that this process is equivalent to accepting parameter proposals in proportion to their likelihood value, given the observed data, as is done in traditional rejection sampling.

¹Code available at github.com/weakly-informative/BA-MBA

In practice, the simulated data virtually never exactly match the observed data and very few samples from the posterior distribution would be acquired. This problem can be circumvented by loosening the acceptance condition to accept samples whose simulations yield results similar enough to the observed data. For this purpose, a dissimilarity function d and a scalar $\varepsilon > 0$ are defined such that a parameter proposal with respective simulation result x_θ is accepted if $d(x_\theta, x_0) \leq \varepsilon$. This function is often defined in terms of summary statistics $s(x_\theta)$ and $s(x_0)$. For example, d could be defined as the Euclidean distance between $s(x_\theta)$ and $s(x_0)$. The aforementioned relaxation results in samples being drawn from an *approximate* posterior instead of the actual posterior distribution, hence the name approximate Bayesian computation.

For a comprehensive introduction to ABC, see Marin et al. (2012). More recent developments are reviewed in Lintusaari et al. (2017). In the following numerical illustrations, the inputs for meta-analysis are posterior distributions for the separate local studies obtained using ABC. These likelihood-free inferences are implemented using the ELFI open-source software package (Lintusaari et al., 2018).

3.2 Example 1: MA(q) process

The moving-average process of order q , abbreviated MA(q), is a standard example in the literature on likelihood-free inference due to its simple structure but fairly complex likelihood and non-trivial relationship between parameters and observed data. In our first example, we use simulated data from a MA(q) process of order $q = 2$. Assuming zero mean, the process $(y_t)_{t \in \mathbb{N}^+}$ is defined as

$$y_t = \epsilon_t + \theta^{(1)}\epsilon_{t-1} + \theta^{(2)}\epsilon_{t-2}, \quad (7)$$

where $(\theta^{(1)}, \theta^{(2)}) \in \mathbb{R}^2$ and $\epsilon_s \sim \mathcal{N}(0, 1)$, $s \in \mathbb{Z}$.² The quantity of interest for which we conduct inference is $\theta = (\theta^{(1)}, \theta^{(2)})$. Following Marin et al. (2012), we use as prior for θ a uniform distribution over the set

$$\mathcal{T} \subset \mathbb{R}^2 \triangleq \{(\theta^{(1)}, \theta^{(2)}) \in \mathbb{R}^2 \mid -(\theta^{(2)} + 1) < \theta^{(1)} < \theta^{(2)} + 1, -1 < \theta^{(2)} < 1\},$$

which, by restriction of the parameter space, imposes a general identifiability condition for MA(q) processes. Inference for θ is then conducted using ABC with rejection sampling, taking as summary statistics the empirical autocovariances of lags one and two, denoted as $\hat{\gamma}_1$ and $\hat{\gamma}_2$, respectively. Euclidean distance of 0.1 is used as the acceptance threshold. Note that in this example, and in contrast to the next example in Section 3.3, posterior sampling from the current model is also possible using MCMC, although evaluation of the likelihood is difficult due to the unobserved ϵ_s .

To illustrate our meta-analysis framework, we first sample $J = 12$ realizations of θ using the following generating process:

$$\theta^{(1)} \sim \text{Unif}(0.4, 0.8), \quad \theta^{(2)} \sim \mathcal{N}(-0.4 + \theta^{(1)}, 0.04^2). \quad (8)$$

²To keep the notation consistent with Sections 2 and 4, we use subscripts of θ to index studies, while different dimensions of θ are indexed using superscripts.




Figure 2: Posterior samples for the parameters $(\theta_j^{(1)}, \theta_j^{(2)})$ of $J = 12$ MA(2) processes, obtained using rejection-sampling ABC. The yellow dots denote the true generating parameter values.

Given each realization $\theta_j = (\theta_j^{(1)}, \theta_j^{(2)})$, $j = 1, \dots, J$, we then generate a series of 10 data points, $(y_{j,1}, \dots, y_{j,10})$, according to (7). For each time-series, we independently conduct likelihood-free inference as described above, generating 1000 samples from the posterior. The computed posterior distributions along with their corresponding true parameter values are shown in Figure 2. It can be seen that the very limited information given by the data in each of the analyses leaves the posteriors with a considerable uncertainty.

For meta-analysis, we first specify a model for the study-specific effects $\theta_1, \dots, \theta_J$ as if they were observed quantities from an exchangeable sequence; see (3). As the true generating mechanism of the effects is typically unknown, the model must be specified according to the analyst's judgment. To reflect this, we will here model the generating process as a Gaussian distribution with parameters $\varphi = (\mu, \Sigma_0)$,

$$\theta_j \sim \mathcal{N}_2(\mu, \Sigma_0). \quad (9)$$

For μ and the covariance matrix Σ_0 , we use Gaussian and inverse Wishart priors, respectively,

$$\mu \sim \mathcal{N}_2(m, V) \quad \text{and} \quad \Sigma_0 \sim \mathcal{W}^{-1}(\nu, \Psi), \quad (10)$$

with

$$m = \begin{bmatrix} 1/2 \\ 0 \end{bmatrix}, \quad V = \begin{bmatrix} 0.4 & 0.05 \\ 0.05 & 0.1 \end{bmatrix}, \quad \text{and} \quad \nu = 4, \quad \Psi = \begin{bmatrix} 0.4 & 0.1 \\ 0.1 & 0.2 \end{bmatrix}.$$

The above values were chosen to provide reasonable coverage of \mathcal{T} , the constrained support of θ . Furthermore, ν was chosen as $\dim(\theta) + 2$ to directly yield Ψ as the mean of the inverse Wishart prior on Σ_0 .

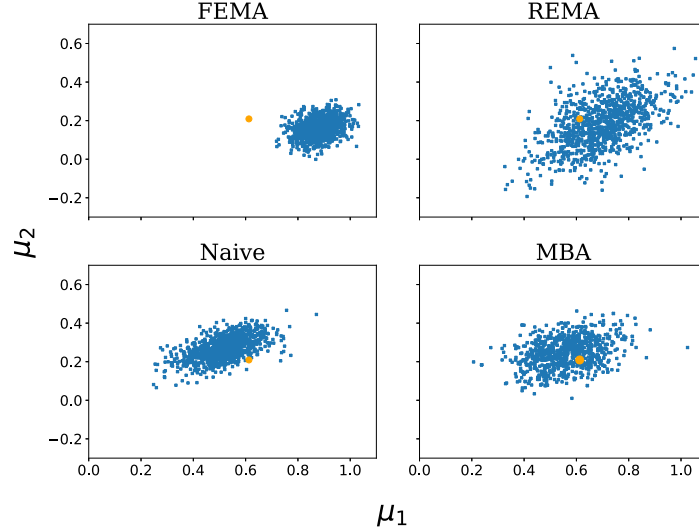


Figure 3: Posterior for the global effect μ using FEMA, REMA, a ‘naive’ model, and MBA. The yellow dot denotes the actual mean of the distribution used to generate the θ_j values for the MA(2) series.

After specifying the assumed generative model for $\theta_1, \dots, \theta_J$ according to (9)–(10), the following step is to incorporate the observed posteriors for each θ_j into the inference. We do this using the computational scheme detailed in Appendix D. Since the observed posteriors can be of arbitrary form, we carry out the computation in two stages. First, we draw samples from an approximate joint distribution, where each observed posterior is approximated with a suitable parametric distribution. Then, in the second stage, we refine the joint distribution using sampling/importance resampling (SIR; Smith and Gelfand, 1992) to correct for discrepancies between the approximated and true forms of the posteriors. In the current example, we apply the MBA updates using bivariate normal approximations to each of the $J = 12$ posteriors.

We compare MBA against results obtained using traditional random-effects meta-analysis (REMA), as specified in Section 2.3. The likelihood of REMA is given by the model

$$\hat{\theta}_j \sim \mathcal{N}_2(\theta_j, \hat{\Sigma}_j), \quad (11)$$

where the effect estimates $\hat{\theta}_j = (\hat{\theta}_j^{(1)}, \hat{\theta}_j^{(2)})$ are computed numerically using conditional sum of squares³ and the study-specific covariance matrices $\hat{\Sigma}_j$ are estimated using bootstrap. The hierarchical distribution on θ_j follows (9) and (10) above. Therefore, the essential difference between REMA and MBA is whether we combine this distribution with a likelihood function based on (11) or with the observed posterior on θ_j . The results of the comparison are presented in Figures 3–5.⁴

³Implemented in the `statsmodels` Python module.

⁴The experiment was repeated with multiple random seeds, yielding similar results.




Figure 4: Updated posteriors for study-specific effects obtained using MBA (red) on top of the ones originally obtained using independent ABC's (blue). Yellow dots denote the true parameter values.




Figure 5: Posterior marginals for study-specific effects obtained using REMA (red) on top of the ones originally obtained using independent ABC's (blue). Yellow dots denote the true parameter values.

Figure 3 shows the posterior distribution (superposterior in the case of MBA) for the global mean effect μ , obtained using four different models: in addition to MBA and REMA, we used fixed-effects meta-analysis (FEMA, which is a special case of REMA,

see Section 2.3), and a ‘naive’ model corresponding to ordinary Bayesian inference using the means of the observed posteriors on θ_j as observed data. In Section 4, we formulate the theoretical framework underlying MBA as an extension of Bayes’ rule from point-valued observations to distributional observations. Notice that the naive model can be thought of as a special case of MBA, where the probability mass of each of the observed posteriors concentrate at a single point, namely the posterior mean.

As expected, FEMA is clearly inappropriate in this situation, and results in a heavily biased and overly confident posterior. The naive model is less biased than FEMA as it makes use of information contained in the observed posteriors, but compared to MBA, it does not properly account for the uncertainty contained in them. Both REMA and MBA result in posteriors with more spread, still assigning reasonably high probability mass to the neighbourhood around the true parameter value. Figure 4 shows updated beliefs for the local effects, obtained using the update rule of (5) in Section 2. The updated beliefs exhibit shrinkage towards the global mean effect and, in this case, concentrate more accurately around the actual local effect values. On the other hand, many of the REMA posteriors for the local effects, shown in Figure 5, are biased and concentrate in regions away from true values.

3.3 Example 2: Tuberculosis outbreak dynamics

We now apply MBA to conduct meta-analysis of parameters regulating a stochastic birth-death (SBD) model proposed by Lintusaari et al. (2019), who used their model in a single-study setting to analyze tuberculosis outbreak data from the San Francisco Bay area, initially reported by Small et al. (1994). The goal of the analysis was to estimate disease transmission parameters from genotype data which, in contrast to outbreak models relying on count data, renders the likelihood function intractable and necessitates the use of likelihood-free inference (Tanaka et al., 2006). Furthermore, such models are often complex in relation to the available data, which may result in poor identifiability, as discussed by Lintusaari et al. (2016). To alleviate the problem, they formulated their model as a mixture of stochastic processes, taking into account the individual transmission dynamics of different subpopulations. In our analyses, we focus on two key parameters of the model, $R_0^{(1)}$ and $R_0^{(2)}$, which are the reproductive numbers for two subpopulations: those that are compliant and non-compliant to treatment, respectively.⁵

For our current experiment, we analyzed three additional data sets using the model of Lintusaari et al. (2019). These data sets reported tuberculosis outbreaks in Estonia (Kr   ner et al., 2001), London (Maguire et al., 2002) and the Netherlands (van Soolingen et al., 1999). For each data set, we independently conducted likelihood-free inference, generating 1000 samples from the posteriors. Following Lintusaari et al. (2019), we used the following six summary statistics for the ABC simulations: the number of observations, the total number of genotype clusters, the size of the largest cluster, the proportion of clusters of size two, the proportion of singleton clusters, and finally, the

⁵Note that Lintusaari et al. (2019) used the notation R_0 and R_1 , respectively. The standard notation for the reproductive number is R_0 .




Figure 6: Posteriors for the reproductive numbers $(R_0^{(1)}, R_0^{(2)})$, individually obtained using ABC in four different studies on tuberculosis outbreak dynamics.

average successive difference in size among the four largest clusters. The original publication additionally used two summary statistics on the observation times of the largest cluster. While these improved model identifiability, such information was not available for the additional data sets analyzed in our current experiment. We used a weighted Euclidean distance as dissimilarity function, with the same weights as Lintusaari et al. (2019).

Figure 6 shows the joint posterior distributions for the parameters $R_0^{(1)}$ and $R_0^{(2)}$, obtained individually for all four geographical locations using ABC. Compared to the San Francisco data, the posteriors computed on the remaining data sets, in particular London and the Netherlands, show symptoms of unidentifiability, with the posterior mass spread along lines between, and including, the boundaries of the parameter space. This can at least partly be attributed to these data sets being less informative than the San Francisco data set. A key question for our experiment then is whether we could borrow strength across the studies to improve the identifiability of the models computed on the remaining data sets. Additionally, it will be of interest to obtain an overall analysis of the central tendency of the reproductive numbers.

As in the previous experiment of Section 3.2, we first define a model for the local effects $\theta_j = (R_{0,j}^{(1)}, R_{0,j}^{(2)})$, and assign priors for the components of $\varphi = (\mu, \Sigma_0)$, where $\mu = (\mu_1, \mu_2)$ is the global mean effect and Σ_0 is the covariance matrix. This results in the model:

$$\theta_j \sim \mathcal{N}_2(\mu, \Sigma_0), \quad \mu_1 \sim \text{Gamma}(a_1, b_1), \quad \mu_2 \sim \text{Gamma}(a_2, b_2), \quad \Sigma_0 \sim \mathcal{W}^{-1}(\nu, \Psi).$$




Figure 7: Posteriors for the reproductive numbers $(R_0^{(1)}, R_0^{(2)})$ updated using MBA (red), plotted on top of the original, individually obtained posteriors (blue).

The hyperparameters were set as follows:

$$a_1 = 0.12, b_1 = 0.36, \quad a_2 = .030, b_2 = 0.05, \quad \text{and} \quad \nu = 4, \quad \Psi = \begin{bmatrix} 4 & -0.1 \\ -0.1 & 0.01 \end{bmatrix}.$$

We then follow the same procedure from the previous experiments to sample from the MBA-updated effects via SIR.

Note that due to the indirect nature in the relationship between the infection data and the parameters of interest, using the data to directly construct an estimator for the parameters of interest would be difficult. While this does not pose a challenge in our framework, it renders the application of traditional meta-analysis approaches infeasible.

Figure 7 shows the updated beliefs for the local effects $\theta_1, \dots, \theta_J$, after borrowing strength across the individual studies. While the updated beliefs clearly retain some of the individual characteristics of their original counterparts (e.g. a similar covariance structure), they exhibit a much more identifiable behavior. The superposterior of the overall mean of the reproductive numbers is shown in Figure 8.

In the supplementary material (Appendix E) we provide a synthetic version of these tuberculosis experiments, again showing that sharing statistical strength across studies yields local posteriors with improved identifiability compared to their original counterparts.




Figure 8: MBA joint superposterior for the overall mean effect $\mu = (\mu_1, \mu_2)$.

4 Bayesian inference with distributional observations

Overall, our approach to conduct meta-analysis from a set of related posteriors, MBA, can be seen more broadly as a framework for Bayesian inference with distributional observations. This section develops a foundation for this framework, which serves as the theoretical justification for MBA. As we will see, this framework is very general and encompasses the standard settings of point-valued observations as special cases, given the appropriate restrictions.

We start our presentation by treating $\theta_1, \dots, \theta_J$ as directly observable, point-valued random quantities. Then, we modify Bayes' rule to incorporate distributions Π_1, \dots, Π_J over $\theta_1, \dots, \theta_J$ by substituting the conventional likelihood factors with their expected values, arriving at the measure-theoretic version of (1).

Remark. *We switch notations slightly in this section: Instead of the probability density $q(\varphi)$ for φ we shall use the probability measure $Q(B)$ for sets B of possible values of φ ; and accordingly, the measure-theoretic style shall be used for integrals. The fact that $Q(B) = \int_B Q(d\varphi) = \int_B q(\varphi)d\varphi$ may help readers to switch back and forth between the measure-theoretic notations and their counterparts in terms of densities, when they exist. A similar remark applies to all other variables and parameters.*

To gain intuition, let us first suppose that $\theta_1, \dots, \theta_J$ are observable and exchangeable random quantities. Following standard theory, de Finetti's representation theorem (e.g. Schervish, 1995) states that if $\theta_1, \theta_2, \dots$ is an infinitely exchangeable sequence of random quantities taking values in a Borel space (Θ, \mathcal{A}) , then there exists a probability measure Q on a Borel space (Φ, \mathcal{B}) such that the joint distribution of the subsequence $\theta_1, \dots, \theta_J$, i.e. the *predictive distribution*, has the 'mixture' form

$$\mathbb{P}(\theta_1 \in A_1, \dots, \theta_J \in A_J) = \int_{\Phi} \prod_{j=1}^J P_{\varphi}(A_j) Q(d\varphi). \quad (12)$$

It is understood that the equation holds for any sets $A_1, \dots, A_J \in \mathcal{A}$. The details of the theorem are omitted for simplicity, but the essence of de Finetti's representation is establishing that exchangeability of $\theta_1, \theta_2, \dots$ is equivalent to conditional independence of these quantities given the value φ of a random parameter on a suitable space Φ . The

family $\{P_\varphi | \varphi \in \Phi\}$ consists of probability measures on Θ , indexed by the parameter φ . Assuming that these measures have density functions $p(\cdot | \varphi) := dP_\varphi / d\lambda$ with respect to a reference (Lebesgue or counting) measure λ , and Q has density q , one may rewrite (12) in the following (arguably friendlier) form:

$$\mathbb{P}(\theta_1 \in A_1, \dots, \theta_J \in A_J) = \int_{\Phi} \prod_{j=1}^J \left[\int_{A_j} p(\theta_j | \varphi) \lambda(d\theta_j) \right] q(\varphi) d\varphi. \quad (12')$$

For reference, the corresponding joint density, if it exists, then is as follows:

$$p(\theta_1, \dots, \theta_J) = \int_{\Phi} \prod_{j=1}^J p(\theta_j | \varphi) q(\varphi) d\varphi. \quad (13)$$

In the above standard setting, one may think of Q as being a prior distribution on the parameter φ , and the Bayesian learning process works through updating Q conditional on observed data. Following the above equations, the posterior distribution of φ , given observed values $\theta_1 = t_1, \dots, \theta_J = t_J$, has the form

$$Q(B | t_1, \dots, t_J) = \frac{\int_B \prod_{j=1}^J p(t_j | \varphi) Q(d\varphi)}{\int_{\Phi} \prod_{j=1}^J p(t_j | \varphi) Q(d\varphi)}. \quad (14)$$

The equation holds for any set $B \in \mathcal{B}$, where \mathcal{B} is the Borel σ -algebra on Φ . Let us emphasize that, as a preparatory step to build intuition, we have up until this point in the current section assumed that the θ_j are directly observable, point-valued random quantities. This is in contrast with standard notational conventions, where θ usually denotes an unobservable parameter (cf. the hierarchical meta-analysis model described in Section 1, where the point-valued observables are the summary statistics D_j).

We now consider a setting where the observables are distributions Π_1, \dots, Π_J on the values of $\theta_1, \dots, \theta_J$ respectively. This setting therefore departs from the standard setting built for point-valued observations, and requires a different formulation. We propose a simple modification as the basis to build an update rule for the new setting, substituting the likelihood factors $p(t_j | \varphi)$ of the standard setting with the *expected likelihood factors* $\int_{\Theta} p(\theta_j | \varphi) \Pi_j(d\theta_j)$, where the expectation is with respect to the observed distributions. We argue that these factors are reasonable for the setting we are considering, where the observables are distributions; not only because they are computable on the basis of the observed distributions, but also because they lead to an update rule that retains some basic properties of the classical Bayesian update rules, see Section 4.1.

The proposed expected likelihood factors now lead to an update of the form

$$Q^*(B | \Pi_1, \dots, \Pi_J) = \frac{\int_B \prod_{j=1}^J \left[\int_{\Theta} p(\theta_j | \varphi) \Pi_j(d\theta_j) \right] Q(d\varphi)}{\int_{\Phi} \prod_{j=1}^J \left[\int_{\Theta} p(\theta_j | \varphi) \Pi_j(d\theta_j) \right] Q(d\varphi)}, \quad (15)$$

for $B \in \mathcal{B}$, where we write $Q^*(\cdot | \Pi_1, \dots, \Pi_J)$ to denote conditioning on distributions, in analogy with conditioning on fully observed point values. We give more context for

this choice of notation below in Section 4.1. We refer to the measure defined by (15) as a *superposterior distribution* to distinguish it from (14) and to convey the notion of it being the result of an update on top of probability distributions Π_j . It is easy to see that (14) emerges as a special case of the superposterior by setting Π_j to be δ_{t_j} , the Dirac measure centered at t_j . This yields

$$\int_{\Theta} p(\theta_j|\varphi)\delta_{t_j}(d\theta_j) = p(t_j|\varphi),$$

such that $Q^*(\cdot|\delta_{t_1}, \dots, \delta_{t_J}) = Q(\cdot|t_1, \dots, t_J)$. Equation (15) also induces a joint distribution on the space $\Phi \times \Theta^J$, which we can marginalize with respect to Q to obtain a predictive distribution:

$$\mathbb{P}^*(\theta_1 \in A_1, \dots, \theta_J \in A_J) = \frac{\int_{\Phi} \prod_{j=1}^J [\int_{A_j} p(\theta_j|\varphi)\Pi_j(d\theta_j)]Q(d\varphi)}{\int_{\Phi} \prod_{j=1}^J [\int_{\Theta} p(\theta_j|\varphi)\Pi_j(d\theta_j)]Q(d\varphi)}. \quad (16)$$

Throughout this work, we assumed that every Π_j is a probability measure. It is interesting to note, however, that if we allow Π_j to be the Lebesgue (or counting) measure λ for all j , which corresponds to having a uniformly distributed—possibly improper—belief about the value of θ_j , then the updated measure $Q^*(\cdot|\Pi_1, \dots, \Pi_J)$ in (15) equals the prior probability measure Q . Moreover, with this choice of Π_j , (16) reduces to the standard predictive distribution in (12).

4.1 Theoretical properties

Two well-known properties of standard Bayesian inference, which are of practical relevance in our meta-analysis setting, are: (i) order-invariance in successive posterior updates, and (ii) posterior concentration. The former ensures that inferences conditional on exchangeable data are coherent. The latter tells us that the posterior distribution becomes increasingly informative about the quantity of interest, as we accumulate more observations. We now briefly discuss these properties in the context of the framework introduced above. Formal proofs are given in Appendix A and Appendix B.

Order-invariance in successive updates

Under the exchangeability assumption, standard Bayesian inference can be constructed as a sequence of successive updates, invariant to the order in which data are processed. It is straightforward to verify that the superposterior update defined in (15) shares this property of being order-invariant. The proof of the following proposition is given in the supplementary material (Appendix A).

Proposition 4.1. *The superposterior update defined in (15) is invariant to permutations of the indices $1, \dots, J$.*

Concentration in the limit of infinite observations

Asymptotic theory states that if a consistent estimator of the true value (or an optimal one in terms of KL divergence) of the parameter φ exists, then the posterior distribution in (14) concentrates in a neighborhood of this value, as $J \rightarrow \infty$ (e.g. Schervish, 1995). Here we discuss conditions under which the same property holds for the measure $Q^*(\cdot|\Pi_1, \dots, \Pi_J)$, defined in (15).

Our strategy is to first formulate a conceptual model for the observed distributions $\Pi_1, \dots, \Pi_J \in \mathcal{P}(\Theta)$; where the notation $\mathcal{P}(\Theta)$ stands for the space of probability measures on Θ . To this end, consider the following:

$$\varphi \sim Q \tag{17a}$$

$$\Pi_j \sim H_\varphi^{(j)}. \tag{17b}$$

Thus, $H_\varphi^{(j)}$ is the distribution of Π_j , given φ . We assume that these distributions are defined by their densities $h_j(\cdot|\varphi) := dH_\varphi^{(j)}/d\kappa$ with respect to a dominating measure κ on $\mathcal{P}(\Theta)$. Furthermore, assuming integrability, we take $h_j(\Pi_j|\varphi) := C_j \int_\Theta p(\theta_j|\varphi) \Pi_j(d\theta_j)$, where the C_j are the normalizing factors such that h_j integrates to 1. Then the probability measure $Q^*(B|\Pi_1, \dots, \Pi_J)$ defined in (15) takes the form

$$Q^*(B|\Pi_1, \dots, \Pi_J) = \frac{\int_B \prod_{j=1}^J h_j(\Pi_j|\varphi) Q(d\varphi)}{\int_\Phi \prod_{j=1}^J h_j(\Pi_j|\varphi) Q(d\varphi)}. \tag{18}$$

An essential difference between the standard posterior distribution in (14) and that of (18) is that, in the former, the observations are conditionally i.i.d., while in the latter they are conditionally independent but *non-identically* distributed. Note that the observed distributions Π_j might in general be the outputs of Bayesian analyses with different prior assumptions and design choices across the various studies. Furthermore, in our setting, the meta-analyst may not have access to the exact computational details of each study, and indeed, only access to the posterior distribution produced by each study is required. Hence, it is not reasonable to postulate homogeneity across the various studies. Nonetheless, from a practical point of view, we assume that the posteriors in each study target the same set of parameters. Another important difference to highlight is that (18) is based on *distributional* observations, as our setting considers measure-valued observations as the inputs, whereas (14) is based on point-valued observations, as per the standard setting.

One can interpret (18) as giving a representation of $Q^*(B|\Pi_1, \dots, \Pi_J)$ which is sufficient to establish concentration of this distribution in the limit $J \rightarrow \infty$. These results are given in the supplementary material (Appendix B, where Section B.1 covers discrete parameter spaces, and Section B.2 covers general parameter spaces).

For the case of discrete parameter spaces, the proof follows the basic structure found in many sources; interested readers are referred to Gelman et al. (2013, App. B, p. 588) and Bernardo and Smith (1994, Sec. 5.3, p. 286). Note however that these references

assumed i.i.d. point-valued observations; whereas our setting deals with measure-valued observations which are independent but non-identically distributed. A key step in proving the concentration result is to establish a limit for the sums of log-likelihood ratios $\sum_{j=1}^J \log(h_j(\Pi_j|\varphi_0)/h_j(\Pi_j|\varphi))$ for $\varphi \neq \varphi_0$, where φ_0 is a distinguishable parameter value which meets some technical requirements. This limit is established using a version of the strong law of large numbers that is valid for non-identically distributed random variables (cf. Sen and Singer, 1993, Theorem 2.3.10). For continuous parameter spaces the argument is more intricate and requires the technical assumptions to be adjusted in some ways; however, the basic building blocks of the proof are analogous to those of the discrete case. Interested readers are referred to Schervish (1995, Theorem 7.80) for the result for continuous parameter space and i.i.d. point-valued observations. Our result extends this theorem to the setting of observations which are measure-valued and independent but non-identically distributed.

In conclusion, if the superposterior distribution $Q^*(\cdot|\Pi_1, \dots, \Pi_J)$ can be written in the form of (18) and some technical conditions are met, then concentration when $J \rightarrow \infty$ can be established, using similar arguments to those that establish the concentration results for classical posterior distributions.

5 Related work

To the best of our knowledge, meta-analysis frameworks for combining posterior distributions have not been presented before. Nonetheless, similar elements can be found in previous works.

5.1 Bayesian meta-analysis

In the context of conventional Bayesian meta-analysis, Lunn et al. (2013) introduced a two-stage computational approach for fitting hierarchical models to individual-level data. In the first stage, posterior samples for local parameters are generated independently for each study, enabling study-specific complexities to be dealt with separately. In a second stage, the samples are used as proposals in Metropolis-Hastings updates within a Gibbs algorithm to fit the full hierarchical model. While the paper focuses on a specific computational approach for hierarchical models, involving no propagation of local prior knowledge into the joint model, the idea of utilizing independently computed posterior samples to fit a joint model is common to our framework.

In more recent work, Rodrigues et al. (2016) developed a hierarchical Gaussian process prior to model a set of related density functions, where grouped data in the form of samples assumed to be drawn under each density function are available. Despite a superficial similarity between their work and ours, the inferential goals in these works are very different. Specifically, the former is concerned with nonparametric estimation of group-specific densities, which is shown to be useful when the sample sizes in some or all of the groups are small. Thus, in cases where the number of posterior samples per study is limited (e.g. due to computational reasons), the method of Rodrigues et al.

(2016) could be used within our framework to provide density estimates for the initial beliefs.

5.2 Frequentist meta-analysis

In the frequentist paradigm, Xie et al. (2011) introduced a general framework for the combination of confidence distributions, a concept loosely related to Bayesian posteriors and deeply rooted in the frequentist paradigm (see Schweder and Hjort, 2002). In the same realm, Liu et al. (2015) combine analyses from related results using a product of confidence distributions, under the assumption that known linear functions map the local effects to the overall effect.

5.3 Bayesian updating using uncertain evidence

Our work builds on the interpretation that each $\pi_j(\theta_j)$ is a probabilistic representation of belief (Bernardo and Smith, 1994), which reflects uncertainty about the value of the corresponding local effect θ_j . Updating prior knowledge (in our case regarding φ) subject to uncertain or ‘soft’ evidence, represented as probability distributions on the local effects instead of point estimates, has been extensively studied as a philosophical topic in both statistics and artificial intelligence. The most well-known of such update rules, Jeffrey’s rule of conditioning (Diaconis and Zabell, 1982; Jeffrey, 2004; Smets, 1993; Zhao and Osherson, 2010), computes the updated probability for an event as a weighted average of the posterior probabilities under all possible values of the evidence. Due to its construction, Jeffrey’s rule is applicable in simple discrete cases but becomes computationally infeasible for more complex models with continuous variables. It is also well known that Jeffrey’s update rule is in general not invariant to the order in which data are processed, unlike our proposed update, as discussed in Section 4.1.

5.4 Distributed Bayesian inference

We finally note that while our framework can be described as an approach to combine (information provided by) a set of pre-computed posteriors, it is substantially different from works in the field of *Embarrassing Parallel* (or divide-and-conquer) inference (e.g. Neiswanger and Xing, 2017; Wang et al., 2018). In these works, given a partitioned dataset $\mathcal{X} = \cup_{j=1}^J \mathcal{X}_j$, a posterior $p(\theta|\mathcal{X})$ is approximated by first independently computing a *subposterior* $p(\mathcal{X}_j|\theta)p(\theta)^{1/J}$ for each data partition, and then combining these as $p(\theta|\mathcal{X}) \propto \prod_{j=1}^J p(\mathcal{X}_j|\theta)p(\theta)^{1/J}$. In contrast, the goal of our framework is to use posteriors computed over *distinct* quantities $\theta_1, \dots, \theta_J$ related through exchangeability, to provide information about a shared global quantity φ .

6 Conclusion and future work

The natural outcome of a Bayesian analysis is a posterior distribution over quantities of interest. Meta-analysis methods combine results from multiple related statistical studies

into a consensus analysis, and to the best of our knowledge all previous approaches to conduct meta-analyses have used summary statistics (or individual-level data) as the observable inputs. In this paper, we have developed a framework that uses distributions as the observable inputs to conduct meta-analyses, thus enabling the combination of posteriors from multiple related Bayesian analyses. The framework builds on standard Bayesian inference over exchangeable quantities, by treating each observed posterior as data observed with uncertainty. We have also shown how the framework can be used to update independently computed, study-specific posteriors *post-hoc* through sharing of statistical strength. This is analogous to how estimates for related parameters share strength between each other in hierarchical models (e.g. Gelman et al., 2013).

In many fields, it has become common practice to store data sets in dedicated repositories to be reused for the benefit of the entire research community. Given the view taken in this work, that posterior distributions can be seen as a special kind of observational data, we believe that in many cases it would be equally beneficial to make full posterior distributions available for reuse. This would enable posteriors from potentially time-consuming and costly Bayesian analyses to be used as a basis for new studies. Indeed, even if the original data, the model and the code implementing it were available, reproducing posterior distributions could require a substantial computational effort. While in this work we have used posterior samples as a representation for study-specific posteriors, the optimal format for storing and sharing posterior distributions is a question for future research. In addition to ways of making posteriors publicly available, more research is needed on developing methods to make appropriate use of the information they provide. The current work is a first step in this direction and our hope is that it will inspire other researchers to make further advances to this end.

Furthermore, we believe exploring MBA as a framework for Bayesian inference using uncertain evidence is a promising direction for future work. For instance, in the context of multiple imputations, MBA could be used to directly incorporate the posterior predictive distributions for the missing values as beliefs, avoiding repeated MCMC runs for multiple imputed datasets (Zhou and Reiter, 2010). Given that MBA accommodates the combination of posteriors and point estimates in a single meta-analysis, another relevant direction for future work is to extend the analysis to these settings; however further research would be required for a careful study of the properties of such hybrid combinations of posteriors and point estimates.

Funding

This work was supported by the Research Council of Finland (Flagship programme: Finnish Center for Artificial Intelligence FCAI, 359207), EU Horizon 2020 (European Network of AI Excellence Centres ELISE, grant agreement 951847), UKRI Turing AI World-Leading Researcher Fellowship (EP/W002973/1), Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro FAPERJ (SEI-260003/000709/2023), São Paulo Research Foundation FAPESP (2023/00815-6), Conselho Nacional de Desenvolvimento Científico e Tecnológico CNPq (404336/2023-0). We also acknowledge the computational resources provided by the Aalto Science-IT Project from Computer Science IT.

Supplementary Material

Supplementary Material (DOI: [10.1214/24-BA1465SUPP](https://doi.org/10.1214/24-BA1465SUPP); .pdf).

References

- Albert, I., S. Donnet, C. Guihenneuc-Jouyaux, S. Low-Choy, K. Mengersen, and J. Rousseau (2012). Combining expert opinions in prior elicitation. *Bayesian Analysis* 7(3), 503–532. [MR2981623](#). doi: <https://doi.org/10.1214/12-BA717>. 5
- Bernardo, J. M. and A. F. M. Smith (1994). *Bayesian Theory*. John Wiley & Sons, Inc. [MR1274699](#). doi: <https://doi.org/10.1002/9780470316870>. 20, 22
- Blomstedt, P., Mesquita, D., Rivasplata, O., Lintusaari, J., Sivula, T., Corander, J., and Kaski, S. (2024). “Supplementary Material for “Meta-analysis of Bayesian analyses”” *Bayesian Analysis*. doi: <https://doi.org/10.1214/24-BA1465SUPP>. 7
- Diaconis, P. and S. L. Zabell (1982). Updating subjective probability. *Journal of the American Statistical Association* 77(380), 822–830. [MR0686405](#). doi: <https://doi.org/10.1080/01621459.1982.10477893>. 22
- Gelman, A., J. Carlin, H. Stern, D. Dunson, A. Vehtari, and D. Rubin (2013). *Bayesian Data Analysis* (Third ed.). London: Chapman & Hall/CRC. [MR3235677](#). doi: <https://doi.org/10.1201/b16018>. 1, 20, 23
- Higgins, P. T., S. G. Thompson, and D. J. Spiegelhalter (2009). A re-evaluation of random-effects meta-analysis. *Journal of the Royal Statistical Society. Series A* 172(1), 137–159. [MR2655609](#). doi: <https://doi.org/10.1111/j.1467-985X.2008.00552.x>. 1
- Jeffrey, R. (2004). *Subjective Probability: The Real Thing*. Cambridge University Press. [MR2063965](#). doi: <https://doi.org/10.1017/CB09780511816161>. 22
- Krüüner, A., S. E. Hoffner, H. Sillastu, M. Danilovits, K. Levina, S. B. Svenson, S. Ghebremichael, T. Koivula, and G. Källénus (2001). Spread of drug-resistant pulmonary tuberculosis in Estonia. *Journal of Clinical Microbiology* 39(9), 3339–3345. doi: <https://doi.org/10.1128/JCM.39.9.3339-3345.2001>. 14
- Kuikka, S., J. Vanhatalo, H. Pulkkinen, S. Mäntyniemi, and J. Corander (2014). Experiences in Bayesian inference in Baltic salmon management. *Statistical Science* 29(1), 42–49. [MR3201845](#). doi: <https://doi.org/10.1214/13-STS431>. 5
- Lintusaari, J., P. Blomstedt, T. Sivula, M. Gutmann, S. Kaski, and J. Corander (2019). Resolving outbreak dynamics using approximate Bayesian computation for stochastic birth-death models [version 2; peer review: 2 approved]. *Wellcome Open Research* 4(14). 14, 15
- Lintusaari, J., M. U. Gutmann, R. Dutta, S. Kaski, and J. Corander (2017). Fundamentals and recent developments in approximate Bayesian computation. *Systems Biology* 66(1), e66–e82. 10

- Lintusaari, J., M. U. Gutmann, S. Kaski, and J. Corander (2016). On the identifiability of transmission dynamic models for infectious diseases. *Genetics* 202(3), 911–918. doi: <https://doi.org/10.1534/genetics.115.180034>. 14
- Lintusaari, J., H. Vuollekoski, A. Kangasrääsiö, K. Skytén, M. Järvenpää, P. Marttinen, M. U. Gutmann, A. Vehtari, J. Corander, and S. Kaski (2018). ELFI: Engine for likelihood-free inference. *Journal of Machine Learning Research* 19(16), 1–7. MR3862423. 10
- Liu, D., R. Y. Liu, and M. Xie (2015). Multivariate meta-analysis of heterogeneous studies using only summary statistics: Efficiency and robustness. *Journal of the American Statistical Association* 110(509), 326–340. MR3338506. doi: <https://doi.org/10.1080/01621459.2014.899235>. 22
- Lunn, D., J. Barrett, M. Sweeting, and S. Thompson (2013). Fully Bayesian hierarchical modelling in two stages, with application to meta-analysis. *Journal of the Royal Statistical Society. Series C* 62(4), 551–572. MR3083911. doi: <https://doi.org/10.1111/rssc.12007>. 21
- Maguire, H., J. W. Dale, T. D. McHugh, P. D. Butcher, S. H. Gillespie, A. Costetsos, H. Al-Ghusein, R. Holland, A. Dickens, L. Marston, P. Wilson, R. Pitman, D. Strachan, F. A. Drobniewski, and D. K. Banerjee (2002). Molecular epidemiology of tuberculosis in London 1995–7 showing low rate of active transmission. *Thorax* 57(7), 617–622. doi: <https://doi.org/10.1136/thorax.57.7.617>. 14
- Marin, J.-M., P. Pudlo, C. P. Robert, and R. J. Ryder (2012). Approximate Bayesian computational methods. *Statistics and Computing* 22(6), 1167–1180. MR2992292. doi: <https://doi.org/10.1007/s11222-011-9288-2>. 10
- Neiswanger, W. and E. Xing (2017). Post-inference prior swapping. In *Proc. 34th Int. Conf. Machine Learning*, Volume 70, pp. 2594–2602. 22
- Normand, S.-L. T. (1999). Meta-analysis: formulating, evaluating, combining, and reporting. *Statistics in Medicine* 18, 321–359. doi: [https://doi.org/10.1002/\(SICI\)1097-0258\(19990215\)18:3<321::AID-SIM28>3.0.CO;2-P](https://doi.org/10.1002/(SICI)1097-0258(19990215)18:3<321::AID-SIM28>3.0.CO;2-P). 8
- Riley, R. D., P. C. Lambert, and G. Abo-Zaid (2010). Meta-analysis of individual participant data: rationale, conduct, and reporting. *BMJ* 340. 2
- Rodrigues, G., D. J. Nott, and S. Sisson (2016). Functional regression approximate Bayesian computation for Gaussian process density estimation. *Computational Statistics & Data Analysis* 103, 229–241. MR3522629. doi: <https://doi.org/10.1016/j.csda.2016.05.009>. 21
- Schervish, M. J. (1995). *Theory of Statistics*. New York: Springer. MR1354146. doi: <https://doi.org/10.1007/978-1-4612-4250-5>. 17, 20, 21
- Schweder, T. and N. L. Hjort (2002). Confidence and likelihood. *Scandinavian Journal of Statistics* 29(2), 309–332. MR1909788. doi: <https://doi.org/10.1111/1467-9469.00285>. 22
- Sen, P. K. and J. M. Singer (1993). *Large Sample Methods in Statistics: An Introduction*

- with Applications*. Boca Raton: Chapman & Hall/CRC. MR1293125. doi: <https://doi.org/10.1007/978-1-4899-4491-7>. 21
- Small, P. M., P. C. Hopewell, S. P. Singh, A. Paz, J. Parsonnet, D. C. Ruston, G. F. Schecter, C. L. Daley, and G. K. Schoolnik (1994). The epidemiology of tuberculosis in San Francisco – a population-based study using conventional and molecular methods. *New England Journal of Medicine* 330(24), 1703–1709. doi: <https://doi.org/10.1056/NEJM199406163302402>. 14
- Smets, P. (1993). Jeffrey’s rule of conditioning generalized to belief functions. In *Proc. 9th Conf. Uncertainty in Artificial Intelligence (UAI-93)*, pp. 500–505. MR1229424. doi: <https://doi.org/10.1016/B978-1-4832-1451-1.50065-2>. 22
- Smith, A. F. M. and A. E. Gelfand (1992). Bayesian statistics without tears: A sampling–resampling perspective. *The American Statistician* 46(2), 84–88. MR1165566. doi: <https://doi.org/10.2307/2684170>. 12
- Tanaka, M. M., A. R. Francis, F. Luciani, and S. A. Sisson (2006). Using approximate Bayesian computation to estimate tuberculosis transmission parameters from genotype data. *Genetics* 173(3), 1511–1520. doi: <https://doi.org/10.1534/genetics.106.055574>. 14
- Turner, B. M. and T. Van Zandt (2014). Hierarchical approximate Bayesian computation. *Psychometrika* 79, 185–209. MR3255116. doi: <https://doi.org/10.1007/s11336-013-9381-x>. 5
- van Soolingen, D., M. W. Borgdorff, P. E. W. de Haas, M. M. G. G. Sebek, J. Veen, M. Dessens, K. Kremer, and J. D. A. van Embden (1999). Molecular epidemiology of tuberculosis in the Netherlands: A nationwide study from 1993 through 1997. *The Journal of Infectious Diseases* 180(3), 726–736. doi: <https://doi.org/10.1086/314930>. 14
- Wang, D., Z. Zeng, and Q. Liu (2018). Stein variational message passing for continuous graphical models. In *Proc. 35th Int. Conf. Machine Learning*, Volume 80, pp. 5219–5227. 22
- Xie, M., K. Singh, and W. E. Strawderman (2011). Confidence distributions and a unifying framework for meta-analysis. *Journal of the American Statistical Association* 106(493), 320–333. MR2816724. doi: <https://doi.org/10.1198/jasa.2011.tm09803>. 22
- Yedidia, J. S., W. T. Freeman, and Y. Weiss (2001). Generalized belief propagation. In *Advances in Neural Information Processing Systems 13*, pp. 689–695. MIT Press. 7
- Zhao, J. and D. Osherson (2010). Updating beliefs in light of uncertain evidence: Descriptive assessment of Jeffrey’s rule. *Thinking & Reasoning* 16(4), 288–307. doi: <https://doi.org/10.1080/13546783.2010.521695>. 22
- Zhou, X. and J. P. Reiter (2010). A note on Bayesian inference after multiple imputation. *The American Statistician* 64(2), 159–163. MR2757007. doi: <https://doi.org/10.1198/tast.2010.09109>. 23