

Prior knowledge elicitation: The past, present, and future

Petrus Mikkola^{*}, Osvaldo A. Martin^{*,||}, Suyog Chandramouli^{*,†}, Marcelo Hartmann[†], Oriol Abril Pla[†], Owen Thomas[‡], Henri Pesonen[‡], Jukka Corander^{‡,†,**}, Aki Vehtari^{*}, Samuel Kaski^{*,¶,††}, Paul-Christian Bürkner^{§,††} and Arto Klami^{†,††,‡‡}

Abstract. Specification of the prior distribution for a Bayesian model is a central part of the Bayesian workflow for data analysis, but it is often difficult even for statistical experts. Prior elicitation transforms domain knowledge of various kinds into well-defined prior distributions, and offers a solution to the prior specification problem, in principle. In practice, however, we are still fairly far from having usable prior elicitation tools that could significantly influence the way we build probabilistic models in academia and industry. We lack elicitation methods that integrate well into the Bayesian workflow and perform elicitation efficiently in terms of costs of time and effort. We even lack a comprehensive theoretical framework for understanding different facets of the prior elicitation problem.

Why are we not widely using prior elicitation? We analyze the state of the art by identifying a range of key aspects of prior knowledge elicitation, from properties of the modelling task and the nature of the priors to the form of interaction with the expert. The existing prior elicitation literature is reviewed and categorized in these terms. This allows recognizing under-studied directions in prior elicitation research, finally leading to a proposal of several new avenues to improve prior elicitation methodology.

Keywords: prior elicitation, prior distribution, informative prior, Bayesian workflow, domain knowledge.

1 Introduction

Bayesian statistics uses probabilistic models, formalized as a set of interconnected random variables following some assumed probability distributions, for describing observations. Designing a suitable model for a given data analysis task requires both significant statistical expertise and domain knowledge, and is typically carried out as an iterative

^{*}Helsinki Institute of Information Technology, Department of Computer Science (PM,OAM,AK,SK) and Department of Communications and Networking (SC), Aalto University, Finland

[†]Helsinki Institute of Information Technology, Department of Computer Science (MH,OA,AK,SC) and Department of Mathematics and Statistics (JC), University of Helsinki, Finland

[‡]Institute of Basic Medical Sciences, University of Oslo (JC,HP) and Akershus Universitetssykehus (OT), Norway

[§]Cluster of Excellence SimTech, University of Stuttgart, Germany

[¶]Department of Computer Science, University of Manchester, UK

^{||}Instituto de Matemática Aplicada San Luis, CONICET-UNSL, Argentina

^{**}Parasites and Microbes, Wellcome Sanger Institute, UK

^{††}Equal contribution

^{‡‡}Corresponding author, arto.klami@helsinki.fi

process that involves repeated testing and refinement. This process is called the *Bayesian workflow*; see [Gelman et al. \(2020\)](#) for a recent detailed formalization partitioning the process into numerous sub-workflows focusing on different facets of the process, such as model specification, inference and model validation.

We focus on one central part of that Bayesian workflow: the choice of prior distributions for the parameters of the model. In particular, we discuss approaches to *eliciting* knowledge from a domain expert to be converted into prior distributions suitable for use in a probabilistic model, rather than assuming the analyst can specify the priors directly. The fundamental goal of this *expert knowledge or prior elicitation* process (defined in Section 2.1) is to help practitioners design models that better capture the essential properties of the modelled problem; good elicitation tools could also help in the additional goal of fostering wide-spread adoption of probabilistic modelling by reducing the required statistical expertise. An ideal prior elicitation approach would simultaneously make model specification faster, easier, and better at representing the knowledge of the expert. It is hoped that availability of good prior elicitation tools would qualitatively transform the process of prior specification within the Bayesian modelling workflow, analogously to what *probabilistic programming* languages and their efficient model-agnostic algorithms have done for model specification and inference (e.g., [Stan Development Team, 2021](#); [Salvatier et al., 2016](#); [Ge et al., 2018](#)).

Prior elicitation has a long history dating back to the 1960s ([Winkler, 1967](#)), and excellent textbook accounts ([O’Hagan et al., 2006](#)), surveys and reviews ([Garthwaite et al., 2005](#); [O’Hagan, 2019](#)) are available. Despite the established problem formulation and broad scientific literature on methods for eliciting priors in different special cases – often for some particular model family – we are still lacking practical tools that would routinely be used as part of the modelling workflow. While a few actively developed tools for interactive prior elicitation exist, exemplified by SHELF ([Oakley and O’Hagan, 2019](#)) and `makemyprior` ([Hem et al., 2021](#)), their active user-base remains a tiny fraction of people regularly applying probabilistic models. Instead, most practitioners use rather ad hoc procedures to specify and modify the priors (e.g. [Sarma and Kay, 2020](#)), building on personal expertise and experience, ideally learned by following literature on prior recommendations ([Stan Development Team \(2021\)](#), logistic regression: [Gelman et al. \(2008\)](#); [Ghosh et al. \(2018\)](#), hierarchical models: [Gelman \(2006\)](#); [Simpson et al. \(2017\)](#); [Chung et al. \(2015\)](#), Gaussian random fields: [Fuglstad et al. \(2019\)](#), and autoregressive processes: [Sørbye and Rue \(2017\)](#)).

We discuss reasons for the still limited impact of prior elicitation research on prior specification in practice, and propose suggestions for a range of research directions that need to be pursued to change the situation. Our main claim is that we are still fairly far from having practical prior elicitation tools that could significantly influence the way probabilistic models are built in academia and industry. To improve over the current state, coordinated research involving expertise from multiple disciplines is needed. This paper is both our call for these experts to join their efforts, as well as a concrete guide for future research. Consequently, the paper is written both for people already developing prior elicitation techniques and for people working on specific complementary problems, who we are encouraging to contribute to the common goal. For people looking for

practical methods for prior elicitation in their own modelling problems, we unfortunately cannot yet provide very concrete solutions, but we are looking for your feedback on the requirements and desired goals.

As will be clarified later, several interconnected elements hinder the uptake of prior elicitation methods. Some of these are purely *technical* properties of the elicitation algorithms, relating to limited scope in terms of models that prevents their use in general probabilistic programming, or ability to only address univariate priors, sequentially, rather than jointly eliciting all priors of a model. Some are more *practical*, such as many of the approaches still being too difficult for non-statistical experts to use, and lack of good open source software that integrates well with the current probabilistic programming tools used for other parts of the modelling workflow. Finally, some of the aspects are more *societal*: The concrete value of prior elicitation has not yet been adequately demonstrated in highly visible case studies, and hence end-users do not know to request better approaches and decision-makers have not invested resources for them.

Critically, these issues are highly interconnected. For building large-scale demonstrations of the practical value of prior elicitation in visible applications, we would already need to have high-quality software that integrates with existing modelling workflows, as well as elicitation methods capable of efficiently eliciting priors for models of sufficient complexity. Given that the field is currently falling short of achieving any of these aspects, we argue that significant coordinated effort is needed before we can make concrete recommendations on best practices for elicitation in any given instance. We can largely work in parallel towards mitigating these issues, but it is important to do this in a coordinated manner, typically so that researchers with complementary scientific expertise work together to address the most closely connected elements. For instance, an ideal team for designing the software tools would combine at least computer engineers, statisticians, interface designers and cognitive scientists, to guarantee that the most important aspects for all dimensions are accounted for.

To proceed towards practical recommendations, we start by identifying seven key dimensions that characterize the prior elicitation challenge and possible solutions for it, to provide a coherent framework for discussing the matter. We inspect prior elicitation from the perspectives of (1) properties of the prior distribution itself, (2) the model family and the prior elicitation method’s dependence on it, (3) the underlying elicitation space, (4) how the method interprets the information provided by the expert, (5) computation, (6) the form and quantity of interaction with the expert(s), and (7) the assumed capability of the expert, both in terms of their domain knowledge and statistical understanding of the model. We discuss all of these fundamental dimensions in detail (Section 2.3), identifying several practical guidelines on how specific characteristics for each of them influence the desired properties for the elicitation method. We also provide a review of existing elicitation methods to highlight gaps in the available literature, but for more comprehensive reviews at earlier stages of the literature, we recommend consulting O’Hagan et al. (2006) and Garthwaite et al. (2005).

Building on this framework, we proceed to make recommendations for future research, by characterizing in more detail the current blockers listed above, and outlining our current suggestions on what kind of research is needed to resolve the issues. These

recommendations are necessarily on a relatively high abstraction level, but we hope they still provide a tangible starting point for people coming from outside the current prior elicitation research community. In particular, we discuss easy-to-use software that integrates with open probabilistic programming platforms as a necessary requirement for practical impact, already outlining a possible architecture and key components for such a system. We emphasize the need for considerably extended user evaluation for verifying that the methods have practical value.

2 Prior Elicitation

2.1 What is prior elicitation?

Specifying prior probability distributions over variables of interest (such as model’s parameters) is an essential part of Bayesian inference. These distributions represent available information regarding values of the variables prior to considering the current data at hand. *Prior elicitation* is one way to specify priors and refers to the process of eliciting the subjective knowledge of domain experts in a structured manner and expressing this knowledge as prior probability distributions (Garthwaite et al., 2005; O’Hagan et al., 2006). This involves not only actually gathering the information from an expert, but also any computational methods that may be needed to transform the collected information into well-defined prior probability distributions.

It should be noted that while prior elicitation is the focus of our article, it is only one of many ways to specify priors. Analysts may directly specify informative prior distributions based on a variety of sources including relevant literature or databases when the parameters have “fairly concrete real-world referents” (Gelman and Shalizi, 2013). For instance, in medicine data-based priors have been adopted in large scale (Bartoš et al., 2021), while there are situations where expert-knowledge-based priors are more appropriate, such as with “parameter settings that are unverifiable from the data to hand”, as discussed by Dallow et al. (2018). Another example is that when the aim of inference is to select between different theoretical models with differing assumptions, priors can be defined to encode the theoretical assumptions about model parameters (Lee and Vanpaemel, 2018).

Besides encoding subjective knowledge, there are other grounds for specifying priors. For instance, priors can be chosen such that they affect the information in the likelihood as weakly as possible (*noninformative priors*), yield smoother and more stable inferences (*regularizing priors*), or yield ‘asymptotically acceptable’ posterior inference (*reference priors*) (Gelman et al., 2017; Kass and Wasserman, 1996). While we acknowledge the validity of these approaches as well, we do not discuss them in more detail in this article due to our specific goal of investigating the state of prior elicitation, not prior specification in general. However, we repeat the general observation that flat priors (such as $N(0, 10^6)$) sometimes used by practitioners should be avoided, due to problems in posterior inference (Carlin, 2000; van Dongen, 2006; Gelman, 2006; Gelman et al., 2013, 2017; Gelman and Yao, 2020; Smid and Winter, 2020).

A good prior elicitation procedure takes into account that the distribution being

elicited is part of a model and cannot be viewed in isolation. The Bayes’s rule connects the prior $p(\theta)$ to the posterior $p(\theta|y)$ within the context of the likelihood $p(y|\theta)$,

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)}, \quad (2.1)$$

where the observables and the parameters are denoted by y and θ , respectively. The goal of prior elicitation is to elicit $p(\theta)$ from an expert. In line with [Gelman et al. \(2017\)](#), we note that the likelihood $p(y|\theta)$ partially determines the scale and the range of reasonable values for θ . In that respect, prior elicitation differs from the elicitation for evidence-based decision-making (e.g. [Kennedy et al., 2008](#); [Brownstein et al., 2019](#)) or expert systems (e.g. [Studer et al., 1998](#); [Wilson and Corlett, 2005](#)).

A common elicitation process involves two persons, called expert and analyst. We follow the convention that the expert is referred to as a female and the analyst as a male ([Oakley and O’Hagan, 2007](#)), and use the term *analyst* instead of facilitator to emphasize that the analyst can play many roles simultaneously ([O’Hagan et al., 2006](#), Section 2.2.1), for instance, as a statistician and a facilitator. The *facilitator* is an expert in the process of elicitation. He can take an active role such as manage dialogue between the expert(s) or a more passive role such as assisting in the elicitation between the expert and an elicitation algorithm. Not all elicitation methods require a human facilitator but instead he/it is built into the elicitation software (see an interesting alternative definition by [Kahle et al., 2016](#)). The *expert* refers to the domain expert, who is also called as a substantive expert. She has relevant knowledge about the uncertain quantities of interest, such as the model parameters or observables. For more about the definition and recruitment of the experts, see ([Bolger, 2018](#)).

2.2 Why isn’t the use of prior elicitation widespread (yet)?

Specifying good priors can have significant effect on the outcome of the whole modelling process, and support is clearly needed ([Robert, 2007](#); [O’Hagan, 2019](#)), yet prior elicitation techniques are not routinely used within practical Bayesian workflows. The most natural explanation for this is that the current solutions are simply not sufficient for the needs of the people building statistical models and doing practical data analysis. We are not aware of structured literature looking into these aspects in a systematic way, and hence we provide here our evaluation of the main reasons prior elicitation has not yet entered daily use in the statistical modelling community. The goal here is to provide a high-level overview to the main issues we have identified based on both the scientific literature and our experiences while interacting with the modelling community, in particular the user bases of Stan ([Stan Development Team, 2021](#)), brms ([Bürkner, 2017](#)) and PyMC ([Salvatier et al., 2016](#)). Not all claims of this subsection are supported by direct scientific evidence.

As briefly mentioned in the Introduction, we believe the reasons for limited use of prior elicitation are multifaceted and highly interconnected, consisting of three primary reasons:

- **Technical:** We do not know how to design accurate, computationally efficient, and general methods for eliciting priors for arbitrary models.
- **Practical:** We lack good tools for elicitation that would integrate seamlessly to the modelling workflow, and the cost of evaluating elicitation methods is high.
- **Societal:** We lack convincing examples of prior elicitation success stories, needed for attracting more researchers and resources.

By the *technical* dimension we refer to the quality and applicability of the prior elicitation methods and interfaces, for instance in terms of what kinds of models and priors are supported, and how accurate and efficient the algorithms are. An ideal solution would work in general cases, provide an easy interface for the expert to provide information, accurately reproduce the true knowledge of an expert, and be computationally efficient and reliable to be incorporated into the modelling workflow. In Section 3 we will cover the current literature and discuss the limitations of the current technical solutions, effectively concluding that we do not yet have prior elicitation techniques that would reach a sufficient level of technical quality.

By the *practical* dimension we refer to concrete tools ready to be used by practitioners. On a rough level, a prior elicitation method consists of some interface for interacting with the expert and the computational algorithm for forming the prior. Often the interfaces proposed for the task have been fairly general, but the majority of the research on the computational algorithms has been dedicated to methods that are only applicable for specific models or forms of priors. Their practical value remains limited. Even though some examples of model-agnostic elicitation methods exist and some of the initiatives have been developed for extended periods of time, we are still nowhere near a point where prior elicitation tools would routinely be used as part of the modelling process. Besides the technical reasons mentioned above, one major reason is that the tools have not been integrated as parts of the broadly used modelling ecosystems, but rather as isolated tools with their own interface conventions, modelling languages, and internal data formats. To put it briefly, a person building a model e.g. in `Stan` cannot launch an elicitation interface to elicit priors for their specific model, and in the extreme case there might not even exist any tools applicable to their model. In Section 4.5, we will outline directions for overcoming this practical challenge.

Another practical issue concerns evaluation of prior elicitation methods. Even though the basis of evaluating the elicitation methodologies is well established (see Section 4.4), the practical value of prior elicitation is extremely difficult and costly to evaluate. Already isolated studies demonstrating e.g. improved task completion time, compared to manual prior specification, for some prototypical model require careful empirical experimentation with human users. While this is a common practice in human computer interaction research, for statisticians it requires quite notable additional effort and expertise. More importantly, for the real cases of interest the evaluation setup is unusually complex because the modelling process itself is a highly complex iterative process that requires statistical expertise and takes a long time, possibly weeks or months. Any empirical evaluation of the value of prior elicitation requires enrolling high-level experts who are

tasked to carry out complex operations with systems that are unfamiliar to them, and possible significant individual differences in the way models are built necessitate large user bases for conclusive evidence. This can only be done once the practical software is sufficiently mature, and even then is both difficult and expensive. The problem is naturally not unique to prior elicitation, but instead resembles e.g. the cost of evaluating the effect of new medical practices that require medical professionals testing new procedures that may also result in worse treatments, or evaluation of new educational policies and practices. However, justifying the cost is often easier for these tasks that are considered critically important for the society.

Following the above discussion on cost of evaluation, we believe that there is a significant *societal* argument explaining the limited use. As detailed in this article, the task is challenging and consequently requires significant resources spanning several scientific fields, combining fundamental statistical and algorithmic research with cognitive science and human-computer interaction for forming the solid basis with high-quality software integration and costly evaluation. This requires significant resources, yet the current research is driven solely by academia and the field has remained somewhat small. To some extent this can be attributed to the long history of avoiding strong subjective priors in quest for objective scientific knowledge, but we argue that lack of broad interest more specifically on prior elicitation is largely because the value of prior elicitation has not been concretely demonstrated in breakthrough applications of societal importance. Without such demonstrations, the level of interest for these tools will remain low outside the statistics research community. However, already isolated examples of significant scientific or economical breakthroughs building on explicit use of prior elicitation could lead to increase in both research funding (e.g. in form of public-private partnerships for applying the technology) and in particular in interest for open source software development.

This argumentation, unfortunately, is very circular in nature. To boost interest in developing better prior elicitation methods, we would need a high-profile demonstration of their value, but establishing that demonstration would require access to high-quality solutions that integrate well with the modelling tools. However, it is important to realize that the demonstrator can likely be done well before having a robust general-purpose solution. Instead, it is sufficient to have proper software and interface integration of prior elicitation with one modelling ecosystem that is already used for addressing societally important modelling questions, combined with elicitation algorithms that work for the specific types of models needed. For instance, Bayesian models developed within the **Stan** ecosystem played a significant role in modelling the effect of various interventions had on the spread of COVID-19 (Flaxman et al., 2020), and demonstrating the value of prior elicitation in such a context would likely have been sufficient for raising the awareness of this research direction.

2.3 Prior elicitation hypercube

The interdisciplinary nature of the prior elicitation problem, and therefore scattered coverage of the topic (see Section 3), makes it difficult to obtain an overall perspective to the current state of research. To provide a frame of reference, we identify seven key




Figure 1: **Prior Elicitation Hypercube**. The seven dimensions (D1-D7) of the hypercube.

dimensions that characterize the prior elicitation problem. Together the dimensions form a *prior elicitation hypercube*, depicted in Figure 1, that both helps discuss the current literature in a more structured manner and enables identifying understudied directions. The first two dimensions (D1-D2) cover the Bayesian model itself (prior and likelihood). Dimensions D3-D5 specify key proprieties of an elicitation algorithm, such as in which space the elicitation is conducted (D3), how the expert’s input is modelled (D4), and how to deal with the computational issues (D5). The last dimensions D6-D7 cover what is assumed about the expert(s) and the interaction with them. In Section 3, the current prior elicitation literature is reviewed and categorized in terms of these dimensions.

D1: Properties of the prior distribution itself. Two properties of the prior distribution have attained considerable attention in the literature: dimensionality and parametric vs nonparametric nature. Dimensionality is about the number of parameters: is the prior univariate or multivariate? Eliciting multivariate (joint) distributions is a more complex task than eliciting univariate (marginal) distributions (O’Hagan et al., 2006). It is not enough to elicit the prior in a parameter-by-parameter manner because it is the joint behavior that affects inferences, and hence it is the joint distribution that must be considered (Gelman et al., 2017). Maybe because of the challenge of eliciting multivariate priors, univariate elicitation has been studied more, even though most models have more than one parameter and hence multiparameter prior elicitation is really needed (Gelman et al., 2020, Section 7.3).

The second property is about whether the prior comes from some parametric family or is nonparametric. The main strand of the prior elicitation literature is about the elicitation of parametric prior distributions, in which case the follow-up question is to which parametric family the prior belongs. The choice of the family is closely connected to the likelihood/model (see Section 3.4), since the natural conjugate family is often considered. On the other end, there is an important line of research on nonparametric prior elicitation that has been built upon Gaussian processes (Oakley and O’Hagan, 2007).

D2: The model family and the method’s dependence on it. The underlying probabilistic model and the data analysis task in which it is applied, significantly impact the choice of the prior elicitation method. A bulk of the prior elicitation research addresses the elicitation of parameters of some specific model or model class (Section 3.4). In our literature review (Section 3), we found a relatively small number of model-agnostic prior elicitation methods, but to name some, we refer the reader to Gelfand et al. (1995); Oakley and O’Hagan (2007); Hartmann et al. (2020). To promote adoption of prior elicitation in applications, it is highly desirable that a prior elicitation method is not model-specific, or at least is applicable to a wide range of models, and we strongly encourage research in this direction as acknowledged earlier by Kadane and Wolfson (1998). Furthermore, the underlying data analysis task may indicate which parameters are of interest, and thus need to be elicited (in the context of a chosen Bayesian model), and which may be less relevant (Clemen and Reilly 2001, p.292; Stefan et al. 2020).

D3: Elicitation space. Prior elicitation is about eliciting expert knowledge to form a prior distribution for model parameters. Hence, it is not surprising that the majority of prior elicitation research is focused on querying values of parameters, or quantities directly related to parameters, from the experts (Section 3.1). In this case, we say that the underlying elicitation space is the *parameter space*. This implies that the expert has to have at least some intuition about the meaning of the parameters (interpretability) and about their natural scales. However, this cannot be assumed in all cases. The elicitation of parameters of Bayesian neural networks serves as an extreme example, but the topic is profound, and will be discussed in Section 3.2. In many cases it may be more beneficial to query the expert about something else, such as model observables. In this case we say that the underlying elicitation space is the *observable space*. The *model observables* are variables (e.g. model outcomes) that can be observed and directly measured, in contrast to *latent variables* (e.g. model parameters) that only exist within the context of the model and are not directly observed. Kadane and Wolfson (1998) made a similar dichotomy where they called elicitation in the parameter space *structural elicitation*, and elicitation where the expert makes judgments about “the dependent variable given various values of the predictor variables”, *predictive elicitation*. Predictive elicitation is closely related to elicitation in the observable space but it is not the same, because the latter does not require the model to have both dependent and independent variables (e.g. Coolen, 1992; Hughes and Madden, 2002; Gaoini et al., 2009), and does not assume a “regression likelihood” (Kadane and Wolfson, 1998, p.5) or the prior predictive distribution (3.1) in general. For instance, an ‘elicitation likelihood’ can be used for connecting the expert’s knowledge on observables to the parameters of interest (see Section 4.1).

D4: Elicitation model. There are fundamental differences between elicitation methods in terms of how the information provided by the expert is interpreted. Since early prior elicitation research (Winkler, 1967; Bunn, 1978), the dominant approach has been “fitting a simple and convenient distribution to match the elicited summaries” (O’Hagan et al., 2006). This ‘*fitting approach*’ does not assume any specific mechanism on how the expert data is generated, and for instance, inconsistencies in the data are reconciled by least-square minimization. *Overfitting* in elicitation means eliciting more summaries than needed to fit a parametric distribution (O’Hagan et al., 2006; Hosack

et al., 2017), in which case inconsistencies may appear. Overfitting itself is desirable because it allows for imprecision in the elicited summaries, and the fitted compromise prior may be expected in practice to yield a more faithful representation of the expert’s knowledge (O’Hagan et al., 2006).

There is an alternative to the fitting approach, and to how to deal with inconsistencies. The elicitation of an expert’s knowledge can be treated as any other Bayesian inference problem where the analyst’s posterior belief about the expert’s knowledge is updated in the light of received expert data (Lindley et al., 1979; Gelfand et al., 1995; Gosling, 2005; Oakley and O’Hagan, 2007; Daneshkhah et al., 2006; Gosling et al., 2007; Oakley et al., 2010; Moala and O’Hagan, 2010; Micallef et al., 2017; Hartmann et al., 2020). This standpoint is similar to supra-Bayesian pooling found in the literature of aggregating knowledge of multiple experts (Section 3.6). We follow the latter terminology, even if there is only a single expert to be elicited, and say that such an elicitation method follows the *supra-Bayesian approach*. In this approach, inconsistencies in the elicited data are accounted for by a noise mechanism built into the elicitation likelihood.

D5: Computation. Computation is needed in many parts of an elicitation algorithm, such as in constructing prior from the elicited data and in active elicitation (Section 3.5), and the computational aspects need to be accounted for in practical tools. One-shot (D6) elicitation methods that follow the fitting approach (D4) and solely operate in the parameter space are often computationally efficient and can easily be incorporated into a practical workflow. In contrast, iterative (D6) predictive (D3) elicitation methods that operate in both spaces and require repeated computation of the prior predictive distribution require considerably more attention in terms of computational efficiency, both because of increased computational cost and the need for fast response time for convenient user experience.

D6: The form and quantity of interaction with the expert(s). The sixth dimension is about the interaction between the expert(s) and the analyst. On the one hand, the form of *assessment tasks* that an expert performs (and similar aspects relating to the interaction modality with a single expert) is important. On the other hand, if there is more than one expert, the format in which the experts interact is also important. For instance, the behavioral aggregation method used in the SHELF protocol (Oakley and O’Hagan, 2019) encourages the experts to discuss their opinions, and to settle upon group consensus judgments, to which a single prior distribution is fitted (O’Hagan, 2019). Eliciting the knowledge of a group of experts, and how to combine the elicited information into a single *aggregate distribution*, is a well established topic.

Concerning a single expert, there are choices to be made about the interaction modality of the elicitation. The expert can be either queried in a one-shot manner (*one-shot elicitation*), or iteratively where the expert’s input affects what is queried next (*iterative elicitation*). For instance, a prior elicitation algorithm that exploits active elicitation (Section 3.5) is iterative. We distinguish iterative elicitation from *interactive elicitation* that entails interaction with the elicitation system (Kadane et al., 1980), such as the system updating a visualization of a prior distribution based on a slider position controlled by the expert (Jones and Johnson, 2014). It is not obvious at all in which form the information should be elicited from the expert. Several things need to be taken

simultaneously into account, such as what assessment tasks are informative, easy for the expert, and computationally feasible. Fortunately, there exist research on different assessment tasks, which are discussed in Section 3.1. Since the assessment task should also consider psychological and cognitive aspects of a person being elicited, this topic is also related to the next dimension.

D7: Capability of the expert in terms of their domain knowledge and statistical understanding of the model. If the expert has only vague domain knowledge, the elicitation algorithm should validate the provided information, for instance, by using ‘seed variables’ as in Cooke’s method (Cooke, 1991). On the other hand, if the expert has no solid statistical training, she may not be able to provide reliable probabilistic assessments. In that case, we can resort to assessment tasks that do not require probabilistic input, such as querying likely hypothetical samples (Casement and Kahle, 2018; Casement, 2017). Even if the expert has both excellent statistical and domain knowledge, she may be inclined to commit to popular cognitive biases and to use cognitive shortcuts (heuristics) in her reasoning, as well documented by Tversky and Kahneman (1974). This line of research is known as *heuristics and biases* in prior elicitation, and it is intrinsically connected to psychology (Hogarth, 1975).

3 Current Main Lines of Research

Prior elicitation has a long history dating back to the 1960s when Edwards et al. (1963) introduced “Problem of prior probabilities” to the psychological community, followed by a prior elicitation article by Winkler (1967). The interdisciplinary topic spans several fields including statistics, psychology, economics, decision sciences, and more recently, machine learning. Similarly, there are numerous applications in many fields including clinical and pharmaceutical research (Best et al., 2020; Alhussain and Oakley, 2020a,b; Montague et al., 2019; Grigore et al., 2013; Wolpert, 1989; Haakma et al., 2011; Tompsett et al., 2020; Legedza and Ibrahim, 2001), management sciences (Galway, 2007; Kallen and Kuniewski, 2009), environmental sciences (Wolfson et al., 1996; Coles and Tawn, 1996; Hirsch et al., 1998; Choy et al., 2009; Hammond et al., 2001; Al-Awadhi and Garthwaite, 2006; León et al., 2003), social sciences (Gill and Walker, 2005; Lombardi and Nicoletti, 2012; Del Negro and Schorfheide, 2008), business (Crosby, 1980; Bajari and Ye, 2003; Ming et al., 2016) and physics (Craig et al., 1998). This complicates systematic literature review, as prior elicitation is addressed from different perspectives, in different contexts, and under different terminology. In decision sciences, elicitation research has been conducted using the phrases ‘expert judgement’, ‘expert knowledge elicitation’ and ‘probability encoding’ (Spetzler and Holstein, 1975; Dias et al., 2018; Böcker et al., 2010; Hanea et al., 2021). In statistics, the main branch of the prior elicitation research is presented in (O’Hagan et al., 2006; Garthwaite et al., 2005; O’Hagan, 2019). For other relevant reviews, see (Chesley, 1975; Wolfson, 1995; Jenkinson, 2005; EFSA, 2014; Oakley and O’Hagan, 2019; Stefan et al., 2020), and in clinical research (Johnson et al., 2010a; Azzolina et al., 2021).

In this section, we present the current main lines of research in prior elicitation through the lens of the prior elicitation hypercube (Section 2.3). The literature review was

done in two stages. For forward reference searching, we used the Google Scholar search engine. The search strings used were “ALLINTITLE: PRIOR ELICITATION”, “ALLINTITLE: PRIOR ASSESSMENT”, and “ALLINTITLE: (ELICITATION OR ELICITING) (DISTRIBUTION OR DISTRIBUTIONS)”. The searches were conducted on March 2021, produced in total $133 + 1000 + 182 = 1315$ results. Our inclusion criteria were (i) published work, (ii) available, and (iii) pertinent to methodological prior elicitation. In exceptional cases, we included work that did not meet all the criteria (i)-(iii). The criterion (ii) means work which can be found in the public web. The two most common reasons of failing the criterion (iii) were that the work was an application of an existing prior elicitation method, and the work did not address the elicitation at all (e.g. papers about prior specification). In the backward reference searching stage, we identified relevant research lines and references on prior elicitation by starting search from the references in the papers found in the first stage. The majority of relevant references were found in the backward reference searching stage. Indeed, the used search strings turned out to reach only a small fraction of the overall prior elicitation research.

3.1 Elicitation in parameter space

This section discusses prior elicitation in the space of parameters (D3). The goal is to form $p(\theta)$ ¹ by directly eliciting information on $p(\theta)$ from the expert. However, many of the methods can be equivalently applied to eliciting information on $p(y)$, if combined with a suitable algorithm to construct $p(\theta)$ based on the obtained information, discussed in Section 3.2.

The first three subsections are organized according to the prior assumption (D1): univariate prior, multivariate prior, and some prior-specific elicitation methods. An introduction of elicitation in univariate parameter spaces can be found in (Oakley, 2010), and in multivariate parameter spaces in (Daneshkhah and Oakley, 2010). The last subsection discusses a popular technique for evaluating the performance of a prior elicitation algorithm.

Within the subsections themselves, we further organize the content according to the types of assessment tasks (D6). The basic assessment task involves assessing quantiles in all three well-know elicitation protocols: Sheffield Elicitation Framework (SHELF) (Oakley and O’Hagan, 2019), Cooke protocol (Cooke, 1991), and Delphi protocol (as in EFSA (2014)). Specifically, SHELF incorporates many common assessment tasks such as (i) plausible limits, (ii) median, (iii) quartiles, (iv) tertiles, (v) roulette, (vi) quadrant probabilities, and (vii) conditional range. The tasks (ii)-(v) amount to asking different types of quantiles that are particularly relevant when the prior is univariate, and they are discussed next. Subsequently, the tasks (vi)-(vii) that are relevant for a multivariate prior, are discussed.

¹Recall the notation in Eq. (2.1).

Univariate prior

Let us consider assigning a probability distribution for a scalar θ . Directly asking the full density function $p(\theta)$ is challenging (O’Hagan, 1988; Goldstein and Wooff, 2007). For this reason, a major strand in the prior elicitation literature approaches this problem by asking certain descriptive elements of the distribution of θ from the expert. The descriptive elements, known as *summaries*, commonly include quantiles such as quartiles and medians, moments such as means, and limit values such as lower and upper bounds of θ . Prior elicitation reduces to eliciting a number of summaries of $p(\theta)$, and if well-chosen, the summaries should be enough to identify $p(\theta)$ (O’Hagan et al., 2006).

What summaries to elicit? There are plenty of psychological studies of people’s ability to estimate statistical quantities such as measures of central tendency. For an overview of relevant research in this area, see (Garthwaite et al., 2005, Section 2.2), (O’Hagan et al., 2006, Section 5.2), and (O’Hagan, 2019, Section 6). The literature broadly agrees on a few aspects. Mode and median are preferred over mean, since there is strong evidence that when the population distribution is highly skewed then people’s assessments of the mean are biased toward the median (e.g. Peterson and Miller, 1964). The higher-order moments (than expectation) should not be asked (Kadane and Wolfson, 1998). For example, Garthwaite et al. (2005) write about estimating the variance, “people are poor both at interpreting the meaning of ‘variance’ and at assigning numerical values to it” (Peterson and Beach, 1967).

There are two popular elicitation methods that differ by the type of elicited summary. The first is the *variable interval method* (or *V-method*, Spetzler and Holstein, 1975), which asks quantiles from the expert. The second is the *fixed interval method* (or *P-method*, Spetzler and Holstein, 1975), which asks probabilities. For an overview, see (Oakley, 2010). In decision sciences, the methods are referred to as probability encoding methods (Haggerty, 2011; Spetzler and Holstein, 1975; Hora, 2007). The methods were first introduced in the context of a temperature forecasting application (Murphy and Winkler, 1974). Through experimentation involving 103 judges, Abbas et al. (2008) showed “slight but consistent superiority” for the variable interval method along several dimensions such as accuracy and precision of the estimated fractiles. It is also possible to jointly ask both quantiles and probability, a *quantile-probability tuple* (Perepolkin et al., 2021), in which case the method is referred to as *PV-method* (Spetzler and Holstein, 1975).

Variable interval method as presented by Oakley (2010) and Raiffa (1968) consists of eliciting the median, lower and upper quartiles, and possibly minimum and maximum of θ . Oakley (2010) divided the elicitation into seven steps. In the beginning, the analyst elicits the (i) median, (ii) lower quartile, and (iii) upper quartile from the expert. It is possible to present assessment tasks in a form of gambles when the expert is not comfortable with directly providing quantiles. Then, the analyst asks (iv) the expert to reflect on her choices to check for consistency. Next, he (v) fits a parametric distribution by minimizing a sum of squared errors between the elicited quartiles and the quartiles of the cumulative distribution function of a parametric prior. The sixth step is ‘feedback stage’ that involves (vi) presenting the fitted distribution back to the expert and allowing

the expert to modify her original judgements. Similar quantile-based elicitation has been studied in (Garthwaite and O’Hagan, 2000; Dey and Liu, 2007).

Fixed interval method as presented by O’Hagan (1998) and Oakley (2010) consists of eliciting lower and upper bounds, mode, and five probabilities of θ . For instance, the probabilities $P(\theta_{min} < \theta < \theta_{mode})$ and $P(\theta_{min} < \theta < (\theta_{mode} + \theta_{min})/2)$ are elicited. Then, the analyst fits a parametric distribution by minimizing a sum of squared differences between elicited probabilities and the corresponding probabilities implied by the prior distribution. A default choice for prior distribution is a beta distribution with possibly scaled support $(\theta_{min}, \theta_{max})$. Prior work on this method can be found in (Phillips and Wisbey, 1993; O’Hagan, 1995).

A closely related method is the *histogram method*, which is popular in the expert knowledge elicitation applications (Smith and Mandac, 1995; Forbes et al., 1994; Coolen et al., 1992; Önkal and Muradoglu, 1996; Roosen and Hennessy, 2001; Lahiri et al., 1988; Pattillo, 1998; Grisley and Kellogg, 1983). The method is simple, consisting of a few steps: The analyst asks the minimum and the maximum of θ . Then, he splits the obtained range into a finite number of sub-intervals of equal length. Finally, he asks the expert to assign probabilities to each interval. The elicited histogram can be either used as a nonparametric prior (see Section 3.3) or a parametric prior distribution can be fitted by using some optimization procedure. An alternative, facilitated approach, is to ask the expert to construct the histogram by the aid of *gambling devices* such as dice, coins, and roulette wheels. Gore (1987) introduced the *roulette method* in which the expert allocates chips to the bins. The probability of a parameter lying in a particular bin is interpreted as the proportion of chips allocated to that bin. Johnson et al. (2010b) verified ‘feasibility, validity, and reliability’ of this method in an experiment with 12 subjects. The roulette method is supported in many elicitation tools, such as SHELF (Oakley and O’Hagan, 2019) and MATCH (Morris et al., 2014).

Multivariate prior

Let us consider a vector of parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_d)$ to which we want to assign a joint probability distribution. Given that there are available methods for eliciting univariate θ , a popular strategy is to elicit $\boldsymbol{\theta}$ by multiple independent univariate elicitation. For this, $\boldsymbol{\theta}$ can be transformed into a new vector $\boldsymbol{\phi}$ in which the coordinates are independent. Alternatively, the analyst may elicit the conditional distributions of $\theta_i | \boldsymbol{\theta}_{-i}$ for all $i \in \{1, \dots, d\}$ with hypothetical values for the $\boldsymbol{\theta}_{-i}$. In some special cases it is possible to elicit marginals by using univariate elicitation methods, and then complete the multivariate prior by additional elicited information such as covariances and correlations between parameters. A problem is that correlations are difficult to assess (Jennings et al., 1982; Wilson, 1994). For an overview of eliciting covariances and correlations, and in general the dependence, see (Werner et al., 2017) and (Kurowicka and Cooke, 2006).

In the context of multivariate elicitation, the process of transforming $\boldsymbol{\theta}$ into a new parameter vector with independent coordinates is called *elaboration* (Oakley and O’Hagan, 2019). The independence here refers to ‘subjective independence’ meaning that the additional information about one parameter will not affect the expert’s opinion

about the others. For instance, if θ_1 and θ_2 are treatment effects of two drugs, then the analyst may consider a transformation $\phi = (\theta_1, \theta_2/\theta_1)$. The relative effect θ_2/θ_1 is likely to be subjectively independent of θ_1 , since the expert can consider separately the magnitude of θ_1 , and how much more or less effective is θ_2 relative to θ_1 . An illustration of the transformation approach is presented in (O'Hagan, 1998).

Elicitation based on Gaussian copula allows the elicitation of a multivariate $p(\boldsymbol{\theta})$ to be decomposed into the elicitation of the marginals $p(\theta_i)$ and a so-called copula density. The method is based on the Sklar's theorem (1959) and the fact that the joint density $p(\boldsymbol{\theta})$ can be written as a product of the marginals $p(\theta_i)$ and the copula density, given that the marginal densities and the copula density are differentiable (Sklar, 1959). If the multivariate Gaussian copula is assumed with the density parameterized by a correlation matrix, then the elicitation reduces to elicitation of the marginals and the correlation matrix of the Gaussian copula. A detailed description of the method can be found in (Clemen and Reilly, 1999); see also (Dall'aglio et al., 1991) and (Kurowicka and Cooke, 2006). Clemen et al. (2000) compared six different methods for eliciting correlation in two experimental studies. The main result was that directly asking the expert to report a correlation is a reasonable approach. The elicitation of correlations in the context of multivariate normal distribution has been studied by Gokhale and Press (1982) and Dickey et al. (1985).

One option is to directly elicit the joint or conditional probabilities of $p(\boldsymbol{\theta})$. Fackler (1991) introduced a multivariate elicitation method based on quadrant probability or what he called the median deviation concordance (MDC) probability. MDC is defined as the probability that two variables will both fall either above or below their respective median. Abbas et al. (2010) proposed using isoprobability contours for multivariate elicitation. An isoprobability contour of a joint cumulative distribution function $F(\theta_1, \dots, \theta_d)$ is the collection of parameters $\boldsymbol{\theta}$ that have the same cumulative probability. They motivated the use of isoprobability contours by avoiding elicitation of the association or correlations between the parameters.

Beta and Dirichlet prior

The beta distribution typically describes the distribution of a proportion or a probability. It is the conjugate prior of the binomial sampling model in which case it describes the prior for the success probability. The elicitation of a beta prior has received considerable attention from the very beginning of the prior elicitation research. Winkler (1967) essentially compared four elicitation techniques for a beta prior in a binomial sampling model. An overview of the elicitation methods for the beta distribution can be found in (Hughes and Madden, 2002) with a summary in (Jenkinson, 2005, Section 2.3.1).

The methods differ with respect to which quantities are elicited and the procedure how the elicited values are converted into the parameters of a beta distribution. The methods of (Weiler, 1965; Fox, 1966; Gross, 1971; von Holstein, 1971; Duran and Booker, 1988; Pham-Gia et al., 1992; Enøe et al., 2000) start by asking a location (mode, median or mean) of the beta distribution from the expert. A typical follow-up assessment task involves some dispersion measure or confidence interval for the asked quantity. For

instance, [Gross \(1971\)](#) asked the mean μ and the expert’s subjective probability that this mean lies in the interval $(0, K\mu)$, where $0 < K < 1$ is given by the analyst. In contrast, [Pham-Gia et al. \(1992\)](#) elicited the median and the mean deviation about the median. It is also possible to first ask the confidence interval, and then match the center of the interval with the mean of the beta distribution, as demonstrated by [Joseph et al. \(1995\)](#); [Cuevas \(2012\)](#).

The methods of Equivalent Prior Samples (ESP) and Hypothetical Future Samples (HFS) ([Good and England, 1965](#); [Winkler, 1967](#); [Bunn, 1978](#)) assume a binomial sampling model, the methods are model-specific (D2). Both methods first require the expert to assess the mean of the beta prior. Then, HFS requires the expert to update her prior mean given an imaginary sample, and ESP requires the expert to report the corresponding sample size. Both methods require expert feedback (mean of parameter) in the space of parameters (D3), but EFS also asks a sample size, which is an observable quantity. So, EFS can be classified as a *hybrid prior elicitation method*. Similarly, the beta prior elicitation methods introduced by [Chaloner and Duncan \(1983\)](#); [Gavasakar \(1988\)](#) assume a binomial sampling model, being also model-specific. In contrast, the assessment tasks contain observables solely, such as asking a number of successes given a fixed number of trials. These types of methods are discussed in Section 3.2.

Dirichlet distribution is a multivariate generalization of the beta distribution. Since the marginal distributions are betas, many elicitation methods for the Dirichlet distribution are based on univariate elicitation of the beta marginals. However, the problem is that the set of elicited beta marginal distributions would almost certainly not be consistent with any Dirichlet distribution. For instance, in order for the elicited beta marginals to be consistent with a Dirichlet distribution, the two parameters of a beta marginal have to sum to the same value for all beta marginals, and the expectations of the beta marginals have to sum to one. To overcome these issues, [Zapata-Vázquez et al. \(2014\)](#) and [Elfadaly and Garthwaite \(2017\)](#) suggest reconciling a suitable Dirichlet by specifying its parameters in terms of the parameters of “corrected beta marginals”. [Elfadaly and Garthwaite \(2013\)](#) proposed to use conditional corrected beta marginals. [Van Dorp and Mazzuchi \(2004\)](#) studied how the parameters of beta marginals and Dirichlet extension can be recovered given a number of quantile constraints. Based on this work, [Srivastava et al. \(2019\)](#) considered elicitation of cumulative distribution function plots for the univariate marginal beta distributions for constructing a Dirichlet prior. [Evans et al. \(2017\)](#) let the expert state a lower or upper bound on each marginal probability that they are “virtually certain” of.

The Dirichlet distribution is a conjugate prior of the categorical distribution and the multinomial distribution. Hence, the elicitation is often presented in the context of a multinomial model. For an overview, see ([Wilson et al., 2021](#)) or ([Elfadaly, 2012](#)). Alternative approaches, which offer greater flexibility to the Dirichlet prior, are also introduced such as elicitation of generalized Dirichlet distribution ([Elfadaly and Garthwaite, 2013](#)), mixture of Dirichlets ([Regazzini and Sazonov, 1999](#)), multivariate copulas ([Elfadaly and Garthwaite, 2017](#)) and vines ([Wilson, 2018](#)). The methods of [Dickey et al. \(1983\)](#) and [Chaloner and Duncan \(1987\)](#) consider assessment tasks on the model observables (see Section 3.2).

Scoring rules

Scoring rules are a class of devices for eliciting and evaluating probabilities. The concept appears in the classical work by [Savage \(1971\)](#), [De Finetti \(1974\)](#), and [Murphy and Winkler \(1970\)](#). The key idea behind scoring rule based elicitation is that the assessment tasks are formulated so that they encourage the expert to provide careful assessment from which subjective 'true' probabilistic judgment can be recovered. Indeed, in early work by [Brier et al. \(1950\)](#), the idea of scoring was presented as a verification and 'bonus' system for forecasters to provide accurate predictions. [Matheson and Winkler \(1976\)](#) write: "in terms of elicitation, the role of scoring rules is to encourage the assessor to make careful assessments and to be 'honest', whereas in terms of evaluation, the role of scoring rules is to measure the 'goodness' of the probabilities". Hence, the scoring rule methods can be used either for elicitation or for evaluating the performance of an elicitation method. The focus in this subsection is on the former, whereas the latter is relevant for Section 4.4. See [Lindley \(1987\)](#) for a philosophical justification and [Winkler et al. \(1996\)](#); [Gneiting and Raftery \(2007\)](#) for a review of scoring rules.

The scoring rules are derived from the expected utility hypothesis ([von Neumann et al., 1944](#)). The hypothesis states that an agent (expert) chooses between risky prospects or 'lotteries' by selecting the one with the highest expected utility. For the sake of illustration, let the expert's subjective win probability be p , so her loss probability will be $1 - p$. The analyst asks the expert to report her assessment for win probability, which is denoted by q . If the expert is risk-neutral (explained later) and the analyst gives win reward $\log(q)$ and loss reward $\log(1 - q)$ to the expert, then the expert reports her true subjective win probability, that is $q = p$. This result can be extended to continuous lotteries, where the expert reports a probability density $q(x)$ that reflects her subjective probability density $p(x)$ over possible outcomes x ([Matheson and Winkler, 1973, 1976](#); [Holt Jr, 1979](#)). If we think of the parameter of interest θ as the outcome x , then the scoring rule methods can be applied to prior elicitation.

The aforementioned term risk-neutral means that the underlying scoring rule implies that the expert is assumed to only care about the expected value of the payoff, not about how risky it is. In practice, it is difficult to choose the scoring rule that reliably reflects the expert's risk preference, and in principle that should be elicited also ([Kadane and Winkler, 1988](#)). In our example, the analyst assumed a logarithmic scoring rule, but other scoring rules are also possible, such as quadratic and spherical ([von Holstein, 1970](#); [Murphy and Winkler, 1970](#); [Matheson and Winkler, 1976](#)). [Karni \(2009\)](#) and [Qu \(2012\)](#) proposed a scoring based elicitation mechanism that allows incentives to be set at any desired level.

A drawback of scoring rule based elicitation systems is that they often require the expert to report a full probability distribution at once, which is a challenging task. This problem is mitigated in the ELI method ([van Lenthe, 1993a,b](#)) that is a graphical computerized interactive elicitation technique, where scoring functions are displayed along with subjective probability density functions. The technique also graphically depict the consequences of a probability assessment, which resulted in a better calibration and a higher accuracy of the probability assessments in a study with 192 subjects ([van Lenthe, 1994](#)). The scores that allow elicitation by asking quantiles instead of a full

density function have been studied as well, see (Cervera and Munoz, 1996; Gneiting and Raftery, 2007; Schervish et al., 2012; Ehm et al., 2016). Osband and Reichelstein (1985), Lambert et al. (2008), Gneiting (2011), and Steinwart et al. (2014) studied, in general, when a property of a probability distribution is *elicitable* in scoring rule based elicitation. A property is elicitable if there exists a scoring function such that minimization of the associated risks recovers the property, or equivalently there exists a non-negative scoring function whose ‘predictions that are further away from the property have a larger risk’ (Theorem 5, Steinwart et al., 2014).

3.2 Elicitation in observable space

Prior elicitation methods that let the expert provide her assessment in the space of model observables date back to the 1980s (Kadane et al., 1980; Kadane, 1980; Winkler, 1980). The idea of deriving a prior for the parameters from a prior on the observables can be seen to be as old as Bayesian statistics itself (Stigler, 1982). In the prior elicitation literature, the idea of conducting the elicitation in the space of observables instead of parameters originates from the question about what is the essence of the model parameters, which is a controversial topic (De Finetti, 1937; Geisser, 1971; Briggs, 2016; Billheimer, 2019). A natural requirement for a successful elicitation is that the expert is able to provide meaningful information on the asked quantities. When the asked quantity is a model parameter, the expert must know how to interpret it and have some idea of its natural scale and magnitude. In simple cases (e.g. the model is a Bernoulli trial) direct elicitation of the parameters can be justified on the grounds of the *operational interpretation of the model parameters* in which the parameters are interpreted in terms of a (simple) limiting average of observables (Bernardo and Smith, 1994). Kadane et al. (1980) took a position that:

Without touching the philosophical issue of the sense in which parameters may be said to “exist”, we acknowledge that even experienced observers have trouble answering questions on their beliefs about them.

Later, Kadane and Wolfson (1998) concluded that the choice of the elicitation space depends on the nature of the problem and whether the parameters have intrinsic meaning to the expert. For instance, applications in economics may benefit from elicitation methods that operate in the parameter space, since economists are “quite used to thinking parametrically”. However, in general, Kadane and Wolfson (1998) state that there has been some agreement in the statistical literature (e.g. Winkler, 1986; Chaloner, 1996) on that the “experts should be asked to assess only observable quantities”. The argument is that the expert does not need to understand what a parameter is. Kadane and Wolfson (1998) also noted that the expert does not usually consider the correlation between regression coefficients, so elicitation in the observable space is particularly suited since it allows an indirect assessment of the correlations. Wolfson (2001) argued that even though the methods in the observable space can be more difficult to implement, they can be more appropriate in terms of allowing the expert to communicate both her knowledge and uncertainty within the constraints of the model. Akbarov (2009) argued that elicitation in the observable space can be made model-agnostic (D2), since it need

not rely on any assumptions about the family of distributions for the prior and sampling distributions, it can be designed to work with any prior predictive distribution. However, we found that there are relatively few model-agnostic prior elicitation methods that assume expert’s feedback in the space of observables (Hartmann et al. (2020) compared to Table 1).

A challenge in doing elicitation in the observable space is the difficulty to separate two sources of randomness: due to θ and due to y (Garthwaite et al., 2005; Perepolkin et al., 2021). There are some solutions to that, for instance, Kadane and Winkler (1988) proposed to the expert to consider only the mean \bar{y} , and then conduct other assessment tasks to elicit the randomness in y . In linear regression, the latter amounts to eliciting the measurement error. Stefan et al. (2020) pointed out that obtaining unique prior distributions from the elicited data patterns on model observables becomes difficult for complex models, as for models with highly correlated parameters, the different prior parameter combinations can lead to similar model predictions on observables (Gutenkunst et al., 2007). For other arguments against and in favor of both elicitation schemes, see the discussed references and (Gaoini et al., 2009; Choy et al., 2009; James et al., 2010; Al-Labadi et al., 2018).

The main research line on elicitation in the observable space falls under the title of *predictive elicitation* (see discussion in D3). It assumes a setting where the expert is asked about the median, quantiles, mode, or expectation of the response variable (denoted by y) at various design points (denoted by \mathbf{x} , design matrix by \mathbf{X}), and the underlying model comes from the family of the generalized linear models (Kadane et al., 1980; Oman, 1985; Garthwaite and Dickey, 1988; Ibrahim and Laud, 1994; Bedrick et al., 1996; Chen and Ibrahim, 2003; Denham and Mengersen, 2007; Elfadaly and Garthwaite, 2011; Garthwaite et al., 2013; Elfadaly and Garthwaite, 2015). The goal of the elicitation is to construct priors for the model hyperparameters. For a detailed review, see Garthwaite et al. (2005); here we only discuss hyperparameters concerning regression coefficients. Given the model error variance σ^2 , the regression coefficients are normally distributed with an unknown mean vector \mathbf{b} and an unknown covariance matrix $\sigma^2\mathbf{R}$. The elicitation of \mathbf{b} can be as easy as setting $\mathbf{b} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}_{.50}$ where $\mathbf{y}_{.50}$ is a vector of elicited medians (Kadane et al., 1980). However, the elicitation of \mathbf{R} is difficult, since the matrix must be positive-definite and there can be many elements in the matrix. Kadane et al. (1980), Garthwaite and Dickey (1988), and (Garthwaite et al., 2013) proposed an elicitation procedure where a structured set of sequential elicitation queries ensures that \mathbf{R} is a positive-definite matrix. Oman (1985) and Ibrahim and Laud (1994) proposed setting $\mathbf{R} = c(\mathbf{X}'\mathbf{X})^{-1}$ where c is a constant that is either estimated by using empirical Bayes or provided by the expert. Oman (1985) argued that this approach is similar to a two-stage prior specification approach presented by Schlaifer and Raiffa (1961) and Tiao and Zellner (1964), in that the first stage is a “mental experiment” where the design matrix \mathbf{X} (which reflects the relations among explanatory variables) is specified by the expert. However, neither of the methods (Oman, 1985; Ibrahim and Laud, 1994) uses the expert knowledge on the relationship between y and \mathbf{x} to infer \mathbf{R} , and thus the methods are not recommended for the elicitation of \mathbf{R} .

Complementing the main research line on elicitation in the observable space, Garthwaite and Dickey (1992), Laud and Ibrahim (1995), and Chen et al. (1999) studied feature

selection in regression models (see also Ibrahim and Chen, 2000; Ibrahim et al., 2015, in a context of power priors). The elicitation method of Garthwaite and Dickey (1992) is hybrid in that it combines elicitation data from both the parameter and observable spaces (see also Good and England, 1965; Denham and Mengersen, 2007; Casement and Kahle, 2018). Another elicitation problem that has got some attention is the prior elicitation of a multivariate normal sampling model. Al-Awadhi et al. (1997; 1998) assumed the natural conjugate prior (normal inverse-Wishart) and Garthwaite and Al-Awadhi (2001) assumed a non-conjugate prior (normal generalized inverse-Wishart). Natural conjugate prior forces a dependence between the mean and the covariance, so Garthwaite and Al-Awadhi (2001) proposed assessment tasks that allow the expert to quantify separately assessments about each of these parameters. Assessment tasks include conditional and unconditional quantiles where the conditions were specified by hypothetical data. Al-Awadhi and Garthwaite (2001) compared the approaches of Al-Awadhi and Garthwaite (1998) to Garthwaite and Al-Awadhi (2001), and concluded that the independence between mean vector and covariance matrix (which holds in Garthwaite and Al-Awadhi, 2001) leads to better performance in terms of obtained scoring rule values. Furthermore, there are prior elicitation methods for special applications, where the expert’s assessment task contains assessing observables, such as in survival analysis (Coolen, 1992; Wang and Zhou, 2009), risk control (Hosack et al., 2017), item response theory (Tsutakawa, 1984), and flash flood prediction (Gaoini et al., 2009).

Many authors have pointed out that the *prior predictive distribution*,

$$p(y) = \int p(y|\theta)p(\theta)d\theta, \quad (3.1)$$

gives rise to an integral equation when the distributions concerning the observables y are known (Aitchison and Dunsmore, 1975; Winkler, 1980; Wolpert et al., 2003; Gribok et al., 2004; Akbarov, 2009; Jarociński and Marcet, 2019). Suppose that $p(y)$ is elicited from the expert, the likelihood $p(y|\theta)$ is specified by the analyst, and the analyst is looking for an unknown prior $p(\theta)$. Then, the equation (3.1) is known as a Fredholm integral equation of the first kind, which is a well-known example of linear ill-posed problems. Additional regularity assumptions are needed in order to solve the prior from this equation. For instance, Gribok et al. (2004) proposed Tikhonov regularization that imposes smoothness constrains on the prior density, which is a natural restriction because probability density functions are typically smooth and differentiable. Alternatively, the analyst may assume a parametric prior $p(\theta|\lambda)$ where λ are its hyperparameters, in which case the prior predictive distribution reads as $p(y|\lambda) = \int p(y|\theta)p(\theta|\lambda)d\theta$, and the problem reduces to finding the optimal hyperparameters. For instance, Percy (2002; 2003; 2004) considered beta, gamma and normal priors, and demonstrated how the problem reduces to solving a system of nonlinear equations.

The prior predictive distribution provides a link between the model observables and the parameters, and the aforementioned integral equation is only one way to use that link in prior elicitation. In a recent work, Hartmann et al. (2020) assumed that the expert’s assessment task consists of assessing (prior predictive) probabilities of observables falling in certain regions of the observables space. The elicited probabilities were treated as noisy data with a suitable likelihood (see supra-Bayesian approach in D4). The prior

predictive distribution can also be used in prior elicitation via computer simulations. Monte Carlo integration of the integral (3.1) involves sampling first $\theta' \sim p(\theta)$ and then $y' \sim p(y|\theta')$. [Wesner and Pomeranz \(2021\)](#) considered comparing y' to some reference value y^* , and then back-tracking θ' that produced y' close to y^* . [Thomas et al. \(2020\)](#) considered two assessment tasks: verisimilitude judgements (is y' a credible draw from reality?) and pairwise judgements (for given y' and y'' , which one is more realistic?). The authors trained a Gaussian process classifier on this binary label data to recover the expert’s prior distribution on realistic parameters. [Chandramouli and Shiffrin \(2016\)](#) and [Chandramouli \(2020\)](#) proposed a framework where “data priors” are specified over the space of predictive distributions. Priors on model parametrizations are informed by how well their predictions approximate the different predictives.

3.3 Nonparametric priors

How should the functional form of prior distribution be chosen? Given that the goal of prior elicitation is to faithfully represent the belief of the person being elicited ([Garthwaite et al., 2005](#); [O’Hagan et al., 2006](#)), the family of parametric prior distribution or the form of nonparametric prior should be chosen to reflect this belief. However, also other objectives and practical aspects influence the choice of prior family, for example a specific parametric prior may be easier to use with probabilistic inference engines. For instance, the natural conjugate family is often chosen for computational easiness ([Garthwaite et al., 2005](#); [Gelman et al., 2013](#)), although also ‘piecewise conjugate priors’ ([Meeden, 1992](#)) and ‘mixture of natural conjugate priors’ ([Dalal and Hall, 1983](#)) have been proposed. Even if we agree that the choice of prior should be based on available information, and not for instance computational convenience, there is still some ambiguity in specifying the prior family. *Conservative methods* have been designed to avoid arbitrary choices among all possible prior distributions whose features can fit with the elicited information from the expert ([Robert, 2007](#); [Press, 2002](#); [Bousquet, 2008](#)). For instance, the prior choice can be based on the maximum entropy principle ([Jaynes, 1957](#)), leading to *maximal entropy priors*. Such priors are chosen to be the distribution with the largest entropy within the set of priors satisfying given constraints ([Zellner, 1977, 1991, 1996](#); [Bagchi, 1986](#); [Robert, 2007](#)). In our elicitation context, ‘the constraints’ correspond to the elicited information from the expert. Abbas has studied both univariate ([Abbas, 2003](#)) and multivariate ([Abbas, 2006](#)) cases. On the other hand, *non-conservative methods* incorporate additional information into the prior by selecting a functional form for the prior that reflects assumptions about the expert’s knowledge distribution. For instance, [Oakley and O’Hagan \(2007\)](#) introduced a nonparametric supra-Bayesian elicitation framework where the expert’s knowledge (a Gaussian process model) is assumed to have a specific prior mean function that differs from the zero-mean function.

Elicitation of a parametric prior has a limitation: it forces the expert’s knowledge distribution to fit the parametric family ([Oakley and O’Hagan, 2007](#)). Even though it is not reasonable to assume that there exists a ‘true’ prior distribution waiting to be elicited but rather a ‘satisfying’ prior distribution ([Winkler, 1967](#)), the fitted parametric prior may fall too far away from a distribution that reflects the expert’s knowledge distribution ‘satisfactorily’ well. A nonparametric prior, such as a histogram or a Gaussian process,

typically makes fewer assumptions than parametric priors, and thus provides more flexibility for finding a ‘satisfying’ prior.

Some early prior elicitation methods can be regarded as nonparametric, such as (Schlaifer, 1969), where the expert is asked to directly draw the density function. Similarly, the histogram method can be used for eliciting a nonparametric prior distribution (Berger, 1985). However, more complex nonparametric prior models can be found in the literature. Oakley and O’Hagan (2007) proposed modelling the expert’s density, $p(\theta)$, as a Gaussian process model (GP) that is completely characterized by its mean and covariance function (Williams and Rasmussen, 2006). The choice of the mean and covariance function provides a way to incorporate assumptions about the expert’s knowledge distribution. For instance, the authors proposed a covariance function that reflects the assumption that for those parameter values where the expert’s density is small, the prior GP variance should be small there too. The assessment tasks involve the expectation and percentiles of the distribution of θ . These elicited summaries are then used for inferring the posterior distribution for the GP model.

Gosling et al. (2007) developed the method of Oakley and O’Hagan (2007) further. The authors raised two deficiencies in the method. First, they pointed out that the analyst’s uncertainty about the tails of the expert’s knowledge density can be underestimated. Second, there are situations where the method assumes that the analyst knows the expert’s true density with certainty. The authors tackled these challenges by assuming the prior mean of the expert’s knowledge distribution to follow Student’s t -distribution. Oakley et al. (2010) and Daneshkhah et al. (2006) coupled the method of Oakley and O’Hagan (2007) with the roulette method (see Section 3.1). This reduces the cognitive burden of the expert, however, at the expense of the method becoming more computationally intensive. Daneshkhah et al. (2006) also investigated the issue of imprecision in the expert probability assessments in Oakley and O’Hagan (2007). The authors modelled the imprecision in the expert’s probability assessments by assuming that the elicited probabilities are contaminated by additive noise (Gaussian or uniform). They concluded that this “has sensible influence on the posterior uncertainty and so the expert density”. An example of a nonparametric elicitation method that does not use GPs is provided by Lins and Souza (2001). The authors represented the prior as a convex combination of distributions fitting the expert’s data.

Moala and O’Hagan (2010) proposed a multivariate extension of Oakley and O’Hagan (2007). In multivariate setting, the probabilities of the form $P(\boldsymbol{\theta} \in A)$, where A is a region of the parameter space, should be considered instead of percentiles. The authors illustrated that good results can be obtained by eliciting marginal probabilities supplemented by a small number of joint probabilities. However, the authors considered only bivariate elicitation, and they did not demonstrate that the method can deal with a higher number of parameters than two. The expert’s nonparametric (GP) model was capable of adapting to a bimodal ground-truth distribution (in a numerical experiment), showing that this nonparametric method is flexible enough to capture a target that is not a simple unimodal distribution. Hahn (2006) introduced a nonparametric elicitation method that can be applied to multivariate elicitation. The method essentially constructs a histogram from the expert’s data that consists of the elicited relative odds of parameters.

3.4 Model-specific methods

We will not address model-specific elicitation methods in detail, since we consider it more relevant to focus on model-agnostic methods such as (Gelfand et al., 1995; Gosling, 2005; Oakley and O’Hagan, 2007, 2019; Hartmann et al., 2020) discussed in the other parts of the article. However, some model-specific methods have been designed for model families that still offer broad flexibility in terms of practical applications, and hence can be valuable in practical modelling tasks. An excellent example is the research line on generalized linear models (GLMs) started by Kadane et al. (1980) and (Bedrick et al., 1996) (more about that in Section 3.2). While GLMs are clearly a limited model family, they are widely used in several fields and good elicitation tools for them can be useful. We refer the reader to Table 1 that categorizes the model-specific methods according to the model, the prior and the elicitation space, as a starting point for investigating the literature on model-specific elicitation methods for different cases.

The abbreviations in Table 1. NLR = normal linear regression, GLM = generalized linear model, LR = logistic regression, MVN = multivariate normal, Dir = Dirichlet, Bin = Binomial, Mult = Multinomial, Gam = Gamma, Poiss = Poisson, BQR = Bayesian quantile regression, AR = autoregressive model, PHR = proportional hazards regression, BN = Bayesian network, CIHM = conditionally independent hierarchical model, H = hierarchical, RE = random effect, PW = piecewise, (N)CP = (natural) conjugate prior, GIW = generalized inverse Wishart, PC = penalized complexity prior (Simpson et al., 2017), CMP = conditional mean prior (Bedrick et al., 1996), O = observable space, P = parameter space, H = hybrid space

3.5 Active elicitation

Human effort is costly, so an elicitation algorithm should assume asking the expert on a budget. Active elicitation helps to make most out of this limited budget, and it is relevant when overfitting (see D4) is considered. Active learning (Cohn et al., 1994; Settles, 2012; Gal et al., 2017) and Bayesian experimental design (Chaloner and Verdinelli, 1995; Ryan et al., 2016) study the selection of informative data-points (or ‘experiments’) to be used in regression or classification tasks. When active learning is applied to prior elicitation, we call it active prior elicitation or just *active elicitation*. While active learning typically chooses new data to maximize information gain about the parameters, active elicitation can use the same principles, but instead querying new data from the expert.

Various strategies have been proposed for active prior knowledge elicitation of the weights of a sparse linear regression model. For instance, Dae et al. (2017) used a strategy that maximizes the expected information gain in the prediction of an outcome (see Seeger, 2008 for a non-elicitation context). The higher the gain, the greater impact the new expert’s input has on the predictive distribution. Sundin et al. (2018) studied this gain in a personalized medicine scenario where a single prediction was considered instead of predicting over a whole training dataset. Micallef et al. (2017) modelled the expert’s binary relevance input on the regression coefficients as a linear bandit algorithm. The authors proposed to select the next query by maximizing a high probability bound for the uncertainty of feature relevance.

Author	Model	Prior	D3
Kadane et al. (1980)	NLR	NCP	O
Oman (1985)	NLR	NCP	O
Garthwaite and Dickey (1988)	NLR	NCP	O
Ibrahim and Laud (1994)	NLR	NCP	O
Bedrick et al. (1996)	GLM	CMP	O
Chen and Ibrahim (2003)	GLM	CP	O
Denham and Mengersen (2007)	GLM	NCP	H
Elfadaly and Garthwaite (2011)	NLR	NCP	O
Garthwaite et al. (2013)	PW-GLM	NCP	O
Elfadaly and Garthwaite (2015)	Gam-GLM	NCP/log-normal	O
Garthwaite and Dickey (1992)	NLR	mixture-NCP	H
Laud and Ibrahim (1995)	NLR	NCP	O
Chen et al. (1999)	LR	custom	O
Leamer (1992)	NLR	NCP	P
Hosack et al. (2017)	GLM	NCP	O
Carlin et al. (1992)	RE-LR	custom	P
Al-Hamzawi et al. (2011)	RE-BQR	power prior	P
Garthwaite and Al-Awadhi (2006)	PW-LR	NCP	O
Kadane et al. (1996)	AR	PW-CP	O
Garthwaite and Dickey (1991)	NLR	NCP	O
Chaloner et al. (1993)	PHR	adjusted-NCP	O
Ibrahim et al. (1999)	PHR	semi-parametric	O
Soare et al. (2016)	NLR	delta	P
Micallef et al. (2017)	NLR	half-normal	P
Afrabandpey et al. (2017)	NLR	NCP	O
Hem et al. (2021)	RE-H	PC/Dir	P
Rizzo and Blackburn (2019)	BN	nodepriors	P
Al-Awadhi and Garthwaite (2006)	PW-LR	NCP+other	O
Kynn (2005)	PW-LR	CMP	O
O’Leary et al. (2008; 2009)	LR	mixture-normal	P
Al-Awadhi and Garthwaite (1998)	MVN	NCP	O
Garthwaite and Al-Awadhi (2001)	MVN	GIW	O
Steffey (1992)	CIHM	custom	P
Elfadaly and Garthwaite (2020)	Mult-LR	logit-normal	O
Bunn (1978)	Mult	Dir	P
Chaloner and Duncan (1987)	Mult	Dir	O
Dickey et al. (1983)	Mult	Dir	H
Regazzini and Sazonov (1999)	Mult	mixture-Dir	O
Wilson (2018)	Mult	D-vine	P
Elfadaly and Garthwaite (2013)	Mult	Dir/Connor-Mosimann	P
Elfadaly and Garthwaite (2017)	Mult	Dir/Gaussian copula	P
Srivastava et al. (2019)	Mult	Dir	P
Moreno and Giròn (1998)	Poiss-Bin	NCP	P
Good and England (1965)	Bin	Beta	P/H
Chaloner and Duncan (1983)	Bin	Beta	O
Gavasakar (1988)	Bin	Beta	O

Table 1: Summary of model-specific elicitation methods.

Active elicitation assumes iterative elicitation (see D6). This enables the theoretical study of the algorithmic convergence in elicitation tasks. However, surprisingly, we are not aware of any results of convergence guarantees in prior elicitation tasks. For instance, in clustering tasks, Balcan and Blum (2008) and Awasthi et al. (2014) provided theoretical guarantees for the convergence of a clustering algorithm where the expert can be queried about split and merge requests. Similarly, Ashtiani et al. (2016) provided convergence guarantees when the expert can be queried about instances belonging to same clusters. Kane et al. (2017) extended the theoretical study to comparison queries.

In active elicitation, modelling of the expert’s behavior becomes important. Dae et al. (2018) showed a significant performance improvement of using a user model, to model the expert’s behavior, in a feature relevance elicitation task for sparse linear regression. The expert is not treated as a passive data source but instead her behavior is modelled to anticipate her input and take into account her biases, see also Micallef et al. (2017); Afrabandpey et al. (2019). Active elicitation could be used to minimize experts’ effort in any type of prior elicitation, in principle. In many works cited here, the expert was shown data during the elicitation, which is an unconventional choice in prior elicitation. The motivation was to further reduce the expert’s effort. Furthermore in some works, the elicitation data was used as separate auxiliary data.

3.6 Multiple experts

Eliciting the knowledge of a group of experts is an important research line. However, in this article, we do not go into this topic in any great depth, but instead we will confine ourselves to providing the key references on the subject. A brief summary of the topic can be found in O’Hagan (2019, Sec. 3), Garthwaite et al. (2005, Sec. 5), and Jenkinson (2005, Sec. 7). We refer the reader to the following papers (Winkler, 1968; Albert et al., 2012; Hanea et al., 2018; Phillips, 1999; Phillips and Wisbey, 1993; French, 2011; Ouchi, 2004; Daneshkhah et al., 2011; Genest and Schervish, 1985; Morris, 1977), and to a recent comparison study (Williams et al., 2021).

Around the elicitation of a group of experts, there have been research on protocols how to coordinate the elicitation process. Four prominent elicitation protocols have been designed for elicitation with multiple experts: Sheffield protocol (Gosling, 2018; Oakley and O’Hagan, 2019), Cooke protocol (Cooke, 1991), Delphi protocol (Brown, 1968), and IDEA protocol (Hanea et al., 2018). An overview of the protocols is provided by EFSA (2014), Dias et al. (2018), and O’Hagan (2019).

Finding a single distribution that combines elicited knowledge from all the experts is known as the problem of aggregation. There are two main approaches. In the *mathematical aggregation*, the experts are elicited separately, and a probability distribution is fitted to each expert’s knowledge. These are then combined into an aggregate distribution using a mathematical formula, a so-called *pooling rule*. A popular pooling rule is *supra-Bayesian pooling* (Morris, 1974, 1977; Keeney and Raiffa, 1976; Lindley, 1985; West, 1988; Genest and Zidek, 1986; Roback and Givens, 2001; Albert et al., 2012). The *behavioral aggregation* encourages the group of experts to discuss their knowledge, and to settle upon group consensus judgments, to which an aggregate ‘consensus’ distribution is fitted.

3.7 Heuristics and biases

Every expert is inclined to commit to cognitive biases and to use cognitive shortcuts (heuristics) when making probabilistic judgments. Tversky’s and Kahneman’s heuristics and biases research program in the 1970s led to a series of published papers on describing difficulties in assessing probabilities. The work was summarized in the seminal paper (Tversky and Kahneman, 1974). O’Hagan (2019) highlighted four relevant heuristics and biases for expert knowledge elicitation: anchoring, availability, range-frequency, and overconfidence, of which the first two already appear in the Tversky’s and Kahneman’s original article (1974). We encourage the reader to take a look at an interesting summary of these provide by O’Hagan (2019), although there has also been some controversy on the topic (Gigerenzer, 1996; Kynn, 2008), such as an allegation that the statistical literature ignores the effect of framing in elicitation.

Every prior elicitation method should address systematic cognitive biases identified in the literature. Indeed, the leading principle in designing the elicitation protocols mentioned in Section 3.6 has been to minimize cognitive biases to which expert probabilistic judgments may be subject (O’Hagan, 2019).

We provide some entry-points to this topic. We refer the reader to O’Hagan (2019, Sec. 2), Garthwaite et al. (2005, Sec. 2), Kynn (2008), Gigerenzer and Todd (1999), and Burgman (2016). See also the discussion on how to deal with imprecision in probability judgements (O’Hagan and Oakley, 2004).

4 Where Should We Be Going?

In this section, we discuss understudied directions in the prior elicitation literature and propose five promising avenues in these directions (Sections 4.1-4.5). They aim to help in solving the technical, practical, and societal challenges described in Section 2.2, and hence we believe research on these avenues will increase the adoption of prior elicitation techniques.

In our literature review (Section 3) we observed that there are regions in the prior elicitation hypercube (Section 2.3) that are well understood. Elicitation in the space of parameters is well covered, and there are widely accepted protocols on how to conduct the elicitation with single (Section 3.1) and multiple experts (Section 3.6). The research on elicitation in the space of observables is also abundant (Section 3.2), but with a severe limitation. Namely, almost all the research is model-specific. The literature is aware of many cognitive biases present in elicitation (Section 3.7), although not all methods take them properly into account. There are many model-specific elicitation methods, and particularly elicitation for (generalized) linear models has received a lot of attention (Section 3.4).

However, there are understudied corners in the hypercube that should be explored more.

Technical solutions: We believe that an elicitation method should support elicitation both in the parameter and observable space, should be model-agnostic, and should be

sample-efficient since human effort is costly. In Section 4.1, we propose an approach for prior elicitation that takes these objectives into account. We also believe that elicitation is easier when the prior is globally joint. These globally joint priors are discussed in Section 4.3, but essentially, they let elicitation be reduced to just a few interpretable hyperparameters.

Practical solutions: To help make model building easier, faster and better in reflecting expert knowledge, we need to integrate prior elicitation into the Bayesian workflow (Section 4.2). And this requires software able to inter-operate with already existing tools for Bayesian modelling, including probabilistic programming languages (Section 4.5). The software needs to support model-agnostic elicitation, otherwise there will be problems with integration into the Bayesian workflow, because a change in the model specification could preclude prior elicitation.

Societal solutions: We emphasize the need for considerably extended user evaluation, required for verifying that the methods have practical value (Section 4.4), and the need of case studies showing the advantages that a careful prior elicitation process can bring to the modelling process.

4.1 Bayesian treatment of the expert in prior elicitation

In this section, we propose a unified approach to prior elicitation that brings together several model-agnostic elicitation methods. The approach allows the expert to provide her response in both the parameter and observable space (D3), and supports sample-efficient elicitation (D6) by treating the expert in a Bayesian fashion.

In the supra-Bayesian approach, elicitation of an expert’s knowledge is treated as any other Bayesian inference problem where the analyst’s posterior belief about the expert’s knowledge is updated in the light of received expert data (see the discussion in D4). We propose viewing the prior elicitation event itself as an interplay of the expert and the analyst with the following characteristics:

Analyst Poses queries to the expert and gathers the expert’s input into a dataset \mathcal{D} . The analyst’s goal is to infer the distribution of the parameters θ , conditional on the expert’s input data, $p(\theta|\mathcal{D})$.

Expert Based on her domain expertise, the expert answers to the analyst’s queries. The expert’s input is modelled through the *user model* $p(z|q)$ that is the conditional probability of the expert’s input z given the analyst’s query q . That is, \mathcal{D} consists of N samples $(z_i, q_i)_{i=1}^N$, and all the q_i are treated as fixed.

Expert data can be provided in multiple elicitation spaces, all of which can be combined to derive a single prior within the user model. For example, we can elicit expert data in both the observable space (data \mathcal{D}_Y) and in the parameter space (data \mathcal{D}_Θ). The analyst’s goal is then to infer the distribution of the parameters conditional on the expert’s input data, that is $p(\theta|\mathcal{D}_Y, \mathcal{D}_\Theta)$. We assume that the analyst updates his knowledge according

to Bayes’ rule. Hence, he treats elicitation as a posterior inference problem,

$$p(\theta|\mathcal{D}_Y, \mathcal{D}_\Theta) = \frac{p(\mathcal{D}_Y|\theta)p(\mathcal{D}_\Theta|\theta)p(\theta)}{p(\mathcal{D}_Y, \mathcal{D}_\Theta)}, \quad (4.1)$$

given the likelihoods $p(\mathcal{D}_Y|\theta)$ and $p(\mathcal{D}_\Theta|\theta)$, and the analyst’s prior belief on the expert’s knowledge $p(\theta)$. In (4.1), we have assumed \mathcal{D}_Y and \mathcal{D}_Θ to be conditionally independent given θ . The likelihoods $p(\mathcal{D}_Y|\theta)$ and $p(\mathcal{D}_\Theta|\theta)$ account for the uncertainty inherent to the elicitation process due to the mechanism how the expert quantifies her knowledge on θ . Hence, the conditional independence assumption essentially states that: given that there exists a fixed parameter vector θ that the expert thinks to be ‘true’, the mechanism how the expert reveals her knowledge on θ is independent between the two elicitation spaces.

Besides the prior $p(\theta)$, the framework only requires specifying $p(z|q, \theta)$ which describes, at individual query q level, how the expert responds if she thinks that θ is true. This $p(z|q, \theta)$ is also the *likelihood for a single data-point* (z, q) , since q is treated as fixed without a probability distribution assigned to it. The user model can be obtained by marginalization, $p(z|q) = \int p(z|q, \theta)p(\theta)d\theta$.

The proposed approach can be readily extended to support both sample-efficient elicitation (via active elicitation) and AI-assisted elicitation.

Active elicitation. How to make the most out of the limited budget of N expert’s inputs? In other words, what is an optimal strategy to select a sequence of queries $(q_i)_{i=1}^N$? This is where the user model comes to play. When the analyst poses a query q , he anticipates that the expert’s input z is distributed according to $p(z|q)$. The analyst applies the user model to choose the most informative queries. For instance, if the analyst wants to maximize the expected information gain of $p(\theta|\mathcal{D})$ with respect to a new query q , then the user model is needed for anticipating the corresponding yet unseen response z , which involves taking expectation over $p(z|q)$.

AI-assisted elicitation. One important thing to note is that the analyst (or here facilitator) need not manually select the next queries, but the whole elicitation process can be supervised by an ‘artificial facilitator’ – an *AI-assistant*. For instance, the AI-assistant can be as simple as consisting only of a user model combined with an active learning criterion for selecting next queries. However, in principle, it is possible to extend the functionalities and capabilities of the AI-assistant to take into account, for instance, the expert’s biases and incapacibilities of providing informative input for some queries.

Through the following examples, we illustrate how the proposed approach brings together prior elicitation methods found in the literature.

Examples.

- *Quantiles with mixture beta assumption* (Gelfand et al., 1995). \mathcal{D} is a set of quantiles of the prior distribution of parameters. The elicitation space is the parameter space, $\mathcal{D} = \mathcal{D}_\Theta$. The likelihood $p(\mathcal{D}_\Theta|\theta)$ equals to Eq. (4) in (Gelfand et al., 1995), and it is derived from a few assumptions, one being that the expert’s input is a transformation of mixture of beta-distributed random variables. The authors proposed using Markov chain Monte Carlo for sampling from the posterior $p(\theta|\mathcal{D}_\Theta)$.

- *Judgements about plausible outcomes* (Hartmann et al., 2020). \mathcal{D} is a set of prior predictive probabilities where the expert provides $P(A_i|\lambda)$ for all $i = 1, \dots, n$, given a partition $\mathbf{A} = \{A_1, \dots, A_n\}$ of the observable space and hyperparameter vector λ of a parametric prior $\pi(\theta|\lambda)$. The elicitation space is the observable space, $\mathcal{D} = \mathcal{D}_Y$. Hartmann et al. (2020) assumed a Dirichlet likelihood and used maximum likelihood estimation to estimate λ .
- *Quantile-probability tuples with indirect Bayesian inference* (Perepolkin et al., 2021). In their simple version (Perepolkin et al., 2021, Sec. 7.3), \mathcal{D} is a set of quantile-probability tuples (see Section 3.1) of the observables. The elicitation space is the observable space, $\mathcal{D} = \mathcal{D}_Y$. The parameters of interest θ can be, for instance, probabilities of fish weights in different categories (small, medium, large, huge). The likelihood $p(\mathcal{D}_Y|\theta)$ is a meta-logistic distribution (Keelin, 2016). The prior $p(\theta)$ is a “flat” QDirichlet prior, but essentially one can think of it as a “uniform” Dirichlet, that is *Dirichlet*(1, 1, 1, 1). For the posterior inference (Perepolkin et al., 2021, Sec. 6), the authors used a Hamiltonian Monte Carlo algorithm in Stan, interfaced via the `rstan` package (Stan Development Team, 2020) in R.

4.2 Bayesian modelling workflow

Having to choose a prior distribution can be portrayed both as a burden and a blessing. We choose to affirm that it is a necessity. If you are not choosing your priors yourself, then someone else is inevitably doing it for you and the automatic assignment of flat priors is not a good idea (Carlin, 2000; van Dongen, 2006; Gelman, 2006; Gelman et al., 2013, 2017; Smid and Winter, 2020). Under some scenarios, we can rely on default priors and default models. For instance, we may simply need to use a given model for routine inference over new datasets. However, having the flexibility to alter model assumptions could be advantageous, and priors are just one form of assumptions. Thus, adopting a Bayesian workflow for prior elicitation should help to reduce the burden and increase the blessing.

We need a Bayesian workflow, rather than mere Bayesian inference, for several reasons (Gelman et al., 2020): Bayesian computation can be challenging and generally requires exploration and iteration over alternative models, including different priors, in order to achieve inference that we can trust. Even more, for complex problems we typically do not know ahead of time what model(s), that is, the combination of prior and likelihood, we want to fit and even if so, we would still want to understand the fitted model(s) and its relation to the data. Such understanding can often best be achieved by comparing inferences from a series of related models and evaluating when and how conclusions are similar or not.

In practice, modelling usually starts with a template model (see discussion by Gelman et al., 2020) with default priors. A need for a more carefully designed prior may be revealed only after careful analysis of the first models. And it may be motivated by unrealistic results, computational problems, or the need for incorporating domain knowledge into a model. In other words, the choice of prior, as with other modelling decisions, is often

informed by iterative model exploration. Prior elicitation is thus a central part of a Bayesian workflow, and is not restricted to happen only at the beginning of the workflow.

A useful workflow does not just follow all pre-described steps, but also omits them when they are unnecessary, in order to help allocate finite resources where they are most needed. For example, for simple parametric models and informative data, the likelihood can dominate the prior and the gain from prior elicitation could be negligible. Thus, in many cases it may be sensible to start with some common default priors or priors weakly informed by some summary statistics of the data (e.g., by centering and normalizing the covariate and target values in regression), and then assess the need for more careful prior elicitation using prior diagnostic tools (Kallioinen et al., 2021).

In that sense, knowing when to perform prior elicitation is central to a prior elicitation workflow. A good general heuristic is “in situations where prior information is appreciable, and the data are limited” as O’Hagan et al. (2006) have put it. Then, the question whether we should perform prior elicitation can be reformulated into: Is it worthwhile to spend resources to incorporate domain knowledge? Or more nuanced: How much information do we need to gather, and how accurate should that information be? In many instances, getting the order of magnitude right and/or obtaining a prior that works to remove nonsensical outcomes may be sufficient. Furthermore, the level of accuracy does not need to be the same for all the parameters in a model, as refining a few or even just one prior can translate into considerably better inference. There are more nuanced scenarios where informative priors via prior elicitation are crucial.

For example, there can be gaps in time-series data in which case the expert may provide structural information in a form of a prior distribution that helps to fill gaps in the posterior distribution, or the expert knowledge may help to extrapolate from one group in the data to another (e.g., see Siivola et al., 2021).

In line with the current literature, we have so far discussed prior elicitation with regard to the choice of distributions and their parameters. This definition can be naturally extended to prior elicitation over models, which could provide a new sub-field for prior elicitation or a sister field of model elicitation. As evaluating over the entire range of conceivable models is unfeasible, answering questions such as: “Is a linear model adequate?”, “Do we need to extrapolate and perform predictions outside the observed domain?”, and similar ones would help us to narrow down options and save resources. Restricting the search to a few options early on will help, even if we later choose to expand the set of models.

Finally, a prior elicitation workflow should include one step to assess that the incorporated information is actually useful and an evaluation of the sensitivity of the results to the prior choice, including possible prior-data conflicts (Depaoli et al., 2020; Gelman et al., 2020; Lopes and Tobias, 2011; Al-Labadi and Evans, 2017; Evans and Moshonov, 2006; Reimherr et al., 2021; Berger, 1990; Berger et al., 1994; Canavos, 1975; Hill and Spall, 1994; Skene et al., 1986; Jacobi et al., 2018; Roos et al., 2015; Pérez et al., 2006; Giordano et al., 2018; Bornn et al., 2010; Ho, 2020; Kallioinen et al., 2021).

4.3 Developing better priors

One direction to improve prior elicitation is to develop priors for which elicitation is easier *per se*. In this context, 'easier' can mean one of at least three perspectives: (a) easier to understand for experts (D7), (b) computationally easier (D5), and/or (c) leaving fewer degrees of freedom, that is, fewer hyperparameters to elicit. Perspective (a) is especially relevant for direct elicitation in the parameter space, while perspective (b) is mostly relevant for indirect elicitation in the observable space due to computational requirements of the translation procedure to the parameter space ((D3); see Section 3.2). Both of these perspectives tend to go hand in hand with the perspective (c) because fewer required choices often make the priors easier to understand for experts due to reduced cognitive load, and reduce computational requirements due to a lower-dimensional target space of the translation. Accordingly, if we focus on (c), we can have the justified hope that other advantages will naturally follow in the process.

Reducing the number of hyperparameters comes with the initial (model-building) choice of what matters to be elicited and what is acceptable to just fix to a constant or forced to be of the same value (equality constraint). This line of reasoning leads to the notion of *joint* hyperparameters where the individual priors all depend on a much smaller (or highly structured) set of hyperparameters, jointly shared across parameters. Any kind of *hierarchical prior* follows this logic by design (Bürkner, 2017). For example, consider a simple hierarchical linear model across observations i with intercepts a_j varying across a total number of J groups:

$$\begin{aligned} y_i &\sim \text{normal}(\mu_i, \sigma) \\ \mu_i &= a_{j[i]} \\ a_j &\sim \text{normal}(a, \tau) \\ a &\sim \text{normal}(\mu_a, \sigma_a) \\ \tau &\sim \text{Gamma}(\alpha_\tau, \beta_\tau) \end{aligned}$$

Focusing on the priors for a_j , we have essentially reduced the problem of finding a total of J priors, each with one or more hyperparameters, to just choosing four hyperparameters, namely the location μ_a and scale σ_a of the normal prior on the joint mean a as well as the shape α_τ and rate β_τ of the Gamma prior on the joint standard deviation τ .² However, such hierarchical priors are only *locally joint* in the sense that they do not encompass all or even most parameters but only a subset. This becomes apparent if we extend the above model by additional additive terms, for example,

$$\mu_i = a_{j[i]} + b_i + c_i + d_i,$$

with each term having their own mutually independent set of parameters and corresponding hyperparameters.

²In hierarchical models, it is common to also call a and τ 'hyperparameters' although they are not set by the analyst or expert but rather estimated from the data along with other model parameters. To avoid confusion, we continue to restrict the use of 'hyperparameters' to parameters chosen in the elicitation process, which are thus fixed during model fitting.

It would be desirable to develop priors that are *globally joint* in that they span most or even all parameters leaving just a few hyperparameters to choose. With the purpose of preventing overfitting and facilitating variable selection in high-dimensional linear regression models on comparably sparse data, several *hierarchical shrinkage priors* have been developed that fulfil these properties (Bhattacharya et al., 2015; Piironen et al., 2017; Zhang et al., 2020). However, they do not yet generalize much beyond linear regression settings and their usefulness in the context of prior elicitation has not been studied so far. If we can extend these priors to more complicated models and find parameterizations with intuitive hyperparameters, such globally joint priors could prove extremely valuable in making prior elicitation more practical and widely applicable.

4.4 Evaluating prior elicitation

When any new prior elicitation method is proposed, a natural question that arises is whether it works as desired. Similarly, when a variety of prior elicitation methods are available for a given context, the practitioner wonders which one is better. Such questions concern the evaluation of prior elicitation methods. There are multiple desiderata for prior elicitation. Johnson et al. (2010a), for instance, categorize these into (i) validity – whether the elicitation captures the true belief of the expert, (ii) reliability – whether repeated elicitations reproduce the same priors, (iii) responsiveness – whether the elicitation is sensitive to changes in beliefs, and (iv) feasibility, which refers to the costs or resources required for elicitation. Many of these desiderata may seem as being at odds with each other, but they are all relevant for the eventual goal of supporting the building of good models with available resources.

In an ideal scenario, any researcher or user of prior elicitation methods would easily be able to compare the pros and cons of existing off-the-shelf methods for her problem, or even test new ones in small-scale user studies. So far, there have been very few projects where multiple prior elicitation methods have been empirically compared (Winkler, 1967; Johnson et al., 2010b; Grigore et al., 2016), but these have been in very application specific contexts. There is a need for more general and standard validation paradigms for prior elicitation, and the prior elicitation field has no equivalents to practices such as using benchmark datasets for comparing machine-learning algorithms, e.g. Deng et al. (2009); LeCun et al. (2010). We think this is a particularly challenging topic to work on because we also lack good metrics for evaluation. The simple metrics that have been widely used in this context may not be a valid measure of the quantities we care about. For example, (i) an expert’s subjective feedback about elicited priors may be subject to the kind of biases that also distort their priors, (ii) task completion time is considered to be a proxy for cognitive effort, but the elicitation may be finished inaccurately and in a hurried manner due to the cognitive strain it produces, and so on. Prior elicitation metrics can be potentially improved by incorporating research from areas such as Psychology and Human Computer Interaction. Improved metrics, increased comparative work and the development of standardized validation paradigms or platforms would be essential as the prior-elicitation field makes more progress. In addition, many proposed evaluation metrics are model-specific, but we also need more general methods that can be used across the board in a model-agnostic manner.

Among the different criteria for prior elicitation, assessing faithfulness, accuracy, or validity may be the hardest. From this perspective, the aim of prior elicitation is to accurately capture subjective knowledge of experts/users. However, there are many sources of distortions in priors elicited by an expert including their cognitive biases while making judgments in uncertain settings, and measurement noise introduced by the prior elicitation method, for example, by eliciting probability distributions over discretized intervals (Miller III and Rice, 1983; Parmar et al., 1994; Tan et al., 2003), especially when there are a smaller number of intervals or bins. A promising empirical approach to evaluating faithfulness of prior elicitation would involve validating elicited priors against an expected ground truth. For instance, one could train participants on data produced by a specified model with specified priors, and see how well the true parameter priors are recovered by the elicitation methods. Such methods could be the basis for developing test-beds for prior elicitation evaluation.

Model-specificity and training efforts in paradigms to evaluate faithfulness can also be bypassed by comparing elicited results against a 'gold standard' model-agnostic method which is known to have higher accuracy. While the nature of such baseline methods would be a topic of future research, there may be some viable candidates. A very promising perspective in psychology treats human judgements as a result of sampling from their subjective probabilities. This viewpoint has been successfully applied in the Markov chain Monte Carlo with people (MCMCP) approach (Sanborn and Griffiths, 2008; Sanborn et al., 2010) and its variants (Hsu et al., 2012; León-Villagr a et al., 2020; Harrison et al., 2020) to elicit beliefs about how stimuli from a multidimensional stimulus space (e.g., n-dimensional stick figures) maps onto a target category (e.g, 'shape of a cat'). In MCMCP participants take the place of an MCMC acceptance function, and repeatedly accept or reject proposals regarding the category membership of the sampled stimuli. The adaptive nature of MCMC ensures that proposals are over time increasingly sampled from parts of the stimulus space representing the participants' subjective representation of the category. The participants' prior beliefs are then constructed as the stationary distribution of the Markov chain that their judgments eventually converge to. The performance of MCMCP and its variants, on natural categories as well as trained artificial categories make us believe that similar sampling-based methods may have promise in the prior elicitation field both, for obtaining faithful priors, and for acting as model-agnostic baseline methods in paradigms that assess faithfulness.

When evaluating the accuracy of prior elicitation, we may also want to consider the effect of the elicited prior on the predictions or decisions made based on the model. In some scenarios, it is possible that even coarse elicitation processes can obtain practically useful information and further refinement of the elicitation may not bring additional benefits. Also, even if there is a significant bias in the elicited prior, that bias may have negligible effect on the end result. It can thus be useful to evaluate sensitivity and robustness of inference with respect to the elicited prior and its potential aspects that are difficult to elicit. For example, it is difficult for humans to estimate tiny probabilities, which is reflected in the difficulties of determining the tail shape of the elicited prior. A bias in the elicited prior and too thin tails can lead to strong prior sensitivity or prior-data conflict (Al-Labadi and Evans, 2017; Evans and Moshonov, 2006; Kallioinen et al., 2021; B urkner, 2021). On the other hand, thick tailed priors may lead to ignoring the otherwise correctly elicited prior information.

4.5 Software for Prior Elicitation

There are several desirable features that a software for prior elicitation could have. In addition to being open source it should have a simple and intuitive interface suitable for non-specialists. At least one part of such an interface should be visual to enable better input from humans and to perform validation of the proposed priors, and some level of interactive visualization capability would further help to obtain information from experts. Furthermore, switching between different types of visualization (kernel density estimates plots, quantile dotplots, histograms, etc.) would also be valuable as would be the possibility to add user-defined transformations before visualization. For example, [Sarma and Kay \(2020\)](#) describe how different visualizations could lead to different strategies for prior elicitation, and that most participants in their study primarily used a combination of strategies for determining their choice of priors. In addition, research shows that even people with statistical training can have problems correctly interpreting probability densities (Section 3.1), and so alternative representations like quantile plots may be preferred ([Kay et al., 2016](#)).

Prior elicitation software could be written agnostic of the underlying programming language, or at least interoperable with as many languages as possible, in order to avoid duplication of efforts. Building on top of already present open source libraries related to Bayesian workflow and uncertainty visualization like `ggdist` ([Kay, 2021](#)), `Bayesplot` ([Gabry et al., 2019](#); [Gabry and Mahr, 2021](#)) and `ArviZ` ([Kumar et al., 2019](#)) would help to achieve this goal. Moreover, working on top of such libraries could help to maintain modularity, which is especially desirable at the present state of development of the software for prior elicitation. Modularity would also help to reduce computational costs, if experimentation with visualizations and transformations can be made independent of the model. By specifying each task as distinctly as possible and dividing work, the community can generate and maintain software more easily, while at the same time encouraging research in prior elicitation on one or several dimensions of the research hypercube.

Given that we still need more research to assert which elicitation space (D3) is more appropriate for a given research problem, building software in a modular fashion should allow users to switch between the parameter and the observable space as needed. Similarly, the type of assessment task (D6) should be something that can be chosen by the user (e.g. as in `SHELF` or `MATCH`). It is also important to develop software that supports model-agnostic prior elicitation (D2), otherwise there will be problems with integration into the Bayesian workflow (Section 4.2) because a change in the model specification could preclude prior elicitation.

5 Conclusion

This paper covered the state of the prior elicitation today, focusing on discussing reasons for the somewhat limited impact the research has had on practice. We identified bottlenecks at different levels and argued that significant coordinated effort covering several scientific disciplines will be needed to transform the current state and make prior

elicitation routine part of the practical modelling workflow. In summary, we make the following concrete calls to arms:

1. **We need to focus on elicitation matters that answer to the needs of practical modelling workflow.** Compared to past research, the efforts should be re-directed more towards (a) elicitation methods that are agnostic of the model and prior, (b) elicitation strategies (e.g. active elicitation) that are efficient from the perspective of the modeler and compatible with iterative model-building, and (c) formulations that make elicitation of multivariate priors easier, for instance by designing hierarchical priors that are simpler to elicit.
2. **We need better open software** that integrates seamlessly into the current modelling workflow, and that is sufficiently modular so that new elicitation algorithms can be quickly taken into use and evaluated in concrete modelling cases. The elements not specific to elicitation algorithms (e.g. visualization of the priors, the language used for specifying the models and desired prior families) should be implemented using existing libraries whenever possible, and the tools should be open source.
3. **We need cost-efficient and well-targeted evaluation techniques** for supporting development of new methods and validating their relative quality and value in practical tasks. In ideal case, we would like to see a testbed for prior elicitation techniques that enable easy evaluation of alternative methods in varying situations with feasible experimentation cost, as well as practical ways of collecting information about efficiency of elicitation methods in real use cases.
4. **We need spearhead examples** that clearly demonstrate the value of prior elicitation in applications of societal interest to increase enthusiasm beyond the current niche.

For the first two we already outline concrete directions in Section 4. We hypothesize that addressing all four foci will transform the status of prior elicitation, by providing the required infrastructure, public interest and funding for speeding up future development.

References

- Abbas, A. E. (2003). “Entropy Methods For Univariate Distributions in Decision Analysis.” *AIP Conference Proceedings*, 659(1): 339–349. [21](#)
- (2006). “Entropy methods for joint distributions in decision analysis.” *IEEE Transactions on Engineering Management*, 53(1): 146–159. [21](#)
- Abbas, A. E., Budescu, D. V., and Gu, Y. (2010). “Assessing joint distributions with isoprobability contours.” *Management Science*, 56(6): 997–1011. [15](#)
- Abbas, A. E., Budescu, D. V., Yu, H.-T., and Haggerty, R. (2008). “A comparison of two probability encoding methods: Fixed probability vs. fixed variable values.” *Decision Analysis*, 5(4): 190–202. [13](#)

- Afrabandpey, H., Peltola, T., and Kaski, S. (2017). “Interactive Prior Elicitation of Feature Similarities for Small Sample Size Prediction.” In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, 265–269. ACM. Conference on User Modeling, Adaptation and Personalization, UMAP ; Conference date: 09-07-2017 Through 12-07-2017.
URL <http://www.um.org/umap2017/> 24
- (2019). “Human-in-the-loop Active Covariance Learning for Improving Prediction in Small Data Sets.” In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 1959–1966. International Joint Conferences on Artificial Intelligence Organization.
URL <https://doi.org/10.24963/ijcai.2019/271> 25
- Aitchison, J. and Dunsmore, I. R. (1975). *Statistical prediction analysis*. Cambridge University Press. 20
- Akbarov, A. (2009). “Probability elicitation: Predictive approach.” Ph.D. thesis, University of Salford. 18, 20
- Al-Awadhi, S. (1997). “Elicitation of prior distributions for a multivariate normal distribution.” Ph.D. thesis, University of Aberdeen. 20
- Al-Awadhi, S. A. and Garthwaite, P. H. (1998). “An elicitation method for multivariate normal distributions.” *Communications in Statistics - Theory and Methods*, 27(5): 1123–1142. 20, 24
- (2001). “Prior distribution assessment for a multivariate normal distribution: An experimental study.” *Journal of Applied Statistics*, 28(1): 5–23. 20
- (2006). “Quantifying expert opinion for modelling fauna habitat distributions.” *Computational Statistics*, 21(1): 121–140. 11, 24
- Al-Hamzawi, R., Yu, K., and Pan, J. (2011). “Prior elicitation in Bayesian quantile regression for longitudinal data.” *Journal of Biometrics & Biostatistics*. 24
- Al-Labadi, L., Baskurt, Z., and Evans, M. (2018). “Statistical Reasoning: Choosing and Checking the Ingredients, Inferences Based on a Measure of Statistical Evidence with Some Applications.” *Entropy*, 20(4).
URL <https://www.mdpi.com/1099-4300/20/4/289> 19
- Al-Labadi, L. and Evans, M. (2017). “Optimal robustness results for relative belief inferences and the relationship to prior-data conflict.” *Bayesian Analysis*, 12(3): 705–728. 30, 33
- Albert, I., Donnet, S., Guihenneuc-Jouyaux, C., Lowchoy, S., L. Mengersen, K., and Rousseau, J. (2012). “Combining Expert Opinions in Prior Elicitation.” *Bayesian Analysis*, 7: 503–532. 25
- Alhussain, Z. A. and Oakley, J. E. (2020a). “Assurance for clinical trial design with normally distributed outcomes: Eliciting uncertainty about variances.” *Pharmaceutical Statistics*, 19(6): 827–839.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/pst.2040> 11

- (2020b). “Assurance for clinical trial design with normally distributed outcomes: Eliciting uncertainty about variances.” *Pharmaceutical Statistics*, 19(6): 827–839.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/pst.2040> 11
- Ashtiani, H., Kushagra, S., and Ben-David, S. (2016). “Clustering with Same-Cluster Queries.” In Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc. 25
- Awasthi, P., Balcan, M., and Voevodski, K. (2014). “Local algorithms for interactive clustering.” In Xing, E. P. and Jebara, T. (eds.), *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, 550–558. Beijing, China: PMLR.
URL <https://proceedings.mlr.press/v32/awasthi14.html> 25
- Azzolina, D., Berchiolla, P., Gregori, D., and Baldi, I. (2021). “Prior Elicitation for Use in Clinical Trial Design and Analysis: A Literature Review.” *International journal of environmental research and public health*, 18(4): 1833. 11
- Bagchi, P. (1986). “Assessment of Prior Distributions based on Information Theory.” *Journal of Information and Optimization Sciences*, 7(2): 117–128. 21
- Bajari, P. and Ye, L. (2003). “Deciding Between Competition and Collusion.” *The Review of Economics and Statistics*, 85(4): 971–989.
URL <https://doi.org/10.1162/003465303772815871> 11
- Balcan, M.-F. and Blum, A. (2008). “Clustering with interactive feedback.” In *International Conference on Algorithmic Learning Theory*, 316–328. Springer. 25
- Bartoš, F., Gronau, Q. F., Timmers, B., Otte, W. M., Ly, A., and Wagenmakers, E.-J. (2021). “Bayesian model-averaged meta-analysis in medicine.” *Statistics in Medicine*.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.9170> 4
- Bedrick, E. J., Christensen, R., and Johnson, W. (1996). “A New Perspective on Priors for Generalized Linear Models.” *Journal of the American Statistical Association*, 91(436): 1450–1460. 19, 23, 24
- Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis*. Springer, New York. 22
- (1990). “Robust Bayesian analysis: sensitivity to the prior.” *Journal of statistical planning and inference*, 25(3): 303–328. 30
- Berger, J. O., Moreno, E., Pericchi, L. R., Bayarri, M. J., Bernardo, J. M., Cano, J. A., De la Horra, J., Martín, J., Ríos-Insúa, D., Betrò, B., et al. (1994). “An overview of robust Bayesian analysis.” *Test*, 3(1): 5–124. 30
- Bernardo, J. M. and Smith, A. F. (1994). *Bayesian theory*. Chichester: John Wiley. 18
- Best, N., Dallow, N., and Montague, T. (2020). “Prior elicitation.” *Bayesian methods in pharmaceutical research*, 87–109. 11
- Bhattacharya, A., Pati, D., Pillai, N. S., and Dunson, D. B. (2015). “Dirichlet–Laplace

- priors for optimal shrinkage.” *Journal of the American Statistical Association*, 110(512): 1479–1490. [32](#)
- Billheimer, D. (2019). “Predictive Inference and Scientific Reproducibility.” *The American Statistician*, 73(sup1): 291–295.
URL <https://doi.org/10.1080/00031305.2018.1518270> [18](#)
- Böcker, K., Oakley, J., Daneshkhah, A., and Peters, J.-P. (2010). *Rethinking Risk Measurement and Reporting: Uncertainty, Bayesian analysis and expert judgement: PART II EXPERT JUDGEMENT*. Rethinking Risk Measurement and Reporting. Risk Books. [11](#)
- Bolger, F. (2018). *The Selection of Experts for (Probabilistic) Expert Knowledge Elicitation*, 393–443. Cham: Springer International Publishing.
URL https://doi.org/10.1007/978-3-319-65052-4_16 [5](#)
- Bornn, L., Doucet, A., and Gottardo, R. (2010). “An efficient computational approach for prior sensitivity analysis and cross-validation.” *Canadian Journal of Statistics*, 38(1): 47–64. [30](#)
- Bousquet, N. (2008). “Eliciting vague but proper maximal entropy priors in Bayesian experiments.” *Statistical Papers*, 51(3): 613–628. [21](#)
- Brier, G. W. et al. (1950). “Verification of forecasts expressed in terms of probability.” *Monthly weather review*, 78(1): 1–3. [17](#)
- Briggs, W. (2016). *Uncertainty: the soul of modeling, probability & statistics*. Springer. [18](#)
- Brown, B. B. (1968). “Delphi process: a methodology used for the elicitation of opinions of experts.” Technical report, Rand Corp Santa Monica CA. [25](#)
- Brownstein, N. C., Louis, T. A., O’Hagan, A., and Pendergast, J. (2019). “The Role of Expert Judgment in Statistical Inference and Evidence-Based Decision-Making.” *The American Statistician*, 73(sup1): 56–68. PMID: 31057338. [5](#)
- Bunn, D. W. (1978). “Estimation of a Dirichlet prior distribution.” *Omega*, 6(4): 371–373. [9](#), [16](#), [24](#)
- Burgman, M. A. (2016). *Trusting judgements: how to get the best out of experts*. Cambridge University Press. [26](#)
- Bürkner, P.-C. (2021). “Specifying Priors in a Bayesian Workflow.”
URL https://paul-buerkner.github.io/data/prior_specification_bayesian_workflow.pdf [33](#)
- Bürkner, P.-C. (2017). “brms: An R Package for Bayesian Multilevel Models using Stan.” *Journal of Statistical Software*, 80(1): 1–28. [5](#), [31](#)
- Canavos, G. C. (1975). “Bayesian estimation: A sensitivity analysis.” *Naval Research Logistics Quarterly*, 22(3): 543–552. [30](#)
- Carlin, B. P. (2000). *Bayes and empirical Bayes methods for data analysis*. Chapman & Hall/CRC. [4](#), [29](#)

- Carlin, B. P., Kass, R. E., Lerch, F. J., and Huguenard, B. R. (1992). “Predicting working memory failure: A subjective Bayesian approach to model selection.” *Journal of the American Statistical Association*, 87(418): 319–327. 24
- Casement, C. (2017). “Graphical Methods in Prior Elicitation.” Ph.D. thesis, Baylor University. 11
- Casement, C. J. and Kahle, D. J. (2018). “Graphical prior elicitation in univariate models.” *Communications in Statistics - Simulation and Computation*, 47(10): 2906–2924. 11, 20
- Cervera, J. and Munoz, J. (1996). “Proper scoring rules for fractiles.” *Bayesian Statistics*, 5: 513–519. 18
- Chaloner, K. (1996). “Elicitation of prior distributions.” *Bayesian biostatistics*, 141: 156. 18
- Chaloner, K., Church, T., Louis, T. A., and Matts, J. P. (1993). “Graphical Elicitation of a Prior Distribution for a Clinical Trial.” *Journal of the Royal Statistical Society. Series D (The Statistician)*, 42(4): 341–353.
URL <http://www.jstor.org/stable/2348469> 24
- Chaloner, K. and Duncan, G. T. (1987). “Some properties of the Dirichlet-multinomial distribution and its use in prior elicitation.” *Communications in Statistics - Theory and Methods*, 16(2): 511–523. 16, 24
- Chaloner, K. and Verdinielli, I. (1995). “Bayesian experimental design: A review.” *Statistical Science*, 10(3): 273–304. 23
- Chaloner, K. M. and Duncan, G. T. (1983). “Assessment of a Beta Prior Distribution: Pm Elicitation.” *Journal of the Royal Statistical Society: Series D (The Statistician)*, 32(1-2): 174–180. 16, 24
- Chandramouli, S. H. (2020). “Accommodation of Scientific Intuition Within Statistical Inference.” Ph.D. thesis, Indiana University. 21
- Chandramouli, S. H. and Shiffrin, R. M. (2016). “Extending bayesian induction.” *Journal of Mathematical Psychology*, 72: 38–42. 21
- Chen, M.-H. and Ibrahim, J. G. (2003). “CONJUGATE PRIORS FOR GENERALIZED LINEAR MODELS.” *Statistica Sinica*, 13(2): 461–476.
URL <http://www.jstor.org/stable/24307144> 19, 24
- Chen, M.-H., Ibrahim, J. G., and Yiannoutsos, C. (1999). “Prior elicitation, variable selection and Bayesian computation for logistic regression models.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(1): 223–242. 19, 24
- Chesley, G. R. (1975). “Elicitation of Subjective Probabilities: A Review.” *The Accounting Review*, 50: 325–337. 11
- Choy, S. L., O’Leary, R., and Mengersen, K. (2009). “Elicitation by design in ecology: using expert opinion to inform priors for Bayesian statistical models.” *Ecology*, 90(1): 265–277. 11, 19

- Chung, Y., Gelman, A., Rabe-Hesketh, S., Liu, J., and Dorie, V. (2015). “Weakly Informative Prior for Point Estimation of Covariance Matrices in Hierarchical Models.” *Journal of Educational and Behavioral Statistics*, 40(2): 136–157. 2
- Clemen, R. T., Fischer, G. W., and Winkler, R. L. (2000). “Assessing Dependence: Some Experimental Results.” *Management Science*, 46(8): 1100–1115.
URL <http://www.jstor.org/stable/2661587> 15
- Clemen, R. T. and Reilly, T. (1999). “Correlations and Copulas for Decision and Risk Analysis.” *Management Science*, 45(2): 208–224.
URL <http://www.jstor.org/stable/2634871> 15
- (2001). *Making Hard Decisions with DecisionTools*. Duxbury/Thomson Learning. 9
- Cohn, D., Atlas, L., and Ladner, R. (1994). “Improving generalization with active learning.” *Machine learning*, 15(2): 201–221. 23
- Coles, S. G. and Tawn, J. A. (1996). “A Bayesian Analysis of Extreme Rainfall Data.” *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 45(4): 463–478.
URL <http://www.jstor.org/stable/2986068> 11
- Cooke, R. M. (1991). *Experts in Uncertainty: Opinion and Subjective Probability in Science (Environmental Ethics and Science Policy)*. Oxford University Press. 11, 12, 25
- Coolen, F. (1992). *Elicitation of expert knowledge and assessment of imprecise prior densities for lifetime distributions*. Memorandum COSOR. Technische Universiteit Eindhoven. 9, 20
- Coolen, F., Mertens, P., and Newby, M. (1992). “A Bayes-competing risk model for the use of expert judgment in reliability estimation.” *Reliability Engineering & System Safety*, 35(1): 23–30. 14
- Craig, P. S., Goldstein, M., Seheult, A., and Smith, J. (1998). “Constructing partial prior specifications for models of complex physical systems.” *Journal of the Royal Statistical Society: Series D (The Statistician)*, 47(1): 37–53.
URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/1467-9884.00115> 11
- Crosby, M. A. (1980). “Implications of prior probability elicitation on auditor sample size decisions.” *Journal of Accounting Research*, 585–593. 11
- Cuevas, J. R. T. (2012). “Eliciting beta prior distributions for binomial sampling.” *Rev. Bras. Biom*, 30(1): 159–172. 16
- Dae, P., Peltola, T., Soare, M., and Kaski, S. (2017). “Knowledge elicitation via sequential probabilistic inference for high-dimensional prediction.” *Machine Learning*, 106(9): 1599–1620. 23
- Dae, P., Peltola, T., Vehtari, A., and Kaski, S. (2018). “User Modelling for Avoiding Overfitting in Interactive Knowledge Elicitation for Prediction.” In *23rd International Conference on Intelligent User Interfaces*, IUI ’18, 305–310. New York, NY, USA:

- Association for Computing Machinery.
URL <https://doi.org/10.1145/3172944.3172989> 25
- Dalal, S. and Hall, W. (1983). “Approximating priors by mixtures of natural conjugate priors.” *Journal of the Royal Statistical Society: Series B (Methodological)*, 45(2): 278–286. 21
- Dall’aglio, G., Kotz, S., and Salinetti, G. (1991). *Advances in Probability Distributions with Given Marginals*. Mathematics and Its Applications. Springer, Dordrecht. 15
- Dallow, N., Best, N., and Montague, T. H. (2018). “Better decision making in drug development through adoption of formal prior elicitation.” *Pharmaceutical Statistics*, 17(4): 301–316.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/pst.1854> 4
- Daneshkhah, A., Oakley, J., and O’Hagan, A. (2006). “Nonparametric prior elicitation with imprecisely assessed probabilities.” Technical report, Citeseer. 10, 22
- Daneshkhah, A. and Oakley, J. E. (2010). “Eliciting multivariate probability distributions.” *Rethinking risk measurement and reporting*, 1: 23. 12
- Daneshkhah, A., Revie, M., Bedford, T., Walls, L., and Quigley, J. (2011). *Eliciting Subjective Probability Distributions from Groups*. American Cancer Society.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470400531.eorms0293> 25
- De Finetti, B. (1937). “La prévision: ses lois logiques, ses sources subjectives.” In *Annales de l’institut Henri Poincaré*, volume 7, 1–68. 18
- (1974). *Theory of Probability*. Wiley, New York. 17
- Del Negro, M. and Schorfheide, F. (2008). “Forming priors for DSGE models (and how it affects the assessment of nominal rigidities).” *Journal of Monetary Economics*, 55(7): 1191–1208. 11
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). “Imagenet: A large-scale hierarchical image database.” In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee. 32
- Denham, R. and Mengersen, K. (2007). “Geographically assisted elicitation of expert opinion for regression models.” *Bayesian Analysis*, 2(1): 99 – 135.
URL <https://doi.org/10.1214/07-BA205> 19, 20, 24
- Depaoli, S., Winter, S. D., and Visser, M. (2020). “The Importance of Prior Sensitivity Analysis in Bayesian Statistics: Demonstrations Using an Interactive Shiny App.” *Frontiers in Psychology*, 11. 30
- Dey, D. K. and Liu, J. (2007). “A quantitative study of quantile based direct prior elicitation from expert opinion.” *Bayesian Analysis*, 2(1): 137 – 166.
URL <https://doi.org/10.1214/07-BA206> 14
- Dias, L. C., Morton, A., and Quigley, J. (2018). *Elicitation : The Science and Art of*

- Structuring Judgement*, volume 261 of *International Series in Operations Research & Management Science*. Springer, Cham. 11, 25
- Dickey, J., Jiang, J., Kadane, J., of New York at Albany, S. U., and of America, U. S. (1983). “Bayesian methods for multinomial sampling with missing data using multiple hypergeometric functions.” *State University of New York at Albany Department of Mathematics and Statistics: Research Report*, 6. 16, 24
- Dickey, J., Lindley, D., and Press, S. (1985). “Bayesian estimation of the dispersion matrix of a multivariate normal distribution.” *Communications in Statistics - Theory and Methods*, 14(5): 1019–1034.
URL <https://doi.org/10.1080/03610928508828960> 15
- Duran, B. and Booker, J. (1988). “A Bayes sensitivity analysis when using the beta distribution as a prior.” *IEEE Transactions on Reliability*, 37(2): 239–247. 15
- Edwards, W., Lindman, H., and Savage, L. J. (1963). “Bayesian statistical inference for psychological research.” *Psychological review*, 70(3): 193. 11
- EFSA, E. F. S. A. (2014). “Guidance on Expert Knowledge Elicitation in Food and Feed Safety Risk Assessment.” *EFSA Journal*, 12(6): 3734. 11, 12, 25
- Ehm, W., Gneiting, T., Jordan, A., and Krüger, F. (2016). “Of quantiles and expectiles: consistent scoring functions, Choquet representations and forecast rankings.” *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 78(3): 505–562.
URL <http://www.jstor.org/stable/24775351> 18
- Elfadaly, F. and Garthwaite, P. H. (2011). “Eliciting a prior distribution for the error variance in normal linear models.” Technical report, Tech. Rep. 11/07, Department of Mathematics and Statistics, Open University. 19, 24
- Elfadaly, F. G. (2012). “Elicitation of Subjective Probability Distributions.” Ph.D. thesis, The Open University UK. 16
- Elfadaly, F. G. and Garthwaite, P. H. (2013). “Eliciting Dirichlet and Connor–Mosimann prior distributions for multinomial models.” *Test*, 22(4): 628–646. 16, 24
- (2015). “Eliciting prior distributions for extra parameters in some generalized linear models.” *Statistical Modelling*, 15(4): 345–365. 19, 24
- (2017). “Eliciting Dirichlet and Gaussian copula prior distributions for multinomial models.” *Statistics and Computing*, 27(2): 449–467.
URL <http://oro.open.ac.uk/49480/> 16, 24
- (2020). “On quantifying expert opinion about multinomial models that contain covariates.” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 183(3): 959–981.
URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/rssa.12546> 24
- Enøe, C., Georgiadis, M. P., and Johnson, W. O. (2000). “Estimation of sensitivity and specificity of diagnostic tests and disease prevalence when the true disease state is unknown.” *Preventive veterinary medicine*, 45(1-2): 61–81. 15

- Evans, M., Guttman, I., and Li, P. (2017). “Prior Elicitation, Assessment and Inference with a Dirichlet Prior.” *Entropy*, 19: 564. [16](#)
- Evans, M. and Moshonov, H. (2006). “Checking for prior-data conflict.” *Bayesian analysis*, 1(4): 893–914. [30](#), [33](#)
- Fackler, P. L. (1991). “Modeling Interdependence: An Approach to Simulation and Elicitation.” *American Journal of Agricultural Economics*, 73(4): 1091–1097.
URL <http://www.jstor.org/stable/1242437> [15](#)
- Flaxman, S., Mishra, S., and et al., A. G. (2020). “Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe.” *Nature*, 584: 257–261. [7](#)
- Forbes, R., Sanson, R., and Morris, R. (1994). “Application of subjective methods to the determination of the likelihood and consequences of the entry of foot-and-mouth disease into New Zealand.” *New Zealand Veterinary Journal*, 42(3): 81–88. [14](#)
- Fox, B. L. (1966). *A Bayesian Approach to Reliability Assessment*. Santa Monica, CA: RAND Corporation. [15](#)
- French, S. (2011). “Aggregating expert judgement.” *Revista de la Real Academia de Ciencias Exactas, Fisicas y Naturales. Serie A. Matematicas*, 105(1): 181–206. [25](#)
- Fuglstad, G.-A., Simpson, D., Lindgren, F., and vard Rue, H. (2019). “Constructing Priors that Penalize the Complexity of Gaussian Random Fields.” *Journal of the American Statistical Association*, 114(525): 445–452. [2](#)
- Gabry, J. and Mahr, T. (2021). “bayesplot: Plotting for Bayesian Models.” R package version 1.8.0.
URL <https://mc-stan.org/bayesplot/> [34](#)
- Gabry, J., Simpson, D., Vehtari, A., Betancourt, M., and Gelman, A. (2019). “Visualization in Bayesian workflow.” *J. R. Stat. Soc. A*, 182: 389–402. [34](#)
- Gal, Y., Islam, R., and Ghahramani, Z. (2017). “Deep bayesian active learning with image data.” In *International Conference on Machine Learning*, 1183–1192. PMLR. [23](#)
- Galway, L. A. (2007). “Subjective probability distribution elicitation in cost risk analysis: A review.” Technical report. [11](#)
- Gaoini, E., Dey, D., and Ruggeri, F. (2009). *Bayesian modeling of flash floods using generalized extreme value distribution with prior elicitation*. University of Connecticut, Department of Statistics. [9](#), [19](#), [20](#)
- Garthwaite, P. H. and Al-Awadhi, S. A. (2001). “Non-conjugate prior distribution assessment for multivariate normal sampling.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(1): 95–110. [20](#), [24](#)
- (2006). “Quantifying opinion about a logistic regression using interactive graphics.” *Statistics Group*, 6. [24](#)
- Garthwaite, P. H., Al-Awadhi, S. A., Elfadaly, F. G., and Jenkinson, D. J. (2013). “Prior

- distribution elicitation for generalized linear and piecewise-linear models.” *Journal of Applied Statistics*, 40(1): 59–75. 19, 24
- Garthwaite, P. H. and Dickey, J. M. (1988). “Quantifying Expert Opinion in Linear Regression Problems.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 50(3): 462–474.
URL <http://www.jstor.org/stable/2345708> 19, 24
- (1991). “An elicitation method for multiple linear regression models.” *Journal of Behavioral Decision Making*, 4(1): 17–31.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/bdm.3960040103> 24
- (1992). “Elicitation of Prior Distributions for Variable-Selection Problems in Regression.” *The Annals of Statistics*, 20(4): 1697–1719.
URL <http://www.jstor.org/stable/2242364> 19, 20, 24
- Garthwaite, P. H., Kadane, J. B., and O’Hagan, A. (2005). “Statistical Methods for Eliciting Probability Distributions.” *Journal of the American Statistical Association*, 100: 680–701. 2, 3, 4, 11, 13, 19, 21, 25, 26
- Garthwaite, P. H. and O’Hagan, A. (2000). “Quantifying Expert Opinion in the UK Water Industry: An Experimental Study.” *Journal of the Royal Statistical Society: Series D (The Statistician)*, 49(4): 455–477.
URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/1467-9884.00246> 14
- Gavasakar, U. (1988). “A Comparison of Two Elicitation Methods for a Prior Distribution for a Binomial Parameter.” *Management Science*, 34(6): 784–790.
URL <http://www.jstor.org/stable/2632131> 16, 24
- Ge, H., Xu, K., and Ghahramani, Z. (2018). “Turing: a language for flexible probabilistic inference.” In *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain*, 1682–1690.
URL <http://proceedings.mlr.press/v84/ge18b.html> 2
- Geisser, S. (1971). “The Inferential Use of Predictive Distributions.” *Foundations of Statistical Inference*, eds. Godambe and Sprott, 456–466. 18
- Gelfand, A. E., Mallick, B. K., and Dey, D. K. (1995). “Modeling Expert Opinion Arising as a Partial Probabilistic Specification.” *Journal of the American Statistical Association*, 90: 598–604. 9, 10, 23, 28
- Gelman, A. (2006). “Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper).” *Bayesian analysis*, 1(3): 515–534. 2, 4, 29
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). “Bayesian Data Analysis.” 4, 21, 29
- Gelman, A., Jakulin, A., Pittau, M. G., and Su, Y.-S. (2008). “A weakly informative default prior distribution for logistic and other regression models.” *The Annals of*

- Applied Statistics*, 2(4): 1360 – 1383.
URL <https://doi.org/10.1214/08-A0AS191> 2
- Gelman, A. and Shalizi, C. R. (2013). “Philosophy and the practice of Bayesian statistics.” *British Journal of Mathematical and Statistical Psychology*, 66(1): 8–38. 4
- Gelman, A., Simpson, D., and Betancourt, M. (2017). “The Prior Can Often Only Be Understood in the Context of the Likelihood.” *Entropy*, 19(10): 555. 4, 5, 8, 29
- Gelman, A., Vehtari, A., Simpson, D., Margossian, C. C., Carpenter, B., Yao, Y., Kennedy, L., Gabry, J., Bürkner, P.-C., and Modrák, M. (2020). “Bayesian Workflow.” *arXiv*. ArXiv: 2011.01808.
URL <http://arxiv.org/abs/2011.01808> 2, 8, 29, 30
- Gelman, A. and Yao, Y. (2020). “Holes in Bayesian statistics.” *Journal of Physics G: Nuclear and Particle Physics*, 48(1): 014002. 4
- Genest, C. and Schervish, M. J. (1985). “Modeling Expert Judgments for Bayesian Updating.” *The Annals of Statistics*, 13: 1198–1212. 25
- Genest, C. and Zidek, J. V. (1986). “Combining probability distributions: A critique and an annotated bibliography.” *Statistical Science*, 1(1): 114–135. 25
- Ghosh, J., Li, Y., and Mitra, R. (2018). “On the Use of Cauchy Prior Distributions for Bayesian Logistic Regression.” *Bayesian Analysis*, 13(2): 359 – 383.
URL <https://doi.org/10.1214/17-BA1051> 2
- Gigerenzer, G. (1996). “On narrow norms and vague heuristics: A reply to Kahneman and Tversky.” 26
- Gigerenzer, G. and Todd, P. M. (1999). “Simple heuristics that make us smart.” 26
- Gill, J. and Walker, L. D. (2005). “Elicited Priors for Bayesian Model Specifications in Political Science Research.” *The Journal of Politics*, 67(3): 841–872.
URL <http://www.jstor.org/stable/10.1111/j.1468-2508.2005.00342.x> 11
- Giordano, R., Broderick, T., and Jordan, M. I. (2018). “Covariances, Robustness, and Variational Bayes.” *Journal of Machine Learning Research*, 19(51): 1–49.
URL <http://jmlr.org/papers/v19/17-670.html> 30
- Gneiting, T. (2011). “Making and Evaluating Point Forecasts.” *Journal of the American Statistical Association*, 106(494): 746–762.
URL <http://www.jstor.org/stable/41416407> 18
- Gneiting, T. and Raftery, A. E. (2007). “Strictly proper scoring rules, prediction, and estimation.” *Journal of the American statistical Association*, 102(477): 359–378. 17, 18
- Gokhale, D. V. and Press, S. J. (1982). “Assessment of a Prior Distribution for the Correlation Coefficient in a Bivariate Normal Distribution.” *Journal of the Royal Statistical Society: Series A (General)*, 145(2): 237–249. 15
- Goldstein, M. and Wooff, D. (2007). *Bayes Linear Statistics: Theory and Methods*. Wiley Series in Probability and Statistics. Wiley. 13

- Good, I. J. and England, C. B. (1965). “The estimation of probabilities.” 16, 20, 24
- Gore, S. (1987). “Biostatistics and the Medical Research Council.” *MRC News*, 35: 19–20. Cited By 50. 14
- Gosling, J. (2005). “Elicitation: A nonparametric view.” Ph.D. thesis, University of Sheffield. 10, 23
- Gosling, J. P. (2018). “SHELF: The Sheffield Elicitation Framework.” In Dias, L. C., Morton, A., and Quigley, J. (eds.), *Elicitation*, International Series in Operations Research & Management Science, chapter 0, 61–93. Springer. 25
- Gosling, J. P., Oakley, J. E., and O’Hagan, A. (2007). “Nonparametric elicitation for heavy-tailed prior distributions.” *Bayesian Anal.*, 2(4): 693–718.
URL <https://doi.org/10.1214/07-BA228> 10, 22
- Gribok, A. V., Urmanov, A. M., Hines, J. W., and Uhrig, R. E. (2004). “Backward specification of prior in bayesian inference as an inverse problem.” *Inverse Problems in Science and Engineering*, 12(3): 263–278. 20
- Grigore, B., Peters, J., Hyde, C., and Stein, K. (2013). “Methods to elicit probability distributions from experts: a systematic review of reported practice in health technology assessment.” *Pharmacoeconomics*, 31(11): 991–1003. 11
- (2016). “A comparison of two methods for expert elicitation in health technology assessments.” *BMC medical research methodology*, 16(1): 1–11. 32
- Grisley, W. and Kellogg, E. D. (1983). “Farmers’ subjective probabilities in Northern Thailand: an elicitation analysis.” *American Journal of Agricultural Economics*, 65(1): 74–82. 14
- Gross, A. J. (1971). “The Application of Exponential Smoothing to Reliability Assessment.” *Technometrics*, 13(4): 877–883. 15, 16
- Gutenkunst, R. N., Waterfall, J. J., Casey, F. P., Brown, K. S., Myers, C. R., and Sethna, J. P. (2007). “Universally sloppy parameter sensitivities in systems biology models.” *PLoS computational biology*, 3(10): e189. 19
- Haakma, W., Bojke, L., Steuten, L., and IJzerman, M. (2011). “Pmd4 expert elicitation to populate early health economic models of medical diagnostic devices in development.” *Value in Health*, 14(7): A244–A245. 11
- Haggerty, R. P. (2011). “Eliciting continuous probability distributions.” Ph.D. thesis, University of Illinois at Urbana-Champaign. 13
- Hahn, E. D. (2006). “Re-examining informative prior elicitation through the lens of Markov chain Monte Carlo methods.” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 169(1): 37–48.
URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-985X.2005.00381.x> 22
- Hammond, T. R., Swartzman, G. L., and Richardson, T. S. (2001). “Bayesian estimation

- of fish school cluster composition applied to a Bering Sea acoustic survey.” *ICES Journal of Marine Science*, 58(6): 1133–1149. 11
- Hanea, A., McBride, M., Burgman, M., and Wintle, B. (2018). “Classical meets modern in the IDEA protocol for structured expert judgement.” *Journal of Risk Research*, 21(4): 417–433. 25
- Hanea, A. M., Nane, G. F., Bedford, T., and French, S. (2021). *Expert Judgement in Risk and Decision Analysis*, volume 293. Springer Nature. 11
- Harrison, P., Marjeh, R., Adolfi, F., van Rijn, P., Anglada-Tort, M., Tchernichovski, O., Larrouy-Maestri, P., and Jacoby, N. (2020). “Gibbs sampling with people.” *Advances in Neural Information Processing Systems*, 33. 33
- Hartmann, M., Agiashivili, G., Bürkner, P.-C., and Klami, A. (2020). “Flexible Prior Elicitation via the Prior Predictive distribution.” In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence*, Conference on Uncertainty in Artificial Intelligence. United States: AUAI Press / Association for Uncertainty in Artificial Intelligence. Uncertainty in artificial intelligence, UAI ; Conference date: 04-08-2020 Through 07-08-2020.
URL <http://auai.org/uai2020/> 9, 10, 19, 20, 23, 29
- Hem, I. G., Fuglstad, G.-A., and Riebler, A. (2021). “Makemyprior: Intuitive Construction of Joint Priors for Variance Parameters in R.” *arXiv:2105.09712 [stat]*. 2, 24
- Hill, S. D. and Spall, J. C. (1994). “Sensitivity of a Bayesian analysis to the prior distribution.” *IEEE transactions on systems, man, and cybernetics*, 24(2): 216–221. 30
- Hirsch, K. G., Corey, P. N., and Martell, D. L. (1998). “Using Expert Judgment to Model Initial Attack Fire Crew Effectiveness.” *Forest Science*, 44(4): 539–549.
URL <https://doi.org/10.1093/forestscience/44.4.539> 11
- Ho, P. (2020). “Global Robust Bayesian Analysis in Large Models.” 30
- Hogarth, R. M. (1975). “Cognitive Processes and the Assessment of Subjective Probability Distributions.” *Journal of the American Statistical Association*, 70(350): 271–289.
URL <http://www.jstor.org/stable/2285808> 11
- Holt Jr, C. A. (1979). “Elicitation of Subjective Probability Distributions and von Neumann-Morgenstern Utility Functions.” 17
- Hora, S. C. (2007). *Eliciting probabilities from experts*. Cambridge University Press. 13
- Hosack, G. R., Hayes, K. R., and Barry, S. C. (2017). “Prior elicitation for Bayesian generalised linear models with application to risk control option assessment.” *Reliability Engineering & System Safety*, 167: 351 – 361. Special Section: Applications of Probabilistic Graphical Models in Dependability, Diagnosis and Prognosis. 9, 20, 24
- Hsu, A., Martin, J., Sanborn, A., and Griffiths, T. (2012). “Identifying representations of categories of discrete items using Markov chain Monte Carlo with people.” In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 34. 33

- Hughes, G. and Madden, L. (2002). “Some methods for eliciting expert knowledge of plant disease epidemics and their application in cluster sampling for disease incidence.” *Crop Protection*, 21(3): 203–215. 9, 15
- Ibrahim, J. G. and Chen, M.-H. (2000). “Power Prior Distributions for Regression Models.” *Statistical Science*, 15(1): 46–60.
URL <http://www.jstor.org/stable/2676676> 20
- Ibrahim, J. G., Chen, M.-H., Gwon, Y., and Chen, F. (2015). “The power prior: theory and applications.” *Statistics in Medicine*, 34(28): 3724–3749.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.6728> 20
- Ibrahim, J. G., Chen, M.-H., and MacEachern, S. N. (1999). “Bayesian variable selection for proportional hazards models.” *Canadian Journal of Statistics*, 27(4): 701–717.
URL <https://onlinelibrary.wiley.com/doi/abs/10.2307/3316126> 24
- Ibrahim, J. G. and Laud, P. W. (1994). “A predictive approach to the analysis of designed experiments.” *Journal of the American Statistical Association*, 89(425): 309–319. 19, 24
- Jacobi, L., Joshi, M. S., and Zhu, D. (2018). “Automated sensitivity analysis for Bayesian inference via Markov chain Monte Carlo: Applications to Gibbs sampling.” *Available at SSRN 2984054*. 30
- James, A., Choy, S. L., and Mengersen, K. (2010). “Elicitor: an expert elicitation tool for regression in ecology.” *Environmental Modelling & Software*, 25(1): 129–145. 19
- Jarociński, M. and Marcet, A. (2019). “Priors about observables in vector autoregressions.” *Journal of Econometrics*, 209(2): 238 – 255.
URL <http://www.sciencedirect.com/science/article/pii/S0304407619300065> 20
- Jaynes, E. T. (1957). “Information theory and statistical mechanics.” *Physical review*, 106(4): 620. 21
- Jenkinson, D. (2005). “The Elicitation of Probabilities - A Review of the Statistical Literature.” 11, 15, 25
- Jennings, D. L., Amabile, T. M., and Ross, L. (1982). *Informal covariation assessment: Data-based versus theory-based judgments*, 211–230. Cambridge University Press. 14
- Johnson, S. R., Tomlinson, G. A., Hawker, G. A., Granton, J. T., and Feldman, B. M. (2010a). “Methods to elicit beliefs for Bayesian priors: a systematic review.” *Journal of clinical epidemiology*, 63(4): 355–369. 11, 32
- Johnson, S. R., Tomlinson, G. A., Hawker, G. A., Granton, J. T., Grosbein, H. A., and Feldman, B. M. (2010b). “A valid and reliable belief elicitation method for Bayesian priors.” *Journal of Clinical Epidemiology*, 63(4): 370–383.
URL <https://www.sciencedirect.com/science/article/pii/S0895435609002339> 14, 32
- Jones, G. and Johnson, W. O. (2014). “Prior elicitation: Interactive spreadsheet graphics with sliders can be fun, and informative.” *The American Statistician*, 68(1): 42–51. 10

- Joseph, L., Gyorkos, T. W., and Coupal, L. (1995). “Bayesian estimation of disease prevalence and the parameters of diagnostic tests in the absence of a gold standard.” *American journal of epidemiology*, 141(3): 263–272. 16
- Kadane, J. and Wolfson, L. J. (1998). “Experiences in elicitation.” *Journal of the Royal Statistical Society: Series D (The Statistician)*, 47(1): 3–19. 9, 13, 18
- Kadane, J. B. (1980). “Predictive and structural methods for eliciting prior distributions.” *Bayesian Analysis in Econometrics and Statistics*. 18
- Kadane, J. B., Chan, N. H., and Wolfson, L. J. (1996). “Priors for unit root models.” *Journal of Econometrics*, 75(1): 99–111.
URL <https://www.sciencedirect.com/science/article/pii/0304407695017712> 24
- Kadane, J. B., Dickey, J. M., Winkler, R. L., Smith, W. S., and Peters, S. C. (1980). “Interactive Elicitation of Opinion for a Normal Linear Model.” *Journal of the American Statistical Association*, 75: 845–854. 10, 18, 19, 23, 24
- Kadane, J. B. and Winkler, R. L. (1988). “Separating Probability Elicitation from Utilities.” *Journal of the American Statistical Association*, 83(402): 357–363. 17, 19
- Kahle, D., Stamey, J., Natanegara, F., Price, K., and Han, B. (2016). “Facilitated prior elicitation with the wolfram CDF.” *Biometrics & Biostatistics International Journal*, 3. 5
- Kallen, M. and Kuniewski, S. (2009). “On some elicitation procedures for distributions with bounded support—with applications in PERT.” *Risk and Decision Analysis in Maintenance Optimization and Flood Management*, 21. 11
- Kallioinen, N., Paananen, T., Bürkner, P.-C., and Vehtari, A. (2021). “Detecting and diagnosing prior and likelihood sensitivity with power-scaling.” *arXiv preprint arXiv:2107.14054*. 30, 33
- Kane, D. M., Lovett, S., Moran, S., and Zhang, J. (2017). “Active classification with comparison queries.” In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, 355–366. IEEE. 25
- Karni, E. (2009). “A mechanism for eliciting probabilities.” *Econometrica*, 77(2): 603–606. 17
- Kass, R. E. and Wasserman, L. (1996). “The Selection of Prior Distributions by Formal Rules.” *Journal of the American Statistical Association*, 91(435): 1343–1370.
URL <http://www.jstor.org/stable/2291752> 4
- Kay, M. (2021). *ggdist: Visualizations of Distributions and Uncertainty*. R package version 2.4.1.
URL <https://mjskay.github.io/ggdist/> 34
- Kay, M., Kola, T., Hullman, J. R., and Munson, S. A. (2016). “When (ish) is my bus? user-centered visualizations of uncertainty in everyday, mobile predictive systems.” In *Proceedings of the 2016 chi conference on human factors in computing systems*, 5092–5103. 34

- Keelin, T. W. (2016). “The Metalog Distributions.” *Decision Analysis*, 13(4): 243–277. [29](#)
- Keeney, R. L. and Raiffa (1976). *Decisions with multiple objectives: preferences and value trade-offs*. New York: John Wiley. [25](#)
- Kennedy, M., Anderson, C., O’Hagan, A., Lomas, M., Woodward, I., Gosling, J. P., and Heinemeyer, A. (2008). “Quantifying Uncertainty in the Biospheric Carbon Flux for England and Wales.” *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 171(1): 109–135.
URL <http://www.jstor.org/stable/30130733> [5](#)
- Kumar, R., Carroll, C., Hartikainen, A., and Martin, O. (2019). “ArviZ a unified library for exploratory analysis of Bayesian models in Python.” *Journal of Open Source Software*, 4(33): 1143.
URL <https://doi.org/10.21105/joss.01143> [34](#)
- Kurowicka, D. and Cooke, R. M. (2006). *Uncertainty analysis with high dimensional dependence modelling*. John Wiley & Sons. [14](#), [15](#)
- Kynn, M. (2005). “Eliciting expert knowledge for Bayesian logistic regression in species habitat modelling.” Ph.D. thesis, Queensland University of Technology. [24](#)
- (2008). “The ”heuristics and biases” bias in expert elicitation.” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 171: 239–264. [26](#)
- Lahiri, K., Teigland, C., and Zaporowski, M. (1988). “Interest rates and the subjective probability distribution of inflation forecasts.” *Journal of Money, Credit and Banking*, 20(2): 233–248. [14](#)
- Lambert, N. S., Pennock, D. M., and Shoham, Y. (2008). “Eliciting Properties of Probability Distributions.” In *Proceedings of the 9th ACM Conference on Electronic Commerce, EC ’08*, 129–138. New York, NY, USA: Association for Computing Machinery.
URL <https://doi.org/10.1145/1386790.1386813> [18](#)
- Laud, P. W. and Ibrahim, J. G. (1995). “Predictive Model Selection.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1): 247–262.
URL <http://www.jstor.org/stable/2346098> [19](#), [24](#)
- Leamer, E. E. (1992). “Bayesian Elicitation Diagnostics.” *Econometrica*, 60(4): 919–942.
URL <http://www.jstor.org/stable/2951572> [24](#)
- LeCun, Y., Cortes, C., and Burges, C. (2010). “MNIST handwritten digit database.” [32](#)
- Lee, M. D. and Vanpaemel, W. (2018). “Determining informative priors for cognitive models.” *Psychonomic Bulletin & Review*, 25(1): 114–127. [4](#)
- Legedza, A. T. R. and Ibrahim, J. G. (2001). “Heterogeneity in phase I clinical trials: prior elicitation and computation using the continual reassessment method.” *Statistics in Medicine*, 20(6): 867–882.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.701> [11](#)

- León, C. J., Vázquez-Polo, F. J., and González, R. L. (2003). “Elicitation of expert opinion in benefit transfer of environmental goods.” *Environmental and Resource Economics*, 26(2): 199–210. 11
- León-Villagrà, P., Otsubo, K., Lucas, C., and Buchsbaum, D. (2020). “Uncovering Category Representations with Linked MCMC with People.” In *CogSci*. 33
- Lindley, D. V. (1985). “Reconciliation of discrete probability distributions.” *Bayesian statistics*, 2(375-390): 375–390. 25
- (1987). “The Probability Approach to the Treatment of Uncertainty in Artificial Intelligence and Expert Systems.” *Statistical Science*, 2(1): 17–24.
URL <http://www.jstor.org/stable/2245603> 17
- Lindley, D. V., Tversky, A., and Brown, R. V. (1979). “On the Reconciliation of Probability Assessments.” *Journal of the Royal Statistical Society. Series A (General)*, 142(2): 146–180. 10
- Lins, G. and Souza, F. (2001). “A Protocol for the Elicitation of Prior Distributions.” 265–272. 22
- Lombardi, M. J. and Nicoletti, G. (2012). “Bayesian prior elicitation in DSGE models: Macro- vs micropriors.” *Journal of Economic Dynamics and Control*, 36(2): 294 – 313. 11
- Lopes, H. F. and Tobias, J. L. (2011). “Confronting prior convictions: On issues of prior sensitivity and likelihood robustness in Bayesian analysis.” *Annu. Rev. Econ.*, 3(1): 107–131. 30
- Matheson, J. E. and Winkler, R. L. (1973). “The Elicitation of Continuous Probability Distributions.” 17
- (1976). “Scoring Rules for Continuous Probability Distributions.” *Management Science*, 22(10): 1087–1096.
URL <http://www.jstor.org/stable/2629907> 17
- Meeden, G. (1992). “An elicitation procedure using piecewise conjugate priors.” In *Bayesian Analysis in Statistics and Econometrics*, 195–206. Springer. 21
- Micallef, L., Sundin, I., Marttinen, P., Ammad-ud din, M., Peltola, T., Soare, M., Jacucci, G., and Kaski, S. (2017). “Interactive Elicitation of Knowledge on Feature Relevance Improves Predictions in Small Data Sets.” In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*, IUI ’17, 547–552. New York, NY, USA: Association for Computing Machinery.
URL <https://doi.org/10.1145/3025171.3025181> 10, 23, 24, 25
- Miller III, A. C. and Rice, T. R. (1983). “Discrete approximations of probability distributions.” *Management science*, 29(3): 352–362. 33
- Ming, Z., Ling, X., Bai, X., and Zong, B. (2016). “The Bayesian Reliability Assessment and Prediction for Radar System Based on New Dirichlet Prior Distribution.” *Journal of Physics: Conference Series*, 679: 012039.
URL <https://doi.org/10.1088/1742-6596/679/1/012039> 11

- Moala, F. and O’Hagan, A. (2010). “Elicitation of multivariate prior distributions: A nonparametric Bayesian approach.” *Journal of Statistical Planning and Inference*, 140: 1635–1655. [10](#), [22](#)
- Montague, T. H., Price, K. L., and Seaman, J. W. (2019). “The Practice of Prior Elicitation.” In *Bayesian Applications in Pharmaceutical Development*, 61–92. Chapman and Hall/CRC. [11](#)
- Moreno, E. and Giròn, J. (1998). “Estimating with incomplete count data A Bayesian approach.” *Journal of Statistical Planning and Inference*, 66(1): 147–159.
URL <https://www.sciencedirect.com/science/article/pii/S0378375897000736> [24](#)
- Morris, D. E., Oakley, J. E., and Crowe, J. A. (2014). “A web-based tool for eliciting probability distributions from experts.” *Environmental Modelling & Software*, 52: 1–4.
URL <https://www.sciencedirect.com/science/article/pii/S1364815213002533> [14](#)
- Morris, P. A. (1974). “Decision analysis expert use.” *Management Science*, 20(9): 1233–1241. [25](#)
- (1977). “Combining Expert Judgments: A Bayesian Approach.” *Management Science*, 23(7): 679–693. [25](#)
- Murphy, A. H. and Winkler, R. L. (1970). “Scoring rules in probability assessment and evaluation.” *Acta psychologica*, 34: 273–286. [17](#)
- (1974). “Credible Interval Temperature Forecasting: Some Experimental Results.” *Monthly Weather Review*, 102(11): 784 – 794.
URL https://journals.ametsoc.org/view/journals/mwre/102/11/1520-0493_1974_102_0784_citfse_2_0_co_2.xml [13](#)
- Oakley, J. E. (2010). “Eliciting univariate probability distributions.” *Rethinking risk measurement and reporting*, 1: 155–177. [12](#), [13](#), [14](#)
- Oakley, J. E., Daneshkhah, A., and O’Hagan, A. (2010). “Nonparametric prior elicitation using the Roulette method.” Technical report, School of Mathematics and Statistics, University of Sheffield, UK. [10](#), [22](#)
- Oakley, J. E. and O’Hagan, A. (2007). “Uncertainty in Prior Elicitations: A Nonparametric Approach.” *Biometrika*, 94. [5](#), [8](#), [9](#), [10](#), [21](#), [22](#), [23](#)
- (2019). “SHELF: The Sheffield Elicitation Framework (Version 4.0). School of Mathematics and Statistics, University of Sheffield, UK (<http://tonyohagan.co.uk/shelf>).” [2](#), [10](#), [11](#), [12](#), [14](#), [23](#), [25](#)
- O’Hagan, A. (1988). *Probability: methods and measurement*. London: Chapman and Hall. [13](#)
- (1995). “HMIP assessment of Nirex proposals performance assessment project (phase 2): elicitation of hydrogeological data. Contractor report to Her Majesty’s Inspectorate of Pollution. Report TR-Z2-17.” Technical report, Radioactive Waste Management. [14](#)

- (1998). “Eliciting expert beliefs in substantial practical applications.” *Journal of the Royal Statistical Society: Series D (The Statistician)*, 47(1): 21–35. [14](#), [15](#)
- (2019). “Expert Knowledge Elicitation: Subjective but Scientific.” *The American Statistician*, 73(sup1): 69–81. [2](#), [5](#), [10](#), [11](#), [13](#), [25](#), [26](#)
- O’Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J., Oakley, J. E., and Rakow, T. (2006). *Uncertain Judgements: Eliciting Experts’ Probabilities*. John Wiley & Sons, Ltd. [2](#), [3](#), [4](#), [5](#), [8](#), [9](#), [10](#), [11](#), [13](#), [21](#), [30](#)
- O’Hagan, A. and Oakley, J. E. (2004). “Probability is perfect, but we can’t elicit it perfectly.” *Reliability Engineering & System Safety*, 85: 239–248. [26](#)
- O’Leary, R. A., Choy, S. L., Murray, J. V., Kynn, M., Denham, R., Martin, T. G., and Mengersen, K. (2009). “Comparison of three expert elicitation methods for logistic regression on predicting the presence of the threatened brush-tailed rock-wallaby *Petrogale penicillata*.” *Environmetrics*, 20(4): 379–398.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/env.935> [24](#)
- Oman, S. D. (1985). “Specifying a Prior Distribution in Structured Regression Problems.” *Journal of the American Statistical Association*, 80(389): 190–195. [19](#), [24](#)
- Önkal, D. and Muradoglu, G. (1996). “Effects of task format on probabilistic forecasting of stock prices.” *International Journal of Forecasting*, 12(1): 9–24. [14](#)
- Osband, K. and Reichelstein, S. (1985). “Information-eliciting compensation schemes.” *Journal of Public Economics*, 27(1): 107–115.
URL <https://EconPapers.repec.org/RePEc:eee:pubeco:v:27:y:1985:i:1:p:107-115> [18](#)
- Ouchi, F. (2004). “A literature review on the use of expert opinion in probabilistic risk analysis.” [25](#)
- O’Leary, R., Mengersen, K., and Low Choy, S. (2008). “A mixture model approach to representing simple expert information as priors in logistic regression.” Technical report, Technical report, School of Mathematical Sciences, Queensland University. [24](#)
- Parmar, M. K., Spiegelhalter, D. J., Freedman, L. S., and Committee, C. S. (1994). “The CHART trials: Bayesian design and monitoring in practice.” *Statistics in medicine*, 13(13-14): 1297–1312. [33](#)
- Pattillo, C. (1998). “Investment, uncertainty, and irreversibility in Ghana.” *Staff Papers*, 45(3): 522–553. [14](#)
- Percy, D. F. (2002). “Bayesian enhanced strategic decision making for reliability.” *European Journal of Operational Research*, 139(1): 133–145. [20](#)
- (2003). “Subjective reliability analysis using predictive elicitation.” In *Mathematical and Statistical Methods in reliability*, 57–72. World Scientific. [20](#)
- (2004). “Subjective priors for maintenance models.” *Journal of quality in maintenance engineering*. [20](#)

- Perepolkin, D., Goodrich, B., and Sahlin, U. (2021). “Hybrid elicitation and indirect Bayesian inference with quantile-parametrized likelihood.”
URL osf.io/paby6 13, 19, 29
- Pérez, C., Martín, J., and Rufo, M. J. (2006). “MCMC-based local parametric sensitivity estimations.” *Computational Statistics & Data Analysis*, 51(2): 823–835. 30
- Peterson, C. and Beach, L. R. (1967). “Man as an intuitive statistician.” *Psychological bulletin*, 68(1): 29. 13
- Peterson, C. and Miller, A. (1964). “Mode, median, and mean as optimal strategies.” *Journal of Experimental Psychology*, 68: 363–367. 13
- Pham-Gia, T., Turkkan, N., and Duong, Q. (1992). “Using the mean deviation in the elicitation of the prior distribution.” *Statistics & Probability Letters*, 13(5): 373 – 381.
URL <http://www.sciencedirect.com/science/article/pii/016771529290110Q>
15, 16
- Phillips, L. D. (1999). *Group Elicitation of Probability Distributions: Are Many Heads Better Than One?*, 313–330. Boston, MA: Springer US.
URL https://doi.org/10.1007/978-1-4615-5089-1_17 25
- Phillips, L. D. and Wisbey, S. (1993). “The elicitation of judgmental probability distributions from groups of experts: a description of the methodology and records of seven formal elicitation sessions held in 1991 and 1992.” Technical report. 14, 25
- Piironen, J., Vehtari, A., et al. (2017). “Sparsity information and regularization in the horseshoe and other shrinkage priors.” *Electronic Journal of Statistics*, 11(2): 5018–5051. 32
- Press, S. J. (2002). *Prior Distributions*, chapter 5, 73–116. John Wiley & Sons, Ltd.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470317105.ch5>
21
- Qu, X. (2012). “A mechanism for eliciting a probability distribution.” *Economics Letters*, 115(3): 399–400.
URL <https://www.sciencedirect.com/science/article/pii/S0165176511006203> 17
- Raiffa, H. (1968). *Decision analysis : introductory lectures on choices under uncertainty..* New York: Random House. 13
- Regazzini, E. and Sazonov, V. V. (1999). “Approximation of Laws of Multinomial Parameters by Mixtures of Dirichlet Distributions with Applications to Bayesian Inference.” *Acta Applicandae Mathematica*, 58: 247–264. 16, 24
- Reimherr, M., Meng, X.-L., and Nicolae, D. L. (2021). “Prior sample size extensions for assessing prior impact and prior-likelihood discordance.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 30
- Rizzo, D. B. and Blackburn, M. R. (2019). “Harnessing Expert Knowledge: Defining Bayesian Network Model Priors From Expert Knowledge Only—Prior Elicitation for the Vibration Qualification Problem.” *IEEE Systems Journal*, 13(2): 1895–1905. 24

- Roback, P. J. and Givens, G. H. (2001). “SUPRA-BAYESIAN POOLING OF PRIORS LINKED BY A DETERMINISTIC SIMULATION MODEL.” *Communications in Statistics - Simulation and Computation*, 30(3): 447–476. [25](#)
- Robert, C. P. (2007). *From Prior Information to Prior Distributions*, 105–163. New York, NY: Springer New York.
URL https://doi.org/10.1007/0-387-71599-1_3 [5](#), [21](#)
- Roos, M., Martins, T. G., Held, L., and Rue, H. (2015). “Sensitivity analysis for Bayesian hierarchical models.” *Bayesian Analysis*, 10(2): 321–349. [30](#)
- Roosen, J. and Hennessy, D. A. (2001). “Capturing experts’ uncertainty in welfare analysis: An application to organophosphate use regulation in US apple production.” *American Journal of Agricultural Economics*, 83(1): 166–182. [14](#)
- Ryan, E. G., Drovandi, C. C., McGree, J. M., and Pettitt, A. N. (2016). “A review of modern computational algorithms for Bayesian optimal design.” *International Statistical Review*, 84(1): 128–154. [23](#)
- Salvatier, J., Wiecki, T. V., and Fonnesbeck, C. (2016). “Probabilistic programming in Python using PyMC3.” *PeerJ Computer Science*, 2: e55. [2](#), [5](#)
- Sanborn, A. and Griffiths, T. L. (2008). “Markov chain Monte Carlo with people.” In *Advances in neural information processing systems*, 1265–1272. [33](#)
- Sanborn, A., Griffiths, T. L., and Shiffrin, R. M. (2010). “Uncovering mental representations with Markov chain Monte Carlo.” *Cognitive psychology*, 60(2): 63–106. [33](#)
- Sarma, A. and Kay, M. (2020). “Prior Setting in Practice: Strategies and Rationales Used in Choosing Prior Distributions for Bayesian Analysis.” In *Conference on Human Factors in Computing Systems*, 1–12. [2](#), [34](#)
- Savage, L. J. (1971). “Elicitation of Personal Probabilities and Expectations.” *Journal of the American Statistical Association*, 66(336): 783–801. [17](#)
- Schervish, M. J., Kadane, J. B., and Seidenfeld, T. (2012). “Characterization of proper and strictly proper scoring rules for quantiles.” *Preprint, Carnegie Mellon University, March*. [18](#)
- Schlaifer, R. (1969). *Analysis of Decisions Under Uncertainty*. McGraw-Hill. [22](#)
- Schlaifer, R. and Raiffa, H. (1961). *Applied statistical decision theory*. [19](#)
- Seeger, M. W. (2008). “Bayesian Inference and Optimal Design for the Sparse Linear Model.” *Journal of Machine Learning Research*, 9(26): 759–813.
URL <http://jmlr.org/papers/v9/seeger08a.html> [23](#)
- Settles, B. (2012). *Active learning*. Morgan & Claypool. [23](#)
- Siivola, E., Weber, S., and Vehtari, A. (2021). “Qualifying drug dosing regimens in pediatrics using Gaussian processes.” *Statistics in Medicine*, 40(10): 2355–2372.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.8907> [30](#)

- Simpson, D., Rue, H., Riebler, A., Martins, T. G., and Sørbye, S. H. (2017). “Penalising Model Component Complexity: A Principled, Practical Approach to Constructing Priors.” *Statistical Science*, 32(1): 1 – 28.
URL <https://doi.org/10.1214/16-STS576> 2, 23
- Skene, A., Shaw, J., and Lee, T. (1986). “Bayesian modelling and sensitivity analysis.” *Journal of the Royal Statistical Society: Series D (The Statistician)*, 35(2): 281–288. 30
- Sklar, M. (1959). “Fonctions de repartition an dimensions et leurs marges.” *Publ. inst. statist. univ. Paris*, 8: 229–231. 15
- Smid, S. C. and Winter, S. D. (2020). “Dangers of the defaults: A tutorial on the impact of default priors when using Bayesian SEM with small samples.” *Frontiers in Psychology*, 11: 3536. 4, 29
- Smith, J. and Mandac, A. M. (1995). “Subjective versus objective yield distributions as measures of production risk.” *American Journal of Agricultural Economics*, 77(1): 152–161. 14
- Soare, M., Ammad-Ud-Din, M., and Kaski, S. (2016). “Regression with $n \rightarrow 1$ by Expert Knowledge Elicitation.” In *2016 15th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 734–739. 24
- Spetzler, C. S. and Holstein, C.-A. S. S. V. (1975). “Probability Encoding in Decision Analysis.” *Management Science*, 22(3): 340–358. 11, 13
- Srivastava, R., Upadhyay, S. K., and Shukla, V. K. (2019). “Subjective elicitation of hyperparameters of a conjugate Dirichlet prior and the corresponding Bayes analysis.” *Communications in Statistics - Theory and Methods*, 48(19): 4874–4887.
URL <https://doi.org/10.1080/03610926.2018.1500600> 16, 24
- Stan Development Team (2020). “RStan: the R interface to Stan.” R package version 2.21.2.
URL <http://mc-stan.org/> 29
- (2021). “Stan Modeling Language Users Guide and Reference Manual, Version 2.28.”
URL <http://mc-stan.org/> 2, 5
- Stefan, A., Evans, N., and Wagenmakers, E.-J. (2020). “Practical challenges and methodological flexibility in prior elicitation.” *Psychol Methods*. 9, 11, 19
- Steffey, D. (1992). “Hierarchical bayesian modeling with elicited prior information.” *Communications in Statistics - Theory and Methods*, 21(3): 799–821. 24
- Steinwart, I., Pasin, C., Williamson, R., and Zhang, S. (2014). “Elicitation and identification of properties.” In *Conference on Learning Theory*, 482–526. PMLR. 18
- Stigler, S. M. (1982). “Thomas Bayes’s Bayesian Inference.” *Journal of the Royal Statistical Society. Series A (General)*, 145(2): 250–258.
URL <http://www.jstor.org/stable/2981538> 18
- Studer, R., Benjamins, V., and Fensel, D. (1998). “Knowledge engineering: Principles

- and methods.” *Data & Knowledge Engineering*, 25(1): 161–197.
 URL <https://www.sciencedirect.com/science/article/pii/S0169023X97000566> 5
- Sundin, I., Peltola, T., Micallef, L., Afrabandpey, H., Soare, M., Mamun Majumder, M., Daeae, P., He, C., Serim, B., Havulinna, A., Heckman, C., Jacucci, G., Marttinen, P., and Kaski, S. (2018). “Improving genomics-based predictions for precision medicine through active elicitation of expert knowledge.” *Bioinformatics*, 34(13): i395–i403.
 URL <https://doi.org/10.1093/bioinformatics/bty257> 23
- Sørbye, S. H. and Rue, H. v. (2017). “Penalised Complexity Priors for Stationary Autoregressive Processes.” *Journal of Time Series Analysis*, 38(6): 923–935.
 URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/jtsa.12242> 2
- Tan, S.-B., Chung, Y.-F. A., Tai, B.-C., Cheung, Y.-B., and Machin, D. (2003). “Elicitation of prior distributions for a phase III randomized controlled trial of adjuvant therapy with surgery for hepatocellular carcinoma.” *Controlled clinical trials*, 24(2): 110–121. 33
- Thomas, O., Pesonen, H., and Corander, J. (2020). “Probabilistic elicitation of expert knowledge through assessment of computer simulations.” *arXiv preprint arXiv:2002.10902*. 21
- Tiao, G. C. and Zellner, A. (1964). “Bayes’s theorem and the use of prior knowledge in regression analysis.” *Biometrika*, 51(1/2): 219–230. 19
- Tompsett, D., Sutton, S., Seaman, S. R., and White, I. R. (2020). “A general method for elicitation, imputation, and sensitivity analysis for incomplete repeated binary data.” *Statistics in Medicine*, 39(22): 2921–2935. 11
- Tsutakawa, R. K. (1984). “Estimation of two-parameter logistic item response curves.” *Journal of Educational Statistics*, 9(4): 263–276. 20
- Tversky, A. and Kahneman, D. (1974). “Judgement under Uncertainty: Heuristics and Biases.” *Science*, 185: 1124–1131. 11, 26
- van Dongen, S. (2006). “Prior specification in Bayesian statistics: Three cautionary tales.” *Journal of Theoretical Biology*, 242(1): 90–100.
 URL <https://www.sciencedirect.com/science/article/pii/S0022519306000609> 4, 29
- van Dorp, J. and Mazzuchi, T. (2004). *Parameter specification of the Beta distribution and its Dirichlet extensions utilizing quantiles*, 283–318. Marcel Dekker, Inc. 16
- van Lenthe, J. (1993a). “ELI: An Interactive Elicitation Technique for Subjective Probability Distributions.” *Organizational Behavior and Human Decision Processes*, 55(3): 379–413.
 URL <https://www.sciencedirect.com/science/article/pii/S074959788371037X> 17
- (1993b). “A blueprint of ELI: a new method for eliciting subjective probability

- distributions.” *Behavior Research Methods, Instruments, & Computers*, 25(4): 425–433. [17](#)
- (1994). “Scoring-Rule Feedforward and the Elicitation of Subjective Probability Distributions.” *Organizational Behavior and Human Decision Processes*, 59(2): 188–209. [17](#)
- von Holstein, C.-A. S. S. (1970). “Measurement of subjective probability.” *Acta Psychologica*, 34: 146–159. [17](#)
- (1971). “The effect of learning on the assessment of subjective probability distributions.” *Organizational behavior and human performance*, 6(3): 304–315. [15](#)
- von Neumann, J., Morgenstern, O., and Rubinstein, A. (1944). *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press.
URL <http://www.jstor.org/stable/j.ctt1r2gkx> [17](#)
- Wang, Z. and Zhou, H. (2009). “A general method of prior elicitation in Bayesian reliability analysis.” In *2009 8th International Conference on Reliability, Maintainability and Safety*. IEEE. [20](#)
- Weiler, H. (1965). “The Use of Incomplete Beta Functions for Prior Distributions in Binomial Sampling.” *Technometrics*, 7(3): 335–347. [15](#)
- Werner, C., Bedford, T., Cooke, R. M., Hanea, A. M., and Morales-Nápoles, O. (2017). “Expert judgement for dependence in probabilistic modelling: A systematic literature review and future research directions.” *European Journal of Operational Research*, 258(3): 801–819. [14](#)
- Wesner, J. S. and Pomeranz, J. P. F. (2021). “Choosing priors in Bayesian ecological models by simulating from the prior predictive distribution.” *Ecosphere*, 12(9): e03739.
URL <https://esajournals.onlinelibrary.wiley.com/doi/abs/10.1002/ecs2.3739> [21](#)
- West, M. (1988). “Combining Expert Opinion.” *Bayesian Statistics*, 375–390. [25](#)
- Williams, C. J., Wilson, K. J., and Wilson, N. (2021). “A comparison of prior elicitation aggregation using the classical method and SHELF.” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, n/a(n/a).
URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/rssa.12691> [25](#)
- Williams, C. K. and Rasmussen, C. E. (2006). *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA. [22](#)
- Wilson, A. G. (1994). *Cognitive factors affecting subjective probability assessment*. Institute of Statistics and Decision Sciences, Duke University. [14](#)
- Wilson, J. R. and Corlett, N. (2005). *Knowledge Elicitation: Methods, Tools and Techniques*. CRC press. [5](#)
- Wilson, K. J. (2018). “Specification of Informative Prior Distributions for Multinomial

- Models Using Vine Copulas.” *Bayesian Analysis*, 13(3): 749 – 766.
URL <https://doi.org/10.1214/17-BA1068> 16, 24
- Wilson, K. J., Elfadaly, F. G., Garthwaite, P. H., Oakley, J. E., et al. (2021). “Recent Advances in the Elicitation of Uncertainty Distributions from Experts for Multinomial Probabilities.” *Expert Judgement in Risk and Decision Analysis*, 19–51. 16
- Winkler, R. (1986). “On ”good probability appraisers”.” 18
- Winkler, R. L. (1967). “The assessment of prior distributions in Bayesian analysis.” *Journal of the American Statistical Association*, 62(319): 776–800. 2, 9, 11, 15, 16, 21, 32
- (1968). “The Consensus of Subjective Probability Distributions.” *Management Science*, 15(2): B61–B75.
URL <http://www.jstor.org/stable/2628853> 25
- (1980). “Prior information, predictive distributions, and Bayesian model-building.” 18, 20
- Winkler, R. L., Munoz, J., Cervera, J. L., Bernardo, J. M., Blattenberger, G., Kadane, J. B., Lindley, D. V., Murphy, A. H., Oliver, R. M., and Ríos-Insua, D. (1996). “Scoring rules and the evaluation of probabilities.” *Test*, 5(1): 1–60. 17
- Wolfson, L. (2001). “Elicitation of Probabilities and Probability Distributions.” In Smelser, N. J. and Baltes, P. B. (eds.), *International Encyclopedia of the Social & Behavioral Sciences*, 4413–4417. Oxford: Pergamon. 18
- Wolfson, L. J. (1995). “Elicitation of priors and utilities for Bayesian analysis.” Ph.D. thesis, Carnegie Mellon University. 11
- Wolfson, L. J., Kadane, J. B., and Small, M. J. (1996). “Bayesian Environmental Policy Decisions: Two Case Studies.” *Ecological Applications*, 6(4): 1056–1066.
URL <http://www.jstor.org/stable/2269590> 11
- Wolpert, R., Ickstadt, K., and Hansen, M. B. (2003). “A Nonparametric Bayesian Approach to Inverse Problems.” 20
- Wolpert, R. L. (1989). “Eliciting and combining subjective judgments about uncertainty.” *International Journal of Technology Assessment in Health Care*, 5(4): 537–557. 11
- Zapata-Vázquez, R. E., O’Hagan, A., and Bastos, L. S. (2014). “Eliciting expert judgements about a set of proportions.” *Journal of Applied Statistics*, 41(9): 1919–1933.
URL <https://doi.org/10.1080/02664763.2014.898131> 16
- Zellner, A. (1977). “Maximal data information prior distributions.” *Aykae A, Brumat C (eds) New developments in the applications of Bayesian methods*, Amsterdam. 21
- (1991). *Bayesian Methods and Entropy in Economics and Econometrics*, 17–31. Dordrecht: Springer Netherlands.
URL https://doi.org/10.1007/978-94-011-3460-6_2 21

— (1996). “Models, prior information, and Bayesian analysis.” *Journal of Econometrics*, 75(1): 51–68. [21](#)

Zhang, Y. D., Naughton, B. P., Bondell, H. D., and Reich, B. J. (2020). “Bayesian Regression Using a Prior on the Model Fit: The R2-D2 Shrinkage Prior.” *Journal of the American Statistical Association*, 0(0): 1–13.

URL <https://doi.org/10.1080/01621459.2020.1825449> [32](#)

Acknowledgments

This work was supported by the Academy of Finland (Flagship program: Finnish Center for Artificial Intelligence FCAI, Grant 292334), by the Technology Industries of Finland Centennial Foundation, by the Jane and Aatos Erkko Foundation, European Research Council grant 742158 (SCARABEE, Scalable inference algorithms for Bayesian evolutionary epidemiology), and by the UKRI Turing AI World-Leading Researcher Fellowship, EP/W002973/1. Partially funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2075 - 390740016.