Computational Statistics (2006) 21:431–444
DOI 10.1007/s00180-006-0004-1

ORIGINAL PAPER

On the influence of the prior distribution in image reconstruction

Holger Rootzén · Jonny Olsson

Published online: 19 August 2006

Abstract Two measures of the influence of the prior distribution $p(\theta)$ in Bayes estimation are proposed. Both involve comparing with alternative prior distributions proportional to $p(\theta)^s$, for $s \geq 0$. The first one, *the influence curve for the prior distribution*, is simply the curve of parameter values which are obtained as estimates when the estimation is made using $p(\theta)^s$ instead of $p(\theta)$. It measures the overall influence of the prior. The second one, *the influence rate for the prior*, is the derivative of this curve at $s = 1$, and quantifies the sensitivity to small changes or inaccuracies in the prior distribution. We give a simple formula for the influence rate in marginal posterior mean estimation, and discuss how the influence measures may be computed and used in image processing with Markov random field priors. The results are applied to an image reconstruction problem from visual field testing and to a stylized image analysis problem.

Keywords Gibbs distribution · Sensitivity analysis · Glaucoma diagnosis · Visual field test

1 Introduction

In this paper we propose two simple measures of how much a Bayes estimate is determined by the prior distribution and to what degree it is determined by the

Research supported by the Bank of Sweden Tercentenary Foundation.

H. Rootzén (✉)
Department of Mathematical Statistics, Chalmers University of Technology,
412 96 Göteborg, Sweden
e-mail: rootzen@math.chalmers.se

J. Olsson
Statistical Norm, Scheelevägen 17, 223 70 Lund, Sweden

data. We also show how the measures can be used to make simple diagnostic plots.

The methods are designed for situations where: (i) first one fixed prior is determined, from knowledge of the application at hand and perhaps also Bayesian robustness methods (cf. references below), (ii) this fixed prior is then used many times to analyze new data sets, (iii) the problem is high-dimensional, often involving image processing and Markov random field modeling, and (iv) simple evaluation free of choices of parameters is important.

New data sets (e.g., different images) may contain quite varying amounts of information, and the information content may very well be spatially varying. Hence some estimated images may largely be determined by the observations, while other images – or at least parts of images – may mainly reflect the prior opinion and perhaps even be in conflict with the information available. Our aim is to find ways to distinguish between these cases.

We were led to the present research by an important medical problem, glaucoma diagnosis. This application is discussed in Sect. 4. In it the data sets are "images" consisting of results of visual field examinations on glaucoma suspects, and the estimate, i.e., the reconstructed image, is used for diagnosis. This image is modeled as a Markov random field, and one fixed prior has been determined from extensive normal and clinical data sets. This fixed prior is then used to estimate the true visual field for each new patient. The estimation method (see Olsson and Rootzén 1994; Bengtsson et al. 1998) is implemented in a computerized instrument called a field analyzer, and is routinely used by ophthalmologists and opticians in many thousands of clinics around the world. It was deemed useful to know if a diagnosis is based almost exclusively on measurements made on the patient, or if not, the extent to which it was determined by the prior distribution.

Although, this research originated from an effort to improve glaucoma diagnosis, we believe that the methods could be used also in other image reconstruction and high-dimensional problems.

The formal setup is as follows. The parameter of interest θ varies in a subset Θ of R^n, and the prior distribution has a density $p(\theta)$ with respect to a measure μ which is concentrated on Θ. Further, μ is assumed to be a "flat" prior which represents ignorance of the actual value of the parameter. We also assume that some observation x, e.g., an observed image, related to the parameter is available, and that the conditional likelihood of x given θ is $l(x|\theta)$. The posterior distribution then has the density

$$p(\theta|x) = \frac{p(\theta)l(x|\theta)}{\int p(\theta)l(x|\theta)\mu(\mathrm{d}\theta)}$$

with respect to μ. We consider point estimates, $\hat{g}$, of some parameter function $g(\theta) = (g^{(1)}(\theta), \ldots, g^{(k)}(\theta)) \in R^k$.

Important examples are marginal posterior mean (MPM) estimates, for which $\hat{g}$ is the expected value of $g(\theta)$ under the distribution $p(\theta|x)$, and

maximum aposteriori (MAP) estimates, which select a value θ which maximizes $p(\theta|x)$ and then uses $\hat{g} = g(\hat{\theta})$.

The literature on Bayesian robustness, e.g., Berger (1990), Bose (1994), Delampady and Dey (1994), O'Hagan (1994) and Ruggeri and Wasserman (1993), and the references therein, consider the effect of varying the prior distribution in some neighborhood of the $p(\theta)$ of interest. Problems studied include finding a most sensitive direction and deriving bounds on the resulting variation of the estimates. Typically prior distributions contain "unknown" parameters whose values are chosen using prior experience, perhaps including formal estimation procedures from sets of training data. Aims in the cited literature include providing help with these choices of parameter values, and warnings of situations where small parameter changes may give widely varying results. An overview of the literature is Insua and Ruggeri (2000).

To summarize, our aim is to suggest measures of the influence of the prior which:

1. Show if the estimate is determined mainly by the prior or by the data,
2. Are simple, and low-dimensional, enough to be displayed and understood in high-dimensional, e.g., image analysis, problems,
3. Are computationally feasible in high-dimensional image analysis applications.

A method for calculating the influence for each one of the prior, the likelihood function and the data were provided by Clarke and Gustafson (1998). Carlin et al. (1995) partitioned priors into those leading to accepting or rejecting an hypothesis. Bayesian robustness methods discussed in the literature consider variations in the full neighborhood of the parameter values actually used. Image processing applications often use many parameters, e.g., in our Example 2 we have 148 parameters in the prior distribution. It is simply computationally impossible to investigate how variations of all of those affect the result and, in fact, none of the methods suggested in the literature seems practically useful in such situations. Furthermore, even if the computational problems could be overcome, there remain formidable difficulties in displaying and interpreting such high-dimensional information.

In this paper we try to achieve our aims by varying the prior along a one-dimensional curve which has a special significance. Specifically, the influence of the prior distribution will be measured by comparing with alternative priors with density $p_s(\theta)$ proportional to $p(\theta)^s$. In particular, if $\int p(\theta)^s \mu(\mathrm{d}\theta) < \infty$, we may take

$$p_s(\theta) = \frac{p(\theta)^s}{\int p(\theta)^s \mu(\mathrm{d}\theta)},$$

and otherwise p_s is an improper prior. In either case, the posterior density associated with p_s is

$$p_s(\theta|x) = \frac{p(\theta)^s l(x|\theta)}{\int p(\theta)^s l(x|\theta)\mu(\mathrm{d}\theta)}, \tag{1}$$

provided the integral exists.

Now, as discussed above, assume a fixed estimation method, e.g., MPM or MAP has been selected. The method is supposed to be applicable to the posterior distributions $p_s(\theta|x)$ of (1), and then to yield the estimate $\hat{g}_s$ (i.e., $\hat{g}_s$ is calculated from $p_s(\theta|x)$ in the same way as $\hat{g}$ is calculated from $p(\theta|x)$).

The first measure of the influence of the prior, here termed *the influence curve for the prior distribution* is simply the entire curve $\hat{g}_s; s \geq 0$. It measures the overall influence of the prior distribution.

One rationale for using the influence curve is that it interpolates between $s = 0$, which corresponds to a non-informative prior, $s = 1$ which is the initially specified prior distribution, and the case "$s = \infty$" where all prior mass is concentrated at the global maximum (or maxima) of $p(\theta)$. Some comments can be added for the important special case when the prior distribution belongs to an exponential family, with density $p(\theta) = h(\theta)\exp\{z^T\phi - k(\phi)\}$, where $\phi = (\phi_1, \ldots, \phi_m)$ are the canonical parameters and $z = z(\theta) = (z_1(\theta) \ldots z_m(\theta))$ are the canonical statistics, and where $h(\theta)$ and $k(\phi)$ are functions of the indicated arguments. Changing the prior to $p_s(\theta)$ then changes the values of the original canonical parameters by a factor s. If these original canonical parameters are assumed fixed the new canonical parameters takes values along a straight line, $(s\phi_1, \ldots, s\phi_m)$, when s is varied. If instead s would be assumed fixed then we could get back to the original canonical parameters by dividing the new canonical parameters with s. In addition a new canonical statistic $\ln(h(\theta))$ with corresponding canonical parameter is s is introduced if (and only if) $h(\theta)$ depends on θ. If this is so then the dimension, m, of the $p(\theta)$ exponential family increases to $m + 1$ in the $p_s(\theta)$ exponential family.

The second influence measure, called *the influence rate for the prior distribution* is the k-dimensional vector

$$\dot{g}_1 = \frac{\mathrm{d}}{\mathrm{d}s}\hat{g}_s|_{s=1}.$$

Clearly $\dot{g}_1$ is a summary of the behavior of $\hat{g}_s$ close to $s = 1$. It measures the sensitivity of estimates to small changes or inaccuracies in the prior distribution. If the influence rate is large, the estimate is sensitive to small changes in the prior distribution, and should be viewed with some caution.

An alternative that takes random variation into account is the *standardized influence rate*, $\dot{g}_{\mathrm{norm}}$, with elements,

$$\dot{g}_{\mathrm{norm},i} = \frac{\dot{g}_{1,i}}{s_i(\dot{g}_{1,i})} \quad \text{or} \quad |\dot{g}_{\mathrm{norm},i}| = \frac{|\dot{g}_{1,i}|}{s_i(\dot{g}_{1,i})} \tag{2}$$

where $s_i(\dot{g}_{1,i})$ is the standard deviation of the i-th element, $\dot{g}_{1,i}$ of $\dot{g}_1$. The advantages of $\dot{g}_{\text{norm},i}$ and $|\dot{g}_{\text{norm},i}|$ are that they give a measure on a standard scale. However, $\dot{g}_1$ itself is a measure in the relevant physical units.

A further motivation for these influence measures is that they are computable in many complex image processing applications. An advantage of the influence rate over the influence function is that it may be easier to compute and to present and interpret in high-dimensional problems. In image processing $\dot{g}_1$ or $\dot{g}_{\text{norm}}$ provides a diagnostic image. However, as discussed above, the influence curve may also be visualized in such situations, e.g., by making the image reconstruction for each of a discrete grid of s-values, and then presenting the resulting series of pictures.

Obvious further influence measures are the distribution valued curve $p_s(\theta|x)$; $s \geq 0$ or some (functional) derivative of $p_s(\theta|x)$ at $s = 1$. Clearly these contain more information than $\hat{g}_s$ or $\dot{g}_1$. However, high-dimensional distributions are difficult to present and to understand, and typically are hard to compute.

The plan of the paper is as follows. In Sect. 2 we discuss the influence measures for MPM estimation and the results are illustrated in the very simple one-dimensional normal case. In Sect. 3 Markov random field models for the prior are discussed. In Sect. 4 we consider an image reconstruction problem from visual field testing in some detail, and show how the influence measures may be computed and interpreted in this application. In addition, we illustrate the use of the standardized influence rates in a simple synthetic image. Section 5 summarizes our conclusions.

2 Influence measures for MPM estimation

The MPM estimation method uses the estimate $\hat{g} = \int g(\theta)p(\theta|x)\mu(\mathrm{d}\theta)$, and hence

$$\begin{aligned}\hat{g}_s &= \int g(\theta)p_s(\theta|x)\mu(\mathrm{d}\theta) \\ &= \frac{\int g(\theta)p(\theta)^s l(x|\theta)\mu(\mathrm{d}\theta)}{\int p(\theta)^s l(x|\theta)\mu(\mathrm{d}\theta)}.\end{aligned} \tag{3}$$

Throughout this section we will assume that the two integrals in expression (3) converge for $s = 1$ so that $\hat{g} = \hat{g}_1$ is well defined by the last expression in (3), that $p(\theta) > 0$, $\theta \in \Theta$ and that $\mu(\{\theta; l(x|\theta) > 0\}) > 0$ for the observation x under consideration. The influence rate for the prior then has the following simple form.

Theorem 1 (i) *The set of $s \geq 0$ for which $\hat{g}_s$ exists is an interval with end points a, b satisfying $a \leq 1 \leq b$.*

(ii) *$\hat{g}_s$ is infinitely differentiable for $s \in (a, b)$, and the derivatives may be obtained by differentiating under the integral signs.*

(iii) *If $a < 1 < b$ then the i-th component of $\dot{g}_1$ is*

$$\begin{aligned}(\dot{g}_1)_i &= \int g^{(i)}(\theta)\log(p(\theta))p(\theta|x)\mu(\mathrm{d}\theta) \\ &\quad - \int g^{(i)}(\theta)p(\theta|x)\mu(\mathrm{d}\theta)\int \log(p(\theta))p(\theta|x)\mu(\mathrm{d}\theta) \\ &= C(g^{(i)}(\theta), \log(p(\theta))),\end{aligned}$$

where $g^{(i)}(\theta)$ is the i-th component of $g(\theta)$ and where the covariance C is computed with respect to the posterior distribution $p(\theta|x)$.

Proof (i) By the definition (3) $\hat{g}_s$ exists if and only if both $\int p(\theta)^s l(x|\theta)\mu(\mathrm{d}\theta)$ and $\int g(\theta)p(\theta)^s l(x|\theta)\mu(\mathrm{d}\theta)$ exist. Since $p(\theta)^s$ may be written as $\exp\{s\log(p(\theta))\}$ it follows from standard results about one-dimensional exponential families (Lehmann 1959, Sect. 2.7) that the set of $s \geq 0$ for which the integrals exist are intervals. The intersection of two intervals is an interval, and since $\hat{g}_1$ is assumed to be well defined, both intervals contain $s = 1$, and hence the intersection is non-empty with end points a, b satisfying $a \leq 1 \leq b$.

(ii) Again it is a standard result for one-dimensional exponential families that the derivatives of $\int g(\theta)p(\theta)^s l(x|\theta)\mu(\mathrm{d}\theta)$ and of $\int p(\theta)^s l(x|\theta)\mu(\mathrm{d}\theta)$ exist and may be obtained by differentiation under the integral signs. The same result then holds for their ratio $\hat{g}_s$, since $\int p(\theta)^s l(x|\theta)\mu(\mathrm{d}\theta) > 0$.

(iii) The first equality follows from (ii) and a straightforward calculation. The second equality is by definition. □

Example 1 Normal distributions. As a very simple example we compute $\hat{g}_s$ and $\dot{g}_1$ in a one-dimensional case with a normal prior and a normal likelihood. The example is only intended as an introduction and illustration of the concepts and is not representative of the applications we have in mind.

Let $\Theta = R$, assume $\mu(\mathrm{d}\theta) = \mathrm{d}\theta$ and that the prior is normal with mean μ and variance τ^2, i.e., that $p(\theta) = \frac{1}{\sqrt{2\pi\tau^2}}\mathrm{e}^{-\frac{1}{2\tau^2}(\theta-\mu)^2}$. It follows at once that

$$p_s(\theta) = \frac{1}{\sqrt{2\pi\tau^2/s}}\mathrm{e}^{-\frac{1}{2\tau^2/s}(\theta-\mu)^2}.$$

Assume also that the likelihood function is $l(x|\theta) = \frac{1}{\sqrt{2\pi\sigma^2}}\mathrm{e}^{-\frac{1}{2\sigma^2}(x-\theta)^2}$, where x is the (fixed) observation, and $\sigma^2 > 0$ is known. Straightforward calculations (cf. Berger (1980, pp 93–94)) show that $p_s(\theta|x)$ is normal, with mean

$$\nu_s = \frac{\tau^2/s}{\sigma^2 + \tau^2/s}x + \frac{\sigma^2}{\sigma^2 + \tau^2/s}\mu. \tag{4}$$

Hence both the MAP and MPM estimates of $g(\theta) \equiv \theta$ are $\hat{g}_s = \nu_s$, which is a weighted average of the observation x and the prior mean μ, with weights

proportional to τ^2/s and σ^2, respectively. This, of course, explicitly shows how the prior distribution influences the estimate.

Further, by easy computation,

$$\dot{g}_1 = C(g(\theta), \log(p(\theta)) = C\left(\theta, -\frac{1}{2\tau^2}(\theta-\mu)^2\right)$$
$$= \frac{\sigma^2/\tau^2}{(1+\sigma^2/\tau^2)^2}(\mu - x), \tag{5}$$

where the covariance C is computed with respect to the posterior distribution $p(\theta|x)$. (This is equally easy to obtain by differentiation of (4)). The influence curve and influence rate are shown in Fig. 1 for two illustrations.

As an aside, it is easy to see that the influence curve for an i.i.d. sample of size n is obtained by replacing σ^2 by σ^2/n and x by the mean of the observations, $\bar{x}$, in Eq. (5).

In this simple example, the influence measures have the properties one would like them to have. By (5) the influence rate is large, i.e., the estimate is sensitive to small changes in the prior, if μ is far from the observation x, and σ^2 and τ^2 are of comparable sizes. This is precisely as it should be: if x and μ are close, then the amount of confidence in the prior distribution does not make much difference. Further, if $\sigma^2 << \tau^2$ then the estimate is mostly decided by the observation, and if $\sigma^2 >> \tau^2$ it is mainly determined by the prior, and in neither case is the precise amount of confidence in the prior crucial.

The influence curve also contains this information. In addition it makes it possible to differentiate between the last two cases, i.e., when the estimate is mainly decided by the observation and when it is mainly decided by the prior.

3 Markov random field priors

The main goal of this study was to find tools for quantifying the influence of the prior in image processing with Markov random field prior distributions. In this case the prior is a Gibbs distribution of the form

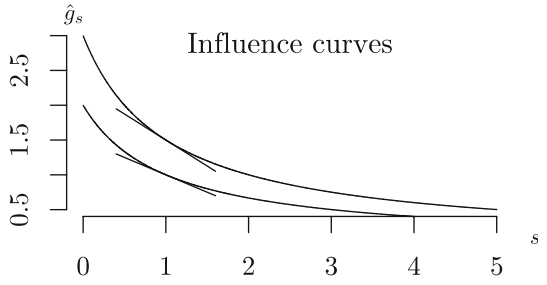

Fig. 1 Two influence curves (4) of Example 1. The range of the s-variable is from 0 to 5. The parameter values were $\sigma = \tau = 1$, $\mu = 0$. The observations are $x = 2$ (*bottom*) and $x = 3$ (*top*), respectively. The numerical values of the influence rates (5), –0.5 and –0.75, are the slopes of the two short tangential lines

$$p(\theta) = \frac{1}{\varphi} \mathrm{e}^{-H(\theta)}, \tag{6}$$

where φ is a normalizing constant, called the partition function, and $H(\theta)$ is the energy function, cf. Kindermann and Snell (1980). It then follows from (3) that

$$\begin{aligned} \hat{g}_s &= \frac{\int g(\theta) \mathrm{e}^{-sH(\theta)} l(x|\theta) \mu(\mathrm{d}\theta)}{\int \mathrm{e}^{-sH(\theta)} l(x|\theta) \mu(\mathrm{d}\theta)} \\ &= \frac{E_{p(\theta|x)} g(\theta) \mathrm{e}^{(1-s)H(\theta)}}{E_{p(\theta|x)} \mathrm{e}^{(1-s)H(\theta)}}, \end{aligned} \tag{7}$$

with $E_{p(\theta|x)}$ denoting expectation under the posterior distribution. Further, since $\log(p(\theta)) = -H(\theta) - \log \varphi$ for the Gibbs distribution, it follows from the theorem that the influence rate is given by

$$(\dot{g}_1) = -C(g^{(i)}(\theta), H(\theta)), \quad i = 1, \ldots, k, \tag{8}$$

where the covariance C is computed with respect to the posterior distribution $p(\theta|x)$. In typical image processing applications φ is very hard to compute, but instead it is possible to simulate (using Markov Chain Monte Carlo, MCMC) from the posterior distribution, see e.g., Winkler (1995). It is then in principle possible to compute $\hat{g}_s$ and $\dot{g}_1$ by simulation, since they are expectations with respect to the posterior distribution of functions which do not involve φ. In practice, the straightforward simulation of the denominator and numerator in (7) is usually not feasible because of the large variation in $\exp\{(1-s)H(\theta)\}$. However, instead it is possible to use the fact that $p_s(\theta)$ has the same form as $p(\theta)$, and to compute $\hat{g}_s$ in the same way as $\hat{g}$. On the other hand, it is often possible to use (8) to evaluate $\dot{g}_1$ by simulation from $p(\theta|x)$.

To calculate the standardized influence rate (2) the standard deviations, $s_i(\dot{g}_{1,i})$, are needed. Under the assumption that the specified model is correct, samples of values of $\dot{g}_{1,i}$ can be calculated, e.g., by first using MCMC to generate a θ and then generate a random x given θ, and finally using (8) to calculate one value of $\dot{g}_{1,i}$. Repeating this procedure will give a sequence of independent $\dot{g}_{1,i}$ values, from which $s_i(\dot{g}_{1,i})$ can be calculated.

As mentioned in the introduction, Berger (1990) considers the sensitivity to replacing the prior $p(\theta)$ with ϵ-contamination priors

$$\{\pi(\theta) = (1-\epsilon)p(\theta) + \epsilon q(\theta)\},$$

for $q(\theta)$ in some class of probability measures. However, it may be hard to find suitable q-s and computations involving the π-s may be very difficult. Ruggeri and Wasserman (1993) studied local sensitivity (infinitesimal changes in the prior) through the norm of the Fréchet derivative of the posterior, and hence their article is closely related to the influence rate. However, again, that seems considerably more difficult to compute in the situation of interest in the present paper.

4 Applications

Example 2 Sensitivity to the prior in visual field testing. Peripheral visual testing, known as perimetry, is an important and widely applied medical procedure. The most common reason for perimetric testing is to detect and monitor glaucoma, a disease in which the patient may gradually go blind.

In a visual field test small spots of light are flashed briefly – one at a time and with waiting time in between – at various points of a regular test point pattern on a hemispherical screen. The patient is instructed to look straight at the center of the screen and to press a button whenever he or she sees a flash of light. In this way one obtains observations, with values "seen" or "not seen", for each stimulus presentation. The "not seen" response occurs if after a certain waiting time the patient has not responded. Computer algorithms which determine the test point sequence, intensities, waiting time for response and when to stop at each test point are quite sophisticated, see Bengtsson et al. (1997). The aim of the visual field test is to determine the "visual threshold" at the tested locations.

In Olsson and Rootzén (1994) we studied Bayes estimation of thresholds, using a Markov random field prior, as discussed in the introduction. In applying this technique the same prior is used for any subject. As an example of the results of Sects. 1 and 2, we exhibit the influence measures for a simulated visual field test.

First a brief discussion of the model. For a complete account see Olsson and Rootzén (1994). The prior distribution was 148-dimensional with parameter $\theta = (t_1, \ldots, t_{74}, s_1, \ldots, s_{74})$ (74 test points with two parameters for each point) where t_i is the "visual threshold" of point i and s_i is an unobservable binary "defect status" variable associated with the test point ($s_i = 1$ corresponds to a "normal" point, and $s_i = -1$ to a "defective" point). In the interior the eight nearest points were taken to be defect status neighbors and the four nearest were taken to be threshold neighbors. The energy function (6) was defined as

$$H(\theta) = -\sum_{i=1}^{74} \sum_{j \in N_i} \beta_{ij} s_i s_j + \sum_{i=1}^{74} \ln(\sigma_i(s_i)) + \frac{1}{2} \left(\frac{t_1 - \mu_1(s_1)}{\sigma_1(s_1)}, \ldots, \frac{t_{74} - \mu_{74}(s_{74})}{\sigma_{74}(s_{74})} \right) A \begin{pmatrix} \frac{t_1 - \mu_1(s_1)}{\sigma_1(s_1)} \\ \vdots \\ \frac{t_{74} - \mu_{74}(s_{74})}{\sigma_{74}(s_{74})} \end{pmatrix}.$$

Here N_i are the defect status neighbors of site i. The matrix A has ones in the diagonal, and has off-diagonal (i, j)-th element equal to $-c/4$ if point j is among the threshold neighbors of point i, and equal to zero otherwise. $|c| < 1$ is a suitably chosen constant determining the degree of threshold correlations. The hyper-parameters $\mu_i(s_i)$ and $\sigma_i(s_i)$ are the expected threshold values and the standard deviations conditional on defect status, respectively, and the hyper-parameters β_{ij} determine the dependence between the defect status variables,

with values choosen such that defective areas resemble glaucomatous visual field loss.

The likelihood for the observed data (described at the start of this example) is a product of the probabilities of "seen" and "not seen" responses, one for each stimulus presentation. The probability of a "seen" is $\mathrm{FP} + (1 - \mathrm{FN} - \mathrm{FP})\Phi(\frac{d - t_i}{\xi_i})$, where FP is the probability of a false positive response (i.e., a response recorded as "seen" due to the patient mistakenly pressing the response button), FN is the probability of a false negative response (the opposite type of mistake) and ξ_i are constants, Φ is the standard normal distribution function and d is the brightness level of the stimulus. The probability of a "not seen" is naturally one minus the probability of a "seen".

We simulated a test on a glaucomatous visual field. Figure 2 shows the influence curves of the threshold estimates t_i, calculated using MPM estimation with 5,000 scans of a Gibbs sampler for the simulation of the posterior distribution.

A few interesting test points numbered 1–5 are listed and commented on in Table 1. Typically 2–6 stimulus presentations are made at each test point. Obviously, in visual field testing the likelihood functions differ substantially between test points since the number of stimulus presentations and the sequence of stimulus intensities differ. The test point labeled 1 has a greater than average number of stimulus presentations (with 5 of 11 "seen") and the influence rate is small as expected. Test points 2, 4 and 5 have the steepest influence curves (and correspondingly high influence rates). This is as could be expected since the threshold value was low and only three stimulus presentations were made at

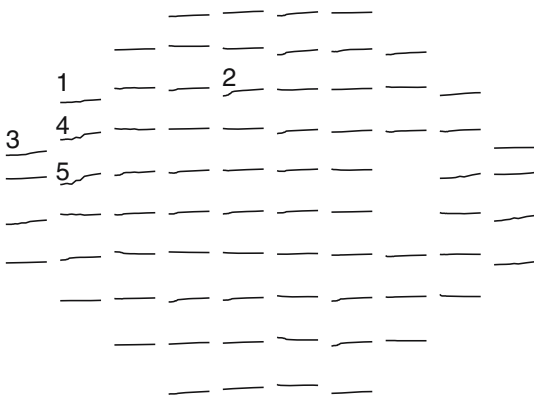

Fig. 2 A graph of 74 influence curves for the prior, $\hat{g}_s$ versus s [see (3)]. The curves are arranged according to the spatial location of the 74 test points of a so-called $30 - 2$ peripheral visual field test pattern. The pattern is centered at the fovea. The spacing between test points is 6 degrees. The four most central points have x and y coordinates ± 3 degrees. This is one of the most common test patterns used by ophthalmologists. The two missing lines correspond to the location of the blind spot. Coordinate axes for the curves (not displayed) are the same for all 74 influence curves. The range of the y-axis (threshold estimate, $\hat{g}_s$) is from -5 to 40, and for the x-axis (s) the range is from 0.1 to 5. A low threshold estimate indicates reduced sensitivity and may be due to glaucoma. However, curves with steep slopes indicate uncertainty in the estimate. Comments on the numbered test points can be found in Table 1. The graph is based on a simulated visual field test

Table 1 The five labeled test points of Fig. 2 with number of stimulus presentations (column 2), number of "seen" stimulus presentations (column 3), influence rate and a comment on each test point

Test points	# Stim.	# Stim. seen	Influence rate, $\dot{g}_1$	Comments on selected influence curves in Fig. 2	
				Influence of prior	Threshold
1	11	5	0.1	Small influence	Low
2	3	1	1.0	Some influence	Slightly low
3	5	0	0.4	Small influence	Very low
4	3	1	1.0	Some influence	Slightly low
5	3	1	2.0	Large influence	Slightly low

these test points. However, three stimulus presentations per test point is quite common and did not itself produce a high influence rate. Thus we conclude that at test points 1 and 3 the threshold estimate was mainly determined by the data, while at test points 2, 4 and 5 the prior influenced the estimates.

Our influence measures may be useful for perimetric testing in three different ways. First, simulation experiments like the one above alleviated our fears that the prior influenced the estimates too much. Secondly, it is desirable for perimetric tests to be short. The influence measures can indicate test points that most need additional stimulus presentations and hence limit the influence of the prior without overly prolonging the test. Finally, they have the potential to be a diagnostic aid by telling doctors that some of the threshold estimates, those which have a large sensitivity to the prior, should be regarded with some caution. However, for routine use of perimetry for diagnosis, information like that presented in Fig. 2 has to be further compressed, perhaps to the extent of just indicating a few points where threshold estimates may be less reliable. We intend to pursue this issue further in the future.

Example 3 Normal prior, normal measurement error, and standardized influence rates. As an illustrative example of the standardized influence rate in a case when the prior model is poorly specified in parts of the image we now consider a synthetic square image consisting of 16×16 pixels. As prior model we used a conditional auto regression with $\theta \sim N(0, \Sigma_\theta)$, for the prior covariance matrix $\Sigma_\theta = 10(I - 0.24B)^{-1}$, and where the i,j-th element of B is one if i and j are neighbors, and zero otherwise.

The likelihood function was obtained by specifying the measurement error, ϵ, to be additive normal and independent between pixels, $\epsilon \sim N(0, D)$, with elements of the diagonal measurement covariance matrix D set to 10 (large measurement error) in columns 1–4 and 9–12 or 0.1 (small measurement error) in columns 5–8 and 13–16. Thus, the likelihood function $l(x|\theta)$ in this example is much less complicated than in Example 2.

We simulated an observed image using the following procedure. A realization, θ, from the prior distribution $p(\theta)$ was simulated. We wanted to see

what happens when the prior is poorly specified in parts of the image. Hence, a perturbation of rows 5–8 and 13–16, in the form of a checkerboard pattern with values -5 or 5, was added to θ, to produce the true image, $\tilde{\theta}$. Thus, this $\tilde{\theta}$ is rather unlikely under the prior. Finally, the observed image X was obtained by adding simulated values of the measurement errors described above to the true image, i.e., the observed image was obtained as $X = \tilde{\theta} + \epsilon$.

The standard normal theory estimate of θ using the prior distribution and likelihood function described above is $\hat{\theta} = \Sigma_\theta(\Sigma_\theta + D)^{-1}X$. We computed the standardized influence rates for this estimator, using $p(\theta)$ and $l(x|\theta)$ (disregarding the perturbation), see Fig. 3. One can note that the areas where the prior is poorly specified are clearly noticeable in the figure. In this simple example it was also possible to calculate the standard deviations $s_i(\dot{g}_{1,i})$ – needed for the standardized influence rate – explicitly, so the simulation techniques of Sect. 4 were not needed.

The squares (5–8, 5–8), (5–8, 13–16), (13–16, 5–8) and (13–16, 13–16) contain the pixels which were assumed both to have a small measurement error and for which the observed image was perturbed by the checkerboard pattern. Small changes in the prior can have large effects in these areas. The areas can be reasonably well identified from the grayscale image, Fig. 3, of the standardized influences rates. One can also see some signs of "edge effects" in the image.

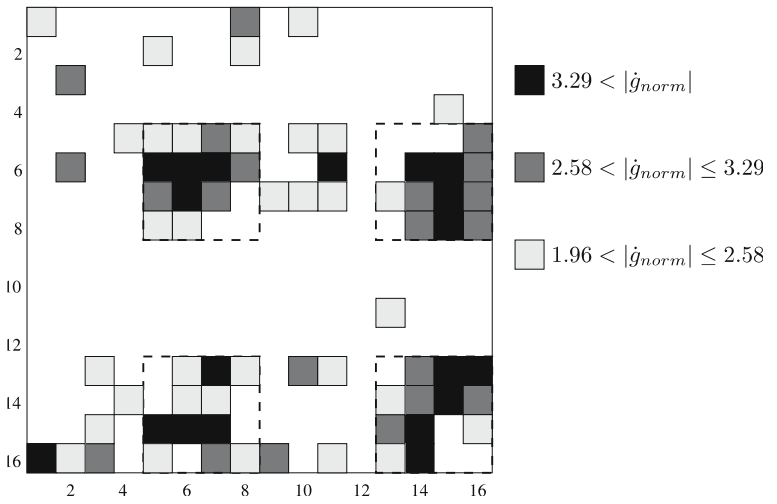

Fig. 3 Grayscale representation of absolute standardized influence rates, $|\dot{g}_{\text{norm}}|$, of Example 3, based on comparisons with normal quantiles. Pixels are white, light gray, dark gray or black, if $|\dot{g}_{\text{norm}}|$ is smaller than 1.96, between 1.96 and 2.58, between 2.58 and 3.29, and larger than 3.29, respectively. Four perturbed areas with low measurement errors, consisting of 4×4 pixels each, are indicated by dashed lines. There are a number of pixels with large values of $|\dot{g}_{\text{norm}}|$ in the perturbed areas of rows 5–8 and 13–16, most of them inside those four areas

5 Conclusions

There is a substantial amount of literature on Bayesian robustness, dealing with many different questions. The present paper is focused on one specific problem in the area: suppose one has decided to use the same Bayesian analysis on a number of similar data sets, e.g., on visual field tests from a number of subjects. How can one then disentangle the influence on estimates of (i) the prior distribution and (ii) the particular data set under analysis? The need was for methods which can be routinely used, without having to set parameters each time a new data set is analysed. The methods have to be simple enough to be displayed and understood in high-dimensional problems, e.g., in image analysis. They also have to be computationally feasible in such applications. To the best of our knowledge, no such methods exist in the literature.

In this paper we have introduced two new influence measures tailored to this problem. The first one, *the influence curve for the prior*, shows how much estimates change when the prior is changed along a curve which interpolates in a "canonical way" between a flat prior, the prior in actual use, and a completely concentrated prior. This measure gives a rather detailed understanding of the extent to which estimates are determined by the prior.

However, in high-dimensional problems the influence curve may contain more information than can be digested. Hence we introduced a second influence measure, *the influence rate*, which is the value of the derivative of the influence curve at the actual prior used. It measures the sensitivity to infinitesimal changes in the prior distribution. For some applications normalized versions of the influence rate may also be useful.

In the first example we illustrated what these new measures mean in the simplest possible situation. The second example concerned a serious real problem encountered in visual field testing for glaucoma diagnosis. In this example the measures were able to identify locations in the visual field where estimates of the "visual threshold" were mostly determined by our prior model and less by the subject's responses. The methods added to the understanding of the very complex algorithms used in visual field testing and can be used to improve test strategies. They are also a potentially useful tool to help doctors evaluate diagnostic uncertainty. However, more empirical work remains to be done before the latter possibility can be exploited in routine perimetric testing.

The third example concerned a stylized image analysis problem, with poorly specified prior in parts of the image. For this example the influence rate was conveniently displayed as a grayscale image. It was able to correctly identify areas where the prior distribution had a large influence.

Acknowledgment We would like to thank anonymous referees and an associate editor for useful advice which improved the paper considerably.

References

Bengtsson B, Heijl A, Olsson J, Rootzén H (1997) A new generation of algorithms for computerized threshold permetry, SITA. Acta Ophthalmol Scand 75:368–375

Bengtsson B, Heijl A, Olsson J (1998) Evaluation of a new threshold visual field strategy, SITA, in normal subjects. Acta Ophthalmol Scand 76:165–169

Berger JO (1980) Statistical decision theory, foundations, concepts, and methods. Springer, Berlin Heidelberg New York

Berger JO (1990) Robust Bayesian analysis: sensitivity to the prior. J Stat Plan Inference 25:303–328

Bose S (1994) Bayesian robustness with more than one class of contaminations. J Stat Plan Inference 40:177–187

Carlin BP, Chaloner KM, Louis TA, Rhame FS (1995) Elicitation, monitoring, and analysis for an AIDS clinical trial, Lecture Notes in Statistics 105, 48–89, Springer, Berlin Heidelberg New York

Clarke B, Gustafson P (1998) On the overall sensitivity of the posterior distribution to its input. J Stat Plan Inference 71:137–150

Delampady M, Dey DK (1994) Bayesian robustness for multiparameter problems. J Stat Plan Inference 40:375–382

Insua RI, Ruggeri F (eds) (2000) Robust Bayesian analysis. Lecture Notes in Statistics, 152, Springer, Berlin Heidelberg New York

Kindermann R, Snell JL (1980) Markov random fields and their applications. American Mathematical Society, Providence

Lehmann EL (1959) Testing statistical hypothesis. Wiley, New York

O'Hagan A (1994) Kendall's advanced theory of statistics. Vol. 2B, Bayesian inference. University Press, Cambridge

Olsson J, Rootzén H (1994) An image model for quantal response analysis in perimetry. Scand J Stat 21:375–387

Ruggeri F, Wasserman L (1993) Infinitesimal sensitivity of posterior distributions. Can J Stat 21:195–203

Winkler G (1995) Image analysis, random fields and dynamic Monte Carlo methods. Springer, Berlin Heidelberg New York