

On the influence of the prior distribution in image reconstruction

Holger Rootzén¹ and Jonny Olsson²

¹Dept of Mathematical Statistics, Chalmers University of Technology, S-412 96 GÖTEBORG, SWEDEN

²Dept of Mathematical Statistics, Lund University, Box 118, S-221 00 LUND, SWEDEN

Summary

Two measures of the influence of the prior distribution $p(\theta)$ in Bayes estimation are proposed. Both involve comparing with alternative priors 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: Influence of the prior distribution, Bayes estimation, image processing, visual field

AMS 1991 classification: Primary 62F15; Secondary 62M30, 68U10.

Research supported by the Bank of Sweden Tercentenary Foundation.

1 Introduction

In this paper we propose two simple measures of how much a Bayes estimate is determined by the prior distribution and how much it is determined by the 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 image may largely be determined by the observations, while other images — or at least part of images — may mainly reflect the prior opinion and perhaps even be in conflict with the available information. Our aim is to find ways to distinguish between those cases.

We were lead to the present research by an important medical problem, Glaucoma diagnosis. This application is discussed in Section 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) and Bengtsson et al (1998)) is implemented in a computerized machine, a field analyser, and is routinely used by ophthalmologists and opticians in many thousand 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.

However, even if this research originated from an effort to improve Glaucoma treatment, 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 supposed to be a “flat” prior which represents ignorance of the actual value of the parameter. We also suppose 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 density

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

with respect to μ . We consider point estimates, \hat{g} , of some parameter function

$g(\theta) = (g^{(1)}(\theta), \dots, g^{(k)}(\theta)) \in \mathbb{R}^k$.

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 a posteriori (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 training sets of data. Aims in the cited literature include providing help with these choices of parameter values, and warning against situations where small parameter changes may give widely varying results.

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

1. shows 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 is provided by Clarke and Gustafson (1998). Carlin et al (1995) partitions priors into those leading to accepting or rejecting an hypothesis. Bayesian robustness methods discussed in the literature considers 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. Further even if the computational problems could be overcome, there remains 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 $\int p(\theta)^s \mu(d\theta) < \infty$, we may take

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

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

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

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* simply is 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 $p(\theta)$.

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

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

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. Sometimes a *standardized influence rate*, \dot{g}_{norm} , with elements,

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

where $s_i(\dot{g}_{1,i})$ is the standard deviation of the i th element, $\dot{g}_{1,i}$, of the influence rate could be used to test for deviations from the prior model. The advantage of $\dot{g}_{1,i}$ and $|\dot{g}_{norm,i}|$ 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}_{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 Section 2 we discuss the influence measures for MPM estimation and the results are illustrated in the very simplest one-dimensional normal case. In Section 3 Markov random field models for the prior are discussed. In Section 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.

2 Influence measures for MPM estimation

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

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

Throughout this section we will assume 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 *Suppose \hat{g}_s given by (3) is well defined. Then, (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

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

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

Proof. (i) By the definition (3) \hat{g}_s exists if and only if both $\int p(\theta)^s l(x|\theta)\mu(d\theta)$ and $\int g(\theta)p(\theta)^s l(x|\theta)\mu(d\theta)$ 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, Section 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 nonempty with endpoints a, b satisfying $a \leq 1 \leq b$.

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

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

Example 1. Normal distributions

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

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

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

Further, assume that the likelihood function is $l(x|\theta) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(x-\theta)^2}$, 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

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

Hence both the MAP and MPM estimates of $g(\theta) \equiv \theta$ is $\hat{g}_s = \nu_s$, 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,

$$\begin{aligned} \hat{g}_1 &= C_{p(\theta|x)}(g(\theta), \log(p(\theta))) = C_{p(\theta|x)}(\theta, -\frac{1}{2\tau^2}(\theta - \mu)^2) \\ &= \dots = \frac{\sigma^2/\tau^2}{(1 + \sigma^2/\tau^2)^2}(\mu - x). \end{aligned} \quad (5)$$

(This is equally easy to obtain by differentiation of (4)). The influence curve and influence rate are shown in Figure 1. for one example.

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 $\sigma^2 \ll \tau^2$ then the estimate is mostly decided by the observation, and if $\sigma^2 \gg \tau^2$ it is mainly determined by the prior, and in neither case the precise amount of confidence in the prior is crucial.

The influence curve also contains this information. In addition it makes it possible to separate between the 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 for this paper is to find tools for understanding 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

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

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

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

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

$$(\dot{g}_1) = -C_{p(\theta|x)}(g^{(i)}(\theta), H(\theta)), \quad i = 1, \dots, k \quad (8)$$

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. 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, straightforward simulation of denominator and numerator in (7) usually is not feasible, because of the large variation of $\exp\{(1-s)H(\theta)\}$. However, instead it is possible to use 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 of (2) the standard deviations, $s_i(\dot{g}_{1,i})$, are needed. Under the assumption that the specified model is correct samples of $\dot{g}_{1,i}$ -values 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 $\dot{g}_{1,i}$. Repeating this procedure will give a sequence of independent $\dot{g}_{1,i}$ -values, from which $s_i(\dot{g}_{1,i})$ are 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. The article by Ruggeri and Wasserman (1993) studies local sensitivity (infinitesimal changes in the prior) through the norm of the Fréchet derivative of the posterior, and hence is closely related to the influence rate. However, again, these 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 eye may gradually go blind. In a visual field test, small spots of light are flashed briefly 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 she/he sees a flash of light. The algorithm which determines the intensities and the number of stimulus presentations at each test point is quite sophisticated. The aim of the visual field test is to determine the “seeing thresholds” at the tested locations.

In Olsson and Rootzén (1994) we study 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 Sections 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 (74 test locations with two parameters for each location) with parameter $\theta = (t_1, \dots, t_{74}, s_1, \dots, s_{74})$ where t_i is the “seeing 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). The energy function (6) was

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

Here N_i are the defect status neighbors to 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. 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.

The likelihood was a product of probabilities of seen and not seen stimuli. The probability of a “seen” is $FP + (1 - FN - FP)\Phi\left(\frac{d - t_i}{\xi_i}\right)$, where FP (the probability of a false positive response), FN (the probability of a false negative response) and ξ_i

are constants, Φ is the standard normal distribution function and d is the brightness level of the stimulus.

We simulated a test on a glaucomatous visual field. Figure 1 shows the influence curves of the threshold estimates t_i , calculated using MPM estimation with 5000 scans of a Gibbs sampler for the simulation of the posterior distribution. A few test points with influence curves which deviated most from the majority were numbered 1–5 and listed in Table 1. 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 in those test points. In test points 1 and 3 which had somewhat smaller influence rates the threshold value was also low but more stimulus presentations were made. However, three stimulus presentations per test point were quite common and did not by itself produce a high influence rate. Thus we conclude that in test points 1 and 3 the threshold estimate was mainly determined by the data, while in test points 2, 3 and 5 the prior did influence 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 to make a short perimetric test. The influence measures can indicate those test points that most needs 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 indicating the uncertainty of threshold estimates. However for routine use of perimetry for diagnosis, information like the one presented in Figure 1 has to be further compressed, perhaps to the extent of just indicating a few points where thresholds estimates may be less reliable. We intend to pursue this issue further in the future.

Test location	# stim	# seen stim	influence rate, \hat{g}_1	Comments to selected influence curves in Fig 1	
				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

Table 1: Number of stimulus presentations, number of seen stimulus presentations, influence rate and a comment to each of the five numbered test points of Figure 1.

Example 3. Normal prior, normal measurement error, and standardized influence rates

As an illustrative example of the standardized influence rate we now consider a synthetic square image consisting of 16×16 pixels. As prior model we used a conditional auto regression with $\theta \sim N(0, \Sigma_\theta)$, $\Sigma_\theta = 10(I - 0.24B)^{-1}$, 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 variance 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.

We simulated an observed image using the following procedure. A realisation, θ , from the prior distribution $p(\theta)$ was simulated. Next, 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, $\tilde{\theta}$. 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 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 rate for this estimator, see Figure 2. In this simple example it was also possible to calculate the standard deviations $s_i(\hat{g}_{1,i})$ explicitly, so the simulation techniques of Section 4 were not needed.

The squares (5–8, 5–8), (5–8, 13–16), (13–16, 5–8) and (13–16, 13–16) contains the pixels which were assumed to have a small measurement error and for which the observed image was perturbed by the checkerboard pattern. Small changes of the prior can have large effects in these areas. The areas can be rather well identified from the grayscale picture, Figure 2, of the standardized influences rates. One can also see some signs of “edge effects” in the picture.

Fig 1 here

Fig 2 here

5 Conclusions

There is a substantial literature on Bayesian robustness, dealing with many different questions. The present paper is focussed 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 from (i) the prior distribution and (ii) the particular data set under analysis? The need was for methods which can be routinely used, without setting parameters each time a new data set is analysed. The methods had to be simple enough to be displayed and understood in highdimensional, e.g. image analysis, problems. They also had 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 highdimensional 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 of the prior distribution. For some applications also normalized versions of the influence rate may be useful.

In a first example we illustrated what these new measures mean in the simplest possible situation. The second example concerned a serious real problem from visual field testing for Glaucoma diagnosis. For this example the measures were able to identify locations in the visual field where estimates of the “seeing 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, can be used to improve test strategies and are potentially useful as a tool to evaluate diagnostic uncertainty.

A third example concerned a stylized image analysis problem. For this example the influence rate was conveniently displayed as a greyscale picture. It was able to correctly identify areas where the prior distribution had a large influence.

Acknowledgement We like to thank anonymous referees and an associate editor for useful advice.

References

- [1] Bengtsson, B., Heijl, A. and Olsson, J. (1998), Evaluation of a new threshold visual field strategy, SITA, in normal subjects. *Acta Ophthalmologica Scandinavica*, 76, 165–169.
- [2] Berger, J. O. (1980), *Statistical decision theory, foundations, concepts, and methods*. Springer-Verlag, New York.
- [3] Berger, J. O. (1990), Robust Bayesian analysis: sensitivity to the prior. *Journal of Statistical Planning and Inference*, 25, 303–328.
- [4] Bose, S. (1994), Bayesian robustness with more than one class of contaminations. *Journal of Statistical Planning and Inference*, 40, 177–187.
- [5] Carlin, B.P., Chaloner, K.M., Louis, T.A. and Rhame, F.S. (1995), Elicitation, monitoring, and analysis for an AIDS clinical trial. *Lecture Notes in Statistics*, 105, 48–89, Springer-Verlag, New York.
- [6] Clarke, B. and Gustafson, P. (1998), On the overall sensitivity of the posterior distribution to its input. *Journal of Statistical Planning and Inference*, 71, 137–150.
- [7] Delampady, M. and Dey, D.K. (1994), Bayesian robustness for multiparameter problems. *Journal of Statistical Planning and Inference*, 40, 375–382.
- [8] Kindermann, R. and Snell, J.L. (1980), *Markov random fields and their applications*, American Mathematical Society, Providence, R.I.
- [9] Lehmann, E.L. (1959), *Testing statistical hypothesis*, Wiley, New York.
- [10] O’Hagan, A. (1994), *Kendall’s advanced theory of statistics, Volume 2B, Bayesian inference*, University Press, Cambridge.
- [11] Olsson, J. and Rootzén, H. (1994), An image model for quantal response analysis in perimetry. *Scandinavian Journal of Statistics*, 21, 375–387.
- [12] Ruggeri, F. and Wasserman, L. (1993), Infinitesimal sensitivity of posterior distributions. *Canadian Journal of Statistics*, 21, 195–203.

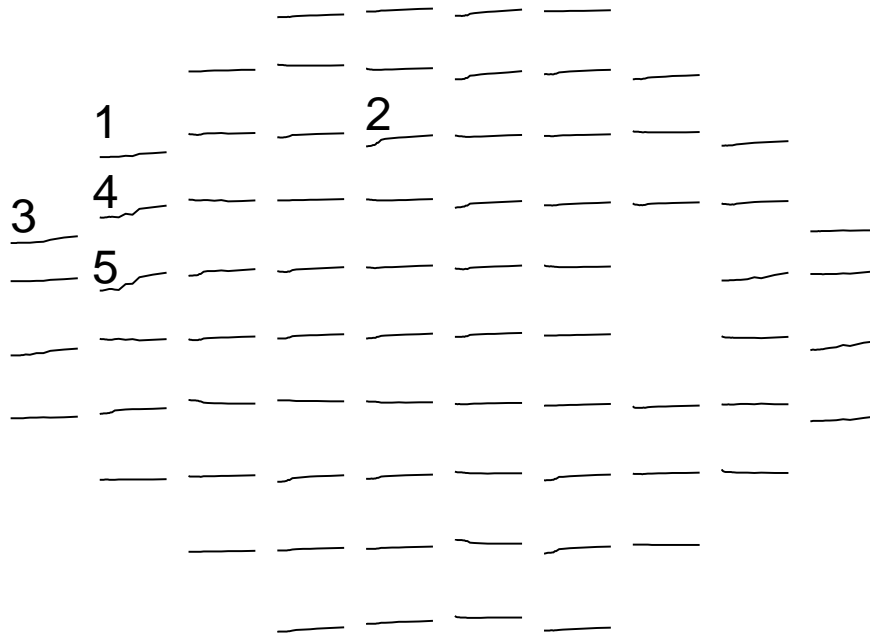


Figure 1: In a visual field test seeing thresholds are estimated at 74 locations in the peripheral visual field. The pattern of locations is shown in the figure, but instead of the seeing threshold, at each location the influence curve for the prior, g_s , as a function of s (see (3)) is plotted. The coordinate axes for the curves are not shown. In the 74 plots g_s is the y-variable, and ranges from -5 to 40, and s is the x-variable and ranges from 0.1 to 5. Comments to numbered test locations are found in Table 1. The plot is based on a simulated visual field test.




Figure 2: Images, grayscaled and rounded to one decimal number, of standardized influence rates, \hat{j}_{norm} , of Example 3. In the grayscaled image, pixels not significant at the 5 % level are shown white and the significant pixels are dark. The darkest pixels are most significant. Disturbed areas with low measurement errors can be identified. Areas within dashed boxes are disturbed areas and areas within framed boxes have low measurement error.

-2.2	-0.8	-0.7	-0.0	0.1	-0.4	0.9	2.9	0.2	-2.3	0.2	-1.3	-0.9	-0.8	0.4	-0.3
1.1	1.9	0.3	-0.6	-2.0	1.1	-1.0	-2.0	-1.7	1.0	0.3	1.1	-0.5	0.7	-1.0	1.4
-0.3	-2.9	0.1	-1.5	-0.2	-0.3	1.4	0.8	1.2	-0.3	-0.3	0.6	0.4	-0.6	0.4	-0.1
-0.3	-0.0	-0.0	-0.0	-0.5	-1.5	-0.2	1.2	-1.0	-0.8	-0.4	1.4	-0.6	0.1	-2.2	-0.6
0.8	-0.5	0.6	-2.2	2.5	-2.1	3.2	-2.4	1.3	-2.3	2.1	0.1	1.7	-0.2	1.7	-3.0
-0.5	3.2	-0.5	1.9	-3.6	3.8	-4.2	3.2	-1.3	1.4	-4.0	0.6	-1.5	5.2	-3.8	2.7
-0.2	-0.8	0.6	-0.8	2.9	-3.8	2.7	-1.8	2.3	-2.1	2.4	-0.7	2.0	-3.2	4.3	-3.0
-0.6	-1.6	-0.7	1.6	-2.2	2.1	-0.9	1.2	-1.4	0.7	-0.2	1.0	-1.7	2.6	-3.6	2.7
0.9	-0.4	-0.4	0.4	-0.2	0.6	-1.5	0.2	0.5	1.0	0.5	-0.4	-0.3	0.1	0.9	-1.5
0.9	-0.5	-1.8	0.6	-0.1	0.9	-0.0	-1.3	-0.4	-0.4	0.1	-0.7	0.6	0.9	-1.7	1.0
-1.0	0.9	-0.3	-0.1	-0.6	-0.6	0.4	0.6	1.2	1.1	0.1	0.9	-2.4	-0.2	-0.8	1.2
0.3	1.7	-0.5	-0.7	0.3	0.3	-1.1	-0.3	-0.6	-0.6	-0.4	-1.6	1.9	-0.3	-0.2	0.5
-1.6	-1.7	-2.1	-0.5	-1.1	-2.3	-4.3	-2.1	-1.6	-3.0	2.1	-0.4	-1.4	-2.9	-3.5	-3.4
-1.5	0.3	0.4	2.0	-1.8	2.3	-2.1	1.8	-0.8	1.3	-0.7	1.6	-2.2	3.0	-3.8	3.1
0.0	0.3	2.3	0.4	3.5	-3.3	4.3	-1.5	1.1	-1.1	1.9	-1.4	3.1	-3.6	0.1	-2.4
-3.5	2.1	-3.1	0.9	-2.5	1.7	-2.6	2.2	-2.7	-0.8	-2.0	1.1	-2.3	3.6	-1.4	0.9