

Estimation of centres and radial intensity profiles of spherical nanoparticles in digital microscopy

Mats Kvarnström¹ and Chris A. Glasbey*²

¹ Department of Mathematical Statistics, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden

² Biomathematics and Statistics Scotland, King's Buildings, Edinburgh, EH9 3JZ, Scotland

Summary

Control of the microscopic characteristics of colloidal systems is critical in a wealth of application areas, ranging from food to pharmaceuticals. To assist in estimating these characteristics, we present a method for estimating the positions of spherical nano-particles in digital microscopy images. The radial intensity profiles of particles, which depend on the distances of the particles from the focal plane of the light microscope and have no closed functional form, are modelled using a local quadratic kernel estimate. We also allow for the case where pixel values are censored at an upper limit of 255. Standard errors of centre estimates are obtained using a sandwich estimator which takes into account spatial autocorrelation in the errors. The approach is validated by a simulation study.

Key words: censored pixels, colloidal suspension, diffusion coefficient, image analysis, kernel regression, rotational symmetry, sandwich estimator, sub-pixel accuracy

1 Introduction

Control of the microscopic characteristics of colloidal systems, primarily the interaction and diffusion coefficients, is critical in a wide range of application areas, including food products and pharmaceuticals. These characteristics determine the macroscopic properties of the system, such as whether the particles will coagulate or remain freely diffusing. For example, in milk, interactions between the small (100 nm to 1 μm in diameter) fat particles and proteins suspended in the fluid determine whether it coagulates into cheese or yoghurt. In pharmaceutical drugs, the active substances must remain stable in tablet or liquid form for several months between production and medication. When finally delivered to the body, these substances have to be released in a controlled fashion in order for them to be effective. The common factor in both these examples is to be able to observe and predict the effect of various modifications on the characteristics of the colloids. However, since the theory for the interactions of colloidal particles fails to predict the behaviour of important suspensions (see, for example, Grier, 1998; Han and Grier, 2005), it has been found necessary to resort to quantitative measurements at a microscopic level to investigate colloidal systems. For a review of the microscopy of colloidal suspensions, see Elliot and Poon (2001) and for a book on colloids in general, see Evans and Wennerström (1999).

Figure 1 shows a single image from a video sequence of a colloidal suspension obtained using a light microscope. The nano-particles are made of latex (polystyrene) and are all identical in size. Their differing appearances are due to variations in their depths relative to the focal plane. The reason for using latex spheres is that suspensions of this kind can be used to emulate most features of more general colloidal systems, by varying the solvent and salt concentrations.

For quantitative measurements of the properties of a colloidal system from images such as Figure 1, we first need to identify the particle centres. Since a typical diffusion of particles is in the order of a few pixels between consecutive

* Corresponding author: email: chris@bioss.ac.uk, Phone: +44 (0)131 650 4899




Fig. 1 An image from a video sequence of diffusing polystyrene particles. The particles are all 494 nm in diameter, and the image consists of 512×512 pixels, each 180×180 nm in size. Pixel values range from 0, displayed as black, to 255, displayed as white.

images, sub-pixel accuracy in the position estimation in the plane is preferable if the diffusion coefficient is to be estimated. Although in principle it is possible to estimate the depths of particles from their appearance, this can be done with considerably less precision than estimation of the particle centres in the two dimensions parallel to the image plane. However, provided that diffusion is isotropic, which is the case for the particles in Figure 1, it is sufficient to locate the particles in only two dimensions (Kvarnström, 2005a), so in this paper we only consider these two dimensions. Unfortunately, to estimate interaction coefficients it is crucial to estimate the depth as well, a point we will return to in §7.

It is evident from Figure 1 that we should be able to estimate particle centres by exploiting the rotational symmetry in their appearance. However, our problem is more challenging than previous colloidal studies using digital video microscopy (Crocker and Grier, 1996), where particles were more uniform in appearance because the depth of focus was shallower, at $\pm 0.5 \mu\text{m}$, rather than $\pm 15 \mu\text{m}$ in our application. In their studies, each particle appeared as a bright circular set of pixels. Further, particles were effectively confined to a crystallised structure in their colloidal suspensions. They achieved sub-pixel accuracy by calculating the geometric centre of the brightness-weighted centroid, obtaining standard errors of about 0.1 pixels (where each pixel was 85 nm in length).

Imaging of objects, even those as simple as spheres, is still a topic of theoretical interest in optics research (Ovryn and Izen, 2000). Ray-tracing using Fourier optics (see, for example, a standard textbook on optics such as Hecht, 2002, chapters 9–12), as used with cashmere fibres (Glasbey et al., 1994) does not work here since the particles are similar in size to the wavelength of light. An alternative and more advanced approach to Fourier optics is to use Mie-theory, as Ovryn and Izen (2000) did to predict the appearance of polystyrene spheres of diameter $7 \mu\text{m}$. However, no closed functional form for the radial intensity profiles is known. Therefore, we resort to empirical methods to model the profiles.

A further complication in our application is that pixel values are censored at an upper limit of $T = 255$. Censoring of pixel values occurs when the intensity measured by the camera at a certain location is higher than the digitised image can represent, which is $2^8 - 1 = 255$ for the images considered here. The censoring can therefore be said to be a result of saturation in the imaging process. For estimating particle centres, censoring does not have much effect, since it occurs on an annulus around the true particle centres, so we can simply omit such pixels from the estimation procedure. However, if radial intensity profiles are needed, for example for depth estimation, then the more sophisticated and computationally-intensive method introduced in §6 can be deployed.




Fig. 2 A single particle as it appears at 15 different distances from the focal plane of the microscope, with the number denoting the depth index relative to the focal plane. Each sub-image consists of 27×27 pixels, and the distance in depth between two consecutive sub-images is $0.6 \mu\text{m}$. Censored pixels are displayed as white.

In §2 we propose a model for the image data, and in §3 we fit the model to estimate particle centres. In §4 we give results, which we validate in §5 using a simulation study. In §6 we extend the estimation method to include censored pixels. Finally, in §7 we discuss the methodology and its extensions.

2 Model

In Figure 2 we display a single particle imaged at a range of distances from the focal plane of the microscope. They were produced by letting particles adsorb on the glass surface of the specimen, and then varying the microscope focus. The sub-image labelled 0 corresponds to a particle at the focal plane, and sub-images with negative and positive labels are below and above the focal plane, respectively. Pixels displayed as white are those which are censored at the value of $T = 255$. For this particle, there is censoring at all depths except 0, 6 and 7. At negative depths (i.e. below the focal plane), censoring occurs for pixels close to the centre of the particle, and for particles above the focal plane, on the first fringe. The relative amount of censoring is however not very large; between 0.5 and 4.5% of the $27^2 = 729$ pixels are censored, with the largest number immediately above the focal plane (depths 1, 2, and 3).

We denote a digital image such as Figure 1 by I . This is a matrix of pixel values I_i indexed by pixel locations $i = (i_1, i_2) \in \mathcal{D} \subset \mathbb{Z}^2$, where \mathcal{D} is the set of indices i for which the image is defined. Each pixel value I_i is an integer between 0 and 255. We associate a set $\mathcal{N} \subset \mathcal{D}$ of pixel locations with a particular particle in an image. The choice of neighbourhood shape and size is a compromise between the precision of estimating the particle centre on one hand, and computational effort and interference from nearby particles on the other. One approach is to let \mathcal{N} consist of all pixel locations $i \in \mathcal{D}$ within a specified distance of an initial guess at the particle centre. Throughout this paper, however, for simplicity we take \mathcal{N} to be the pixel locations in a predetermined 27×27 square, as in Figure 2.

We assume that, for a particle centred at $x = (x_1, x_2) \in \mathbb{R}^2$

$$I_i = \begin{cases} f(\|i - x\|) + \epsilon_i & \text{if } f(\|i - x\|) + \epsilon_i < T \\ T & \text{otherwise} \end{cases} \quad \text{for } i \in \mathcal{N}, \quad (1)$$

where f is termed the *radial intensity profile* and $\|i - x\|$ denotes the Euclidean distance between pixel i and x . The noise component, ϵ_i , is assumed to be normally distributed with isotropic covariance

$$\text{Cov}\{\epsilon_i, \epsilon_j\} = \sigma^2 \exp(-\theta \|i - j\|) \quad \text{for } i, j \in \mathcal{N}, \quad (2)$$

with parameters σ^2 and θ . This is a commonly used covariance structure in images and it was here observed to be a reasonable model when the noise structure of background images was examined (that is, images where no particle is present). See also the sample autocovariances compared with the exponential model in Figure 5.

The intensity profile $f : \mathbb{R} \mapsto \mathbb{R}$ is assumed to be a smooth function with at least two continuous derivatives and symmetric in its argument. For estimating f we have chosen to use a local quadratic kernel estimator with a Gaussian kernel function. We recommend Hastie and Tibshirani (1990, chapters 3 and 4) and Fan and Gijbels (1996, in particular chapters 3 and 4) as general references on non-parametric estimation using local polynomial estimators. One benefit over the alternative spline smoother, is that the derivative of f , which is needed when estimating the standard errors (see §3.3), is given "for free" when kernel smoothing is used. Also, compared to a local linear method, we found the local quadratic method to be less sensitive to the choice of bandwidth.

3 Estimation

We estimate the parameters in model (1) by least squares, ignoring both the correlation in the errors and any censored pixels. We do this for computational speed, and since censoring occurs on an annulus around the true centre, censored pixels only have a small effect on the efficiency of the estimators, which also has been verified by simulations (see §7 for a discussion on this). Our primary concern in this section is to estimate the particle centre. We will nevertheless present a more sophisticated estimator in §6 for cases where radial intensity profiles are of specific interest.

For a particle associated with a neighbourhood \mathcal{N} , we use the minimiser of the sum of squares

$$\mathcal{S}(x) = \sum_{i \in \mathcal{N}_u} \left\{ I_i - \hat{f}(\|i - x\|) \right\}^2 \quad (3)$$

with respect to $x \in \mathbb{R}^2$, as an estimate of the particle centre. Here, $\mathcal{N}_u = \{i \in \mathcal{N} : I_i < T\}$ denotes the set of uncensored pixel locations in \mathcal{N} . The estimate \hat{f} of the radial intensity profile f , is calculated at each x by means of a local quadratic kernel estimator.

3.1 Radial intensity function

At each evaluation of $\mathcal{S}(x)$ at a candidate location x , we require the radial intensity function \hat{f} , which we obtain by approximating f locally by a polynomial of degree 2,

$$f(r') \simeq \sum_{k=0}^2 \beta_{rk} (r' - r)^k \quad \text{for } r' \simeq r. \quad (4)$$

For a specified candidate centre, x , to calculate the value of the estimate of the local quadratic kernel estimate at an arbitrary point $r \in \mathbb{R}$, we fit a second-order polynomial by minimising the local sum of squares,

$$\begin{aligned} \mathcal{S}_r = & \sum_{i \in \mathcal{N}_u} K_h(\|i - x\| - r) \left\{ I_i - \sum_{k=0}^2 \beta_{rk} (\|i - x\| - r)^k \right\}^2 \\ & + \sum_{i \in \mathcal{N}_u} K_h(-\|i - x\| - r) \left\{ I_i - \sum_{k=0}^2 \beta_{rk} (-\|i - x\| - r)^k \right\}^2 \end{aligned} \quad (5)$$

with respect to the triple $(\beta_{r0}, \beta_{r1}, \beta_{r2})$. Here K_h denotes the Gaussian kernel with bandwidth h , that is $K_h(z) = \frac{1}{h} \exp\left(-\frac{z^2}{2h^2}\right)$. To ensure that the estimated intensity profile is symmetric in its argument, as specified in the model, we use all pixel values twice, both at radius $+\|i-x\|$ and $-\|i-x\|$. Denote by $(\hat{\beta}_{r0}, \hat{\beta}_{r1}, \hat{\beta}_{r2})$ the minimisers to equation (5) at point r . The estimate $\hat{f}(r)$ at point r , is $\hat{\beta}_{r0}$.

Note that the minimisation of equation (5) is straightforward, since it is a weighted least-squares problem for each value of r . One benefit of using a local quadratic kernel estimator, is that the first derivative of f at each point r can be estimated by

$$\frac{d\hat{f}}{dr}(r) = \hat{\beta}_{r1}. \quad (6)$$

This will be used when estimating the standard errors in §3.3.

For the kernel bandwidth, h , we initially used a value of 0.7 pixels. Then, after fitting the model, we chose a revised value for h using a cross-validation criterion, and refitted the model. Denote by \hat{f}_h^{-i} the local quadratic kernel estimate of f when leaving out the i th data point, and using h as bandwidth. For each h we constructed the cross-validation score

$$CV(h) = \sum_{i \in \mathcal{N}_u} \left\{ I_i - \hat{f}_h^{-i}(\|i-y\|) \right\}^2, \quad (7)$$

and chose the value of h that minimised $CV(h)$. In practice, CV is calculated for a discrete set of values.

3.2 Algorithm

For optimisation of \mathcal{S} we used a greedy search algorithm: compute the value of \mathcal{S} for values x in a 5×5 square mesh with distance Δ to nearest horizontal and vertical neighbour. Find the minimiser x_c among these, and re-calculate \mathcal{S} for a similar mesh, but this time centred on x_c and with inter-distance $\Delta/2$. This is then iterated until the desired precision is reached. A reasonable such limit is $\eta = 0.001$ pixels.

To start the algorithm, we assume a bandwidth $h = 0.7$ pixels and a first candidate particle centre x_c .

1. Let $\Delta = 1$ pixel.
2. Calculate $\mathcal{S}(x)$ for each of the positions x in the 5×5 square mesh centred around x_c as defined above: For each x :
 - (a) Calculate the distances $r_i = \|i-x\|$ for $i \in \mathcal{N}$.
 - (b) For each r_i , estimate $f(r_i)$ by the minimiser $\hat{f}(r_i) = \beta_{r_i0}$ of the least-squares problem (5) above.
 - (c) Calculate the sum of squares $\mathcal{S}(x)$ between $\hat{f}(r_i)$ and I_i for all $i \in \mathcal{N}_u$
3. Let x_c be the minimiser of $\mathcal{S}(x)$ among the 5×5 positions from the previous step.
4. If $\Delta > \eta$: Let $\Delta = \Delta/2$ and go to step 2. Otherwise, proceed.
5. We are finished, let $\hat{x} = x_c$.

Upon termination of the above algorithm, we use cross-validation, as explained in the previous section, to find a new bandwidth h . (A suitable discrete set of bandwidths for the cross-validation score are from 0.3 up to 1.5, in increments of 0.05.) We then re-run the algorithm using \hat{x} as candidate centre and the cross-validation minimiser h as bandwidth.

3.3 Standard errors

Standard errors of particle centre estimates were derived using a sandwich estimator, which takes into account that the errors are correlated. We used

$$\widehat{\mathbf{Var}}\{x\} = (F^T F)^{-1} F^T \Sigma F (F^T F)^{-1} \quad (8)$$

where F is an $|\mathcal{N}_u| \times 2$ matrix with

$$F_{ij} = \frac{\partial \hat{f}(\|i - x\|)}{\partial x_j}$$

and Σ is $|\mathcal{N}_u| \times |\mathcal{N}_u|$ with $\Sigma_{ij} = \mathbf{Cov}\{\epsilon_i, \epsilon_j\}$ (see, for example Owen, 2001, chapter 4.6). However, both F and Σ need to be estimated.

We computed F by first applying the chain rule of differentiation, and then using the estimated derivatives of f (6), giving:

$$F_{ij} = \frac{d\hat{f}(\|i - x\|)}{d\|i - x\|} \frac{\partial \|i - x\|}{\partial x_j} \simeq \beta_{\|i-x\|,1} \frac{x_j - i_j}{\|i - x\|}.$$

To obtain an unbiased estimate of σ^2 , we first note the value of the local quadratic kernel estimator in each r is linear in the pixel values, since it is a weighted least-squares problem. Hence, we can express \hat{f} as

$$\hat{f}(\|i - x\|) = \sum_{j \in \mathcal{N}_u} W_{ij} I_j$$

for each $i \in \mathcal{N}_u$, for some matrix W , termed the *equivalent kernel matrix*. The trace of W , is an estimate of the degrees of freedom for estimation of f (see, Hastie and Tibshirani, 1990, pages 52-53). Therefore an unbiased estimate of σ^2 is given by

$$\hat{\sigma}^2 = \frac{1}{|\mathcal{N}_u| - \text{trace}\{W\}} \sum_{i \in \mathcal{N}_u} \hat{\epsilon}_i^2 \quad \text{where } \hat{\epsilon}_i = I_i - \hat{f}(\|i - x\|),$$

$|\mathcal{N}_u|$ denotes the number of elements in \mathcal{N}_u and $\hat{\epsilon}$ denotes the model residuals.

The sample autocovariance, C_j , for lag $j = (j_1, j_2)$ is defined as

$$C_j = \frac{1}{N_j} \sum \hat{\epsilon}_i \hat{\epsilon}_{i+j} \quad (9)$$

where the summation is over all pixels such that both $i = (i_1, i_2)$ and $i + j = (i_1 + j_1, i_2 + j_2)$ are uncensored, and N_j is the number of elements in the summation. The exponential correlation function (2) was fitted by estimating θ by

$$\hat{\theta} = -\log \left(\frac{C_{01} + C_{10}}{2C_{00}} \right),$$

using the mean of the two sample autocovariances for lags at distance unity.

Finally, we are able to estimate $\mathbf{Var}\{\hat{x}\}$ by replacing F and Σ in equation (8) by their estimates and obtain

$$\widehat{\mathbf{Var}}\{\hat{x}\} = (\hat{F}^T \hat{F})^{-1} \hat{F}^T \hat{\Sigma} \hat{F} (\hat{F}^T \hat{F})^{-1}.$$

In the simulation study presented in §5, this estimate is found to be in good agreement with the root-mean-square errors.




Fig. 3 Fitted models ($\min(\hat{f}, T)$), for the sub-images in Figure 2.

4 Results

In Figure 3 we display the reconstructed versions of the sub-images of Figure 2 after the centres and the intensity functions have been estimated. The agreement with Figure 2 can be seen to be good. The residuals are shown in Figure 4, also displayed in the form of an image. Autocorrelation is evident, which is consistent with the model specification in (2). Cross-validation bandwidths for the different depths ranged from 0.9, for the depths furthest away from the focal plane, down to 0.5 for the depths closest to the focal plane.

In Figure 5 we show the estimated autocovariances (9) for residuals at depth -5 , displayed on the left as an image and on the right as a plot against Euclidean distance. Also shown in the plot is the fitted exponential correlation function. As can be seen, the exponential correlation function gives a reasonable, albeit slightly pessimistic, fit.

We fitted the model to data from a further 5 particles. Figure 6 shows the average standard errors of the particle centre for each of the 15 depths. We see that the precision in the estimation depends on the depth of the particle. Nevertheless, in all cases the standard errors are considerably less than unity, indicating sub-pixel accuracy. Centres appear to be estimated more precisely if particles are close to, but not precisely, in focus. To validate these results we conducted a small simulation study.

5 Simulation Study

We created an artificial image of size 27×27 for a particle at depth z as follows. The true centre for the particle was chosen as the middle pixel of the image, but with an offset chosen uniformly from $[-0.5, 0.5]$ independently for both coordinates. Then the pixel values were assigned the values predicted by the radial intensity profile f at depth z , based on the estimates shown in Figure 3. To this we added zero mean Gaussian noise with exponential covariance as specified in equation (2), via a Cholesky factorisation of the covariance matrix. From consideration of the estimated parameters from




Fig. 4 Residuals ($\hat{\epsilon}$), for the sub-images in Figure 2. (Zero values are displayed as mid-grey, positive residuals as lighter greys, negative residuals as darker greys and censored pixels as black.)




Fig. 5 Estimated autocovariances (C_j) for residuals at particle at depth -5 in Figure 2. On the left they are displayed as an image for each lag j , with the largest value represented as white and the smallest as black. On the right, C_j is plotted as a function of $\|j\|$, together with the fitted exponential correlation function.

background images as well as residuals, the noise parameters we used were $\sigma^2 = 25$ and $\theta = 0.6$, which was considered to be a suitable (i.e. slightly pessimistic) choice of parameters for the simulation. Finally, we censored all pixels exceeding $T = 255$.

We simulated 500 images at each depth, fitted the model as described in §3 and summarised the results. We found that the estimated particle centres were unbiased, and that there was very good agreement between the estimated standard errors and the root-mean-square (RMS) errors of centres. Figure 7 shows a plot of average standard errors and RMS errors, from




Fig. 6 Average standard error of particle centre for 6 particles at each of 15 depths.

which we see that standard errors slightly over-estimate true errors, and estimation to sub-pixel accuracy is confirmed. We also note that the relationship of precision with depth is consistent with that in Figure 6.




Fig. 7 Comparison between the RMS errors (boxes) of the particle centre estimation errors and the mean of the estimated standard errors (stars) for the simulation study.

6 Estimation of radial intensity profiles

So far our focus has been on estimation of particle centres in 2-D. However, if unbiased estimates of radial intensity profiles are needed, for example to be used as templates for depth estimation, then a more sophisticated and computationally-intensive method has to be adopted, that takes account of censored pixels. Therefore, to obtain $\hat{f}(r)$ for an arbitrary point




Fig. 8 Data and estimated radial intensity profile for the particle at depth -5 from Figure 2, both when taking account of censored pixels (dashed) and not (solid).

$r \in \mathbb{R}$, we replace the local sum of squares, \mathcal{S}_r (5) by a local log-likelihood (ignoring spatial autocorrelations):

$$\begin{aligned} \mathcal{L}_r &= \frac{\mathcal{S}_r}{2} + \frac{1}{2} \log \sigma^2 \left\{ \sum_{i \in \mathcal{N}_u} K_h(\|i - x\| - r) + K_h(-\|i - x\| - r) \right\} \\ &\quad - \sum_{i \in \mathcal{N} \setminus \mathcal{N}_u} K_h(\|i - x\| - r) \log \left\{ \Phi \left(\frac{\sum_{k=0}^2 \beta_{rk} (\|i - x\| - r)^k - T}{\sigma} \right) \right\} \\ &\quad - \sum_{i \in \mathcal{N} \setminus \mathcal{N}_u} K_h(-\|i - x\| - r) \log \left\{ \Phi \left(\frac{\sum_{k=0}^2 \beta_{rk} (-\|i - x\| - r)^k - T}{\sigma} \right) \right\}, \end{aligned} \quad (10)$$

which we minimise with respect to $(\beta_{r0}, \beta_{r1}, \beta_{r2})$. Here, $\mathcal{N} \setminus \mathcal{N}_u$ denotes the set difference, i.e. the censored pixel locations in \mathcal{N} , and Φ the standard normal cumulative distribution function. Unlike (5), no closed form solution exists and iterative optimisation methods have instead to be used. The reason behind minimising (10) is that it corresponds to maximising the local likelihood

$$\prod_{i \in \mathcal{N}_u} \left\{ \frac{1}{\sigma} \phi \left(\frac{I_i - \sum_{k=0}^2 \beta_{rk} (\pm \|i - x\| - r)^k}{\sigma} \right) \right\}^{w_i} \prod_{i \in \mathcal{N} \setminus \mathcal{N}_u} \left\{ \Phi \left(\frac{\sum_{k=0}^2 \beta_{rk} (\pm \|i - x\| - r)^k - T}{\sigma} \right) \right\}^{w_i}$$

of observations I_i , that are possibly censored above T , where each observation has a weight $w_i = K_h(\pm \|i - x\| - r)$ depending on distance from the point r , and we use \pm to denote the inclusion of two terms, one with $+$ and one with $-$.

Figure 8 shows pixel values for the particle at depth -5 plotted against distance from the centre, together with the estimated intensity profile, both when taking account of censored pixels (dashed) and when not (solid), that is, when minimising equation (10) or equation (5), respectively. We see that by including censored pixels in the estimator, we have estimated values of f well in excess of $T = 255$. The modified estimator only affects the estimate for values of r corresponding to where there are censored pixels, which in Figure 8 is true for r close to 0. We also note that by reflecting the data around 0, both estimates of f has zero derivative at the origin.

7 Discussion

We have proposed a method for estimating the centres of spherical particles which achieves sub-pixel accuracy, even when the appearance of the particles changes with depth. The simulation study confirmed that the estimated standard errors are

accurate, provided that the assumptions of rotational symmetry and isotropy of the noise process are correct. We therefore believe it to be a versatile tool for automatic measurements in video microscopy of colloidal suspensions.

Estimation of \hat{f} using the algorithm for censored regression in §6 by minimising equation (10) is computationally costly because of censored pixels. Therefore our recommendation is to use the method in §3, ignoring censored pixels, when estimating the particle centre in the plane. A compromise could be to find a preliminary estimate as above and then update this position using the censored version of the local quadratic kernel estimate. Nevertheless, for the degrees of censoring present in the data considered here (up to 5% of the pixels in the neighbourhood), simulation studies has shown that the difference in precision of the particle centre estimate between the two methods is not statistically significant at the 5% level. Therefore we believe that the more elaborate method is not justified.

Apart from estimating the diffusion coefficient of colloidal particles, their interaction is also of interest. Unfortunately, this cannot be estimated from 2-D locations, but also requires depth information. In Kvarnström (2005b), the intensity profiles of colloidal particles at different known depths are estimated using the method in §6. These intensity profiles are then stored as templates for particle appearance at different depths. By comparing the appearance of a particle at an unknown depth with these profiles, the depth of the particle can be estimated. It is of course crucial that the intensity profiles used as templates are estimated taking into account censored pixels.

Kvarnström (2005b) demonstrates how the methodology presented here can be used for tracking of colloidal particles observed in video microscopy. For each image, the previous positions of the particles are used as candidate particle positions. Also, by suitably adjusting the neighbourhoods \mathcal{N} of each particle, according to where the other particles in the image are located, position estimation is possible even when particles occlude each other. Finally the diffusion coefficient can be estimated from the particle centres.

Acknowledgements MK's work was supported by funds from the Swedish Research Council. CAG's work was supported by funds from the Scottish Executive Environment and Rural Affairs Department. We thank Mats Rudemo and Aila Särkkä for encouraging and guiding this work and Lennart Lindfors for supplying the data.

References

- Crocker, J. and Grier, D. (1996). Methods of digital video microscopy for colloidal studies. *Journal of Colloid and Interface Science*, 179:298–310.
- Elliot, M. S. and Poon, W. C. K. (2001). Conventional optical microscopy of colloidal suspensions. *Advances in Colloid and Interface Science*, 92:133–194.
- Evans, D. F. and Wennerström, H. (1999). *The Colloidal Domain. Where Physics, Chemistry, Biology, and Technology meet*. Wiley-VCH, New York, second edition.
- Fan, J. and Gijbels, I. (1996). *Local Polynomial Modelling and Its Applications*. Chapman and Hall, London.
- Glasbey, C. A., Hitchcock, D. A., Russel, J. F., and H., R. (1994). Towards the automatic measurement of cashmere-fibre diameter by image analysis. *Journal of the Textile Institute*, 85:301–307.
- Grier, D. (1998). Colloids: A surprisingly attractive couple. *Nature*, 393:621.
- Han, Y. and Grier, D. (2005). Configurational temperatures and interactions in charge-stabilized colloid. *Journal of Chemical Physics*, 122:064907 (14 pages).
- Hastie, T. and Tibshirani, R. (1990). *Generalized Additive Models*. Chapman and Hall, London.
- Hecht, E. (2002). *Optics*. Addison-Wesley, San Francisco, California, fourth edition.
- Kvarnström, M. (2005a). Estimation of the diffusion coefficient in a mixture model. *Kwantitatieve Methoden*, 72:1–23.
- Kvarnström, M. (2005b). *Position estimation and tracking in colloidal particle microscopy*. PhD thesis, Mathematical Sciences, Chalmers University of Technology, Sweden, <http://www.math.chalmers.se/Stat/Research/Preprints/>.
- Ovryn, B. and Izen, S. (2000). Imaging of transparent spheres through a planar interface using a high-numerical-aperture optical microscope. *Journal of the Optical Society of America A. Optics, Image Science and Vision*, 17:1202–1213.
- Owen, A. (2001). *Empirical Likelihood*. Chapman and Hall/CRC, London.