

# Error analysis for small angle neutron scattering datasets using Bayesian inference

Charles R. Hogg\*, Joseph B. Kadane†, Jong Soo Lee‡ and Sara A. Majetich§

**Abstract.** We present a Bayesian methodology for extracting correlation lengths from small-angle neutron scattering (SANS) experiments. For demonstration, we apply the technique to data from a previous paper, which investigated the presence of dipolar ferromagnetism in assemblies of ferromagnetic Co nanoparticles. Bayesian analysis confirms the presence of multiparticle dipolar domains even at zero magnetic field, but higher-field correlation lengths were found to be much smaller than previously believed, yielding new information on the maximum lengthscale which the instrument can reliably probe. We use two complementary types of graph to visualize the results. Plots of standardized residual distributions show quality of fit, and guide model refinement. These principles can be applied to other types of sample, and even to other small-angle scattering techniques.

## 1 Introduction

A bar magnet is perhaps the most familiar magnetic object. As illustrated in Figure 1, if placed in a magnetic field, the magnet rotates to line up with that field. Once aligned, it moves in the direction where the field increases most rapidly. It also generates its own magnetic field, whose pattern is shown in Figure 1(b). The bar magnet is a good example of a more general class of magnetic objects, called magnetic dipoles (Figure 1(c)), which behave in this way.

Despite their familiarity, important questions remain about their behaviour when large numbers of magnetic dipoles interact. One key goal is to learn whether, in regular assemblies of dipoles, large regions spontaneously order to share the same direction, a phenomenon known as dipolar ferromagnetism (Luttinger and Tisza 1946). Though a vague sense that “magnets prefer to align” might make this seem trivial, Figure 2 shows that this is far from true. Dipoles lined up along the axis indeed prefer to align, but adjacent dipoles prefer to *anti*-align, and the preference for dipoles at other angles can be anywhere in-between.

Monodomain magnetic nanoparticles are an ideal test system to investigate this phenomenon. Each nanoparticle consists of aligned atomic dipoles so that it behaves effectively as a single giant dipole (Stoner and Wohlfarth 1948). Additionally, their

---

\*Department of Physics, Carnegie Mellon University, Pittsburgh, PA, <mailto:chogg@andrew.cmu.edu>

†Department of Statistics, Carnegie Mellon University, Pittsburgh, PA, <mailto:kadane@stat.cmu.edu>

‡Department of Statistics, Carnegie Mellon University, Pittsburgh, PA and Department of FREC, University of Delaware, Newark, DE, <mailto:jslee@udel.edu>

§Department of Physics, Carnegie Mellon University, Pittsburgh, PA., <mailto:sara@cmu.edu>

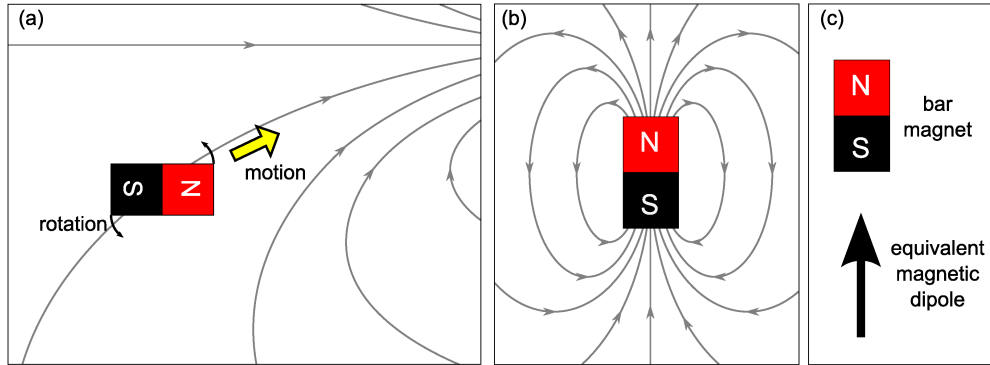


Figure 1: Basic concepts relating to bar magnets and magnetic dipoles. (a) When placed in a magnetic field (represented by the grey lines), the magnet rotates until the lines point from the south to the north pole. It also moves in the direction where the lines are densest, which is where the field is strongest. (b) The magnetic field of a bar magnet. (c) A bar magnet and its equivalent magnetic dipole representation: an arrow pointing from the south to the north pole.




Figure 2: Demonstration that the preferred orientation of one dipole (B) with respect to another (A) depends on the angle between them. Dipole A is shown in the lower left, along with its magnetic field. Along A's dipole axis (position B1), the dipoles preferentially align. Perpendicular to this axis (position B2), B prefers to *anti-align* with A. At intermediate angles (e.g. position B3), the preferential alignment may be anywhere in-between.

highly regular size allows them to self-assemble into ordered two- and three-dimensional structures (Murray et al. 1995; Talapin et al. 2001; Narayanan et al. 2004), as shown in Figure 3.




Figure 3: A picture of a thin cobalt magnetic nanoparticle crystal, taken with a transmission electron microscope (TEM). This particular crystal is only one monolayer thick, so that individual particles are easily distinguished.

The homogeneous regions where nanoparticles are magnetized in the same direction are called *dipolar domains* (Yamamoto et al. 2008), and the size of these domains is the main quantity we are interested in. An example is given in Figure 4. Although both samples have zero average magnetization, the domains in (a) are twice as large as domains in (b). Instruments which measure only the total magnetization, called magnetometers, could not distinguish between these samples. However, scattering techniques can determine the average domain size.




Figure 4: Two hypothetical magnetic samples, identical except for the domain structure. A magnetometer would measure both to be demagnetized, but the correlation length in (a) is twice as large as in (b).

The geometry of a general scattering experiment is shown in Figure 5. The sample is placed in a beam of radiation, whose component particles<sup>1</sup> it deflects (or *scatters*)

<sup>1</sup>All radiation is composed of particles, according to the “particle-wave duality” in quantum me-

at an angle  $\theta$ . The scattering of each individual particle is a stochastic event, whose governing distribution is determined by the size of the ordered regions within the sample. When many particles have been scattered, the intensity pattern built up on the detector can be analyzed using theoretical models, yielding information about the parameters characterizing the sample.




Figure 5: Schematic of a general transmission-geometry scattering experiment. (a) The incident beam comes in from the left and strikes the sample, which deflects the radiation at an angle  $\theta$  towards a two-dimensional detector. Here, a single scattering event is shown, but the detector records the total amount of radiation scattered into all areas, as indexed by the scattering vector  $\vec{Q}$ . (b) The definition of  $\vec{Q}$ : the wavevector  $\vec{k}$  has the direction of the incident neutron, and magnitude  $2\pi/\lambda$ , with  $\lambda$  the wavelength of the radiation. (c) A sample pattern of scattered radiation which might be observed on a detector.

The connection between characteristic sizes and scattered radiation can be briefly illustrated by explaining Bragg's Law (Kittel 2004), the most basic scattering relation. Atoms in a crystal are arranged in planes, each of which reflects a small amount<sup>2</sup> of the incoming radiation, as in Figure 6. These reflected rays undergo *interference* when they recombine at the detector, where a high intensity signal occurs only if they all have the same phase. Rays reflected from deeper planes must travel a correspondingly longer distance,  $\Delta L = 2d \sin \theta$  (see Figure 6); accordingly, their phase is more advanced by an amount  $2\pi\Delta L/\lambda$ , with  $\lambda$  the wavelength of the radiation. Since a phase difference of  $2\pi$  makes no difference, the first bright spot occurs when  $\Delta L = \lambda$ , i.e.

$$\lambda = 2d \sin \theta. \quad (1)$$

Equation 1 shows the connection between a measured size,  $d$ , and the angle  $\theta$  where radiation is strongly scattered. For many scattering experiments,  $\lambda$  is held constant, and the scattered intensity is measured as a function of  $\theta$  (Hammouda 2008). Since the right side of Equation 1 must *also* be constant, structures with larger sizes  $d$  must scatter at smaller angles  $\theta$ . Though many systems exist whose scattering is not described

chanics.

<sup>2</sup>Since only a small fraction of the incident radiation is scattered, we neglect higher-order corrections accounting for multiple scattering.

by Bragg’s law, this inverse relationship between sizes and scattering angles is quite generally true, as Figure 7 shows schematically.




Figure 6: An illustration of Bragg’s Law for scattering. The incident beam comes in from the left and is partially reflected from each plane of atoms. The path lengths of rays 1 and 2 differ by the amount  $\Delta L$ , as shown in the figure.

## 1.1 Correlation lengths

The sizes measured by Bragg’s Law correspond to distances between nearest neighbours. By contrast, we’re interested in measuring *regions* of order, which often extend across many neighbours. The size of these ordered regions is measured by a quantity known as the “correlation length”. The term, ‘correlation,’ is well-established in both the physics and statistics communities, but the meanings differ.

In physics, correlation always refers to systems possessing some kind of order, and the *correlation length*  $\xi$  measures how large the ordered regions tend to be (Yeomans 1992). In the case of our dipolar domains, the ordering is found in the normalized<sup>3</sup> dipole orientation  $\vec{S}$  on different nanoparticles. This ordering is measured as a function of separation by means of the *correlation function*,  $G(r)$ :

$$G(r) = \langle \vec{S}_i \cdot \vec{S}_j \rangle_r \quad (2)$$

where the angle brackets denote averaging over all pairs of nanoparticles,  $i$  and  $j$ , separated by a distance  $r$ .

The behavior of  $G(r)$  relates to  $\xi$  as follows. When  $r \ll \xi$ , most pairs of locations separated by  $r$  are within the same domain, so  $G(r) \sim 1$ . On the other hand, when the separation  $r$  is large ( $r \gg \xi$ ), the magnetization in remote regions is just as likely to point one way as the opposite way, so  $G(r) \rightarrow 0$ . The correlation length  $\xi$  therefore sets

<sup>3</sup>By “normalized,” we mean that  $|\vec{S}| = 1$ .

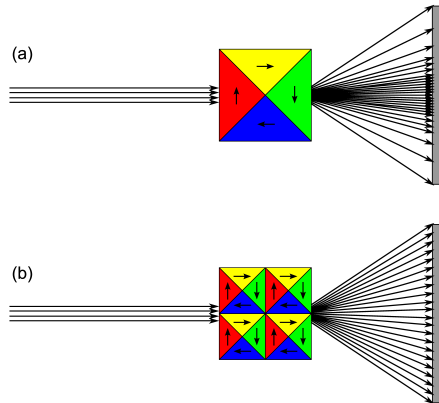


Figure 7: An illustration of the relationship between the scattering angle  $\theta$  and the size of ordered regions. In both cases, the sample is placed in the path of a neutron beam, and it scatters the neutrons into a detector. The sample in (a) has larger ordered regions than (b), and consequently scatters proportionately more neutrons at *smaller* angles.

the scale of how quickly  $G(r)$  goes to zero, and it can be obtained by fitting to the rate of damping of  $G(r)$ .

Some systems, such as ours, require more than one correlation length for their complete description (Bernhoeft 1999). Within each dipolar domain, smaller regions of inhomogeneities (Figure 8) are found, where clusters of nanoparticles deviate slightly from the average domain orientation (Michels et al. 2003). We denote the “domain size” correlation length as  $\xi_S$ , and the “inhomogeneity” correlation length as  $\xi_L$ . These correlation lengths are the parameters of greatest interest.

## 1.2 Small-angle neutron scattering

Neutrons have several key advantages for studies of condensed matter, particularly of magnetic materials (Hammouda 2008; Squires 1997). They are uncharged, which allows them to penetrate very deeply through most matter. “Slow” neutrons — those travelling less than roughly  $v = 1$  km/s — have wavelengths<sup>4</sup>,

$$\lambda = \frac{h}{mv}, \quad (3)$$

suitable for scattering studies on structures of current interest (typically a few tens of nm, where 1 nm is one billionth of a meter). Finally, each neutron also possesses a magnetic moment, which enables it to interact magnetically with the sample.

In small-angle neutron scattering (SANS), theory predicts the relative scattered

<sup>4</sup>Here,  $m$  is the mass of the neutron, and  $h$  is Planck’s constant.

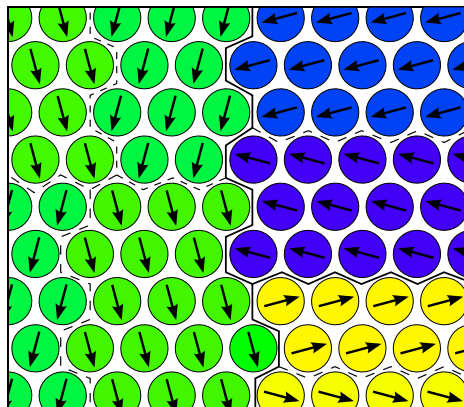


Figure 8: An example configuration of dipoles in a nanoparticle array. The dipole orientation is shown both by the central arrow and the color of the nanoparticle. Here, three dipolar domains are shown, separated by thick solid lines: the downward domain occupies the left half, and the right half supports a leftward domain on top and a rightward domain on the bottom. Within each domain are shown smaller regions of inhomogeneities, bounded by thin dashed lines, where subgroups of nanoparticles deviate slightly from the average direction in the domain.

intensity as a function of the scattering vector  $\vec{Q}$ , which was defined in Figure 5(b). SANS scattering tends to be elastic, meaning that  $|\vec{k}'| = |\vec{k}|$ , and  $|\vec{Q}|$  can be evaluated based purely on geometry:

$$|\vec{Q}| = \frac{4\pi \sin(\theta/2)}{\lambda} \approx \frac{2\pi\theta}{\lambda} \quad (4)$$

A key result of scattering theory is that the intensity  $I(\vec{Q})$  is simply related to  $G(\vec{r})$  by a Fourier transform (Van Hove 1954). This shows the connection between  $I(\vec{Q})$  (the quantity measured experimentally) and  $\xi$  (the quantity of interest).

The two types of regions discussed in Section 1.1 each give rise to a distinctive type of scattering. The correlation function for dipolar domain-like ordering decays like  $\exp(-r/\xi_S)/r$  (Sachan et al. 2008), and its Fourier transform has a Lorentzian-Squared (“S”) lineshape:

$$I_S(Q) \propto \frac{1}{(\kappa_S^2 + Q^2)^2}, \quad (5)$$

where we define the parameter  $\kappa_S = 1/\xi_S$  for convenience. Similarly, the inhomogeneities yield a correlation function like (Sachan et al. 2008)  $\exp(-r/\xi_L)$ , which Fourier transforms to a Lorentzian (“L”) lineshape:

$$I_L(Q) \propto \frac{1}{\kappa_L^2 + Q^2}, \quad (6)$$

again defining  $\kappa_L = 1/\xi_L$ . Adding in appropriate scaling constants  $\Sigma_S$  and  $\Sigma_L$ , the total magnetic scattering at low  $Q$  becomes

$$I(Q) = \frac{\Sigma_S}{(\kappa_S^2 + Q_i^2)^2} + \frac{\Sigma_L}{\kappa_L^2 + Q_i^2}. \quad (7)$$

Changes in the ratio of  $\Sigma_L$  to  $\Sigma_S$ , from one run to another, indicate changes in the relative prominence of domains versus the inhomogeneities within them.

## 2 The SANS Instrument

We now describe the SANS instrument, shown in Figure 9, in greater detail. The workings of each major component are briefly explained, with emphasis on their contributions to uncertainty. Readers desiring a more in-depth description are referred to (Glinka et al. 1998).

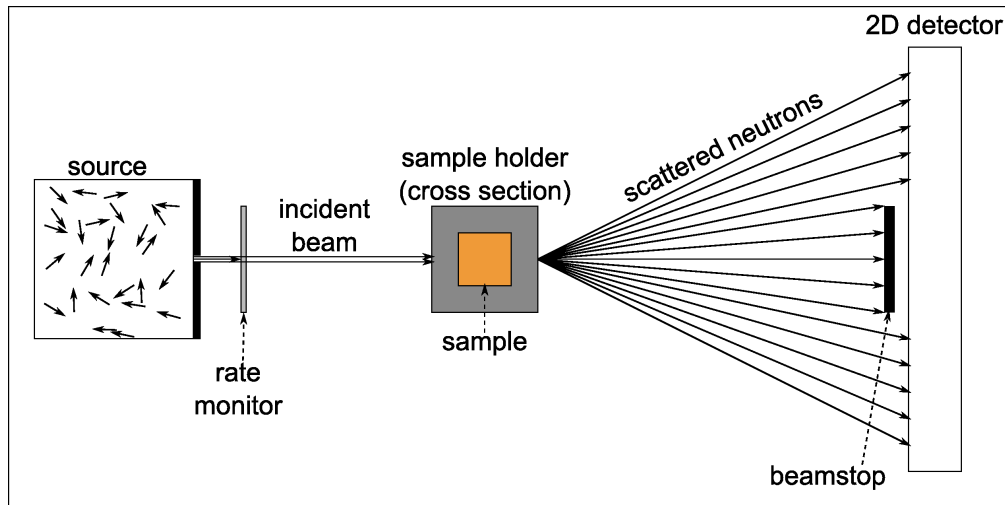


Figure 9: The SANS instrument. Here, we show an abstracted view for simplicity, which includes only the conceptual parts necessary to understand our model.

### 2.1 Instrument Components

#### Neutron source

The *neutron source* as we model it has two main stages: production, and moderation. Production takes place in a nuclear fission reactor, where neutrons are liberated during the splitting of heavy atomic nuclei. Moderation sets the wavelength distribution of the resulting beam, both by changing the speeds<sup>5</sup> of the neutrons, and by filtering out

<sup>5</sup>From Equation 3, the wavelength of the neutron is directly related to its speed.



neutrons of certain speeds. The end result is that the beam has some measurable wavelength distribution  $\varphi(\lambda)$  (Hammouda 2008, Chapter 12), where  $\varphi(\lambda) d\lambda$  is proportional to the amount of neutrons in the beam having wavelengths between  $\lambda$  and  $\lambda + d\lambda$ . Typically, this distribution is reported in terms of the peak wavelength  $\lambda_+$ , and a relative wavelength spread  $\sigma_\lambda/\lambda_+$ .

### Rate Monitor

The output from the source varies, necessitating a *rate monitor* to measure it. A thin  $^{235}\text{U}$  plate, which has a low probability to capture each passing neutron, is inserted immediately after the source. Captured neutrons are counted for a predetermined time, and the totals are recorded in a datafile. The monitor averages roughly  $10^6$  counts per minute.

### Sample

The sample must be held in the path of the beam by a sample holder. The neutrons scattered from this assembly constitute the measured signal. Though the holder is necessary, its presence complicates data interpretation, as discussed in Section 2.2.

### Beam Stop

Because the transmitted beam is very intense, and would quickly damage the detector, a beam stop is inserted which blocks the central portion of the detector. This unfortunately renders the lowest- $Q$  range inaccessible, but is necessary to protect the detector. Datapoints close to the beamstop edge should be viewed with suspicion, and are typically discarded.

### Neutron Detector

The two-dimensional neutron detector consists of a  $128 \times 128$ -pixel grid, where each pixel records the cumulative number of neutrons passing by its position. The grid is comprised of horizontal and vertical wires, immersed in a high-pressure mixture of  $^3\text{He}$  and  $\text{CF}_4$  gases. Passing neutrons have some probability to ionize this gas, and the resulting charge is detected on the nearest wires: one horizontal, one vertical. Their intersection defines the location of a pixel, whose counter is then incremented by one.

Each pixel  $i$  detects only some fraction of the neutrons which impinge upon it. This fraction, known as the *detection efficiency*  $\eta_i$ , must be carefully measured before the data can be quantitatively analyzed. Typically, this is done by staff scientists on a regular basis, and the latest measured efficiencies are distributed to users.

The expected number of counts also depends on the solid angle  $\beta_i$  which the pixel covers, i.e. the apparent size of the pixel as viewed from the sample. Since all pixels have the same area, solid angle is primarily determined by the distance from the sample

to the detector. We define the detection *capacity*  $\gamma_i \equiv \eta_i \beta_i$  as the product of detection *efficiency* and solid angle.

We point out that each pixel corresponds to a given scattering angle  $\theta$ , determined by its displacement from the center and the detector’s distance from the sample. However, the neutron count from each pixel is recorded at a specific  $\vec{Q}$ -value. Equation 4 shows that mapping  $\theta$  onto  $Q$  is unique only if the wavelength  $\lambda$  is precisely defined. The fact that the wavelength spread  $\sigma_\lambda/\lambda_+$  is nonzero means that each pixel actually contains a *probability distribution* of  $Q$ -values, an effect known as smearing.

## 2.2 Instrument configurations

The sample signal is but one of three contributions to the measured  $I(\vec{Q})$ . By reconfiguring the instrument, the other two can be measured. We show these configurations schematically in Figure 10, explaining below the contributions they account for.

The experimental room contains additional sources of neutrons which are independent of the beam configuration, known collectively as the “background”. This background rate is measured by blocking the beam completely (Figure 10(a), “BGR”) and counting neutrons. An empty sample holder (Figure 10(b), “EMP”) is next inserted, letting us account for neutrons scattered by the holder. Finally, the sample is added inside the holder (Figure 10(c), “SAM”), contributing both nuclear and magnetic scattering to the measured signal.

The sample transmits undisturbed only some fraction of neutrons, and the holder’s contribution is smaller by this amount. This fraction can be measured in *transmission mode*, which only counts the neutrons in the transmitted beam: it is the ratio of the intensity with the sample in the holder, to that with the sample removed.

### Magnetic scattering

Accounting for the above effects permits isolation of the sample signal, which is the sum of nuclear and magnetic scattering. To separate them, note that magnetic scattering along the local magnetization direction is always zero (Squires 1997, chap. 7). By applying a magnetic field large enough to saturate the sample, we obtain a purely nuclear signal along the field direction. Since nuclear scattering is isotropic, the remaining magnetic signal can be extracted in all directions.

## 3 SANS error analysis

This paper presents an alternative technique for error analysis of SANS data. We will give an overview of the traditional method, before describing how we have applied Bayesian techniques to SANS.

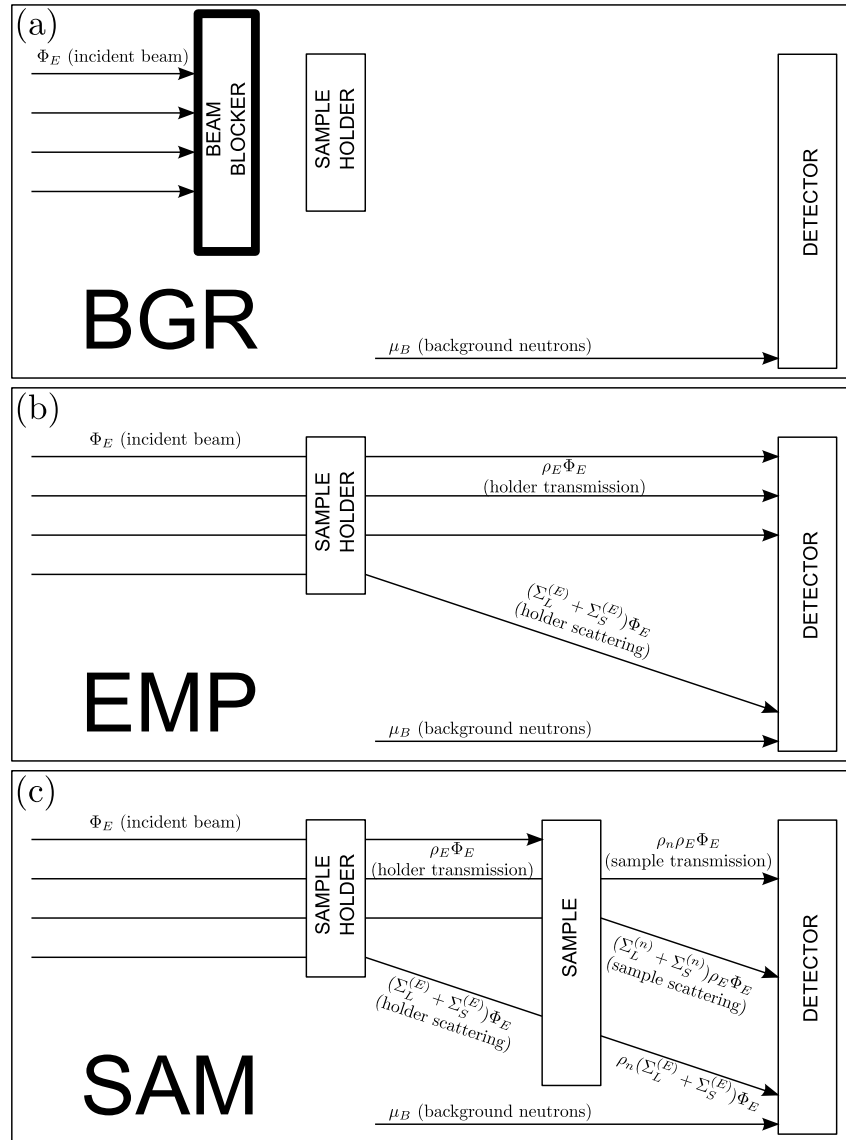


Figure 10: A schematic showing the three contributions to the signal, and how they are measured by including them one by one. (a) The beam is blocked, and only the background contributes. (b) Scattering and transmission are measured for the empty sample holder. (c) The sample is inserted into the holder, both attenuating the holder's signal and adding its own contribution. Neutrons scattered by both sample and holder are neglected in the single-scattering approximation. (We emphasize that this is a *schematic*; the paths show the history of the neutrons, and are *not* intended to represent trajectories.)

### 3.1 Traditional error analysis

Existing tools for SANS error analysis have relied upon traditional statistical methods. These tools are mature and refined through years of heavy use with a variety of users studying diverse types of systems. We illustrate using the package written by Steven Kline at NIST (Kline 2006), where we performed our experiments.

Two separate steps are involved in analysis of SANS data: reduction, and fitting. The former involves using the configurations described in Section 2.2 to compensate for the undesired contributions. This procedure is inherently tied to the specific instrument where the data was taken: reduction procedures are not generic. Traditional fitting procedures, by contrast, can be quite generic, since the *reduced* data is expected to have all major instrument-specific effects accounted for.

Figure 11 shows a diagram of the traditional reduction process. We applied it to our data, and then converted the resulting 2-D function of  $\vec{Q}$  to a 1-D function of  $Q$ , by averaging a narrow range of angles perpendicular to the applied magnetic field. We measured the nuclear signal as described in Section 2.2, and subtracted it off to yield the magnetic data, which we fit to the sum of a Lorentzian and Lorentzian-squared.

#### Data fitting

The “best fit” is traditionally decided by minimizing the  $\chi^2$  per degree of freedom, as follows. Each datapoint  $y_i$  is assigned a standard error  $s_{y_i}$ , and associated with a coordinate  $x_i$ . The  $y_i$  are assumed to be based on a model function  $f(x)$ , such that the deviations  $(f(x_i) - y_i)$  should be small compared to  $s_{y_i}$ . These deviations are called *standardized residuals* when normalized by  $s_{y_i}$ .  $\chi^2$  is the sum of squares of these standardized residuals,

$$\chi^2 = \sum_{i=1}^N \left( \frac{f(x_i) - y_i}{s_{y_i}} \right)^2, \quad (8)$$

and the curve which minimizes  $\chi^2$  is taken to be the “best fit” for a given model. The model itself may be checked by examining the standardized residuals and looking for trends. Figure 12 shows the results of applying this method to our data at  $H = 0$ ; the  $y_i$  are shown in black, and  $f(x)$  is the green curve.

### 3.2 The Bayesian alternative

Bayesian analysis is an alternative methodology; (Lindley 2006) gives a very readable general introduction, and (Agostini 2003) introduces Bayesian methods from a physicist’s perspective.

We distinguish two types of quantities. *Data* (denoted by Latin letters) are directly observed in the experiment, and *parameters* (denoted by Greek letters) are unobserved quantities which affect the distribution of data. The sole reason for collecting the former is to learn about the latter. We use  $\vec{y}$  as shorthand for the set of all data values, and  $\vec{\alpha}$

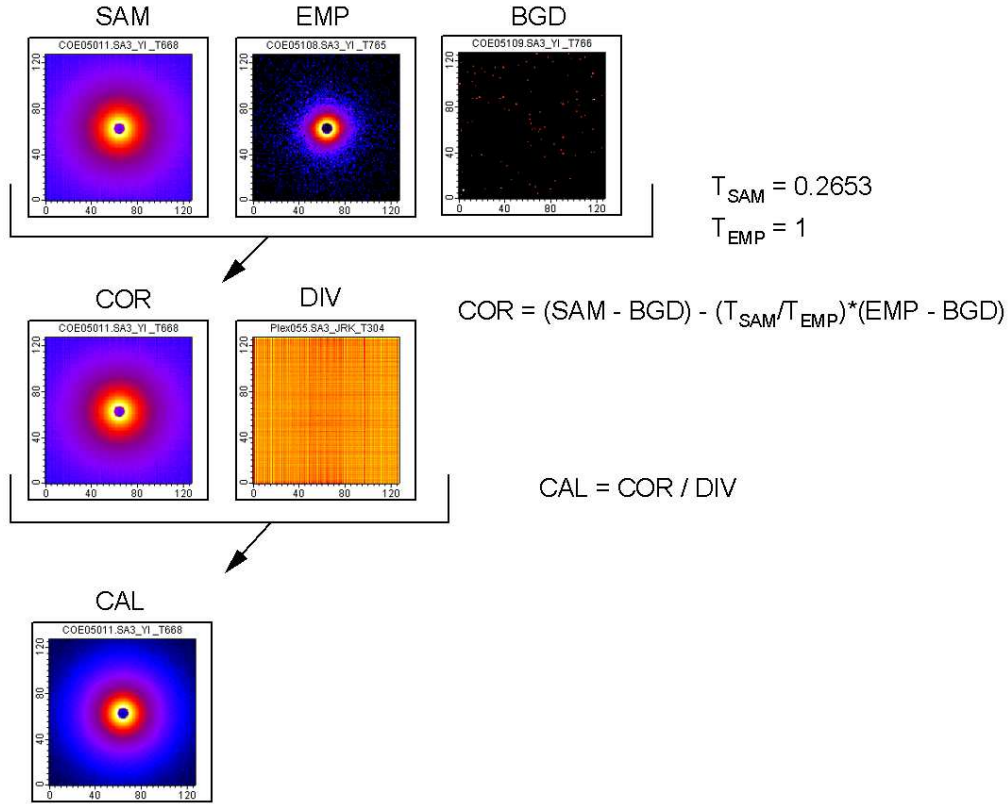


Figure 11: The datafiles used in our Bayesian analysis, processed here using traditional analysis. Note the horizontal and vertical lines visible in COR, which are removed by the DIV correction.

for the set of all parameter values.

The information about  $\vec{\alpha}$  yielded by  $\vec{y}$  is encoded in the posterior probability distribution  $P(\vec{\alpha}|\vec{y})$ . Bayesian analysis relates this to two other distributions through Bayes' rule,

$$P(\vec{\alpha}|\vec{y}) = \frac{P(\vec{y}|\vec{\alpha})P(\vec{\alpha})}{\int P(\vec{y}|\vec{\alpha}')P(\vec{\alpha}') d\vec{\alpha}'} \quad (9)$$

The likelihood,  $P(\vec{y}|\vec{\alpha})$ , gives the probability of obtaining the observed data, given a particular set of values for the parameters. The prior,  $P(\vec{\alpha})$ , summarizes all knowledge of the sample before the experiment was performed. These two distributions constitute the model. When combined with the observed data  $\vec{y}_0$ , they contain all statistical information about the parameters  $\vec{\alpha}$ .

Calculations using  $P(\vec{\alpha}|\vec{y})$  involve integrals in a high-dimensional space (i.e. parameter space). We performed these integrals using a Markov chain Monte Carlo (MCMC),




Figure 12: The results of traditional fitting methods applied to data from this experiment (Sachan et al. 2008, Figure 3). The red Lorentzian-squared curve corresponds to scattering from the *dipolar domains*, and the blue Lorentzian corresponds to magnetic inhomogeneity scattering. Standardized residuals plotted at bottom show that the line of best fit agrees with most datapoints within standard errors.

written in the R programming language (R Development Core Team 2008). The resulting chain of MCMC steps was analyzed using the `boa` package, short for Bayesian Output Analysis (Smith 2007). Both the code and the datafiles used in our analysis are available online (Hogg 2009). We ran for  $10^6$  steps at each field, then used `boa` to discard steps until the remaining chain represented equilibrium. The final MCMC chains contained 4973 steps at  $H = 0$  T, 2813 steps at  $H = 0.2$  T, and 10987 steps at  $H = 1.0$  T. Execution took roughly 4 days on our computer cluster, but recoding in C++ could yield a significant advantage in speed.

We now describe our model in greater detail.

Abbreviation	Name	Purpose
BGR	Background	Measure background neutron rate
EMP	Empty holder	Measure sample holder scattering
SAM-n	Sample (high field) scattering	Measure the nuclear (i.e. nonmagnetic) contribution to the scattering signal
<i>SAM</i>	<i>Sample (low field) scattering</i>	<i>Contains the signal of greatest interest</i>
EMP-t	Empty cell transmission	Measured sample transmission ratio at different fields
SAM-n-t	Sample (high field) transmission	
SAM-t	Sample (low field) transmission	

Table 1: A list of the instrument configurations.

## Data

Most data analyzed in this paper comes directly as a raw number of neutron counts in some detector. Associated with instrument configuration X, we have a count  $N_X^i$  for each pixel  $i$  of the detector, and a monitor count  $M_X$ . The different configurations are listed in Table 1.

The only data not in the form of a raw number of neutron counts is the detector efficiency data (Section 2.1). The efficiency  $e_i$  of the  $i$ th pixel is based on raw data from an isotropic scatterer, but the final values have been processed by instrument scientists. We note in passing that the posterior distribution on the detector efficiencies could be measured once by instrument scientists, and distributed to users desiring to use Bayesian analysis.

## Parameters

Our model requires 16 parameters to describe our system. Of these, the magnetic correlation lengths  $\xi_S$  and  $\xi_L$  are the most important. A complete list of parameters is given in Table 2.

Several additional parameters would be required for complete rigor, but have been approximated as constants for this preliminary version. The relative detection efficiency  $\eta_i$  of pixel  $i$  is not known with certainty, but assumed to be equal to the *measured* efficiency  $e_i$ . Similarly, the instantaneous reactor output during datafile  $v$ ,  $\Omega_v$ , is assumed to be equal to the number of monitor counts  $M_v$ . In both cases, the error introduced by these simplifications is negligible compared with the uncertainty from other sources: using Poisson statistics, it is less than 0.01 for  $\eta$ , and less than 0.001 for  $\Omega$ . Nevertheless, it is important to verify with each new experiment that the counts are high enough to make these errors negligible compared to other sources of uncertainty.

Name	Description
$\sigma_G$	Spread of the transmitted beam
$\mu_B$	Normalized, per-pixel mean rate of background neutrons
$\Phi_E$	Normalized rate of incident neutrons reaching the sample holder, which are either scattered or transmitted
$\rho_E$	Fraction of neutrons which the sample holder transmits
$\rho_n$	The transmission of the sample in high field
$\rho$	The transmission of the sample in low field
$\Sigma_L^{(E)}$	Fraction of neutrons which the sample holder scatters with a Lorentzian signal
$\Sigma_L^{(n)}$	Fraction of neutrons undergoing nuclear Lorentzian scattering
$\Sigma_S^{(n)}$	Fraction of neutrons undergoing nuclear Lorentzian-squared scattering
$\Sigma_L$	Fraction of neutrons undergoing magnetic Lorentzian scattering
$\Sigma_S$	Fraction of neutrons undergoing magnetic Lorentzian-squared scattering
$\kappa_E$	Inverse correlation length for Lorentzian scattering from sample holder
$\xi_L^{(n)}$	Nuclear Lorentzian correlation length
$\xi_S^{(n)}$	Nuclear Lorentzian-squared correlation length
$\xi_L$	Magnetic Lorentzian correlation length (average size of magnetic inhomogeneities)
$\xi_S$	Magnetic Lorentzian-squared correlation length (average size of dipolar domains)

Table 2: A list of parameters in our model, and a brief description of the role of each



## Likelihood

We make the assumption that the data are *conditionally* independent from one another, *given* the parameters. It should be noted that this is equivalent to assuming that there are no unaccounted-for parameters that influence multiple pixels. Under this assumption, the likelihood can be decomposed into a product of likelihood functions for each individual neutron count — whether from the rate monitor, or a pixel in the two-dimensional detector.

We can simplify considerably by grouping pixels receiving the same number of incident neutrons per solid angle. Because of this assumption, each such group may be treated as a giant pixel  $P$ , whose detection capacity  $\gamma_P$  (Section 2.1) is the sum  $\sum_{i \in P} \gamma_i$  of capacities of the pixels which comprise it. For example, the background neutron rate is modeled as independent of location on the detector, so we can replace the separate counts from all pixels by the single count  $N_{\text{BGR}} = \sum_{i=1}^{N_{\text{pixel}}} N_{\text{BGR}}^{(i)}$ . Another example is the signal at the bottom of Figure 11, which only depends on the pixel's distance  $Q$  from the center; here, our groups would be the rings of pixels which have the same  $Q$ . Subject to these reductions, the data we consider are

$$\vec{y} = \left\{ N_{\text{BGR}}; \vec{N}_{\text{EMP}}; \vec{N}_{\text{SAM-n}}; \vec{N}_{\text{SAM}}; \vec{N}_{\text{EMP-t}}; \vec{N}_{\text{SAM-n-t}}; \vec{N}_{\text{SAM-t}}; \right. \\ \left. M_{\text{BGR}}; M_{\text{EMP}}; M_{\text{SAM-n}}; M_{\text{SAM}}; M_{\text{EMP-t}}; M_{\text{SAM-n-t}}; M_{\text{SAM-t}}; \vec{e} \right\} \quad (10)$$

where  $\vec{N}_{\text{EMP}}$  is shorthand for the set of all grouped neutron counts in the EMP configuration, and similarly for  $\vec{N}_{\text{SAM}}$  and  $\vec{N}_{\text{SAM-n}}$ .

With this notation, we use our assumption of conditional independence to write the explicit form of the likelihood:

$$P(\vec{y}|\vec{\alpha}) = [P(N_{\text{BGR}}|\vec{\alpha})] \times [P(N_{\text{EMP}}|\vec{\alpha})] \times [P(N_{\text{SAM-n}}|\vec{\alpha})] \times [P(N_{\text{SAM}}|\vec{\alpha})] \times \\ [P(N_{\text{EMP-t}}|\vec{\alpha})] \times [P(N_{\text{SAM-n-t}}|\vec{\alpha})] \times [P(N_{\text{SAM-t}}|\vec{\alpha})] \quad (11)$$

Each factor in Equation 11 corresponds to one of the seven datafiles used in fitting. Factors corresponding to the measured detector efficiencies  $\vec{e}$ , or to any of the  $M_v$ , are missing because we have approximated them as constants.

Each neutron count  $N$  is modeled with an underlying Poisson distribution (Hengartner 2008), with mean  $\nu$ :

$$N|\nu \sim \frac{\nu^N e^{-\nu}}{N!} \quad (12)$$

Here,  $\nu$  is proportional to the *normalized* mean  $\mu$ , but also to the detection efficiency  $\eta$ , the solid angle  $\beta$  which the pixel covers, and reactor output  $\Omega$ , i.e.  $\nu = \mu\eta\beta\Omega$ . The sum of independent Poisson-distributed random variables is another Poisson distribution, with the aggregate mean given by the sum of the individual means, but  $\eta$ ,  $\beta$ , and  $\Omega$  are the same for neutrons of all sources. Accordingly, the factor corresponding to each datafile  $v$  can be described simply by giving the form of this aggregate mean,  $\mu_v$ .

The functional forms for the  $\mu_v$  are complicated by an additional source of experimental uncertainty, known as *instrumental smearing*. As discussed more fully in

Appendix 5, neutrons at a given pixel correspond not to a single  $Q$ -value, but to a distribution of  $Q$ -values. It is the *unsmear*ed mean  $\bar{\mu}_v$  which we model, but the *smear*ed mean  $\bar{\mu}_v^{(\text{Sm})}$  which we measure, where the bar indicates that background neutrons are excluded because they are not smeared, i.e.

$$\mu_v = \mu_B + \bar{\mu}_v^{(\text{Sm})}. \quad (13)$$

We now give the associated unsmear

ed mean  $\bar{\mu}$  for each of the seven factors listed in Equation 11. The corresponding likelihood factor is given by:

$$P(N|\mu(\vec{\alpha}), \eta, \beta, \Omega) = \frac{(\mu(\vec{\alpha})\eta\beta\Omega)^N \exp[-\mu(\vec{\alpha})\eta\beta\Omega]}{N!}, \quad (14)$$

here, the functions  $L(Q; \kappa)$ ,  $S(Q; \kappa)$ , and  $G(Q)$  are defined precisely in Appendix 5:

$$\begin{aligned} \bar{\mu}_{\text{BGR}} &= 0 \\ \bar{\mu}_{\text{EMP}} &= \Phi_E \left[ \Sigma_L^{(E)} L(Q; \kappa_E) + \Sigma_S^{(E)} S(Q; 1/L_t) \right] \\ \bar{\mu}_{\text{SAM-n}} &= \Phi_E \left\{ \rho_n \left[ \Sigma_L^{(E)} L(Q; \kappa_E) + \Sigma_S^{(E)} S(Q; 1/L_t) \right] + \right. \\ &\quad \left. \rho_E \left[ \Sigma_L^{(n)} L(Q; 1/\xi_L^{(n)}) + \Sigma_S^{(n)} S(Q; 1/\xi_S^{(n)}) \right] \right\} \\ \bar{\mu}_{\text{SAM}} &= \Phi_E \left\{ \rho \left[ \Sigma_L^{(E)} L(Q; \kappa_E) + \Sigma_S^{(E)} S(Q; 1/L_t) \right] + \right. \\ &\quad \left. \rho_E \left[ \Sigma_L^{(n)} L(Q; 1/\xi_L^{(n)}) + \Sigma_S^{(n)} S(Q; 1/\xi_S^{(n)}) + \Sigma_L L(Q; 1/\xi_L) + \Sigma_S S(Q; 1/\xi_S) \right] \right\} \\ \bar{\mu}_{\text{EMP-t}} &= \Phi_E \rho_{\text{att}} \rho_E G(Q) \\ \bar{\mu}_{\text{SAM-n-t}} &= \Phi_E \rho_{\text{att}} \rho_E \rho_n G(Q) \\ \bar{\mu}_{\text{SAM-t}} &= \Phi_E \rho_{\text{att}} \rho_E \rho G(Q). \end{aligned} \quad (15)$$

## Priors

We decompose the parameters into disjoint independent subsets. The functional form of the prior on each of these subsets is given, along with a brief justification of why we believe its parameters are independent from all other parameters.

## Background neutrons

Background neutrons are completely described in our model by the parameter  $\mu_B$ , the mean number of background neutrons arriving at each pixel per monitor count. We do not expect this rate to be affected by the experimental setup in any way, so the assumption of independence is well-justified. We turn to previous runs to elicit a prior: for each BGR file, the total number of neutrons detected, divided by the monitor counts for that file, gives an approximation for  $\mu_B$ . We can calculate the sample mean and variance of  $\mu_B$  values obtained from several such runs, and use these as the mean and

Filename	BGR Neutron Counts	Monitor Counts	Count time	$\mu_B$ Estimate
AUG07154.ASC	1486	$23.91 \times 10^6$	20 min.	$37.9 \times 10^{-10}$
AUG07156.ASC	8094	$23.92 \times 10^6$	20 min.	$206.5 \times 10^{-10}$
JUN08007.ASC	738	$6.02 \times 10^6$	5 min.	$74.8 \times 10^{-10}$

Table 3: Total neutron counts, compared to monitor counts, for blocked-beam files from different runs. The first two were taken on the same day with the detector at different locations, and the third was taken almost a year later. Variation arises because the background depends on the location of the detector in the room, and also on what other neutron experiments are being run at the same time. Despite agreeing only to within an order of magnitude, the background is overall a very small effect, and minimally affects the parameters of greatest interest.

variance for our prior distribution. The specific form chosen is a Gamma distribution, since its domain is the same as for  $\mu_B$ :  $(0, \infty)$ . The results are shown in Table 3. The rightmost column was obtained by dividing the total number of neutron counts by the monitor counts for that datafile, then further dividing by the number of pixels (i.e.  $2^{14}$ ). The mean and standard deviation for our prior distribution are estimated using the sample mean and sample standard deviation of these values:

$$s_{\text{BGR}}^{\text{old}} = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{x})^2}. \quad (16)$$

Values for  $\mu_{\text{BGR}}^{\text{old}}$  and  $s_{\text{BGR}}^{\text{old}}$ , along with all other values needed to describe our priors, are found in Table 4.

### Beam spread

Our transmitted beam is fit to a Gaussian form (Hammouda 2008). Since the beam is centred around  $Q = 0$ , only the spread  $\sigma_G$  is needed to characterize its shape. We expect that the transmitted beam will be insignificant outside the beamstop, whose edge  $Q_b$  should therefore be at least  $2\sigma_G$  from the center. We use a Gaussian (normal) prior having mean  $Q_b/2$  and standard deviation  $Q_b/4$ .

Independence is justified because the intrinsic spread of the beam does not depend on the sample holder, sample, or background rate.

### Incident non-absorbed flux

We define  $\Phi_E$  as the rate of neutrons incident on the sample holder, considering only neutrons which contribute to some measured signal in some way. Other incident neutrons are either absorbed or scattered at wide angles; since they have no effect on any measured signal, we exclude them from  $\Phi_E$ . We have not performed an absolute cali-

bration in this experiment, so we do not know the *magnitude* of  $\Phi_E$ . To express this, we choose a prior which is uniform in log-space, subject to a cutoff  $X$  representing the maximum order of magnitude we will probe:

$$P(\Phi_E) = \frac{1}{\Phi_E} \Theta(X - |\log_{10}(\Phi_E)|), \quad (17)$$

where the theta-function  $\Theta(x)$  is 1 if its argument is positive and zero otherwise.

All other parameters relating to the sample or holder are expressed as *fractions* of the flux incident upon them; hence, we may treat  $\Phi_E$  as an independently adjustable measure of this flux.

### Partitioning of neutrons for sample holder alone

Consider all neutrons incident on the bare sample holder which are either scattered at small angles or transmitted. The relative fractions transmitted and scattered are intrinsic properties of the holder, and hence independent of all other parameters. Calling the fraction transmitted  $\rho_E$ , the fraction undergoing Lorentzian scattering  $\Sigma_L^{(E)}$ , and the fraction undergoing Lorentzian-squared scattering  $\Sigma_S^{(E)}$ , we have

$$\rho_E + \Sigma_L^{(E)} + \Sigma_S^{(E)} = 1, \quad (18)$$

since we have disregarded all other neutrons. Because this constraint leaves only two free parameters, we only explicitly specify  $\rho_E$  and  $\Sigma_L^{(E)}$ . Our prior on these parameters is a Dirichlet distribution, governed by the hyperparameters  $A_\rho^{(E)}$ ,  $A_{\Sigma_L}^{(E)}$ , and  $A_{\Sigma_S}^{(E)}$ .

### Partitioning of neutrons for sample (both high- and low-field)

These parameters partition the neutrons incident on the sample according to the type of interaction they experience (i.e. magnetic Lorentzian-squared scattering, undisturbed transmission, nuclear Lorentzian scattering, etc.). We divide them into three groups: nuclear scattering, SAM, and SAM-n. Our strategy is to first specify the prior for the nuclear scattering, which contributes to both SAM and SAM-n configurations. Parameters in the remaining two groups are independent of each other, *given* values for the nuclear scattering parameters, so we can specify these two groups separately. All of these parameters are intrinsic properties of the sample, justifying our assumption of independence from all other parameters.

The prior for the nuclear scattering parameters is Dirichlet, governed by  $A_{\Sigma_L}^{(n)}$ ,  $A_{\Sigma_S}^{(n)}$ , and  $A^{(n)}$ . The remaining two priors are *scaled* Dirichlets, on the domain  $(0, F)$ , where we define  $F = 1 - (\Sigma_L^{(n)} + \Sigma_S^{(n)})$  for convenience. Conservation of neutrons is expressed for SAM-n as

$$\Sigma_L^{(n)} + \Sigma_S^{(n)} + \rho_n < 1, \quad (19)$$

and for SAM as

$$\Sigma_L^{(n)} + \Sigma_S^{(n)} + \rho + \Sigma_L + \Sigma_S < 1. \quad (20)$$

The total prior for these sample partitioning parameters is thus

$$\begin{aligned}
 P(\Sigma_L^{(n)}, \Sigma_S^{(n)}, \rho_n, \Sigma_L, \Sigma_S, \rho) = & D(\Sigma_L^{(n)}, \Sigma_S^{(n)}, F; A_{\Sigma_L}^{(n)}, A_{\Sigma_S}^{(n)}, A^{(n)}) \times \\
 & D(\rho_n/F, 1 - \rho_n/F; A_\rho^{(n)}, A_\delta^{(n)}) \times \\
 & D(\Sigma_L/F, \Sigma_S/F, \rho/F; A_{\Sigma_L}, A_{\Sigma_S}, A_\rho).
 \end{aligned} \tag{21}$$

### Sample holder correlation lengths

We now turn our attention to correlation length-describing parameters. The normalization described in Appendix 5 means that each of these should be *a priori* independent of its corresponding  $\Sigma$ , and we have only to consider prior knowledge on any possible relationship among the correlation lengths.

The prior for  $\kappa_E$  should be uniform in  $\xi_L^{(E)} = 1/\kappa_E$ , since any correlation length is as likely as any other; hence,

$$P(\kappa_E) = \frac{1}{\kappa_E^2}. \tag{22}$$

We simulate the inverse length  $\kappa_E$  because that is what appears directly in  $L(Q)$ .

The sample holder also has a Lorentzian-squared signal which is important at lower  $Q$ . Because the holder is made from highly crystalline aluminum, we expect this correlation length to be longer than this SANS instrument can probe. This limitation is the transverse coherence length  $L_t$ . Setting  $\xi_S^{(E)} = L_t$  saves computation time without compromising accuracy, as we verified by checking the fits at low  $Q$ .

### Nuclear correlation lengths

The nuclear correlation lengths  $\xi_S^{(n)}$  and  $\xi_L^{(n)}$  describe the positional arrangement of the elements in our sample. The nuclear data is known to be well-fit by the sum of a Lorentzian and Lorentzian-squared, and that like any other correlation lengths they are constrained to be less than  $L_t$ . Accordingly, the prior for each is uniform within the region  $(0, L_t)$ .

### Magnetic correlation lengths

The magnetic correlation lengths are similar to their nuclear counterparts, except that we identify  $\xi_L$  as the size of magnetic inhomogeneities within dipolar domains of size  $\xi_S$  (Michels and Weissmuller 2008). This implies that  $\xi_S > \xi_L$  always. Additionally, because our nanoparticles are single-domain, the domains can never be smaller than the diameter  $d_{NP}$  of a single nanoparticle. This allowed region roughly has the shape of a triangle with the bottom corner snipped off, as shown in Figure 13. We choose our prior to have uniform probability density inside this region and zero outside.

These are the parameters of greatest interest.

Category	Hyperparameter	Value	Description
Experimental setup	$Q_b$	$0.04 \text{ nm}^{-1}$	$Q$ -value of the edge of the beamstop
	$\lambda$	$0.5 \text{ nm}$	Mean wavelength of incident neutrons
	$L_t$	$1000 \text{ nm}$	Transverse coherence length of neutron beam
	$\rho_{\text{att}}$	$0.0003$	Transmission ratio of attenuator (inserted during transmission mode)
	$d_{\text{NP}}$	$8 \text{ nm}$	Diameter of a single Co nanoparticle
Priors	$\mu_{\text{BGR}}^{\text{old}}$	$106 \times 10^{-10}$	Prior mean of $\mu_B$
	$s_{\text{BGR}}^{\text{old}}$	$89 \times 10^{-10}$	Prior standard deviation of $\mu_B$
	$X$	10	Cutoff order of magnitude for $\Phi_E$
	$A_{\Sigma_L}^{(E)}$	2	Dirichlet portion for Lorentzian scattering from sample holder
	$A_{\Sigma_S}^{(E)}$	2	Dirichlet portion for Lorentzian-squared scattering from sample holder
	$A_\rho^{(E)}$	16	Dirichlet portion for transmission through sample holder
	$A_{\Sigma_L}^{(n)}$	2	Dirichlet portion for nuclear Lorentzian scattering from sample
	$A_{\Sigma_S}^{(n)}$	2	Dirichlet portion for nuclear Lorentzian-squared scattering from sample
	$A^{(n)}$	6	Dirichlet portion for neutrons not nuclearly scattered
	$A_{\Sigma_L}$	3	Dirichlet portion for magnetic Lorentzian scattering from sample
	$A_{\Sigma_S}$	3	Dirichlet portion for magnetic Lorentzian-squared scattering from sample
	$A_\rho$	4	Dirichlet portion for neutrons transmitted by sample in low field
	$A_\delta$	2	Dirichlet portion for remaining neutrons in low field (absorbed, high-angle scattered, etc.)
	$A_\rho^{(n)}$	4	Dirichlet portion for neutrons transmitted by sample in high field
$A_\delta^{(n)}$	8	Dirichlet portion for remaining neutrons in high field (absorbed, high-angle scattered, etc.)	

Table 4: List of hyperparameters which characterize our experimental setup or govern the shape of our prior distributions, along with the values used in our analysis.




Figure 13: Probability distribution for the length scales of the Lorentzian and Lorentzian-squared terms,  $\xi_L$  and  $\xi_S$ .  $L_t$  is the transverse coherence length of the beam, here 1000 nm; correlation lengths larger than this cannot be distinguished by the instrument.  $D_{\text{particle}}$  is the diameter of a single particle, here 8 nm. The probability distribution is flat within the shaded region, and zero outside.

## 4 Results and Discussion

The data we fit are identical to those in (Sachan et al. 2008). The same datafiles were used in both cases. Pixels have been grouped according to sector averaging within a  $\pm 5^\circ$  range around the given direction: horizontal for the pure nuclear signal, and vertical for the mixed nuclear-plus-magnetic signal. These pixels were subsequently binned to match the  $Q$ -values from (Kline 2006), with the center of each  $Q$ -bin placed at the corresponding  $Q$ -value from the traditional analysis.

The general picture which emerges is qualitative reproduction of general trends, with significant quantitative differences. For direct comparison, Bayesian analysis can easily reproduce the traditional style of uncertainty reporting, where the parameter mean is given along with standard deviations. However, the richness of information available in  $P(\vec{\alpha}|\vec{y})$  enables novel forms of presentation, capable of conveying deep insight into uncertainty and correlation at a glance.

Among the parameters listed in Table 2,  $\xi_S$  and  $\xi_L$  are the two of overwhelmingly greatest interest. Henceforth, our discussion concerns not the full posterior  $P(\vec{\alpha}|\vec{y})$ , but the marginal posterior  $P(\xi_L, \xi_S|\vec{y})$ , with all other parameters integrated out.

### 4.1 Separate marginal distributions on $\xi_L$ and $\xi_S$

Table 1 of (Sachan et al. 2008) gave values for four different field configurations. Because the final two gave very similar results, we focus on the first three. We calculated uncertainty estimates of one standard deviation based on our MCMC results. For ease

of comparison, the results from the original paper (“T.A.”, for “traditional analysis”), are presented alongside the updated uncertainties (“B.A.”, for “Bayesian analysis”), in Table 5.

Parameter	Method	$H = 0$	$H=0.2$ T	$H=1.0$ T
$\xi_L/\text{nm}$	T.A.	$37 \pm 8$	$9.1 \pm 0.3$	$1.4 \pm 0.1$
	B.A.	$50 \pm 10$	$11.8 \pm 0.4$	$1.68 \pm 0.04$
$\xi_S/\text{nm}$	T.A.	$72 \pm 9$	$91 \pm 4$	1000
	B.A.	$64 \pm 1$	$102 \pm 10$	$504 \pm 34$

Table 5: Traditional fit results (“T.A.”) presented alongside new Bayesian results (“B.A.”). Bayesian uncertainty estimates represent one standard deviation of the posterior distribution.

The agreement is best at  $H = 0$ , which is the configuration of greatest interest for proving the existence of domains. Here the Bayesian results overlap the traditionally obtained values, but with significantly smaller uncertainty for  $\xi_S$ . The fact that  $\xi_S(H = 0) \gg d_{\text{NP}}$  indicates magnetic correlations extending over multiple particles, providing strong supporting evidence for the presence of dipolar domains.

At higher fields, the agreement is more qualitative, and only the general trends in  $\xi_S(H)$  are reproduced. In particular, correlations in the apparently-saturated sample are closer to 500 nm than the nominal limit of 1000 nm. This shorter correlation length may indicate that the sample is not fully saturated, but we feel this is unlikely, because (Sachan et al. 2008) shows negligible change from 1 T to 5 T. A more likely explanation is that the limit of what SANS can probe (i.e. the transverse coherence length  $L_t$  is lower than we previously believed.

## 4.2 Joint posterior distribution on $(\xi_L, \xi_S)$

A plot of the full joint posterior distribution  $P(\xi_L, \xi_S | \vec{y})$  is perhaps the most informative way to present data. We used the `hexbin` package (Carr et al. 2008), without smoothing, to estimate the posterior density by counting the number of MCMC steps within each bin. The results are shown in Figure 14.

The most probable regions are darker. Moreover, correlations which are not straightforward to obtain in traditional analysis are readily apparent in these plots. For instance, there is a high positive correlation between  $\xi_S$  and  $\xi_L$  at 0.2 T: in this regime, larger domains are particularly likely to be found containing larger inhomogeneous regions. This qualitative visual observation is supplemented by Table 6, which lists the calculated correlation between the parameters at each field.

If we were to represent *traditional* uncertainty graphically, summarizing the two separate means and standard deviations, it might take the form of a box, with side lengths given by the standard deviation. Depending on the number of standard deviations included, we would be more or less “confident” to find the true parameter values within






Figure 14: The joint posterior distribution  $P(\xi_L, \xi_S | \vec{y})$ , plotted for (a)  $H=0$ , (b)  $H=0.2$  T, (c)  $H = 1.0$  T. The red  $\times$  marks the location of the best fit according to traditional analysis, and the box represents the region within one standard deviation of the traditional fit. The diagonal line in (a) represents the constraint  $\xi_S > \xi_L$ ; recall that the prior is uniform above and to the left of this line. Note that in (c), the red bar has been moved down from its actual location of  $\xi_S = 1000$  nm, for ease of visualization.

	$H = 0$	$H=0.2$ T	$H=1.0$ T
$\text{Cor}(\xi_L, \xi_S)$	0.384	0.781	0.031

Table 6: Measured correlations between  $\xi_L$  and  $\xi_S$  at different applied field strengths.

the box, and more or less surprised to find them outside. It should be noted, however, that the confidence bounds in the traditional analysis are frequency proportions of a procedure, implemented in only a single instance. Thus their correct interpretation is that, “repeating the procedure used to calculate these bounds many times, 68% of the time the bounds calculated as these were would contain the parameter value.” Additionally, reporting parameters separately from one another always implies an assumption of independence, and this traditional method therefore cannot capture the correlations which are so evident in, say, Figure 14b.

Note that Bayesian analysis has done more than simply refine the uncertainty estimates. As Figure 14 starkly shows, it has shifted the estimates of the parameters, to the point that the traditionally obtained values do not overlap for  $H > 0$ . In part, this may be due to neglecting high- $Q$  datafiles in the present work:  $\xi_L$  is always shifted in the same direction, and the missing datafiles contain a  $Q$ -range which strongly constrains the Lorentzian. But the main reason they differ is that a proper accounting of uncertainty can have profound effects on the parameter values extracted from experimental data.

### 4.3 Representing fits graphically: “family of curves” and standardized residual distributions

Though the joint posterior plots of Figure 14 are informative, they give no indication of how well these parameters fit the data, only that they fit better than the others which were explored. We consider the traditional method of representing variation in graphs, then present an attractive Bayesian alternative.

Traditionally, datapoints are plotted along with error bars of one standard deviation, which represent the variation one might expect if the experiment were repeated numerous times. Automated fitting routines then determine the “best fit” values of the parameters, and the curve described by these parameters is plotted with the data, hopefully passing nearer each datapoint than its associated error bar. Two things are peculiar here. First, plotting a single curve gives no indication when other parameter values might fit just as well. Second, associating uncertainty with data is misplaced, since it inevitably refers to quantities which *might* have been observed, but in fact were not. (If the experiment is repeated, all repetitions may be analyzed as part of the data. If not, then statements about uncertainty should depend only on what was *actually* observed.)

Bayesian analysis assigns no uncertainty to datapoints. Instead, variation may be shown by taking a random subset of MCMC steps, plotting the curves corresponding to each, and plotting the datapoints on top. As a bonus, correlations between *datapoints* are automatically accounted for. (By contrast, two adjacent datapoints with errorbars implies they are just as likely to vary oppositely as together, which is generally untrue.)

In this case, we see that the shape of the curve is quite tightly constrained, giving no visual indication of the degree to which the parameters vary. Significant magnification is required before the curves can be distinguished, suggesting the need for alternative plots to show goodness of fit.

Standardized residuals fit this role, and can also show all MCMC steps simultaneously, like the joint posterior plots of Figure 14. The residual  $\Delta_i^s$  for pixel  $i$  at MCMC step  $s$  is  $(N_i - \nu_i^s)$ , where  $N_i$  is the observed number of neutrons, and  $\nu_i^s$  the number expected based on the parameter values at  $s$ . Considering all MCMC steps leads to a residual *distribution* at each  $Q$ , shown in Figure 16.

The residual distributions at all fields exhibit wide variation at low  $Q$ , and become narrow at higher  $Q$ . This trend is a consequence of the higher counts at low  $Q$ . The variation expected from Poisson statistics,  $\sqrt{\nu_i^s}$ , is plotted in the middle row of Figure 16. Dividing each residual (top) by the expected variation (middle) yields the *standardized residuals* (bottom),  $\delta_i^s = (N_i - \nu_i^s)/\sqrt{\nu_i^s}$ . These plots are the closest Bayesian analogue to the bottom of Figure 12, but we caution against *direct* comparison, because different data are being fit: traditional analysis processes the data first, while Bayesian analysis models all contributions and fits the unaltered data.

The zero-field data lie within  $2\sigma$  of most fits, but display clear systematic trends as a function of  $Q$ . We attribute these trends to the extra contributions, and suggest

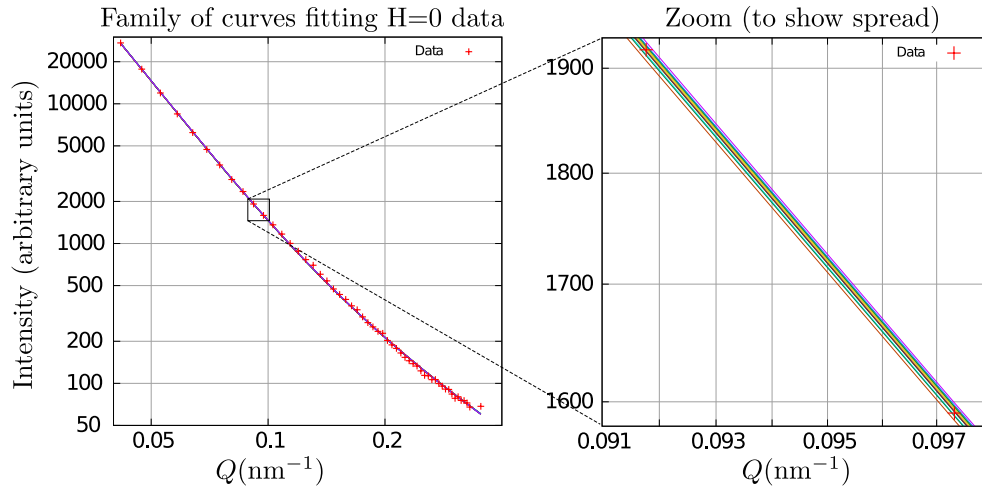


Figure 15: The “family of curves” corresponding to a randomly chosen subset of 16 MCMC steps from the chain. In this case, the difference between curves is indistinct when the plot is fully zoomed out, so a central portion has been expanded to show the variation. (Note that in contrast to Figure 12, these data include contributions from nuclear scattering, background, and the sample holder, making direct comparison difficult.)

that refining the models for nuclear and holder scattering may mitigate these effects. At intermediate field ( $H = 0.2$  T), only the last several points disagree significantly. The fits systematically and increasingly underestimate the data, but since data near the detector edges are inherently less trustworthy, we expect the inclusion of higher- $Q$  datafiles to improve agreement in this region. The datapoints at the highest field ( $H = 1.0$  T) agree well all around, generally within  $1\sigma$  of most fits. Note that at all fields, the residual distributions widen at lowest  $Q$ , even *after* standardization.

## 5 Conclusions

In conclusion, we have applied Bayesian analysis to data from a recently published SANS paper (Sachan et al. 2008). Our analysis supports the main conclusions, that dipolar domains existed in a Co colloidal crystal even at zero field, and grew larger as higher fields were applied. Our new approach has put our uncertainty analysis on firmer footing and yielded quantitative results.

More broadly, small-angle neutron scattering datasets can be analyzed using a Bayesian approach, which yields the conditional posterior probability  $P(\vec{\alpha}|\vec{y})$ . Correlations between parameters can be made obvious by a glance at a plot of this posterior distribution. Furthermore, since the MCMC explores all of parameter space, it often




Figure 16: Residual distributions at each field. Each distribution is binned, with darker bins having higher density, and black representing the most populous bin for each plot. The top row shows a distribution at each  $Q$  of non-normalized residuals, i.e. the difference between the number of neutrons observed and the number expected for each MCMC step. The expected variation is plotted in the middle row as the square root of the fit function at each step. (Discontinuities correspond to changes in the number of grouped pixels.) The bottom row plots distributions on *standardized* residuals, equivalent to dividing the top row by the middle row.

finds possibilities in obscure regions, which traditional analysis might easily miss. Modeling the system in detail requires eliciting reasonable priors for the parameters, and often reveals important details about the system or key gaps in knowledge even before the MCMC is run. Bayesian analysis is more computationally intensive than conventional forms of data fitting, due to the need to run a Monte Carlo simulation. Finally, the lack of a mature, versatile analysis package constitutes an additional barrier to scientists who might otherwise make use of Bayesian analysis. We hope this work will serve as the first step to the creation of just such a tool.

## Appendix A: Normalization conventions

Our  $\Sigma$  parameters are intended to represent the total scattering from a given functional form (i.e. Lorentzian or Lorentzian-squared). It is therefore important to normalize these functional forms, much more so than for traditional fitting, where no model gives meaning to the magnitudes of the relative scaling factors.

We begin by considering the Lorentzian function; the discussion for the Lorentzian-squared is similar. When we say that the Lorentzian  $L(Q; \kappa)$  should be *normalized*, we mean that it should represent the stochastic scattering pattern for an *individual* neutron. In other words, since the neutron must be scattered with some  $\vec{Q}$ , integration over all allowed  $\vec{Q}$ -values should yield unity:

$$\int_R L(Q; \kappa) d\vec{Q} = 1, \quad (23)$$

where  $R$  stands for the region of integration.

This equation is required to hold for all  $\kappa$ . The integral of the non-normalized Lorentzian,

$$\tilde{L}(Q; \kappa) = \frac{1}{Q^2 + \kappa^2}, \quad (24)$$

typically depends on  $\kappa$ ; hence, the normalization factor must be some function of  $\kappa$  and of the region  $R$  of integration, i.e.

$$L(Q; \kappa) = F_L(\kappa; R) \tilde{L}(Q; \kappa) = \frac{F_L(\kappa; R)}{Q^2 + \kappa^2}, \quad (25)$$

where  $F_L$  is the normalizing factor for the Lorentzian function.

Since SANS is an elastic technique, we normalize over all elastically scattered vectors. This region takes the form of a sphere of radius  $2\pi/\lambda$ , whose center is displaced from the origin by this same amount. For radiation of wavelength  $\lambda$  scattered at an angle  $\theta$ , we have

$$Q_{\text{elastic}} = \frac{4\pi \sin(\theta/2)}{\lambda}. \quad (26)$$

In order to turn our integral from  $d\vec{Q}$  into  $dQ$ , we must weight the integrand by the amount of  $\vec{Q}$ -space available at each value of  $Q$ . The amount of  $\vec{Q}$ -space available

on the sphere at an angle  $\theta$  is proportional to  $(4\pi^2/\lambda)\sin\theta$ . Solving the geometry to express everything in terms of  $Q$ , we find that the  $Q$ -dependent weighting factor is

$$W(Q) = 2\pi Q \sqrt{1 - \left(\frac{Q\lambda}{4\pi}\right)^2}, \quad (27)$$

Note that this reduces to  $2\pi Q$  at small angles, as it must, before the curvature of the sphere distorts this factor.

Putting it all together, we find

$$\int_R L(Q; \kappa) d\vec{Q} = \int_0^{4\pi/\lambda} \frac{W(Q)F_L(\kappa; R)}{Q^2 + \kappa^2} dQ = 1, \quad (28)$$

which implies that

$$F_L(\kappa; R) = \left[ \int_0^{4\pi/\lambda} \frac{W(Q)dQ}{Q^2 + \kappa^2} \right]^{-1}. \quad (29)$$

After carrying out this integration, along with a similar one for the Lorentzian-squared  $S(Q; \kappa)$ , we find the following normalized forms of the functions, defining the dimensionless variable  $\tilde{\kappa} = \kappa\lambda/4\pi$  for convenience:

$$\begin{aligned} L(Q; \kappa) &= \left( \frac{1}{2\pi [(1 + \tilde{\kappa}^2) \operatorname{arcsinh}(\tilde{\kappa}^{-1}) - 1]} \right) \frac{1}{Q^2 + \kappa^2} \\ S(Q; \kappa) &= \left( \frac{8\pi\sqrt{1 + \tilde{\kappa}^2}}{\lambda^2 [\tilde{\kappa}^{-2}\sqrt{1 + \tilde{\kappa}^2} - \operatorname{arcsinh}(\tilde{\kappa}^{-1})]} \right) \frac{1}{(Q^2 + \kappa^2)^2} \end{aligned} \quad (30)$$

It is these forms which we have used in our simulation.

## Appendix B: Smearing corrections

Equation 4 makes the connection between the scattering vector  $\vec{Q}$ , used in theory, and the deflection angle  $\theta$  which is experimentally measured. This connection is one-to-one as long as only neutrons of a single wavelength are used. However, no single wavelength gives enough neutrons to yield a detectable signal in a reasonable amount of time; in other words, SANS is a *flux-limited* technique. In practice, we are forced to allow a distribution of wavelengths, which means that each deflection angle  $\theta$  corresponds to a *distribution* of  $Q$ -values. This effect is known as smearing, and it distorts the measured signal.

Traditional analysis uses one of two methods to account for smearing. The first is to “desmear” the data, by attempting to invert the function. Desmearing has the advantage that it only needs to be performed once, thus saving computation time. However, smearing is not strictly invertible.

The alternative is to smear the model function before fitting it to the measured data. Smearing at every step incurs a considerable computational cost, but leaves the experimental data inviolate. This latter option is very commonly done in traditional analysis; in Bayesian analysis, it is the *only* option.

The technique for smearing correction is outlined in (Kline 2006). The fit function,  $I(Q)$ , is a continuous function determined by the values of its associated fit parameters (i.e.  $\xi_L, \xi_S, \dots$ ). To calculate the *smear*ed intensity  $I_s(Q_i)$ , at the  $i$ th  $Q$ -value, one takes a weighted average of the unsmeared intensity,

$$I_s(Q_i) = \int_0^{\infty} R(Q, Q_i) I(Q) dQ \quad (31)$$

where  $R(Q, Q_i)$  is the resolution function of the instrument at the point  $Q_i$ . This resolution function is well approximated as a Gaussian

$$R(Q, Q_i) \equiv \frac{f_s}{(2\pi\sigma_Q^2)^{1/2}} \exp\left[-\frac{(Q - \bar{Q})^2}{2\sigma_Q^2}\right], \quad (32)$$

characterized by a mean  $\bar{Q}$ , standard deviation  $\sigma_Q$ , and a parameter  $f_s$  which represents the fraction not shadowed by the beamstop (typically,  $f_s = 1$  for all but the lowest- $Q$  points). Each of these parameters varies with the nominal  $Q$  (i.e. the angle  $\theta$ ), and these values are measured by the instrument scientists and provided to the users.

To improve execution time, we have replaced the integral in Equation 31 by a sampled Riemannian sum. A series of  $Q$ -values is generated, consisting of the original  $Q$ -values plus  $F$  more between each pair. The unsmeared fit function is evaluated at each  $Q$ -value in this new series. To calculate the *smear*ed intensity at the original  $Q$ -values, a cutoff  $C$  is supplied by the user, such that only points between  $(\bar{Q}_i - C\sigma_Q)$  and  $(\bar{Q}_i + C\sigma_Q)$  are considered. These points are averaged according to

$$I_s(Q_i) = K_i \sum_{j: |Q_j - \bar{Q}_i| < C\sigma_Q} I(Q_j) \exp\left[-\frac{(Q_j - \bar{Q})^2}{2\sigma_Q^2}\right], \quad (33)$$

where the normalizing factor

$$K_i = f_s \left[ \sum_{j: |Q_j - \bar{Q}_i| < C\sigma_Q} \exp\left[-\frac{(Q_j - \bar{Q})^2}{2\sigma_Q^2}\right] \right]^{-1} \quad (34)$$

ensures that the weights sum to  $f_s$ .

The results presented in this paper correspond to  $F = 5$  and  $C = 3$ .

## References

Agostini, G. D. (2003). *Bayesian Reasoning in Data Analysis: A Critical Introduction*. World Scientific Publishing Company. 12

- Bernhoeft, N. (1999). “Geometrical effects in diffraction analysis.” *Acta Crystallographica. Section A, Foundations of Crystallography*, 55(Part 2 Number 2): 274–288. [6](#)
- Carr, D., ported by Nicholas Lewin-Koh, and Maechler, M. (2008). *hexbin: Hexagonal Binning Routines*. R package version 1.18.0. [24](#)
- Glinka, C. J., Barker, J. G., Hammouda, B., Krueger, S., Moyer, J. J., and Orts, W. J. (1998). “The 30 m Small-Angle Neutron Scattering Instruments at the National Institute of Standards and Technology.” *J. Appl. Cryst.*, 31(3): 430–445. [8](#)
- Hammouda, B. (2008). “The SANS Toolbox.” [http://www.ncnr.nist.gov/staff/hammouda/the\\_SANS\\_toolbox.pdf](http://www.ncnr.nist.gov/staff/hammouda/the_SANS_toolbox.pdf). [4](#), [6](#), [9](#), [19](#)
- Hengartner, N. (2008). “Statistical Analysis of Neutron Scattering Experiments.” 5th LANSCE Neutron School 2008. [17](#)
- Hogg, C. R. (2009). “Bayesian SANS MCMC.” <http://lib.stat.cmu.edu/modules.php?op=modload&name=Downloads&file=index&req=viewdownload&details&lid=931&tttitle=SANS>. [14](#)
- Kittel, C. (2004). *Introduction to Solid State Physics*. Wiley, 8 edition. [4](#)
- Kline, S. R. (2006). “Reduction and analysis of SANS and USANS data using IGOR Pro.” *Journal of applied crystallography*, 39: 895–900. [12](#), [23](#), [31](#)
- Lindley, D. V. (2006). *Understanding Uncertainty*. Wiley-Interscience, 1 edition. [12](#)
- Luttinger, J. M. and Tisza, L. (1946). “Theory of Dipole Interaction in Crystals.” *Physical Review*, 70(11-12): 954. [1](#)
- Michels, A., Viswanath, R. N., Barker, J. G., Birringer, R., and Weissmüller, J. (2003). “Range of Magnetic Correlations in Nanocrystalline Soft Magnets.” *Physical Review Letters*, 91(26): 267204. [6](#)
- Michels, A. and Weissmüller, J. (2008). “Magnetic-field-dependent small-angle neutron scattering on random anisotropy ferromagnets.” *Reports on Progress in Physics*, 71(6): 066501. [21](#)
- Murray, C. B., Kagan, C. R., and Bawendi, M. G. (1995). “Self-Organization of CdSe Nanocrystallites into Three-Dimensional Quantum Dot Superlattices.” *Science*, 270(5240): 1335–1338. [3](#)
- Narayanan, S., Wang, J., and Lin, X. M. (2004). “Dynamical Self-Assembly of Nanocrystal Superlattices during Colloidal Droplet Evaporation by in situ Small Angle X-Ray Scattering.” *Physical Review Letters*, 93(13): 135503. [3](#)
- R Development Core Team (2008). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.  
URL <http://www.R-project.org> [14](#)



- Sachan, M., Bonnoit, C., Majetich, S. A., Ijiri, Y., Mensah-Bonsu, P. O., Borchers, J. A., and Rhyne, J. J. (2008). “Field evolution of magnetic correlation lengths in epsilon-Co nanoparticle assemblies.” *Applied Physics Letters*, 92(15): 152503–3. 7, 14, 23, 24, 27
- Smith, B. J. (2007). “boa: An R Package for MCMC Output Convergence Assessment and Posterior Inference.” *Journal of Statistical Software*, 21(11): 1–37. 14
- Squires, G. L. (1997). *Introduction to the Theory of Thermal Neutron Scattering*. Dover Publications. 6, 10
- Stoner, E. C. and Wohlfarth, E. P. (1948). “A mechanism of magnetic hysteresis in heterogeneous alloys.” *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences (1934-1990)*, 240(826): 599–642. 1
- Talapin, D. V., Shevchenko, E. V., Kornowski, A., Gaponik, N., Haase, M., Rogach, A. L., and Weller, H. (2001). “A New Approach to Crystallization of CdSe Nanoparticles into Ordered Three-Dimensional Superlattices.” *Advanced Materials*, 13(24): 1868. 3
- Van Hove, L. (1954). “Correlations in Space and Time and Born Approximation Scattering in Systems of Interacting Particles.” *Physical Review*, 95(1): 249.  
URL <http://link.aps.org/abstract/PR/v95/p249> 7
- Yamamoto, K., Majetich, S. A., McCartney, M. R., Sachan, M., Yamamuro, S., and Hirayama, T. (2008). “Direct visualization of dipolar ferromagnetic domain structures in Co nanoparticle monolayers by electron holography.” *Applied Physics Letters*, 93(8): 082502–3. 3
- Yeomans, J. M. (1992). *Statistical Mechanics of Phase Transitions*. Oxford University Press, USA. 5

