

PAPER • OPEN ACCESS

## Fermi surface determination from momentum density projections

To cite this article: Michael Leitner *et al* 2016 *New J. Phys.* **18** 063033

View the [article online](#) for updates and enhancements.

### Related content

- [Spin-dependent three-dimensional electron momentum density studies in ferromagnetic Co by means of full-scale use of a two-dimensional angular correlation of polarized positron annihilation radiation](#)  
H Kondo, T Kubota, H Nakashima et al.
- [Application of the Lock-Crisp-West theorem to vanadium and paramagnetic chromium](#)  
A K Singh and R M Singru
- [Reconstruction techniques in angular correlation of positron annihilation experiments](#)  
L M Pecora

### Recent citations

- [Lattice harmonics expansion revisited](#)  
G. Kontrym-Sznajd and A. Holas



## PAPER

## Fermi surface determination from momentum density projections

## OPEN ACCESS

RECEIVED  
16 March 2016REVISED  
6 May 2016ACCEPTED FOR PUBLICATION  
13 May 2016PUBLISHED  
29 June 2016Michael Leitner<sup>1,2</sup>, Josef Andreas Weber<sup>1</sup> and Hubert Ceeh<sup>1</sup><sup>1</sup> Physics Department, Technische Universität München, James-Franck-Str. 1, 85748 Garching, Germany<sup>2</sup> Heinz Maier-Leibnitz-Zentrum, Technische Universität München, Lichtenbergstr. 1, 85748 Garching, GermanyE-mail: [michael.leitner@frm2.tum.de](mailto:michael.leitner@frm2.tum.de)**Keywords:** positron annihilation, Fermi surface, reconstruction, electronic structureSupplementary material for this article is available [online](#)

Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

**Abstract**

The problem of determining a metal's Fermi surface from measured projections of the electron or electron/positron momentum densities, such as those obtained by Compton scattering or angular correlation of positron annihilation radiation, respectively, is examined in a Bayesian formulation. A consistent approach with an explicit treatment of the Fermi surface already at the reconstruction stage is presented, and its advantages compared to previous practice are discussed. A validation of the proposed method on simulated data shows its systematic accuracy to be very satisfactory and its statistical precision on modest experimental data to be surprisingly good.

**1. Introduction**

The classical approach for experimentally determining the Fermi surface of metallic systems is to exploit the connection of its shape to quantum oscillations, such as in the de Haas-van Alphen effect or in the Shubnikov-de Haas effect. These techniques, while proven to be very powerful for obtaining quantitative information on the dimensions of the Fermi surface with high precision, vitally depend on long electron scattering lengths, and therefore are applicable only at cryogenic temperatures and at vanishing occupational disorder, which excludes interesting cases such as, for instance, the insulator/metal-transition in FeSi around room temperature (Tomczak *et al* 2012). Apart from this restriction with respect to situations that can be studied, the task of assigning the measured extremal orbits to actual features of the Fermi surface becomes challenging for systems with more complicated multi-sheet Fermi surfaces (e.g., Brasse *et al* 2013).

Experimental methods that measure occupied electron momentum densities work equally well for disordered states (including both site occupational disorder and temperature) and therefore hold the promise of determining the Fermi surface from the discontinuities in the occupations, as long as the concept of a well-defined Fermi surface is meaningful at all (Dugdale 2014). These methods comprise Compton scattering (Cooper 1985) and angular correlation of positron annihilation radiation (ACPAR, also ACAR; Bisson *et al* 1982). However, in these techniques the primary experimental data are plane projections (Compton scattering, early positron annihilation) or line projections (recent positron annihilation setups), from which the three-dimensional momentum density has to be computationally determined. Also, in positron annihilation experiments the sampled two-photon momentum density differs from the underlying ideal electron momentum density due to positron wave-function effects and electron-positron correlations, although the position of the discontinuity due to the Fermi surface will remain unaffected.

The problem of reconstructing densities from projections has seen much attention due to its relevance for medical and technological imaging. For used approaches in the specific case of momentum densities in solids see, e.g., the recent review by Kontrym-Sznajd (2009). In short, in the majority of previous works densities are reconstructed in a first step either by discretizing an analytical inverse of the Radon transform (Radon 1917) or by expanding both the measured projections as well as the unknown densities into basis functions with convenient transformation behaviour, and in a second step the Fermi surface is determined from the densities. These approaches are direct methods in the sense that the results are derived by applying a sequence of explicitly defined transformations to the data. As a consequence, their computational complexity is modest, which

historically was a reason for their adoption, and for diagnostical applications still is (though recently more involved methods are being taken up also in this field (Fleischmann and Boas 2011), due to their better performance at low radiation doses or in situations where, e.g., metal implants render some scan directions inaccessible). However, experimental methods for determining electron momentum densities are typically limited by the count rate. Therefore it should be the power of an analysis method rather than the computational effort that dictates which method is to be preferred.

In this paper, we will give a Bayesian formulation of the data analysis problem. For illustration, we will concentrate on the case of two-dimensional ACPAR, although our approach is equally applicable to Compton scattering. We will show that the formulation corresponds to a regularized inverse problem, and we will illustrate how its solution can be practicably obtained. The main features of our method lie in the avoidance of systematic errors thanks to a consistent description of the whole problem, and quantitative results due to an explicit parametrization of the Fermi surface. We will discuss the associated advantages in comparison with previous approaches, and we will demonstrate our method's power by applying it to simulated data. This will allow us to conclude that an ACPAR experiment with moderate statistics and resolution is able to determine the shape of the Fermi surface *quantitatively* with an accuracy that is comparable to quantum-oscillatory methods, but at elevated temperatures and in the presence of disorder.

## 2. Definition of the problem

Here we will propose a specific formulation of the data analysis problem and connect it to the physical aspects of the situation. Further, we will discuss our approach in comparison to those used previously.

### 2.1. Bayesian formulation

For the problem at hand, the basic unknown quantity is the three-dimensional momentum density  $\rho(\mathbf{p})$ . For Compton scattering, this concerns the actual electron momentum density (Fourier components of the electron states), while in ACPAR the positron wave-function and electron-positron correlation effects modify the probed density, which is then termed two-photon momentum density. In either case, the occupied density is a smooth function, apart from steps when crossing sheets of the Fermi surface. In the extended-zone scheme, the density will decay towards high  $\mathbf{p}$  and display the point group symmetry of the crystal, while the Fermi surface features have the crystal's space group periodicity.

The experimental data  $y$  are given by a set of one- or two-dimensional spectra, corresponding to plane or line projections of the underlying three-dimensional occupied densities. The uncertainties of the data follow Poissonian statistics, in particular the noise for distinct data points is independent, and the noise probability distribution is defined by the point-wise expectation (and can be estimated from the measured signal).

The relation between these two quantities is given by a linear operator  $\mathbf{P}$ , which specifies how the expected value of the measured signal (i.e., before Poisson quantization) results from a given, but experimentally *a priori* unknown, momentum density. We will call it the projection operator, as its function is essentially to integrate over the transversal or longitudinal momentum components for the distinct experimental orientations of the crystal. Additionally, it also takes the detector efficiency and momentum sampling functions, as well as the smearing of the spectra due to finite resolution into account, which will be discussed in more detail in section 3.2.

Thus, our initial formulation for the problem to be solved is the following: what is the posterior probability distribution  $p_{\text{post}}$  for  $\rho$ , given an outcome of the experiment  $y$  and considering prior knowledge or assumptions, quantified in the prior distribution  $p_{\text{prior}}(\rho)$ ? By Bayes' formula, the answer is formally given by

$$p_{\text{post}}(\rho|y) = \frac{p_{\text{like}}(y|\rho)p_{\text{prior}}(\rho)}{p(y)}, \quad (1)$$

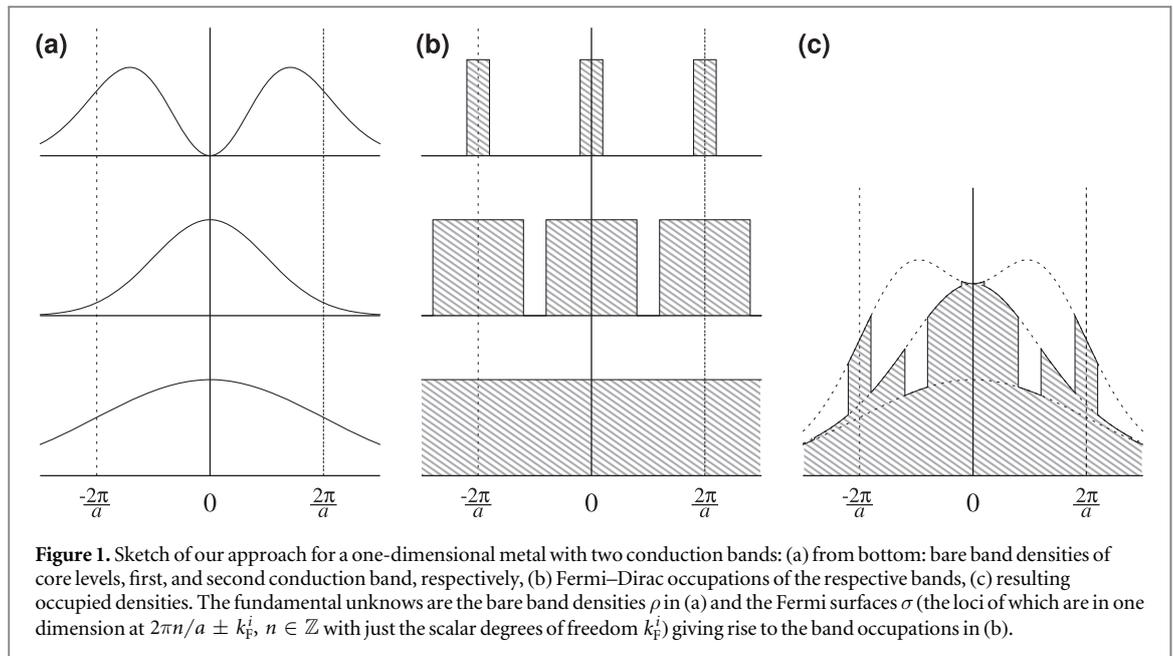
where  $p_{\text{like}}(y|\rho)$  is the so-called likelihood function, which is the modelled probability for observing the actual experimental outcome  $y$  in a repeated experiment under the assumption of an underlying fixed momentum density  $\rho$ . In the present case it is just

$$p_{\text{like}}(y|\rho) = \prod_i p_{\text{Poisson}}(y_i; (\mathbf{P}(\rho))_i) =: P_{\text{Poisson}}(y; \mathbf{P}(\rho)), \quad (2)$$

where the measured spectra are treated as a vector  $(y)_i$  and

$$p_{\text{Poisson}}(k; \lambda) = \frac{\lambda^k}{k!} e^{-\lambda}, \quad (3)$$

is the familiar expression for the Poisson distribution, with  $P_{\text{Poisson}}$  its version for vector arguments.



**Figure 1.** Sketch of our approach for a one-dimensional metal with two conduction bands: (a) from bottom: bare band densities of core levels, first, and second conduction band, respectively, (b) Fermi–Dirac occupations of the respective bands, (c) resulting occupied densities. The fundamental unknowns are the bare band densities  $\rho$  in (a) and the Fermi surfaces  $\sigma$  (the loci of which are in one dimension at  $2\pi n/a \pm k_{\text{F}}^i$ ,  $n \in \mathbb{Z}$  with just the scalar degrees of freedom  $k_{\text{F}}^i$ ) giving rise to the band occupations in (b).

The challenge to the physicist lies in formulating an expression for  $p_{\text{prior}}(\rho)$  that takes into account the available understanding of the problem and therefore ensures a physically meaningful result. A first step towards this goal is the strict requirement for  $\rho$  to conform to the point symmetry group of the crystal. Instead of encoding this via a  $\delta$ -like part in  $p_{\text{prior}}(\rho)$ , it is more efficient to describe the densities only by their values in the point symmetry group’s irreducible wedge, with appropriate continuation over all of the momentum space. From here on,  $\rho$  is to be understood in this sense.

The defining idea of our proposed approach, which is illustrated in figure 1, actually follows just from letting modelling be guided by the physical picture: the expected behaviour of the occupied density (smooth variations apart from jumps at the Fermi surface) is due to its being composed of contributions from the distinct conduction bands, each multiplied by a Fermi–Dirac occupation function with a practically discontinuous jump, sitting atop the contribution from the core states. The bare band densities, i.e., before taking occupation into account, will in fact be smooth, which can be derived from any simply tractable band structure model, be it nearly free electrons, tight binding, or the effective potentials in density functional theory. Instead of considering the total occupied density with jumps at the Fermi surface as the fundamental unknown, as has hitherto been done almost exclusively, the most natural formulation therefore is to have as free parameters both (1) a parametrization of the Fermi surface sheets in each conduction band (and spin channel in the case of magnetic ordering) and (2) the bare densities of the respective conduction bands as well as the summed contribution from all core bands. The prior distribution  $p_{\text{prior}}(\rho)$  can then be used to favour smooth band densities with, e.g., additional positivity constraints, or include assumptions on the decay towards high momenta encoded in an entropy prior, and to consider also information on the Fermi surface shapes, if available from prior experiments or calculations.

For parametrizing the Fermi surfaces, we use smooth functions that conform to the space group symmetry and define the Fermi surfaces as their level sets. Obviously, also this approach is inspired by the physical picture, where the role of above-mentioned smooth functions is fulfilled by the conduction band dispersions.

With these observations, we can rewrite equation (1) as

$$p_{\text{post}}(\rho, \sigma|y) \propto P_{\text{Poisson}}(y; \mathbf{A}_{\sigma}(\rho)) p_{\text{prior}}(\rho, \sigma) \quad \text{with} \quad \mathbf{A}_{\sigma} = \mathbf{P}\mathbf{X}_{\sigma}, \quad (4)$$

where  $\rho$  is now to be understood as the bare band densities and the action of  $\mathbf{X}_{\sigma}$  for given Fermi surface parameters  $\sigma$  is to multiply the bare band densities by the occupations and sum over the band index, i.e., it is again a linear operator.

If the primary quantity to be determined is the Fermi surface, equation (4) can be marginalized over  $\rho$  to give the posterior distribution of the Fermi surface parameters

$$p_{\text{post}}(\sigma|y) = \int d\rho p_{\text{post}}(\rho, \sigma|y). \quad (5)$$

Note that as we will elaborate below, the shape of  $p_{\text{post}}(\rho, \sigma|y)$  with respect to  $\rho$  does not vary much as a function of  $\sigma$ , so for practical purposes  $p_{\text{post}}(\sigma|y)$  is proportional to the maximum of  $p_{\text{post}}(\rho, \sigma|y)$  for given  $\sigma$ .

## 2.2. Comparison to previous approaches

The main points that distinguish our approach from the majority of those used previously are the following: first, it is formulated as a general problem of Bayesian inference instead of as a recipe of transforms obtained from analytical manipulations, and second, the Fermi surface is treated explicitly during the reconstruction instead of determined afterwards from the reconstructed densities. Here we will discuss the implications of these differences.

The Radon transform is a bijection (in particular, it is invertible) between suitably regular  $n$ -dimensional function spaces (Radon 1917). Specifically, it corresponds to the relation between a pointwise defined function and its integrals over all lines in the plane (in two dimensions), or its integrals over all planes in space (in three dimensions). This has the consequence that in two-dimensional ACPAR, where the accessible data would in principle be the integrals over all lines in space (a four-dimensional manifold), typically only projections within a certain plane of rotation have been considered, which reduces the three-dimensional problem to independent two-dimensional problems that can be solved by direct algorithms for the two-dimensional inverse Radon transform (the literal meaning of tomography). While in the case of medical imaging such a sectioning approach is appropriate for minimizing the necessary radiation doses due to the elongated shape of the human body, this does not apply for projections of momentum densities in reciprocal space and thereby corresponds to neglecting potentially independent information. In fact, due to the crystal's point group symmetry taking only projections within a common rotation plane already implies information about out-of-plane projections in many cases, which a pure sectioning approach does not take into account. In a recent proposal (Kontrym-Sznajd *et al* 2004) this deficiency is addressed by re-parametrizing the reconstructed slices in terms of basis functions with the appropriate symmetry, but it is not known to which extent this can recover the information that was initially available, considering the correct symmetry, or whether the filtered reconstruction still reproduces the measured projections. In addition, the direct approaches typically need evenly and densely spaced projections in the plane with equal statistical precision, where the angular increment corresponds to the smallest resolvable features. In contrast, our proposed formulation allows an arbitrary number of projections with arbitrary orientation and noise level to be used, and takes into account all available information considering the prescribed symmetry. An early suggestion by Pecora (1987) to solve directly for the coefficients of three-dimensional basis functions with appropriate symmetry, dealing with the above-mentioned dimensional problem in a least-squares sense, should in principle be of equivalent quality in terms of these criteria, but apparently has seen only limited use. Note that it was shown in above-quoted reference that it is not merely a welcome option to be able to use arbitrary orientations, but that using low-symmetry projections is actually indicated for optimal results in reconstruction.

Direct transform methods are derived from analytical inversions of the mathematical Radon transform. As such, they cannot cope with experimental subtleties such as finite resolutions. Therefore, the resolution is typically deconvolved from the measured spectra before reconstruction, if it is considered at all. In most cases, maximum-entropy regularization is used to solve this ill-posed problem, either explicitly (Fretwell *et al* 1995) or implicitly (Gerhardt *et al* 1998), and sometimes even more arbitrary methods are used (Chiba *et al* 2007). In contrast, aggregating all experimental complications (e.g., resolution, momentum sampling, and centring) into the forward operator as proposed here allows us to solve the inverse problem in a single step, subject to prior assumptions (equivalent to regularization) on the fundamental physical quantities, and thereby to rule out a compounding of the potentially conflicting regularization biases introduced at sequential steps. A very welcome additional consequence of our not touching the experimental spectra at all is that the noise statistics of the spectra remain unadulterated, specifically the Poissonian behaviour with independence between neighbouring pixels is conserved, which allows us to propagate the uncertainty of the data consistently into an estimated error of the final reported quantities.

The effect of experimental noise on the reconstructed densities is a critical issue for the conventional methods. Especially with the direct transform methods that rely on the central slice theorem and interpolation (Kondo *et al* 1993), the reconstruction is typically very ill-defined around the origin due to the conflicting information, and there is no obvious way of countering this problem. Such an effect is present also for basis expansion methods (Samsel-Czekała and Boguszewicz 2005). In general, the regularity of the reconstructions is controlled by the number of considered basis functions in the case of expansion methods, and by the use of an appropriate filter function in the various filtered transform methods, which in either case is rather opaque to the user. In contrast, in our formulation regularization by way of the prior distribution is explicit, and no noise artefacts appear in the reconstruction. Note that an explicit regularization functional is also the only practicable way to have non-linear biases such as a non-negativity constraint (as already observed by Pylak *et al* 2011).

In most published works, the identification of the Fermi surface is done subsequent to and independent from the reconstruction of the density. The simplest option is to transform the density from  $\mathbf{p}$ - to  $\mathbf{k}$ -space, i.e., subject it to the so-called LCW folding (Lock *et al* 1973), and define the Fermi surface as an iso-density contour, e.g., according to a maximum gradient criterion (Biasini *et al* 2002). The problem with such an approach in

positron annihilation experiments is that due to the inequivalence between the measured two-photon momentum density and the electron momentum density, filled bands give rise to a non-constant background (Lock and West 1975), so that for finite resolution the Fermi surface does not strictly correspond to any iso-density contour. Edge-detection or enhancement methods (e.g., Dugdale *et al* 1994, O'Brien *et al* 1995) should be able to obviate this issue, but also in this case finite experimental resolution will tend to smooth regions of the Fermi surface with high curvature. Because of these systematic uncertainties, reconstructed Fermi surfaces are typically reported only as three-dimensional renderings (e.g., Weber *et al* 2015), in contrast to the precise statements, for instance in terms of Fourier coefficients, in classical quantum oscillatory works (Roaf 1962).

Due to the explicit treatment of the Fermi surface during reconstruction, our proposed method does not suffer from the above-mentioned effects. We are aware of only two comparable proposals: Biasini (2000) determined the parameters of a model Fermi surface by minimizing the deviation from the measured spectra after LCW folding, which is essentially the same idea as our proposal, only with constant band densities in  $\mathbf{k}$ -space. This idea of reconstruction by a piecewise-constant function with explicit treatment of the step manifold has also been suggested in the context of diagnostical imaging (Ramlau and Ring 2007). Second, Laverock *et al* (2010) propose to fit the LCW-folded spectra by calculated band densities, where the free parameters correspond to a state-dependent annihilation enhancement and energy shifts of the rigid bands. Clearly, in such an approach the freedom in the shape of the Fermi surface is very restricted, and the results will depend on the correctness of the electronic structure used as input.

The only disadvantage of our proposed approach concerns the increased numerical effort compared to direct methods, although this has ceased to be relevant with today's computing power as we will show below. Also, no well-tested software packages exist yet, but for momentum density reconstructions—different from medical imaging—typically custom implementations are used in any case, which especially in the case of orthogonal expansions can easily become more involved than our implementation.

### 3. Implementation

Due to the large dimensionality of the problem defined by equation (4), an efficient way to arrive at its solution is imperative for its practicability. We will give a detailed discussion of our implementation below, while for its numerical aspects we refer to the supplementary material.

#### 3.1. Parametrization of the solution space

In our formulation, both the band densities  $\rho$  and the Fermi surface sheets  $\sigma$  are explicit degrees of freedom. As mentioned above, the physical picture suggests to describe the Fermi surface sheets as the level sets of auxiliary smooth functions with the appropriate symmetry. For this purpose we use a Fourier description, where the required reciprocal-space translation symmetry is enforced by considering only those Fourier coefficients that correspond to real-space lattice vectors. Note that this is formally identical to a tight-binding description, and in fact it has been shown that it can reproduce the experimental Fermi surfaces in the noble metals with only a few free parameters (Roaf 1962). In contrast, for systems that conform rather to the free-electron picture, such as Al, different models will probably be more efficient in describing the Fermi surfaces (Ashcroft 1963).

Thanks to the explicit treatment of the Fermi surface, the band densities  $\rho$  will be smooth and can therefore be described with comparatively low resolution. Here we use quadratic B-splines. The point symmetry is fulfilled by considering only coefficients within the irreducible wedge, with appropriate continuation over all of the reciprocal space. For reasons of efficiency we employ a non-constant sampling density, with higher resolution at low momentum. This is justified by the observation that the contribution at high momenta is mainly due to the core electrons, which will not be influenced much by the crystal structure and therefore be nearly isotropic.

We want to emphasize that the choice of basis functions is not essential to the idea of the method. Considerable mathematical effort has been expended in deriving expansions that fulfill orthogonality relations of some sense under idealized projection operations (Louis 1984). In principle, such parametrizations could be used also in our proposal, where we expect that a sufficient description could be attained already at lower degrees of the expansion compared to previous implementations, as the parametrization needs to capture only the variation of the band densities excluding the Fermi surface steps. As an additional advantage, the system of linear equations to be solved would be better conditioned due to the approximate orthogonality properties (strict orthogonality is destroyed by experimental complications such as a finite resolution). However, such basis functions are typically non-vanishing nearly everywhere in reciprocal space, corresponding to full instead of sparse projection operators, which would constitute a serious drawback as we will discuss below.

### 3.2. Matrix form of the operators

The projection operator  $\mathbf{P}$  specifies how a given three-dimensional momentum density leads to the set of two- or one-dimensional spectra in the specified orientations by line or plane projections. For the construction of this matrix we compute for each voxel the set of pixels or bins in each spectrum that can potentially have an overlap with the projection of the voxel, and set the respective entry to the proportion received by the corresponding pixel. We find this proportion by an approach based on a look-up table with linear interpolation, pregenerated from a high-resolution projection of a single voxel cube. Given the finite experimental resolution (discussed below), this approach can be considered as equivalent to the exact solution.

The optimal choice for the discretization of the three-dimensional momentum density is somewhat finer than the experimental resolution; a coarser representation will lead to artefacts, while a finer discretization will be numerically more expensive without any effect on the spectra after the application of resolution smearing. The same reasoning holds for the discretization of the spectra. As a consequence, each voxel will contribute to only a few pixels in each spectrum, allowing us to precompute and store the projection matrix in sparse format.

By setting the resolution kernel to zero once its value has dropped below some threshold, also the resolution smearing operator has the form of a sparse matrix. The final step in obtaining the spectra from a given momentum density is the pointwise multiplication by the momentum sampling function, which corresponds to a diagonal matrix. In principle, the projection operator  $\mathbf{P}$  is then the product of these three precomputed sparse matrices, although due to the involved dimensionalities it is more efficient to utilize associativity of matrix multiplication in the further steps and never compute the product of the resolution smearing matrix and the projection matrix proper explicitly.

For fixed Fermi surface parameters  $\sigma$ , the occupation operator  $\mathbf{X}_\sigma$ , which computes the momentum density from the band densities, is represented as the horizontal concatenation of diagonal matrices, with the entries corresponding to the respective occupations. Here special care has to be taken with the voxels at the Fermi surface; for a continuous variation of the resulting spectra with the shape of the Fermi surface, the occupations have to be computed according to the proportions of the voxel within the Fermi surface (at the relevant temperatures the width of the Fermi surface compared to the voxel size is negligible). We accomplish this by the Gilat–Raubenheimer method (Gilat and Raubenheimer 1966a, 1966b), i.e., by linearizing the variation of the band energy within the voxel and computing the enclosed volume explicitly.

Summing up, the linear operator  $\mathbf{A}_\sigma$  that relates the band density parameter vector  $\rho$  to the spectra for a given shape of the Fermi surface  $\sigma$  is the matrix product  $\mathbf{P}\mathbf{X}_\sigma$  of the matrices discussed above, modified by an additional degree of freedom corresponding to a constant background intensity.

### 3.3. Algebraic solution of the Bayesian problem

The fact that allows for an efficient computation of the posterior density (4) lies in the observation that  $p_{\text{like}}(\rho, \sigma|y)$  can be approximated very accurately by a Gaussian distribution for fixed  $\sigma$ . Specifically, by equating the measured experimental spectra with their true value (i.e., the expected value before quantization) and expanding the logarithm of the Poissonian distribution around the maximum we have

$$\log(p_{\text{like}}(\rho, \sigma|y)) = -\frac{1}{2}(\mathbf{A}_\sigma\rho - y)^\top \mathbf{W}(\mathbf{A}_\sigma\rho - y) + \text{const}, \quad (6)$$

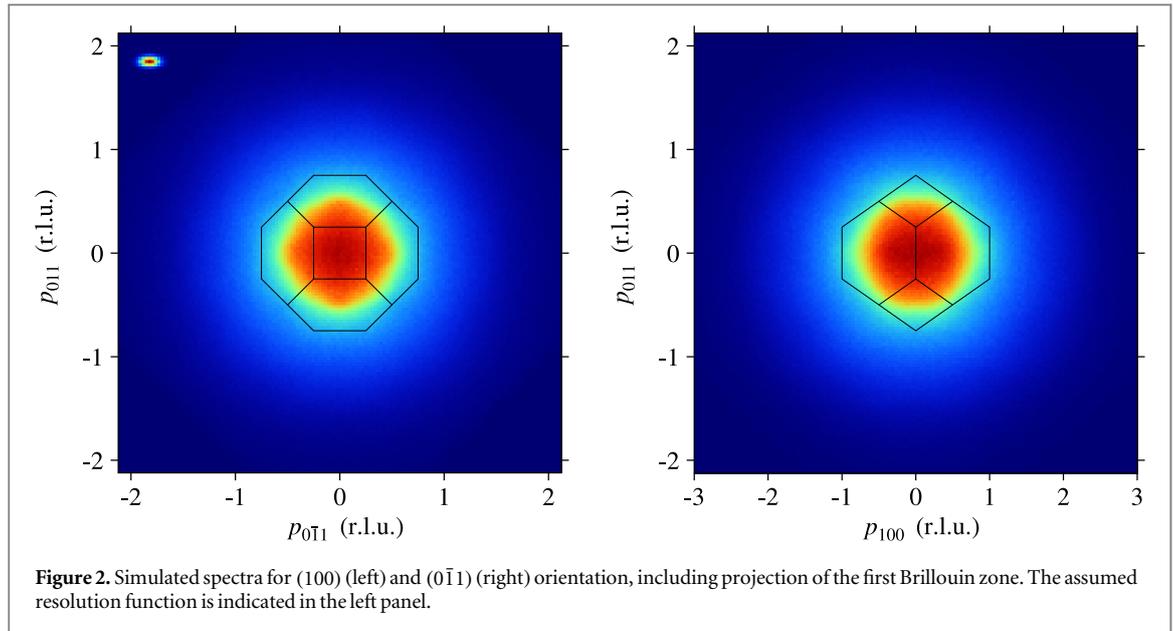
where the weighting matrix  $\mathbf{W}$  is the inverse of the covariance matrix. This initial problem can be iteratively refined by expanding around the updated expected values. In our tests, such an iterative approach converged rapidly (within one iteration), and even the initial distribution was a faithful representation of the final result, as the counts per pixel were not too low and the residuals were small.

To guarantee smooth reconstructed band densities, we formulate our prior distribution as a Gaussian distribution of the square norm of the band densities' second derivatives plus an analogous contribution from the first derivatives to favour monotonicity. A non-constant weighting of these norms can be used to penalize the same density curvatures more if they occur at high momenta, where the absolute values of the densities are smaller. Our choice is functionally equivalent to Tikhonov regularization (Tikhonov 1963), i.e., a positive semidefinite quadratic form of the parameters, when searching for the maximum a posteriori estimate.

As both  $p_{\text{like}}$  and  $p_{\text{prior}}$  are multivariate Gaussian distributions, so is their product  $p_{\text{post}}$ . As a consequence, the maximum a posteriori estimate  $\rho^*(\sigma)$  for given  $\sigma$  can be obtained by solving the system of linear equations

$$\left(\mathbf{A}_\sigma^\top \mathbf{W} \mathbf{A}_\sigma + \sum_i \lambda_i \mathbf{D}_i^\top \mathbf{D}_i\right) \rho = \mathbf{A}_\sigma^\top \mathbf{W} y, \quad (7)$$

where  $\mathbf{D}_i$  are the matrices that compute the (optionally weighted) derivatives of the band densities from the parameters and  $\lambda_i$  the corresponding regularization parameters. Also the marginal posterior distribution with respect to  $\sigma$  as defined by (5) follows easily as



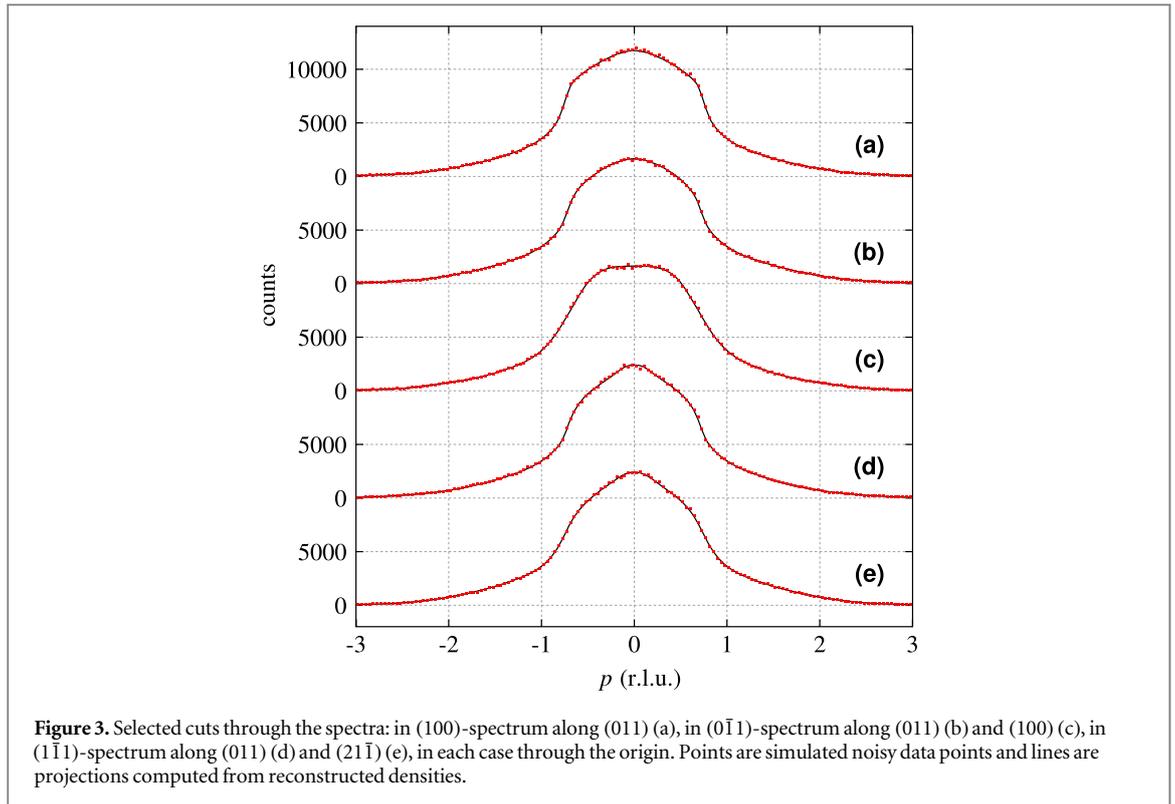
$$p_{\text{post}}(\sigma|y) = p_{\text{post}}(\rho^*(\sigma), \sigma|y) \sqrt{\det \left( \mathbf{A}_\sigma^\top \mathbf{W} \mathbf{A}_\sigma + \sum_i \lambda_i \mathbf{D}_i^\top \mathbf{D}_i \right)}. \quad (8)$$

#### 4. Application to model data

Most previous proposals for algorithms to reconstruct three-dimensional momentum densities have been validated only by comparing the results on experimental data to those of other algorithms, if at all (e.g., Kontrym-Sznajd *et al* 2008, Pylak *et al* 2011). Obviously this is not satisfactory, as specific aspects of the data could have been missed simultaneously by both the tested and the benchmark methods. In other cases the algorithms have been applied to synthetic data obtained by comparatively simplistic models, but also there the comparison has been done only in qualitative terms (e.g., Pecora 1987, Kontrym-Sznajd *et al* 2004, 2008). As we claim here to be able to reconstruct Fermi surfaces quantitatively, we have to substantiate this claim by demonstrating both its fidelity (the magnitude of introduced systematic errors) and statistical performance (the propagation of experimental noise to the resulting dimensions) on realistic data for which the correct solution is known.

For this purpose we chose the system of copper with its prototypical and well-known Fermi surface. We computed the electronic structure of Cu self-consistently in the generalized gradient approximation with the PBE exchange-correlation functional (Perdew *et al* 1996) by the density functional code `ABINIT` (Gonze *et al* 2009). With the converged density we computed both the electron wave functions and energies on a fine mesh and the  $\Gamma$ -point positron wave function (Barbiellini *et al* 1995). Thanks to the plane-wave formulation used in the `ABINIT` code, the electron-positron momentum densities could be conveniently derived in a custom implementation corresponding to the independent particle model. From the three-dimensional density plus some realistic background we computed the corresponding ACPAR spectra of  $144^2$  pixels for (001), (110) and (111) orientation, each with  $25 \times 10^6$  counts distributed according to Poissonian statistics. We chose a discretization of 24 pixels per reciprocal lattice constant and assumed an anisotropic Gaussian resolution function with  $2 \times 1$  pixel standard deviation. Two of the resulting spectra are given in figure 2. With the actual lattice constant of Cu, our chosen resolution would correspond to  $1.32 \times 0.66$  mrad<sup>2</sup> FWHM, representative of a typical experimental situation (Leitner *et al* 2012).

We reconstructed the density in a volume of  $144^3$  voxels at the discretization of 24 voxels per reciprocal lattice constant. We want to emphasize that the coincidence of pixel and voxel size is by no means necessary; specifically for an actual experiment it will be beneficial to choose the voxel size as an integer fraction of the reciprocal lattice constant, while the discretization of the spectra is often established by the apparatus. As Cu has only a single band crossing the Fermi energy, we considered a fully occupied core level electron density and a single conduction band. We parametrized the space of band densities by cubic B-splines with an increasing sampling density towards small  $\mathbf{p}$ , corresponding to 385 degrees of freedom per band, and we described the single Fermi surface sheet by a five-parameter Fourier expansion (where the  $\langle 110 \rangle$  coefficient is fixed to 1 as it



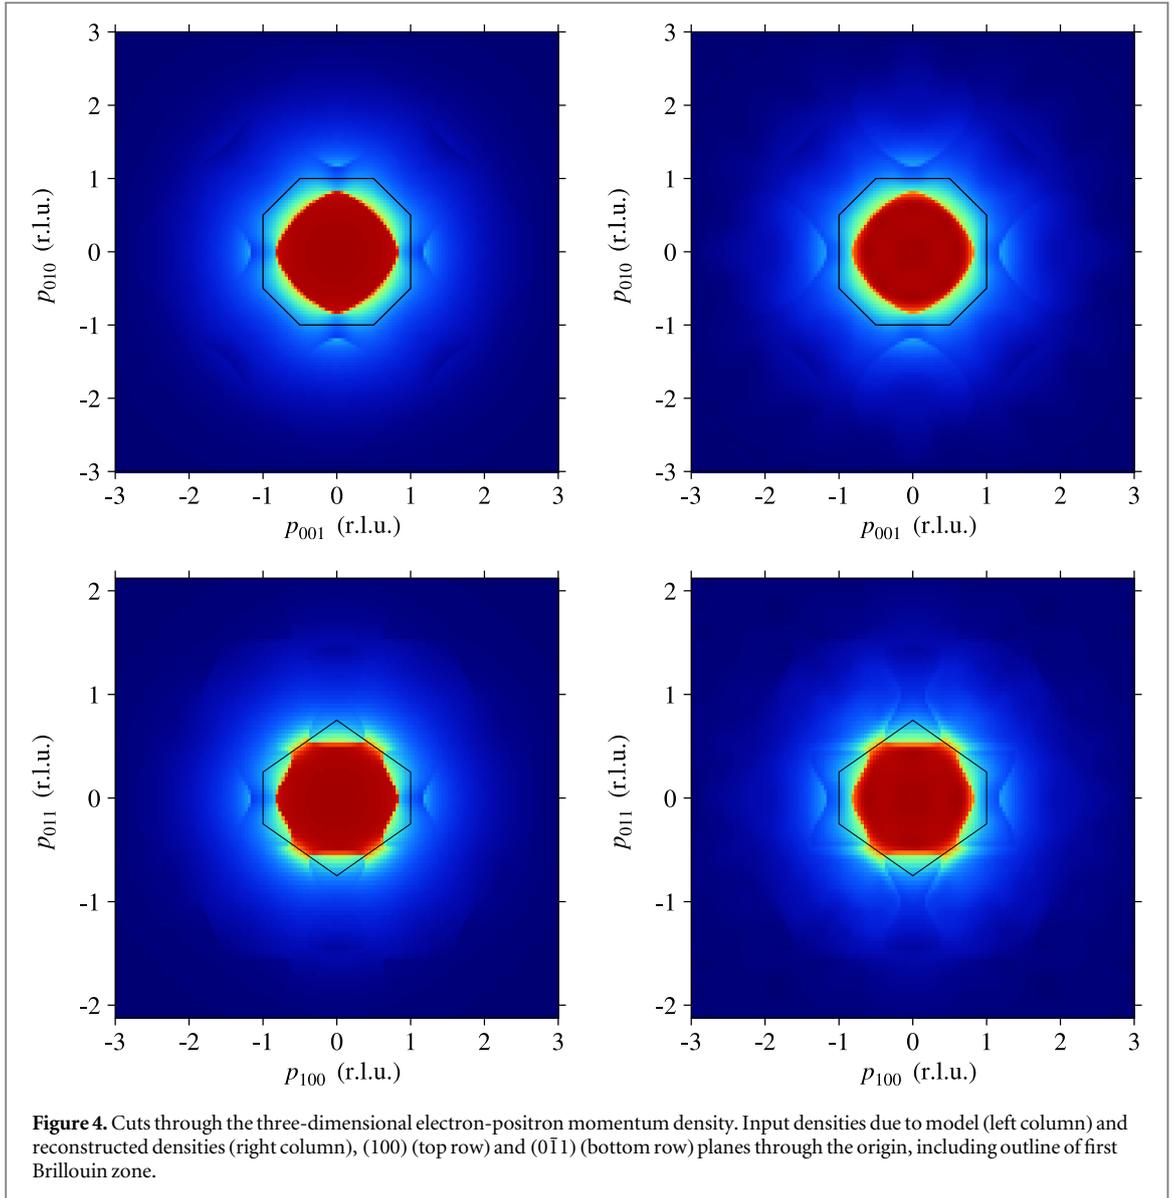
corresponds to a trivial scaling of the energy range, and the  $\langle 000 \rangle$  coefficient is chosen relative to the Fermi energy so as to constrain the occupied volume to half the Brillouin zone). We also added a term to the prior distribution to favour the decay of the coefficients with interaction range, because, as already noted by Roaf (1962), the Cu Fermi surface changes only by very small amounts under certain modifications of the parameters.

Maximizing the posterior probability given by equation (8) for the simulated spectra gives reconstructed spectra that, apart from the missing noise, are visually identical to the input spectra. This is substantiated by the one-dimensional cuts through the spectra shown in figure 3 and the corresponding reduced  $\chi^2$  value of 1.047. The degree of achieved faithfulness to the data obviously depends on the choice of the regularization parameters  $\lambda_j$  in equation (7). Here we used the smallest values that still suppress visually noticeable artefacts (that would correspond to the reconstruction of experimental noise) in the reconstructed densities.

A comparison of input and reconstructed density is given in figure 4. Here significant systematic differences can be discerned. Specifically, the original density is virtually constant within the Fermi surface in the first Brillouin zone and decays quite abruptly through the necks into the second Brillouin zone. In contrast, the tendency of the reconstruction towards smooth variations in the band densities leads to a decrease already within the Fermi surface in the first Brillouin zone and to higher contributions from outer zones. Apart from that, the main features are clearly reconstructed in a qualitatively correct way.

For assessing the algorithm's performance in determining the Fermi surface in quantitative terms, we focus on three specific features: the extent of the Fermi surface along (100), along (110), and the radius of the (111)-neck (its deviation from a circular shape is negligible). In the input data, these three dimensions are 1.063, 0.946 and 0.203, respectively, measured in units of  $r_f$ , the radius of the free-electron Fermi sphere. In the reconstructions, the corresponding values are 1.063(3), 0.949(1) and 0.182(6), denoting the average and the standard deviations of the maximum-a-posteriori values for different realizations of the counting noise. For a single realization of the counting noise, the posterior probability distribution defined by equation (8) is to a good approximation a Gaussian distribution, with a covariance matrix that essentially corresponds to above-quoted standard deviations, which would allow us to estimate the statistical uncertainties of the dimensions obtained from an experiment.

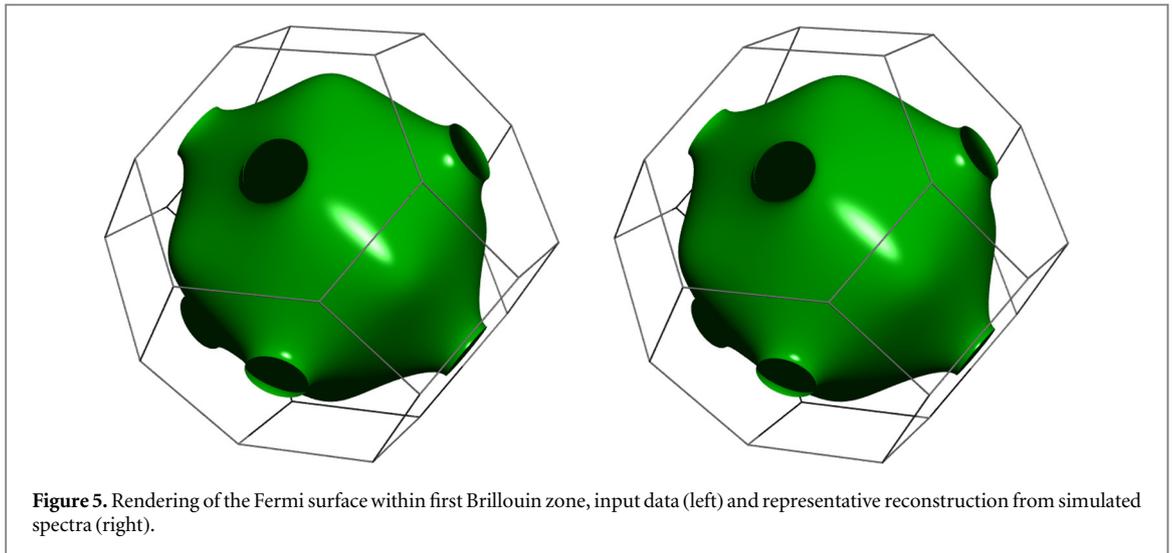
The small statistical uncertainty of the results indicates that our approach can also be used for qualitative statements. For instance, constraining the Fermi surface parameters in the reconstruction so that the necks along (111) become disconnected gives a maximum-a-posteriori probability that is smaller by 140 orders of magnitude, i.e., in a fictitious experiment this scenario could essentially be ruled out.



## 5. Discussion

The fact that our reconstruction reproduces the measured spectra perfectly within the noise shows that it considers all the information present in the data. Within these boundary conditions it yields the reconstruction most probable according to the prior assumptions. We think that our assumptions of smoothly varying band densities are arguably the soundest on physical grounds, and definitely the most transparent and easiest to adjust compared to those inherent to direct methods.

Our results show that three ACPAR spectra with moderate statistics and unexceptional experimental resolution suffice for our data analysis method to give statistically very well-defined results on the Fermi surface dimensions. In terms of accuracy it has to be noted that specifically the radius of the (111)-necks is underestimated by 10%. However, in the overall picture this corresponds only to an error of 2% of the mean Fermi surface radius, and all other regions of the Fermi surface are determined still much more accurate (see the juxtaposition of the input and reconstructed Fermi surface renderings in figure 5). For comparison, the Fermi surfaces of Ag and Au differ much more from the actual Fermi surface of Cu than the reconstruction does with a (111)-neck of 0.137 in Ag and a (100)-radius of 1.135 in Au (Roaf 1962). Also, if a plausible model for the electronic structure in a given system is available, an analysis as presented here can be done on the model and the experimental results can be corrected in first approximation for the systematic effects displayed by the model reconstructions. Note further that also in quantum-oscillatory methods the Fermi surface has to be reconstructed from the measured data, guided by some plausibility assumptions.



In the interpretation of an actual experiment it would probably be too optimistic to expect a  $\chi^2$  as low as reported here due to experimental imperfections such as additional contributions to the spectra, e.g., from surface positronium ejection, or slight misorientations. However, due to the adaptability of our formulation such effects can be included at the cost of a few additional free parameters. This is in contrast to direct methods with, e.g., the strict assumption of crystal symmetry in the plane of sample rotation.

It has to be admitted that after the alkali metals, the noble metals with Cu as the example chosen here have the simplest Fermi surface and are therefore probably the easiest systems to consider. For systems with multiple Fermi surface sheets separated only by a small distance, resolution effects will render the problem ill-defined. Due to an analogous reasoning also large real space cells and consequently small reciprocal space features increase the difficulty in solving a system. However, these limitations obviously apply equally to any method of data interpretation.

The last point we want to stress is that in principle any additional information can be considered in the prior probability assumptions. For example, in the case at hand the ill-defined (111)-neck radius could be fixed to the value given by de Haas-van Alphen-measurements, where it corresponds to a prominent and well-defined frequency (Shoenberg 1962), while the shape of the Fermi surface belly is more directly encoded in the positron annihilation data.

## 6. Conclusion

Here we have presented a new approach on obtaining the shape of the Fermi surface from Angular Correlation of Positron Annihilation Radiation, or equivalently Compton scattering, data. We have pointed out the advantages of a unified formulation as an inverse problem in the Bayesian setting instead of the conventional sequential steps. A test with simulated data has demonstrated that modest requirements on the data statistics lead to statistically well-defined results on the Fermi surface dimensions, and we have discussed the small systematic biases introduced by the prior assumptions for the example of copper. These insights open the way for obtaining Fermi surface information with a quality that is comparable to quantum-oscillatory methods under conditions such as high temperature and occupational disorder, where the classical methods cannot be applied. For instance, even for prototypical elemental systems there is as yet no experimental information on possible quantitative reconstructions of the Fermi surface in the percent range at finite temperatures, which probably cannot be excluded *a priori* except for the most prototypical free-electron-like metals.

## Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft (DFG) through TRR 80. We thank Stephen Dugdale for encouraging words on our approach.

## References

- Ashcroft N W 1963 *Phil. Mag.* **8** 2055–83
- Barbiellini B, Puska M J, Torsti T and Nieminen R M 1995 *Phys. Rev. B* **51** 7341–4
- Biasini M 2000 *Physica B* **275** 285–94
- Biasini M, Ferro G, Kontrym-Sznajd G and Czopnik A 2002 *Phys. Rev. B* **66** 075126
- Bisson P E, Descouts P, Dupanloup A, Manuel A A, Perreard E, Peter M and Sachot R 1982 *Helv. Phys. Acta* **55** 100–21
- Brasse M, Chioncel L, Kuneš J, Bauer A, Regnat A, Blum C G F, Wurmehl S, Pfeleiderer C, Wilde M A and Grundler D 2013 *Phys. Rev. B* **88** 155138
- Chiba T, Toyama T, Nagai Y and Hasegawa M 2007 *Phys. Stat. Sol. (c)* **4** 3993–6
- Cooper M J 1985 *Rep. Prog. Phys.* **48** 415
- Dugdale S B 2014 *Low Temp. Phys.* **40** 328–38
- Dugdale S B, Alam M A, Fretwell H M, Biasini M and Wilson D 1994 *J. Phys.: Condens. Matter* **6** L435–43
- Fleischmann D and Boas F E 2011 *Eur. Radiol.* **21** 510–7
- Fretwell H M, Dugdale S B, Alam M A, Biasini M, Hoffmann L and Manuel A A 1995 *Europhys. Lett.* **32** 771
- Gerhardt U, Marquardt S, Schroeder N and Weiss S 1998 *Phys. Rev. B* **58** 6877–91
- Gilat G and Raubenheimer L J 1966a *Phys. Rev.* **144** 390–5
- Gilat G and Raubenheimer L J 1966b *Phys. Rev.* **147** 670 (erratum)
- Gonze X et al 2009 *Comput. Phys. Comm.* **180** 2582–615
- Kondo H, Kubota T, Nakashima H, Murakami Y and Tanigawa S 1993 *Phys. Stat. Sol. (b)* **177** 345–60
- Kontrym-Sznajd G 2009 *Low Temp. Phys.* **35** 599–609
- Kontrym-Sznajd G, Samsel-Czekała M and Biasini M 2008 *Appl. Phys. A* **91** 131–6
- Kontrym-Sznajd G, Samsel-Czekała M, Biasini M and Kubo Y 2004 *Phys. Rev. B* **70** 125103
- Laverock J, Haynes T D, Alam M A and Dugdale S B 2010 *Phys. Rev. B* **82** 125127
- Leitner M, Ceeh H and Weber J A 2012 *New J. Phys.* **14** 123014
- Lock D G, Crisp V H C and West R N 1973 *J. Phys. F* **3** 561–70
- Lock D G and West R N 1975 *Appl. Phys.* **6** 249–56
- Louis A 1984 *SIAM J. Math. Anal.* **15** 621–33
- O'Brien K M, Brand M Z, Rayner S and West R N 1995 *J. Phys.: Condens. Matter* **7** 925–38
- Pecora L M 1987 *IEEE Trans. Nucl. Sci.* **NS-34** 642–50
- Perdew J P, Burke K and Ernzerhof M 1996 *Phys. Rev. Lett.* **77** 3865–8
- Pylak M, Kontrym-Sznajd G and Dobrzyński L 2011 *Appl. Phys. A* **104** 587–92
- Radon J 1917 *Ber. Verh. Sachs. Ges. Wiss. Leipzig, Math. Phys. Kl.* **69** 262–77
- Ramlau R and Ring W 2007 *J. Comput. Phys.* **221** 539–57
- Roaf D J 1962 *Phil. Trans. A* **255** 135–52
- Samsel-Czekała M and Boguszewicz L 2005 *Phys. Stat. Sol. (b)* **242** 1279–85
- Shoenberg D 1962 *Phil. Trans. A* **255** 85–133
- Tikhonov A 1963 *Sov. Math. Dokl.* **5** 1035
- Tomczak J M, Haule K and Kotliar G 2012 *Proc. Natl Acad. Sci. USA* **109** 3243–6
- Weber J A, Bauer A, Böni P, Ceeh H, Dugdale S B, Ernsting D, Kreuzpaintner W, Leitner M, Pfeleiderer C and Hugenschmidt C 2015 *Phys. Rev. Lett.* **115** 206404