

1.1. OVERVIEW AND PRINCIPLES

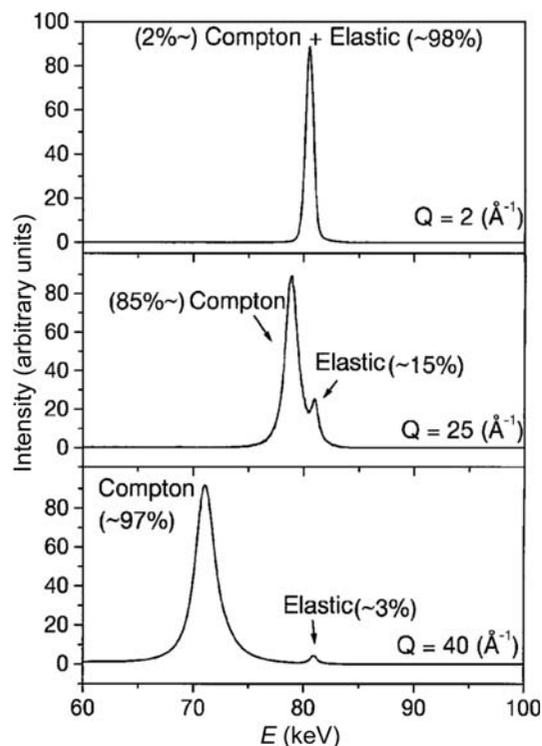


Figure 1.1.24

Spectrum from an energy-resolving detector that shows the elastic and Compton signals as a function of scattering vector Q . [Reprinted with permission from Petkov *et al.* (2000). Copyright (2007) by the American Physical Society.]

the non-energy-resolved magnetic diffuse scattering signal. Magnetic PDF is now possible (Frandsen & Billinge, 2015) as described in Chapter 5.7, as well as reciprocal-space studies of magnetic diffuse scattering (Paddison & Goodwin, 2012).

1.1.5.3.4. Incoherent scattering

Incoherent scattering does not contain any structural information, and cannot be used to study structure in a diffraction experiment since the intensities of the scattered waves do not depend on the position of the scatterers. This does not mean that all incoherent scattering intensity is useless. The fluorescence intensity is incoherent, but may be used in EXAFS experiments to yield structural information. This is because coherent scattering of the photoexcited electron during an absorption event modulates the absorption cross section and therefore the incoherent fluorescence intensity, so a coherent scattering process leaves a measurable response in an incoherent intensity. Incoherent scattering can also be used to measure excitations, although all momentum-transfer information is lost so it is not possible to measure, for example, dispersions of excitations such as phonons and magnons. Even if the scattering process is incoherent, the energy exchanged between the probe and the sample can be measured by the change in wavelength of the scattered wave, and the amplitude of the scattering at each energy transfer is proportional to the density of states of the excitation being probed. In the case of neutrons, the very large incoherent cross section for scattering by hydrogen ($\sim 100\times$ the scattering cross section of most atoms) provides a strong signal for studying low-probability inelastic scattering events. Measuring inelastic scattering from powders can be a rapid way of determining the density of states of phonons, magnons and so on, which is very useful for determining the thermodynamic properties of materials, even though it is less precise than measurement of the full set of dispersion curves.

Another type of incoherent scattering that can be observed in X-ray experiments is Compton scattering (Compton, 1923; Cooper *et al.*, 2004), which is an inelastic incoherent process where the scattering atom recoils during the scattering event. An example of Compton scattering measured in the spectrum from an energy-resolving detector is shown in Fig. 1.1.24.

The Compton scattering is strong in this experiment because the incident X-ray energy is high (80 keV) and the sample is a low-atomic-number alumina-silicate glass. Both the high X-ray energy and the low atomic numbers of the atoms in the sample increase the Compton cross section with respect to the coherent elastic scattering. As the magnitude of the scattering vector, Q , is increased the Compton scattering moves to lower energy and increases in intensity, but the elastic line stays fixed in energy and its intensity decreases because of form-factor and Debye–Waller effects. Momentum as well as energy is conserved in this process and the Compton scattering can be used to measure the momentum distribution of electrons in a material, although this kind of experiment is not widespread these days.

Elastic incoherent scattering provides no information about the sample, and simply degrades the signal-to-noise ratio of the measurement. As such, it is just inconvenient and cannot be easily removed. Monotonic Laue diffuse scattering originates from different chemical species with different scattering powers residing on different sites in the crystal, and when a destructive interference condition is satisfied the resulting intensity does not go to zero but is proportional to $[f_i(h) - f_j(h)]^2$ (Warren, 1990). In pure elements in an X-ray experiment, the atoms on every site are the same and there is no Laue diffuse scattering. This is not true in neutron experiments where different nuclei have different scattering powers and most elements contain a range of isotopes in their natural form (the ‘natural abundance’; Squires, 1996). This results in Laue diffuse scattering even in an element, although it is normally not referred to in these terms but is encompassed by a so-called ‘incoherent neutron cross section’ that is defined and tabulated (see Table 4.4.4.1 in *International Tables for Crystallography*, Volume C) for each element. This is not the only source of incoherent scattering in neutron diffraction, since the scattering power also depends on the relative orientation of the neutron and nuclear spins. In general these spins are all orientationally disordered (and fluctuating) and the result is an additional scattering-event-dependent contribution to the incoherent scattering from the sample, again encompassed by the ‘incoherent neutron cross section’ of the element. Where necessary, it may be possible to make isotopically enriched samples for neutron experiments so that the proportion of isotopes with large incoherent scattering cross sections is minimized (or the isotopes are removed altogether), and the range of isotopes can also be reduced, which further reduces the incoherent component of the signal. However, the cost and difficulty of doing this means that it is rarely done.

1.1.6. Local and global optimization of crystal structures from powder diffraction data

1.1.6.1. Rietveld refinement

More than 40 years have passed since the publication of the pioneering papers by Hugo Rietveld (Rietveld, 1967, 1969), in which he described a method for the refinement of crystal structures from neutron powder diffraction data. Neutron data sets from reactor sources were more amenable than X-ray data sets to this method because the line profiles are quite Gaussian.

1. INTRODUCTION

However, it was not long before the method was extended to X-ray powder diffraction. The quality of the data and the computation power available these days have allowed the technique to develop enormously, to the point that even the (successful) Rietveld refinement of small protein structures from synchrotron powder diffraction data is now possible (see Chapter 7.1). Another development is the extension of the Rietveld method towards parametric refinement on large numbers of complimentary data sets with various as-yet unexplored new applications. Rietveld refinement is so important it is described in detail in Chapter 4.7, but we describe a number of important fundamentals of the method here by way of introduction.

The basic idea behind the Rietveld method is simple: Instead of extracting the integrated intensities of Bragg peaks and fitting models to these, as would be done in single-crystal and early powder diffraction studies, the full powder pattern, for example available as step-scanned intensity data, is fitted using a model whose parameters are refined using a least-squares procedure. The model parameters are varied in such a way as to minimize the sum of the squares of the difference between the n observed Y_{obs_i} and n calculated $Y_{\text{calc}_i}(\{p\})$ step-scan intensities in the powder pattern, where the latter are calculated from a model containing a set of parameters $\{p\}$. The function that is minimized is usually the profile-weighted residual function, or R factor, given by

$$R_w = \sum_{i=1}^n w_i (Y_{\text{obs}_i}(2\theta) - Y_{\text{calc}_i}(2\theta; \{p\}))^2. \quad (1.1.91)$$

The weight w_i is derived from the variance of the values of Y_{obs_i} , while all covariances between different Y_{obs_i} values are assumed to be zero.

The calculated intensity Y_{calc_i} is expressed by combinations of mostly nonlinear and analytic or non-analytic functions as

$$Y_{\text{calc}_i} = \sum_{ph=1}^{\text{phases}} \left(S_{ph} \sum_{hkl(ph)} \left(K_{hkl(ph)} |F_{hkl(ph)}|^2 \Phi_{hkl(ph)}(2\theta_i - 2\theta_{hkl(ph)}) \right) \right) + b_i(\text{obs}). \quad (1.1.92)$$

The outer sum runs over all phases ph present in the powder pattern, while the inner sum runs over all reflections hkl of a phase ph that contribute to the intensity at the position i in the powder pattern. A scaling factor S_{ph} is assigned to the reflection intensities for each phase; the scaling factor is proportional to the weight fraction of the phase. $K_{hkl(ph)}$ represents the product of various correction factors to the square of the structure-factor amplitudes, $|F_{hkl(ph)}|^2$, which may depend on the diffraction geometry and/or individual reflections. The value of the profile function $\Phi_{hkl}(2\theta_i - 2\theta_{hkl})$ is given for the profile point $(2\theta_i - 2\theta_{hkl})$ relative to the position of the Bragg reflection hkl . The observed background at position i in the powder pattern is denoted as $b_i(\text{obs})$. Parameters in the model such as atomic positions, lattice parameters and experimental factors that affect peak shape and background are varied, using a least-squares approach, until the agreement between the calculated and measured diffraction profiles is optimized. In a least-squares approach, optimization consists of minimizing a cost function that is the weighted sum of the squared differences. This is a refinement method: a good initial guess at, or knowledge of, the structure is required and this model is refined by small adjustments.

This approach requires the modelling of the *entire* powder pattern. To simplify this complex task, the information content of

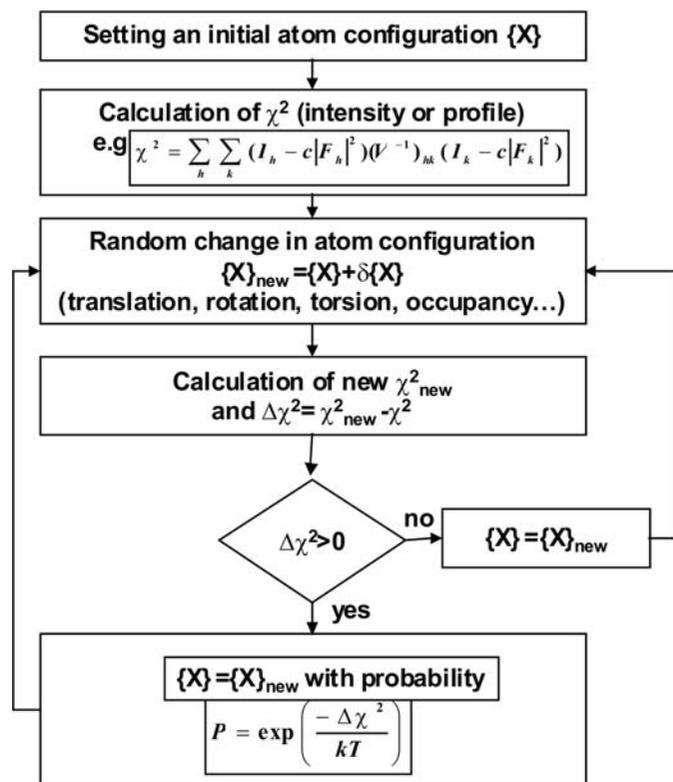


Figure 1.1.25

Flow diagram of a simulated-annealing procedure used for structure determination from powder diffraction data (from Mittemeijer & Welzel, 2012). F in the double sum is the structure factor from the structural model at each step of the optimization. Each sum runs over all reflections. h and k are summation indices representing hkl and $h'k'l'$, respectively.

the powder pattern can be divided into several parts (Fig. 1.1.1), allowing the separation of groups of parameters with respect to their origin:

- the peak intensity $|F_{hkl(ph)}|^2$ – the time- and space-averaged crystal structure and geometrical contributions;
- the peak position – crystallographic lattice and symmetry, and instrumental contributions;
- the peak shape $\Phi_{hkl}(2\theta_i - 2\theta_{hkl})$ – microstructural parameters and instrument profile;
- the background $b_i(\text{obs})$ – local structure and instrumental conditions.

Each part contains contributions from the sample and the instrument.

Rietveld refinement is a nonlinear least-squares process and requires starting values for all parameters. It is generally implemented with a local, rather than a global, optimizer and it is important for the starting parameters to be close to those of the actual solution to ensure that it is in the valley in parameter space that contains the global minimum. It is usual to guide the refinement into the (relatively narrow) range of convergence by hand by adding the parameters to the refinement sequentially. In this sense, Rietveld refinement takes some time to learn, but with care it can provide robust quantitative structures and a wealth of information can be extracted from the data.

Of course, there is no reason (other than computational efficiency) why the minimization algorithm could not be a more robust global optimizer, and this is now starting to be implemented in modern Rietveld codes. The most common and most easily implemented global optimizer, though one of the least efficient, is the Metropolis or simulated-annealing (SA) algo-

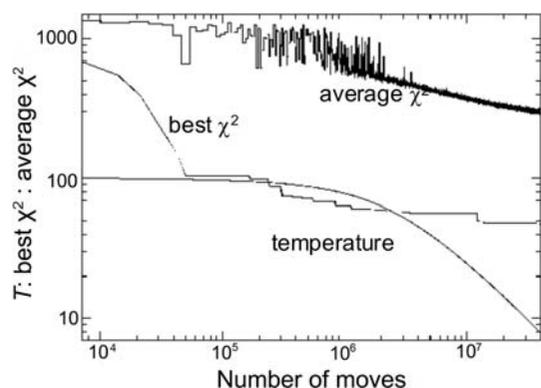


Figure 1.1.26

χ^2 (cost function) and ‘temperature’ dependence of the number of moves during a simulated-annealing run. [From Mittemeijer & Welzel (2012). Copyright Wiley-VCH Verlag GmbH & Co. KGaA. Reproduced with permission.]

rithm. The most usual implementation is actually as a ‘regional’ optimizer where the updates to parameters such as atomic position are constrained to be not too far from the previous values in such a way that the algorithm makes a random walk through the parameter space. This algorithm can avoid being trapped in a local minimum by ‘walking uphill’, since changes to the parameters that produce a worse agreement may be accepted with a probability based on the Boltzmann criterion, $\exp(-\Delta R/kT)$. The temperature in this expression is fictitious (*i.e.*, it does not refer to any real temperature) and ΔR is the change in the agreement produced by the trial update. The temperature plays the role of tuning the probability of accepting a bad move. It is initially chosen to have a high value, giving a high probability of escaping a minimum and allowing the algorithm to explore more of the parameter space. Later in the run the temperature is lowered, trapping the solution into successively finer valleys in the parameter space until it settles into (hopefully) the global minimum (Fig. 1.1.26). The calculation of R can be based on the entire profile, or on integrated intensities. For the latter, the correlation between partially or fully overlapping reflections must be taken into account (as shown schematically in Fig. 1.1.25).

A flow diagram of a typical SA algorithm as used for structure determination from powder diffraction data is shown in Fig. 1.1.25. Parameters that can be varied during the SA runs include internal and external degrees of freedom like translations (fractional coordinates or rigid-body locations), rotations (Cartesian angles, Eulerian angles or quaternions, describing the orientation of molecular entities), torsion angles, fractional occupancies, displacement parameters *etc.* Fig. 1.1.26 shows the results of a typical simulated-annealing run in which the cost function, χ^2 , falls dramatically in the first few thousand moves, indicating that the scattering is dominated by the positioning of heavier atoms or globular molecules. Several million trial structures are usually generated before a minimum can be reached. At the end of the simulated-annealing run, Rietveld refinement is used to find the bottom of the global minimum valley.

Special algorithms are not usually used to prevent close contact of atoms or molecules during the global-optimization procedure, as in general these have not been found to be necessary, as the fit to the intensities alone quickly moves the molecules to regions of the unit cell where they do not grossly overlap with neighbouring molecules. A subsequent Rietveld refinement in which only the scale and overall displacement parameters are refined will immediately show whether further

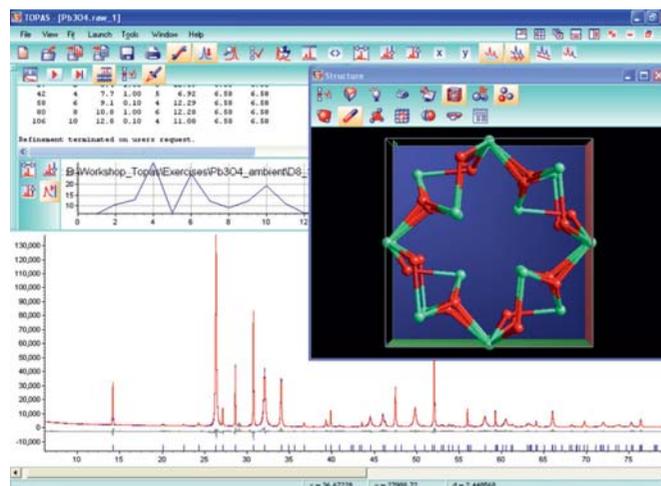


Figure 1.1.27

Screen shot (*TOPAS 4.1*; Bruker-AXS, 2007) of a simulated-annealing run on Pb_3O_4 measured with a D8 advance diffractometer in Bragg-Brentano geometry. [From Mittemeijer & Welzel (2012). Copyright Wiley-VCH Verlag GmbH & Co. KGaA. Reproduced with permission.]

refinement of bond lengths and bond angles is necessary. Since unconstrained refinement often results in severe distortions from the ideal molecular geometry, either rigid bodies or soft constraints on bond lengths, the planarity of flat groups and bond angles can be used to stabilize the refinement. Another advantage of the simulated-annealing technique is that hydrogen atoms can often be included at calculated positions from the beginning if their relative position with respect to other atoms can be anticipated, which is often the case for molecular structures.

For inorganic crystal structures in particular, the identification of special positions or the merging of defined rigid bodies is useful during the final stages of structure solution. This can be accomplished by a so-called ‘occupancy-merge’ procedure as proposed by Favre-Nicolin & Černý (2004; see also Chapter 4.5). Here, the occupancies of the sites are modified as a function of the fractional coordinates, *i.e.* they are changed when the atoms get ‘too close’ to a special position. The sites are thought of as spheres with a radius r . In this way any number of sites can be merged when their distances are less than $2r$. As an example, the crystal structure solution of minium (Pb_3O_4) is shown in Fig. 1.1.27. In this example, special positions are identified when two oxygen or lead atoms approach within a distance less than the sum of their respective merging radii, which is estimated as 0.7 \AA . The occupancies of the sites then become: $1/(1 + \text{intersection fractional volumes})$.

The power of the Rietveld approach lies in its ability to extract the maximum information from the region of the data where peaks overlap. Since peak overlap is a significant problem even at moderate d -spacings, this method revolutionized powder diffraction to the point where the quantitative results are often trusted more than those coming from refinements of single-crystal data, since they are less sensitive to factors such as extinction that can affect single-crystal structure refinements. Single-crystal data are still preferred for structure solution, but Rietveld refinement is often the method of choice for obtaining the fine quantitative details of the structure after a solution has been found. However, the Rietveld method has also opened the door to using powder data for structure solution. In structure-resolution methods, the structure factors are calculated from the intensities of all the available peaks, and algorithms are used to find the missing phases for each of these peaks and therefore the positions of the atoms in the unit cell. As mentioned above, full

1. INTRODUCTION

profile fitting following the Rietveld method can be carried out without a model, where the ‘parameters’ are the Bragg-peak intensities themselves; this is known as Pawley or Le Bail refinement, depending on details of the approach used (see Chapter 3.5). This allows more accurate determination of the structure factors from Bragg peaks in regions where there is significant peak overlap.

These days, with high-quality data from synchrotron X-ray sources and excellent algorithms (either direct methods or global-optimization methods in direct space), determination of even quite complex crystal structures from powder diffraction data is becoming a routine method in almost all branches of natural sciences and engineering. The success rate mainly depends on three parameters: the choice of measurement device, how well the pattern profile is described and how good the structure-solving algorithm is. It is becoming increasingly evident that the use of highly monochromatic parallel-beam synchrotron radiation is a huge advantage for obtaining accuracy in the atomic parameters, which allows for the interpretation of bonding and reaction mechanisms. In some cases, even details like rotational disorder can be extracted from powder diffraction data if maximum-entropy methods are combined with high-resolution synchrotron data.

1.1.6.2. Local structure refinement

As described in Section 1.1.5.3.2, similar full-profile-fitting strategies are now also carried out on total-scattering data that include diffuse-scattering intensity residing in what used to be considered as the ‘background’. This is either done by taking a structural model, which may be similar to the crystal model used in the Rietveld method (but the crystallographic symmetry of the model could also be reduced) or be a discrete cluster or molecule. As with the Rietveld method, structural parameters are varied in such a way as to obtain a good fit of the calculated function to the measured one. These methods go beyond the average structure and yield information about the local structure in the material, which may be different from the long-range ordered (LRO) crystal structure (or indeed there may be no LRO structure, as is the case in liquids and glasses). They are becoming more popular as data quality and computational power increase.

Solving the structures of nanoparticles from PDF data is less well developed, although it has been demonstrated for some simple structures such as C_{60} and simple inorganic crystalline compounds. We expect that this will grow in importance in the coming years, following the trend of the Rietveld method and structure solution from powders.

1.1.6.3. Parametric Rietveld refinement

The conventional approach to analysing a set of powder patterns is to treat each powder pattern independently, thus refining the entire set of all relevant parameters for each pattern separately. Further analysis of the values of these parameters, for example fitting with empirical or physics-based functions such as fitting the temperature dependence of the ADPs with a Debye model, is then performed after the Rietveld refinements. Alternatively, all powder patterns can be subjected to refinement simultaneously, which allows the refinement of the functional dependence of external variables instead of deriving the parameters of the function from the individual Rietveld refinements afterwards. This so-called parametric or surface Rietveld refinement was first introduced by Stinton & Evans (2007). Parametric refinement offers several advantages over the traditional

sequential refinement approach because the correlation between parameters and the final standard uncertainty can be reduced by introducing simple and physically meaningful constraints and restraints. Furthermore, it is possible to refine noncrystallographic parameters such as rate constants or temperatures directly from Rietveld refinement (Stinton & Evans, 2007). Of course, introducing external constraints in this way may introduce bias into the refinement if the constraint is not valid. For example, if there is anharmonicity in the motion and the temperature dependence of the ADPs does not follow the Debye law, carrying out a parametric refinement where the Debye law is presumed will result in biased refinements. However, with careful application, this is a potentially powerful approach to maximizing the quantitative information available from powder data in complex systems. In the following, the basic concept of parametric refinement is illustrated with several examples.

If we assume a set of p_{\max} powder patterns from a single sample that have been measured as a function of the value of an external variable, *e.g.* time, temperature or pressure, equation (1.1.92) can be formally written for each powder pattern separately:

$$\begin{aligned} Y_{\text{calc},i,\text{pattern}(1)} &= \text{function}(p_{1,\text{pattern}(1)}, p_{2,\text{pattern}(1)}, \dots, p_{m,\text{pattern}(1)}) \\ Y_{\text{calc},i,\text{pattern}(2)} &= \text{function}(p_{1,\text{pattern}(2)}, p_{2,\text{pattern}(2)}, \dots, p_{m,\text{pattern}(2)}) \\ &\vdots \\ Y_{\text{calc},i,\text{pattern}(p_{\max})} &= \text{function}(p_{1,\text{pattern}(p_{\max})}, p_{2,\text{pattern}(p_{\max})}, \dots, p_{m,\text{pattern}(p_{\max})}). \end{aligned} \quad (1.1.93)$$

If a functional dependency of some of the parameters p on external variables T exists, these parameters may be expressed as functions of these variables, for example T . This functional relationship can be used to constrain together the p parameters for individual patterns measured at different temperatures, drastically reducing the number of global parameters. Equation (1.1.93) can thus be written as

$$\begin{aligned} Y_{\text{calc},i,\text{pattern}(1)} &= \text{function}(p_{1,\text{pattern}(1)}, p_{2,\text{pattern}(1)} = f(T_1, T_2, \dots, T_l), \dots, p_{m,\text{pattern}(1)}) \\ Y_{\text{calc},i,\text{pattern}(2)} &= \text{function}(p_{1,\text{pattern}(2)}, p_{2,\text{pattern}(2)} = f(T_1, T_2, \dots, T_l), \dots, p_{m,\text{pattern}(2)}) \\ &\vdots \\ Y_{\text{calc},i,\text{pattern}(p_{\max})} &= \text{function}(p_{1,\text{pattern}(p_{\max})}, p_{2,\text{pattern}(p_{\max})} = f(T_1, T_2, \dots, T_l), \\ &\quad \dots, p_{m,\text{pattern}(p_{\max})}). \end{aligned} \quad (1.1.94)$$

The cost function (1.1.91) to be minimized changes accordingly:

$$R_w = \sum_{\text{pattern}=1}^{p_{\max}} \left(\sum_{i=0}^{n-1} \left(w_{i,\text{pattern}} (Y_{\text{obs},i,\text{pattern}} - Y_{\text{calc},i,\text{pattern}})^2 \right) \right). \quad (1.1.95)$$

1.1.7. Outlook

As is evident from the above, the information content in a powder diffraction pattern is enormous. This chapter gives only an overview of the types of information about materials that can be obtained from powder diffraction data, and the various approaches mentioned here are described in greater detail in the rest of this volume. The powder community is growing, as is the number of applications of powder diffraction in all the materials