Universiteit Leiden
The Netherlands

EDIFF: a program for automated unit-cell determination and indexing of electron diffraction data

Jiang, L.H.; Georgieva, D.; Abrahams, J.P.

Citation

Jiang, L. H., Georgieva, D., & Abrahams, J. P. (2011). EDIFF: a program for automated unit-cell determination and indexing of electron diffraction data. *Journal Of Applied Crystallography*, *44*, 1132-1136. doi:10.1107/S0021889811030731

Version:	Publisher's Version
License:	Licensed under Article 25fa Copyright Act/Law (Amendment Taverne)
Downloaded from:	https://hdl.handle.net/1887/3620576

Note: To cite this publication please use the final published version (if applicable).

Journal of
Applied
Crystallography
ISSN 0021-8898

Received 2 December 2010
Accepted 31 July 2011

EDIFF: a program for automated unit-cell determination and indexing of electron diffraction data

Linhua Jiang,[a,b]* Dilyana Georgieva[a] and Jan Pieter Abrahams[a]

[a]Gorlaeus Laboratory, Faculty of Science, Leiden University, The Netherlands, and [b]Software Engineering Institute, East China Normal University, Shanghai, People's Republic of China. Correspondence e-mail: dr.jianglh@gmail.com

EDIFF is a new user-friendly software suite for unit-cell determination of three-dimensional nanocrystals from randomly oriented electron diffraction patterns with unknown independent orientations. It can also be used for three-dimensional cell reconstruction from diffraction tilt series. In neither case is exact knowledge of the angular relationship between the patterns required. The unit cell can be validated and the crystal system assigned. *EDIFF* can index the reflections in electron diffraction patterns. Thus, *EDIFF* can be employed as a first step in reconstructing the three-dimensional atomic structure of organic and inorganic molecules and of proteins from diffraction data. An example illustrates the viability of the *EDIFF* approach.

© 2011 International Union of Crystallography
Printed in Singapore – all rights reserved

1. Introduction

The first step in crystallographic data analysis is determining the unit-cell parameters of the crystal. The conventional method for unit-cell determination by electron microscopy requires diffraction patterns from the main crystallographic zones and knowledge of the crystal orientation associated with each of these zones. Finding such zones requires careful orientation of the crystal, a process that is made easier if the diffraction patterns show Kikuchi lines when using a convergent beam. From the spacings between the diffraction spots and the angular relationship between the patterns, the unit-cell parameters can then be calculated. If the diffraction is too weak to display Kikuchi lines in convergent-beam diffraction mode, deeper crystallographic knowledge and an understanding of three-dimensional reciprocal space are required in order to tilt the crystal to the correct or desired orientation. In any case, this method of unit-cell determination requires accurate angular information on the three-dimensional rotation angles. The tilt directions are read from the specimen holder, and well known formulae (Zou *et al.*, 2004) can then be used for calculating the angular relationship between the diffraction patterns, assuming that all were recorded from the same crystal. This remains a difficult procedure, and often X-ray powder diffraction and electron diffraction are combined in the final stage to validate and refine the cell parameters.

Several groups have developed methods for the data processing of misoriented electron diffraction patterns. The electron diffraction tomography technique (Kolb *et al.*, 2007, 2008) and the diffraction rotation technique (Zhang *et al.*, 2010) have recently attracted major interest. In diffraction tomography, patterns are acquired in STEM (scanning tunning electron microscopy) mode in order to reduce beam damage and data are collected every 1° by tilting the crystal stepwise. The technique is suited to collecting data from beam-sensitive materials, where the nature of the sample does not allow long exposures. The diffraction rotation technique aims to collect a complete data set of diffraction patterns in a fashion similar to X-ray crystallography. In this method, the beam is tilted continuously by up to a few degrees, after which the crystal is rotated and the beam is tilted back by the same angle.

New methods for processing the data and determining the unit cell of three-dimensional nanocrystals from randomly oriented electron diffraction patterns with unknown and independent orientations have already been presented by our group (Jiang, Georgieva, Zandbergen & Abrahams, 2009; Jiang, Georgieva, Ijspeert & Abrahams, 2009). Here, a new program, *EDIFF*, based on this algorithm is described. It can process data acquired from exact crystallographic zones or mistilted patterns recorded on image plates, CCD cameras or hybrid pixel detectors. The program system can be used for analysing data from crystals with small or large cell parameters. It can be employed for all crystal systems and usually performs well even if the angular relationship between the diffraction patterns is unknown.

2. Methods

2.1. Electron diffraction data recording

It is convenient (but not mandatory) when processing data with *EDIFF* to record diffraction patterns from a single crystal by tilting the beam or rotating the crystal. However, in the case of very radiation-sensitive crystals (for instance protein crystals), patterns can also be acquired from several different randomly oriented crystals. In the latter case, there are two aspects that require special attention: (i) impurities (*i.e.* crystals not sharing the same crystallographic features as the sample studied) must be recognized and discarded from the data analysis, and (ii) if the crystals have a preferred orientation on the grid, the sample must be tilted in order to record a sufficient number of different orientations. Crystals can also be rotated in continuous mode, similar to X-ray data acquisition, or stepwise, as in the tomography technique.

It is of prime importance to sample reciprocal space sufficiently densely. It is also essential that the diffraction data are recorded using the same camera length, and special care is taken with the accuracy of the eucentric height (z height) during data acquisition.

The minimum number of diffraction patterns required for successful unit-cell determination depends on the number of free parameters of the unit cell, which is determined by the crystal system. For instance, for cubic cells (one free parameter), ideally a single

diffraction pattern should be enough. In more realistic cases, with low diffraction quality and random tilt angles of the acquired data, more data will always make the algorithm more robust. For an unknown crystal system, where we have to assume space group $P1$ in the worst case, more than 30 diffraction patterns with random tilt angles of selected nanocrystals may be required.

2.2. *AMP*: pre-processing electron diffraction data

First, the data need to be pre-processed by the program *AMP* (autocorrelation mapping program), which is used to remove the background noise from the diffraction patterns and for selecting reflection spots. The algorithm is described in detail by Jiang, Georgieva, Ijspeert & Abrahams (2009). The program has a graphical user interface (GUI), shown in Fig. 1.

AMP expects electron diffraction patterns of about 1024×1024 pixels (though they need not be exactly this size) and can process images in JPEG and TIF formats. It is essential that all the patterns within a data set are acquired using the same voltage, camera length and digitization parameters. If these fundamental microscope parameters are changed during a session, the patterns must be rescaled.

The design of *AMP* allows the user to make the program more or less permissive in finding the central (nondiffracted) beam and the centre of the diffraction pattern automatically. A beam-stop removal tool is also available and can cope with any given shape of beam stop. This is especially useful when the beam centre is unknown.

The diffraction patterns may suffer from high noise levels. For most cases, the standard background removal tools are sufficient. However, if 'salt-and-pepper' noise is present in the diffraction patterns, there is an option for further smoothing of the patterns. In this case a minor loss of data may occur.

There is also an information box present, which provides feedback on the performance and progress of the program (see Fig. 1). It is possible to view the intermediate output images while running the program, which helps the user to evaluate whether the patterns contain sufficient diffraction information or are too weak, and whether the data are being processed properly.

Four output files are generated for each diffraction pattern:

(*a*) `<image-name>.ctr.png`: the centred background-removed diffraction pattern;

(*b*) `<image-name>.ctr.pks`: the peak positions of the centred background-removed diffraction pattern;

(*c*) `<image-name>.atc.jpg`: the autocorrelation map;

(*d*) `<image-name>.atc.plt`: the peak positions of the autocorrelation map.

2.3. *EDIFF*: finding unit cells

`EDIFF.exe` is the main program of the electron-diffraction (*EDIFF*) software package. It finds and optimizes the unit-cell parameters, and fits and indexes the diffraction patterns. *EDIFF* also has a graphical user interface (GUI); a snapshot is shown in Fig. 2.

The input data set of *EDIFF* is the pre-processed output data from *AMP* (see above), not the original electron diffraction data. All the input is stored in one directory and typically four files for each diffraction pattern are fed into *EDIFF*: `<image-name>.atc.plt`, `<image-name>.atc.jpg`, `<image-name>.ctr.pks` and `<image-name>.ctr.png`. For unit-cell determination only, just the `<image-name>.atc.plt` and `<image-name>.atc.jpg` files are required. The other two files, `<image-name>.ctr.pks` and `<image-name>.ctr.png`, are used for indexing of the reflections.

Two diffraction spots, together with the beam or image centre point, form a triangle which we call a 'facet'. The facet defined by the two shortest vectors (corresponding to the two reflection spots closest to the centre) is called the 'main facet'. In short, every reflection pair defines a facet. The main facet defines the smallest repeat unit of the two-dimensional lattice defined by the low-resolution spacings (but not any higher-order Laue zone that may be visible at higher resolution). The facets found from all the diffraction patterns contained in a data set are used to find the cell parameters. For a more detailed explanation, please refer to Jiang, Georgieva, Zandbergen & Abrahams (2009).

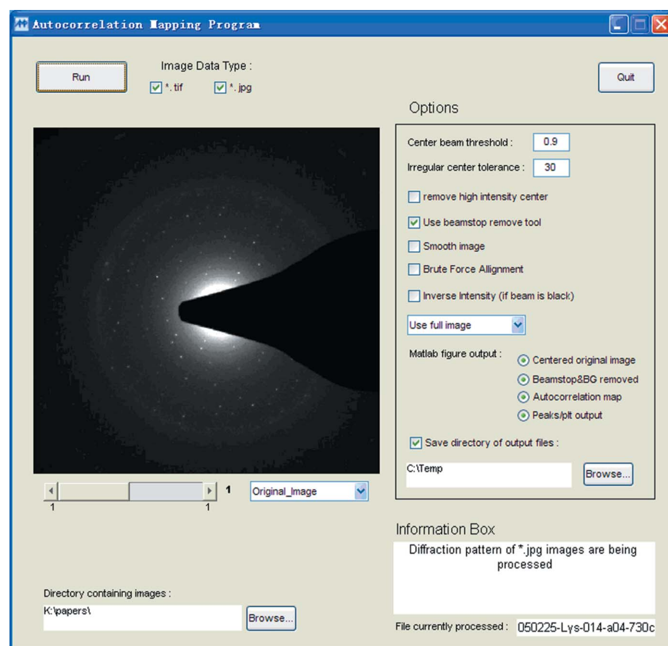

Figure 1
The GUI of the pre-processing program *AMP* (Version 1.8). The program calculates autocorrelation patterns and extracts peak positions from diffraction patterns.

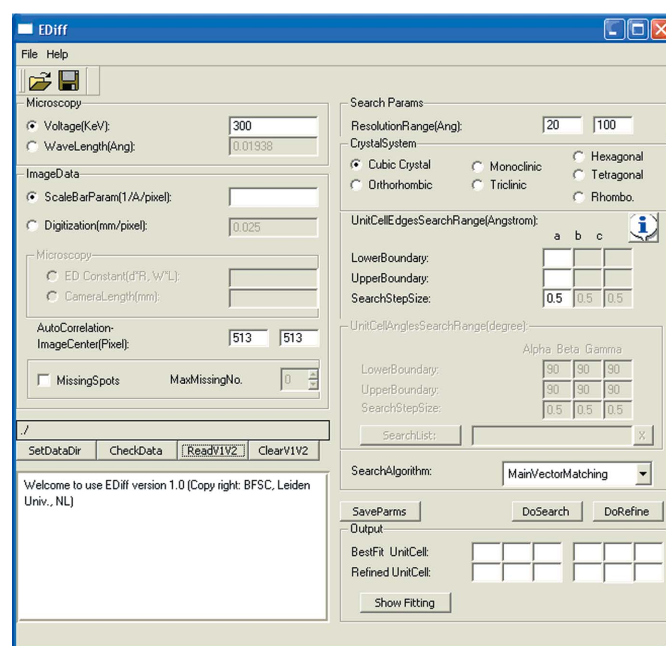

Figure 2
The main GUI of the program *EDIFF* (Version 1.0).

The accelerating voltage of the microscope (in keV), the camera length and the digitization parameters should also be supplied. The wavelength will be automatically calculated from the voltage. By its very nature, the image centre of the autocorrelation pattern in pixels is always unknown; it can only be at the origin. Through correlation analysis of the autocorrelation pattern with the diffraction pattern, the centre of the background-corrected electron diffraction pattern is found. For files generated by the pre-processing program *AMP*, the centre of the autocorrelation pattern is always (513, 513) when analysing a 1024 × 1024 pixel image, and this value can be used by default.

EDIFF first finds the facets formed by the shortest vectors in the diffraction patterns. For structures with a primitive Bravais lattice, the procedure is straightforward. However, it is more complicated for centred structures. For this case an option (MaxMissingNo.) is implemented to indicate how many spots are allowed or expected to be absent between the spots of the main vectors and the centre, if this is known. The parameter is used for two of the different methods that can be selected, MainVectorMatching and FullVectorMatching, which allow extra information on the symmetry of the crystal to be included. If this parameter is not used (MaxMissingNo. is zero), only 'prime' index reflections will be considered in the unit-cell searching. 'Prime' means that there is no common factor for the *hkl* index, *e.g.* the index 543 is prime, but 642 is not, because the *h*, *k* and *l* values all have 2 as a common divisor.

After the diffraction data have been pre-processed with *AMP*, the user can inspect the output. The purpose is to check whether the peak positions of the auto-correlation patterns and those of the background-removed diffraction patterns coincide. *EDIFF* displays the main facet (also called the 'main vectors', V1 and V2) for each autocorrelation pattern. In Fig. 3, the peaks of the spots within the resolution range are visualized as crosses, while the points of the simulated lattice generated from this main facet are indicated by small circles. If the crosses fit the circles, then the diffraction pattern is well oriented and the main facet has been correctly found by the program. It is important to set properly the resolution range for the unit-cell parameter search. This resolution range, in ångström, is used for finding the unit-cell parameters and for the indexing calculation. Only the peak positions within this range will be used for the calculation. If a very large resolution range is selected, a large number of spots is used in the calculation and this may result in lengthy calculations. If the resolution range is too narrow, the loss of essential information may prevent the finding of the correct answer. In practice, a resolution range that is a little wider than that defined by the main facet has been found to be acceptable. As a rule of thumb, a resolution range from half of the expected smallest unit-cell dimension to double the largest unit-cell dimension gives satisfactory results.

Seven different crystal systems are supported in the unit-cell parameter determination (see Fig. 2). If the crystal system is known beforehand, its constraints can be applied in the exhaustive search of edges and angles, which can dramatically speed up the calculations and improve their precision. If the crystal system is unknown, the triclinic crystal system can be chosen for searching without special constraint.

Four different search algorithms can be employed in the search for the cell parameters:

(*a*) UniqueFacetMatching: the friendliest algorithm. Verification of the data is not needed (except perhaps for finding the proper resolution range). It runs in an automated mode and gives quick results in cases of good diffraction data.

(*b*) MainVectorMatching: requires 'verification' of the data. In this case, the main vectors often have to be selected manually and the facets of each diffraction pattern that are to be used in the calculation must be saved explicitly. This option is useful for very noisy and/or marginal data.

(*c*) FullVectorMatching: uses all vectors within the resolution range. This algorithm can be rather slow. It is especially useful when the unit cell is approximately known, for instance for final refinement or comparison.

(*d*) BrightestSpotsMatching: a variation of MainVectorMatching, this is especially useful for thin nanocrystals with a large unit cell (see §2.4 below).

Parameters that can affect the performance of the algorithms can be set in a special window of the GUI (see Fig. 4). The default parameters were found to work well for most of the test data sets. Most of the names of the parameters are self-explanatory:

MaxFilesReadIn: the maximum number of `*.plt` data files that are read into memory;

MinPattFitRatio (for all patterns): the minimum good fit ratio of the total number of patterns;

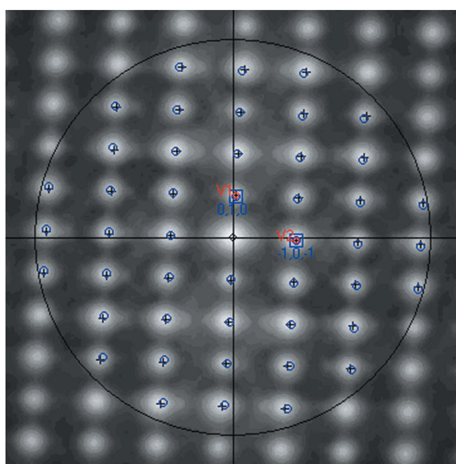

Figure 3
Pattern fitting of the main facet. The background is the autocorrelation pattern, small crosses are the peaks in the background and small blue circles are the simulated fitting pattern. Red circles labelled V1 and V2 indicate the two main vectors defining the two-dimensional lattice. Blue squares with indices show the fitting results corresponding to the two main vectors.

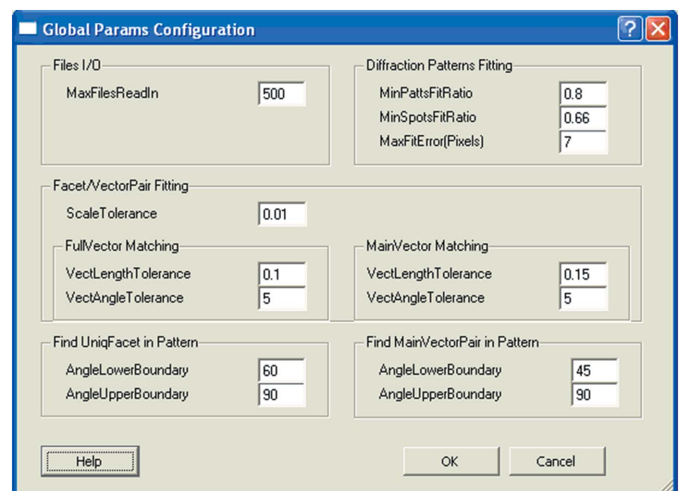

Figure 4
More detailed parameter configuration in *EDIFF* (Version 1.0).

MinSpotsFitRatio (for one pattern): the minimum spot fit ratio to detect correct fitting;

MaxFitError (for one spot): the maximum distance, in pixels, between fitted and real spots;

ScaleTolerance: the maximum value of the scale tolerance;

VectLengthTolerance: the vector length tolerance;

VectAngleTolerance: the vector angle tolerance (in degrees);

AngleLowerBoundary: the angular lower boundary (in degrees);

AngleUpperBoundary: the angular upper boundary (in degrees).

Two functions were implemented to calculate these parameters. The FacetFitting function is used to calculate the residual of fitting two facets, which is used to judge the similarity of the two facets. For the FullVectorMatching method, a vector length tolerance of 10% is usually selected, or 15% for the MainVectorMatching method. In both cases, the default value for the vector angle tolerance is 5°. The Lattice2MainFacet function is used to generate main facets from electron diffraction patterns for the Unique Facet Matching algorithm calculation. The angle of the facet is limited by the AngleLowerBoundary and AngleUpperBoundary parameters. The percentage of the minimum number of spots fitted for one pattern, as well as the minimum good fit ratio of the total number of patterns, can be chosen or varied (see Fig. 4.)

Upon completion of the search procedure, a list of possible cell parameters or solutions is presented and the cell parameters can be further refined, starting from the cell parameters with the best fit.

2.4. BrightestSpotsMatching

The BrightestSpotsMatching algorithm is a variation of the MainVectorMatching algorithm. It was designed for thin nanocrystals with large unit cells. Plate-like nanocrystals often adopt a preferred orientation on the grid. If this is the case, the information on the unit-cell dimension in the direction of the beam is not (well) determined. In order to obtain sufficient information about the third dimension, the crystal must be tilted from the main zone. Moreover, if the unit cell is large, the reflections will often be elongated in the direction normal to the plane, especially if the electron beam converges or diverges at the back focal plane (exact parallel-beam illumination can be hard to obtain), owing to the 'spike function' (or 'rocking curve' in X-ray terms). Therefore, the positions of the reflections present in the diffraction patterns may not represent the centroid in φ of the reflection. This may compromise the other algorithms, as may the fact that information on the unit-cell dimension normal to the plane may be missing (a few ultra-high tilted diffraction patterns may still have some useful information, but may be of very poor quality because of the high tilt).

The BrightestSpots algorithm abandons the idea of using the main facets, and instead uses the two brightest spots in the diffraction pattern (not the autocorrelation pattern). The brightest spots are most likely to have their centroids closest to the Ewald sphere and can be further from the centre, thus containing higher index information, also in the direction normal to the plane of the crystal. The main practical difference is that the two brightest spots are set as the new main facet (V1 and V2). This algorithm was also found to give reliable results when the experimental data were very noisy and many randomly oriented diffraction patterns were collected.

2.5. Pattern fitting of the autocorrelation pattern to verify the unit-cell parameters

After the unit-cell parameters have been found, their accuracy can be visually inspected by fitting the auto-correlation patterns with the lattice derived from the refined unit cell. Effectively, the PatternFitting algorithm indexes the auto-correlation pattern (see Fig. 3).

This PatternFitting algorithm works as follows. Firstly, it finds the facet in a simulated three-dimensional lattice that best fits the main facet of an autocorrelation pattern, then it cuts through the three-dimensional model lattice using the plane defined by the best fitting facet to generate a two-dimensional diffraction simulation.

In Fig. 3, crosses mark the peaks of the spots of the autocorrelation pattern and blue circles mark the model lattice. If the unit cell is correct, the crosses and blue circles should overlap well. If most of the autocorrelation patterns (>90%) fit well, the unit cell is most probably correct. The outliers can be caused by the inclusion of some deviant crystals in the data set, by poor crystals with streaked rocking curves, by unfortunate orientations or by unexpected failures to index the pattern properly. Only an experienced eye can tell. However, it is important to ascertain that information about all the cell parameters is present in the original data, *i.e.* the reciprocal lattice is sampled sufficiently well.

If the main facet (V1 and V2) does not fit well (*i.e.* upon a visual check) and there are uncertainties about the correctness of the indices, the GUI allows the user to inspect all other potential fits using a spin-box. The fitted facets are sorted according to how well they fit the diffraction data.

We have also implemented a function called MosaicType, which encodes different models for simulating the mosaic spread of the diffraction pattern. It defines whether a reflection off the Ewald sphere should be shown in the simulated diffraction pattern or not. The criteria can be (i) the angular error of a reflection; (ii) the absolute reciprocal distance between a reflection and the Ewald sphere; (iii) side elongation: the function simulates the spot elongation along the main direction of the unit cell that lies closest to the direction of the electron beam and calculates the reciprocal distance to the Ewald sphere. This function was found to be useful in the case of very thin plate-like crystals that have characteristics of two-dimensional crystals.

2.6. Indexing of electron diffraction patterns from a crystal with known cell parameters

Once the unit-cell parameters have been inferred, the diffraction data can be indexed.

The main difference between PatternFitting (Fig. 3) and IndexingRefinement (Fig. 5) is that in the former an autocorrelation pattern is indexed and in the latter a diffraction pattern is indexed. The default background of the IndexingRefinement window is the diffraction pattern, while the default background for PatternFitting is the autocorrelation pattern.

After having found the unit-cell parameters, the global ResolutionRange (set in the *EDIFF* main window) can be increased

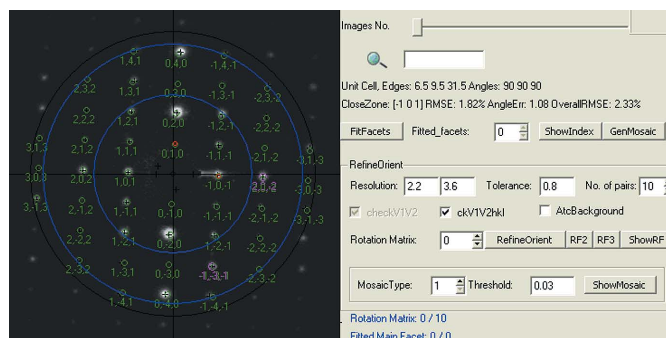

Figure 5
Indexing refinement of the diffraction pattern. Red spots indicate the main facet and purple spots denote a matching spot pair in a higher-resolution area.

for IndexingRefinement. The SearchAlgorithm should be changed to MainVectorMatching, irrespective of which algorithm was used for obtaining the unit-cell parameters. It is necessary to select a main facet on which the indexing refinement will be based.

First, a rough fit is performed. This is the same as in the PatternFitting procedure: the program finds the best-fitting facet in a simulated three-dimensional unit-cell model for the main facet (V1 and V2), then cuts through the three-dimensional model lattice along the plane defined by the selected facet in order to generate a simulated two-dimensional diffraction pattern (shown as small blue circles). If the diffraction pattern is taken directly from a main zone, this provides accurate indexing. However, it is more usual for the experimental diffraction pattern to be tilted away from a main zone. In order to find the exact orientation of an individual diffraction pattern (so as to index the reflections correctly), it would be necessary to identify the misorientation.

Potential orientations that the program has found are stored in a list of rotation matrices, sorted according to their quality of fit. It is possible to inspect all these orientations and check how well they match the diffraction pattern by selecting their sequence number in the RotationMatrix spin-box. Most of the time, the orientations highest on the list fit best. It is sometimes possible to find more than one orientation or indexing solution that fits well. This depends on the unit-cell parameters, and unfortunate combinations can exist in which it is not possible to distinguish on the basis of the positions of the diffraction spots alone. In this case, intensity information must be included in addition to three-dimensional merging and scaling of the diffraction data. This is not yet implemented.

3. Results and discussion

EDIFF has been tested on crystals of different compounds. Only in very difficult cases (poor sampling of the three-dimensional reciprocal lattice, poor streaked diffraction patterns) did we fail to find satisfactory solutions. Here, we discuss a single straightforward example to illustrate the operation of the program. The example is based on experimental data collected from penicillin G potassium. The crystals are known to have an orthorhombic unit cell (a = 6.342, b = 9.303 and c = 30.015 Å), as measured using X-ray crystallography. The sample is beam-sensitive and tilt series could not be acquired from a single nanocrystal. Instead, random orientation data were collected from different crystals. The diffraction patterns used for the analysis were of comparable quality regarding signal-to-noise ratio (diffraction reflections could be clearly resolved) and resolution. Calibration of the camera length was performed using the ring pattern of an aluminium standard sample.

For the pre-processing of the data, the program *AMP* was run with default values. Smoothing of the pattern was not applied in this case.

In total, 13 diffraction patterns of penicillin G potassium were acquired (Jiang, Georgieva, Zandbergen & Abrahams, 2009), each of different unknown crystal settings. The data were analysed with the three search modes. The autocorrelation maps generated from the electron diffraction patterns were of good quality, allowing straightforward determination of the main vectors. Therefore, the data were subjected first to analysis with the MainVectorMatching mode (the search with this mode is much quicker than that with the other two). In the next step, the FullVectorMatching and UniqueFacetMatching search modes were employed. The results are given in Table 1.

From each search mode, *EDIFF* suggests five different solutions with the lowest residuals; the first three are presented in Table 1. The results for two of the cell dimensions (6.3 and 9.3 Å) were always found, irrespective of the algorithm used. Information about the third, longest, axis was rather scarce in the diffraction patterns, but still present, allowing the correct solution to be found.

Table 1
Output of the searches performed for finding the cell parameters of penicillin G potassium using the three different algorithms implemented in *EDIFF*.

Search mode	MainVectorMatching	FullVectorMatching	UniqueFacetMatching
Unit-cell parameters (Å, °)	6.3 × 9.5 × 29.2; 3 × 90	6.3 × 9.3 × 29.6; 3 × 90	6.4 × 9.4 × 30; 3 × 90
	6.3 × 9.5 × 29.1; 3 × 90	6.3 × 9.3 × 29.7; 3 × 90	6.4 × 9.4 × 29.9; 3 × 90
	6.3 × 9.6 × 28.6; 3 × 90	6.3 × 9.3 × 29.6; 3 × 90	6.4 × 9.4 × 30.1; 3 × 90

The pattern fitting to the autocorrelation pattern was also satisfactory, as seen in Fig. 3.

After initial unit-cell determination, the data were subjected to indexing refinement. The PatternFitting procedure indexes the main facet and thus finds a rough orientation of a diffraction pattern. Only low-resolution information is used in this PatternFitting routine, as the main facet contains two spots close to the centre. In order to find an accurate orientation (which leads to a refined indexing solution), reflections to a higher resolution need to be included (Fig. 5).

Different types of crystal data were used to test the program, including both organic and inorganic nanocrystals: lysozyme, penicillin, mayenite, aspirin, ferric phosphate and oxacillin. Most of the samples used to test the program were cubic or orthorhombic crystals. To test the robustness of the software, the angles were allowed to vary on cubic and orthorhombic crystals (choosing 'triclinic' or 'monoclinic' for the crystal system) and angular values of 90.0 (5)° were obtained.

Empirically, three factors are of prime importance for the accuracy of the results obtained when using our approach. These are the proper calibration of the camera length, the accuracy of digitization of the data (which will directly influence the d spacings in the electron diffraction patterns or the length of the main vectors) and the completeness of the data.

4. Conclusions

A method for automated unit-cell determination and indexing of electron diffraction data has been implemented in a computer program called *EDIFF*. As shown in the example, the method provides a practical way of determining the unit cell of a crystalline phase from electron diffraction data, and it is especially suited for beam-sensitive materials.

This work was supported in part by The Netherlands Organization for Scientific Research. The authors thank Dr R. A. G. de Graaff for revising the manuscript. For access to the program, please contact JPA (abrahams@chem.leidenuniv.nl). *EDIFF* is free for academic research and a licence fee is required for commercial use.

References

Jiang, L., Georgieva, D., Ijspeert, K. & Abrahams, J. P. (2009). *CISP'09, 2nd International Congress on Image and Signal Processing*, Tianjin, People's Republic of China, 17–19 October 2009. doi:10.1109/CISP.2009.5301421.
Jiang, L., Georgieva, D., Zandbergen, H. W. & Abrahams, J. P. (2009). *Acta Cryst.* D**65**, 625–632.
Kolb, U., Gorelik, T., Kübel, C., Otten, M. T. & Hubert, D. (2007). *Ultramicroscopy*, **107**, 507–513.
Kolb, U., Gorelik, T. & Otten, M. T. (2008). *Ultramicroscopy*, **108**, 763–772.
Zhang, D., Oleynikov, P., Hovmöller, S. & Zou, Z. (2010). *Z. Kristallogr.* **225**, 94–102.
Zou, X. D., Hovmöller, A. & Hovmöller, S. (2004). *Ultramicroscopy*, **98**, 187–193.