



Sharing powder diffraction raw data: challenges and benefits

Miguel A. G. Aranda

J. Appl. Cryst. (2018). **51**, 1739–1744



IUCr Journals
CRYSTALLOGRAPHY JOURNALS ONLINE

Copyright © International Union of Crystallography

Author(s) of this paper may load this reprint on their own web site or institutional repository provided that this cover page is retained. Republication of this article or its storage in electronic databases other than as specified above is not permitted without prior permission in writing from the IUCr.

For further information see <http://journals.iucr.org/services/authorrights.html>



Sharing powder diffraction raw data: challenges and benefits

Miguel A. G. Aranda*

ALBA Synchrotron, Carrer de la Llum 2-26, Cerdanyola del Vallès, Barcelona 08290, Spain. *Correspondence e-mail: migarcia@cells.es

Received 25 July 2018

Accepted 2 November 2018

Edited by A. J. Allen, National Institute of Standards and Technology, Gaithersburg, USA

Keywords: raw data; powder diffraction; data processing.

Supporting information: this article has supporting information at journals.iucr.org/j


Scientific data are as important as scientific publications. If this statement holds true, why are we not routinely sharing scientific data? The tools are now out there, for instance Zenodo and related repositories. It could be a lack of motivation of researchers derived from an apparent lack of short-term reward. Here the author will try to show the importance of sharing ready-to-analyse raw powder diffraction data with immediate benefits for authors and for the wider community. Moreover, it is speculated that sharing curated scientific data may have more important medium-term benefits, including credibility and not least reproducibility. Raw data sharing is coming.

1. Introduction

It is nine years since *Nature* dedicated an editorial and a section to preserving research data and making it accessible (Nature Editorial, 2009). Initially, sharing raw data was conceived mainly for helping experiment replication and data analysis improvement. However, the arrival of artificial intelligence and machine-learning tools makes sharing scientific data even more important as new (unexpected by the original research teams) correlations could emerge when interrogating many related and shared data sets (Warren, 2018).

Warren (2018) in his article based on the Fred Kavli Distinguished Lectureship in Materials Science stated ‘... The existing publication paradigm is an accident of history. If you were designing the scientific publication system today instead of letting it evolve over more than 500 years, it might be a little bit different, optimized toward better scientific outcomes.’ I cannot agree more, and to me this is the main explanation of the fact that, as of today, most researchers are not releasing the raw data associated with their publications. It may also justify why well reputed journals are not requesting the sharing of raw data associated with a publication either.

Funding agencies are starting to request research data sharing, although this seems to be still in its infancy. As an example, the European Commission has very recently launched Recommendation C(2018)2375, adopted on 25 April 2018, on access to and preservation of scientific information (<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32018H0790>). Open data is just a subset of a much larger framework of ‘open science’, which has several main pillars: (i) open access (for publications); (ii) open data (for replication and new learning); (iii) open science evaluation (for improving metrics and impacts); and (iv) open science tools (for repositories, data services *etc.*). An overview of open science is



available at several sources, including Wikipedia (https://en.wikipedia.org/wiki/Open_science) and the European Commission (<https://ec.europa.eu/research/openscience/index.cfm>).

From now on I will focus just on scientific data sharing. The International Union of Crystallography (IUCr) has been leading efforts for many years through its journals such as *Acta Crystallographica Section C* in the publishing and sharing of data (both reduced and derived) linked to the scientific publication narrative in words. In the crystal structure determination field, the raw data are the diffraction images/patterns, the reduced data are the structure factors, and the derived data are the atomic coordinates, atomic displacement parameters and atom occupancies (Kroon-Batenburg *et al.*, 2017). Building on all this, with the advent of new, huge, digital storage opportunities, there have been several recent reports underlining the importance of sharing raw diffraction data [see the article by Helliwell *et al.* (2017), and references therein, which provides a wide range of case studies in crystallography and related fields]. These case studies span many of the IUCr Commissions and fall into three general categories: those where data sharing is beneficial, those where data preservation is important in allowing further progress and those where the absence of data is a significant problem. Furthermore, the necessity of having accurate metadata associated with the raw data for their correct processing has been highlighted (Kroon-Batenburg *et al.*, 2017).

Concerning powder diffraction data, the International Centre for Diffraction Data (ICDD) has been collecting these types of data since 1941 in evolving formats (Bruno *et al.*, 2017). The data are stored in the Powder Diffraction File database, housing more than 890 000 (reduced) diffraction patterns and 330 000 crystal structures in the 2016 release. A timely development for quite some years has been the archiving of more than 11 000 (unreduced) raw powder diffraction patterns. These data can be used for different purposes. For instance, the crystal structure of trandolapril was solved from its archived raw powder diffraction data (Reid *et al.*, 2016), highlighting the utility of raw data deposition in the Powder Diffraction File.

In this context, the IUCr has over many years developed tools to facilitate the sharing of diffraction data in general and powder diffraction data in particular. For instance, the original development of the CIF format (Hall *et al.*, 1991; Bernstein *et al.*, 2016), already adopted in 1990 for storing and distributing crystallographic derived data, has evolved to be able to archive and share diffraction data (<https://www.iucr.org/resources/cif>). This mechanism is being extended and updated and now also includes a dictionary specifically for powder diffraction, pdCIF (see https://www.iucr.org/resources/cif/dictionaries/cif_pd).

Data sharing goes beyond replicating and improving analysis results. It can yield new science as machine-learning tools will obtain new outcomes when multiple data sets are investigated, our science's version of big data. However, there are intermediate stages as new databases will be created just by sharing curated raw data in appropriate repositories. For instance, the proposal for constructing an international X-ray

absorption fine structure (XAFS) database (Asakura *et al.*, 2018) is very interesting. This database could be developed by sharing curated raw XAFS data recorded at synchrotrons, ideally in cross-validated beamlines. In my opinion, to have such a common international database should be a priority.

2. Some definitions and scope

Raw data is a very difficult term to define, as even the first data set store out of a detector (normally termed 'primary data') can be already processed by the firmware of the detectors. These 'primary data' can be corrected (if needed, for instance dark and flat fields), and then pre-processed and post-processed to give the final 'ready-to-analyse scientific data'. In some disciplines these are termed reduced data but this is not the case in other fields. In single-crystal work, for example, there are a variety of benefits of not just taking the predicted diffraction spot positions and their intensities; the whole diffraction image instead is preserved. Thus far it is deemed reasonable to assume that the detector firmware corrections are well established and acceptable and what we might call the very primary data before those corrections are applied need not be preserved.

Let us now consider other fields. A full pipeline for data processing (and analysis) for synchrotron X-ray full-field tomographic microscopy has been recently reported (Marone *et al.*, 2017). For tomographic work, including X-ray diffraction computed tomography, key steps include the sinogram generator in the pre-processing step and the reconstruction module in the processing step. Note that tomographic beamlines can generate more than ten terabits of 'raw and processed data' in a single day. Then, data analysis takes over, with visualization, segmentation, understanding *etc.*

In the powder diffraction field, the terms raw data, reduced data and derived data are still under debate, and here I give my view. On the one hand, raw data could be considered any patterns (processed at different levels) which still keep their (intensity of scattered photons *versus* diffraction angle) data-point character (see Fig. 1). On the other hand, there are many types of derived data depending upon the application: (i) atomic parameters in structure determination; (ii) phase contents in quantitative phase analysis; (iii) average coherent diffraction domain size and microstrains in microstructural analysis; and so on. Reduced data could be exemplified by a list of diffraction peak positions and intensities. A similar approach could be taken for the X-ray absorption field, where raw data could be considered any pattern conserving its transmitted/emitted photons *versus* energy data-point character. Therefore, incident intensity, I_0 , correction and energy calibration are processing steps yielding still (processed) raw data.

If raw data are in fact several (related) data sets (for instance, primary, pre-processed, post-processed), a key question arises: which 'raw data' should be shared by a meticulous researcher? There is no community-agreed answer but I offer my opinion. Firstly I state a caveat that concerns me as current Scientific Director of a large facility (the ALBA

synchrotron). Data sharing is not free and someone (probably the funding agencies rather than the individual researcher) will pay the costs. The more raw/scientific data we store, the more funding will be needed. So, I think that we have to pay special attention, at least at large facilities, with regard to which raw data to store and how we store them. Furthermore, most science budgets are capped, and this is an additional reason to be efficient. Every large facility is developing a data policy plan and the ALBA data policy was approved in July 2017; it can be found at <https://www.cells.es/en/users/call-information>. At this point it should be mentioned that many data sets may never be used in publications. Archived raw data (and metadata) will be made openly available after a number of years of embargo (three for most synchrotron facilities in Europe). This embargo time could be extended if a justified request by the researchers is issued. Some data sets labelled as tests, alignment *etc.* will not be archived and therefore these data will not be made openly available.

To make things a little easier for facilities but perhaps not straightforward for researchers, the HDF5 file format is becoming a standard, where data sets with different levels of processing are stored in the same file with a hierarchical structure (Könnecke *et al.*, 2015). The outcome of one experiment can be several data sets (for instance in HDF5 file format) with subsets. It could be possible to make openly available some data sets, associated with a given publication, while keeping other data sets under embargo. However, it is not foreseen that a subset of data within a data set would be published. Therefore, researchers must be aware of, and aim at having, the proper granularity when acquiring raw data. So, coming back to our question, an obvious answer could be 'primary raw data', in other words the images/data directly out of the detector. These data should be archived with all associated metadata for data processing and sample characterization. However, are these data ready to be analysed by peer scientists and to be easily interrogated by artificial intelligence and machine-learning tools? My answer, today, is still no. Therefore, I advocate compulsory sharing of processed powder diffraction data associated with publications, knowing that some flaw(s) in the processing steps could exist. There are reasons for this opinion, the most important being to facilitate use of the data. Primary raw data, and processed data at different levels, could be stored by the facilities (at their data centres or in the cloud) and retrieved on demand when needed (for free?). Additional views in this complex issue are given elsewhere [for instance Kroon-Batenburg *et al.* (2017)].

When the primary raw data have been archived by a facility, the link between the experimental data (primary raw data and metadata) and the publication can be made by providing the doi of the experimental data in the publication. The final goal of the facilities is to provide the doi to the user as soon as the data are produced, but this is still not implemented. For example, ESRF is currently archiving experimental data for macromolecular crystallography beamlines with their associated doi (<https://doi.esrf.fr/>). See, for instance, <https://doi.org/10.15151/ESRF-ES-86533633> for the output of one experiment with the data set(s) archived in 2018 and so under

embargo until 2021. However if only primary data are shared (which will require processing steps that do not need to be necessarily user friendly for every synchrotron user), the risk of a very limited outcome from this (big) effort should not be underestimated. This could lead to a poor perception of scientific research by the layperson and concerns from funding agencies. At this stage, however, it is also fair to say that the new digital storage centres have only recently been made available, and researchers have not taken advantage of them much as yet. Thus the coming decade, say, is in effect a pilot period of gaining much more experience of the opportunities available.

Therefore, the scope of this communication is now restricted to sharing of ready-to-analyse, fully processed when needed, powder diffraction data. A ready-to-analyse data set is defined as data that do not need further treatment/processing in order to be analysed by the software of the powder diffraction scientific community (Rietveld, pair distribution function, auto-indexing *etc.*).

3. Sharing raw powder diffraction data: what raw data?

We have to lead by example. Since mid-2017 our research group in eco-cements has been openly sharing all scientific diffraction data at Zenodo prior to the submission of all of our manuscripts. Therefore, normally in the supplementary information, we describe every file deposited and give the doi link to the Zenodo-deposited diffraction data sets, including sample description and experimental conditions. For instance, we studied gels in cements by the synchrotron radiation pair distribution function, PDF (Cuesta *et al.*, 2017), where 12 processed raw diffraction data sets were deposited. These data sets included ten patterns for cement pastes, plus the data for the empty capillary for data pre-analysis, and the pattern for the nickel sample employed as standard for data analysis assessment, which was recorded under the same experimental conditions (<https://doi.org/10.5281/zenodo.890584>). These diffraction data were collected at the MYTHEN-II detector system of the MSPD beamline at the ALBA synchrotron. Our detector contains six modules with 1280 channels each (7680 data points). The standard protocol at ALBA for PDF data acquisition (a macro) is to collect 72 patterns at different starting angular positions, with more acquisitions at higher angles, taking 30 s per position. Considering the recording time and motor movement, a data set takes 37 min. 'Primary data' are stored on the hard disk from the six modules, merged by the firmware, for every starting position. Then the 72 patterns are merged together, with local software, to give a total diffraction raw powder pattern. To improve the statistics, and to mitigate problems due to software failures and beam dumps, five patterns are usually collected and merged (185 min). The final fully processed 'ready-to-analyse' powder diffraction pattern is a single ascii file after all this data processing. I advocate, at present, that this is the scientific data to be deposited and shared and not the 360 (72 × 5) 'primary' raw powder diffraction data sets or the five processed powder diffraction data sets.

A summary of the information-processing scheme for synchrotron powder diffraction, from primary raw data to derived data, is given in Fig. 1. The primary data strongly depend on the detector used and, for just one point detector, raw data processing is probably not needed. For several point detectors, just merging of the primary data is likely to be required. On the other hand, for 2D detectors, several processing steps are needed for attaining validated raw powder diffraction patterns (see Fig. 1). There is an extensive tradition in the crystallography field of sharing derived data through a number of databases. The description of these databases is out of the scope of this manuscript and the reader is referred to specific publications (Gražulis *et al.*, 2009; Hellenbrandt, 2014; Glasser, 2016). However, the situation is evolving and the Cambridge Structural Database (Groom *et al.*, 2016), initially focused on sharing derived data (crystal structures with carbon–carbon bonds), is moving forward by also archiving diffraction data in CIF format. Therefore, it could be a natural evolution that the Crystallography Open Database (Gražulis *et al.*, 2009), also initially centred on sharing derived data, could evolve by sharing raw diffraction data.

A second example is a study comparing the accuracy in Rietveld quantitative phase analysis, RQPA, by using strictly monochromatic Mo and Cu radiation, with synchrotron patterns as benchmarks (León-Reina *et al.*, 2016). This work was the basis for a chapter in *International Tables for Crystallography*, Vol. H (León-Reina *et al.*, 2019), in relation to which we have deposited 81 processed raw powder diffraction data sets from three different diffractometers: laboratory Cu $K\alpha_1$, laboratory Mo $K\alpha_1$ and synchrotron radiation (<https://doi.org/10.5281/zenodo.1291899>). In this case, and for

the synchrotron diffraction patterns, a MYTHEN-II detector system was also used with an RQPA data acquisition protocol: data were collected from three angular positions with a total acquisition time of 20 min. Three patterns, taken at different positions along the capillaries, were collected for each sample to ensure the homogeneity of the filling of the capillary. For the synchrotron data, every deposited diffraction pattern came from nine ‘primary’ data sets.

A third powder diffraction example is even more tricky (Cuesta *et al.*, 2018). In this case, PDF data for cement pastes were recorded at the ID15A beamline (ESRF synchrotron), which is equipped with a Pilatus3 X CdTe 2M hybrid photon-counting two-dimensional detector. In this case eight processed raw diffraction data sets were deposited: six patterns for cement pastes, one pattern for the empty capillary and one pattern for the nickel standard (<https://doi.org/10.5281/zenodo.1255629>). For the PDF study, the following data acquisition protocol was followed: Eight images were collected for each paste, with an acquisition time of 8 s per image, requiring approximately 1 min per pattern with very good statistics. The images were added together with outlier elimination, to remove artefacts from, for example, cosmic rays and decays of thorium contained in the granite upon which the diffractometer is mounted. This detector has 24 modules and each module is composed of six wire-bonded submodules, which makes the pixel size in the bonding region of the submodules three times larger. The detector has an energy cut-off which is typically set at half the incident energy to avoid over- or undercounting photons of wafer borders. The consequence of this (firmware-based) cut-off is to eliminate the vast majority of sample fluorescence (and part of the Compton contribution), particularly when working at high

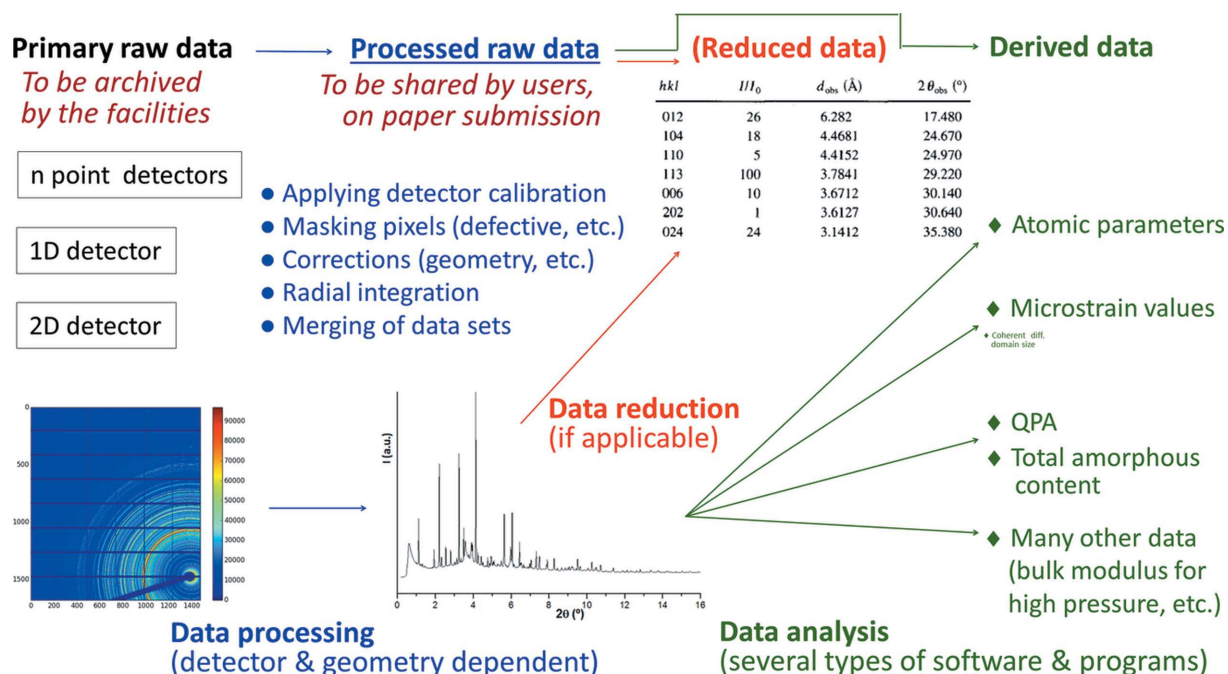


Figure 1 Summary of the information-processing scheme for synchrotron powder diffraction.

energy. These ‘primary data’ (quite clean already) were treated with local software by applying a mask to avoid collecting problematic regions (see bottom-left corner of Fig. 1): (i) the gaps between the modules; (ii) the edge pixels of the modules; (iii) the low-resolution (large) pixels due to the wire bonding of the submodules; and (iv) the dead and defective pixels. These ‘raw data’ were further processed by correcting (i) polarization of the X-rays and (ii) detector geometry, response and transparency. Finally, all the processed images were radially integrated with outlier elimination to yield one-dimensional powder diffraction patterns and merged for a given sample (see Fig. 1). Taking into account this powder diffraction data processing pipeline, I advocate that the final data set for each sample (fully processed) is the raw data to be deposited (and shared) in association with the publication. This is a good example to illustrate that a large facility should also archive the ‘primary’ raw data, with all associated metadata, which would allow reproducible (and improved when possible) future data (re)processing. As of today, ESRF saves all on-site-processed data at this beamline. The great advantage of HDF5 files is that metadata can be conveniently stored with the raw and processed data (*i.e.* information on how it was processed), which is critical for later (re)processing, if/when needed.

The final example is the RQPA study of NIST reference Portland clinker SRM 2686a, which is the sample used in the ASTM C1365 test method for Rietveld quantitative phase analysis validation at cement companies and laboratories (García-Maté *et al.*, 2019). In a first version of this work, three raw powder diffraction data sets from three different diffractometers – laboratory Cu $K\alpha_1$, laboratory Mo $K\alpha_1$ and synchrotron radiation – were deposited. In a second version, after addressing the referee’s comments, 11 raw powder patterns were shared, nine coming from a Cu $K\alpha_1$ diffractometer. In a third version, new powder patterns were uploaded (<https://doi.org/10.5281/zenodo.1318501> and <https://doi.org/10.5281/zenodo.1318499>). All versions of the submission are permanently accessible at Zenodo. Sharing raw data from laboratory sources is normally smoother as the level of data processing is much lower and so defining ‘raw powder diffraction data’ is easier.

4. Benefits of sharing raw powder diffraction data

Here I discuss (possible) benefits of sharing raw diffraction data for researchers/readers and for authors.

4.1. Benefits for researchers/readers

The first and obvious benefit is a better understanding of the samples. We (readers) have seen many powder pattern figures where the interesting region for the reader (which does not necessarily have to be the same as that for the authors) is not highlighted or even displayed, although it was collected. Sharing raw powder diffraction data removes this obstacle in a very time- and cost-efficient way.

A second advantage is the training capabilities in the specific scientific field and technique. Irrespective of the information to be obtained – (i) crystal structure, (ii) micro-structure, (iii) phase analysis, (iv) unit-cell evolution along a series, *etc.* – young researchers could be specifically trained with data sets in their respective techniques and scientific fields.

Databases for given techniques or even sample system types will naturally emerge. A smart re-indexing of deposited raw diffraction data could make researcher tasks easier with field-tailored raw-data databases.

Finally, new big-data research, probably using artificial intelligence and machine-learning tools, will extract new information from the deposited raw powder diffraction data by establishing correlations with other properties and performance within data sets.

4.1.1. Benefits for authors. Credibility is the major keyword here as a benefit. In a time-restrained situation, I speculate that peers will be tempted to read publications with shared raw data. On the other hand, in the future, publications not sharing raw data could be undervalued.

Metrics is another area of benefit. I am not aware of a well established correlation between downloads, readings and citations. However, I would intuitively think that more readings could lead to more citations for the publications. In any case, the doi for the deposited raw data can be added to selected applications (for instance Google Scholar) and the citations and the downloads could be evaluated. Accountability is also coming.

Analysing raw data from other research groups, in a user-friendly way, could foster collaborations.

Better research may also come about through knowledge exchange. The benefits described above are speculative, but the following scenario has already happened. Concerning the manuscript about the RQPA study of NIST reference Portland clinker SRM 2686a (García-Maté *et al.*, 2019), a reviewer downloaded the data deposited with Zenodo. He/she visualized the diffraction patterns and analysed them with his/her software. Two quite interesting comments were made. In Fig. 2 (top), I compare a Cu $K\alpha_1$ Rietveld plot similar to the one we produced for the submitted article, highlighting the fit and the contribution of the main phases. In Fig. 2 (bottom), I show the plot produced by the reviewer, by downloading our raw diffraction data and then pointing out details about a minor phase in the sample that we did not describe in our manuscript. This phase comes from the hydration and carbonation of the cement. Following his/her advice, we collected three new patterns (deposited with Zenodo in the second version of the manuscript), and we show that this phase is not present in any pattern. So, it was concluded that it was an isolated occurrence.

A second outcome of the review was a query about the RQPA of our initial deposited data with their control file, which highlighted minor differences. In order to address these differences, we reported in the second version full details of the peak shape description as the fitting of the diffraction peaks plays a minor, but not negligible, role in the outcome of a Rietveld quantitative phase analysis.




Figure 2
Laboratory $\text{Cu K}\alpha_1$ pattern for NIST reference Portland clinker SRM 2686a. (Top) Powder pattern figure after García-Maté *et al.* (2019). (Bottom) Powder data downloaded by the reviewer, visualized and then annotated, requesting additional information from us, the authors.

5. Credibility section

5.1. Competing Interests

The author declares no competing interests.

5.2. Raw data sharing

All ‘ready-to-analyse’ raw powder diffraction data described in this communication have been deposited with Zenodo at the links given in §3.

5.3. Weakest point(s) self-assessment

I do not have a sound ground to advise which ‘raw data’ should be shared. This would require a powder diffraction community discussion led by the IUCr’s Commission on Powder Diffraction. Suffice to say, I see that there are advantages and disadvantages for the different options. I advocate in detail here that ‘ready-to-analyse’ fully processed powder diffraction data should be routinely deposited and shared rather than the data before the instrument, *i.e.* detector, corrections. This would then be the same approach as the current single-crystal diffractionist approach, where the firmware corrections are fully trusted and it is after those firmware corrections that the raw diffraction images are recommended to be preserved.

5.4. Data accountability

Synchrotron powder diffraction data were recorded at the BL04-MSPD beamline, ALBA synchrotron, and the ID15A beamline, ESRF. Laboratory powder diffraction data were

recorded at the diffractometers of the SCAI central services, University of Malaga, Spain.

Acknowledgements

I am thankful for fruitful discussions about raw data sharing with John Helliwell, as well as his critical reading of this manuscript. I acknowledge information and discussion about synchrotron raw data processing with Francois Fauth (ALBA synchrotron) and Gavin Vaughan and Armando Solé (ESRF synchrotron). I finally acknowledge collaborators in ecements (Ana Cuesta, Angeles G. De la Torre, Isabel Santacruz and Marta García-Maté) for their meticulous diffraction work.

References

- Asakura, K., Abe, H. & Kimura, M. (2018). *J. Synchrotron Rad.* **25**, 967–971.
- Bernstein, H. J., Bollinger, J. C., Brown, I. D., Gražulis, S., Hester, J. R., McMahon, B., Spadaccini, N., Westbrook, J. D. & Westrip, S. P. (2016). *J. Appl. Cryst.* **49**, 277–284.
- Bruno, I., Gražulis, S., Helliwell, J. R., Kabekkodu, S. N., McMahon, B. & Westbrook, J. (2017). *Data Sci. J.* **16**, 1–17.
- Cuesta, A., Zea-Garcia, J., De la Torre, A. G., Santacruz, I. & Aranda, M. A. G. (2018). *Adv. Cem. Res.* <https://doi.org/10.1680/jadcr.18.00097>.
- Cuesta, A., Zea-Garcia, J. D., Londono-Zuluaga, D., De la Torre, A. G., Santacruz, I., Vallcorba, O. & Aranda, M. A. G. (2017). *Crystals*, **7**, 317.
- García-Maté, M., Álvarez-Pinazo, G., León-Reina, L., De la Torre, A. G. & Aranda, M. A. G. (2019). *Cem. Concr. Res.* **115**, 361–366.
- Glasser, L. (2016). *J. Chem. Educ.* **93**, 542–549.
- Gražulis, S., Chateigner, D., Downs, R. T., Yokochi, A. F. T., Quirós, M., Lutterotti, L., Manakova, E., Butkus, J., Moeck, P. & Le Bail, A. (2009). *J. Appl. Cryst.* **42**, 726–729.
- Groom, C. R., Bruno, I. J., Lightfoot, M. P. & Ward, S. C. (2016). *Acta Cryst.* **B72**, 171–179.
- Hall, S. R., Allen, F. H. & Brown, I. D. (1991). *Acta Cryst.* **A47**, 655–685.
- Hellenbrandt, M. (2014). *Crystallogr. Rev.* **10**, 17–22.
- Helliwell, J. R., McMahon, B., Guss, J. M. & Kroon-Batenburg, L. M. J. (2017). *IUCrJ*, **4**, 714–722.
- Könnecke, M., Akeroyd, F. A., Bernstein, H. J., Brewster, A. S., Campbell, S. I., Clausen, B., Cottrell, S., Hoffmann, J. U., Jemian, P. R., Männicke, D., Osborn, R., Peterson, P. F., Richter, T., Suzuki, J., Watts, B., Wintersberger, E. & Wuttke, J. (2015). *J. Appl. Cryst.* **48**, 301–305.
- Kroon-Batenburg, L. M. J., Helliwell, J. R., McMahon, B. & Terwilliger, T. C. (2017). *IUCrJ*, **4**, 87–99.
- León-Reina, L., Cuesta, A., García-Maté, M., Álvarez-Pinazo, G., Santacruz, I., Vallcorba, O., De la Torre, A. G. & Aranda, M. A. G. (2019). *International Tables for Crystallography*, Vol. H, ch. 3.10. In the press.
- León-Reina, L., García-Maté, M., Álvarez-Pinazo, G., Santacruz, I., Vallcorba, O., De la Torre, A. G. & Aranda, M. A. G. (2016). *J. Appl. Cryst.* **49**, 722–735.
- Marone, F., Studer, A., Billich, H., Sala, L. & Stämpf, M. (2017). *Adv. Struct. Chem. Imag.* **3**, 1–11.
- Nature Editorial (2009). *Nature*, **461**, 145.
- Reid, J. R., Kaduk, J. A. & Vickers, M. (2016). *Powder Diffr.* **31**, 205–210.
- Warren, J. A. (2018). *MRS Bull.* **43**, 452–457.