JOURNAL OF APPLIED CRYSTALLOGRAPHY

ISSN 1600-5767

Serial electron crystallography: merging diffraction data through rank aggregation

Stef Smeets* and Wei Wan

Department of Materials and Environmental Chemistry, Stockholm University, Svante Arrhenius väg 16 C, Stockholm, 106 91, Sweden. *Correspondence e-mail: stef.smeets@mmk.su.se

Received 25 January 2017
Accepted 19 April 2017

Edited by A. Barty, DESY, Hamburg, Germany

Keywords: serial crystallography; electron diffraction; data merging; reflection ranking.

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

Serial electron crystallography is being developed as an alternative way to collect diffraction data on beam-sensitive polycrystalline materials. Merging serial diffraction data from a large number of snapshots is difficult, and the dynamical nature of electron diffraction prevents the use of existing methods that rely on precise measurement of kinematical reflection intensities. To overcome this problem, an alternative method that uses rank aggregation to combine the rankings of relative reflection intensities from a large number of snapshots has been developed. The method does not attempt to accurately model the diffraction intensity, but instead optimizes the most likely ranking of reflections. As a consequence, the problem of scaling individual snapshots is avoided entirely, and requirements for the data quality and precision are low. The method works best when reflections can be fully measured, but the benefit over measuring partial intensities is small. Since there were no experimental data available for testing rank-based merging, the validity of the approach was assessed through a series of simulated serial electron diffraction datasets with different numbers of frames and varying degrees of errors. Several programs have been used to show that these rank-merged simulated data are good enough for *ab initio* structure determination using several direct methods programs.

1. Introduction

Many industrially and commercially relevant materials, such as catalysts, pharmaceuticals, minerals and semiconductors, are synthesized and used in polycrystalline form. The crystals that are formed are typically not large enough for routine single-crystal analysis. Nevertheless, owing to the extensive advances made in structure determination from X-ray powder diffraction (XRPD) data, the structures of many polycrystalline materials can now be determined (David & Shankland, 2008). In cases where XRPD data reach their limit, and data from a single crystal are required, some researchers have turned to electron crystallography. Three-dimensional electron diffraction data are increasingly being used for structure determination of nanocrystalline materials (Kolb *et al.*, 2011; Yun *et al.*, 2015). However, such data prove inadequate to describe a bulk material and can be difficult to collect on beam-sensitive materials.

Serial crystallography is an emerging technique that may offer new opportunities for studying polycrystalline materials. In a serial crystallography experiment, X-ray diffraction patterns (snapshots) are collected on a large number of randomly oriented crystals. Each crystal is exposed only once, and data are collected before radiation damage can occur. By combining diffraction data from these snapshots, a complete dataset can be obtained and used for structure determination

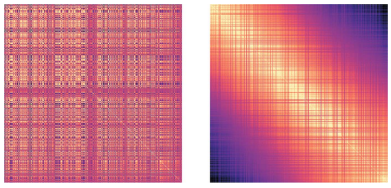

(Boutet *et al.*, 2012). Serial crystallography experiments are normally performed at centralized facilities, initially at X-ray free-electron lasers, making use of the ultrafast femtosecond pulses (Chapman *et al.*, 2011), and later also at synchrotron facilities with high-brilliance beamlines (Stellato *et al.*, 2014). To date, nearly all the activity to develop serial crystallography methods has been in the realm of structural biology (Martin-Garcia *et al.*, 2016), although recently, there have also been some studies anticipating serial crystallography in the domain of materials science (Dejoie *et al.*, 2015; Zhang *et al.*, 2015).

Our research efforts focus on the development of serial crystallography with electron radiation to characterize the structures of polycrystalline materials, particularly those that only form nano-sized crystals and are sensitive to radiation damage. Transmission electron microscopes offer an interesting alternative for serial crystallography, because (1) electrons diffract much more strongly than X-rays and high-quality electron diffraction patterns can be obtained from crystals of a few tens of nanometres in size, (2) crystals can be observed directly in imaging mode, eliminating the randomness typically associated with serial crystallography experiments, so that useful information can be extracted from nearly every frame, and (3) there is an electron microscope available in many laboratories. Modern electron microscopes are computer controlled, such that the entire data collection process can be fully automated. This enables electron diffraction data to be collected on a large number of crystals without human intervention.

A serial electron diffraction dataset contains electron diffraction patterns collected from crystals in random orientations. The reflections can be indexed following the same approach developed for serial X-ray diffraction data (Dejoie *et al.*, 2015) and the intensities extracted. However, scaling and merging of serial diffraction intensities is in general difficult. Diffraction intensities are affected by variations in diffraction volumes, crystal quality, Debye–Waller factors, flux of the incident beam and reflection partiality. Our first intuition was to look to the serial crystallography community, who have devoted great attention to merging snapshot data by modelling the diffraction processes that relate the structure factors to the intensities of the spots observed in the diffraction pattern (Kabsch, 2014; Ginn *et al.*, 2015; Uervirojnangkoorn *et al.*, 2015; Sauter, 2015; White *et al.*, 2016). Over the past five years, several unique ways of merging serial diffraction data have emerged (Kirian *et al.*, 2010; Kroon-Batenburg *et al.*, 2015; Zander *et al.*, 2016; Dilanian *et al.*, 2016). These approaches have been developed with large datasets of macromolecular structures in mind, and rely on having access to precise kinematical diffraction intensities. They appear to be less well suited for the smaller datasets we are targeting and electron diffraction intensities that are dynamical in nature. The dynamic effects arise from multiple scattering of electrons, and are strongly dependent on the size of the crystal and its orientation with respect to the incident beam. This exacerbates the problem of scaling and merging, because the structure factors are no longer directly proportional to the square root of the observed intensities.

Although desirable, it is not necessary to accurately determine the structure factors for structure determination purposes. Dynamical electron diffraction intensities obtained by using the electron diffraction tomography methods have been used to solve many complex structures (Yun *et al.*, 2015), and, despite high R values, atomic positions can be determined well within 0.1 Å accuracy through refinement. The completeness of the electron diffraction data is of greater importance than the determination of accurate structure factor amplitudes, provided that experimentally strong reflections correspond to those with high structure factor amplitudes (Klein, 2013). This principle has been applied for structure determination using XRPD data. In *EXPO*, reflections are ranked as 'strong', 'medium' or 'weak' to partition overlapping reflections, and all possible combinations are tried iteratively (Altomare *et al.*, 2003). The idea is that, once the right combination is found, the structure will be solved in a straightforward manner. Structure refinement is possible even without knowledge of the exact intensity values. Eggeman & Midgley (2012) designed a metric that uses the intensity ranks, rather than the intensity values (*e.g. via* R_1 or R_{wp}), as the objective function during the least-squares minimization. Their reasoning is that the rank is less affected by dynamic effects than are the intensities.

Here we describe an alternative approach for merging serial diffraction data that does not rely on having kinematical reflection intensities. Our method uses rank aggregation to optimize the most likely ranking of reflections. Because no comparisons between frames are made, the problem of scaling is avoided entirely, and the requirements for the data quality and precision are greatly reduced. The algorithm is applied to simulated serial electron diffraction data in order to show its effectiveness. Errors of different levels are introduced into the simulated data and the recovered diffraction intensities are used for structure determination. Important parameters that affect data quality are discussed.

2. Merging serial electron diffraction data by reflection ranking

A serial crystallography experiment can be seen as a collection of approximate rankings, given by the relative intensities of the observed reflections within a snapshot. The idea is that, although electron diffraction intensities may be measured off the Bragg condition and affected by dynamical scattering, the observed intensity for a strong reflection is likely to be higher than that of a weak reflection, such that the overall ranking of the observed reflections is roughly maintained. By viewing it this way, the merging of reflection data from snapshot frames taken from crystals with random orientations can be described as a ranking problem.

Ranking, or rank aggregation, is one of the classic problems in computer science and concerns the problem of reconstructing a ranking between items by combining different sources of information. A particularly well known example is the PageRank algorithm used by Google to rank web pages (Page *et al.*, 1999). Our method for merging snapshot data,

which we call SerialMerge, is based on the SerialRank algorithm (Fogel *et al.*, 2014), which applies a seriation algorithm for ranking a set of n items given pairwise comparisons between these items (*i.e.* if A > B and B > C then A > B > C). This sets it apart from conventional algorithms that achieve the same by deriving a score for each item. In the context of merging diffraction data, the merged intensity value could be seen as the crystallographic equivalent of the score in a ranking problem. For serial snapshot data, this means that the intensity value for a reflection has to be compared across many frames and that all factors affecting the scale must be taken into consideration. By considering only the pairwise comparisons between reflections, the intensity values need to be consistent only within one frame. In this way, any variations from frame to frame can be ignored. The SerialRank algorithm is simple, but robust for merging reflection intensities, and can recover the ranking of reflections even when some intensities are corrupted or missing, provided that the source of errors is random.

2.1. Ranking algorithm

The description of the algorithm here is adapted from the work of Fogel *et al.* (2014), and the terminology is for diffraction experiments. We first consider the simple case where only the reflections in a single diffraction snapshot are compared, and then we expand this to the general case when any number of snapshots are merged. For a single snapshot, start with a list of n reflections $(h_1, h_2, h_3, \ldots, h_n)$, with intensity $(I_1, I_2, I_3, \ldots, I_n)$. If the snapshot contains symmetry-equivalent reflections, their intensities are merged and the mean intensity is used. C is an $n \times n$ matrix of pairwise comparisons between the n individual elements, in this case reflections. Elements of C are defined as

$$C_{i,j} = \begin{cases} 1 & \text{if } I_i > I_j \\ 0 & \text{if } I_i = I_j \\ -1 & \text{if } I_i < I_j \end{cases} \quad (1)$$

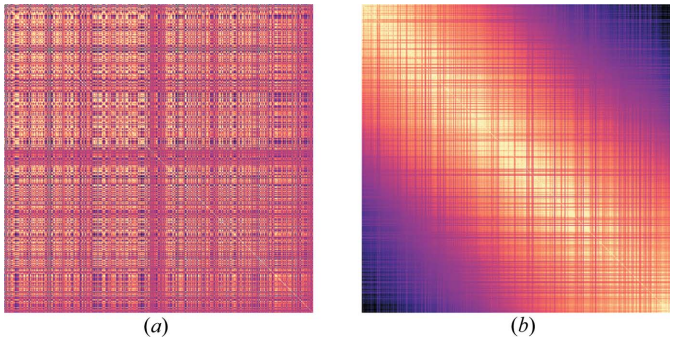

(*a*) (*b*)

Figure 1
Similarity matrix S (*a*) before and (*b*) after applying the retrieved ranking using the SerialMerge algorithm. The x and y axes correspond to the 437 observed reflections for SSZ-55 (N_{frames} = 1000, N_{upsets} = 0.0, full intensities, see §2.4). The colour brightness of each element corresponds to the similarity between two reflections.

for every value of $i \neq j$, where $0 < i, j \leq n$. C is antisymmetric, so that $C_{i,j} + C_{j,i} = 0$, and all elements on the diagonal $C_{i,i} = 1$.

As suggested by Fogel *et al.* (2014), this model can be directly extended to the setting where multiple comparisons are available for each pair, so that $C_{i,j}$ becomes a fractional value. Thus, for the general case with a large number of snapshots, we redefine n to be the total number of unique (symmetry-equivalent) reflections observed over all snapshots, giving a list of reflections $(h_1, h_2, h_3, \ldots, h_n)$. Then, let each element of C, $C_{i,j}$ be the difference between the number of occurrences of $I_i > I_j$ and $I_i < I_j$, normalized to the number of times this comparison is made, so that $-1 \leq C_{i,j} \leq 1$, *i.e.*

$$C_{i,j} = \frac{1}{m}\left(\sum_{k=1}^{m}[I_{k,i} > I_{k,j}] - \sum_{k=1}^{m}[I_{k,i} < I_{k,j}]\right), \quad (2)$$

where m corresponds to the total number of pairwise comparisons between h_i and h_j (with intensities I_i and I_j) that occur over all snapshots. The similarity matrix S is formally defined as

$$S_{i,j} = \sum_{k=1}^{n}\left(\frac{1 + C_{i,k}C_{j,k}}{2}\right), \quad (3)$$

which is essentially the dot product of C with its transpose:

$$S = \tfrac{1}{2}\left(n\mathbf{1}\mathbf{1}^{\mathrm{T}} + CC^{\mathrm{T}}\right). \quad (4)$$

We compute the Laplacian matrix of S:

$$L_S = \operatorname{diag}(S\mathbf{1}) - S. \quad (5)$$

The ranking is then given by the smallest eigenvector of L_S with a nonzero eigenvalue, the so-called Fiedler vector of S. The retrieved ranking is used to sort the reflection list. Fig. 1 shows an example of S before and after applying the retrieved ranking to simulated snapshot data of a zeolite, SSZ-55.

The SerialRank algorithm has the property that the reverse order is also a solution. We assume that the observed dataset, if merged by taking the mean of all individual intensities observed for each symmetry-equivalent reflection, gives an indication of the general trend. We calculate the Kendall rank correlation coefficient (τ) between the ranking from the SerialRank algorithm and that from taking the mean. The Kendall coefficient counts the number of agreeing pairs minus the number of disagreeing pairs normalized to the total number of pairs, so that it takes a value between −1 and 1. A value of 1 indicates identical ordering between two datasets, 0 indicates no correlation and −1 that the order is reversed. Therefore, if $\tau < 0$, we simply reverse the order of the retrieved ranking. As the Kendall coefficient is a measure of rank correlation that can be used to determine the similarity of the rankings between different datasets, we will also use it to evaluate the quality of the obtained ranking.

2.2. Assigning intensities to reflections

The intensity information is lost when applying the SerialMerge algorithm. To be able to use the data in structure determination programs, we should assign intensity values to

the reflections. We have already established that simply being able to classify reflections as strong, medium or weak is sufficient for structure determination. Here we apply a merging strategy that simply takes the mean of all individual intensities observed for each symmetry-equivalent reflection over all snapshots. The intensity values are then sorted in descending order and mapped to the ranked list of reflections output by the SerialMerge algorithm. Alternatively, a histogram of intensities from a similar crystal structure or a powder pattern could be used as a source of intensities.

2.3. Implementation

We implemented our algorithm in Python 2.7, making use of the numpy (1.11.3) (http://www.numpy.org/) and pandas (0.19.2) (http://pandas.pydata.org/) libraries. The algorithm itself spans 25 lines of code and is based on the code provided by Fogel *et al.* (2014). The Python implementation consists of two steps, (1) to construct the C matrix, which scales linearly with the number of frames provided, and (2) the ranking algorithm, which scales exponentially with the number of unique reflections (*i.e.* the size of the C matrix). This means that the algorithm is efficient with any number of frames up to around 10 000 unique reflections, and takes seconds to run on an Intel Core i7-6700K @ 4.00 GHz (Fig. S1). With larger datasets, the size of the full S matrix becomes large enough that the eigenvalue decomposition of its Laplacian becomes inefficient. Fogel *et al.* (2014) noted that with larger datasets it is still possible to use SerialRank by storing the C matrix in sparse format, skipping the calculation of S and L_S as they will not be sparse, and passing the eigenvalue solver a function that directly computes the Laplacian of S.

2.4. Performance evaluation using simulated data

We initially attempted to apply the algorithm to individual frames of a rotation electron diffraction (RED; Wan *et al.*, 2013) dataset, which was collected from a single crystal by rotating it in the electron beam around an arbitrary axis. The algorithm was unable to retrieve an accurate ranking, mainly because of the high correlation between individual frames. The RED frames share a common rotation axis and therefore the total number of unique pairwise comparisons between the reflections is small. For example, if two reflections occur in one frame, the chance is high that both also occur in the next and/or previous frames, adding redundancy instead of new information. The result is a very sparse comparison matrix C, where each observed pairwise comparison has been measured relatively frequently. The SerialMerge algorithm requires the opposite: the denser the comparison matrix, the better. This can only be achieved by collecting data from as many randomly oriented crystals as possible, to maximize the number of unique observed pairwise comparisons. Our work on the development of methods for automated data collection of electron diffraction snapshots from randomly oriented crystals is currently ongoing and we do not yet have access to experimental data. Here we perform our tests using simulated data instead. For this purpose, published structure models of six high-silica zeolites covering a range of complexities and symmetries were used (Table 1).

Table 1
Published cell parameters of the six zeolite structures (Smeets *et al.*, 2014, 2016).

Sample	Space group	a (Å)	b (Å)	c (Å)	α (°)	β (°)	γ (°)	V (Å^3)	$N_{\rm atoms}$
SSZ-45	*Fmmm*	13.719	35.224	22.136				10696.96	28
SSZ-53	$C2/c$	4.9980	33.7801	21.0997		90.77		3561.98	25
SSZ-55	$C222_1$	12.9212	21.2154	5.0912				1395.65	11
SSZ-56	$P2_1/m$	13.9387	19.9122	12.3285		106.70		3277.51	45
SSZ-58	*Pmma*	25.1484	12.5186	12.8678				4051.08	38
SSZ-59	$P\bar{1}$	5.0033	12.6831	14.7108	103.29	90.85	100.63	891.30	24

Simulated data represent the perfect scenario where the random orientations are equally distributed in reciprocal space. In practice, crystals typically tend to have a preferred orientation. This is a common problem in electron microscopy that can be dealt with experimentally, for example, by including a small rotation of the beam or sample stage, or through sample preparation by using ultramicrotomy. With these measures taken, there could still be a missing cone in the data. However, it should be noted that, as long as the data are somewhat randomly sampled around the preferred orientation, the quality of the merging is not affected. Only the completeness of the data will be lower. Rotation data collected on single crystals, particularly those with low symmetries, often have a missing wedge, and this is rarely detrimental for structure determination.

The theoretical intensities were calculated using the program *CrystDiff*, developed in our laboratory to simulate electron diffraction patterns from a crystal with arbitrary orientation. The program calculates the structure factor, as well as the kinematical intensity of each observed diffraction spot, taking into account the excitation error (deviation from the Bragg condition). This allows us to test the effect of having both fully and partially observed reflection intensities. Experimentally, reflections are rarely in exact Bragg conditions when the data are collected with a still electron beam and crystal. If it is desirable, fully observed reflections can be obtained by precessing the electron beam (Vincent & Midgley, 1994). For each crystal structure, we generated the reflection intensities for 10 000 randomly oriented snapshots up to 1 Å resolution.

Several parameters that may affect data quality were examined, including the following:

(1) The effect of scaling, for example, because of variations in the diffraction volumes and the flux of the incident beam. For each snapshot, the reflection intensities were multiplied by a random number between 1 and 50.

(2) The effect of having fully or partially observed reflection intensities.

(3) The effect of the number of frames, $N_{\rm frames}$, used for merging. Each dataset consists of 100 to 1000 frames selected randomly from the 10 000 snapshots.

(4) The effect of inaccuracies and errors in the data, such as dynamic effects, and other systematic or non-systematic factors influencing the intensities. This is achieved by

Table 2
Parameters for generating the different test datasets.†

Scaling	Random (1–50)
Intensities	Partial, full
N_{frames}	100, 200, 500, 1000
N_{upsets} (%)	0, 10, 20, 30, 40, 50, 60, 80

† N_{frames} is the number of random snapshot frames used for merging. N_{upsets} is the percentage of the reflections randomly selected and their intensities shuffled.

introducing upsets to the simulated reflection intensities, where for every set of reflection intensities, a certain percentage of the reflections were randomly sampled and their intensities shuffled around. The primary goal here is to perturb the ranking of the reflections within each set. The number of upsets is denoted by N_{upsets}, chosen from 0.0 to 0.8, where $N_{\text{upsets}} = 0.8$ means 80% of the reflection intensities are shuffled around.

The parameters chosen for generating the test datasets are summarized in Table 2. All combinations of the parameters were generated using a Python script, and the data were fed into the SerialMerge algorithm. The output of the SerialMerge algorithm is a set of reflection indices, ranked in descending order from strong to weak.

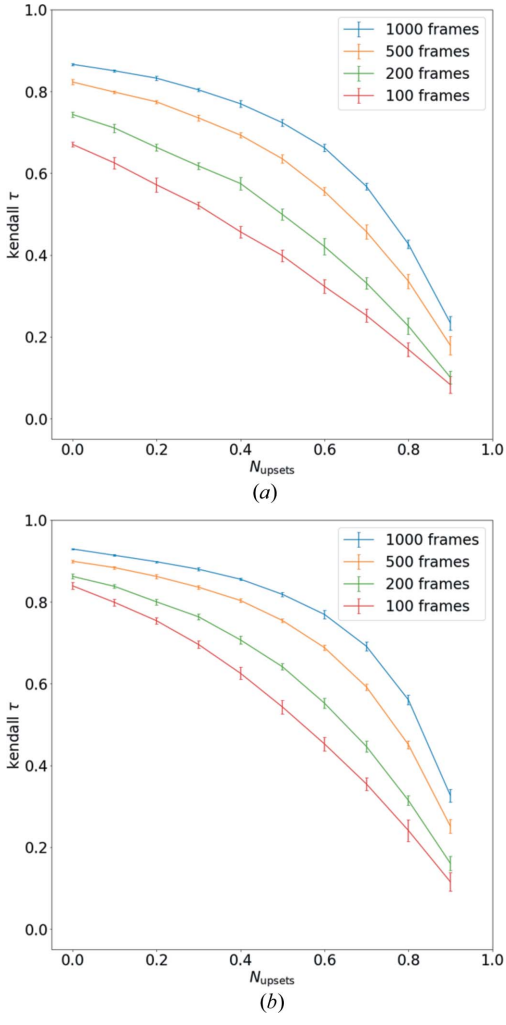

Figure 2
Kendall correlation coefficients showing the correlation between the reflection rank from merging simulated data for SSZ-45, with (*a*) partial and (*b*) full intensities, and the true rank generated from the structural model. Each value of Kendall τ corresponds to an average of 10 runs and the spread of the values is shown as error bars.

3. Results and discussion

3.1. Quality of the obtained ranking

We evaluate the quality of the rank obtained using the SerialRank algorithm by calculating its Kendall correlation coefficient (τ) with the true rank. We find that for all test cases the retrieved ranking after merging shows significant correlation with the true ranking. Both increasing the number of frames and reducing the number of upsets improve the correlation. This is highlighted in Fig. 2, which shows the trend in the Kendall coefficients between the true ranking and the retrieved ranking for SSZ-45 when varying N_{upsets} and N_{frames} used for merging. Similar trends are observed for the other tests, as shown in Tables S1 and S2, although the exact numbers vary from case to case. The highest correlations are found with $N_{\text{upsets}} = 0$ and $N_{\text{frames}} = 1000$ (the maximum number of frames included in the tests). In terms of the quality of the retrieved ranking, merging 100 frames with $N_{\text{upsets}} = 0$ is roughly equal to merging 1000 frames with $N_{\text{upsets}} = 0.6$. This is illustrated in Fig. 3(*a*), starting with 200 frames and 40% upsets. The ranking can be improved by a similar degree either by increasing the number of frames to 1000 (Fig. 3*b*) or by reducing the number of upsets to 10% (Fig. 3*c*). This indicates that by simply collecting more diffraction patterns the effect of

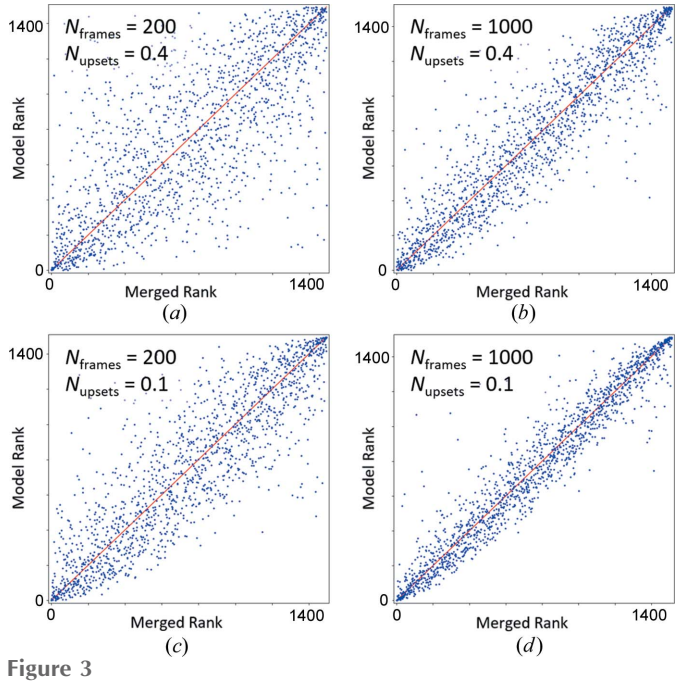

Figure 3
Reflection ranks from the simulated data for SSZ-45 using partial intensities, varying N_{frames} (200, 1000) and N_{upsets} (0.1, 0.4), plotted against the true rank generated from the structural model. The corresponding Kendall coefficients are (*a*) 0.56, (*b*) 0.77, (*c*) 0.70 and (*d*) 0.85.

Table 3
Merging statistics for SSZ-45.†

	Reflections				*C* matrix	
N_{frames}	Observed	Unique	Completeness (%)‡	Redundancy	Density (%)	Redundancy
100	6548	1434	94.8	4.6	17.4	1.2
200	13039	1493	98.6	8.7	27.7	1.4
500	32629	1510	99.8	21.6	49.5	1.9
1000	65278	1513	100.0	43.2	68.9	2.7

† For our tests, the merging statistics depend exclusively on the choice of N_{frames}. The numbers are representative of the values and trends one can expect. ‡ Up to a resolution of 1 Å.

bad data can be compensated considerably. Naturally, the best correlation is found when a large number of frames with a low degree of error are merged (Fig. 3*d*).

The number of frames used for merging essentially acts as an upper limit for the accuracy of the retrieved ranking that can be achieved. Stronger Kendall correlations are found when merging fully measured intensities, although the advantage over using partial intensities is relatively small. The effect of using full intensities is particularly noticeable when merging 100 or 200 frames, but becomes less significant when merging 500 or 1000 frames.

The merging statistics for SSZ-45 (Table 3) show that data collected on 100 randomly oriented crystals are already sufficient to create a dataset with 95% completeness. Increasing the number of frames to 1000 results in only a marginal increase of the completeness, but it increases the density of the *C* matrix from 17.4 to 68.9% and the redundancy of each observed reflection pair from 1.3 to 2.7. This substantially improves the data for merging. For five out of the six test structures, 1000 frames are enough to obtain a density of the *C* matrix of over 65%. Only for SSZ-59, with $P\bar{1}$ space-group symmetry, does the density of the *C* matrix not exceed 31.7% (Table S3). This is reflected in its Kendall coefficients (Tables S1 and S2), which are lower than those for the other test structures.

3.2. Structure determination using obtained intensities

We focus here on the data merged using partial intensities, because these better reflect the real case scenario. Fig. 4 shows the effects of the number of frames and upsets on the reflection intensities by comparing the merged data with the simulated data for SSZ-45. Despite the rudimentary estimation of reflection intensities, the bulk of the reflections are reasonably well represented. Large differences in the absolute values of the intensities occur mostly in the strong reflections. As with the ranking, the accuracy of the merged intensities can be improved to a similar degree by increasing the number of frames (Figs. 4*a* and 4*b*) and reducing the number of errors (Figs. 4*a* and 4*c*) in the data.

We were interested to see how these data would fare in practice. For our initial tests, we used the zeolite-specific program *FOCUS* (Grosse-Kunstleve *et al.*, 1997; Smeets *et al.*, 2013). The major advantage of using *FOCUS* is that it includes a built-in framework search, which enables it to classify and group solutions with the same crystal structure. This makes it ideally suited to run and evaluate the quality of a large number of datasets automatically. The success rate of finding the solution (the number of correct solutions found in a fixed number of runs) can be used as an indicator of the data quality. The specific details of the structure solution tests using *FOCUS* can be found in Table S4.

We found that *FOCUS* is able to deal with the merged data very well. In nearly all cases where $N_{upsets} \leq 0.4$, *FOCUS* was able to retrieve the correct framework structure. In general, the higher the frame number and the lower the upsets, the higher the chance that *FOCUS* finds the solution. Only a few of the datasets (SSZ-59, SSZ-56) failed to produce a structure when a low number of frames were used (100, 200). In both cases, increasing the number of frames resolved this.

FOCUS relies on *a priori* information about zeolites to supplement the diffraction data, and, while useful for testing, does not represent the general use case. Therefore, a second round of tests was performed using the generally applicable structure solution programs *SHELXS* (Sheldrick, 2008), *SHELXT* (Sheldrick, 2015), *SIR2014* (Burla *et al.*, 2015) and the charge-flipping algorithm (*olex.solve*) implemented in *OLEX2* (Dolomanov *et al.*, 2009). On the basis of the results above, we choose the datasets for all six structures with 500 frames, partial intensities and $N_{upsets} = 0.3$. Table 4 shows how the algorithms compare for these data. Default parameters were used where possible. Structure determination was considered successful only if the complete structure (>90% of the atoms located) could be recovered from the electrostatic potential map.

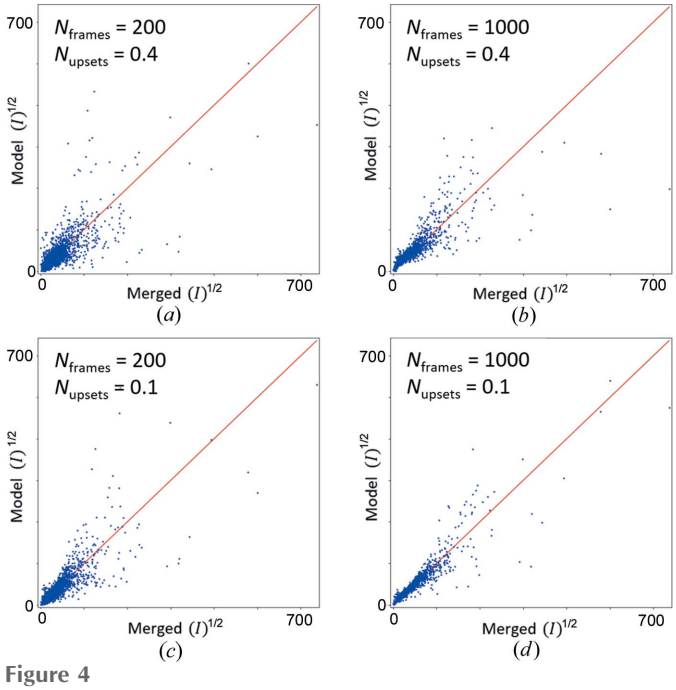

Figure 4
Merged intensities from the simulated data for SSZ-45 using partial intensities, varying N_{frames} (200, 1000) and N_{upsets} (0.1, 0.4), plotted against the theoretical values generated from the structural model.

Table 4
Structure solution results obtained using merged datasets from 500 simulated frames of randomly oriented crystals, random scale (1–50), N_{upsets} = 0.3, using partial intensities.

Sample	*FOCUS*	*olex.solve*	*SHELXS*	*SHELXT*	*SIR2014*
SSZ-55	Yes	No	No	Yes	No
SSZ-59	Yes	No	Yes	Yes†	Yes
SSZ-53	Yes	Yes	Yes	Yes	No
SSZ-45	Yes	No	Yes	Yes†	Yes
SSZ-58	Yes	Yes	Yes	Yes	No
SSZ-56	Yes	Yes	Yes	Yes†	Yes

† Only after the correct symmetry was specified *via* the '`-s`' flag.

For all tests, at least one of the programs was able to produce a fully interpretable structure solution. The charge-flipping algorithm and *SHELXS* tended to give the cleanest solutions. *SIR2014* and *SHELXT* produced somewhat messier solutions, often misassigning Si and O atoms and containing spurious *Q* peaks. We found *SHELXT* to be the most effective with our data. That said, in *SHELXT* the structure is determined in *P*1, and only in the final step is the most likely space group decided and applied. With our merged data, we found the algorithm to be too assertive in changing the space group, resulting in messy or incorrect solutions for SSZ-45, SSZ-56 and SSZ-59. In these cases, a clean solution could only be obtained by forcing the correct space group *via* the '`-s`' flag. The algorithm in *SIR2014*, on the other hand, had no issues finding the correct structure in these three cases. Of the six structures, SSZ-55 proved to be the most difficult to solve. Only *FOCUS* and *SHELXT* were able to retrieve its structure reliably. The noncentrosymmetric space group of SSZ-55 was likely to have complicated the structure determination process. By reducing N_{upsets} to 0.2 it became possible to solve the structure using *olex.solve*, and by further reducing it to 0.1 with *SHELXS*.

4. Conclusion

We have shown that ranking can be used as an effective method for merging snapshot serial electron diffraction data. The potential for using rank aggregation for merging stems from the fact that no comparisons between frames are made and data only need to be consistent within one frame. The scaling problem common to other serial data merging algorithms is thus avoided entirely. Instead of modelling the diffraction processes that relate the structure factors to the intensities of the spots observed in the diffraction pattern, our method retrieves the most likely reflection ranking. As a result, it is tolerant to errors in the diffraction intensities, which are inevitable in electron diffraction because of dynamical scattering. Diffraction intensities are mapped to the reflection ranking to construct a dataset that can be used for *ab initio* determination of complex zeolite structures. It has been shown that in general the quality of the merged intensities improves as the number of frames increases and the errors in the reflection intensities decrease. Measuring full intensities, for example by using precession electron diffraction, leads to a small advantage for datasets consisting of a small number of frames compared to measuring partial ones, but the difference is insignificant for datasets with a large number of frames. It is worth noting that, although we developed our ranking algorithm with serial electron diffraction in mind, it is generally applicable and has the potential to be applied to any problem that involves merging serial snapshot data.

The Python implementation of the algorithm, the simulated data used and the simulation software *CrystDiff* have been provided in the supporting information and are also freely available from the link http://github.com/stefsmeets/serialmerge.

Funding information

Funding for this research was provided by: Swiss National Science Foundation (award No. 165282); Swedish Research Council (VR); Swedish Governmental Agency for Innovation Systems (VINNOVA); Knut & Alice Wallenberg Foundation through the project grant 3DEM-NATUR.

References

Altomare, A., Caliandro, R., Cuocci, C., Giacovazzo, C., Moliterni, A. G. G. & Rizzi, R. (2003). *J. Appl. Cryst.* **36**, 906–913.
Boutet, S. *et al.* (2012). *Science*, **337**, 362–364.
Burla, M. C., Caliandro, R., Carrozzini, B., Cascarano, G. L., Cuocci, C., Giacovazzo, C., Mallamo, M., Mazzone, A. & Polidori, G. (2015). *J. Appl. Cryst.* **48**, 306–309.
Chapman, H. N. *et al.* (2011). *Nature*, **470**, 73–77.
David, W. I. F. & Shankland, K. (2008). *Acta Cryst.* A**64**, 52–64.
Dejoie, C., Smeets, S., Baerlocher, C., Tamura, N., Pattison, P., Abela, R. & McCusker, L. B. (2015). *IUCrJ*, **2**, 361–370.
Dilanian, R. A., Williams, S. R., Martin, A. V., Streltsov, V. A. & Quiney, H. M. (2016). *IUCrJ*, **3**, 127–138.
Dolomanov, O. V., Bourhis, L. J., Gildea, R. J., Howard, J. A. K. & Puschmann, H. (2009). *J. Appl. Cryst.* **42**, 339–341.
Eggeman, A. S. & Midgley, P. A. (2012). *Acta Cryst.* A**68**, 352–358.
Fogel, F., d'Aspremont, A. & Vojnovic, M. (2014). *arXiv*:1406.5370 [Cs, Stat].
Ginn, H. M., Messerschmidt, M., Ji, X., Zhang, H., Axford, D., Gildea, R. J., Winter, G., Brewster, A. S., Hattne, J., Wagner, A., Grimes, J. M., Evans, G., Sauter, N. K., Sutton, G. & Stuart, D. I. (2015). *Nat. Commun.* **6**, 6435.
Grosse-Kunstleve, R. W., McCusker, L. B. & Baerlocher, Ch. (1997). *J. Appl. Cryst.* **30**, 985–995.
Kabsch, W. (2014). *Acta Cryst.* D**70**, 2204–2216.
Kirian, R. A., Wang, X., Weierstall, U., Schmidt, K. E., Spence, J. C. H., Hunter, M., Fromme, P., White, T., Chapman, H. N. & Holton, J. (2010). *Opt. Express*, **18**, 5713–5723.
Klein, H. (2013). *Z. Kristallogr. Cryst. Mater.* **228**, 35–42.
Kolb, U., Mugnaioli, E. & Gorelik, T. E. (2011). *Cryst. Res. Technol.* **46**, 542–554.
Kroon-Batenburg, L. M. J., Schreurs, A. M. M., Ravelli, R. B. G. & Gros, P. (2015). *Acta Cryst.* D**71**, 1799–1811.

Martin-Garcia, J. M., Conrad, C. E., Coe, J., Roy-Chowdhury, S. & Fromme, P. (2016). *Arch. Biochem. Biophys.* **602**, 32–47.
Page, L., Brin, S., Motwani, R. & Winograd, T. (1999). *The PageRank Citation Ranking: Bringing Order to the Web*. Stanford InfoLab. Stanford University, California, USA.
Sauter, N. K. (2015). *J. Synchrotron Rad.* **22**, 239–248.
Sheldrick, G. M. (2008). *Acta Cryst.* A**64**, 112–122.
Sheldrick, G. M. (2015). *Acta Cryst.* A**71**, 3–8.
Smeets, S., McCusker, L. B., Baerlocher, C., Elomari, S., Xie, D. & Zones, S. I. (2016). *J. Am. Chem. Soc.* **138**, 7099–7106.
Smeets, S., McCusker, L. B., Baerlocher, C., Mugnaioli, E. & Kolb, U. (2013). *J. Appl. Cryst.* **46**, 1017–1023.
Smeets, S., Xie, D., McCusker, L. B., Baerlocher, C., Zones, S. I., Thompson, J. A., Lacheen, H. S. & Huang, H.-M. (2014). *Chem. Mater.* **26**, 3909–3913.
Stellato, F. *et al.* (2014). *IUCrJ*, **1**, 204–212.
Uervirojnangkoorn, M., Zeldin, O. B., Lyubimov, A. Y., Hattne, J., Brewster, A. S., Sauter, N. K., Brunger, A. T. & Weis, W. I. (2015). *eLife*, **4**, e05421.
Vincent, R. & Midgley, P. A. (1994). *Ultramicroscopy*, **53**, 271–282.
Wan, W., Sun, J., Su, J., Hovmöller, S. & Zou, X. (2013). *J. Appl. Cryst.* **46**, 1863–1873.
White, T. A., Mariani, V., Brehm, W., Yefanov, O., Barty, A., Beyerlein, K. R., Chervinskii, F., Galli, L., Gati, C., Nakane, T., Tolstikova, A., Yamashita, K., Yoon, C. H., Diederichs, K. & Chapman, H. N. (2016). *J. Appl. Cryst.* **49**, 680–689.
Yun, Y., Zou, X., Hovmöller, S. & Wan, W. (2015). *IUCrJ*, **2**, 267–282.
Zander, U., Cianci, M., Foos, N., Silva, C. S., Mazzei, L., Zubieta, C., de Maria, A. & Nanao, M. H. (2016). *Acta Cryst.* D**72**, 1026–1035.
Zhang, T., Jin, S., Gu, Y., He, Y., Li, M., Li, Y. & Fan, H. (2015). *IUCrJ*, **2**, 322–326.