

High-Resolution Cathodoluminescence Hyperspectral Imaging of Nitride Nanostructures

Paul R. Edwards,^{1,*} Lethy Krishnan Jagadamma,¹ Jochen Bruckbauer,¹ Chaowang Liu,² Philip Shields,² Duncan Allsopp,² Tao Wang,³ and Robert W. Martin¹

¹Department of Physics, SUPA, University of Strathclyde, Glasgow G4 0NG, United Kingdom

²Department of Electronic and Electrical Engineering, University of Bath, Bath BA2 7AY, United Kingdom

³EPSRC National Centre for III-V Technologies, Department of Electronic and Electrical Engineering, University of Sheffield, Sheffield S1 3JD, United Kingdom

Abstract: Hyperspectral cathodoluminescence imaging provides spectrally and spatially resolved information on luminescent materials within a single dataset. Pushing the technique toward its ultimate nanoscale spatial limit, while at the same time spectrally dispersing the collected light before detection, increases the challenge of generating low-noise images. This article describes aspects of the instrumentation, and in particular data treatment methods, which address this problem. The methods are demonstrated by applying them to the analysis of nanoscale defect features and fabricated nanostructures in III-nitride-based materials.

Key words: cathodoluminescence, hyperspectral imaging, gallium nitride, principal component analysis, multivariate statistical analysis, SEM

INTRODUCTION

The analysis of cathodoluminescence (CL), the emission of light from a material under irradiation with an energetic electron beam, is well established as a spectroscopic characterization tool in the study of light-emitting semiconductors. Its primary advantage over the analogous methods of photoluminescence and electroluminescence spectroscopies lies in its ability to confine the injection of charge carriers within a substantially sub- μm -scale volume, allowing the spatial resolution of the measurement to be well below the diffraction limit imposed by any far-field collection optics. Using a scanning electron microscope (SEM) with a field-emission source, and limiting the spread of the beam within the material by selecting a low accelerating voltage, panchromatic (i.e., spectrally unresolved) images with resolutions down to 20 nm have been demonstrated (Norman, 2000).

A major challenge in using low accelerating voltages—particularly when coupled with the low beam current required for a small spot size—is in detecting the much-reduced light intensity. This is compounded by the fact that a near-surface region of the sample is being probed in such cases, with the result that the light output is further reduced by the nonradiative loss of carriers through surface recombination. The task is made more difficult still if the light is also to be spectrally dispersed into multiple channels, as required in CL *hyperspectral imaging*.

The extension of CL to the hyperspectral imaging (or spectral mapping) mode has greatly enhanced the power of the technique by allowing the full set of spectroscopic and microscopic information to be recorded together (Christen et al., 1991). By acquiring a spectrum at each pixel in a

two-dimensional (2D) scan, a multidimensional dataset can be assembled that is both spatially and spectrally resolved. Subsets of these data can then be extracted, such as images showing wavelength shifts or variations in peak width, which would not have been possible using conventional monochromatic or panchromatic CL imaging modes.

In recent work (Liu et al., 2009; Bruckbauer et al., 2011; Edwards et al., 2011), we presented CL hyperspectral imaging results in which variations over length scales approaching 20 nm were observed. In this article, we more fully describe the steps we have taken to overcome the high noise level inherent in nanoscale CL hyperspectral images, including aspects of both instrumentation and data treatment. We illustrate these with results acquired by applying CL hyperspectral imaging to the analysis of GaN-based nanostructures.

MATERIALS AND METHODS

Samples

InGaN/GaN Multiple Quantum Wells

InGaN/GaN multiple quantum wells (MQWs) were grown on *c*-sapphire substrates by metal-organic chemical vapor deposition (MOCVD). The growth of a high-temperature AlN buffer layer and a 1- μm -thick GaN layer was followed by ten periods of $\text{In}_x\text{Ga}_{1-x}\text{N}/\text{GaN}$ layers and finally a GaN capping layer. Such structures form the basis of the active region in most GaN-based light-emitting devices.

Nanopyramids

Arrays of quantum-well-containing nanopyramids were fabricated by selective area growth (SAG) of GaN. The SAG template consisted of an SiO_2 layer on *c*-plane GaN, grown on a *c*-sapphire substrate. An hexagonal array of circular holes was patterned into the SiO_2 using nanoimprint lithog-

raphy and CHF_3 plasma etching, with a pitch of 450 nm. Subsequent growth of Si-doped GaN (including a single QW of $\text{In}_x\text{Ga}_{1-x}\text{N}/\text{GaN}$) using MOCVD resulted in the formation of hexagonal nanopillars. Such structures offer a promising route to the fabrication of high-efficiency light-emitting devices, through the use of active layers grown on lattice planes other than the strongly polar and piezoelectric (0001) basal plane (Liu et al., 2009). Examples of two such samples are used in this study, using different growth times and temperatures.

CL Detection

This work was carried out in a modified FEI Sirion 200 (FEI Company, Hillsboro, OR, USA) field emission gun SEM at room temperature. This section describes aspects of the instrument design and relates these to the specific challenges encountered in high-resolution CL hyperspectral imaging.

The optical arrangement most commonly used for CL collection is a 90° off-axis semiparaboloidal mirror placed between the sample and the microscope pole piece. While this allows emitted rays to be intercepted over almost the whole hemisphere, there are a number of drawbacks inherent to this geometry.

1. The increased working distance required and the presence of a large obscuring optical element compromise the imaging performance of the microscope.
2. The minimal clearance between sample and mirror generally prevents the sample from being tilted to different angles during the CL measurement. This restriction to the plan-view geometry is a particular problem when analyzing three-dimensional structures.
3. The optical invariant *étendue* imposes a trade-off between the optical magnification and the collection solid angle when the emitted light is being imaged to the entrance slit of a spectrometer. In practice, this results in either a loss of spectral resolution or very restricted field of view if attempts are made to focus the full 2π steradians to an f -matched image at the spectrometer (Edwards & Martin, 2011).

Light collection using an objective with its axis perpendicular to the electron beam overcomes the first two of these limitations. By placing the optics remote from the sample, an arbitrarily short working distance can be used, keeping the electron spot small even when using low accelerating voltages. The sample needs to be tilted when this geometry is used, and this tilt angle can be varied. Line-of-sight is maintained for all detectors above and to the tilt side of the sample; this includes annular backscattered electron and both through-the-lens and off-axis secondary electron detectors, and potentially also electron backscattered diffraction and energy-dispersive X-ray detectors. CL images are obtained by scanning the electron beam over the stationary sample, resulting in a field of view that will be limited by the optical magnification and the spectrograph slit width. If a reasonable field of view and spectral resolution are re-




Figure 1. Schematic layout of the high-resolution CL hyperspectral imaging system.

quired, then the system *étendue* will already be limiting the effective collection solid angle, so the small numerical aperture of the optics will not cause significant additional signal loss. While a $25\text{-}\mu\text{m}$ -wide slit is used at the entrance of the spectrometer (to match the pixel size of the detector), the imaged light spot is smaller than this; inserting the slit therefore does not decrease the signal throughput, but rather only acts to block any indirectly scattered light. Figure 1 shows a schematic layout of this CL system.

Spectrally dispersing the collected luminescence into multiple channels before detection inevitably increases signal noise when compared with a single detector (e.g., a photomultiplier tube). The use of a cooled electron-multiplying charge-coupled device allows the recording of lower noise spectra than a conventional camera array. This is a particular advantage when working at high magnifications, when the intolerance to even a few nanometers of image drift necessitates the recording of the individual spectra as quickly as possible.

Data Analysis

The optimum microscope parameters for acquiring high-resolution CL include a low electron energy (to restrict the beam/sample interaction volume), a low beam current (to keep the spot size small), and a short spectral acquisition time (to avoid any sample/image drift). With all but the brightest samples, these factors combine to limit the signal-to-noise ratio of any given spectral channel, and this needs to be considered when extracting 2D images from the data cube. For example, while the simplest way to extract maps of a peak's height and position is to plot the number of counts of the most intense pixel and the wavelength of the corresponding channel respectively, this will produce images dominated by noise. For this reason, the most effective methods of extracting subsets of the multidimensional dataset are those that take multiple channels into consideration at the same time; such methods include calculating statistical moments, least-squares peak fitting, and multivariate statistical analysis.

Statistical Moments

One set of techniques for extracting spectral peak characteristics is one that is used to quantify discrete statistical distributions generally: *moments*. If a given range of a

spectrum has n channels with corresponding energies e_i and intensities I_i , then the mean or centroid value (first moment) μ will be given by the expectation calculation:

$$\mu = \langle e \rangle = \frac{\sum_{i=1}^n I_i e_i}{\sum_{i=1}^n I_i}, \quad (1)$$

where the denominator term is the total integrated intensity across this range. Higher order (k 'th) moments are generally given as *central moments*, or moments about the mean:

$$\mu_k = \langle (e - \mu)^k \rangle = \frac{\sum_{i=1}^n I_i (e_i - \mu)^k}{\sum_{i=1}^n I_i}. \quad (2)$$

The second central moment, or variance, describes the width of the peak and provides a statistical alternative for quantifying a peak width instead of calculating, for example, the raw full-width at half-maximum (FWHM). It is also possible—although less common—to apply the same principle to extract higher order central moments. *Skewness*, the third central moment, provides a statistical measure of the asymmetry of a distribution, which for a spectrum can result, for example, from the presence of longitudinal optical phonon replicas on the long wavelength side of an emission peak. *Kurtosis*, the fourth central moment, is related to the “peakedness” of a distribution; mapping this parameter could be used to show variations in the degree of inhomogeneous broadening, the calculated kurtosis value approaching zero as a peak gets closer to a Gaussian profile. Press et al. (2002) describe in more detail the concept of moments of distributions and their calculation.

Peak Fitting

One limitation of the use of these statistical functions is that they rely on peaks being spectrally separated. Peak fitting, on the other hand, allows overlapping peaks to be deconvolved. Line profile functions can be chosen that are appropriate to the underlying physical mechanism involved (e.g., Gaussian, Lorentzian, Voigt, Pekarion, etc.), which are then fitted in turn to each spectrum in the dataset using a nonlinear least-squares (NLLS) optimization routine. The resultant arrays of fit parameters are then plotted to produce maps showing variations in the peak characteristics.

Principal Component Analysis

One further approach for extracting information from hyperspectral images is provided by the set of techniques referred to as multivariate statistical analysis. One of the simplest of these is *principal component analysis* (PCA), a method for identifying correlations in multidimensional data that is widely used across the experimental and social sciences. Its use in simplifying hyperspectral images in the field of remote sensing is well established, and in microscopy variations on this technique have been used effectively

to identify and quantify the contribution of peaks in X-ray fluorescence hyperspectral images (Kotula et al., 2003), and also combined hyperspectral CL/X-ray datasets (Edwards et al., 2007).

The fundamental principle of PCA lies in the identification of correlations in the variation of the data along each of its different dimensions, and using these to re-express the data in terms of a new set of directions in which the variations are now separated. In the context of hyperspectral images, these dimensions can be considered to be the spectral channels: in other words, each spectrum consisting of n channels can equivalently be described as a point in n -dimensional space. Initially, each of the m spatial pixels in the hyperspectral image is described by a linear combination of n monochromatic wavelengths; together these data are then a set of m points within the n -dimensional space of the values for each of the spectral channels. After PCA, the same data will be described as coordinates with respect to a new set of n orthonormal axes (principal basis vectors). These span the same space as the original spectral channels but are chosen such that each successive axis (component) accounts for the maximum possible variance remaining in the data.

Conceptually, the first step in calculating these principal components is the construction of a variance-covariance matrix. As with any variance calculations, this first requires calculation of the data mean for each dimension (spectral channel) i :

$$\mu_i = \frac{\sum_{k=1}^n x_{ik}}{m}. \quad (3)$$

The $n \times n$ elements c_{ij} of the covariance matrix \mathbf{C} are then constructed using:

$$c_{ij} = \frac{\sum_{k=1}^m (x_{ik} - \mu_i)(x_{jk} - \mu_j)}{m}, \quad (4)$$

where \mathbf{C} is symmetric ($c_{ij} = c_{ji}$) and the diagonal values c_{ii} are the variances of x_{ik} . Subsequent eigen decomposition of this covariance matrix results in a matrix of eigenvectors $\mathbf{V}[n, n]$:

$$\mathbf{V}^T \mathbf{C} \mathbf{V} = \mathbf{\Lambda}, \quad (5)$$

which can now be used as a new set of n basis vectors with which to describe the mean-adjusted data, and which are conventionally referred to as the principal component *loadings*. The elements λ_{ii} of the diagonal matrix $\mathbf{\Lambda}[n, n]$ are the associated eigenvalues, which quantify the fraction of the total dataset variance that is now accounted for by the i 'th component. Inspection of these eigenvalues allows us to determine the number of components r to retain in subsequent analysis: we retain a subset of $\mathbf{V}[n, n]$ as $\mathbf{W}[n, r]$ and discard the remaining $(n - r)$ components. Since the raw spectra in \mathbf{W} can be hard to interpret, it is often beneficial to carry out further orthogonal rotations of this

subset of eigenvectors; we employ the commonly used varimax rotation, which has the effect of minimizing overlap between the spectra.

In addition to the benefit of allowing rapid visualization of the dominant spectral features within these large datasets, PCA also provides the opportunity to compress, or noise-filter, the data. This is achieved by reconstructing an approximation \mathbf{X}' to the original data using only our reduced rank of r principal components:

$$\mathbf{X}' = \mathbf{WH} + \boldsymbol{\mu}, \quad (6)$$

where $\boldsymbol{\mu}$ is a matrix in which each column contains the mean spectrum $\boldsymbol{\mu}$. In this equation, $\mathbf{H}[r, m]$ is the array of principal component *scores*; these describe how much each of the new basis vectors contributes to a given data point and are calculated by projecting the original mean-adjusted data onto the loadings:

$$\mathbf{H} = \mathbf{W}^T(\mathbf{X} - \boldsymbol{\mu}). \quad (7)$$

In practice, since it is not necessary to evaluate all n eigenvectors, iterative algorithms can be used to significantly shorten the calculation time. In this work we have used the NIPALS (nonlinear iterative partial least-squares) approach, which calculates each eigenvector-eigenvalue pair in order but otherwise yields identical results to the full eigen decomposition of the covariance matrix.

Due to our choice of matrix in equation (4), in which the covariance in only spectral dimensions is considered and the m hyperspectral image pixels are treated as spatially uncorrelated data points, this analysis will be limited to looking only at spectral correlations in the hyperspectral images. This results in the PCA loadings and scores being spectra and images, respectively. However, Keenan (2009) has extended such techniques to include the analysis of correlations in the spatial domain, and this method has recently been applied to electron energy loss spectroscopy hyperspectral images by Guiton et al. (2011).

RESULTS AND DISCUSSION

Surface Defects in MQW Samples

Figure 2 shows a spectrum and images extracted from a CL hyperspectral image obtained from the surface of a GaN layer containing multiple $\text{In}_x\text{Ga}_{1-x}\text{N}/\text{GaN}$ quantum wells. Figure 2b shows a secondary electron image of the sample surface (acquired before the CL map), which shows the presence of inverted pyramidal holes, or “V-pits”. For the CL, a 5-keV, ~ 130 -pA electron beam was used, with 10-nm scan steps and an acquisition time of 40 ms/pixel. The images are extracted by fitting a Gaussian function to each spectrum in the scan, and plotting the resultant fit parameters as 2D maps. An example of a typical spectrum and the fitted peak are shown in Figure 2a. Despite the challenging signal-to-noise ratio, clear variations in the data are seen on sub- μm length scale. Dark spots in the emission intensity map (Fig. 2c) are seen to correlate with the presence of inverted pyramidal holes (“V-pits”) in the surface; these pits

are mainly due to threading dislocations that propagate from the lowest InGaN QWs and that act as centers for nonradiative carrier recombination. The corresponding peak energy map (Fig. 2d) shows that the emission from the pits is blue shifted with respect to the surrounding areas, which may result from QWs on the pyramidal facets being thinner or having a lower InN fraction than those grown on the planar material, as well as a contribution from the reduction in the electric field (Bruckbauer et al., 2011). The energy map also shows longer-range shifts in the emission wavelength over the field of view, suggesting a variation in the elastic strain resulting from the presence of these dislocations.

Figure 3 shows profiles extracted from the intensity and energy maps, along the same line (labeled A to B in Fig. 2c), which intersects two V-pits. While the FWHM of these features is around 50 nm, this corresponds approximately to the actual feature sizes as measured in the secondary electron image and so only provides an upper limit for the CL resolution.

Nanopyramids

Figure 4 shows a secondary electron image of an array of InGaN-QW-containing GaN nanopyramids grown on a SAG template. A CL hyperspectral image was acquired from the central $1.5 \times 1.5 \mu\text{m}$ square of this image using a 5-keV, ~ 2 -nA beam. The spectra showed three peaks: GaN near-band-edge emission at 3.4 eV, the QW peak at 2.9 eV, and a defect-related yellow band at 2.2 eV. Figure 5 shows a comparison of maps extracted from the 2.9-eV QW peak using different methods. Figure 5a shows the integrated intensity over the 2.6–3.2-eV range, comparable to a conventional monochromatic CL image. Figure 5b plots the raw energy value of the highest intensity spectral channel over this range; it can be seen that this is dominated by noise. Two alternative methods of plotting the emission energy are shown in Figure 5d, which shows the result of fitting a Gaussian peak with phonon satellites to the data, and Figure 5c, which shows the centroid energy calculated over the same spectral range. By making use of multiple spectral channels, both of these maps show a marked improvement in the signal-to-noise ratio, such that it is now possible to clearly distinguish emission energy shifts of only a few meV over length scales of < 100 nm. (The small energy offset between these two images is due to the asymmetry of the QW peak: the map from the peak fitting is showing the energy of the zero-phonon emission line, while the centroid map shows a mean that is calculated from this peak *and* all phonon replicas.) This simultaneous combination of spectral and spatial resolution is invaluable for characterizing such nanostructured materials during their development for applications in light-emitting devices (Liu et al., 2009).

The final example is a higher magnification CL hyperspectral image from another nanopyramid sample that incorporates an additional 110-nm-thick *p*-GaN capping layer. Following Monte Carlo simulations of the beam/sample interaction (Drouin et al., 2007), the energy of the beam has




Figure 2. (a) Typical spectrum from a CL hyperspectral image of a MQW sample, with a Gaussian fit; (b) secondary electron image; and fitted parameters of (c) peak height and (d) energy from the same data. The crosshair on image c shows the position from which spectrum (a) was extracted, and the line labeled A to B shows the position from which the linescans in Figure 3 were extracted.




Figure 3. Linescans extracted from the maps of fitted emission intensity and energy in Figures 2c and 2d, respectively. The labels A and B correspond to the points marked on Figure 2c.




Figure 4. Secondary electron image of an array of GaN nanopyrramids containing an InGaN/GaN quantum well.

been increased to 6 keV so that it just penetrates through to the QW region. The spectra again feature three dominant peaks from the GaN band edge, InGaN QW, and yellow band, and in this case we demonstrate the use of principal component analysis to deconvolve these contributions.

Figure 6 shows the mean spectrum and the scores of the first three principal components calculated from this data using the approach described in the Principal Component Analysis section, followed by a varimax rotation. Since we note that the mean spectrum itself can *in this case* be well described by a combination of the retained PCA scores, we have chosen to project the full non-mean-adjusted data onto these basis spectra [i.e., we disregard μ in equation (7)]. This amounts to an offset in the resultant images and thus avoids negative values and simplifies qualitative interpretation. These resultant spatial distributions of the spectral components are shown in the correspondingly labeled images in Figure 7. The three spectra are each clearly dominated by one of the expected luminescence peaks, with the QW emission originating primarily from the pyramid




Figure 5. Subsets of a single CL hyperspectral image from the central region of the nanopillar array in Figure 4, showing (a) the integrated intensity of the 2.6–3.2-eV QW emission band; (b) the raw peak energy value of this band; (c) the energy obtained by NLLS peak fitting; and (d) the centroid position calculated for this range. The nonluminescent regions of the energy maps have been masked in black for improved clarity.




Figure 6. The first three principal component loadings calculated from the CL hyperspectral image of an InGaN/GaN QW nanopillar.

facets and the GaN band edge and yellow band from the area between the nanostructures. However, closer inspection of the spectra reveals that the component with the 2.2-eV yellow band also has a feature coinciding with the 3.4-eV near-band-edge peak. This feature is positive on the low energy side of the 3.4-eV peak and negative on the high energy side; adding this contribution to the peak will




Figure 7. Images of scores corresponding to the first three principal component loadings shown in Figure 6. (a–c) Images are dominated by the 3.4-eV GaN near-band-edge, 2.8-eV InGaN QW peak, and 2.2-eV yellow band, respectively.

therefore have the effect of shifting it to a lower energy. Similarly, the emission from the InGaN/GaN QW is also associated with a feature near the main GaN peak, this time being negative on either side of the peak position and




Figure 8. (a) Map of the centroid energy of the 2.5–3.0-eV QW band, generated using the three principal components shown in Figures 6 and 7; (b) linescan along the line marked on this map, showing a small shift in emission energy; and (c) the full spectrum along the same line.

positive near the peak. This will have the effect of narrowing the peak width. From this analysis, therefore, we can note that the GaN band-edge peak is narrower in those areas where the QW peak is more intense and is red-shifted in those areas where more yellow band is emitted.

Similar (although less pronounced) effects can be observed in the 2.5–3.0-eV region of the spectrum. While the QW emission is mostly accounted for by one PCA component, there are small positive and negative contributions from the other two that describe a subtle shift in the emission. If a centroid energy is now calculated over this spectral range

using only the mean spectrum and first three principal components rather than the whole dataset, a low-noise map of the emission energy of the QW peak can be produced, as seen in Figure 8a. This was calculated using equation (1), in which the raw CL intensities I_i are now substituted with values approximated using equation (6) (and with scores \mathbf{H} now calculated from the mean-adjusted data).

As observed for the previous nanopillar structures, this shows a small blue shift in the emission as the apex is approached. This shift can be seen quantitatively in the associated linescan in Figure 8b and can be confirmed by plotting out the full spectrum extracted from original data along this same line, as in Figure 8c.

SUMMARY

The simultaneously high spectral and spatial resolution of CL hyperspectral imaging is proving invaluable in the characterization of nanoscale light-emitting semiconductor structures. However, the increased noise that comes from spectrally dispersing the luminescence, while at the same time keeping acquisition parameters that are consistent with high spatial resolution, leads to a decreased signal-to-noise ratio. Effective extraction of images showing variations in aspects of the emission such as peak energy therefore favors data analysis methods that make full use of the many spectral channels available, including statistical moments, peak fitting, and multivariate statistical analysis. By continuing to investigate and develop these methods as much as the instrumentation, we aim to further extend the capabilities of this technique as applied to the study of nanoscale semiconductor defects and devices.

ACKNOWLEDGMENTS

The authors would like to thank the U.K. Engineering and Physical Sciences Research Council (EPSRC) and the University of Strathclyde for financial support.

REFERENCES

- BRUCKBAUER, J., EDWARDS, P.R., WANG, T. & MARTIN, R.W. (2011). High resolution cathodoluminescence hyperspectral imaging of surface features in InGaN/GaN multiple quantum well structures. *Appl Phys Lett* **98**, 141908.
- CHRISTEN, J., GRUNDMANN, M. & BIMBERG, D. (1991). Scanning cathodoluminescence microscopy—A unique approach to atomic-scale characterization of heterointerfaces and imaging of semiconductor inhomogeneities. *J Vac Sci Technol B* **9**, 2358–2368.
- DROUIN, D., COUTURE, A.R., JOLY, D., TASTET, X., AIMEZ, V. & GAUVIN, R. (2007). CASINO V2.42—A fast and easy-to-use modeling tool for scanning electron microscopy and microanalysis users. *Scanning* **29**, 92–101.
- EDWARDS, P.R. & MARTIN, R.W. (2011). Cathodoluminescence nano-characterization of semiconductors. *Semicond Sci Technol* **26**, 064005.
- EDWARDS, P.R., MARTIN, R.W. & LEE, M.R. (2007). Combined cathodoluminescence hyperspectral imaging and wavelength dispersive X-ray analysis of minerals. *Am Miner* **92**, 235–242.

- EDWARDS, P.R., SLEITH, D., WARK, A.W. & MARTIN, R.W. (2011). Mapping localized surface plasmons within silver nanocubes using cathodoluminescence hyperspectral imaging. *J Phys Chem C* **115**, 14031–14035.
- GUITON, B.S., IBERI, V., LI, S., LEONARD, D.N., PARISH, C.M., KOTULA, P.G., VARELA, M., SCHATZ, G.C., PENNYCOOK, S.J. & CAMDEN, J.P. (2011). Correlated optical measurements and plasmon mapping of silver nanorods. *Nano Lett* **11**, 3482–3488.
- KEENAN, M.R. (2009). Exploiting spatial-domain simplicity in spectral image analysis. *Surf Interface Anal* **41**, 79–87.
- KOTULA, P.G., KEENAN, M.R. & MICHAEL, J.R. (2003). Automated analysis of SEM X-ray spectral images: A powerful new micro-analysis tool. *Microsc Microanal* **9**, 1–17.
- LIU, C., ŠATKA, A., JAGADAMMA, L.K., EDWARDS, P.R., ALLSOPP, D., MARTIN, R.W., SHIELDS, P., KOVAC, J., UHEREK, F. & WANG, W. (2009). Light emission from InGaN quantum wells grown on the facets of closely spaced GaN nano-pyramids formed by nano-imprinting. *Appl Phys Express* **2**, 121002.
- NORMAN, C. E. (2000). Challenging the spatial resolution limits of CL and EBIC. *Solid State Phenom* **78–79**, 19–25.
- PRESS, W.H., TEUKOLSKY, S.A., VETTERLING, W.T. & FLANNERY, B.P. (2002). Moments of a distribution: Mean, variance, skewness, and so forth. In *Numerical Recipes in C++: The Art of Scientific Computing*, 2nd ed., Chap. 14.1, pp. 615–620. Cambridge, UK: Cambridge University Press.