A method for cytometric image parameterization

Patrick M. Pilarski and Christopher J. Backhouse

Department of Electrical and Computer Engineering, University of Alberta
Edmonton, Alberta, Canada, T6G 2V4.
Email: {pilarski, chrisb}@ualberta.ca

Abstract: New advances in wide-angle cytometry have allowed researchers to obtain micro- and nano-structural information from biological cells. While the complex two-dimensional scattering patterns generated by these devices contain vital information about the structure of a cell, no computational analysis methods have been developed to rapidly extract this information. In this work we demonstrate a multi-agent computational pipeline that is able to extract features from a two-dimensional laser scattering image, cluster these features into spatially distinct regions, and extract a set of parameters relating to the structure of intensity regions within the image. This parameterization can then be used to infer medically relevant properties of the scattering object.

OCIS codes: (100.5010) Pattern recognition and feature extraction; (100.2960) Image analysis; (100.0100) Image processing; (150.0150) Machine vision; (290.3200) Inverse scattering; (100.3190) Inverse problems; (000.1430) Biology and medicine; (170.1530) Cell analysis; (170.3880) Medical and biological imaging; (170.4580) Optical diagnostics for medicine.

References and links

1. K. Singh, C. Liu, C. Capjack, W. Rozmus, and C. J. Backhouse, "Analysis of Cellular Structure by Light Scattering Measurements in a New Cytometer Design Based on a Liquid-Core Waveguide," IEE Proc.-Nanobiotechnol. **151**, 10–16 (2004).
2. K. Singh, X. Su, C. Liu, C. Capjack, W. Rozmus, and C. J. Backhouse, "A Miniaturized Wide-Angle 2D Cytometer," Cytometry A **69A**, 307–315 (2006).
3. P. L. Gourley and R. K. Naviaux, "Optical Phenotyping of Human Mitochondria in a Biocavity Laser," IEEE J. Quantum Electron. **11**, 818–826 (2005).
4. V. P. Maltsev, "Scanning flow cytometry for individual particle analysis," Rev. Sci. Instrum. **71**, 243–255 (2000).
5. P. L. Gourley, "Biocavity laser for high-speed cell and tumour biology," J. Phys. D-Appl. Phys. **36**, R228–R239 (2003).
6. P. Chacon, F. Moran, J. F. Diaz, E. Pantos, and J. M. Andreu, "Low-resolution structures of proteins in solution retrieved from X-ray scattering with a genetic algorithm," Biophys. J. **74**, 2760–2775 (1998).
7. P. Chacon, J. F. Diaz, F. Moran, and J. M. Andreu, "Reconstruction of protein form with X-ray solution scattering and a genetic algorithm," J. Mol. Biol. **299**, 1289–1302 (2000).
8. J. D. Watson and F. H. C. Crick, "Molecular Structure Of Nucleic Acids - A Structure For Deoxyribose Nucleic Acid," Nature **171**, 737–738 (1953).
9. K. A. Sem'yanov, P. A. Tarasov, J. T. Soini, A. K. Petrov, and V. P. Maltsev, "Calibration-free method to determine the size and hemoglobin concentration of individual red blood cells from light scattering," Appl. Opt. **39**, 5884–5889 (2000).
10. N. Ghosh, P. Buddhiwant, A. Uppal, K. Majumder, H. S. Patel, and P. K. Gupta, "Simultaneous determination of size and refractive index of red blood cells by light scattering measurements," Appl. Phys. Lett. **88**, 084101 (3 pages) (2006).
11. Z. Ulanowski, Z. Wang, P. H. Kaye, and I. K. Ludlow, "Application of neural networks to the inverse scattering problem for spheres," Appl. Opt. **37**, 4027–4033 (1998).

12. R. Drezek, A. Dunn, and R. Richards-Kortum, "Light scattering from cells: finite-difference time-domain simulations and goniometric measurements," Appl. Opt. **38**, 3651–3661 (1999).
13. R. Drezek, A. Dunn, and R. Richards-Kortum, "A pulsed finite-difference time-domain (FDTD) method for calculating light scattering from biological cells over broad wavelength ranges," Opt. Express **6**, 147–157 (2000), http://www.opticsinfobase.org/abstract.cfm?URI=oe-6-7-147.
14. C. Liu, C. E. Capjack, and W. Rozmus, "3-D simulation of light scattering from biological cells and cell differentiation," J. Biomed. Opt. **10**, 014007 (12 pages) (2005).
15. K. Sem'yanov and V. P. Maltsev, "Analysis of sub-micron spherical particles using scanning flow cytometry," Part. Part. Syst. Charact. **17**, 225–229 (2000).
16. N. Richard, M. Dojat, and C. Garbay, "Automated segmentation of human brain MR images using a multi-agent approach," Artif. Intell. Med. **30**, 153–175 (2004).
17. J. Liu, Y. Y. Tang, and Y. C. Cao, "An evolutionary autonomous agents approach to image feature extraction," IEEE Trans. Evol. Comput. **1**, 141–158 (1997).
18. M. Schmidt, "Automated Brain Tumor Segmentation," Ph.D. thesis, University of Alberta (2005).
19. C. E. Priebe, D. J. Marchette, and G. W. Rogers, "Segmentation of random fields via borrowed strength density estimation," IEEE Trans. Pattern Anal. Mach. Intell. **19**, 494–499 (1997).
20. E. G. P. Bovenkamp, J. Dijkstra, J. G. Bosch, and J. H. C. Reiber, "Multi-agent segmentation of IVUS images," Pattern Recogn. **37**, 647–663 (2004).
21. E. Duchesnay, J. J. Montois, and Y. Jacquelet, "Cooperative agents society organized as an irregular pyramid: A mammography segmentation application," Pattern Recogn. Lett. **24**, 2435–2445 (2003).
22. M. P. Wachowiak, R. Smolikova, Y. F. Zheng, J. M. Zurada, and A. S. Elmaghraby, "An approach to multimodal biomedical image registration utilizing particle swarm optimization," IEEE Trans. Evol. Comput. **8**, 289–301 (2004).
23. N. Pal and S. Pal, "A review on image segmentation techniques," Pattern Recognit. **26**, 1277–1294 (1993).
24. L. G. Shapiro and G. C. Stockman, *Computer Vision* (Prentice Hall, 2001).
25. J. M. Liu and Y. Y. Tang, "Adaptive image segmentation with distributed behavior-based agents," IEEE Trans. Pattern Anal. Mach. Intell. **21**, 544–551 (1999).
26. J. M. Liu, H. Jing, and Y. Y. Tang, "Multi-agent oriented constraint satisfaction," Artif. Intell. **136**, 101–144 (2002).
27. D. K. Panjwani and G. Healey, "Markov Random-Field Models For Unsupervised Segmentation Of Textured Color Images," IEEE Trans. Pattern Anal. Mach. Intell. **17**, 939–954 (1995).
28. T. Ojala and M. Pietikainen, "Unsupervised texture segmentation using feature distributions," Pattern Recogn. **32**, 477–486 (1999).
29. A. K. Jain and K. Karu, "Learning texture discrimination masks," IEEE Trans. Pattern Anal. Mach. Intell. **18**, 195–205 (1996).
30. S. Lee and M. M. Crawford, "Unsupervised multistage image classification using hierarchical clustering with a Bayesian similarity measure," IEEE Trans. Image Process. **14**, 312–320 (2005).
31. L. Itti, C. Koch, and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," IEEE Trans. Pattern Anal. Mach. Intell. **20**, 1254–1259 (1998).
32. D. Walther, L. Itti, M. Riesenhuber, T. Poggio, and C. Koch, "Attentional Selection for Object Recognition - a Gentle Way," in *Proceedings of Biologically Motivated Computer Vision, Second International Workshop* (Tubingen, Germany, 2002), pp. 472–479.
33. A. Sha'ashua and S. Ullman, "Structural Saliency: The Detection of Globally Salient Struc-tures Using a Locally Connected Network," in *Proceedings of The International Conference on Computer Vision* (Tarpon Springs, Florida, 1988), pp. 321–327.
34. M. Meister and M. Berry, "The Neural Code of the Retina," Neuron **22**, 435–450 (1999).
35. X. L. Wu, "Image-Coding By Adaptive Tree-Structured Segmentation," IEEE Trans. Inf. Theory **38**, 1755–1767 (1992).
36. M. G. H. Omran, A. Salman, and A. P. Engelbrecht, "Dynamic clustering using particle swarm optimization with application in image segmentation," Pattern Anal. Appl. **8**, 332–344 (2006).
37. C. Bourjot, V. Chevrier, and V. Thomas, "How Social Spiders Inspired an Approach to Region Detection," in *Proceedings of International Conference on Autonomous Agents and MultiAgent Systems* (Bologne, Italy, 2002), pp. 426–433.
38. Y. Wang and B. Yuan, "Fast method for face location and tracking by distributed behaviour-based agents," IEE Proc.-Vis. Image Signal Process. **149**, 173–178 (2002).
39. T. Mirzayans, N. Parimi, P. Pilarski, C. Backhouse, L. Wyard-Scott, and P. Musilek, "A swarm-based system for object recognition," Neural Netw. World **15**, 243–255 (2005).
40. A. Broggi, M. Cellario, P. Lombardi, and M. Porta, "An evolutionary approach to visual sensing for vehicle navigation," IEEE Trans. Ind. Electron. **50**, 18–29 (2003).
41. A. P. Engelbrecht, *Computational Intelligence: An Introduction* (John Wiley & Sons, 2002).
42. B. Prasad, S. Du, W. Badawy, and K. Kaler, "A real-time multiple-cell tracking platform for dielectrophoresis (DEP)-based cellular analysis," Meas. Sci. Technol. **16**, 909–924 (2005).

43. R. Ghrist and D. Lipsky, "Gramatical Self Assembly for Planar Tiles," in *Proceedings of International Conference on MEMS, NANO and Smart Systems*, W. Badawy and W. Moussa, eds. (Banff, Alberta, 2004), pp. 205–211.

44. R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification* (2nd Ed., Wiley Interscience, New York, 2001).

45. I. H. Witten and E. Frank, *Data Mining: Practical Machine Learning Tools and Techniques* (Morgan Kaufmann, 2005).

46. J. K. Udupa, V. R. LeBlanc, Z. G. Ying, C. Imielinska, H. Schmidt, L. M. Currie, B. E. Hirsch, and J. Woodburn, "A framework for evaluating image segmentation algorithms," Comput. Med. Imaging Graph. **30**(2), 75–87 (2006).

47. L. Bergman, A. Verikas, and M. Bacauskiene, "Unsupervised colour image segmentation applied to printing quality assessment," Image Vision Comput. **23**, 417–425 (2005).

48. J.R. Taylor, *An introduction to error analysis* (2nd Ed., University Science Books, Sausalito, California, 1997).

49. V. Navalpakkam and L. Itti, "Modeling the influence of task on attention," Vision Res. **45**, 205–231 (2005).

50. P.M. Pilarski, V.J. Sieben, C. Debes Marun, and C.J. Backhouse, "An artificial intelligence system for detecting abnormal chromosomes in malignant lymphocytes," in *Proceedings of Canadian Society for Immunology, Annual Conference* (Halifax, Canada, 2006), pp. 126.

1. Introduction

There is a great need for methods to extract and recognize patterns in cellular scattering images [1–5]. Scattering patterns contain vital information about the scattering source, and their interpretation facilitates diagnostic techniques ranging from the analysis of protein and DNA structure from x-ray diffraction [6–8], to the assessment of cell health based on patterns of laser light scattered by cellular components [1–4]. In perhaps the best known example, Watson and Crick used information from patterns seen in two-dimensional x-ray scatter plots to infer the double-helix nature of DNA [8]. In assessing cellular structure, Sem'yanov *et al.* and Ghosh *et al.* recognized regular patterns in one-dimensional cell scattering plots, and were able to use a parameterization of these patterns to extract microstructural cell information [9–11].

Scattering pattern analysis techniques are especially crucial in light of new medically-relevant optical analysis methods–specifically the development of the wide-angle cytometer. Wide-angle cytometry devices are rapid, cost effective systems able to capture two-dimensional scattering patterns from a single cell or particle suspended within a fluidic channel. In these devices, laser light is propagated through a cellular body, where it scatters and is collected by a digital imaging device (as described by Singh *et al.* [1,2]). A schematic diagram of a wide-angle cytometer is shown in Fig. 1.

Building on traditional cytometry schemes–which typically only capture scattered light at a few fixed angles or an angular slice–these label-free (i.e. non-fluorescent) detection devices provide extensive information about the internal structure of cells and are highly relevant to medical diagnostic practices [1, 2]. It is important to be able to rapidly ascertain small deviations in cell structure, as the structure of a cell can be an indicator for the progression of diseases (such as cancer) in patients [3, 5]. However, to infer cell structure from two-dimensional scattering plots, a method must be developed to extract and parameterize intensity patterns in
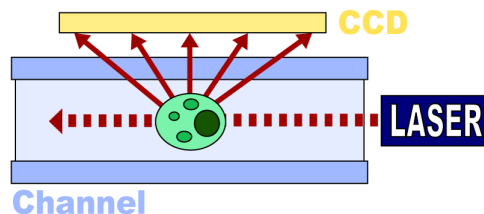

Fig. 1. Schematic diagram of a wide-angle cytometer. It includes a fluidic channel, a laser source, and a two-dimensional charge-coupled device (CCD).
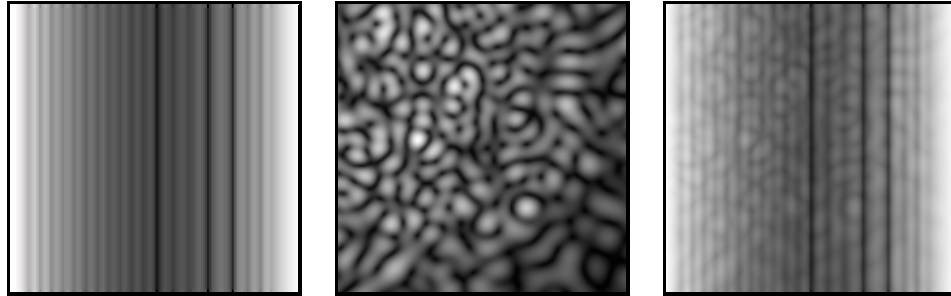
Fig. 2. Simplified example images containing features known to be present in experimental and numerically simulated scattering patterns: a series of vertical intensity bands, like those found in micro-structural scattering (left), and a number of randomly-placed high-frequency intensity regions, characteristic of nano-structural Rayleigh scattering (middle). Varying levels of high– and low– frequency intensity variation may be present in a single image, leading to complex, information-rich image structures (right). These simulated images were generated by the methods explained in the Sec. 3

cytometric scattering images. This is a previously unsolved problem, and we present here a methodology for this.

Previous work has shown that when light scatters through the cellular body it generates a complex and information-rich pattern of overlapping intensity regions. These regions are created by interfering waves propagating through a variety of cellular structures with differing size and optical properties [5]. Based on our current understanding of the scattering mechanisms present in biological cells (as indicated experimentally [1, 2, 5, 9] and through numerical simulation [12–14]), these two-dimensional scattering images are typically comprised of a set of scattering bands of varying intensity and width, with a number of additional high-frequency intensity regions (e.g. resembling those in Fig. 2). For examples of experimentally acquired scattering signatures, please see the recent work of Singh *et al.* [1, 2].

Scattering intensity contributions in cells typically come from three sources: large cell structures with diameter d greater than the incident wavelength λ (geometric scattering, $d > \lambda$, on the order of micrometers), cell structures slightly less than the wavelength of incident light (Mie scattering, $\lambda/15 < d < \lambda$), and very small organelles (Rayleigh scattering, sizes on the order of nanometers, $d < \lambda/15$) [5]. These lead to three general image cases.

In the first case (geometric scattering, and Mie scattering as d approaches λ), the scattered light will form large regular intensity bands, which–in the case of our wide-angle cytometers–appear as vertical stripes in captured wide-angle scattering images [2]. While bands may arc at low scattering angles (as shown by the images of Singh *et al.* [2]), they appear approximately linear over smaller solid angles–particularly in the side-scattering region (e.g. Fig. 2, left). These larger intensity bands are most prominent (e.g. highest intensity) in the forward and back scatter regions of a 180×180 degree scattering image, and are primarily due to the geometry of the cell wall and the larger organelles within the cell [2, 4, 5, 14].

In the second case, combining the influence of both large and medium-sized microstructural elements (e.g. both geometric scatterers and larger Mie scatters), a scattering image may contain bands that vary greatly in intensity along their length. Interference can lead to lighter or darker regions positioned within the intensity band structure.

For cellular scattering, the presence of smaller micro– and nano-scale cellular structures (like the mitochondria, which are primarily responsible for scattering at large angles [3]) will lead to a set of small randomly-distributed intensity regions. The number, frequency, and size of the

regions relates to the internal complexity of the cell. This is a result of the third case: Rayleigh scattering (and also Mie scattering where d approaches $\lambda/15$). Intensity contributions from spatially distributed organelles will constructively and destructively interfere to create a number of high-frequency intensity regions (e.g. Fig. 2, middle).

The end result is a complex scattering pattern that is comprised of interfering contributions from high-frequency intensity components and a series of vertical intensity bands (such as in Fig. 2, right), and which indicates the detailed internal morphology of the cellular body. The combination of image cases one+two, one+three, or one+two+three will all lead to images similar to the one presented in Fig. 2, right. We have observed this complex structure from both our work with wide-angle cytometers and numerical Finite Difference Time Domain (FDTD) simulations.

Computational methods have done little to take advantage of this rich image structure. One of the major factors inhibiting the development of wide-angle diagnostic devices is the computational effort needed to analyze the scattered light signatures. To deduce cellular information from scattered laser light we must somehow solve the inverse scattering problem for light through biological cells. This inverse scattering problem involves recreating the geometric parameters of a cell based on the observed path of light propagating through its cytoplasm. This is a largely unsolved problem, and any direct mathematical methods are either computationally intractable and/or lead to non-unique solutions [4]. While numerous attempts have been made to simulate the effects of scattering in cellular bodies, a method for quickly inferring the geometric structure of a cell based solely on its light scattering data still eludes researchers [4].

Given the challenge of solving the inverse problem for scattering from a living cell, the literature to date has focused on the empirical classification of cells based on their scatter at a few specific angles or an angular slice through the center of the full two-dimensional scattering pattern (commonly called the "indicatrix"). It is evident from the rich structure of the scatter patterns (along both the ϕ and θ axis) that there is far more information present than is contained in simple angular slices.

Techniques have been developed to address this problem by mathematically calculating the potential scattering patterns of cells [12–14]. In these 'forward methods', hypothetical cell geometries are used to generate simulated scattering signatures, which are then compared to experimental results. Further work has been done to use these calculated scattering patterns with evolving computer programs (such as genetic algorithms and neural networks) to interpret scattering data from crystals, proteins [6, 7], and single cells [11]. These methods largely involve the creation and verification of multiple potential structures (e.g. "generate and test" through repeated FDTD simulations [14]). These scattering simulations may take days to complete, and require the use of large parallel-processing computer arrays.

As shown by the work of Sem'yanov *et al.*, Ghosh *et al.*, and Ulanowski *et al.*, a more computationally tractable method is to effect a 'parametric solution' to the inverse scattering problem [4, 9–11, 15]. In this two-step method (parameterization and pattern recognition), they parameterize some aspect of a scattering pattern and use a set of mathematical relations [4, 9], fast Fourier transforms [10], or standard data mining techniques [11] to relate the extracted parameters to the initial structure of the scattering source. This process is rapid by comparison to forward methods, and may allow a degree of structural generalization that alleviates some of the problems caused by non-unique forward solutions.

Extracting viable parametric information from information-rich wide-angle scattering signatures presents a number of computational challenges. Because of complex cellular geometries, intensity bands may partially overlap in some places, the maximum intensity of each band may differ greatly from that of its neighbours, and the ambient background intensity is not consistent over the entire image. In addition, band boundaries are smooth gradients, not sharp intensity

level transitions. These characteristics make it quite difficult to extract relevant features from an image and group them into meaningful regions.

While researchers have addressed the individual components that make up this high-level segmentation problem (e.g. feature detection/extraction, connected component labeling, noise rejection, region clustering), to the best of our knowledge no groups have developed a way to extract and analyze the full range of information present in two-dimensional cytometric scattering images. This problem involves partitioning two-dimensional scattering images into spatially distinct regions and extracting high-level semantic information (i.e. image parameters) from the detected regions. In this work we integrate and extend upon several tested image segmentation and computer vision techniques to enhance the diagnostic capacity of wide-angle cytometry systems through the automated parameterization of scattering plots.

1.1. Recent segmentation work

Computer vision and image segmentation lie at the heart of most medical image analysis schemes [16–22]. These are widely studied areas of research that are covered extensively in the literature. For the interested reader, Shapiro and Stockman (computer vision), and Pal and Pal (image segmentation) provide excellent reviews of the relevant background literature [23, 24].

While there are many possible methods to segment wide-angle scattering images, after surveying the body of current segmentation literature we chose to design our system within the framework of a multi-agent image processing environment (described below) due to its demonstrated power, flexibility, and novelty. Multi-agent segmentation systems (such as that of Liu *et al.* [17, 25, 26]) have been thoroughly tested in a number of image processing situations, and demonstrate comparable or superior performance when compared to traditional methods. In addition, the distributed nature of multi-agent systems is a benefit for future hardware implementation. As such, they provide a solid basis for the development of a cytometric image processing pipeline.

Cytometric image parameterization is primarily a high-level segmentation problem. A number of effective algorithms developed to subdivide an image into its component parts, using everything from texture information [19, 24, 27–29] and Markov Random Fields [30] (shown to be computationally demanding [23, 30]), to standard image processing techniques [23, 24], models based on the human visual processing system [31–34], and complex image processing networks [16, 20, 21, 35].

In addition, a large body of recent image segmentation work relies on the use of multi-agent swarms, including particle swarm optimizations [22, 36], evolutionary autonomous agents [17, 25, 26, 37–39], and ant-colony optimizations [40]. These multi-agent ('swarm') systems are composed of a population of autonomous or semi-autonomous 'agents' that collaborate (directly, indirectly, and/or competitively) to achieve a common goal. In this context, an agent is defined as a independent computational unit with a set of internal states and action rules; an agent's future behaviour depends on its current condition, programmed rules, and the state of its environment [41]. (Multi-agent systems are finding widespread use in engineering and computer science applications, ranging from process optimization to computer vision, population modeling to industrial control; Engelbrecht provides a good introduction to this topic [41].)

All of these segmentation algorithms have one thing in common: they attempt to break a complex image into a set of smaller regions, where each region is homogeneous with respect to one or more parameters (e.g. contrast, intensity, texture) [23]. The effectiveness of each method varies depending on the size, texture, orientation, and shape of features contained in an image; no single approach will work well for every image [23]. In most cases, image sub-division is a two stage process–an image is segmented into smaller sections which are then clustered into groups based on some similarity metric [27, 30] (i.e. the split-and-merge or region-growing

approach, recently used for tracking cells in diagnostic images [42]).

Liu *et al.* have recently proposed several interesting agent-based approaches to region detection and segmentation. They demonstrate a segmentation system capable of rapidly labeling homogeneous tissue regions in detailed brain scan images [25], and present several methods to quickly detect edges and track shape information via a swarm of evolving computational instances (agents) [17, 26]. In their swarm intelligence approach to image segmentation, the behavior of an agent is influenced by the agent's relation to other agents and local texture information (contrast threshold, regional pixel intensity deviation, and mean pixel value) contained in the analyzed image [17, 25, 26]. Their methods typically outperform traditional image processing techniques, and are successful over a diverse range of input data. Liu *et al.*'s method has distinct advantages in that it is highly parallel (a benefit for future hardware implementations), has proved successful in complex medical imaging environments, and facilitates a distributed feature clustering procedure.

Localized action and communication are the key components of most agent-based systems. Bourjot and colleagues have recently shown that a multi-agent system, based on the web-spinning behavior of social spiders, can effectively classify regions of homogeneous color in photographic images [37], and ant colony optimizations have been used in autonomous vehicle navigation to detect roadways in low-contrast environments [40]. The work of Ghrist and Lipsky with self-assembling tile structures demonstrates an effective method for high-level organization with no centralized control [43], and Omran *et al.* further show how particle swarm optimizations can dynamically cluster information during image segmentation [36]. The distributed shape classification of Mirzayans *et al.* [39], and Wang and Yuan's agent-based face identification [38] also use local neighbourhoods to detect prominent features.

We use components of these successful swarm / image processing techniques ([23, 28, 38, 39, 42]) to complement the approach of Liu *et al.* and refine our system for use in a scattering analysis situation. We have also developed a set of unique algorithms to fully parameterize the detected image features in a way amenable to detailed scattering analysis.

Unlike most previous swarm segmentation work, our system does not involve agent reproduction or movement; the added complexity of agent dynamics, agent evolution, and agent fitness evaluation (with the additional possibility of incomplete feature detection) offsets any noticeable improvement for our particular application.

1.2. *Computational challenges*

To parameterize scattering images we need to be able to detect continuous intensity regions and characterize them with respect to their spatial orientation within the image, their intensity profile, and their relationship to other parts of the image. This allows us to numerically represent the low and high frequency intensity structures present in scattering images (as described above).

The complex image texture in cytometric scattering images makes simple feature detection problematic [27]. It is not possible to simply extract contiguous regions–corresponding to intensity bands–based solely on the raw intensity of the pixels (e.g basic threshold-based region/edge detection [23]); the high intensity point of one band may be equal in value to the background intensity at another point in the image. Feature detection methods based on local information have proved useful in solving this problem [23]: compensation techniques such as adaptive thresholding [23, 24], and the contrast-based thresholding in Liu *et al.*'s "Local Stimulus" [17] have been effective at reducing the effect of differing background levels. In these systems an image is divided up into sections and the detection threshold is set independently for each region. Due to the success of this approach (as described in recent work [17, 23–25]), our feature detection method uses adaptive thresholding (within the framework of Liu *et al.*'s "Local Stimulus" [17])

to compensate for varying background intensity.

Another challenge is "region bridging", defined as the partial overlap of two intensity regions along a small portion of their boundary. In some circumstances (e.g. low-resolution input data and/or input images that contain greatly varying band width due to complex scattering source structure) small groups of high intensity pixels will form in regions of overlap between two distinct regions. This can cause two separate intensity bands to be classified as one continuous region, greatly (and erroneously) altering the final parameter extraction.

Wang and Yuan demonstrate an effective method for separating partially blended (i.e. weakly-connected) regions based on the number of shared boundary pixels [38]; only if the number of pixels linking two regions is greater than a set threshold will two regions be merged into a single region. Wang and Yuan's technique effects a specialized form of the "erosion" and "opening" operators, commonly used to separate weakly connected image regions in binary morphology problems [24]. We use a similar bridge-eroding method in the feature detection and clustering stages of our pipeline to mitigate the effect of region bridging.

In addition to the problems of feature detection and clustering, there is the additional challenge of extracting a relevant numerical parameter set from the segmented images (i.e. extracting "region properties" [24]). Contiguous and homogeneous regions must be extracted as numerical entities for later parametric analysis. We use a form of localized communication (based on the widely-used classical connected component labeling [24]) to organize the detected image features into a set of regions. These regions are parsed to extract a set of useful image parameters. As shown by the work of Sem'yanov *et al.*, Ulanowski *et al.*, and Maltsev, once an input image has been reduced to parametric form it is possible to infer some information regarding the internal structure of the scattering source from the extracted parameter values; mathematical relations and supervised learning algorithms were previously used to determine cell size and hemoglobin concentration from the parametric profile of scattering intensity slices [4, 9, 11].

In this paper we present a computational intelligence parameterization method as the first step in a parametric solution to the inverse scattering problem for laser light through a biological cell. Our method combines and builds upon a series of successful image processing methods (image segmentation [23, 28, 42], multi-agent systems [17, 25, 38, 39, 43], and computer vision [24]) to identify and group samples of local texture information into high-level patterns (e.g. semantic information such as intensity band location and structure). While our system is designed for cytometry problems involving vertical intensity bands with randomly distributed high-frequency components, its modular analysis pipeline, numerical representation of regions, and independent parameterization routine (all described in the following sections) give it the flexibility to be easily adapted for a variety of other cytometric image analysis situations (e.g. those with different band structures/orientations and/or arbitrarily shaped intensity regions).

To the best of our knowledge, our technique is the first computational system designed to comprehensively parameterize wide-angle scattering signatures. We show that our system is able to identify the overall structure and relationships present in a scattering image. The resulting parameterization scheme, built from the numerical characterization of intensity bands and independent intensity blobs, can be used in the identification of cellular morphology. The end goal of this work is to facilitate the rapid division of experimental samples into healthy and diseased categories for expedient medical diagnosis. A pattern recognition system to infer cellular structure from image parameterization values will be presented in future work (in preparation).

2. The computational pipeline

We present a computational intelligence pipeline (called Cythe) to effectively segment, cluster, and parameterize cytometric scattering images. The problem can be described as follows:

1. Given an input scattering image $\mathbf{I}$ of size $u \times v$, where each of the $u \cdot v$ pixels represents an 8-bit grey-scale intensity value, how can we effectively segment the image into its component intensity bands and sub-band regions?

2. Furthermore, once the salient features of the image have been identified, how can we extract relevant parametric information from these features and use this information to categorize the initial input image $\mathbf{I}$?

Previous work has shown viable two-stage image segmentation systems: in stage one all salient pixels are labeled with one or more identifiers; in stage two all labeled pixels are clustered and grouped according to some similarity or congruency metric [23, 24, 27, 29, 30]. In this work we use an additional stage to organize the clustered regions and extract a set of relevant parameters.

After performing an initial user-specified image size reduction, the first stage of our pipeline (*feature detection*) is responsible for creating and fixing a population of computational agents ($\mathbf{A}$) to the salient features of the target image (as in the approach of Liu *et al.* [17, 25]). This stage effectively labels all the pixels corresponding to relevant intensity regions; an explanation of saliency determination will be presented in the following sub-section. Stage two (*feature clustering*) is responsible for clustering the fixed agent population ($\mathbf{A}$) into a list of spatially distinct regional groups ($\mathbf{G}$). The final stage of the pipeline (*post-processing*) removes large-scale image noise, creates a band-like grouping structure from identified regions, and extracts a parametric representation ($\mathbf{P}$) of the input data.

Detailed explanations of each stage are presented in the following subsections, which also describe the parameterization equations and the four major algorithms used to implement the individual stages of the pipeline. These are: the agent_fixation() routine, which is responsible for fixing the agent population to the salient image features; the propagate_id() routine, responsible for clustering agents into connected groups; the scrub() routine, which removes image noise and erroneous groupings; and the join() routine, which joins groups into a band-like structure of super-groups.

An animated example of the complete pipeline is presented in Fig. 3.

2.1. Feature detection

The first stage of the Cythe parameterization pipeline takes the input image $\mathbf{I}$, scales it to user specified dimensions $u \times v$, renders the image as a two-dimensional array [35], and creates an
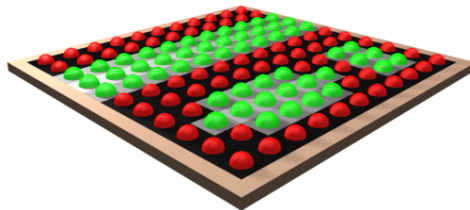

Fig. 3. (2.46 MB) Animated movie of the complete Cythe pipeline processing an example 10 pixel by 10 pixel image. Agents are represented by colored hemispheres. During the ID propagation stage of the pipeline, ID values are indicated by agent radius. The final extracted region groupings are shown by solid objects of uniform color.
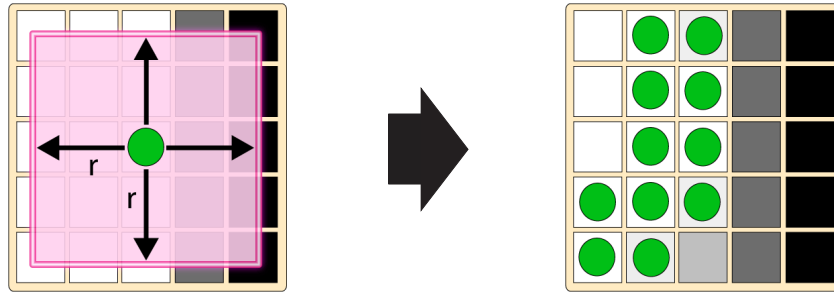
Fig. 4. Agent fixation is determined by comparing the image intensity at an agent's position to the average intensity (μ_a) within its view radius (left). After the agent_fixation() routine, members of the agent population will be fixed on areas of high intensity relative to the local image texture (right - shown here for an agent view radius of $r = 1$). This adaptive process detects edges independent of differing background levels. Pixel color indicates 8-bit intensity, from 0 (black) to 255 (white).

agent population **A** equal in size to the number of pixels in the image grid. A single agent is assigned to every pixel in the image grid. These agents then use the information available in their local neighbourhood to detect features and sort themselves into regions; this is the standard approach used in most agent-based image processing systems [17, 25, 28, 39]. To proceed we must elaborate on several definitions:

Definition (Agent): An agent is a single computational unit that is assigned to a pixel or region of the image grid **I**. Each agent has a number of internal states and potential actions, and can alter these internal states and/or perform actions based on the information present in a localized area of the image grid **I**.

Definition (Agent Neighborhood): The agent neighborhood **N** is a $n \times n$ region of the image grid **I** centered on the agent location x_a, y_a. This region determines where an agent will look for and communicate with other agents (as in Liu *et al.*'s "Neighbouring Region of an Agent" [17]).

Definition (Agent View Radius): The agent view radius **R** is a $(2r+1) \times (2r+1)$ region of the image grid **I** centered on the agent location x_a, y_a (Fig. 4, left). This area helps determine agent feature detection preferences, and the pixels within this area are used in the calculation of Average Pixel Intensity μ_a. This is akin to the image region used in the "local binary pattern and contrast measure" of Ojala and Pietikainen [28] and the area used to acquire local stimulus by Liu *et al.* [17, 25] and Mirzayans *et al.* [39].

Definition (Average Pixel Intensity): This value, denoted μ_a, is the average pixel intensity value that agent *a* observes within its view radius **R**. Average Pixel Intensity is equivalent to the "mean pixel value" component of Liu *et al.*'s texture definition, as used in their multi-agent feature detection routine [17].

During the feature detection stage of the pipeline, each agent calls on a fixation routine–agent_fixation()–to determine its immediate behavior [17, 39]. When the fixation routine is called, the agent will perform one of two actions: an agent will affix to (and therefore identify as a salient region) a pixel at image grid location **I**(x_a, y_a) if the pixel has an intensity value greater than the agent-computed average pixel intensity μ_a, or, if this condition is not satisfied, the agent *a* will be removed from the agent population **A**. In this way, agents are able to detect salient intensity edges in the image **I** independent of differing background intensity values (i.e. a fixation routine based on an agent's relation to its "average pixel intensity" is an adaptive thresholding function, as described by Pal and Pal [23]).

After the entire agent population has been polled for fixation, only agents that reside on
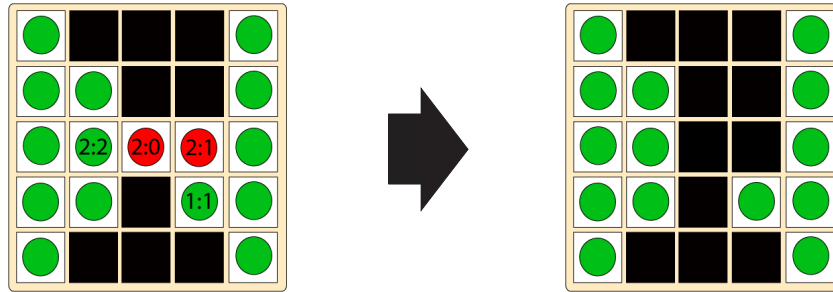
Fig. 5. An example of horizontal bridge removal (before and after removal – left and right respectively), following the agent fixation shown in Fig. 4. Green circles indicate fixed agents. Red circles represent agents that will be removed, severing the connection between the two minimally-connected bands. Numbers inside the pixels represent the ratio of horizontal to vertical neighbours within the 4-neighbourhood of a given agent (H:V).

salient pixels will remain in the population [39] (Fig. 4, right). To aid in effective region segmentation, we then scan the entire fixed agent population and remove all agents with more horizontal neighbours than vertical neighbours. Recalling the vertical nature of the intensity bands present in our scattering images (as discussed in the introductory section), we see that this helps eliminate any horizontal intensity 'bridges'; much like a horizontally-selective version of the "opening" operator used in binary image analysis [24], the removal of these weakly-connected 'bridges' facilitates region discrimination (as shown by Wang and Yuan [38]). An example of the removal process is shown in Fig. 5.

2.2. *Feature clustering*

Once the agent population has completely labeled all relevant pixels in the image grid, a clustering process–propagate_id()–takes over to form the population **A** into a set of spatially distinct regions **G** (i.e. it links all adjacent agents to create spatially-connected sub-regions). propagate_id() is a form of the classical connected components labeling algorithm [24], traditionally used to identify spatially-connected image features. Each time the propagate_id() routine is called, a sweep is done over the entire agent population; each agent in the population polls all other agents in its local neighbourhood (**N**) for their current ID value. Based on its initial scan, an agent records the highest ID value, id_{max}, in its local neighborhood. The agent then re-propagates the value id_{max} to all neighbors with ID values less than id_{max}, and the receiving agents take on the new maximum ID value. The entire agent population is iterated through until no further ID changes are observed [28]. At this point all agents in a separate physical region will share a unique ID number. A single iteration of the propagation process is shown graphically in Fig. 6.

It is important to note that ID propagation occurs in an agent's 4-neighbourhood (i.e. to agents left, right, above, and below the agent, but not on diagonal corners [24]). This aids in band discrimination and removes additional band bridges. Due to the close horizontal proximity of bands in the scaled image **I**, it was found that communication within an agent's full 8-neighbourhood could lead to a number of intensity regions being erroneously grouped into a single region. Allowing diagonal communication between agents did not facilitate any useful connections beyond those gained through purely horizontal and vertical transmission.

Since every agent starts with a unique ID value, the clustering process guarantees that every connected image region will have a common unique identifier [24]; we can now form a set of agent groups (**G**), where each group (**G**) contains a list of spatially-connected agents (**A**$_g$) that
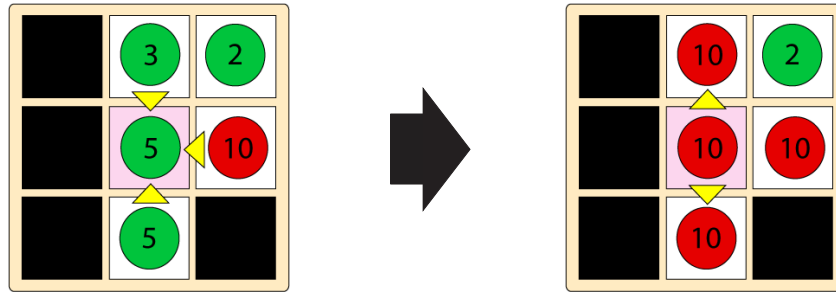
Fig. 6. Two parts of a single propagate_id() cycle for an active agent (center pixel). Initially, the agent surveys its local neighbourhood and records the ID values of its neighbours (left). Seeing there is a higher ID in the area (shown in red), it takes on this ID value and re-broadcasts the new ID to its neighbourhood (right). This leads to an agent neighborhood that is homogeneous with respect to the highest ID value.

share a common ID value ($\mathbf{A}_g$ is a non-overlapping subset of the initial population $\mathbf{A}$).

2.3. Post-processing

As in previous work, a feature detection stage followed by a clustering stage is able to effect image segmentation. However, to utilize (and parameterize) the detected regions in the context of scattering image analysis, we require a third stage to organize and parse the segmentation results.

After the creation of homogeneous ID regions, several post-processing routines take over to remove high-level noise, join vertically correlated regions into a band hierarchy (i.e. create super-groups out of related regions), and extract the final parametric representation of the input file. The first process–scrub()–searches through the list of agent groups $\mathbf{G}$ and removes all groups (and their member agents) smaller than a given percentage of the image size from $\mathbf{A}$ and $\mathbf{G}$ respectively; the removal size threshold θ can be empirically set by human users to match the input image conditions. This method of removing small connected objects was used by Prasad *et al.* to eliminate background noise in their cellular tracking system [42].

Each group that survives the scrub() routine is then analyzed for its dimensions and center point (g_x, g_y). This effects a simple geometric characterization of all surviving groups in $\mathbf{G}$.

Next, horizontally-related regions are connected into band-like structures using the join() routine, a simplified variant of the standard one-dimensional k-means clustering algorithm [44]. As in the k-means algorithm, join() creates list of super-groups and assigns one or more image regions g to each super-group g' based on the horizontal distance $d = |x_{g'} - x_g|$ between the group center and the super-group center. Assignment occurs if d is less than a user defined threshold ϕ (specified as a percentage of the image size), and each group may be assigned to only one super-group. A super-group's center is iteratively re-calculated based on the location of its member sub-groups. The join() process continues until every group has been assigned, clustering image regions with respect to their horizontal proximity. This allows the recognition of vertical bands in a scattering image while still retaining the detailed statistics of each individual sub-group. As such, join() creates a region hierarchy out of the agent population which can be stored at minimal cost for later retrieval and parameter estimation.

2.4. Parameterization

In the last step of the Cythe pipeline, the super-group hierarchy is traversed and cross-referenced with the initial image $\mathbf{I}$ to extract a number of useful global parameters $\mathbf{P}$ (shown in Tab. 1).

Table 1. The set of useful image parameters (**P**).

#	Parameter	Description
1	B	The number of bands in the image
2	BS_{min}	The minimum band spacing
3	BS_{max}	The maximum band spacing
4	BS_{avg}	The average band spacing
5	aBW_{min}	The average over all bands of minimum band width
6	aBW_{max}	The average over all bands of maximum band width
7	aBW_{avg}	The average over all bands of average band width
8	aBW_{dev}	The average over all bands of band width deviation
9	aBI_{min}	The average over all bands of minimum band intensity
10	aBI_{max}	The average over all bands of maximum band intensity
11	aBI_{avg}	The average over all bands of average band intensity
12	aBI_{dev}	The average over all bands of band intensity deviation
13	aBI_{nn}	The average over all bands of nearest-neighbour band intensity deviation

These parameters describe the overall structure and inherent complexity (in terms of spatial frequency components) of the image **I**, and are used to numerically represent the image features generated by light scattering through a biological cell (i.e. the number of regions, their size/shape, their relation to each other, and the variance of region width and intensity).

While finding a direct correlation between scattering signatures and the initial model parameters of a FDTD simulation or the structure of a cell has been shown to be an unsolved problem [4], the parameters in **P** allow us to infer structural information from the presence of intensity regions with varying spatial frequency. The knowledge that certain cellular structures will generate intensity regions of a given spatial frequency allows relationships to be made between the extracted image parameters **P** and the initial layout of the scattering source. From our initial experiments, it is expected that there will be direct correlations between these parameters and the underlying cell model parameters; we have found that this is true for relations between small organelle content and several *aBI* / *aBW* parameters (work in preparation).

In this case, each super-group g' extracted by the Cythe pipeline corresponds to a detected intensity band 'b' in the scattering image. Based on our previous work [1, 2] and a series of FDTD simulation experiments (which demonstrate the presence of vertical intensity bands in our scattering images) we found it most effective to use a band-based parameterization scheme. In this approach, the small high-frequency intensity areas resulting from smaller scattering centers are effectively described by variations to the width and intensity of existing intensity bands (i.e the presence and magnitude of high-frequency intensity fluctuation is indicated by changes to parameters 5-13, Tab. 1). A similar parameter set could be created for images without an observable band-like structure.

These parameters are extracted from the final super-group hierarchy through a series of mathematical operations, shown in Eqs. (1) and (2) below. Every detected super-group is analyzed with Eqs. (1) and (2), and the resulting values are combined into the set of parameters **P**. Width statistics are derived by iterating through the agent population, intensity statistics are derived by taking a single-pixel wide intensity sample down the vertical center line of each super-group, and band spacing statistics are generated by comparing the horizontal centers of all super-groups.

$\min_{x}^{b}(y)$ and $\max_{x}^{b}(y)$ are defined as the minimum and maximum angular values that still contain pixels belonging to band b at the vertical image position y. The function $\text{intensity}(x^{b}, y)$ is the 8-bit intensity value at the horizontal center point x of band b, at the vertical position

y. Set $\mathbf{Y}^b$ is the set of vertical values for band b. The functions min(),max(), and avg() are the standard minimum, maximum, and average operations performed on the list of values for a band. Band spacing (BS) is defined as the distance between the horizontal centers of two neighbouring bands: $|x^b - x^{b+1}|$. Values for the maximum, minimum, and average band spacing are calculated using the standard operations.

$$BW^b(y) = \max_x^b(y) - \min_x^b(y) \tag{1a}$$

$$BW^b_{min} = \min(BW^b(y),\, y \in \mathbf{Y}^b) \tag{1b}$$

$$BW^b_{max} = \max(BW^b(y),\, y \in \mathbf{Y}^b) \tag{1c}$$

$$BW^b_{avg} = \text{avg}(BW^b(y),\, y \in \mathbf{Y}^b) \tag{1d}$$

$$BW^b_{dev} = \frac{1}{|\mathbf{Y}^b|} \sum_{y \in \mathbf{Y}^b} |BW^b(y) - BW^b_{avg}| \tag{1e}$$

$$BI^b(y) = \text{intensity}(x^b, y) \tag{2a}$$

$$BI^b_{min} = \min(BI^b(y),\, y \in \mathbf{Y}^b) \tag{2b}$$

$$BI^b_{max} = \max(BI^b(y),\, y \in \mathbf{Y}^b) \tag{2c}$$

$$BI^b_{avg} = \text{avg}(BI^b(y),\, y \in \mathbf{Y}^b) \tag{2d}$$

$$BI^b_{dev} = \frac{1}{|\mathbf{Y}^b|} \sum_{y \in \mathbf{Y}^b} |BI^b(y) - BI^b_{avg}| \tag{2e}$$

$$BI^b_{nn} = \frac{1}{|\mathbf{Y}^b|} \sum_{y \in \mathbf{Y}^b} |BI^b(y) - BI^b(y-1)| \tag{2f}$$

There is a dramatic increase in the amount of information available when we compare the number of values in this extended parameter set to the number of indicatrix parameters derived from one-dimensional scattering intensity slices. We expect this increase in parametric image information will lead to a corresponding increase in the predictive power of future classification systems. Intensity band relationships (such as band spacing BS, Tab. 1) can be used to predict the nature of larger cell structures [9], while variations in region width and region intensity due to high-frequency image components (Params. 5-13, Tab. 1) may be used to detect the presence and number of micro- and nano- scale cellular organelles (work in preparation).

The final step in any automated diagnostic system is a method to deduce cellular structure from the extracted scattering pattern parameters $\mathbf{P}$. There are a number of potential machine learning approaches that can successfully associate a set of extracted parameters with an initial labeled data set to create a classifier with predictive power [44, 45]. We have developed a pattern recognition system based on the parametrization approach described in this document (in preparation).

3. Analysis methods

We employed two testing methods to verify the validity of the Cythe system: qualitative image analysis, and a quantitative statistical breakdown. For our qualitative analysis we presented the system with images representative of all three cellular scattering cases described in Sec. 1 (e.g. intensity bands with a number of randomly placed intensity regions, as in Fig. 2, right). Due to the difficulty surrounding quantitative segmentation analysis, our statistical breakdown was performed on images containing the first two scattering cases (intensity bands and bands with interference). This is explained below. In both cases our test images closely matched experimental scattering patterns [1, 2] and numerical FDTD simulations [14], both visually and in the magnitude of the output parameters.

In an ideal testing environment we would be able to use FDTD simulations and experimental data to verify the success of our segmentation system. However, to numerically analyze system accuracy it is necessary to identify the 'true' segmentation and parameterization of experimental data. As 'true' image boundaries are subjective in all but the simplest segmentation problems, most segmentation evaluation methods rely on qualitative boundary assessments for comparison values [46]; the few attempts at true quantitative evaluation typically rely on correlation data, and still involve comparisons with a manual (i.e. human) segmentation [23, 30, 46].

Thus, to quantitatively verify the validity of the Cythe extraction pipeline we used a mathematical model to create a set of viable test images. These images contained a fixed number of vertical intensity bands of varying intensity and width, irregularly placed high-frequency intensity components, intensity band overlap, differing background levels, blurring, and poorly defined band boundaries (i.e. qualities we observed in experimental scattering images). Unlike manually measured experimental scattering patterns, these model images were numerous and provided a well defined set of 'true' parameter values (derived directly from our mathematical image model) with which to statistically validate Cythe's parameter extraction.

Despite this, it was still difficult to objectively define the 'true' band width values. As bands are represented in our test images by smooth intensity gradients with no discrete edges, the 'true' band width parameter ($BW^b(y)_{real}$) was measured as the horizontal distance between band points where the pixel intensity was 80% of the band's maximum intensity, relative to a black background. This width most accurately reflected observations about real scattering band width. Because of this approximate edge value definition, the validation data for band width parameters is slightly less precise than for other parameters, as seen in the following section.

These quantitative test images contained a more regular distribution of high-frequency intensity components than was found in experimental images or our qualitative analysis images; high-frequency intensity regions were randomly placed only on intensity bands, as in image cases one and two, Sec. 1. This additional regularity was needed generate reliable true values for band parameterization– images containing a completely random distribution of high-frequency regions (as expected from Rayleigh scattering, image case three) would suffer from the same subjective evaluation problems as real experimental data.

Thus, each quantitative test image consisted of a varying number of Gaussian intensity regions superimposed on a series of vertical intensity bands. Like real scattering patterns, our test images contained bands of varying width and maximum intensity that were placed at intervals across a black background. The intensity profile of individual bands, the size and orientation of Gaussian intensity regions, and the variation of maximum band intensity across the image were picked to match the intensity profiles expected in actual scattering images. Finally, a 5×5 Gaussian blur was applied to the images to smooth out any unrealistic intensity transitions.

These test images were then presented to the Cythe system for analysis. Each test image was processed by the full computational pipeline (i.e. feature extraction, feature clustering, and post-processing) to produce a set of output parameters (P_{cythe}). Another set of parameters were derived directly from the mathematical model used to generate the test images; these parameters (P_{real}) represented the 'true parameter values' used in the creation of the test images. We then inspected how well the true parameters P_{real} matched the parameter values extracted by our pipeline P_{cythe} (i.e. how well they demonstrated a correlated linear relationship that allowed accurate prediction of the true parameter values). Both the true parameter set P_{real} and Cythe parameter set P_{cythe} included all thirteen parameters outlined in Section 2. As explained in the previous section, this band-based parameterization scheme (calculated with Eqs. (1) and (2)) can be used to represent the influence of both large scattering structures and nanostructure-derived high-frequency intensity variation.

To assess the pipeline's ability to detect changes in band width and band intensity, these tests

were performed on 162 sample images. Two different test sets were generated. The first set (T_1: 143 images) was used to determine the system's ability to detect variation in band width and band intensity– parameters primarily influenced by the presence of smaller scattering sources. In this set the number of intensity bands was held constant while the number of Gaussian intensity regions present in the image was varied between zero and fifty. The second set (T_2: 19 images) was used to test the system's ability to detect changes in band structure and spacing, which relate to scattering from larger microstructural cellular objects. In test T_2, the number of intensity bands was varied between two and twenty, while the number of Gaussian regions inserted into the image was held constant.

After each test set the Cythe parameter extractions were compared to the true parameters. System success was determined by measuring how closely the Cythe parameters matched the true parameters, as evaluated with a range of statistical tests for correlation and similarity (described in the following section).

This procedure is similar to the comparison metric of Bovenkamp *et al.*, where the plot of human v.s. machine solutions was compared to a unity slope to determine accuracy [20]. In the absence of any methods to objectively compare and evaluate segmentation schemes [23, 30, 46, 47], this approach allowed a quantitative characterization of system success.

In addition to these quantitative test images, we performed a series of qualitative tests on a range images of containing features from all three scattering image cases presented in Sec. 1 (e.g. Fig. 2, right). This was done to determine the system's ability to remove region bridges, detect regions in noisy images, and join detected regions into a structural hierarchy. For these tests, we generated a test series consisting of images with vertical intensity bands, randomly distributed high-frequency intensity regions, and vertical intensity bands overlayed with a pattern of randomly distributed high-frequency regions of irregular shape and size. Sample images and the corresponding results are presented in the following section.

For these tests the scrub() and join() thresholds (θ, ϕ) were set to 0.0018% and 0.028% respectively. The agent view radius was $\mathbf{R} = 5 \times 5$, and the agent neighbourhood was $\mathbf{N} = 3 \times 3$. Images were reduced to $\mathbf{I} = 125 \times 125$. These values were empirically derived on a small subset of the images, and subsequently used on the full set of test images without modification; one parameter setting performed well for an entire family of images.

4. Results

As a summary of the following performance assessment: the Cythe pipeline was able to detect relevant image features (e.g. intensity band pixels) in test images, remove horizontal region bridges, and cluster the detected features into a set of spatially distinct regions ($\mathbf{G}$). These regions were then used to harvest a parametric representation (P_{cythe}) of the initial image that directly matched the known parameters of the input image (P_{real}). As described earlier, the parameter set $\mathbf{P}$ serves to numerically capture the structure of both the large bands and small high-frequency intensity regions present in scattering images. This section will begin with qualitative verification results, and conclude with a quantitative numerical assessment of Cythe.

4.1. Qualitative assessment

As shown in Fig. 7, the Cythe pipeline was able to affix the agent population ($\mathbf{A}$) to the vertical intensity bands in the test image. A visual comparison of the Cythe labeling (Fig. 7, right) with the initial image (Fig. 7, left) showed that the Cythe extraction matched quite closely with our expectations from test image. We also observed that the system was able to detect the presence and magnitude of width variations in the band structure (Fig. 7, bottom row).

In addition to being able to detect linear bands, Cythe was able to detect small, arbitrarily-shaped intensity regions of varying brightness (Fig. 7, middle). As shown by the difference
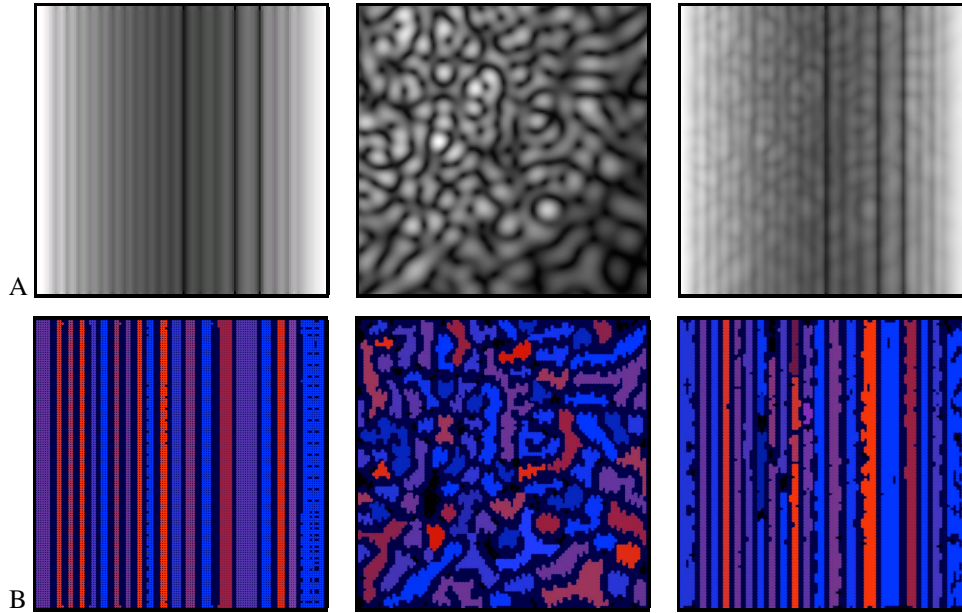
Fig. 7. A visual comparison of Cythe region detection (row B) with the initial test image (row A) for images with vertical intensity bands (left), high-frequency intensity regions (middle), and high-frequency regions overlayed onto a band structure (right, similar to those observed in FDTD simulations). Region color was assigned based on each group's unique ID value; all regions were verified to contain distinct ID values.

between Fig. 7 left and right, Cythe was also able to detect the level of high-frequency variation present in images containing high-frequency components that overlap a pattern of vertical intensity bands. This observation further supports the efficacy of the band-based parameterization scheme **P**. Random intensity regions (like those expected from Rayleigh scattering) were indicated by width and intensity deviations within the detected band structure– their intensity contributed to and noticeably altered the shape of existing bands.

We found that Cythe was able to remove parameter-degrading horizontal intensity bridges and use the clustering stage to group the agent population (**A**) into a set of distinct regions (**G**). The removal of horizontal bridging can be seen in Fig. 8, and the ability to form a population into spatially connected regions can be seen by the homogeneous region colors in Figs. 7 and 8. As shown in the difference between the two agent populations in Fig. 8 (middle and right), we found that horizontal bridges less than three pixels in width were removed during the feature detection stage . In addition, the use of a 4-neighbourhood for communication in the feature clustering stage prevented distinct bands from being classified as a single region due to any remaining weak connections (Figs. 7 and 8, right).

The join() routine constructs a set of vertical bands g' out of the detected image regions **G**. For noisy images this process would not be possible without prior use of the scrub() routine to filter out small unconnected intensity regions. Figure 9 illustrates the use of the join() and scrub() routines in the creation of a vertical band structure for simple images with and without 10% of the images pixels assigned a random 8-bit intensity noise value (i.e. random or independent noise, as expected from dust on a lense or bad CCD pixels). While portions of the agent population affixed to noise-related pixel clusters (9, bottom middle), the scrub() routine removed these small groups and the pipeline identified the same regions found in the noise-free
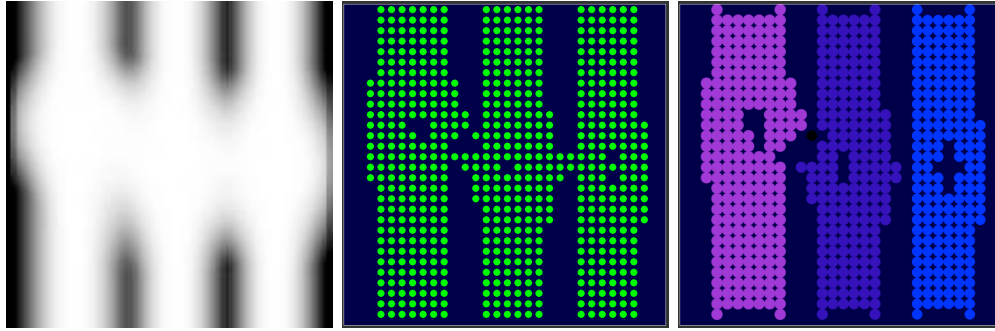
Fig. 8. A visual example of horizontal bridge removal. Left - the initial test image. Middle - the agent population directly after the agent_fixation() routine; there are three bridges at this point. Right - final region identification after post-processing; weak connections between bands did not adversely affect region identification–the two horizontal bridges were removed in the feature detection stage, and the diagonal propagation restriction prevented ID leaking over the remaining bridge (which was subsequently removed by the scrub() routine). Green dots represent fixed agents (middle), and different colors in the clustering image indicate spatially distinct regions (right).

image (9, right). In addition, the detected regions were joined into the same band structure for both the noisy and noise-free image (as shown by the number and horizontal position of the yellow vertical lines, Fig. 9, right). This lead to the same parameters being extracted for both the noisy and noise-free images. Similar performance was observed for Poisson/couting noise, though high levels of Gaussian noise required the use of a larger scrub threshold due to larger detected noise regions. A join threshold of $\phi = 0.08$ was used for the tests in Fig. 9.

In addition to accurately parameterizing our model test images, Cythe was able to extract realistic parameters for a large set of FDTD scattering simulation images containing many arbitrarily-shaped randomly-distributed high-frequency intensity regions, as derived from complex cell structures with varying physical characteristics and organelle distributions (work in preparation).

4.2. Quantitative assessment

We found that the parameters extracted by Cythe from the test images (P_{cythe}) allowed us to accurately predict the true parameters P_{real} extracted from the initial test images. This was statistically determined by calculating the correlation coefficient (r – the amount of covariance in the two populations, a good indicator of segmentation accuracy [23]), the statistical significance of the correlation ($P(r)$ – the probability that correlation value r could have arisen by pure chance for a given sample size), the chi-squared significance ($P(\chi^2)$ – the probability of both input and output variables coming from the same distribution), and the standard error (SE) for each population of input/output variables (all calculated as per [48], using Eq. 12.11 and Sec. 12.4 on Pgs. 268, 271–275 for $P(\chi^2)$, Eq. 9.15 and Sec. 9.4 on Pgs. 217–220 for r / $P(r)$, and Eq. 4.14 on Pg. 102 for SE).

This comparison is shown in tabular form for tests T_1 (band intensity parameters, Tab. 2, and band width parameters, Tab. 3) and T_2 (band number/spacing parameters; Tab. 4). These tables present the statistics for an image reduction size of $\mathbf{I} = 125 \times 125$. From statistical theory [48], r values greater than 0.216 (test T_1, 143 samples) and 0.561 (test T_2, 20 samples) indicate a statistical correlation (i.e. a probability $P(r) < 0.01$ that the correlation score could have originated by pure chance). These threshold r values are based on the sample sizes used in
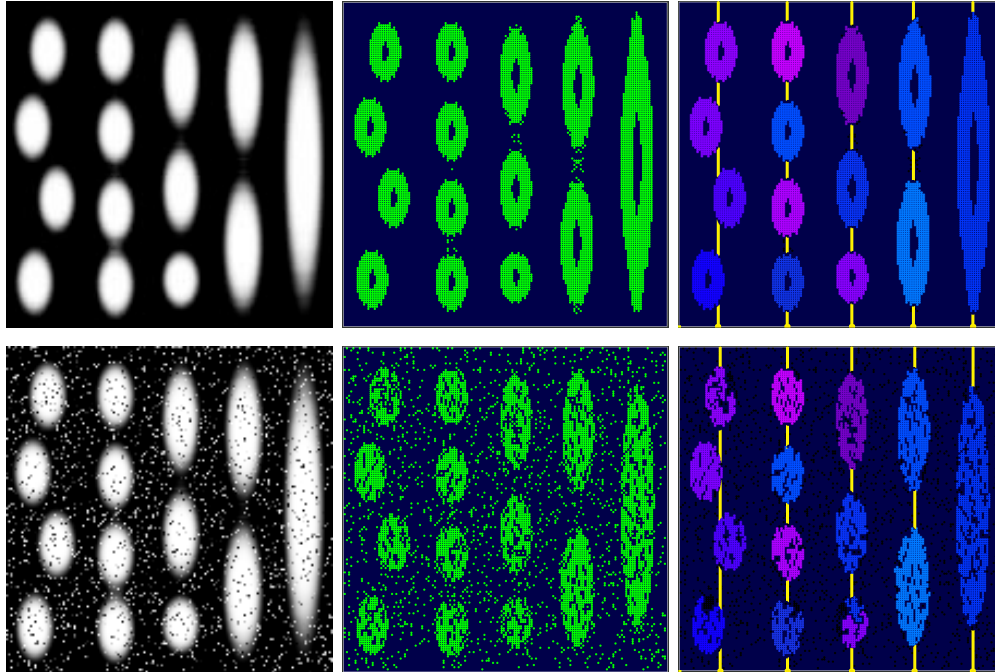
Fig. 9. Extraction of a band hierarchy for a simple noise-free image (top row) and for the same image with 10% of the images pixels assigned a random 8-bit intensity value (i.e. random noise; bottom row): the initial image (left), the agent population after the agent_fixation() routine (middle), and the final regions after post-processing (right). Yellow lines indicate band position ($x_{g'}$), and coloured regions in the post-processing image indicate spatially distinct regions g.

our experiments. As shown in Tabs. 2–4, our derived values are consistently greater than these minimum values for statistical correlation. Similarly, chi-squared significance values approaching $P(\chi^2) = 1.00$ indicate no difference in the distribution of input and output values.

The uncertainty in each parameter was estimated by adding Poisson/counting noise (i.e. each pixel was varied according to a normal distribution equal to the square root of the pixel value) and processing the resulting image by the same method as the test data sets. This was done for 56 images, allowing the extraction of a standard deviation that then allowed the calculation of chi-squared significance values.

As described in Sec. 2, the parameters shown in Tabs. 2 and 3 are used to characterize the intensity contributions from smaller Mie and Rayleigh scattering sources, while the parameters in Tab. 4 characterize the intensity contributions from larger Mie and geometric scattering objects.

With regard to the system's ability to correctly identify deviations in band intensity (test T_1), we found that Cythe was able to identify the magnitude and variance of intensity to a high degree of certainty. At an image size of $\mathbf{I} = 125 \times 125$ Cythe was able to correctly identify the number of bands in every test image. The input and output intensity parameters ($aBI_{min/max/avg/dev/nn}$) showed strong correlation, as indicated by the r, $P(r)$, and $P(\chi^2)$ values (Tab. 2). Standard error for these intensity parameters was less than half an intensity step on a 8-bit intensity scale.

The close relationship between input and output parameters was also evident for band width

Table 2. Statistical analysis for band intensity parameters.

Parameter	Description	r	$P(r)$	$P(\chi^2)$
aBI_{avg}	Avg. Band Intensity Average	0.992	<0.0001	1.000
aBI_{min}	Avg. Band Intensity Minimum	1.000	<0.0001	1.000
aBI_{max}	Avg. Band Intensity Maximum	0.998	<0.0001	1.000
aBI_{dev}	Avg. Band Intensity Deviation	1.000	<0.0001	1.000
aBI_{nn}	Avg. Band Intensity Deviation (NN[1])	1.000	<0.0001	1.000

[1] nearest neighbour.

Table 3. Statistical analysis for band width parameters.

Parameter	Description	r	$P(r)$	$P(\chi^2)$
aBW_{avg}	Avg. Band Width Average	0.386	<0.0001	1.000
aBW_{min}	Avg. Band Width Minimum	0.872	<0.0001	0.528
aBW_{max}	Avg. Band Width Maximum	0.724	<0.0001	1.000
aBW_{dev}	Avg. Band Width Deviation	0.907	<0.0001	0.286

Table 4. Statistical analysis for band number/spacing parameters.

Parameter	Description	r	$P(r)$	$P(\chi^2)$
B	Number of bands	1.000	<0.0001	1.000
BS_{min}	Minimum band spacing	0.995	<0.0001	1.000
BS_{max}	Maximum band spacing	0.993	<0.0001	1.000
BS_{avg}	Average band spacing	1.000	<0.0001	1.000

and band width deviation parameters ($aBW_{min/max/avg/dev}$), as shown by the values in Tab. 3. As explained in the previous section, difficulty defining the 'true' width values in the test images lead to greater variability in the evaluation statistics r and χ^2. While width statistics (Tab. 3) did show lower correlation between input and output values than the other parameters (Tab. 2, Tab. 4), all other values represented an excellent fit. Width values were still well above the thresholds for chance correlation, as indicated by $r > 0.261$, $P(r) \ll 0.01$. Despite having a high degree of correlation, the parameters aBW_{dev} and aBW_{min} exhibited a low $P(\chi^2)$, and further investigation showed that this deviation in input/output distribution similarity was due to a shallow (i.e < 0.5) regression slope between the input and output parameter sets. Considering the lack of 'true-value' precision when quantitatively analyzing spatial parameters in this situation, the set of width statistics in Tab. 3 sufficiently demonstrated a distinct relation between the actual layout of the test images and the Cythe parameter extraction.

In addition to band width and intensity parameters, we observed that the system was able to accurately determine the number and spacing of bands (test T_2). As shown in Tab. 4, the correlation coefficient (r) for each band-structure parameter approached 1.0 (i.e perfect correlation). This indicates a one-to-one correspondence between the input parameters P_{real} and the output parameters P_{cythe}. For the parameters $BS_{min}, BS_{max}, BS_{avg}$ there was a standard error of less than 1.1% of the image width for both reduction levels. There were no band number (B) identification errors in test T_2, and the chi-squared significance test for all parameters in Tab. 4 indicated no statistical difference between the input and output parameters.

With respect to the magnitude of observed values from Eq. 1, we found a typical range of 0.72–9.6 pixels for $aBW_{min/max/avg}$, and 0.0–0.99 pixels for aBW_{dev}. For Eq. 2, intensity parameter values were typically between 127.2–157.2 for $aBI_{min/max/avg}$, 0.59–4.29 for aBI_{nn},

and 0.26–72.7 for aBI_{dev}. Band spacing parameters varied greatly depending on the number of detected intensity bands in a sample image; for our tests we found spacing parameters between 6.46–14.8 (Test T1) and 6.48–72.2 pixels (Test T2). The standard deviation of parameter values observed under conditions with counting noise and random noise was much less than the total parameter range observed during these tests.

As the size of images presented to the system increased (with the Agent View Radius being held constant), we found that Cythe began to identify small erroneous bands within the larger regions of the image. At an image size of $\mathbf{I} = 150 \times 150$ the pipeline incorrectly identified one extra band in 67 of the 143 T_1 tests images. This lead to a noticeable decline in correlation values, and incurred a corresponding increase in standard error. It is apparent that the relationship between image size and Agent View Radius plays a role in feature detection; this will be discussed in the following section.

5. Discussion

5.1. Remarks on feature detection

The success of the Cythe feature detection system is in a large part due to the use of regional texture information to affect agent fixation. We chose to use an adaptive local thresholding method based to its success within other texture-based segmentation problems and its compatibility with agent-based image processing (as shown by a large body of previous work [17, 19, 25, 28, 39]).

Edge detection is by its very nature a local undertaking [23] and thus lends itself well to an agent-based framework. By determining fixation based on an adaptive local threshold (μ_a, the average intensity value within an agent's view radius $\mathbf{R}$), Cythe was able to effectively label all edges irrespective of the differing background intensity levels found in scattering images. By setting the adaptive threshold level greater than the local average (as in the agent_fixation() routine), the system consistently labeled the high-intensity side of all edges, isolating the band regions from the lower-intensity background.

Using an Agent View Radius of $\mathbf{R} = 5 \times 5$, we found $\mathbf{I} = 125 \times 125$ to be the most appropriate size for image reduction. At this image size and view radius, the fixation routine was able to accurately divide the image into spatially distinct regions regardless of differing background levels and gradient slopes. The distance between identified band edges was small enough that the two edge-labeling agent populations for a given band connected along the center of their band. This allowed bands to be detected as continuous units in both our model test images and complex FDTD scattering simulations.

We found that it was important to select a view radius close to the size of target image features in the reduced image $\mathbf{I}$. The two edge regions of a single band may not connect if the Agent View Radius is significantly smaller than the band size. This lead to the identification of extra bands by the clustering stage. By varying the size of the view radius (i.e. the adaptive thresholding region [23]) to match the image reduction level, feature extraction remains accurate at any image size (though increased image size comes with an increased computational cost, as described below). This follows from recent work in image saliency detection and model matching [31, 49].

5.2. Remarks on clustering

The propagate_id() routine was a reliable and effective way to cluster the labeled pixels into contiguous regions. This routine, which was based on the connected components labeling algorithm commonly used in region identification problems [24], provided a simple way to identify groups of connected agents. Much like the self-organizing tile behavior shown by Ghrist and Lipsky [43], our system was able to effectively perform long-range organization through simple

local interactions. In addition, the distributed approach lends itself well to parallelization–one of the major advantages of multi-agent systems [41].

The removal of band bridging (as described in Section 2) was essential in the accurate clustering of spatially distinct regions. The agent fixation stage eliminated direct horizontal communication over bridges by eroding bridging agents (shown above in Fig. 8), while the use of an agents 4-neighbourhood for communication prevented ID propagation over any remaining (weak) junctions between band protrusions. Without the removal of band bridging it was impossible to successfully parameterize complex images. A similar restriction on the union of weakly-connected regions has proved effective in other segmentation and image identification situations [24, 38]. The assumption that there will be no strong horizontal links between bands follows from the structure of experimental scattering images and our understanding of cellular scattering mechanisms.

Using an agent neighbourhood of size $N = 3 \times 3$ further prevented erroneous ID propagation between distinct bands. By only allowing communication between adjacent agents, ID information was not able to travel over gaps between neighbouring intensity bands.

Due to the nature of the local interactions and the multiple sweeps through the agent population, the ID propagation routine was found to be the largest computational component of the parameterization pipeline. As the proper choice of agent view radius and image reduction size decreased the final size of experimental images to approximately $I = 125 \times 125$, scalability was not an issue for our application. In the case of larger images, the use of a union-find structure (described in [24]) in the connected components algorithm would greatly reduce the number of iterations through the agent population (though it would require a higher level of centralized control).

5.3. Remarks on post-processing and parameterization

The join() routine was found to be an effective way to model the structure inherent in scattering images. It has been shown that changes in the relationships between full bands are indicative of large changes in cellular structure [9]. By linking several vertically aligned regions to a single band structure, we were able to analyze the relationship between whole band units while still retaining specific information regarding the variation present in each band and its associated sub-regions.

As shown in the results section, the join() routine managed to consistently group smaller regions into cohesive bands, even in the presence of noise. Noisy images were divided into the same number of bands (i.e. super-groups) as noise-free images (Fig. 9). This is important for the parameterization stage of the pipeline; band discrimination plays a large part in the calculation of band-based statistics, which in turn contain vital information about the nature of the scattering source.

The scrub() routine helped the parameter extraction process by removing any large noise regions that remained after the initial image reduction. By keeping the scrub threshold low, large features were still preserved (e.g. Fig. 9) while small agent clusters were rejected as noise; this parameter can be tuned to the specific nature (ambient noise and feature size) of the images under analysis. However, it should be noted that setting the scrub level too low can cause erroneous band identification, whereby bands of small pixel mass may appear near the edges of each real band. Extra bands will distort the extracted parameter space and should be avoided.

With regard to the selection of the variables for join() and scrub() (i.e θ and ϕ): these values are derived empirically based on observations regarding the size of noise artifacts inherent in a scattering image and the approximate size and frequency of bands in the image. Once selected, the parameters performed effectively on the entire test set. If significant changes are made to

the ambient background noise or the angular range of a scattering image, the system threshold levels may need to be re-calculated.

While there are a number of other potential frequency analysis methods (such as Fast fourier Transforms) which could be used to ascertain spatial frequency information from a scattering image, the Cythe parameterization routine allowed the extraction of a related structure of image regions within the context of a scattering situation (effectively embedding frequency information within an interpretable band-like structure). This relationship information is of use to human observers, both for validating the extracted parameter data and for comparing results with those previously published in the cell scattering literature (which generally reference the number and size of bands, or the angular location and span of intensity band maxima).

5.4. Remarks on image size reduction

Image size reduction was found to play an important role in both the generalization of region boundaries and the rejection of low-level noise, as it influences the degree of abstraction applied to the input image [31, 32]. A similar reduction approach is used in saliency-based vision systems to detect high-level features in natural scenes, where detail (and the associated noise) is sacrificed to rapidly form an accurate structural impression of the image [31, 32].

Appropriate choice of image reduction size depends on the size of the Agent View Radius, and the number, spacing, and width of intensity bands within an image. A large reduction to an image with very narrow bands or important high-resolution features could merge independent intensity regions, or render some relevant features undetectable. Failing to reduce an image with wide bands could lead to erroneous band detection or extra computational cost / increased run times. We found that disparity between image reduction size and Agent View Radius either lead to the identification of too many small intensity regions (e.g. when the true $aBW_{avg} >> \mathbf{R}$) or the grouping of many distinct initial regions into a small number of larger features (e.g. when the true $aBW_{avg} << \mathbf{R}$).

Image reduction was also essential for manageable run times, as un-reduced experimental cytometer images are typically greater than 700 pixels on a side. At an image reduction size of $\mathbf{I} = 100x100$, an entire pipeline run took approximately six seconds. At a reduction size of $\mathbf{I} = 300x300$ or larger, runs lasted two minutes or more. All performance tests were conducted on a Pentium IV desktop computer. The entire Cythe pipeline and all related routines were implemented in the Python programming language.

5.5. Remarks on versatility

Cythe is expandable and may be readily adapted to new scattering situations; once intensity regions are segmented and numerically represented (as in the group list $\mathbf{G}$) it is possible to test for any number of spatial relationships. In addition, the pipeline should be robust to variations in expected image structure. By adjusting the post-processing settings, agent view radius, and image reduction size, Cythe can be made to detect intensity regions with greatly varying geometric properties. Furthermore, due to the local nature of the parameter calculation equations, slight band curvatures should not adversely affect parameter extraction.

In the event that images with different (e.g. non-vertical) spatial relationships need to be analyzed, the join() routine may be modified or replaced to create a different region hierarchy, and orientation-related changes may be made to the band bridge removal process. This would also allow identification of randomly-placed intensity regions, such as would be generated by a field of Rayleigh scatters, without imposing a band structure on the intensity data. The additional analysis of intensity region perimeter and area would allow further distinctions to be made between differing region types (e.g. independent regions and full bands). This would facilitate the parameterization of heterogeneous images consisting of horizontal bands, tightly grouped

region clusters, blobs arranged without a band structure, or any other arbitrary cluster shape.

To this end, we have used Cythe within other applications, including the identification of geometric objects in natural scenes, and the detection of bright fluorescent regions during the genetic analysis of cell populations [50].

The final goal of the Cythe system is the classification of biological samples based on light scattering. Our preliminary results have shown that cell classes (e.g. those with features indicating cell health or malignancy) typically reside at the extremes of the possible parameter space (work in preparation). The difference between cell classes appears to be much greater than the variation due to noise, such as from imperfections in a fluid wave-guide or CCD. Thus, based on the parameter deviation indicated from random and counting noise (as discussed in Sec 4.2), measurement noise should only moderately detract from Cythes classification ability and the correlations we observed between input and output parameters.

6. Conclusions

In this work we present a multi-agent system (Cythe) to parameterize laser scattering images of the kind produced by a wide-angle 2D cytometer. Extending upon a solid base of tested image processing methods, Cythe uses a three stage pipeline of feature detection, feature clustering, and post-processing to create a parametric representation of an input scattering image. The resulting parameter set numerically represents the complex image features created by light scattering through a cellular body. This facilitates a parametric solution to the inverse scattering problem of laser light through a single biological cell.

Comparison of the Cythe-extracted parameter sets with those derived from a mathematical image model show that our pipeline is able to accurately extract information about the structure and variation present in an image. In addition to our model test images, Cythe is able to parameterize complex FDTD scattering images containing a number of randomly distributed high-frequency intensity regions. Cythe was also able to effectively extract information from images without a noticeable band structure, and has been successfully modified to help detect and parameterize fluorescent genetic material in populations of stained cells [50].

This is possible through the combination of an adaptive feature detection system, an agent-based clustering scheme, and a set of post-processing routines that reject noise and extract high-level information about the relationships between image features. Once a parameter set has been extracted from a scattering signature, it is possible to infer cellular structure from regularities in the extracted parameters [1, 9–11]. For example, we have observed correlations between organelle content in simulated cells and several of the intensity and width parameters present in Eqs. 1 and 2 (work in preparation). To date, no other group has developed an computational system to extract detailed parametric information from wide-angle cytometric scattering signatures.

A rapid method to infer cell characteristics from the information contained in two-dimensional light scattering plots is essential to the further development of wide-angle cytometry systems. We have developed a method of predicting micro- and nano-structural cellular information from the parameters generated by Cythe. A full system integrating Cythe with a machine learning classifier to characterize the organelle content of cells will be presented in future work (in preparation).

Acknowledgments