LA-UR-11-00019

Approved for public release;
distribution is unlimited.

Title: Interactive Image Quantification Tools in Nuclear Material Forensics

Author(s): Reid Porter, Christy Ruggiero, Don Hush, Neal Harvey, Patrick Kelly, Wayne Scoggins, Lav Tandon

Intended for: SPIE Electronic Imaging 2011

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Form 836 (7/06)

This page intentionally left blank.

Interactive Image Quantification Tools in Nuclear Material Forensics

Reid Porter*[a], Christy Ruggiero[b], Don Hush[a],
Neal Harvey[a], Patrick Kelly[c], Wayne Scoggins[d], Lav Tandon[e]
[a] Space and Remote Sensing Sciences, [b] Nonproliferation and Sciences, [c] Information Sciences,
[d] Scientific Software Engineering, [e] Actinide Analytical Chemistry
Los Alamos National Laboratory, Los Alamos, New Mexico, USA, 87545

ABSTRACT

Morphological and microstructural features visible in microscopy images of nuclear materials can give information about the processing history of a nuclear material. Extraction of these attributes currently requires a subject matter expert in both microscopy and nuclear material production processes, and is a time consuming, and at least partially manual task, often involving multiple software applications. One of the primary goals of computer vision is to find ways to extract and encode domain knowledge associated with imagery so that parts of this process can be automated. In this paper we describe a user-in-the-loop approach to the problem which attempts to both improve the efficiency of domain experts during image quantification as well as capture their domain knowledge over time. This is accomplished through a sophisticated user-monitoring system that accumulates user-computer interactions as users exploit their imagery. We provide a detailed discussion of the interactive feature extraction and segmentation tools we have developed and describe our initial results in exploiting the recorded user-computer interactions to improve user productivity over time.

Keywords: image quantification, segmentation, interactive machine learning, nuclear material forensics

1. INTRODUCTION

An important problem in many medical and scientific image analysis applications is image quantification: the indirect measurement and characterization of physical objects through images. While relatively well defined, this problem continues to present a challenge since quantities of interest are very application-specific and accurate quantification typically requires a domain expert with a large amount of experience in interpreting the images available.

In the image processing community, sophisticated algorithms are being developed to help provide consistent image quantification, however the focus is typically to provide general purpose tools. The goals of domain experts and image processing experts are complementary: 1) domain experts would like algorithms to reduce tedious and time consuming tasks as well as help improve consistency in ambiguous situations, and 2) algorithm developers would like to automate as much as possible and parameterize their algorithms in a way that is most useful to end users. In this light, the outstanding problem is how to best allocate image quantification tasks between humans and machines so that the domain expert's (or end-user's) productivity is maximized. Productivity is an application-specific quantity, but it typically characterizes the volume of data processed as well as the quality of quantification. For a general discussion of human-computer interaction with particular focus on image quantification see Chapter 3 in [1].

In this paper we describe our recent work in developing an interactive image quantification system for nuclear materials. We outline the key problems in the application in Section 2 and then briefly describe our system-level architecture and key components in Section 3. In Section 4 we discuss pixel-level interactions that have proven useful for image quantification in a large number of applications. In Section 5 we suggest region-level interactions and describe how they can be used to solve some of the specific image quantification tasks found in nuclear material forensics. In Section 6 we describe how machine learning can also be used to learn region-level interactions and present experiments which quantify the potential user productivity gains. We conclude with a summary of our results and discussion of future directions.

*rporter@lanl.gov; phone 1 505 665-6117; www.pooka.lanl.gov

2. NUCLEAR MATERIAL IMAGE QUANTIFICATION

There are two main types of material that are of interest in our application: particles and metals. Some example Scanning Electron Microscopy (SEM) imagery of these two materials is shown in Figure 1. Many other imaging modalities are also used extensively in nuclear material forensics but were not considered in this paper. The main image quantification tasks identified by nuclear material domain experts include:

1. Accurate particle/grain segmentation with minimal user intervention (accurate segmentation allows for the calculation of accurate particle metrics to contrast and compare samples and particles).
2. Annotation and quantification of process indicators (image features such as friable versus non-friable particles, fracture marks, evidence of scintering, porosity of material etc.).

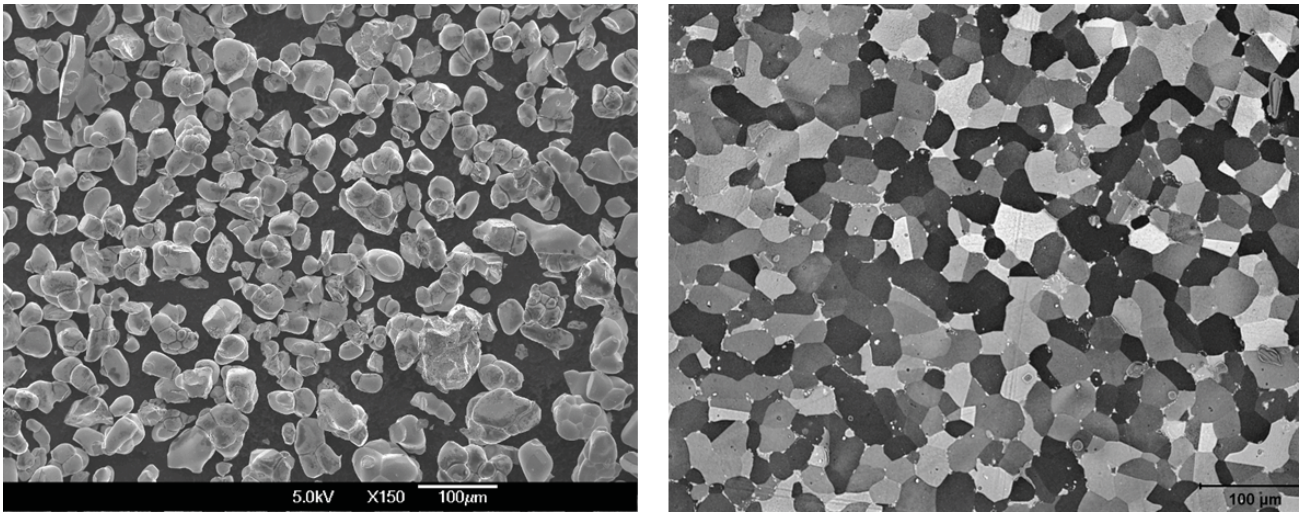

Figure 1. Scanning Electron Microscopy images for Left) Uranium Oxide particles and Right) Plutonium metals.

Similar tasks on similar images have been pursued by researchers in other applications including particles in SEM [2], [3] and optical[4], [5] imagery as well as metal SEM images [6], [7], [8], [9]. Segmentation is often a key step and many papers discuss specific filters and optimizations of standard segmentation algorithms to the specific images and content of interest. In nearly all cases, these papers provide highly automated solutions where the algorithm is highly optimized (by the image processing expert) to obtain good results. In this paper we focus on methods that empower domain experts (or end users) to optimize the algorithm through interactive image quantification systems.

3. INTERACTIVE IMAGE QUANTIFICATION

A large number of image quantification tools have been developed in both commercial and academic communities. Systems in use by nuclear material domain experts span from general purpose tools, such as Adobe Photoshop, through to specialized image acquisition and analysis packages that typically ship with the imaging instrument. In some sense, many of these systems provide end-users with *interactive* image quantification tools, in that users are able to select and apply a number of image processing algorithms (e.g., threshold, smooth, etc.) to solve their image analysis problems. However, we propose that several significant advantages can be gained by incorporating and exploiting interaction to a much greater extent. Specifically, we propose that interaction needs to play a more prominent role in the system-level architecture, and a more integrated role in the choice and application of image quantification tools.

At a system level, there are three main components in a general image quantification system, which are shown in Figure 2. There is a database-type technology for storing and organizing images and associated metadata. There is an image viewer that provides the main interface for users to inspect and annotate images as well as execute image quantification tools. The third component is traditionally given the least attention, and is called workflow tracking. This component captures the sequence of choices that users make during image analysis and records it in the database for future reference. This repository of workflow histories becomes a second dataset which contains a detailed record of how

particular quantification results were obtained (provenance). This record is useful for validation and quality control and several commercial and academic image quantification systems now support this capability in various forms.

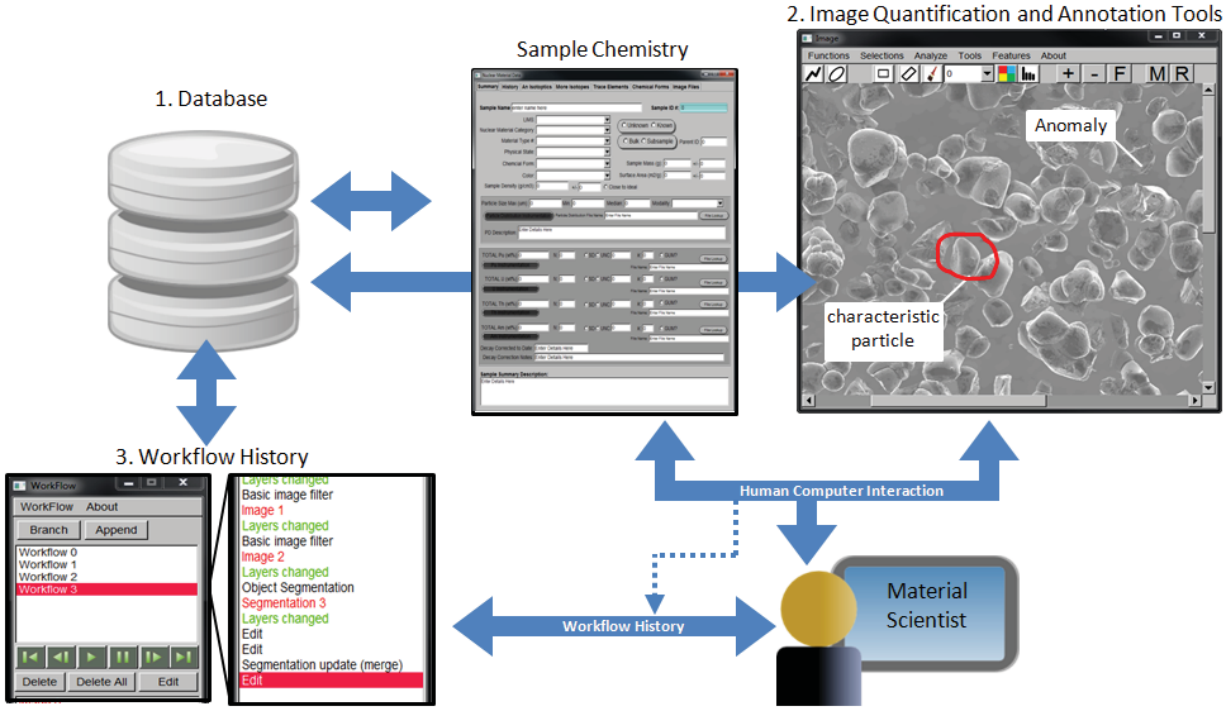

Figure 2. System architecture for an interactive image quantification system showing the key components.

Our objective in implementing a prototype of Figure 2 was to develop a novel combination of image interface, image quantification tools, and workflow tracking with the following properties:

- ***Intuitive Image Quantification:*** users interact directly with the image as much as possible using simple image annotation and mark-up tools to identify important observables that they *see,* rather than trying to quantify their observations by choosing sequences of image processing tools.
- ***Adaptive Image Quantification:*** users are intimately involved in validating and correcting tool errors and this information, which is embedded in the workflow history, is exploited to improve the efficiency and accuracy of the image quantification tools over time.
- ***Knowledge Capture:*** as domain experts use the system to help perform their regular work requirements, domain knowledge is captured explicitly within the workflow tracking system, and implicitly within optimized tools. This knowledge can be shared and used by non-expert users with minimal impact on domain experts.

4. PIXEL-LEVEL INTERACTION

One of the most intuitive ways for users to interact with images is through pixel-level mark-up or annotation. Paint-like drawing tools are provided and a user simply circles, or paints image content that has particular significance. This type of interaction is straightforward to implement since a user is able to identify and color individual pixels within the image regardless of the image content. The interaction can be described by an overlay which has a spatial relationship to the base image but no explicit relationship to image content. Almost all image processing systems provide this functionality and users often quantify their images manually with these tools e.g. manual delineation and classification of image content.

Several researchers have proposed ways to learn and extrapolate pixel-level interactions. In our previous work we developed a machine learning system called GENIE (Genetic Image Exploitation) that learns from the users' annotation of *interesting* versus *non-interesting* to predict future annotations [10]. An example of this system applied to SEM images of nuclear materials is shown in Figure 3. While the system was originally developed for satellite image analysis

the approach is very general purpose and it is now being used in a variety of domains including quantification of digital pathology images [11]. A similar approach describes their interactive tools as *crayons* [12].

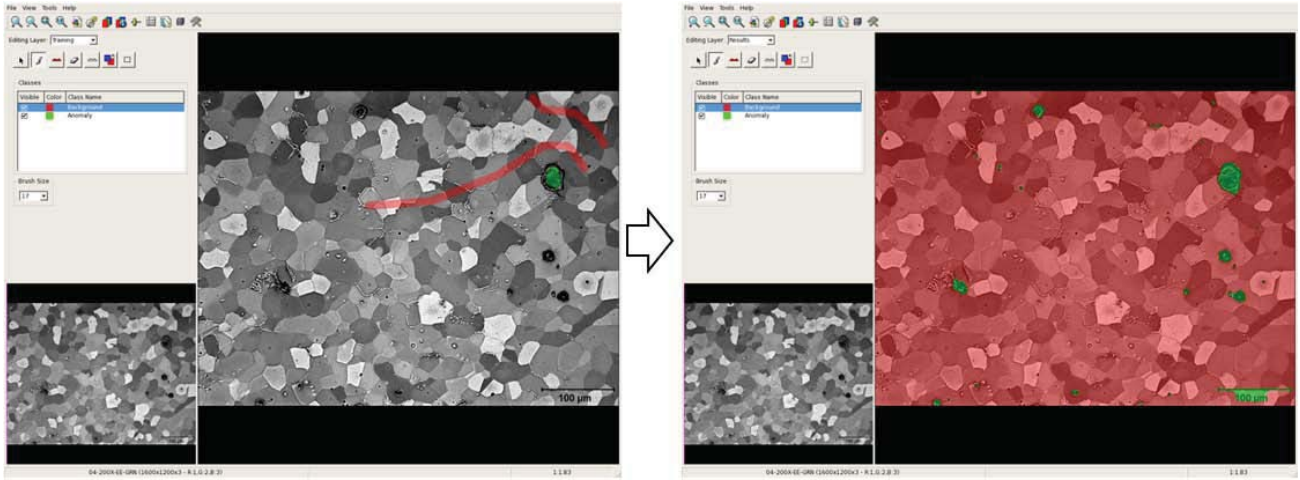

Figure 3. Left) A user uses paint tools to color an anomalous feature of interest green and examples of non-interesting content in red. Right) The GENIE system finds a sequence of image processing operators that can reproduce the annotation, which are then applied across the whole image to predict additional annotations.

Pixel-level classification is not the only way to exploit pixel-level annotations. It is also exploited in some object segmentation algorithms, particularly in medical applications where the there is a single complex object of interest that must be delineated. For example, in active contour approaches a user provides a rough delineation of the object as a starting point for the segmentation algorithm [13]. In other cases, the user can provide *markers*, where pixels belonging to objects and background regions of interest are identified with paint-like tools much like those in Figure 3. These markers are used as seed pixels for graph-cut or watershed segmentation [14], [15].

Pixel-level interaction provides an intuitive and simple way for users to interact with image quantification tools and there are many research questions left to explore. For example, at what time-scale should the user and computer interact? In the examples described, the user and computer take turns during the interaction, but much finer time scales are possible. For example the user and computer may be able to work together to delineate an object of interest in real time [16]. Another open question concerns the complexity of model that should be used to represent how users annotate imagery. In the examples described annotations are assumed independent and even within annotations, pixels are assumed independent. However, as is often the case in hand-written character recognition, the temporal and spatial relationships between and within annotations could also be exploited in some applications [17].

5. REGION-LEVEL INTERACTION

Another way for users to interact with images is through region (or object) level mark-up. In this case users interact with regions defined by image content which means these regions are defined ahead of time using an image segmentation technique. This segmentation is typically error-prone and therefore does not perfectly delineate all regions of interest required by the user. However, it does provide a starting point from which a user can begin to intuitively interact with the image. In our prototype, our initial segmentation is based on a hierarchical watershed, and we allow the user to select the flooding level (with a threshold slider) prior to region-level interaction. This is shown on the left in Figure 4. By interactively adjusting the threshold, the user is able to select a scale appropriate to the content of interest [18].

Given the initial segmentation, the user is then able to select, paint and manipulate the image at a region-level. There are three main types of interaction:

1. Region classification: a user would like to organize the image into different categories of content. On the right in Figure 4 the user is selecting regions to exclude from future analysis, however regions could also correspond to the objects of interest.

2. Region merging: the user selects multiple regions that they would like to manipulate and analyze as a single region. Region merging is usually required in areas of the image which are *over-segmented* in the original

segmentation. On the left in Figure 5 we show an example where a particle has been segmented into multiple regions.

3. Region splitting: a user selects regions that contain multiple image elements that the user would like to treat separately. Region splitting is usually required in areas of the image which are *under-segmented* in the original segmentation. On the right in Figure 5 we show an example where the user is able to selectively split a region. In this case, the user is able to locally adjust the flooding level used in the segmentation with a threshold slider.

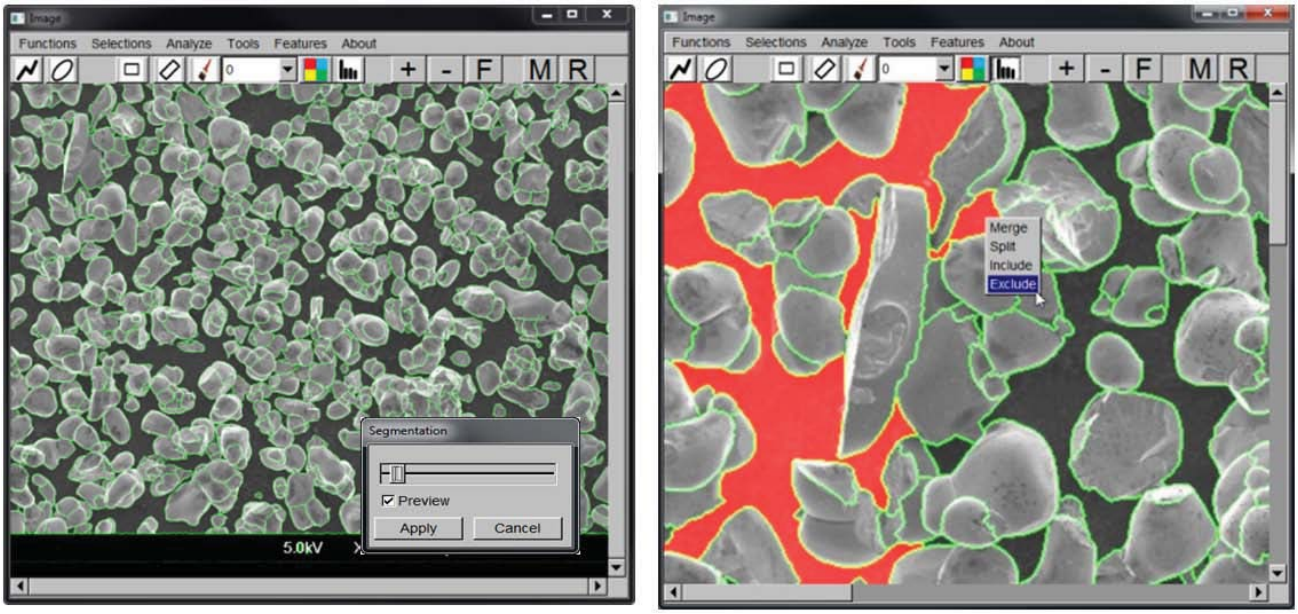

Figure 4. Left) the user interactively sets the flooding level used in the initial watershed segmentation. Right) Region classification: the user is able to select and choose labels for regions in the image.

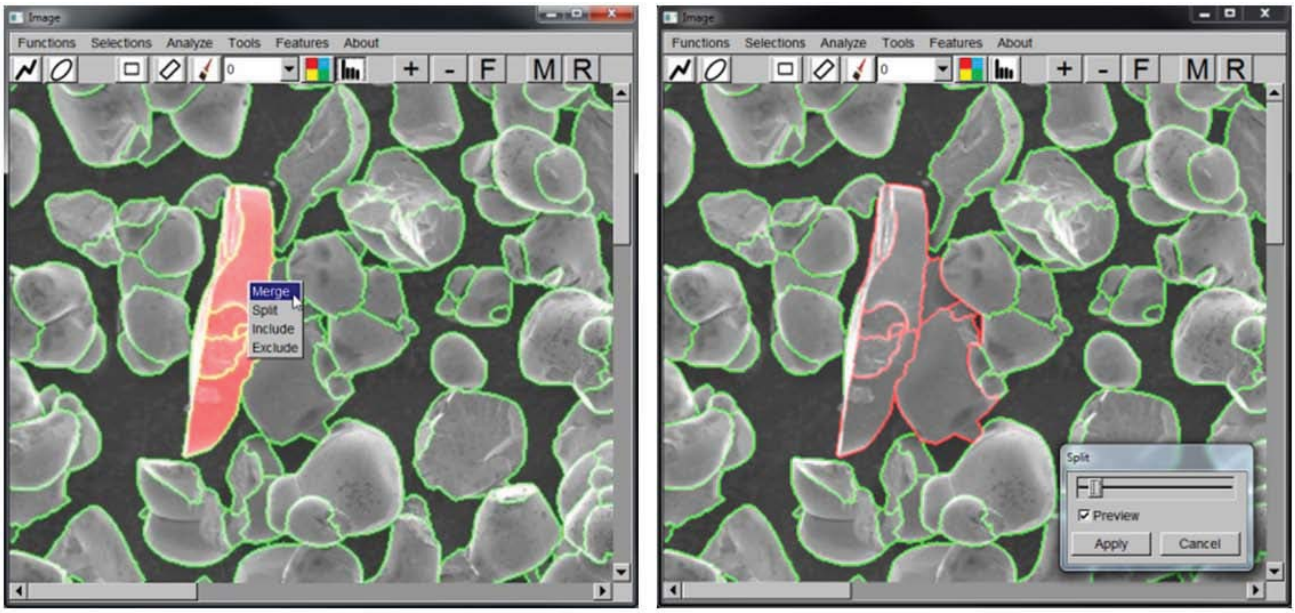

Figure 5. Left) Region merging: the user selects multiple regions that should be treated as a single region. Right) Region splitting: the user selects regions where further segmentation is required and in this case, adjusts the local flooding level used in the watershed algorithm.

6. ADAPTIVE IMAGE QUANTIFICATION

One of the key tasks in our application is to segment and quantify a large number of particles (or grains) so that the population statistics of the particles or grains in the sample can be estimated. This can be accomplished with the region-level interactive tools described in the last section, but it is labor intensive. In this section we describe experiments to quantify how machine learning can help improve the efficiency of the process. In this paper we focus on region classification and region merging. Region splitting and methods to optimize the combination of interactions is a topic of future work.

6.1 Region Classification

To a first approximation, region classification can be posed as a standard classification problem. Regions are assumed to be independent and identically distributed samples from an unknown distribution, and our task is to find a classifier that will predict regions as belonging to one of a number of target classes. However, this problem differs from the standard classification problem in that training data may only be available for one class. For example, in our example in Figure 4, the user begins to *exclude* regions from further analysis. These regions make up the training data or examples of our target class. However the user does not provide examples of which regions should be included, and the remaining regions in the image are some mixture of regions the user would like to exclude and regions the user would like to keep. We call this type of one-side classification problem hemi-supervised learning to differentiate it from semi-supervised learning (which typically assumes training examples from all classes are available).

To investigate the effectiveness of machine learning in this task we manually identified all the regions that should be excluded in sample imagery. A number of basic features for each region are calculated: the gray value mean, and variance as well as the region area and perimeter. For a given number of training samples we apply a machine learning algorithm and evaluate performance on the remaining regions in the image. We compare two different machine learning algorithms for the hemi-supervised problem described:

1. Nearest Neighbor (NN): this one-sided variant of the traditional Nearest Neighbor classifier predicts by simply finding the minimum distance between the feature vector of the test region and the feature vector of all the training samples. We threshold this minimal distance at various points to produce a Receiver-Operator Curve (ROC).
2. Support Vector Machine (SVM): we use the hemi-supervised SVM method we developed in [19]. In this paper we use a Radial Basis Function Kernel with the default parameters found in LibSVM [20]. The hemi-supervised method includes the unlabeled regions within the learning algorithm by assigning different weights to examples. Assuming the user is certain when they exclude regions, we set the SVM cost for errors very high (C = 1e6) and set the weight associated with labeled data very high (rho=1e6) in all our experiments. Although not optimal, we use the real-valued decision values from the SVM to generate ROC curves. This is somewhat justified by our novel method of choosing thresholds for our classifiers which is described next.

To evaluate performance we use a variation of the detection rate. We observe that for a given image, there is a finite amount of work that a user must complete, and this corresponds roughly to the total number of excluded regions in the image which we denote E. We assume a unit of work is required for each region - a user selects each region with a mouse click. Given that a user has already selected some number of regions to exclude, a machine learning algorithm is trained and then applied to the complete image, correctly identifying P of the remaining excluded regions and identifying F false alarms. Assuming a unit of work is also required to discount false alarms (a user deselects the predicted region) the fraction of work required by a user to complete the task compared to the total work is approximately:

$$Work\ Required = \frac{(E - P + F)}{E}$$

On the left in Figure 6 we show an example of this quantity for the two different machine learning algorithms. We used 5 regions as training data. We observed that although not guaranteed, there is often a unique minimum for this curve. In practice, after a machine learning algorithm has been developed, the user is able to move along this curve by adjusting a threshold slider, much like the slider used for segmentation shown on the left in Figure 4. We propose that the interactive slider will allow a user to set the threshold at the point which minimizes the work required.

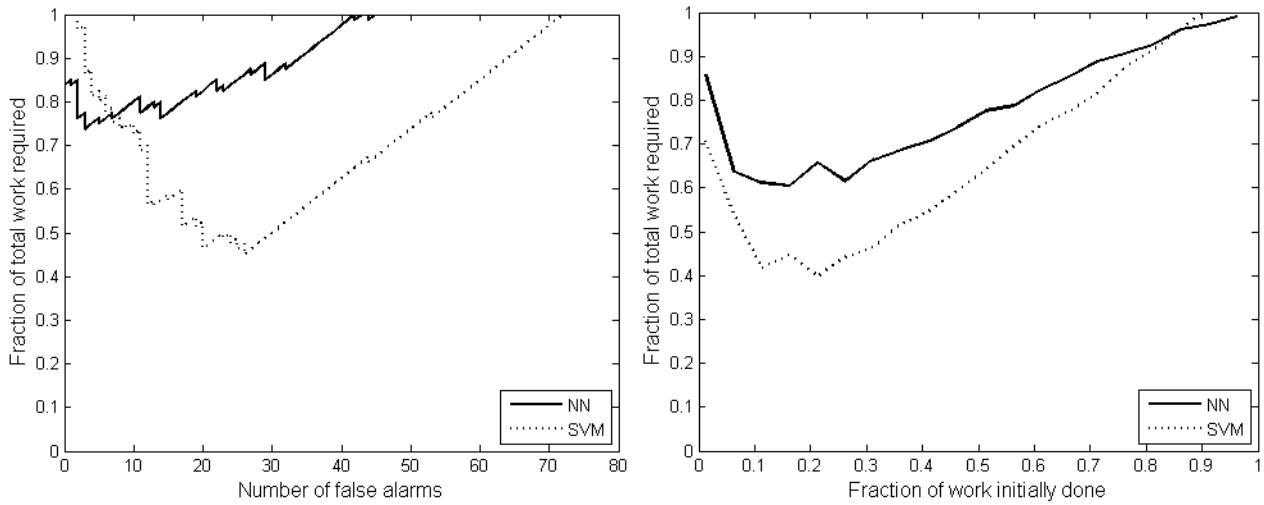

Figure 6. Left) A modified ROC curve used to emphasize work required for Nearest Neighbor (NN) and Support Vector Machine (SVM) based algorithms. Right) The reduction in effort for region classification produced by NN and SVM methods as a function of the amount of training data provided by the user.

Assuming users are able to select the optimal threshold we performed a second experiment, shown on the right in Figure 6, where we plot the fraction of work required versus the number of training samples as a fraction of the total number of regions to be excluded. We observe that the SVM outperforms the NN method regardless of the amount of training data. We also observed that the greatest efficiency gains were made with a very small amount of training data e.g. 2 - 4 regions. By spending their time correcting the classifier output instead of performing the task manually, the user is able to complete the task with approximately 40% of the effort. However, our model for work does not include a cost for the user to validate the correct classifier outputs (only classifier errors are penalized) which means our estimate is somewhat optimistic.

6.2 Region Merging

Region segmentation is similar to region classification in that we are only given positive examples of which regions to merge. We therefore approach region merging as a hemi-supervised learning problem. However unlike region classification, even a first order approximation will require the classifier to consider multiple regions. To investigate this further we over-segmented SEM imagery of particles, and then manually merged regions corresponding to particles of interest. The number of regions involved in each merge event (each training sample) varied from 2 to 8. We calculate features for these merged regions, and in our experiments we used the first six Hu invariants calculated from the spatial moments [21]. For the test set we exhaustively generate all possible combinations of regions (subject to connectivity constraints) up to a maximum number of regions (in our case 10) and calculate the corresponding Hu invariants. We then performed a similar experiment to Region Classification and the results are shown on the left in Figure 7. We observe that the learning algorithms are able to again provide a substantial reduction in effort with very few training examples. However in this case the NN and SVM methods obtain similar performance.

In our experiments we purposely chose images where particles had little overlap and this led to images with sparse connectivity. This allowed us to increase the maximum number of regions considered during prediction and investigate the effect that this number has on classification performance. On the right in Figure 7 we show performance with 4 training samples, at various values of this parameter. Somewhat surprisingly we found that the best performance was obtained when we only considered pair wise merging events. This is in contrast to the training data which had merging events up to size 8. This can be partly attributed to our method of evaluating effort for merges which heavily penalized incorrect merging: If a predicted merge event contained 6 regions, and one of these regions was a false alarm, then the prediction cost 6 units of work due to the added complexity required by the user to split and correct the predicted merge.

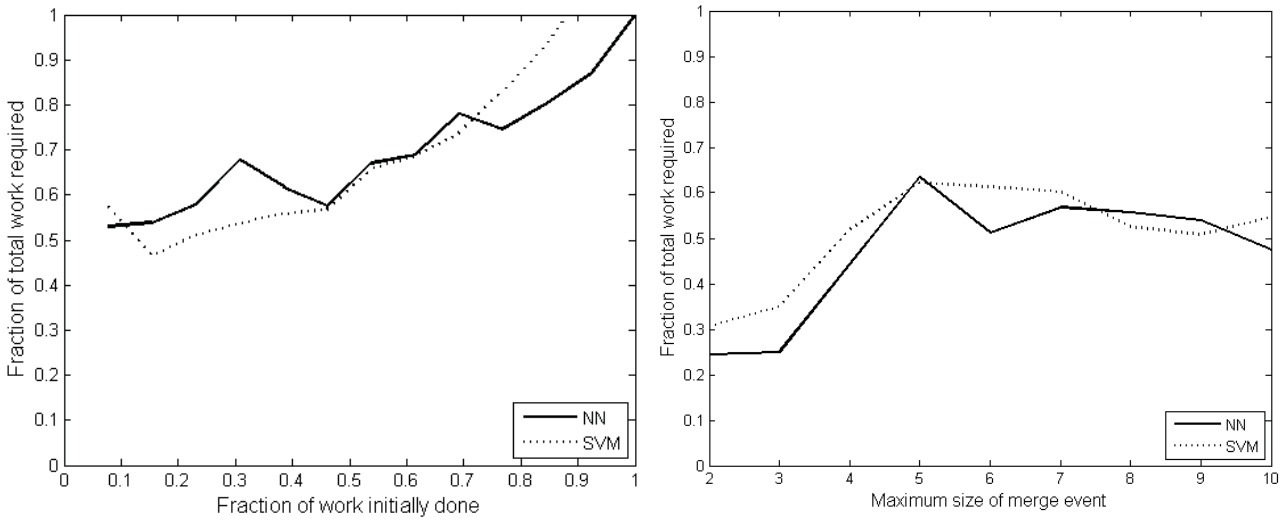

Figure 7. Left) The reduction in effort for region merging produced by NN and SVM based algorithms as a function of the amount of training data. Right) the reduction in effort produced by algorithms when merge predictions are constrained to a maximum number of regions.

7. SUMMARY

Nuclear material forensics suggests a relatively new and unexplored application for interactive image quantification tools: the accurate and consistent segmentation of multiple particles / grains within microscope images. In this paper we have presented a new way to improve the efficiency of this process through region-level interaction and hemi-supervised learning. Our experiments produced promising results on real data and future work will continue to refine the approach by exploring larger classes of images and features.

We have also presented a system level architecture for interactive image quantification that includes workflow tracking as a basic component. This will allow us to explore and exploit user interactions at a system-level. Specifically, we plan to observe how users combine region merging, splitting and classification during image quantification and will investigate how more sophisticated user modeling might help to maximize system-level performance metrics and capture domain knowledge within optimized tools.

ACKNOWLEDGEMENTS

We gratefully acknowledge the support of the U.S. Department of Homeland Security through the Domestic Nuclear Detection Office for this work.

REFERENCES

1. Hahn, H.K., *Morphological Volumetry, Chapter 3*, in *MeVis - Center for Medical Diagnostic Systems and Visualization*. 2005, University of Bremen. p. 37-47.
2. Coster, M. and J.L. Chermant, *Image analysis and mathematical morphology for civil engineering materials.* Cement and Concrete Composites, 2001. **23**(2-3): p. 133-151.
3. Xavier Arnould, M.C., Jean-Louis Chermant, Liliane Chermant, Thierry Chartier, Abder Almoataz, *Segmentation and Grain Size of Ceramics.* Image Analysis and Stereology, 2001. **20**: p. 5.
4. Luc Vincent, E.R.D., *Morphological Segmentation for Texture and Particles*, in *Digital Image Processing Methods*, E.R. Dougherty, Editor. 1994, Marcel Dekker: New York. p. 43 - 102.
5. Wang, W., *Image Classification and Segmentation for Densely Packed Aggregates*. 2007. p. 887-894.
6. Benesova, W., A. Rinnhofer, and G. Jakob. *Determining the average grain size of super-alloy micrographs*. in *2006 International Conference on Image Processing*. 2006. Piscataway, NJ, USA: IEEE.
7. Henault, E., M. Coster, and J.L. Chermant, *Grain boundary reconstruction by automatic image analysis on etch pits.* Materials Research Bulletin, 1991. **26**(7): p. 569-75.

8. Coster, M., et al., *The use of image analysis for sintering investigations: The example of CeO2 doped with TiO2.* Journal of the European Ceramic Society, 2005. **25**(15): p. 3427-3435.
9. Dengiz, O., A.E. Smith, and I. Nettleship, *Grain boundary detection in microstructure images using computational intelligence.* Machine Vision, 2005. **56**: p. 854-866.
10. Perkins, S., et al. *Genie Pro: Robust Image Classification using Shape, Texture, and Spectral Information.* in *Proc. SPIE 5806* 2005.
11. Inc., A.T. *GENIE Histology Pattern Recognition*. 2010; Available from: http://www.aperio.com/imageanalysis/Genie-Histology.asp.
12. Fails, J.A. and J. Dan R. Olsen. *Interactive Machine Learning*. in *Intelligent User Interfaces, IUI '03*. 2003: ACM.
13. Kass, M., A. Witkin, and D. Terzopoulos, *Snakes: active contour models.* International Journal of Computer Vision, 1987. **1**(4): p. 321-331.
14. Boykov, Y. and M.-P. Jolly. *Interactive Graph Cuts for Optimal Boundary & Region Segmentation of Objects in N-D images.* . in *International Conference on Computer Vision*. 2001.
15. Klava, B. and N.S.T. Hirata. *Watershed segmentation: Switching back and forth between markers and hierarchies.* in *Proceedings of ISMM 2007 - 8th International Symposium on Mathematical Morphology*. 2007.
16. Pavlopoulou, C., A.C. Kak, and C. Brodley. *Application of Semi-supervised and Active Learning to Interactive Contour Delineation* in *Proceedings ICML-EDDS Workshop on the Continuum from Labeled Data to Unlabeled Data*. 2003. Washington DC.
17. Caelli, T., A. McCabe, and G. Briscoe, *Shape tracking and production using hidden Markov models*, in *Hidden Markov models: applications in computer vision*. 2002, World Scientific Publishing Co., Inc. p. 197-221.
18. Hahn, H.K. and H.-O. Peitgen. *IWT - Interactive Watershed Transform: A hierarchical method for efficient interactive and automated segmentation of multidimensional grayscale images*. in *Proc. of SPIE* 2003.
19. Hush, D.R., R.B. Porter, and C.E. Ruggiero. *Density-based similarity measures for content based search*. in *Asilomar Conference on Signals, Systems and Computers* 2009. Pacific Grove, CA.
20. Chang, C.-C. and C.-J. Lin, *LIBSVM: a library for support vector machines*. 2001.
21. Ming-Kuei, H., *Visual pattern recognition by moment invariants.* Information Theory, IRE Transactions on, 1962. **8**(2): p. 179-187.