

Automated classification of protein crystallization images using support vector machines with scale-invariant texture and Gabor features

Shen Pan,^a Gidon Shavit,^a Marta Penas-Centeno,^b Dong-Hui Xu,^c Linda Shapiro,^{a,d*} Richard Ladner,^a Eve Riskin,^d Wim Hol^e and Deirdre Meldrum^d

^aDepartment of Computer Science and Engineering, University of Washington, Seattle, USA, ^bDepartment of Computer Science, University of A Coruña, Spain, ^cSchool of Computer Science, Telecommunications and Information Systems, DePaul University, Chicago, USA, ^dDepartment of Electrical Engineering, University of Washington, Seattle, USA, and ^eDepartment of Biological Structure, University of Washington, Seattle, USA

Correspondence e-mail:
shapiro@cs.washington.edu

Received 23 November 2005

Accepted 13 December 2005

Protein crystallography laboratories are performing an increasing number of experiments to obtain crystals of good diffraction quality. Better automation has enabled researchers to prepare and run more experiments in a shorter time. However, the problem of identifying which experiments are successful remains difficult. In fact, most of this work is still performed manually by humans. Automating this task is therefore an important goal. As part of a project to develop a new and automated high-throughput capillary-based protein crystallography instrument, a new image-classification subsystem has been developed to greatly reduce the number of images that require human viewing. This system must have low rates of false negatives (missed crystals), possibly at the cost of raising the number of false positives. The image-classification system employs a support vector machine (SVM) learning algorithm to classify the blocks making up each image. A new algorithm to find the area within the image that contains the drop is employed. The SVM uses numerical features, based on texture and the Gabor wavelet decomposition, that are calculated for each block. If a block within an image is classified as containing a crystal, then the entire image is classified as containing a crystal. In a study of 375 images, 87 of which contained crystals, a false-negative rate of less than 4% with a false-positive rate of about 40% was consistently achieved.

1. Introduction

Protein crystallography provides access to structural information with atomic resolution for macromolecules such as DNA, RNA, saccharides and proteins. This structural information helps biologists to understand the function of the macromolecular systems investigated. Understanding structure–function relationships can then also help projects aiming at blocking or improving the properties of biomacromolecules. One example is ‘structure-guided drug design’ (*e.g.* Hol, 1986; Veerapandian, 1997; Guillemont *et al.*, 2005), where drug designers use detailed knowledge of the shape of key cavities and binding modes of ligands to improve the affinities and pharmacological properties of ‘leads’, *i.e.* of potential pharmaceuticals. In other areas knowledge of, for example, the insulin structure has led to research to obtain variants of human insulin with improved pharmacokinetic properties (see, for example, Vajo *et al.*, 2001).

Crystallographers go through several essential experimental steps. During the protein crystallization stage, proteins are injected into many different solutions in the hope of

crystal growth. The crystallographers then identify the successful experiments that yield mountable crystals. The final stage is the computation of protein structures from the X-ray diffraction patterns of protein crystals.

Currently, during the stage of identification of successful experiments, human crystallographers need to manually examine images taken by microscopic cameras. With the increasing number of crystal-growth experiments, the total is up to millions of images per year. The process is both time-consuming (hence expensive) and prone to human error. Automating this task is therefore an important goal. As part of the capillary-based protein crystallography instrument project, we have developed a new image-classification subsystem to greatly reduce the number of images that require human viewing.

Fig. 1 shows some example images of protein crystallization trials taken by a Structural Genomics of Pathogenic Protozoa (SGPP) RoboDesign microscope camera (RoboDesign International Incorporated, Carlsbad, CA, USA). As can be observed, great variability is present in the images on which we need to perform classification. Different drop shapes, crystal shapes and precipitate clouds in the images make automatic classification a very challenging machine-learning problem.

2. Related literature

Research on this problem is still in its early stages. Only a few papers have appeared on the subject and standard image collections and benchmarks do not yet exist. In early work, Zuk & Ward (1991) used the Hough transform (Shapiro &

Stockman, 2001) to identify straight edges of crystals, but did not attempt to classify precipitates.

Wilson (2002) used the Sobel edge detector to locate drop boundaries and detected objects (connected component of edge pixels) with high circularity. The classification of objects inside the drop boundary was based on features computed from edge pixels. Their results had 86% of true crystals classified as crystals, while 77% of precipitates were classified as precipitates. Jurisica *et al.* (2001) used a custom-built image acquisition and image-processing system to analyze protein crystallization experiments under oil. They detected drop boundaries by fitting a conic curve and classified images based on spectral analysis.

Spraggon *et al.* (2002) applied the Canny edge detector and circle fitting on the largest connected component of edge pixels to drop-boundary detection. Their features included geometric features related to straight lines (detected by the Hough transform) and texture features based on correlations between intensities at various distances and directions. For classification, they used a self-organizing neural network, achieving 25% false-negative and false-positive rates.

Cumbaa *et al.* (2003) described a custom-built system which uses a probabilistic graphical model for drop-boundary detection and for classification. Features are based on correlation filters and the Radon transform, which also detects straight lines (Weisstein, 1999–2005). The classifier achieves a balanced error rate of 15% for binary classification of crystal-positive and crystal-negative. Bern *et al.* (2004) used edge detection followed by dynamic programming curve tracking to determine the drop boundary. They used a decision-tree classifier with hand-crafted thresholds operating on features




Figure 1

Examples of images of crystallization trials taken by an SGPP RoboDesign microscope camera. (a) Clear; (b) clear with an air bubble; (c) crystal with an air bubble; (d) precipitate; (e) precipitate with an air bubble; (f) complex precipitate (or denatured protein).

derived from the Hough transform and from curve tracking. The classification achieves a false-negative rate of 12% and a false-positive rate of 14%.

The best result so far is reported in Zhu *et al.* (2004) using geometric features to characterize local pixel gradients and texture features derived from the gray-level co-occurrence matrix. For classification, a support vector machine (SVM) is compared with an automatic decision-tree classifier, achieving a false-positive rate of 14.6% and a false-negative rate of 9.6%.

The results reported above are not directly comparable with the 40 and 4% reported here because of the different wells used to perform protein crystallization and the different sets of both training and testing images.

3. System overview

The system we have built accomplishes several important goals. First of all, it handles the required throughput (thousands of images per day). Secondly, it correctly rejects the majority of the non-crystal images, maintaining a false-positive rate of less than 40%. This means that three out of five images without crystals do not need to be manually examined. Since the majority of images do not contain crystals, this is a significant time saving. Most importantly, the system consistently identifies images with crystals correctly, with less than 4% error. Finally, it can be integrated into the overall system pipeline illustrated in Fig. 2 and easily used by crystallographers with no special computer skills.

Fig. 2 shows the role our imaging subsystem plays in the overall crystallography system. It receives images taken by a RoboDesign microscopic camera of plates containing results from completed experiments. The classifier then labels the images it receives as positive (containing crystals) or negative (containing no crystals). Images classified as containing no crystals are discarded. Images classified as containing crystals will be transferred to a human viewing station for further confirmation. Based on results from the human viewing station, successful experiments are identified and crystallographers can proceed to the next stage.

Our image-analysis system consists of five stages. During the first stage, the system preprocesses the raw image obtained directly from the microscopic camera. In the second stage, it

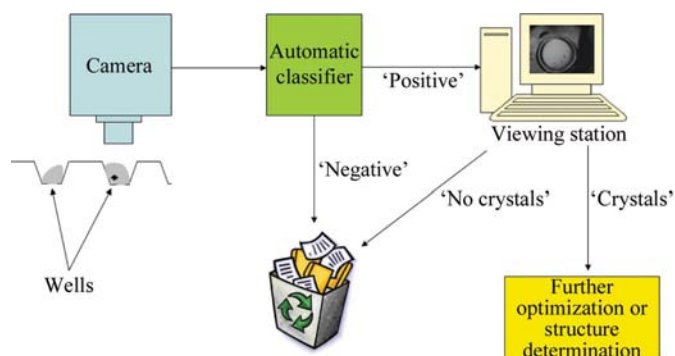


Figure 2
System overview.

performs image segmentation to locate the area of interest and crop the image. In the third stage, it performs feature extraction to obtain features that will serve as numerical input to the classifier. The system then divides the image into small overlapping blocks of size 40×40 pixels in steps of 20 pixels. The size of the blocks was determined experimentally by the average size of crystals on images. It was selected to be large enough to encompass entire crystals, but small enough to not present a huge computational burden. Sliding blocks are used to provide a finer coverage of the image. Because we classify the entire image as containing a crystal if any of the overlapping blocks are classified positive, we bias the classifier to detecting crystals when they occur.

In the fourth stage, for each block, the system computes the numerical representations of features extracted in the third stage. In the fifth and final stage, the system utilizes a trained SVM to classify each block by its numerical feature vector. An image is classified as containing crystals if at least one block in the image is classified positive by the SVM.

4. Image segmentation

Fig. 3 shows enlarged top and side views of a well being imaged. In the images taken by the microscopic camera, only the regions associated with the solution drop inside the well are of interest to the classification step. Several problems arise in the classification step if one tries to classify the entire image. First of all, random noise outside the region of interest is very likely to be classified as crystals. Secondly, boundary edges of the well and of the occasional air bubble inside the drop often show up as high-contrast textures or significant segments in feature extraction, which could potentially mislead the classifier. Finally, computation time is wasted on processing information that should not matter in the first place. Therefore, finding and segmenting the solution drop area precisely is an important task.

Unfortunately, the images can be very noisy owing to clouds of precipitants and other artifacts on the plate outside of the well. The solution drop is sometimes difficult to locate because drop shrinkage owing to evaporation often results in bubbles. Furthermore, there are many variations in image characteristics, such as well position and size, general lighting and focus depth.

To locate the well, we use an algorithm that finds the best approximation circle in the image. We then remove the regions containing bubbles from the area of interest.




Figure 3
The well on a plate.




Figure 4
Illustration of the algorithm for locating the well. (a) Original image; (b) Canny edge-detected image; (c) final image with located well.

4.1. Locating the well

We first examined the possibility of using the classic circle-finding algorithm based on the Hough transform (Shapiro & Stockman, 2001). The algorithm produces accurate well-detection results for clear images. However, it fails to work well for images in which the well boundary is obscured by precipitants inside the well and there is excessive noise. For those images, the algorithm often yields a circle that is off from the center of the well. Therefore, we designed a new algorithm that better suits our application.

The well-detection algorithm is illustrated by example in the images of Fig. 4. The algorithm first preprocesses the image using a median filter to reduce noise. It then applies the Canny edge detector to the image to obtain an edge image. It then fits a circle to the edge pixels in the image using a novel algorithm that takes advantage of the fact that concentric circles of edge pixels are often observed in the images.

The key idea of the new circle-fitting algorithm is illustrated by Fig. 5: a Euclidean (x, y) to polar (r, θ) coordinate transformation about the center of a circle turns the circle into a straight line in the new coordinate system. This idea can be easily used to determine the correctness of the center of the circle determined to be the best approximation of the well. For each possible center, we transform an annulus about it to a rectangle. The rectangle with the straightest lines is selected and the corresponding center is the best candidate for the center of the well. To determine the straightness of the lines in the transformed rectangular image, we compute the variance of the means of pixel intensities of all columns in the transformed image. The higher the variance, the more contrast between the columns and therefore the straighter the lines in the image.




Figure 5
Euclidean (x, y) to polar (r, θ) coordinate transformation of an annulus.

One problem introduced by this algorithm is the selection of the size of the annulus on which we perform the transformation. This is closely related to the question of how much information is needed for each possible center for the algorithm to make the best decision. Owing to random noise in the images, some images work better with a larger annulus, while others work better with a smaller one. The solution to this problem is to record for each possible center only the highest variance computed among a range of annulus sizes.

Since the wells are located within a certain range on all the images, we only need to perform the coordinate transformation and variance computation for a small number of center candidates. The center candidate with the highest computed variance is determined to be the center of the best approximation circle for the well. After the correct center is located, the radius of the circle must be determined. This can be accomplished by finding the radius that makes the circle overlap with the most edge pixels, because it is often the case that well boundaries contain the largest number of adjacent edges in an image.

The results produced by this algorithm are excellent. We tested the algorithm on 514 images of various noise levels and drop shapes. Out of those 514 images, the algorithm produced a slightly off-center circular fit to the well for only five images, a success rate of 99%.

4.2. Locating and removing bubbles

Fig. 6 shows the process of finding the bubbles inside an image. The algorithm first segments the well into regions with gradual changes in pixel intensities by growing regions along the boundary of the well. The region growth is performed by a classic flood algorithm that recursively grows a region pixel by pixel until the image gradient at a pixel exceeds a certain predefined threshold in all directions. The recursion is implemented using a stack for performance. The algorithm then checks to see if regions contain holes (produced by edge pixels) and removes those regions from the list of potential bubbles. Finally, the algorithm thresholds the list of bubbles by pixel-intensity statistics and sizes (bubbles are generally not very large). The remaining regions are then marked as bubbles.

Problems arise in the determination of the thresholds. Currently, the thresholds are set manually through experimenting with the images. Some bubbles in precipitate images do not have well defined boundaries and the algorithm fails to produce bubbles that are visually identifiable. The final product of the well- and bubble-locating routines is a mask of the region of interest.

5. Feature extraction

Feature extraction is the most important stage in the system, because for the SVM classifier to classify well the features must distinguish well between crystals and non-crystals. After much consideration and experimentation, we selected the following features: (i) pixel-intensity statistics, (ii) BlobWorld texture features (UC Berkeley; Carson *et al.*, 2002) and (iii) results from Gabor wavelet decomposition.

Pixel-intensity statistics for a block are the mean and standard deviation of the intensity values in that block. In the following sections, the texture and wavelet features are explained in detail.

5.1. BlobWorld texture features

Although texture is a well studied property, most work extracts texture features either on a global scale or on a uniform texture scale across the whole image. Since texture is a local neighborhood property, texture features computed on the wrong scale would lead to confusion. In order to find an appropriate scale for each pixel, we adapted the method of scale selection based on edge-polarity stabilization, which is derived from the gradient of gray-level intensity (Carson *et al.*, 2002). Polarity is a measure of the extent to which the gradient vectors in a certain neighborhood all point in the same direction. The proper scale is selected based on the derivative of the polarity with respect to scale. Polarity is calculated at every pixel in the image for various scales and each polarity image is smoothed using a Gaussian filter. The best scale is

selected based on the difference between the values of polarity at successive scales (Carson *et al.*, 2002).

After a scale is selected for each pixel, three texture descriptors, which are derived from the gradient of gray-level intensity, are assigned to that pixel. The three texture descriptors are polarity at that scale, anisotropy and normalized texture contrast. Anisotropy measures the degree of local orientation. Texture contrast defines the homogeneity of the pixels within the local region (Carson *et al.*, 2002). The results of a sample texture analysis are shown in Fig. 7.

5.2. Gabor wavelet decomposition

Gabor wavelets (Gabor, 1946) are complex exponential signals modulated by Gaussians and are widely used in computer vision for edge detection (Mehrotra *et al.*, 1992; Wang & Jenkin, 1992), noise filtering (Donoho, 1995; Cristóbal & Navarro, 1994), image compression (Daugman, 1983; Porat & Zeevi, 1988), texture analysis and synthesis (Turner, 1986; Porat & Zeevi, 1988) *etc.* Specifically, Gabor wavelets are suitable for edge detection as a consequence of two important properties: good edge localization (Van Deemter & Du Buf, 2000) and the absence of image-dependent parameter tuning.

In this work, Gabor wavelets have been used for the extraction of the low-level directional primitives present in an image. Precisely, a bank of eight filters centered at frequency $\frac{1}{4}$ and eight orientations ($k\pi/8$, $k = 0, \dots, 7$) has been used. The Gabor decomposition through the bank of wavelets maps the input image into eight resulting images that have been integrated into a unique result, called the response. To compute the response, each pixel (x, y) is assigned an eight-component vector $H(x, y)$ composed of its Gabor decomposition results. The response is then computed through

$$\text{response}(x, y) = 0.5 + \arctan \left[\frac{(|H(x, y)| - |\bar{H}|)}{|\bar{H}|} \right] / \pi, \quad (1)$$

where $|H(x, y)|$ is the modulus of $H(x, y)$, $|\bar{H}|$ is the modulus of the mean of $H(x, y)$ over the whole image and $|\bar{H}|$ is the square

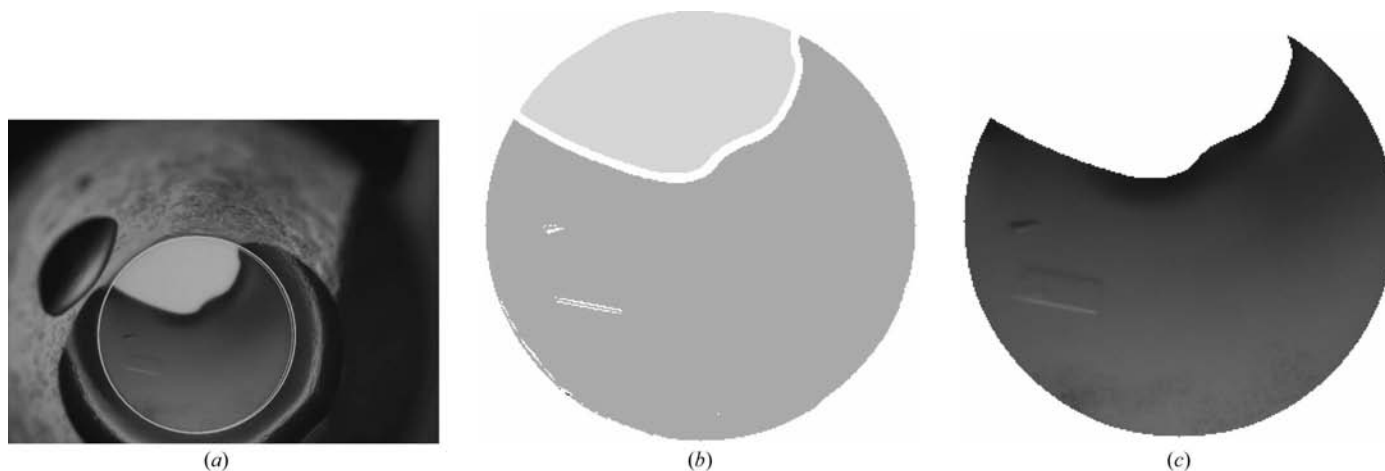


Figure 6

Illustration of the algorithm for locating the bubbles. (a) Original image with well located; (b) bubble regions found inside the well, where each region appears as a different shade of gray and edges are in white; (c) final region of interest.

Table 1
Features.

Pixel intensity	
F1	Mean of pixel intensities
F2	Standard deviation of pixel intensities
Texture	
F3	Maximum of polarity
F4	Minimum of polarity
F5	Mean of polarity
F6	Standard deviation of polarity
F7	Maximum of anisotropy
F8	Minimum of anisotropy
F9	Mean of anisotropy
F10	Standard deviation of anisotropy
F11	Maximum of contrast
F12	Minimum of contrast
F13	Mean of contrast
F14	Standard deviation of contrast
Gabor wavelet decomposition	
F15	Maximum of Gabor response
F16	Minimum of Gabor response
F17	Mean of Gabor response
F18	Standard deviation of Gabor response

deviation of the mean of $H(x, y)$ over the whole image. The response is bounded by the range $[0, 1]$; high response values indicate the presence of a directional primitive in the pixel. The response is the only Gabor feature currently employed in our work.

5.3. Blocking and feature vectors

After feature extraction, we need some way of summarizing the features into numerical values which are the required input to the SVM classifier. The way we accomplish this is to first divide the region of interest into overlapping blocks of size 40×40 pixels in steps of 20 pixels. We then summarize for each block a vector of numerical values from the output of the feature-extraction process. These overlapping blocks are the key to our low false-negative rates. Any crystal that is missed by one block is very likely to be captured by a nearby overlapping block. The feature vector computed for a particular block determines the classification of the block. We have a




Figure 7
Example of BlobWorld texture. (a) Original image; (b) texture image with polarity, anisotropy and contrast values as its color. It is possible to see the different texture of the crystals. One example would be the crystal located near the center of the well showing up as having a different color from its neighborhood in the texture image.

total of 18 feature values for each block (see Table 1). All feature values above are computed on a per-block basis.

6. Classification using support vector machines

Each block feature vector is classified using a pre-trained support vector machine. SVMs are binary classifiers. They first map the feature vector into a very high-dimensional space *via* a non-linear kernel function. They then cut the space in half with a linear hyperplane. This hyperplane is chosen to maximize the margin of separation between positive and negative samples in a training set. The image is classified as containing crystals if any one of its blocks is classified positive by the SVM. Fig. 8 is an overview of the entire process of training and classification. Both training and classification use very similar processes.

6.1. Training support vector machine

For training, we manually label (positive/negative) blocks in selected training images to generate a list of labeled (classified) image blocks; images whose blocks appear in that list are processed, a vector of numeric features is computed for each block on the list and the generated list of labeled vectors is used to train the SVM classifier. The classifier learns to differentiate between positive and negative samples.

6.2. Classification

Classification of new images proceeds similarly: the image is processed as before and feature vectors are computed for all blocks in the region of interest in the image. These vectors are then classified using the previously trained SVM. Fig. 9 shows some classification examples.

7. Results

Experiments were performed over a data set of 375 images annotated by hand as positive or negative. This data set consists of 138 clear images, 150 precipitate images and 87 crystal images. These images were obtained by a SGPP RoboDesign microscope camera. An average of 4–5 blocks were manually selected and labeled from each of the images in the data set. The selected blocks in a positive image contain crystals and are labeled as positive. On the other hand, the selected blocks in a negative image contain non-crystals (either clear or precipitate) and are labeled as negative.

A total of four different experiments of varying sizes of the training set were performed over the data set. In each of the experiments, the full data set was divided randomly into two disjoint data sets: a training set and a test set. For the

four different experiments, the training sets were made up of 50, 60, 70 and 80%, respectively, of the image database and testing was performed on the remainder¹. Furthermore, the percentages of clear (36.8%), precipitate (40%) and crystal (23.2%) images were kept constant in each of the four training sets (and corresponding test sets) to agree with the percentages in the full data set.

In each experiment, ten separate trials were performed and the false-negative rates and false-positive rates were averaged over the ten trials. In each trial, the feature vectors of the manually selected blocks in each image of the training set served as the input to train a classifier. The trained classifier then classified all the images in the test set for that trial. The resulting false-negative rate and false-positive rate for each trial were computed based on the classification results over whole images, rather than the blocks in the images. The results of all four experiments are shown in Table 2. Note that because the number of training blocks per image varies, the average number of training blocks for the ten trials is given for each of the four experiments.

It is easy to see that the false-negative rate and the false-positive rate consistently improved as the size of the training set increased. The best overall results, a false-negative rate of 2.94% and a false-positive rate of 37.68%, were obtained from the experiment with 80% of the data set as the training set. This series of experiments was performed in the hope of determining the best number of training images or blocks that are needed when one is required to use the system in practice

¹ We initially tried experiments with 90% of the data in the training set, but the remaining 10% were not sufficient for a test.

Table 2

Results from four experiments of varying sizes of the training set.

Training-set size (%)	50	60	70	80
No. of training images	187	225	263	300
No. of clear training images	69	80	97	110
No. of precipitate training images	75	90	105	120
No. of crystal training images	43	55	61	70
No. of training blocks (averaged over ten trials)	640	772	897	923
No. of crystal test images	44	32	26	17
False negatives (%)	3.86	3.75	3.42	2.94
False positives (%)	42.50	41.19	40.58	37.68

with thousands of images. Although the data set over which the experiments were performed is too small to accomplish this goal, a trend of diminishing returns is clearly visible from the results.

Finally, Figs. 10 and 11 show some common misclassifications in these experiments. These misclassifications arose from such factors as obscured crystals, crystals mixed with precipitate, small crystal-shaped bubbles, precipitate at the edge of a bubble, precipitate containing edges and impurities that look like crystals.

8. Conclusion and discussion

The results presented here show that the system has an extremely low false-negative rate of 2.9%, while maintaining a tolerable false-positive rate of 37.7%. The false-negative rate is low partly because the system is classifying overlapping blocks of interest in the images so that none of the details of an image is missed. However, this also contributes to the fact that it has a relatively high false-positive rate. This shows that the

features described here are not sufficient to completely distinguish crystal blocks from non-crystal blocks. There is a need for the inclusion of new features that take advantage of subtle cues such as the presence of corners and parallel segments that humans use to classify images. Attempts to include features of perceptual groupings (Penas, 2005) that capture high-level structures composed of line segments were not successful owing to the difficulties involved in summarizing such features numerically. Better summarizing schemes need to be devised in order to best incorporate features from these perceptual groupings into the system. In addition to the inclusion of better features into the system, there is the possibility of developing a better classifier that classifies crystal and non-crystal images into subclasses to help the optimization step of the protein crystallization experiments. Furthermore, optimizations in performance are under consideration since currently the system takes

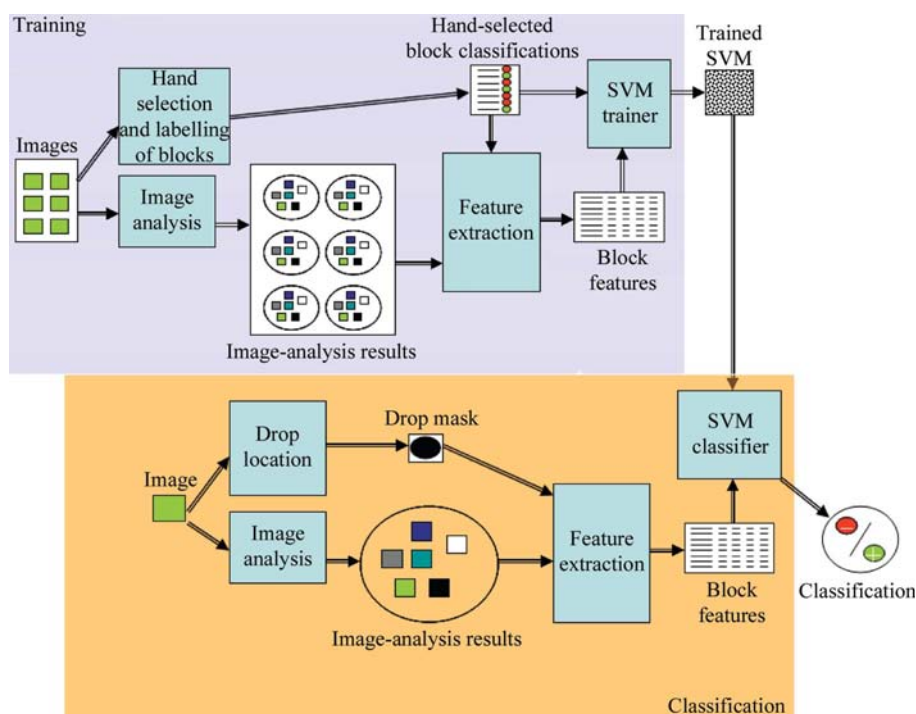


Figure 8
Training/classification pipeline.

approximately 30 s to process an image on a Pentium IV machine. Finally, the most intriguing direction for future work to take is to design a system that helps recognize the trends in

crystallization conditions, basing its decisions on the classification results, so that protein crystallization experiments could be performed much more efficiently.




Figure 9
Classification examples. A green block indicates that block is classified positive. (a) A crystal is correctly classified; (b) multiple crystals are correctly classified; (c) precipitates are correctly classified negative.




Figure 10
Some mistaken classifications, shown as details of larger images. (a) Obscure crystal (at 9 o'clock) classified as non-crystal. (b) Precipitate and crystal mixture classified as non-crystal. (c) Bubble classified as crystal.




Figure 11
More mistaken classifications, shown as details of larger images. (a) Light impurity classified as crystal. (b) Precipitate at the edge of the bubble classified as crystal. (c) Edgy precipitate classified as crystal.

The authors thank Larry DeSoto, Stewart Turley, Mark Robien, Oleksandr Kalyuzhniy, Mark Holl, Shawn McGuire and members of SGPP for useful assistance and data. This work was funded in part by the National Institutes of Health NIGMS 'High-Throughput Capillary-Based Protein Crystallography' (grant No. NIH 5 R01 GM 068878-02).

References

- Bern, M., Goldberg, D., Stevens, R. C. & Kuhn, P. (2004). *J. Appl. Cryst.* **37**, 279–287.
- Carson, C., Belongie, S., Greenspan, H. & Malik, J. (2002). *IEEE Trans. Pattern Anal. Machine Intell.* **24**, 1026–1038.
- Cristóbal, G. & Navarro, R. (1994). *Pattern Recog. Lett.* **15**, 273–277.
- Cumbaa, C. A., Lauricella, A., Fehrman, N., Veatch, C., Collins, R., Luft, J., DeTitta, G. & Jurisica, I. (2003). *Acta Cryst.* **D59**, 1619–1627.
- Daugman, J. G. (1983). *IEEE Trans. Syst. Man Cybern.* **13**, 882–888.
- Donoho, D. (1995). *IEEE Trans. Inform. Theor.* **41**, 613–627.
- Gabor, D. (1946). *J. Inst. Electr. Eng.* **93**, 429–457.
- Guillemont, J. *et al.* (2005). *J. Med. Chem.* **48**, 2072–2079.
- Hol, W. G. J. (1986). *Angew. Chem. Int. Ed.* **25**, 767–778.
- Jurisica, I., Rogers, P., Glasgow, J. I., Fortier, S., Luft, J. R., Wolfley, J. R., Bianca, M. A., Weeks, D. R. & DeTitta, G. T. (2001). *IBM Syst. J.* **40**, 394–409.
- Mehrotra, R., Namaduri, K. R. & Ranganathan, N. (1992). *Pattern Recog.* **25**, 1479–1494.
- Penas, M. (2005). *Perceptual Grouping Techniques for the Detection of Objects in Digital Images*. Department of Computer Science, University of A Coruña, Spain.
- Porat, M. & Zeevi, Y. Y. (1988). *IEEE Trans. Pattern Anal. Machine Intell.* **10**, 452–468.
- Shapiro, L. G. & Stockman, G. C. (2001). *Computer Vision*. Upper Saddle River, NJ, USA: Prentice Hall.
- Spraggon, G., Kreusch, A., Lesley, S. A. & Priestle, J. P. (2002). *Acta Cryst.* **D58**, 1915–1923.
- Turner, M. (1986). *Biol. Cybern.* **55**, 71–82.
- Vajo, Z., Fawcett, J. & Duckworth, W. C. (2001). *Endocr. Rev.* **22**, 706–717.
- Van Deemter, J. H. & Du Buf, J. M. H. (2000). *Int. J. Pattern Recog. Artif. Intell.* **14**, 757–777.
- Veerapandian, P. (1997). *Structure-Based Drug Design*. New York: Marcel Dekker.
- Wang, Z. & Jenkin, M. R. M. (1992). *Advances in Machine Vision: Strategies and Applications*, edited by C. Archibald & E. Petriu, pp. 151–170. Singapore: World Scientific.
- Weisstein, E. (1999–2005). *MathWorld*. <http://mathworld.wolfram.com>.
- Wilson, J. (2002). *Acta Cryst.* **D58**, 1907–1914.
- Zhu, X., Sun, S., Cheng, S. E. & Bern, M. (2004). In *Proceedings of the 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS-04)*. Piscataway, NJ, USA: IEEE.
- Zuk, W. M. & Ward, K. B. (1991). *J. Cryst. Growth*, **110**, 148–155.