

Materials Discovery: Fine-Grained Classification of X-ray Scattering Images

M. Hadi Kiapour Kevin Yager[†] Alexander C. Berg Tamara L. Berg

University of North Carolina at Chapel Hill, NC, USA [†]Brookhaven National Lab, NY, USA

{hadi, aberg, tlberg}@cs.unc.edu

kyager@bnl.gov

Abstract

We explore the use of computer vision methods for organizing, searching, and classifying x-ray scattering images. X-ray scattering is a technique that shines an intense beam of x-rays through a sample of interest. By recording the intensity of x-ray deflection as a function of angle, scientists can measure the structure of materials at the molecular and nano-scale. Current and planned synchrotron instruments are producing x-ray scattering data at an unprecedented rate, making the design of automatic analysis techniques crucial for future research. In this paper, we devise an attribute-based approach to recognition in x-ray scattering images and demonstrate applications to image annotation and retrieval.

1. Introduction

X-ray scattering is a powerful technique for probing the physical structure of materials at the molecular and nano-scale. The technique consists of shining an intense and collimated x-ray beam through a sample of interest, and recording the intensity of x-ray deflection as a function of angle. This redistribution of the intensity of the primary x-ray beam arises from the interference from the myriad of microscopic scattering and diffraction events which occur when the incident x-ray wave interacts with the atoms, molecules, and nano-scale interfaces in the sample [6, 23].

Premier x-ray scattering and diffraction instruments are now housed at “synchrotrons” [21, 1]: cyclic electron accelerators in which strong magnetic fields are used to induce the emission of x-ray beams from the circulating electron beam. Subsequent x-ray optics are used to generate monochromatic, highly collimated and focused beams suitable for demanding x-ray experiments. With the construction of next-generation synchrotrons, the materials science community has access to x-ray beam flux and brightness previously unimaginable, allowing for extremely rapid (millisecond) data collection.

Modern synchrotrons are thus beginning to generate massive quantities of data; far in excess of what human




Figure 1. Images from X-ray Material Dataset. Top row shows small-angle x-ray scattering (SAXS). Bottom row shows wide-angle x-ray scattering (WAXS) imagery.

experimenters can manually interpret and analyze. Simultaneously, modern instruments are being specifically built to operate in high-throughput modes. Many physical systems of interest have large and complicated phase spaces, which require concomitantly broad exploration. Modern materials are frequently blends and composites, exhibiting heterogeneous and hierarchical structuring; the exploration of these vast multi-component parameter spaces remains an outstanding challenge. Modern x-ray detectors include on the order of 1 million 20-bit pixels, and can acquire at 20-200 Hz. Next-generation beamlines are anticipated to generate 50,000 to 1,000,000 images/day (1-4 TB/day), with a synchrotron facility as a whole thereby generating millions of images per day (~ 100 TB/day).

It is thus crucial to automate as much of the x-ray scattering workflow as possible. Whereas bespoke analysis routines can be devised for restricted subsets of experimental results, a bottleneck remains if a human is required to categorize, cluster, and select images.

Currently, the standard workflow in an x-ray scattering experiment consists of an experimental team traveling to a synchrotron beamline, capturing a detailed dataset over several days, and then returning to their home institution with the images for later analysis. Automated analysis can improve upon this model in crucial ways: providing the experimental team with immediate feedback on their measurement results, and enabling facile sharing of data and insights between different research teams. Such tools are at present not available.

To the best of our knowledge, this paper presents the first computer vision based system that takes steps toward automatic, large-scale assessment of x-ray scattering images and toward building general purpose tools for organizing and sharing data between x-ray researchers. Some previous work [25] uses unsupervised clustering of a dataset with a limited number of categories to sort diffraction snapshots into groups; that work does not exploit training on labeled data and also does not exploit compact feature representations, hence uses a cropped version of only the center of an intensity image to make the clustering computationally feasible. In contrast, we consider the application of modern techniques in computer vision to the problem of recognizing appearance characteristics in x-ray scattering images. We take advantage of recent advances in attribute based recognition, e.g. [11, 17, 18, 10, 12, 5, 9, 8, 7] for retrieval, organization, and classification. Attributes can be viewed as an interpretable mid-level representation for recognition tasks. They are particularly useful for tasks that involve end users such as image search [9, 8] where it may be natural for people to specify initial search desiderata or refinements using attributes. Attributes are also composable [11] and can divide the space of images effectively, allowing efficient search. If desired, attributes can be automatically personalized to a particular user [7].

In this paper we begin by demonstrating that a careful selection of computer vision techniques makes it possible to classify x-ray scattering images according to a number of attributes that are relevant for domain experts. This is a challenge as classification schemes have been optimized for realspace imagery (e.g. photographs), whereas x-ray scattering images are more abstract (essentially raw data). Physically, x-ray images are produced by scattering into the far-field, as opposed to focusing onto an image plane, as in typical optical imaging. Thus, scattering images do not share most of the image properties that we take for granted in other computer vision experiments – there is no notion of objects, position has meaning, there is no spectral information, and more.

We go on to consider and analyze scenarios in which this approach could help material scientists using x-ray scattering and diffraction techniques. The *first scenario* takes place during data collection, as thousands of images are collected over hundreds of samples, while varying measurement geometry, as well as physical parameters (temperature, pressure, etc.). We show very high accuracy for predicting attributes that indicate problems with the complex measurement apparatus (e.g. beam blocked, detector saturated) – bearing in mind that the material scientists who use x-ray instruments are not usually experts on the complex and constantly changing beamline equipment. This feedback would allow experimental settings to be adjusted immediately, avoiding costly wastes in measurement time

and effort. The *second scenario* is automatically tagging images with attributes. This would allow quick indexing of potentially relevant attributes or tags that had been specified by other scientists, something that is only possible through expert knowledge today. The *third scenario* is automatically identifying similar x-ray scattering images, possibly gated by specific attributes. For instance, a researcher might want to find previous measurements on similar materials, or identify different materials with similar scattering features. Attributes allow searching for similar samples while specifying what aspects of the images are more important for search.

1.1. X-ray scattering

We use the generic term *x-ray scattering* to describe a variety of related techniques: those which measure diffuse scattering (e.g. from roughness), those which measure diffraction peaks (e.g. from crystal unit cells), those operating in reflection (‘grazing-incidence’) mode, etc. We also consider both small-angle x-ray scattering (SAXS) and wide-angle x-ray scattering (WAXS) results, which probe nano-scale and molecular/atom order respectively.

Modern x-ray scattering instruments typically use two-dimensional x-ray detectors to simultaneously capture a wide angular range of the scattering signal. The resultant two-dimensional arrays of x-ray intensity values are, conceptually, a slice through a three-dimensional “reciprocal-space” whose intensity distribution non-trivially encodes the sample’s microscopic structure. The reciprocal-space is essentially the Fourier-transform of the realspace electron-density distribution in the sample; however x-ray experiments only record the amplitude (and not the phase) of the scattered radiation, making direct data inversion impossible. Instead, x-ray data must be carefully analyzed and iteratively fit to candidate models, in order to extract parameters of physical interest (atom or molecular spacing, crystallographic unit cell, nano-scale shape and periodicity, grain size, orientation distribution, roughness, etc.) [20, 13, 19, 24]. X-ray scattering data from two-dimensional detectors can be thought of as images, which contain a variety of features (spots, rings, halos, diffuse intensity) which result from different kinds of microscopic order. Indeed these datasets are frequently visualized using false-color images, such as those shown in Figure 1. X-ray measurement experts can oftentimes interpret these image features ahead of detailed analysis (e.g. sharp peaks indicate crystalline structures, diffuse halos indicate amorphous order). However, such manual categorization is quickly becoming a bottleneck.

Ultimately, the entire process of measurement, analysis, and exploration is meant to be automated; with the x-ray instrument itself automatically selecting subsequent sample formulations for study. This automated materials discov-

ery will require revolutionary progress in data management such as those described here.

The rest of the paper is organized as follows: In section 2 we describe the dataset and how it is tagged and organized. In section 3, we evaluate attribute classification for x-ray scattering data. In section 4, we demonstrate two applications desired by material science researchers.

2. X-ray Materials Discovery Dataset (XMD)

Our dataset contains images from 13 x-ray scattering measurement runs – a set of related x-ray scattering images collected for closely-related material samples, continuously captured over a short time period. The data were all collected on the X9 beamline at the National Synchrotron Light Source, an endstation capable of both small-angle x-ray scattering (SAXS) and wide-angle x-ray scattering (WAXS); in either transmission (TSAXS/TWAXS) or grazing-incidence (GISAXS/GIWAXS) geometry. The measurement runs we analyze include images acquired in all of these measurement modes. The dataset includes a wide range of different kinds of samples: nano-particles in solution, lithographic gratings, self-assembling polymer films, organic semiconductors, etc. The substantial variability in the samples and measurement modes thus provides a realistic sampling of x-ray scattering data. Example x-ray scattering images are shown in Figure 1. The number of images in each experiment varies between 54 and 618, giving a total of 2832 single-channel images with intensities in the range $[0, 2^{16}]$.

All of the images have been labeled with 98 binary attributes by a domain expert. These attributes represent a diverse set of characteristics ranging from the type of measurement, to appearance based scattering features, to chemical composition and physical properties of the materials. Images are labeled with an average of 11.7 attributes. The attribute vocabulary is organized into groupings by attributes of: **G1** experiments, **G2** instrumentation, **G3** imaging, **G4** scattering features, **G5** samples, **G6** materials, **G7** specific substances. A complete list of attributes along with their group labels is shown in Table 1.

Figure 2 shows the co-occurrence matrix for the 25 most frequent attributes in the dataset. Note that the co-occurrence matrix tends to become very sparse as we increase the number of attributes, while some pairs of attributes have significantly higher co-occurrence rates. For example, *Specular rod* is more likely to co-occur with attribute *Ring: Textured* than with *Ring: Isotropic*. This can be expected on physical grounds: a specular rod only appears during grazing-incidence (reflection-mode) experiments; the thin films and interfaces thereby probed have a higher tendency to be anisotropic or textured. In addition, some attributes are directly dependent on the x-ray scattering technique used in measurement (e.g. SAXS or WAXS)

Attribute	#	Attribute	#
Thin film [G5]	1646	Silicon [G7]	130
Specular rod [G3]	1597	GTSAXS [G1]	127
Beam off image[G2]	1591	MWCNT [G7]	125
Photonics CCD[G2]	1591	Nanoporous [G5]	125
Ordered [G5]	1462	Theta sweep [G1]	109
GIWAXS [G1]	1439	PDMS [G7]	107
MarCCD [G2]	1241	Saturation artifacts [G3]	97
Horizon[G4]	1171	Peaks: Line z [G4]	90
Linear beamstop [G2]	1156	Circular beamstop [G2]	85
Peaks: Isolated[G4]	1099	Peaks: Line xy [G4]	79
GISAXS[G1]	870	Diffuse low-q: Anisotropic [G4]	78
Ring: Oriented z [G4]	856	Many rings [G4]	78
Polymer [G6]	821	Diffuse low-q: Oriented z [G4]	76
Halo: Isotropic [G4]	791	Misaligned [G3]	76
Ring: Isotropic [G4]	604	Beam streaking [G3]	70
Ring: Textured [G4]	528	Diffuse low-q: Oriented xy [G4]	69
Higher orders: 2 to 3 [G4]	513	Blocked [G3]	62
P3HT [G7]	505	Diffuse specular rod [G4]	62
Ring: Oriented xy [G4]	491	Smearred horizon [G4]	55
SiO2 [G7]	467	Symmetry ring: 4-fold [G4]	55
Vertical streaks [G4]	434	Higher orders: 10 to 20 [G4]	53
Single crystal [G5]	430	Ring doubling [G4]	53
Block-copolymer [G6]	416	Halo: Anisotropic [G4]	46
Peaks: Many/field [G4]	396	Powder [G5]	44
Grating [G5]	375	Specular rod peaks [G4]	41
PCBM [G7]	369	AgBH [G7]	40
Diffuse high-q: Isotropic [G4]	357	Ring: Oriented other [G4]	33
Higher orders: 4 to 6 [G4]	351	Peaks: Line [G4]	23
Weak scattering [G3]	318	Diffuse high-q: Oriented z [G4]	20
Rubrene [G7]	266	Bad beam shape [G3]	19
TSAXS [G1]	264	LaB6 [G7]	16
Higher orders: 7 to 10 [G4]	260	Phi sweep [G1]	16
2D detector obstruction [G3]	224	Peak doubling [G4]	15
Bragg rods [G4]	211	Halo: Oriented xy [G4]	14
Ring: Anisotropic [G4]	205	Polycrystalline [G5]	14
Peaks: Along ring [G4]	201	Diffuse high-q: Oriented xy [G4]	11
Amorphous [G5]	197	Direct [G3]	11
Saturation [G2]	193	Object obstruction [G3]	9
PS-PMMA [G7]	190	Peaks: Line other [G4]	9
Composite [G5]	179	Waveguide streaks [G4]	8
Diffuse low-q: Isotropic [G4]	170	Higher orders: 20 or more [G4]	4
Yoneda [G4]	167	Substrate streaks/Kikuchi [G4]	4
Strong scattering [G3]	159	Diffuse low-q: Oriented other [G4]	3
TWAXS [G1]	152	Halo: Spotted [G4]	3
Halo: Oriented z [G4]	148	Diffuse low-q: Spotted [G4]	2
High background [G4]	142	Diffuse high-q: Spotted [G4]	1
Asymmetric (left/right) [G2]	138	Empty cell [G3]	1
Ring: Spotted [G4]	136	Parasitic slit scattering [G3]	1
Superlattice [G6]	136	Point detector obstruction [G3]	1

Table 1. Fine-grained attributes ordered by number of samples.

High-level attributes	Fine-grained attributes
Diffuse high-q	Diffuse high-q: Isotropic, Oriented xy or z, Spotted
Diffuse low-q	Diffuse low-q: (An)Isotropic, Oriented xy or z or other, Spotted
Halo	Halo: (An)Isotropic, Oriented xy or z, Spotted
Higher orders	Higher orders: 2 to 3, 4 to 6, 7 to 10, 10 to 20, 20 or more
Peaks	Peaks: Along ring, Isolated, Line xy or z or other, Many/field
Ring	(An)Isotropic, Oriented xy or z or other, Spotted, Textures

Table 2. Mapping between high-level and fine-grained attributes

which is directly reflected in the value for attribute *MarCCD* (true for SAXS samples and false for WAXS ones). This is particular to the dataset we have chosen to analyze: on this beamline, a MarCCD detector is used for the small-angle measurement, whereas a Photonic Sciences detector is used for the wide-angle measurements. These co-occurrences, however, point to generic kinds of correlations which can be exploited in the automated analysis of data coming from particular beamlines. For instance, while in our dataset *Beam off image* (detector offset with respect to

1	00	1049	794	794	1154	783	852	688	852	638	720	369	626	414	238	140	350	505	406	467	420	00	402	98	241	
2	049	00	1245	1245	934	1245	352	1163	352	1031	346	780	590	665	381	437	383	493	440	460	288	418	310	297	12	
3	794	1245	00	1591	656	1439	00	1171	00	1091	00	834	707	789	487	528	415	505	480	467	00	430	00	299	00	
4	794	1245	1591	00	656	1439	00	1171	00	1091	00	834	707	789	487	528	415	505	480	467	00	430	00	299	00	
5	1154	934	656	656	00	634	806	614	721	539	561	270	593	399	296	143	270	445	351	464	313	00	382	97	264	
6	783	1245	1439	1439	634	00	00	1171	00	1091	00	824	599	707	450	507	412	505	450	467	00	430	00	299	00	
7	852	352	00	00	806	00	00	00	1156	08	870	22	114	02	117	00	98	00	11	00	434	00	416	97	375	
8	688	1163	1171	1171	614	1171	00	00	969	00	00	722	588	637	377	424	360	488	441	458	00	399	00	279	00	
9	852	352	00	00	721	00	1156	00	00	08	870	22	114	02	78	00	98	00	11	00	434	00	416	97	375	
10	638	1031	1091	1091	539	1091	08	969	08	00	08	679	512	588	310	418	337	423	342	413	00	414	00	291	05	
11	720	346	00	00	561	00	870	00	00	22	27	02	28	00	80	00	11	00	434	00	416	30	261	00	00	
12	369	780	834	834	270	624	22	722	22	679	22	00	220	405	293	378	328	270	261	119	00	386	00	261	00	
13	626	590	707	707	593	599	114	588	114	512	27	220	00	454	181	134	213	413	355	466	17	00	27	00	00	
14	414	665	789	789	399	707	02	637	02	588	02	405	454	00	244	253	278	352	345	291	00	200	00	84	00	
15	238	381	487	487	296	450	117	377	78	310	28	293	181	244	00	198	144	156	142	91	08	128	10	143	00	
16	140	437	528	528	143	507	00	424	00	418	00	378	134	253	198	00	186	81	46	129	00	313	00	169	00	
17	350	383	415	415	270	412	98	360	98	337	80	328	213	278	144	186	00	257	239	117	53	98	12	01	37	
18	505	493	505	505	445	505	00	488	00	423	00	270	413	352	156	81	257	00	381	280	00	00	00	00	00	
19	406	440	480	480	351	450	11	441	11	342	11	261	355	345	142	46	239	381	00	207	00	10	00	05	00	
20	467	460	467	467	464	467	00	458	00	413	00	119	466	291	91	129	117	280	207	00	00	00	00	00	00	
21	420	288	00	00	313	00	434	00	434	00	434	00	17	00	08	00	53	00	00	00	00	00	330	00	51	
22	00	418	430	430	00	430	00	399	00	414	00	386	00	200	128	313	98	00	10	00	00	00	00	00	270	00
23	402	310	00	00	382	00	416	00	416	00	416	00	27	00	10	00	12	00	00	00	00	00	330	00	00	00
24	98	297	299	299	97	299	97	279	97	291	30	261	00	84	143	169	01	00	05	00	00	00	270	00	96	
25	241	12	00	00	264	00	375	05	261	00	00	00	00	00	00	00	00	00	00	00	00	51	00	00	96	00
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25

Figure 2. Co-occurrence between attributes. Each entry is the count of samples which has the corresponding pair of attributes. There are 2832 samples in total. To increase the contrast of the off-diagonal elements, the diagonal elements are set to zero. Indices correspond to order of attributes shown in Table 1

direct x-ray beam) never co-occurs with *MarCCD*, *Linear beamstop* always co-occurs with *MarCCD*. The correlations inherent in the tags themselves can thus be exploited in order to hierarchically split datasets. For instance, an initial analysis can classify images as SAXS or WAXS, making tag prediction within those classes simpler and more robust. This procedure can, in principle, be repeated iteratively to a depth limited by the inherent hierarchical correlations of the tag taxonomy.

3. Experiments and Results

In this section, we train and evaluate classifiers to predict each x-ray scattering attribute. In particular, we explore a number of different standard computer vision features with one vs. all linear support vector machines (SVMs) to produce a benchmark on the x-ray materials discovery dataset.

We select several features that are widely used for image classification in computer vision: Downsampled (Tiny Images) [22], GIST [16], Histogram of Oriented Gradients (HOG) [3], Pyramid of HOG (PHOG) [2], Local Binary Patterns (LBP) [15], dense grayscale SIFT [14], and concatenation of PHOG with LBP. Downsampled images are obtained by reducing the image dimensions to 32×32 . GIST, an abstract representation of statistics of oriented structures within the image, is computed in 8 orientations on 4×4 grid at 4 scales from coarse to fine. HOG features, which can encode local shape information from regions within an image as a histogram of gradients, are extracted on 8×8 grid with 31 orientations. PHOG is com-

Attribute	ds	hog	phog	gist	lbp	lbpphog	sift
Thin film	86.8	89.1	90.2	87.8	86.5	89.6	92.1
Specular rod	79.9	80.1	78.7	78.5	81.1	80.2	74.1
Beam off image	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Photonics CCD	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Ordered	62.8	76.3	77.1	82.1	72.5	78.9	72.6
GIWAXS	99.3	99.8	99.8	99.3	99.8	99.8	98.4
MarCCD	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Horizon	85.9	95.3	95.8	95.0	90.1	94.1	92.5
Lin. beamstop	99.3	96.5	95.8	96.1	99.6	96.8	96.1
Peaks: Iso.	57.4	61.0	61.2	63.0	59.2	61.4	58.4
GISAXS	82.3	92.5	93.5	97.5	95.2	94.7	96.3
Ring: Ori. z	58.7	67.7	67.0	58.2	60.5	71.3	74.3
Polymer	64.6	65.6	65.5	66.4	68.8	66.3	64.5
Halo: Iso.	59.4	60.3	55.9	53.3	76.9	69.4	63.9
Ring: Iso.	46.9	79.1	79.0	79.4	74.1	79.4	71.4
Ring: Text.	46.7	48.5	50.7	45.9	40.3	48.5	43.7
High. ord. 2-3	27.7	34.5	29.7	27.1	29.0	35.8	29.3
P3HT	76.4	83.1	84.0	82.8	82.5	82.4	83.1
Ring: Ori. xy	39.9	50.7	49.6	50.6	52.3	46.3	45.6
SiO2	99.8	99.8	99.8	99.9	100.0	99.8	99.9
Vertical streaks	46.7	68.8	67.9	66.1	73.4	69.2	64.2
Single crystal	74.9	88.2	88.1	90.1	88.3	90.0	81.7
Block-copoly.	55.3	93.2	92.0	95.0	95.4	93.6	84.2
Peaks: Many	44.7	49.1	51.9	67.4	49.9	54.3	53.0
Grating	82.1	87.9	89.3	88.3	87.4	87.6	84.5
PCBM	40.5	55.6	56.4	46.9	52.5	57.2	42.9
Diffuse hq: Iso.	39.1	57.9	54.5	49.3	76.0	59.8	55.9
High. ord. 4-6	28.1	30.9	35.5	35.8	31.4	30.7	37.2
Weak scatter.	44.7	37.9	38.2	46.6	42.1	38.2	25.2
Rubrene	52.1	71.2	71.0	75.7	59.4	67.0	53.9
All 98 Attr.	46.1	53.3	54.2	52.0	54.0	55.5	51.9

Table 3. Average Precision for classification of attributes

Feature	No hierarchy	SAXS-WAXS
downsampled	42.3	46.1
hog	49.5	53.3
phog	49.3	54.2
gist	47.2	52.0
lbp	50.2	54.0
lbpphog	51.5	55.5
sift	48.2	51.9

Table 4. Mean Average Precision (mAP) for all 98 attribute classifiers. Right column shows results when the system first tries to separate SAXS and WAXS imagery before identifying detailed attributes.

puted with 31 orientations in 4 layers from coarse to fine. LBP summarizes the local structure in an image by comparing each pixel with its neighborhood which can capture very fine-grained details in the image. For LBP, the uniform LBP was extracted on a 8×8 grid. SIFT descriptors are invariant to scaling transformations and robust to moderate perspective and illumination variations. For SIFT features, we use a bag-of-words model, with a visual vocabulary produced by k-means clustering (k=1000). Prior to computing features, all images are resized to 512×512 and a median filter is applied to remove the extreme pixel values produced by errors in the image acquisition process (a standard practice in x-ray scattering research).

We utilize support vector machines SVMs [4] to learn visual models for attributes on these computer vision features. For each attribute we train a one vs. all linear SVM classifier with parameter C optimized to maximize the AUC mea-




Figure 3. Sample precision-recall curves for three attributes. Left shows attribute *Linear beamstop* with mAP 96.8%. Middle shows attribute *Strong Scattering* with mAP 77.2%, but still allows high precision (100%) at low recall (14%). Finally prediction for *Peaks: Along ring* with mAP 45.7% is generally poor.

sure (area under precision-recall curve) using cross validation. Cross validation must be performed carefully in order to not bias the models unfairly. A given measurement run will inevitably include replicate measurements of the same sample, as well as measurements of highly similar samples. Since the images within a measurement run can be much more similar than across different runs, “k-fold” cross validation will end up choosing unfair validation sets since it may partition images from one measurement run into both train and validation sets. Therefore, instead of randomly partitioning the samples, we cross validate using as many folds as number of runs there are in the training set. In each fold, we train the model on all runs except one and use it to predict on the excluded one. We repeat this cross validation process for all runs in training and choose the parameter that results in the highest overall area under the curve. We note that our validation scheme makes the ultimate classification inherently more challenging; our results are thus inherently conservative. However, our scheme provides the best approximation for how these methods would be used in practice, where data from a new measurement run would be analyzed using classifiers trained on the complete set of past measurement runs.

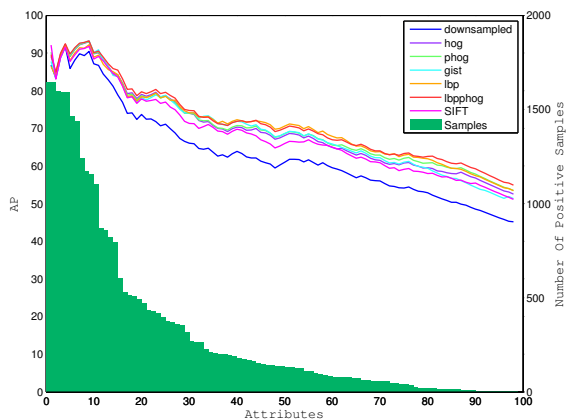


Figure 4. Mean average precision for attribute prediction. Attributes are ordered by the number of positive samples for each.

We also evaluate a combined classifier, training an SVM on the outputs of the individual feature SVMs, but found that it did not increase performance over selecting the best performing feature for each attribute.

We evaluate our classifiers on the x-ray material discovery dataset, containing 13 measurement runs. Similarly as for the cross validation, random sampling of train/test images across runs is not advised. Care must be taken to produce evaluations that are meaningful with regards to generalization capability for future experiments. We take a leave one out approach to evaluation, where we train our SVM classifiers on all measurement runs except for one, and evaluate predictions of attributes on the excluded run. We repeat this process for all 13 measurement runs, and average across runs to evaluate performance.

Another challenge of the x-ray scattering data is that some attributes may be prevalent in many samples, e.g. SAXS vs. WAXS whereas some attributes may be quite sparse, e.g. specific material attributes. Therefore the number of positive samples for an attribute can be quite limited. Finally, the attributes present in a measurement run may vary within the run, e.g. some attributes for a particular sample may appear in some measurements but not in others (e.g. as the angle of the x-ray beam with respect to the sample is varied). A useful prediction tool needs to effectively handle all of these situations.

There are two main x-ray scattering techniques: small-angle and wide-angle (as described earlier). Given the different nature of images obtained by small-angle (SAXS) and wide-angle (WAXS) techniques, we build a two-level classification system, where in the top level we train a classifier that can separate SAXS from WAXS and the bottom level contains classifiers for fine-grained attributes corresponding to each technique. In this manner we establish consistent training sets for fine-grained classifiers. This two-stage classification also improves performance with respect to flat classification schemes, especially when the set of positive samples is very small.

This is an unbalanced classification problem, and classification accuracy becomes a less useful evaluation metric because performance on the large number of negative exam-

ples outweighs performance on the few positive examples. Therefore, we use average precision (AP) for evaluation, which can be more appropriate for unbalanced data.

Results for the 30 most frequent attributes are reported in Table 3. Examples of precision-recall curves for tags are shown in Figure 3. In Table 4 we also compare the overall performance for all features in two scenarios: 1) using two level hierarchical classification system described above and 2) using flat classification without a hierarchy. The results confirm that pre-filtering SAXS and WAXS data before fine-grained classification improves performance for all visual feature types. On average LBPPHOG descriptor performs best on attribute classification. Hence, we use it as the main descriptor for applications introduced in section 4. Due to the imbalanced distribution of samples for each attribute over the experiments, we also show overall performance of our classifiers across attributes sorted by number of positive samples in descending order. Figure 4 shows that performance is strongly aligned with the number of available positive samples in training. Therefore, our performance results should greatly improve if our methods are applied to larger datasets.

Another classification task of great utility to domain experts, is to reliably predict the coarse high-level characteristics present in images, e.g. *Ring* vs *Ring:Textured* or *Ring:Isotropic*. From these characteristics, domain experts can then infer much about the microstructure of the sample. Therefore, we also build higher level attributes by merging fine-grained attributes into the following groups: *Diffuse high-q*, *Diffuse low-q*, *Halo*, *Higher Orders*, *Peaks* and *Ring*. The mapping between fine-grained attributes and high-level attributes is presented in Table 2. There is one important factor for training high-level attributes since they are no longer specific to SAXS or WAXS data and both techniques can have produced the same high-level attributes in the resulting images. Therefore, for high-level attributes we train our classifiers on both SAXS and WAXS, leading to larger positive training sets. Average precisions for classification of high-level attributes is reported in Table 5. Overall, results are best for LBP descriptor and LBPPHOG descriptor obtains a slightly lower performance. Note the significantly higher performance on high-level attributes which is due to more available samples in training and the fact that it is a somewhat easier task compared to fine-grained classification.

4. Applications

We apply our recognition systems to two applications that would be of great use to domain experts in the management of realistic x-ray scattering datasets. The first application is automatic image annotation (Section 4.1) where the goal is to tag images with their appropriate attributes. The second application is image retrieval (Section 4.2) where

Attribute	ds	hog	phog	gist	lbp	lbpphog	sift
Diffuse high-q	31.2	72.7	63.6	58.3	74.2	65.6	52.9
Diffuse low-q	64.7	71.9	61.2	65.0	71.6	70.3	58.4
Halo	70.4	73.6	73.6	73.4	80.6	75.0	74.5
Higher orders	62.4	70.1	69.4	76.4	74.7	74.9	72.6
Peaks	73.3	77.7	79.0	83.9	76.6	80.0	78.5
Ring	70.6	96.1	95.5	94.9	95.4	96.1	92.2
Total Average	62.1	77.0	73.7	75.3	78.8	77.0	71.5

Table 5. Average Precision for classification of high-level attributes

we want to retrieve similar images given a query image, or given a set of query attribute tags.

4.1. Automatic Image Annotation

Given an input image, we would like to automatically tag the image with its relevant tags. Figure 5 shows automatically predicted attribute tags for example SAXS and WAXS images. Users can submit new images and the system will select the best performing classifiers (best performing visual feature classifier for each tag) to predict attributes automatically.

This will be a useful application for scattering systems to generate annotations in real-time. Generally, experiments are performed and then stored without any relevant metadata for later look-up. This application would be useful for scientists conducting experiments because it could quickly and automatically annotate all of the scattering data produced during an experiment. Furthermore, such tags could be used to organize all of the x-ray scattering experiments produced by researchers around the world, connecting the data in ways that is currently infeasible.

4.2. Attribute-based Image Retrieval

We also provide a retrieval based application for our system. This application implements a multi-modal attribute-based retrieval. Users can submit a query image and ask the system to return a ranked list of similar images. Alternatively, users can specify target attributes and retrieve relevant images. Example retrieval results using both image and text-based queries are shown in Figure 6.

The system is highly successful in identifying similar images; e.g. retrieving measurements of the same materials from different measurement runs. The addition of tag constraints provides a notably useful refinement; successfully selecting the subset of images that match. These kinds of retrievals are useful both for identifying whether the existing database of images includes measurements of a particular kind (e.g. whether a particular material has ever exhibited anisotropic scattering pattern), and can also be useful as an aid to researchers in managing large datasets (e.g. allowing them to retrieve an image which they can visually recall, but whose identity they cannot recall). By allowing different experimenters to share datasets, such a system provides

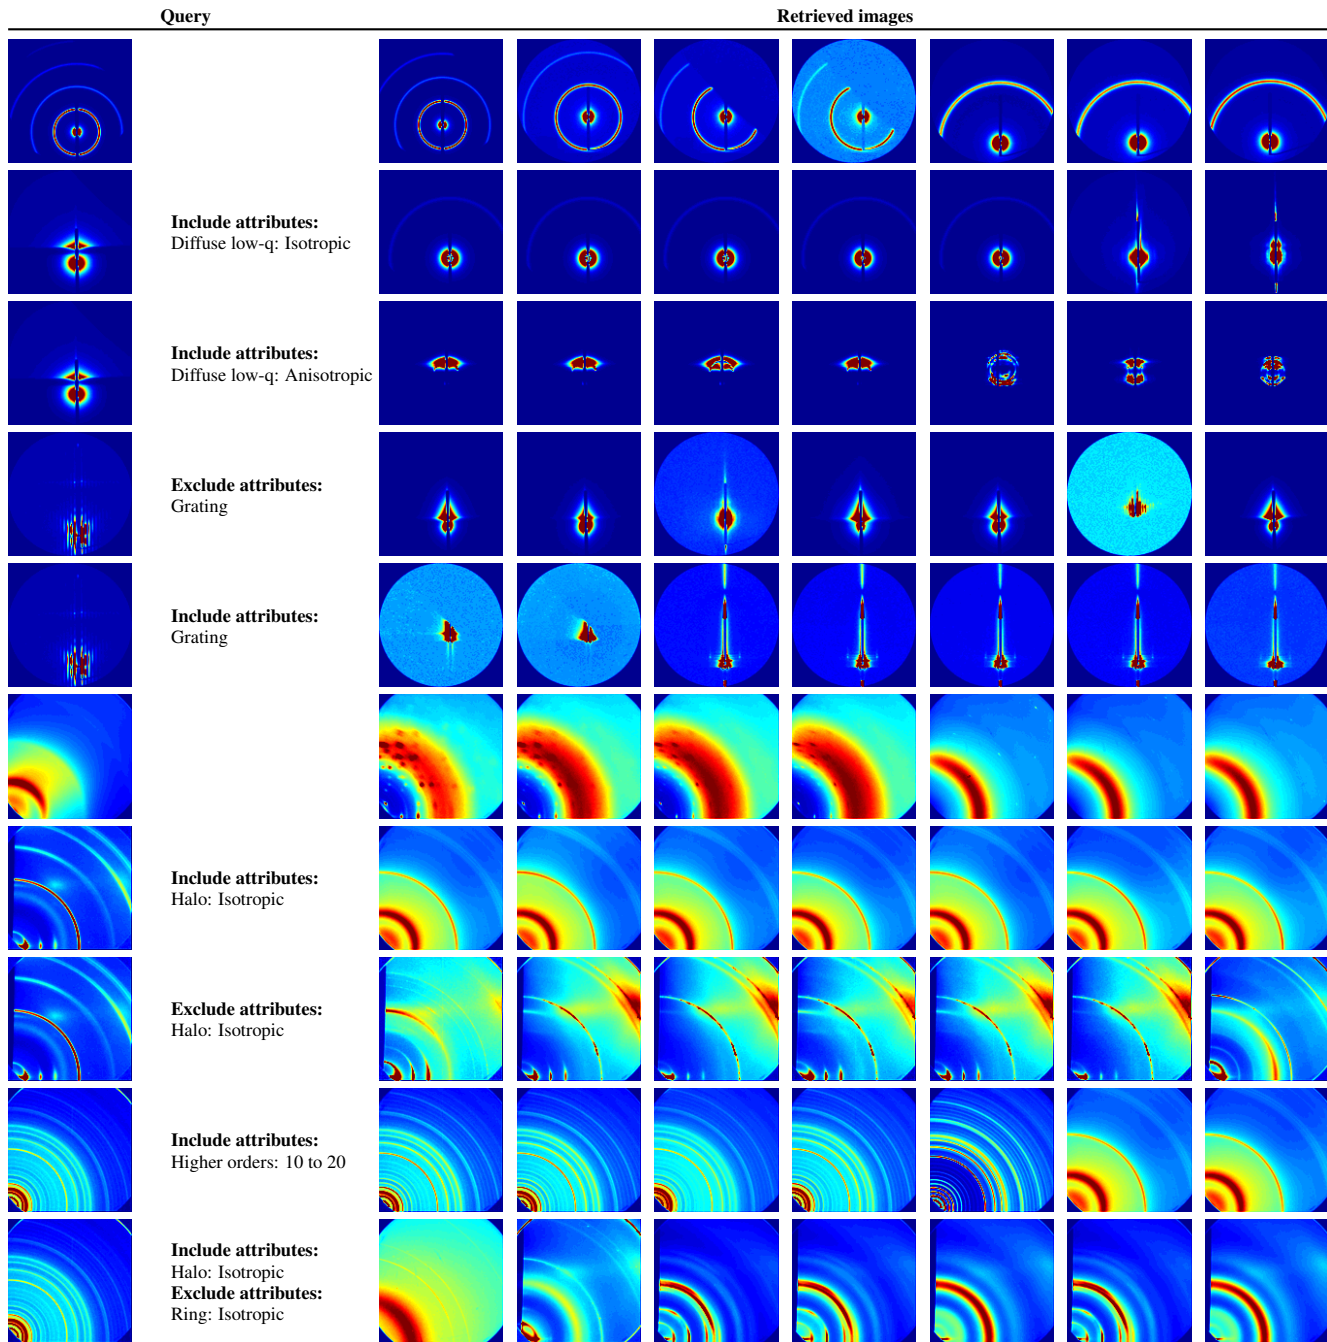


Figure 6. Retrieval examples. Left column shows query image. Second from left column shows optional text filters for retrieval. Other columns show retrieved results.

the opportunity to foster productive collaborations and maximal usage of existing data.

5. Conclusion

We have provided an automatic visual recognition based system for analyzing x-ray scattering data. Our system takes an attribute based approach to recognition. Evaluations show that our methods already perform well at attribute prediction across a number of experimental scenar-

ios, and that performance may increase in the future as more data is labeled. We also demonstrate application to attribute annotation and retrieval. We will release our tools along with the dataset to the general scientific community for use in worldwide efforts on x-ray scattering experiments.

6. Acknowledgment

This work was partially supported by SBU/BNL Seed Grant, The Data Sensorium: Multi-Modal Explorations of




Image	Predicted tags
(a)	Block-copolymer, GISAXS, Linear beamstop, MarCCD, Ordered, Specular rod, Thin Film, Vertical streaks, Weak scattering, Bragg rods, Silicon
(b)	AgBH, Diffuse low-q: Isotropic, Linear beamstop, MarCCD, Ordered, Powder Ring: Isotropic, Saturation, TSAXS
(c)	GISAXS, Grating, Higher orders: 4 to 6, Linear beamstop, MarCCD Peaks: Along ring, Thin film, Vertical streaks
(d)	Blocked, Linear beamstop, MarCCD, TSAXS
(e)	AgBH, Beam off image, Higher orders: 10 to 20, Ordered, Photonics CCD Ring: Isotropic, TWAXS
(f)	Beam off image, GIWAXS, Halo: Isotropic, Horizon, PCBM, Photonics CCD, Ring: Oriented xy, Specular rod, Thin film, Higher orders: 2 to 3
(g)	Beam off image, Halo: Isotropic, Photonics CCD, RingL Isotropic, TWAX, Ordered, Higher orders: 10 to 20, PCBM
(h)	Beam off image, GIWAXS, Horizon, Peaks: Isolated, Peaks: Many / field, Photonics CCD, Ring: Isotropic, Ring: Oriented z, Specular rod, Thin film, SiO ₂ , Saturation

Figure 5. Automatic image annotation. Most predictions are correct, mistakes are highlighted in red.

Scientific Data, has received funding from NSF Career Award 1054133 and NSF Award 1161876. Research was carried out in part at the Center for Functional Nanomaterials and the National Synchrotron Light Source, Brookhaven National Laboratory, which are supported by the U.S. Department of Energy, Office of Basic Energy Sciences, under Contract No. DE-AC02-98CH10886.

References

- [1] D. H. Bilderback, P. Elleaume, and E. Weckert. Review of third and next generation synchrotron light sources. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 38(9):S773, 2005. 1
- [2] A. Bosch, A. Zisserman, and X. Munoz. Image classification using random forests and ferns. *ICCV*, 2007. 4
- [3] N. Dalal and B. Triggs. Histogram of oriented gradients for human detection. *CVPR*, pages 886–893, 2008. 4
- [4] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008. 4
- [5] V. Ferrari and A. Zisserman. Learning visual attributes. In *Advances in Neural Information Processing Systems*, 2007. 2
- [6] A. Guinier. *X-ray Diffraction in Crystals, Imperfect Crystals, and Amorphous Bodies*. Mineola: Dover Publications, ISBN: 0486680118, 1994. 1
- [7] A. Kovashka and K. Grauman. Attribute adaptation for personalized image search. In *ICCV*, 2013. 2
- [8] A. Kovashka and K. Grauman. Attribute pivots for guiding relevance feedback in image search. In *ICCV*, 2013. 2
- [9] A. Kovashka, D. Parikh, and K. Grauman. Whittlesearch: Image search with relative attribute feedback. In *CVPR*, 2012. 2
- [10] N. Kumar, A. Berg, P. Belhumeur, and S. Nayar. Attribute and simile classifiers for face verification. In *ICCV*, 2009. 2
- [11] N. Kumar, A. C. Berg, P. N. Belhumeur, and S. K. Nayar. Describable visual attributes for face verification and image search. In *PAMI*, October 2011. 2
- [12] C. Lampert, H. Nickisch, and S. Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*, 2009. 2
- [13] R. Lazzari, F. Leroy, and G. Renaud. Grazing-incidence small-angle x-ray scattering from dense packing of islands on surfaces: Development of distorted wave born approximation and correlation between particle sizes and spacing. *Phys. Rev. B*, 76:Sep, 2007. 2
- [14] D. G. Lowe. Object recognition from local scale-invariant features. *ICCV*, 1999. 4
- [15] T. Ojala, M. Pietkainen, and M. Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *PAMI*, 2010. 4
- [16] A. Oliva and A. Torralba. Modeling the shape of the scene: a holistic representation of the spatial envelope. *IJCV*, 42(3):145–175, 2001. 4
- [17] D. Parikh and K. Grauman. Interactively building a discriminative vocabulary of nameable attributes. In *CVPR*, 2011. 2
- [18] G. Patterson and J. Hays. Sun attribute database: Discovering, annotating, and recognizing scene attributes. In *CVPR*, 2012. 2
- [19] G. Renaud, R. Lazzari, and F. Leroy. Probing surface and interface morphology with grazing incidence small angle x-ray scattering. *Surface Science Reports*, 64(8):255 – 380, 2009. 2
- [20] S. K. Sinha, E. B. Sirota, S. Garoff, and H. B. Stanley. X-ray and neutron scattering from rough surfaces. *Phys. Rev. B*, 38:2297–2311, Aug 1988. 2
- [21] J. V. Smith. Tutorial review. synchrotron x-ray sources: instrumental characteristics. new applications in microanalysis, tomography, absorption spectroscopy and diffraction. *Analyst*, 120:1231–1245, 1995. 1
- [22] A. Torralba, R. Fergus, and W. T. Freeman. 80 million tiny images: a large dataset for non-parametric object and scene recognition. *PAMI*, 30(11):1958–1970, 2008. 4
- [23] B. E. Warren. *X-ray Diffraction*. Mineola: Dover Publications, ISBN: 0486663175, 1990. 1
- [24] K. G. Yager, Y. Zhang, F. Lu, and O. Gang. Periodic lattices of arbitrary nano-objects: modeling and applications for self-assembled systems. *Journal of Applied Crystallography*, 47(1), Feb 2014. 2
- [25] C. H. Yoon, P. Schwander, et al. Unsupervised classification of single-particle x-ray diffraction snapshots by spectral clustering. *Opt. Express*, 19(17):16542–16549, Aug 2011. 2