

Recognizing Blasting Caps in X-Ray Images*

M.A. Maloof¹, Z. Duric^{1,2}, R.S. Michalski¹, A. Rosenfeld²

¹Machine Learning and Inference Laboratory, George Mason University
Fairfax, VA 22030-4444
<http://www.mli.gmu.edu/>

²Computer Vision Laboratory, University of Maryland
College Park, MD 20742-3275
<http://www.cfar.umd.edu/>

Abstract

This paper presents work in progress on an approach to the problem of recognizing blasting caps in x-ray images. An analysis of functional properties of blasting caps was used to design the representation space, which combines intensity and shape features. Recognition proceeds in two phases. The first phase is a bottom-up process in which low intensity blobs are used as attention-catching devices to generate object hypotheses. The second phase is a top-down process in which object hypotheses are confirmed or rejected by fitting a local model to ribbons surrounding the low intensity blob. The local model is acquired using inductive learning. Flexible matching routines are used during recognition that provide a measure of confidence for the identification. Experimental results demonstrate the ability to learn the relationship between image characteristics and object functionality.

1 Introduction

This paper presents work in progress on an approach to the problem of recognizing blasting caps in x-ray images. This problem is an instance of a class of problems in which a vision system must inspect a sequence of images for known objects. Unfortunately,

the fact that the object is known is often of little or no help. If there is little standardization of the class of known objects, it becomes impractical to attempt to model the objects geometrically. Yet, what often constrains the class of known objects is their functionality [Freeman and Newell, 1971; Stark and Bowyer, 1991; Rivlin et al., 1994]. Consequently, learning can be useful for acquiring the relationship between image characteristics and object functionality [Woods et al., 1995].

Our primary focus is to investigate how vision and learning can be combined to find blasting caps as well as objects that could occlude blasting caps. In a previous study [Maloof and Michalski, 1995], learning was used to acquire concept descriptions of blasting caps. Simple segmentation techniques were used to isolate objects from their background; they were then represented using intensity and geometric features.

For the work presented here, an analysis of functional properties of blasting caps was conducted to design the representation space for learning, which combines intensity and shape features. The first phase of the approach is a bottom-up process in which low intensity blobs, which possibly correspond to heavy metal explosive in the approximate middles of blasting caps, are used as attention-catching devices to generate object hypotheses.

In the second phase, a top-down process, object hypotheses are confirmed or rejected by fitting a local model to ribbons surrounding the low intensity blob. These ribbons possibly correspond to the metal shell of the blasting cap. The local model is acquired using the inductive learning system AQ15c [Wnek et al., 1995] and is represented as a set of decision rules. Flexible matching routines are used during recognition that provide a measure of confidence for the identification. Experimental results demonstrate the ability of an inductive learning system to acquire the relationship between image characteristics and object functionality.

*This research was supported by the Advanced Research Projects Agency under grants F49620-92-J-0549 and F49620-95-1-0462 administered by the Air Force Office of Scientific Research. The development of machine learning systems used in this research was supported in part by the Advanced Research Projects Agency under grant N00014-91-J-1854 administered by the Office of Naval Research, in part by the Office of Naval Research under grant N00014-91-J-1351, and in part by the National Science Foundation under grants IRI-9020266 and CDA-9309725.




Figure 1: The geometry of x-ray imaging [Dance, 1988].

This research provides an opportunity to study the interplay between vision and learning processes [Michalski et al., 1994], especially as it relates to learning object functionality. A vision system capable of reliably recognizing blasting caps or objects that could occlude blasting caps could be used to aid airport security personnel with luggage screening.

2 Preliminaries

2.1 Imaging System and Image Formation

A typical x-ray imaging system consists of an x-ray tube (photon source), an anti-scatter device, and a receptor (photon detector) [Dance, 1988]. The photons emitted by the x-ray tube enter the objects, where they may be scattered, absorbed or transmitted without interaction. The primary photons recorded by the image receptor form the image, but the scattered photons create a background signal (i.e., noise) that degrades contrast. In most cases, the majority of the scattered photons can be removed by placing an anti-scatter device between the objects and the image receptor.

What follows is a simple mathematical model of the imaging process. We start by considering a monochromatic x-ray source that emits photons of energy E and is sufficiently far from the objects (luggage) being inspected that the photon beam can be considered to be parallel (see Figure 1). The incident photon beam is parallel to the z direction and the image is recorded in the xy plane. We assume that each photon interacting with the receptor is locally absorbed and that the response of the receptor is linear, so that the image may be considered as a distribution of absorbed energy. If there are N photons per unit area incident on the object and $I(x, y) dx dy$ is the energy absorbed in area $dx dy$ of the detector,

then

$$I(x, y) = \frac{\exp\left(-\int \mu(x, y, z) dz\right)}{N \varepsilon(E, 0) E (1 + R)} \quad (1)$$

where the line integral is over all materials along the path of the primary photons reaching the point (x, y) , $\mu(x, y, z)$ is the linear attenuation coefficient, $\varepsilon(E, 0)$ is the energy absorption efficiency of the receptor for the photon energy level E at an incident angle of 0, and R is the ratio between the scattered and primary radiation (which is usually very small).

2.2 Imaging Model

Following the description in Section 2.1, we assume orthographic image projection (see Figure 1). The image of the object point (X, Y, Z) is the point (x, y) such that

$$x = sX, \quad y = sY, \quad (2)$$

where s is a constant. The image intensity at the pixel (x, y) is obtained by integrating (1) over the area of the pixel in the image receptor.

2.3 The AQ Learning Method

AQ15c [Wnek et al., 1995] is the latest implementation of the AQ algorithm [Michalski, 1969]. AQ is an inductive learning algorithm that learns decision rules from training examples encoded as VL₁ expressions. VL₁ is an attributional variable-valued logic system and is a subset of the first order variable-valued logic system VL₂. For the attributional case, each training example is represented as a vector of attribute-value pairs annotated with a concept label.

The *representation space* for a problem is the set of all terms and their domains used to encode a problem for learning. Thus, a training example is a point in the representation space. Conceptually, rules carve out decision regions in the representation space that cover training examples in a way that is consistent with their concept labels. Computing the set of maximally general decision rules that cover all positive and no negative training examples is a special case of the general covering problem, which is known to be NP-complete. The AQ algorithm computes quasi-optimal solutions for this problem and guarantees that learned concept descriptions cover all positive and no negative training examples.

During recognition, learned concept descriptions, in this case rules, are matched to observations to deduce class membership. Depending on the generality of the rules induced by AQ, areas of the representation space may not be covered explicitly by any rule. If an observation happens to be taken from an uncovered region, using a *strict matching* method,




Figure 2: Detailed x-ray of a blasting cap.

this observation would remain unclassified — which in certain situations may be desirable. On the other hand, if using a *flexible matching* method, all observations not explicitly covered by a rule can be classified by assigning it the concept label of the closest concept description in the space. Several flexible matching schemes exist and use various metrics when computing the *degree of match*, including syntactic distance, number of satisfied rule conditions, and rule strength.

3 Problem Statement

The intensity in x-ray images is proportional to the number of x-ray photons that pass through objects on their path from the source to the receptor (see Section 2.1). Since different materials have different transparency properties, the intensity of an x-ray image depends on both the thickness and the type of material between the source and the receptor. Moreover, any x-ray photon that is not absorbed by one object on its path can be absorbed by another. Thus, a thick layer of semi-transparent material can have the same effect on the image receptor as a thin layer of opaque material.

Although blasting caps are manufactured objects, there is enough variability in their manufacture that makes a CAD-based recognition system impractical. What is common to all blasting caps, however, is their functionality. Ultimately, blasting caps are defined by their functional properties, not by their shapes.

A typical blasting cap (see Figure 2) consists of a cylindrical metal shell filled primarily with the explosive. In its approximate middle, there is a small globule of heavy metal secondary explosive. Finally,

leg wires from the electric ignitor extend from one of the ends. The most dense (opaque to x-rays) part of a blasting cap is the concentration of the heavy metal explosive, which is approximately centrally symmetric. The leg wires also produce dense features, but are very thin. Finally, the copper or aluminum tube filled with explosive, which is axially symmetric, is typically more dense than the surrounding areas of the luggage.

Regarding images of blasting caps, we begin by considering a generic blasting cap that is not occluded by opaque material. Let l be the length of an approximately cylindrical blasting cap, r be its radius, and σ be the angle between the axis of the cap and the image receptor. Consider the length of the path p of an x-ray photon as it passes through the blasting cap. When $\sigma = 0$, p ranges from $2r$ at the axis to 0 at the occluding contour. In general, p is multiplied by $\sec \sigma$; however, p cannot be longer than l . From (1), we see that the number of photons passing through the blasting cap decreases exponentially as p grows. From (2), we see that the image of a blasting cap is rectangularly shaped; its width is approximately $2rs$, and its length is approximately $ls \sec \sigma$. Its intensity is lowest along the axis of the blasting cap, and highest along the occluding contour, which produces a low contrast boundary. Also, the image of the heavy metal secondary explosive (see Figure 2) appears as a small approximately symmetric blob on the axis of the blasting cap. The center of the blob is nearly opaque and thus its intensity is near zero. The boundary of the blob is lighter, but still has a very low intensity. The leg wires are strong features, but are not clearly visible in the images.¹

Therefore, the strongest feature of a blasting cap is the low intensity blob in the center of a rectangular ribbon of higher intensity. The intensity of both the blob and the ribbon is lowest along the axis of the blasting cap and highest along the occluding contour. Finally, if a blasting cap is occluded by any object, its image will be darker than the image of a blasting cap that is not occluded.

3.1 The Method and Experimental Results

We present a two phase, bottom-up and top-down learning approach for recognizing blasting caps in x-ray images. In the first phase, low intensity blobs, which serve as attention-catching devices, are used to generate object hypotheses. These low-intensity blobs correspond to the secondary high explosive, which is typically a heavy metal compound, located

¹In our examples, image resolution is 565×340 and the leg wires are barely visible. Currently, we are working to obtain images of higher resolution so that the leg wires can be detected.




Figure 3: Sample image used for experimentation.

near the middle of the blasting cap (see Figure 2).

In the second phase, each generated hypothesis spawns a process that attempts to fit a local model to ribbon-like features surrounding the blob. These ribbon features correspond to the metal body of the blasting cap (see Figure 2). The local model is acquired using the inductive learning system AQ15c and captures intensity and geometric features of both the low intensity blob and the surrounding ribbon shape. A flexible matching routine is used to match the local model to the image characteristics, which not only produces an object identification, but also yields a confidence in the identification.

The x-ray images used for experimentation were of luggage containing blasting caps appearing in varying orientations and under varying amounts of clutter, which included clothes, shoes, calculators, pens, bolts, batteries, and the like. The luggage was imaged much as it would be in an airport scenario: flat in relation to the x-ray source, but rotated in the image plane. Five images were selected from a set of 30 which were of low to moderate complexity in terms of clutter and positional variability of the blasting cap. Figure 3 shows one of the images used for experimentation.

Regions of interest were interactively determined, and contained low intensity blobs and ribbons corresponding to positive and negative examples of blasting caps. From each of the 64 selected regions, 27 geometric (e.g., compactness and proximity measures) and intensity-based features (e.g., minimum, maximum, and average) were computed, resulting in 28 blasting cap and 38 non-blasting cap objects. The AQ15c [Wnek et al., 1995] inductive learning system was used to learn concept descriptions of blasting caps and non-blasting caps.

Induced concept descriptions from AQ15c were validated using 100 iterations of 2-fold cross-validation. This validation method involves 100 learning and

Average Predictive Accuracy (%)		
Overall	Correct	83.51±1.3
	Incorrect	16.49±1.3
Blasting Cap	Correct	85.82±2.1
	Incorrect	14.18±2.1
Non-Blasting Cap	Correct	81.19±2.4
	Incorrect	18.81±2.4

Table 1: Summary of quantitative experimental results.




Figure 4: Test image for applying learned concepts.

recognition runs. For each run, the extracted image data was randomly partitioned into a training set and a testing set. After learning from examples in the training set, the induced concepts were tested using examples in the testing set. We can compute the predictive accuracy for each run based on the correct or incorrect classification of the examples in the testing set. The overall predictive accuracy for the experiment is the average of the accuracies computed for each run. These results are summarized in Table 1 and show the average predictive accuracy with a 95% confidence interval for the overall experiment and for each class.

As a qualitative demonstration of the method, learned concepts were also applied to an unseen image. The learned concepts from AQ15c using training data from four images were tested on objects extracted from a fifth, unseen image, which is shown in Figure 4. Objects 1–6 are blasting caps, objects 7–10 are not. Object 5, which is a blasting cap, was mis-classified. All other objects in this image were classified correctly.

4 Conclusions

This paper presented work in progress on the problem of recognizing blasting caps in x-ray images. In the first phase of a two phase learning approach, low intensity blobs were used as attention-catching devices. This bottom-up process was followed by a

top-down recognition process in which a learned local model was matched to ribbon-shaped image regions surrounding a low intensity blob. An analysis of functional properties of blasting caps was used to design the representation space for learning, which combined intensity and geometric features. Experimental results suggest that learning can be used to acquire functional descriptions of objects. This is important for classes of objects for which geometric modeling is impractical.

Future work in this area will involve further automation of the feature extraction process and object labeling functions. In addition, other functional properties present in blasting caps still require exploitation. An example is the presence of leg wires (see Figure 2). Unfortunately, the current image set is not of a resolution that allows for the detection of these functional properties. We hope to acquire additional images that will be better suited for this type of analysis.

Acknowledgments

The authors would like to thank Eric Bloedorn, Ken Kaufman and Janusz Wnek for discussions regarding this work.

References

- Dance, D.R. 1988. Diagnostic radiology with x-rays. In Webb, S., ed., *The Physics of Medical Imaging*, 20–73. IOP Publishing: Philadelphia, PA.
- Freeman, P. and Newell, A. 1971. A model for functional reasoning in design. *Proceedings of the International Joint Conference on Artificial Intelligence*, 621–640.
- Maloof, M.A. and Michalski, R.S. 1995. Learning descriptions of 2D shapes for object recognition in x-ray images. *Proceedings of the 8th International Symposium on Artificial Intelligence*, 124–131.
- Michalski, R.S. 1969. On the quasi-minimal solution of the general covering problem. *Proceedings of the Fifth International Symposium on Information Processing*, A3:125–128.
- Michalski, R.S.; Rosenfeld, A.; and Aloimonos, Y. 1994. Machine vision and learning: research issues and directions. *Reports of the Machine Learning and Inference Laboratory*, MLI 94-6. Machine Learning and Inference Laboratory, George Mason University, Fairfax, VA. Also: CAR-TR-739, Center for Automation Research, University of Maryland, College Park, MD.
- Rivlin, E.; Dickinson, S.J.; and Rosenfeld, A. 1993. Recognition by functional parts. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2–7.
- Stark, L. and Bowyer, K. 1991. Achieving generalized object recognition through reasoning about association of function to structure. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13:1097–1104.
- Wnek, J.; Kaufman, K.; Bloedorn, E.; and Michalski, R.S. 1995. Inductive learning system AQ15c: the method and user's guide. *Reports of the Machine Learning and Inference Laboratory*, MLI 95-4. Machine Learning and Inference Laboratory, George Mason University, Fairfax, VA.
- Woods, K.; Cook, D.; Hall, L.; Bowyer, K.; and Stark, L. 1995. Learning membership functions in a function-based object recognition system. *Journal of Artificial Intelligence Research* 3:187–222.