Tracking Nuclear Material at Low Frame Rate and Numerous False Detections

Paolo Lombardi, Cristina Versino

European Commission – DG Joint Research Centre-IPSC-NUSAF, 21020 Ispra, Italy
{paolo.lombardi, cristina.versino}@jrc.it

Abstract. In Nuclear Safeguards, surveillance cameras monitor the correct processing of nuclear material. Nuclear inspectors are faced with tens of thousands of images to review, of which less than 1% is significant. The Safeguards authorized means of reducing the image set is, at present, a two-frame differencing change-detection filter. We further limit the set by tracking flasks thanks to the timing of events provided by image time-stamps. Traditional visual tracking cannot be applied, owing to the low frame rate, and the need for compatibility with the authorized change detection filter. Our algorithm is based on a hidden semi-Markov model of the nuclear process, and handles multiple flasks and observations available only when the flasks are moved. State occupancy durations are used when selecting the next image to review.

Keywords: Hidden semi-Markov model, tracking, nuclear processing.

1 Introduction

In Europe, there are more than 1000 nuclear sites verified by about 200 inspectors of the EURATOM Safeguards authority. Each surveillance camera acquires about 20,000 images before these are reviewed, of which less than 1% are related to safeguards-relevant events. The remaining are images either with no change or with insignificant events (e.g. moving cranes, trolleys, illumination changes). Inspectors eliminate the no-change images by applying a *scene change detection* algorithm (SCD) based on two-frame differencing [1]. Typically, this operation reduces the original image set from 20,000 to 2,000 images. The latter, reduced image set is manually reviewed and annotated to produce a *review report*.

Nuclear processing being a well structured procedure, we claim that *temporal and historical information* can further reduce the image set to review. Herein we present a second filter, computing after SCD, which performs *tracking* with a motion model *trained on time-stamps of events*.

The specific scenario of nuclear plants poses new challenges to the state-of-the-art visual tracking. The field of view covers all the locations where important processing takes place (Fig. 1). These locations are many meters apart, thus the appearance of a flask, which is the object of interest in image reviews, forcefully changes during the
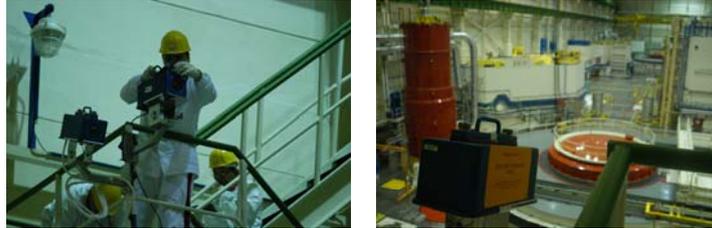
Fig. 1. Inspector setting-up the cameras (left) in a real plant (right) (© D.Calma/IAEA).

process. Furthermore, the acquisition rate is very low (1 frame every 3-10 min). Few images show the flask in any given location during a process cycle.

In such conditions, traditional tracking with templates or adaptable shapes is difficult because: i) widely different appearances of the same objects in different locations; ii) the *low frame rate* does not guarantee the continuity of tracking, and makes dynamic template/shape update difficult. Additionally, iii) any tracking algorithm must be fully compatible with SCD in order to be acceptable by Safeguards authorities, and iv) it must guarantee that, subject to correct parameter setting, a *zero non-detection rate* (or 'false negative rate') is achievable.

The solution discussed in this paper is based on a Markov-model representation of flask processing in a nuclear plant, with a statistical representation of process duration to better predict the next state. The algorithm is proposed for use in *interactive mode* to discover all images referring to *correct* stages of flask processing.

2   Image Review of a Nuclear Process

In a typical nuclear process a flask of fuel enters the *hatch* and reaches the *decontamination* area, whence it is moved to the *pond*. From the pond, the flask moves back to the decontamination area and then exits the scene through the hatch. Table 1 shows the list of Safeguards-relevant events to be annotated by inspectors during reviews.

A plant may be able to support the parallel processing of multiple flasks, so that a flask may begin the processing cycle while the previous is still ongoing. Also, nuclear plants possess a limited number of cranes to move the flasks. This feature limits the number of flasks that can enter a new stage of processing in any given moment, and hence an image can contain a number of safeguards-relevant events at most equal to the number of cranes.

During a review, inspectors first define a set of areas of interest (AOI) on the image plane around the interesting locations. For the example of Table 1, three AOIs would be drawn on hatch (H), decontamination (D), and pond (P). Then, a SCD algorithm selects for review only images with AOIs that exhibited change beyond a threshold. For instance, if an image *k* is labeled [H, P], this means that the AOIs of hatch and pond exhibited sufficient change from *k-1* to *k*. Inspectors browse the labeled batch and annotate only the relevant images with the appropriate event class. *It is important to note that relevant images are annotated as soon as they are reviewed*.

Stage	Event Name	SCD label
1	Flask over hatch (ENTRY)	H
2	Flask in decontamination area (ENTRY)	D
3	Flask over pond (ENTRY)	P
4	Flask over pond (EXIT)	P
5	Flask in decontamination area (EXIT)	D
6	Flask over hatch (EXIT)	H

Table 1. An example of nuclear process composed of six stages, with relative labels.

Fig. 2. Facsimile of a report redacted by inspectors after reviewing an image set.
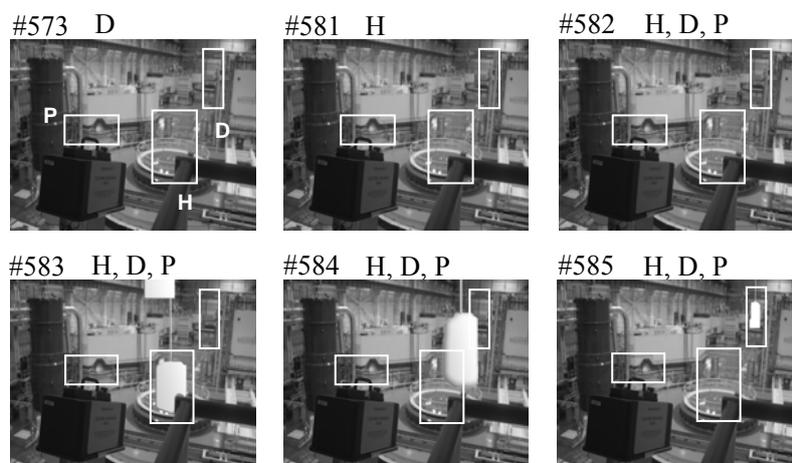
Fig. 3. Facsimile of a sequence selected by SCD during a real review, together with real SCD labels (bottom). The images and AOIs are fictitious (real images are confidential).

A facsimile of inspectors' *report* is shown in Fig. 2 for a multi-flask process: in #611 a second flask enters through the hatch while the first flask is still in decontamination.

With reference to Table 1 and Fig. 2, if SCD had 0% false alarm rate, inspectors would browse a sequence like: #583:H – #585:D – #611:H – #612:D – #1220:P – #1398:P – #1400:D – #1465:P. With thresholds set so as to have zero false negatives, the false alarm rate of SCD is over 95%. Thus, the sequence selected by SCD around #583 is rather like Fig.3, where all but H in #583 and D in #585 are false alarms.

3 A Markov Model for Tracking Flasks

In order to produce the report of Fig. 2, inspectors normally examine every single SCD image. We argue that *tracking the state* of the multi-flask processing would help in further reducing the number of images to review.

Our problem consists in tracking the process state given the sequences of SCD labels and the online confirmation of the events by the inspector. Once the present state is known, one may predict the next event and therefore filter out all images having a SCD label inconsistent with the prediction.

In our low-frame-rate datasets there is meager or no continuity of information (a problem for visual tracking techniques), but there is *recurrence*. A flask undergoes a structured cycle which is the same for all flasks. *Therefore tracking can be performed by modeling these regularities*.

We employ a *hidden semi-Markov model* (HSMM) of the processing performed in a nuclear plant. A discrete hidden Markov model (HMM) is a popular tool to represent time series of events [2] where the duration of state occupancy is intrinsically of geometric distribution. HSMMs relax this constraint by introducing flexible state occupancy distributions to represent sojourn times in non-absorbing states [3]. A semi-Markov model (SMM) is composed of an *embedded* first-order Markov chain X and of discrete distributions of sojourn times S. The embedded chain is described by (T, χ_o), where χ_o is the *initial state* distribution and T is the *transition matrix*, such that $T_{ij} = P(X_{t+1}=j \mid X_t=i)$. For a SMM, $T_{ii}= 0$, $\forall i$. The sojourn time distributions are a set of discrete distributions depending only on the current state, S = {S_i, $\forall i$}. The model is *hidden* if the relation between the state and the observation is probabilistic. The emission distributions for every state are summarized in the *emission matrix* E, $E_{is} = P(O_t=s \mid X_t=i)$, s being an emitted *symbol* (Fig. 4).
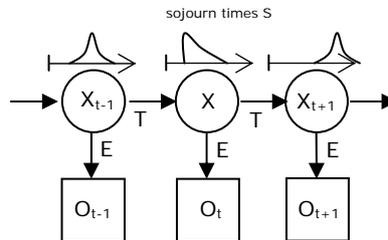
Fig. 4. A hidden semi-Markov model (HSMM) represents events with generic durations.

We define a *plant* by the parameters (*F*, *K*, *N*). *F* is the maximum number of flask processings supported in parallel by the plant. *K* is the number of available cranes to move the flasks around ($K \leq F$). *N* is the number of processing stages that make up the nuclear process. In the example of Table 1, $N = 6$. In realistic plants, $1 \leq F \leq 3$, $K=1$, and $3 \leq N \leq 10$.

3.1 HSMM of a plant

As previously said, a HSMM is completely specified by (T, E, χ_o, S). In our model, time is discretized so that each tick corresponds to a newly acquired image. The ticks are equally spaced in time.

For building the state space, let us first define the *real states* of the plant as *F*-arrays of labels indicating the individual progression of each flask in its processing. For instance, the *real states* of a plant with $F = 3$ and $N = 6$ are [1 1 1], [1 1 2], [1 2 2],… [6 6 6]. We observe that:

1. In practice, flasks cannot be distinguished. Hence a given state coincides with all states that can be expressed as a permutation of its labeling array (e.g. [1 3] is the same as [3 1]). This decreases the number of *real states* the model needs.
2. A correct SCD observation occurs when there is flask movement in an AOI, i.e. when a flask *enters* or *exits* a location. Either ingoing or outgoing movements can be chosen to assign emissions to a state, but the choice must be coherent for all states. We choose to assign an emission distribution to a state on the basis of flasks *entering* that state. Even after this choice, *real states* have in general *F* emission distributions, one for each possible event that may "cause" that state. For instance, consider a plant with $F=2$. Its state [2 3] can be entered from [u 3], $u \neq 2$, thus triggering an emission linked to a flask entering stage 2. Or it can be entered from [2 v], $v \neq 3$, thus triggering an emission typical of a flask entering stage 3. To allow for multiple emission distributions, each *real state* is designed to be represented by *F* *virtual states*.

Given these premises, the size of the state space for a plant (F, K, N) is *M*:

$$M = F \frac{(N + F - 1)!}{F!(N-1)!} = \frac{(N + F - 1)!}{(F-1)!(N-1)!} \tag{1}$$

The transition matrix is a *MxM* square matrix. The size *M* is independent of *K*. However, note that the *K* cranes allow only *K* flasks to change state at any given time. This fact constrains the transition probability T_{ij} to be 0 if two real states *i* and *j* differ for more than *K* single-process states. For example, in a plant with $K = 1$ and $F = 2$, the transition [1 1] $\rightarrow$ [2 2] is impossible. Moreover, all virtual states referring to the same real state have the same transition probabilities towards other real states (equal rows in T). A trivial example of T for a plant with $F=2$, $K=1$, and $N=2$ is given in (2). Note that, with $N=2$, the only possible events are either a flask going from stage 1 to stage 2 or vice versa (T_{ii} being null by definition).

$$
T = \begin{pmatrix}
0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 \\
0.5 & 0 & 0 & 0 & 0 & 0.5 \\
0.5 & 0 & 0 & 0 & 0 & 0.5 \\
0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0
\end{pmatrix}
\begin{matrix}
\text{\small[1 1]} \\
\text{\small[1 1]} \\
\text{\small[1 2]} \\
\text{\small[1 2]} \\
\text{\small[2 2]} \\
\text{\small[2 2]}
\end{matrix}
\quad
E = \begin{pmatrix}
1 & 0 \\
1 & 0 \\
1 & 0 \\
0 & 1 \\
0 & 1 \\
0 & 1
\end{pmatrix}
\tag{2}
$$

Real states

Because the tracking is performed on the labels outputted by SCD, we define the *emission alphabet* of size A as the set of SCD labels associated to the AOIs drawn by inspectors. In the example of Table 1, the emission alphabet is {H, D, P}, and $A = 3$. To give more insight, let us observe that A is the correct size of the emission space because we work under the assumption that SCD has a null false-negative rate – thanks to an "oracle" which has set the correct parameters. If the model is changed to remove this assumption, then an 'idle' symbol representing a missed SCD detection should be introduced.

The matrix E is sized MxA. The emission distributions are 1 for the correct symbol, 0 otherwise, see (2) above. In our case the sequence to be decoded comes from a fully reliable knowledge source (the inspector). As such, the hidden nature of the model lies only in the fact that the same symbol is emitted by a multiplicity of states, but all uncertainty on the "value" of the observed symbol disappears as soon as the inspector confirms or rejects it – as it will be better explained in Section 3.2.

The *initial state* χ_0 is available from the last annotated image of the preceding report: in fact, the surveillance process is continuous and it is split in subsequent batches of images for convenience of inspections (Fig. 2).

The distributions of *sojourn times* S are referred to single-flask processes and in general are non-geometric. Thus, each compound state in a plant has F associated distributions. For instance, state [2 3] is associated to durations of stage 2 and stage 3. We employ parametric distributions of types Gaussian, geometric with shift, and uniform, because: i) the Gaussian pdf captures a mean duration with relatively little divergence; ii) the shifted geometric pdf captures the case of a minimal processing time followed by a decaying duration; iii) the uniform distribution captures the case of the beginning of a new process when the delay is unknown. Table 2 shows an example of S for the plant of Table 1. The unit time is one-frame interval. μ and σ are the mean and standard deviation of a Gaussian, α and β are the parameters of the geometric distribution $\alpha(1-\alpha)^{n-\beta}$, defined for $n \geq \beta$.

Stage	Statistics of duration
1	Geometric, $\alpha = 0.288$, $\beta = 1$
2	Geometric, $\alpha = 0.054$, $\beta = 5$
3	Gaussian, $\mu = 186.9$, $\sigma = 18.4$
4	Geometric, $\alpha = 0.195$, $\beta = 1$
5	Gaussian, $\mu = 183.4$, $\sigma = 11.5$
6	Uniform

Table 2. Descriptive statistics of stage duration S for the process of Table 1.

3.2 Using the model interactively

The general idea is that, given the *history of annotations* produced by the inspector during the review, the (T, E, χ_0, S) model can highlight the next most likely relevant SCD image. The inspector can decide to accept it by marking an annotation with a label corresponding to an admitted event, e.g. H D or P. If she 'scrolls forward' we as-

sume the image was rejected. The interaction is repeated with the next image until the end of the review.

Let us consider a mid-point during the review, when the inspector's past annotations form a sequence ν. In the framework of an official review, we can consider the inspector as a fully reliable knowledge source, so that ν is true with probability 1. This implies that the uncertainty on the timing of events and on the value of symbols in ν is null. HSMM decoding [3] becomes superfluous: we are allowed to use simple HMM decoding [2] to retrieve the *current state* distribution χ, with a decisive advantage in terms of computational complexity.

The next image to present to the inspector is selected by computing the likelihood of every future image selected by SCD given χ and S. This likelihood is given by summing the S distributions of all stages associated to a state i with uniform probability $(1/F)$ and then weighting this sum by the probability of the state χ_i, $\forall i$. The image exhibiting Maximum Likelihood (ML) is selected. If the inspector rejects this candidate, its likelihood is set to zero and the second ML is proposed, etc. When an image is accepted and annotated, its symbol is added to ν and the procedure restarts.

In case of a missed detection, which can happen due to the Bayesian framework of predictions, ν can be either inconsistent or consistent given (T, E, χ_o). In the former case – easily detected because χ is zero –, confirmed symbols are temporarily suppressed in turn, with the heuristic of "last-confirmed suppressed first", and the remaining sequence is re-decoded until a consistent sequence is found. The prediction procedure described above is then applied from that point.

The case of a consistent sequence is treated as correct in all senses. Thus an error (missed detection) giving birth to a consistent sequence is not discovered unless, as the review progresses, the sequence ν becomes inconsistent. This problem unfortunately nullifies any guarantee of null false detection rate of the HSMM.

The guarantee of detecting a missed event given an inconsistent sequence ν instead exists. The duration distributions can be chosen so that all images are assigned a non-null probability of being selected (distributions with support on the whole set). Thus, in the worst case the whole set will be proposed, image by image, until consistency is reestablished.

3.3   Training the model

A nice property of the proposed approach is its automatic learning ability without the need of further intervention by the inspectors. This property plays a more important role than usual, as inspectors would hardly accept the increased overhead of training a computer system on top of their responsibilities and time.

The actual values for elements of T are derived automatically by running the Baum-Welch learning algorithm for HMMs [2] on *previous reports*. Because previous reports contain only *true events* as annotated by the inspectors, the assumption on the distributions of E (Section 3.1) holds true also during training. S are likewise derived by standard parametric fitting on duration data in previous reports.

Note that with our approach we employ for training all and only the information available on previous reports (types of events, their order and the time intervals between).

4 Performance Analysis

Nuclear Safeguards images are confidential in nature. For this reason, we test our techniques on image sequences from [4]. Results are provided for these images (sequences 'A') and for the real Safeguards images (sequences 'B').
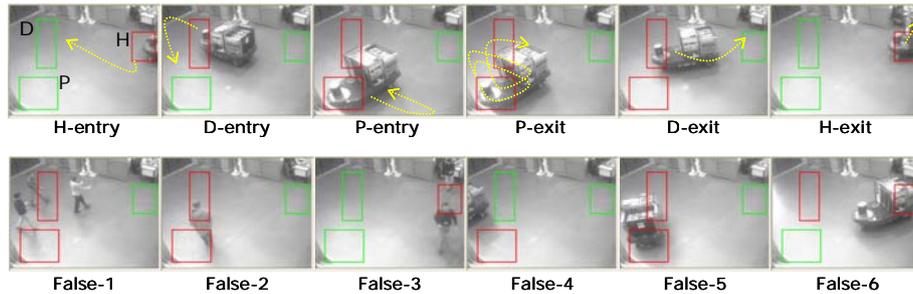

Fig. 5. Images from set A. The first row shows the relevant events in this exercise. The images on the second row are false positives that trigger the SCD filter.

4.1 Sequences A

Sequences A are acquired in a laboratory environment where a mobile robot executes a series of programmed movements (Fig. 5). The setup simulates a plant with $F=1$, $K=1$, $N=6$. The robot moves a flask around a fictitious hatch (H), a decontamination area (D) and a pond (P), each identified by one AOI in Fig. 5. An AOI is drawn in red when SCD is triggered in that area; it is green otherwise. Examples of relevant images that testify the events H-entry, D-entry, P-entry, P-exit, D-exit, H-exit are shown on the first row of Fig. 5; the yellow arrows sketch the path followed by the robot. The flask is not always visible (e.g. due to rotatory movements of the robot after P-exit). The image set contains 1639 time-stamped frames corresponding to five runs of the robot, for a total of 30 relevant events (Table 3).

Due to non-perfectly identical starting positions, the robot reaches the AOIs at different times in different runs, but overall the timing of events has only a limited variability. Thus, this image set captures a salient feature of Safeguards images: *recurrence over time*. On the other hand, many false positives trigger the SCD: people passing by (False 1-3), the shadow projected by the robot on AOIs (False 4) as well as its presence on an AOI during 'flask-processing' stages which are not to be annotated as relevant in a review report (False 5-6).

4.2 Sequences B

Sequences B are *real-situation* images acquired by a safeguards camera in two different plants, and are accompanied by the respective official inspectors' reports. The image sets are characterized in Table 3. The camera has the same bird-view setup as the one reproduced in sequences A. Sets B1-B3 are taken from a plant with $F=1$, whereas

B4 and B5 are special in that they refer to a plant supporting two flask processings in parallel (F=2). Both plants process the flasks following Table 1 (N=6), and they employ a single crane (K=1). As above, SCD data are generated by drawing three AOIs on hatch, decontamination area, and pond. The images of set B are confidential.

Image set	Images	H	D	P	All events
A	1639	10	10	10	30
B1	20160	17	17	17	51
B2	15661	1	1	1	3
B3	16022	-	-	-	-
B4	16020	30	30	30	90
B5	15446	12	12	12	36

Table 3. Number of images and of events on locations H, D and P for image sets A and B.

4.3 Experiments

The SCD algorithm is run on the image sets, with low change-detection thresholds so that the false negative rate is zero. For sequences A, we train the model (T, E, χ_o, S) on two sequences, and test the model on the remaining three. For sequences B, we train the model on two single-flask cycles of B3 (12 total events), and test the same model on all other sequences B. The model for B1 and B2 is induced from the single-flask model by extending it to the case of F=2. The inspectors' reports play the role of the inspector during the interactive phase of the algorithm: an image is accepted if it is annotated on the report, otherwise it is rejected. The descriptive statistics S are either Gaussian, 'delayed' geometric or uniform distributions. The values reported in Table 2 refer to the model trained on B3.

As a measure of performance we compare the false alarms obtained with the SCD filter alone, and with the filter 'SCD + (T, E, χ_o, S) model + the online inspector'. We observe a *null false negative rate* also after our filter. Meanwhile, we observe a *reduction in the false alarms* (Table 4). We observe reductions to 38% of the original number as a least performance, up to a best reduction to 7%. Most interesting is the reduction to 30% or less in B4 and B5, featuring two flasks in parallel.

Set	SCD	SCD+HSMM+INSP	%
A	378	97	25,6
B1	1493	279	18,6
B2	390	147	*37,7*
B3	2411	184	***07,6***
B4	1694	520	30,7
B5	964	244	25,3

Table 4. False alarms generated by two filter configurations, and the reduction apported by HMM expressed as percentage of false alarms in column 3 with respect to column 2.

5 Discussion

We consider the reductions to <40% of the inspectors' effort as a significant result. They translate into a speed-up of 2 to 3 times with respect to current reviews, this way increasing the potential of an endured concentration throughout the review. Meanwhile, the chance of missing an important event remains low. Strictly speaking, our approach does not respect the requirement of zero non-detection in case of errors producing a consistent sequence ν (Section 3.2). However, in our experiments it never occurred that a wrong consistent sequence remained consistent until the end of the review. For this to happen, a full flask cycle must be missed. If at least one event is instead detected, this event will generate an inconsistency in ν and thus will prompt a re-decoding and a retrospective analysis of events previous to the current point of the review, as described in Section 3.2.

A few experiments were made on sequences B also with a non-deterministic E matrix, which takes into account false alarm/false negative rates of the SCD via an extra 'idle' symbol. The results showed that the false alarm rate of >95% misleads the *forward-backward* algorithm into wrongly detecting an overwhelming number of false 'flask processing cycles'. We concluded that the interactive loop with a human operator is necessary, and also it is *for free*: the system adapts by taking advantage of operations the inspector would do in any case to compile the online review report. This "effort-free" adaptive behavior is important for a system to be accepted by "real people" with work practices that cannot be completely redesigned.

Concerning the SCD, it is far from being of state-of-the-art level. However, it is interesting to notice that its possibility to reach a zero non-detection rate when used with conservative SCD thresholds is what makes it attractive for nuclear applications. We are pursuing work on proposing Safeguards authorities new algorithms for motion detection which present this feature accompanied by a lower false alarm rate.

The HMM uses a model of structured processing, thus it is not meant to discover images of *irregular behavior*. On the opposite: irregular behavior would be filtered out by our model. Starting from this observation, we plan to use this model as a basis for a *novelty detection* technique [5] to detect diversions from the approved and safe nuclear processings online.

In our future work, we also plan to test the model on images of plants with $F>2$. Also, we are considering using an emission alphabet constituted by the full label-arrays outputted by SCD instead of single labels.

References

[1] Canberra Aquila – http://www.aquilagroup.com/pdf/canberra/GARS.pdf

[2] L. R. Rabiner, "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition", Proceedings of the IEEE, 77 (2), p. 257–286, February 1989.

[3] Y. Guédon, "Estimating hidden semi-Markov chains from discrete sequences", J. Comput. Graphical Statistics, 12(3), pp. 604-639, 2003.

[4] C. Versino, E. Stringa, J. Goncalves, "Review of Surveillance Images Using Classification Techniques in a Relevance Feedback Context", PETS 2003, Nice, France, 2003.

[5] M. Markou, S. Singh, "Novelty Detection: A Review", Signal Processing, v. 83(12), pp. 2481 – 2497, 2003.