## Optimising the Use of REST for Mine Detection

by Ian G. McLean and Rebecca J. Sargisson, *GICHD*

### Introduction

Remote Explosive Scent Tracing (REST) is a detection technology involving the transfer of odours to an animal detector using filters.[1] Like Remote Scent Tracing (RST), the technology could potentially be used to detect anything that has an odour.[2] REST technology was used originally by Mechem in Mozambique and Angola in the early 1990s.[3] Despite the potential it demonstrated at that time, it received little attention or investment through the late 1990s until a revival of interest occurred in recent years. Currently, it is being used operationally for mine detection in Afghanistan and is likely to be implemented for road clearance in Sudan and Angola by the end of 2004. It is also used operationally for explosives detection at several airports in Europe.

REST is most effective when used as a fast area-reduction system on roads. There is a desperate need for rapid road clearance in several mine-infested countries today. In contrast to technologies such as demining machines, it has a minimal environmental impact. It therefore offers advantages in terms of both recovery of infrastructure and ecological disturbance. However, REST is not a simple technology. It requires considerable animal-training expertise and significant investment in both time and equipment. In training, a REST detector has the following two key objectives:

1. To create high reliability of detection for odour of mine/UXO
2. To minimise the false-alarm rate

If a true positive is missed, then a dangerous item may be left behind. A false alarm creates additional work for a clearance agency. Improved reliability of detection of true positives can be achieved using additional detectors. However, each extra detector is likely to give additional false alarms, raising the overall clearance requirements for an area.

In this article, we explore the following two issues:

1. Although not normally considered in this way, REST is essentially a survey tool. We will explain how it may be used to estimate the true-contamination rate in a suspect area.
2. REST technology involves a trade-off between ensuring high search reliability and minimising clearance requirements. We discuss two ways to reduce false positives without compromising detection reliability.

### Estimating the True Contamination Rate

To be used sensibly, the notion of a “true-contamination rate” requires a link to a geographic scale. In the study of mechanical mine clearance[4] by the Geneva International Center for Humanitarian Demining (GICHD), it was noted that the area of land actually containing mines/UXO in a minefield is very small (about two percent, based on 15 separate measurements). The contamination rate is calculated after clearance by assigning an area of one sq m to a mine and then dividing the number of items found by area searched. The key point is that an area must be assigned to the dangerous item for the calculation to proceed at all.

Retrospective calculations of contamination rates after clearance may be useful for reporting, but they are of little interest for planning. REST technology offers the opportunity to give a prospective estimate of contamination rates (before clearance proceeds) with a higher accuracy than can be achieved using traditional technical survey methods. The estimate will necessarily be linked to the scale of the survey technique though REST operates at a broader scale than one sq m. But REST is not expected to find mines—its key advantage is that it identifies contaminated areas on a geographic scale that is compatible with traditional demining techniques, thus allowing effective planning and deployment of demining resources. For example, REST sampling is typically conducted in intervals of 100–200 m along roads. Using REST, 100 m of road can be cleared by a manual team in less than a day.

A properly run REST analysis system will include ongoing (daily) assessment of two key parameters: the “hit rate” (the identification of true positives) and the “false-alarm rate” (negative areas identified by the detectors as contaminated). Estimation of these parameters can be achieved using test filters made over known sources. If they are representative of the operational filters being analysed concurrently, the values returned from test filters provide accurate estimates of the reliability of the analysis of the operational filters. Something that is named “representative” is an object made containing the same general environmental odours (i.e., made in the same general area). “Positive” objects are made-over mines/UXO that are typical of the area (these are locally found items that have been re-laid in a trial area and left to “soak” for some time).

Thus, the measured hit and false-alarm rates from representative test filters can be used to estimate the number of true positive sectors in an operational area. Using a simple calculation, we can predict the number of filters we expect to be returned as positive:

> number of filters hit = (proportion hit x number of true positive sectors) + (proportion false alarm x number of true negative sectors)

Using this formula, we have developed software allowing a user to estimate a true positive rate by entering four key parameters: hit rate, false-alarm rate, number of filters analysed and number of filters returned as suspect positive. Example results for analysis of 1000 filters and return of 245 suspect positives are in Table 1.

| **Test sequence** | **Hit rate (percent)** | **False alarm rate (percent)** | **No False alarms** | **No true positives found** | **True contamination rate (estimate)** |
|---|---|---|---|---|---|
| **1** | 90 | 20 | 187 | 58 | 64 |
| **2** | 87 | 22 | 212 | 33 | 37 |
| **3** | 93 | 18 | 165 | 80 | 86 |
| **4** | 89 | 24 | 238 | 6 | 7 |
| **5** | 91 | 16 | 142 | 103 | 112 |
| **6** | 80 | 23 | 224 | 21 | 25 |
| **7** | 92 | 19 | 176 | 69 | 75 |
| **8** | 95 | 16 | 102 | 142 | 107 |
| **9** | 84 | 20 | 186 | 59 | 68 |

|  |  |  |  |  |  |
|---|---|---|---|---|---|
| **10** | 90 | 12 | 144 | 101 | 158 |
| **Average** | 89.1 | 19 | 73.9 | N/A | N/A |
| **Standard Error** | 1.40 | 1.15 | 14.19 | N/A | N/A |

Table 1: Projected results from regular testing of operational REST detectors through time. The percent values are cumulative results for separate test sequences. The three right-hand columns give the estimated returns assuming 1000 operational filters tested and 245 returned as suspect positive.

Statistical variance can be assigned to the contamination estimate using a series of test results. In Table 1, the cumulative results for hit and false-alarm rates are given for 10 test sequences, where one test sequence is 10 individual test events on all of the detectors currently being used. The overall datum set is therefore from 100 tests. If four tests are conducted on average each day, then two test sequences are completed each week and the datum set represents testing over five weeks.

The 95-percent confidence interval (CI) is a statistical measure providing a range projection that should encompass 95 percent of the means expected in the system of interest (any statistical text will explain the concept in more detail, or search for “statistical confidence intervals” on the web).[5] The 95-percent CI can be calculated using the descriptive statistics subroutines of any statistical package. Here, we used the descriptive statistics subroutine in Statistica®.

The key measure for estimating the 95-percent CI of the true contamination rate is in the last column. Here, the estimated true contamination rate is given for each set of test returns, as in the example above. From the 10 separate measures, the mean true contamination rate is 73.9 sectors (or filters) and the 95-percent CI (from Statistica®) is 41.8 to 106.0.

Converting to land area given one filter per 100 m of road and a five-m search width, the estimated contaminated area is 36,950 sq m with a range of 20,895 to 53,005 sq m. These values give the mean and range of area requiring clearance if all false alarms are removed from the data, represent the minimum possible clearance requirement from the current REST analysis, and provide an indicative assessment of true contamination in the area.

## Minimising Clearance Resulting From False Alarms

A REST analysis system will always return some false alarms, and therefore some unnecessary clearance requirements. If a suspect road actually has a very low contamination rate and the analysis centre is returning a relatively high rate of false alarms, then most of the clearance will be of false-positive sectors. We suggest here two ways to deal with the problem of false positives in order to optimise clearance requirements without compromising safety.

### *Method 1: Optimising the Number of Detectors*

The hit and false alarm rates on **test** filters can be calculated for individual detectors (dogs or rats are the only detectors in use today), or for the detection system (where the two parameters are calculated as cumulative rates). In Figure 1, the cumulative results for analysis of real test data returned by six Norwegian People's Aid (NPA)-Lubango dogs in May 2004 are ordered from left to right to show the progressive effects on detection success by sequentially adding dogs to the cumulative test results. Figure 1 can alternatively be viewed as sequential removal of dogs when read from right to left. The ordering of dogs is determined from individual hit rates, calculated from test results obtained over a few days.

Looking at Figure 1 from left to right, one can see that the

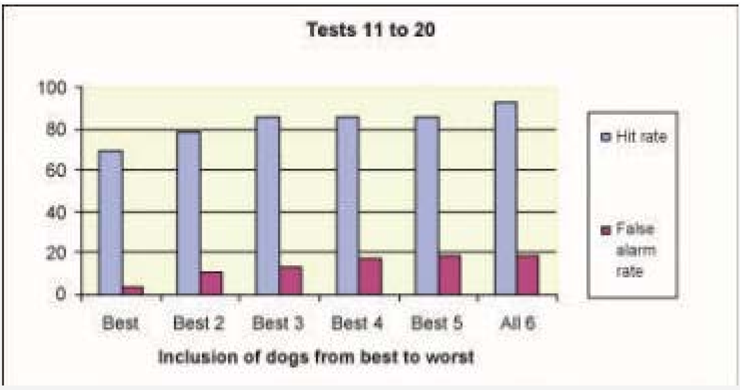

Figure 1: Cumulative hit and false alarm rates based on test data. Dogs were ranked from best to worst according to hit rate and progressively added from left to right to explore the additional effect of each dog on the overall result.

cumulative hit rate for all dogs was 93 percent. Removal of the worst dog results in a seven percent loss of positive filters because this dog uniquely identified one or several positives that all the other dogs missed. Removal of the next two dogs had no effect on the cumulative hit rate because they found no positives uniquely. With respect to the false alarm (FA) rate, all six dogs returned 18.9 percent. Removal of the worst dog had no effect, removal of the second-worst dog reduced the FA rate to 17.9 percent, and removal of the three worst dogs reduced the FA rate to 13.3 percent. It therefore seems reasonable to remove the three worst dogs from the analysis.
This example uses test data both to place the dogs in order and to explore the effects of that ordering on detection success. But where operational filters are being analysed, the test data would only be used to place the dogs in order. Once that order has been determined, then the analysis in terms of costs and benefits would be applied to the results from operational filters. The decision about how many dogs to use would be based on the trade-off between benefits (reduction in clearance) and costs (increased in missed true positives) as each dog was progressively removed.

The specific aim is to minimise the false alarm rate with a minimum of compromise on detection success. In this example, three dogs were removed for a seven percent cost as loss of detected positives and a 5.6-percent benefit as reduction in false alarms. The benefit is more substantial than these values suggest: if the number of operational filters sampled was 1000, the result would be a reduction of 56 sectors (189–133), which is a significant clearance requirement. The cost in terms of number of excluded positives will depend on the true contamination rate.

A key component of this analytical approach is its flexibility. Because the test results are being produced daily, it is possible to produce graphs such as Figure 1 from very short periods—a few days to a week. It is therefore possible to make decisions about the number of dogs to use for the current operational analysis on a very fine time scale. For example, the data may indicate that all six dogs should be used for the operations filters analysed in Week 2, but the four best dogs should be used for the filters analysed in Week 3.

Working to such a narrowly framed time scale requires that test and operations results be entered and updated daily. We have already produced analysis software that supplies the required graphs as the data are entered (Figure 1 was produced by that software).

### *Method 2: Re-sampling Suspect Positive Sectors*

The usual procedure is for the REST analysis team to return a sector identified as suspect positive to the clearance teams, who then search that area—a time-consuming task. However, some of the returned sectors will be false alarms and represent an unnecessary clearance cost. To our knowledge, no REST team has ever re-sampled the returned suspect-positive sectors in an attempt to identify at least some of the false positives. Here, we consider the benefits and costs of such an approach using a simple model. The three key parameters are the same as were described above:

- **The true rate of contamination with mines/UXO** (i.e., the number of sectors that should be returned as positive if perfect detection is achieved).
- **The hit rate on positives given by the detection system**, which is measured weekly or even daily using representative test filters.
- **The false indication rate given by the detection system** also measured as above.

The **benefit** of re-sampling is a reduction in the number of sectors that require clearance, because some of the areas returned as (false) positive in the first analysis will be declared (true) negative after the second analysis. The **cost** of re-sampling is an increased number of missed true positives. The additional sampling and analysis represent an economic cost, but that can be traded off against the benefit of reduced clearance requirements.

In the examples in Figures 2 and 3, it is assumed that 50 km of road was searched. In both figures, adjustment for different lengths of road is easily achieved by ratio calculation. Thus for 100 km, the Y-axis of Figure 2a ranges from 0 to 50, and for Figure 2b, 0 to 10.

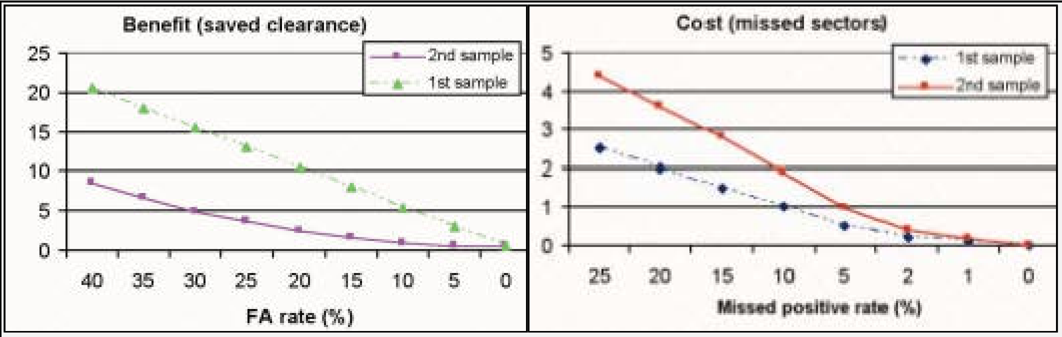

Figure 2: Benefit and cost for re-sampling 50 km of road with a true contamination rate of 0.01 (1 percent). FA = false alarm.

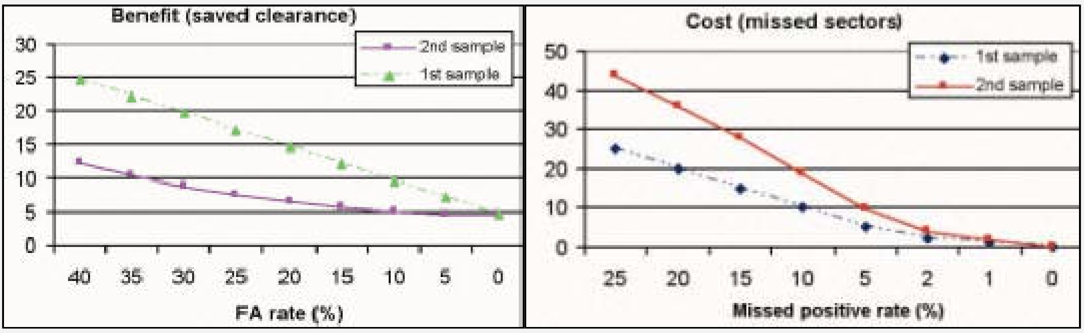

Figure 3: Benefit and cost for re-sampling 50 km of road with a true contamination rate of 0.1 (10 percent). FA = false alarm.

Figure 2 shows benefits and costs for a true contamination rate of 0.01 (one percent = low). Figure 3 shows a true contamination rate of 0.1 (10 percent = medium/high). In each figure, the benefit or cost can be seen in the difference between the two lines. In low-density minefields, there is significant benefit to be gained from re-sampling positives, even at relatively low false-alarm rates. Cost in terms of increased missed true positives is small. At the higher true contamination density, the additional missed sectors become a significant concern at higher missed positive rates, although benefit in terms of saved clearance is still high. Ultimately, the decision to use a re-sampling approach will be made by the project managers, and will be driven by risk-assessment issues. These results suggest that re-sampling offers considerable benefit for very little cost, as long as true contamination rates are low.

Sending the sampling team back to a previously sampled area may be logistically difficult if they have moved on (e.g., to a new camp) in the period between sampling and the return of results from the analysis centre. A rapid turn-around of filters at the analysis centre will help to minimise this problem. The option of sampling two filters at one time and storing the second filter for future analysis, if needed, does not give a truly independent sample and is not recommended.

We note that there is no need for field personnel to calculate true contamination rates in order to use the principles outlined here. The required calculations are built into the analysis software used in the analysis centre, allowing straightforward linking to versions of the graphs.

## Conclusions

Since its earliest use for mine detection, REST analysis has been recognised as a tool with potential for providing rapid area reduction. However, little attention has been paid in the past to optimising its use, or to minimising the clearance costs arising from detection errors.

A part of the optimisation process is to ensure that effective and regular quality assurance (QA) is undertaken during analysis. The results of that QA testing can be used to provide survey information, to adjust the results of analysis and to feed back into clearance requirements. The approach outlined here represents a significant technical refinement of REST analysis over current use, without requiring anything more from an analysis centre than effective QA and regular data entry.

When linked to automated analysis software, which is already available, these refinements indicate that REST is becoming a sophisticated area-reduction tool that should be considered for use by mine-clearance agencies worldwide.

**All figures courtesy of the author.*

### Acknowledgements

The ideas presented here have evolved as a direct result of our involvement with the REST programme run by NPA in Angola. Thanks to Sigbjorn Langvik, Geir Bjersvik, Rune Fjellanger, Andolosi Sanjala, Al Cummings, Johannes Dirscherl, Ian Mansfield and Havard Bach for many useful comments, suggestions and general support. Funding of this work was by NPA and the Department for International Development (UK).

### Endnotes

1. Fjellanger, R. 2003. "The REST concept." *Mine Detection Dogs: Training, Operations and Odour Detection*. GICHD, Geneva, 53–108.
2. McLean, I.G., Bach, H., Fjellanger, R. and Aakerblom, C. 2003. “Bringing the Minefield to the Detector: Updating the REST Concept.” Proc EUDEM-2 SCOT International conference on requirements and technologies for detection, removal and neutralization of landmines and UXO, Brussels, 156–165.
3. Joynt, V. 2003. “The Mechem Explosive and Drug Detection System (MEDDS).” *Mine Detection Dogs: Training, Operations and Odour Detection*. GICHD, Geneva, 165–174.
4. GICHD. 2004. “A Study of Mechanical Application in Demining.” GICHD, Geneva.
5. Simonson, K.M. 1998. “Statistical Considerations in Designing Tests of Mine Detection Systems: I—Measures Related to the Probability of Detection.” Sandia: SAND98-1769/1 and "II—Measures Related to the False Alarm Rate." Sandia: SAND98-1769/2.

## Contact Information

Ian McLean

Researcher, GICHD
7bis Avenue de la Paix
CH-1211 Geneva 1
Switzerland
Tel: +41-22-9061676
Fax: +41-22-9061690
E-mail: i.mclean@gichd.ch
Website: http://www.gichd.ch/

Rebecca Sargisson
Researcher, GICHD
Tel: +41-22-9061658
Fax: +41-22-9061690
E-mail: r_sargisson@yahoo.co.uk