

October 2010

The Effect of Reinforcement Rate Variations on Hits and False Alarms in Remote Explosive Scent Tracing with Dogs

Rebecca Sargisson
University of Waikato

Ian McLean

Follow this and additional works at: <http://commons.lib.jmu.edu/cisr-journal>

 Part of the [Other Public Affairs, Public Policy and Public Administration Commons](#), and the [Peace and Conflict Studies Commons](#)

Recommended Citation

Sargisson, Rebecca and McLean, Ian (2010) "The Effect of Reinforcement Rate Variations on Hits and False Alarms in Remote Explosive Scent Tracing with Dogs," *The Journal of ERW and Mine Action* : Vol. 14 : Iss. 3 , Article 27.
Available at: <http://commons.lib.jmu.edu/cisr-journal/vol14/iss3/27>

This Article is brought to you for free and open access by the Center for International Stabilization and Recovery at JMU Scholarly Commons. It has been accepted for inclusion in Journal of Conventional Weapons Destruction by an authorized editor of JMU Scholarly Commons. For more information, please contact dc_admin@jmu.edu.

The Effect of Reinforcement Rate Variations on Hits and False Alarms in Remote Explosive Scent Tracing with Dogs

Cover Page Footnote

The Effect of Reinforcement Rate Variations on Hits and False Alarms in Remote Explosive Scent Tracing with Dogs

Detection animals offer untapped potential in terms of locating landmines and explosive ordnance in the field and in the laboratory. In this study, the Geneva International Centre for Humanitarian Demining investigated the effect of low, medium, and high levels of reward on the performance of six dogs searching filters for explosive odor.

by Rebecca J. Sargisson [University of Waikato] and Ian G. McLean [Consultant]

Remote Explosive Scent Tracing—Or Odor Capture—is a detection process in which odor is captured on an absorbent filter and analyzed by a detector, such as a dog or rat.^{1,2} The detector works in a safe and controlled environment and is capable of searching large areas of ground in a short period. Odor capture has a wide range of potential applications (for example, the detection of oil-pipeline leaks and the detection of cancer or tuberculosis), but with respect to explosive detection, REST's main value is eliminating road sections that do not contain explosive ordnance, allowing clearance to proceed more rapidly than is possible using most standard detection technologies.

	Indicate	Ignore
Positive	Correct <i>Hit</i>	Incorrect <i>Miss</i>
Negative	Incorrect <i>FA</i>	Correct <i>CR</i>

Table 1: Matrix of outcomes in a REST task.
All graphics and photos courtesy of the authors.

REST will only be used if it can deliver consistently-high detection reliability for filters containing explosive odor (hits on “positive” filters). However, as a key use of REST is for uncontaminated land release, REST must also deliver reliable decisions on filters not containing explosive odor (correct rejection of “negative” filters). A filter analysis produces four possible outcomes (See Table 1), of

which two are undesirable—“miss” and “false alarm.” A miss means that explosive ordnance is undetected, presenting a danger to future land users. A false alarm means unnecessary additional work for the mine-clearance program. Low reliability on either of these outcomes reduces confidence in REST as a detection technology.

The typical procedure is summarized as follows. A team uses a suction pump to vacuum the air over a road section, typically 100 or 200 meters (109 or 218 yards) long and about 5 meters (5 yards) wide. The air is sucked through a filter, and careful records are kept of the road section that each filter represents. The filters are transferred to a laboratory where they are presented to trained detectors (usually dogs or rats) using a standard methodology, such as on the arms of a carousel (Figure 1) or in a line of stands (Figure 2).

The dogs are trained using filters made from controlled odor sources (“benchmark filters”). For training mine detection, most REST agencies plant test minefields, noting each mine’s location, type and depth. Filters can then be made in areas that should be contaminated with explosive odor from a known source, and areas treated as free of explosive odor. With a variety of odor sources used, it is assumed that background odor is consistently variable across filters, and the detectors must therefore use the explosive odor’s presence or absence as the determining variable in their analysis. A key benefit of REST analysis over field-based animal-detection systems is that benchmark filters can be mixed in with operational filters, allowing the continuous monitoring of each detector’s reliability during operational analysis.

All REST agencies use a training system in which hits on positive benchmark filters are reinforced, typically using a toy or food. Correct rejections of negative filters are not reinforced because they do not provide a discrete behavioral unit (the detector moves past the

negative filter without being rewarded for its correct “response”). This training methodology potentially introduces **response bias**, most likely as a tendency to give an indication response on a negative filter (a false alarm). Thus, the training procedure itself may be a source of false alarms, limiting the agency’s ability to attain the objective of minimizing false alarms while maintaining a reliably high hit rate.

Signal-detection theory³ gives the issues and principles discussed above detailed technical analysis, and we use that theory’s language in this paper. With respect to REST’s two objectives of maintaining high hit and low false-alarm rates, the theory distinguishes two processes affecting accuracy:

- **Sensitivity:** The dog’s ability to discriminate between positive and negative filters can be improved in a variety of ways, including increasing the overall reinforcement rate for correct responses.⁴
- **Response bias:** If the training or operational experiences have asymmetries (such as only rewarding responses to positive filters during training or more abundant negative filters than positive filters, which is expected for operational filters), then response asymmetries are also expected.^{5,6} Under signal-detection theory, all filters contain a background odor (noise).

Positive filters should carry an additional odor from the explosive ordnance (signal-plus-noise).^{3,7} A filter’s signal strength can be placed somewhere in the area under two normally distributed Gaussian functions plotting signal intensity as a function of that odor’s probability of being present (Figure 3 on page 66). Signal availability to the left of line “C” will result in an “ignore” response (filter is negative), whereas signal availability to the right of C will result in an “indication” response (filter is positive). Sensitivity (d') is determined by the separation between the peaks. Greater separation should result in greater accuracy because positive filters are less easily confused with negative.



Figure 1: Dog searching filters in a carousel-style presentation system.

Signal-detection theory assumes that each animal responds according to a response criterion (the vertical line C in Figure 3 on page 66). An animal’s responses can become biased toward one response type if more reinforcement is made available for one response type over another or if unequal numbers of positive and negative filters are presented.⁶

Signal-detection theory makes the following predictions:⁸

- If the sensitivity of the detector (d') varies and the response criterion (C) remains constant, hit rate and false-alarm rate should be negatively correlated; i.e., as the functions move apart, hit rate will increase, and false-alarm rate will decrease.
- If the response criterion (C) varies, hit rate should be positively correlated with false-alarm rate. For example, if a detector is biased toward indicating, it will hit more positive filters, but will also indicate more negative filters, creating a high false-alarm rate.
- A strong correlation between hit and false-alarm rate would be a useful finding for REST.
- If hit and false-alarm rates were positively correlated, the relationship between them could be optimized by manipulating reinforcement bias, filter ratios, or the experimental method.
- If hit and false-alarm rates were negatively correlated, the training approach could focus on increasing hit rate, with the desired low false alarm-rate achieved without explicit training.

The present experiment used data from the regular training of six REST dogs in Angola to explore the relationship between hit and false-alarm rates. The overall reinforcement rate for positive-filter hits was manipulated across 28 weeks of a calendar year, according to Table 2 on page 66. The proportion of negative filters was held constant (between 94 and 99 percent of filters presented were negative).

It was expected that hit rate and false-alarm rate would be correlated. Given that only reinforcement for hits was varied, increasing reinforcement availability for hits could have produced a bias toward indicating, producing a positive correlation between hit and false-alarm rate. If, however, the reinforcement-rate manipulation for hits altered the dog’s sensitivity to the signal, we would expect a negative correlation between hit and false-alarm rate. In other words, increasing reinforcement for hits would either have been expected to cause a bias toward indicating or to improve the dog’s ability to discriminate between positive and negative filters.

Method

Subjects. Six male non-neutered dogs, aged between 6½ and 7½ years, with several years of previous REST training participated. Five were Labrador Retrievers (Retzina, Stavros, Tan, Zante and Zulu) and one was a Springer Spaniel (Rusty). Each dog was assigned an experienced Angolan dog handler. The dogs were exercised six days a week by walking and swimming, housed in individual kennels, given free access to water, fed a high-quality dry dog food in sufficient quantities to maintain a healthy weight, and were not food-deprived.



Figure 2: Dog indicating a filter in the line-stand presentation system.

Apparatus. Filters were placed on a carousel apparatus (Figure 1). The carousel was a large stainless-steel wheel, mounted horizontally to the floor, which could be rotated. Filters were mounted horizontally at the ends of 12 arms that were removable for cleaning. The rooms’ walls were concrete block, and tiled floors minimized odor contamination. A stainless-steel screen inside the rooms shielded a supervisor from the searching dog. All other personnel (the dog handler and documenter) watched activities from adjacent rooms through internal one-way glass windows.

The filters were a PVC core wrapped in mosquito netting and housed inside a PVC tube (known as the “Mechem” filter, named for the manufacturer).

Procedure Sampling. Unused filters were contaminated with air to produce positive filters (filters believed to contain the odor from one or more landmines) and negative filters (filters believed to be free of explo-

Weeks	Reinforcement Level	Percentage of Hits Reinforced
6-10	Low	20-30
11-27	Medium	35-50
28-33	High	60-75

Table 2: Experimental conditions.

sive odor but containing other neutral odors from similar locations). Air was added to the filters by placing the filters at the end of a long stainless-steel tube subject to continuous suction via a vacuum-pump machine worn as a backpack. The filter was held close to the ground and swung to the left and the right of the pump operator as he slowly walked a 100-meter distance. Filters were considered positive if the pump operator passed within 1 meter of a buried landmine and negative if no landmines were present within 100 meters of the filter during sampling. The landmines were a range of anti-tank and anti-personnel mines commonly found in Angola. The mines were laid between 0 and 10 centimeters (0-4 inches) beneath the ground surface for a minimum of six months before they were used for sampling. A total of 275 mines were available for sampling. All sampled filters were stored inside small PVC containers, and positive filters were stored separately from negative filters until analysis to avoid odor cross-contamination.

Analysis. The dogs searched filters on the carousel between 8 a.m. and 1p.m., Monday through Friday, taking rest breaks when required. After preparation of the carousel, each dog was brought to the carousel room's door in a sequential but random order. When the dog was calm, the handler instructed the dog to "search," and the dog handler stepped behind a wall out of the dog's view. The dogs walked unaccompanied, off-lead, in an anti-clockwise direction around the carousel, sniffing each filter consecutively. The dog exited the room after it had correct-

ters because the reinforcer for blank runs was not contingent upon a discrete response, such as sitting. Zero to three positive filters were present on the carousel among the remaining negative filters.

After the summer break, training recommenced for all six dogs in Week 2 of 2005 and continued for four weeks before experimental manipulations. At this point, reinforcement frequency for correct indications on positive filters was manipulated by providing a reinforcer, such as a click from the clicker and food or access to a ball, on only **some** correct indications (intermittent reinforcement). This can be contrasted with earlier training stages where reinforcing **every** correct indication is common in order to aid learning (continuous reinforcement). All other variables were held constant, including the number of negative filters available on the carousel, and reinforcement for correct rejections of negative filters.

Table 2 shows the experimental conditions. From Weeks 6 to 10, hit reinforcers were held at a "low" level (20 to 30 percent of hits were reinforced), from Weeks 11 to 27 at a "medium" level (35 to 50 percent) and from Weeks 28 to 33, at a "high" level (60 to 75 percent of hits were reinforced).

Results

A decision for each filter from each dog was obtained. Signal-detection theory terminology was used to define the four analysis results possible for a filter: hit (indication on a positive filter), miss (no indication on a positive filter), false alarm (FA, indication on a negative filter) and correct rejection (CR, no indication on a negative filter). Hits, misses, false alarms, and correct rejections were summed for each week for each dog and used to calculate hit rates [(hits / (hits + misses) * 100] and false-alarm rates [(FAs / (FAs + CRs) * 100].




Figure 3: Hypothetical noise and signal-plus-noise distributions in a sensory discrimination task according to signal-detection theory. The left panel demonstrates discriminability (d') as the distance between the means of the two functions. The right panel illustrates the animal's response criterion (C), which dissects the two functions and can shift to the left and right as a function of response bias.

ly indicated a positive filter by sitting next to it and hearing the conditioned reinforcer (clicker), or when the dog handler called it from the room. Reinforcement was occasionally available for hits (indicating a known positive filter). The reward most often delivered was small pieces of dry dog food and sometimes access to a ball or squeaky toy. A reward was occasionally delivered following a "blank" run (a run containing only negative filters), if the dog correctly ignored all filters. However, the reward may not have acted to reinforce correct responses to negative fil-

Figure 4 shows hit and false alarm rates for all individual dogs, and for the mean across all dogs, as a function of week. When actual reinforcement rates were found to deviate from planned reinforcement rates, these data were removed, and are therefore missing from Figure 4. Pearson correlation coefficients were used to test the relatedness of hit rate to false-alarm rate shown in Figure 4. A significant, negative correlation appeared between mean hit rate and mean false-alarm rate ($r = -.72$, $p = .000$). The correlation between hit and false-alarm rate was also negative

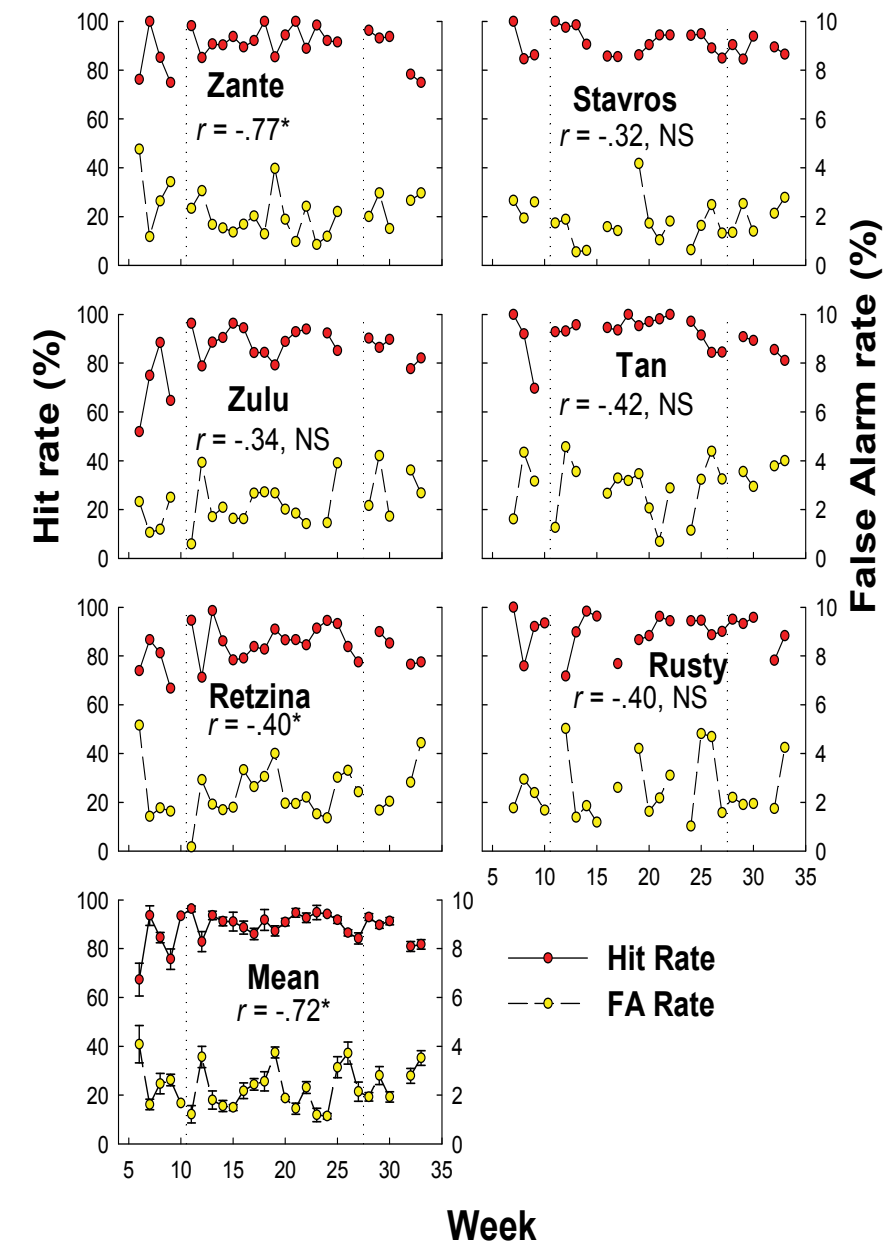


Figure 4: Hit (red circles) and false-alarm (yellow circles) rates calculated as percentages for each week for all six dogs and for the mean across dogs. Vertical dotted lines show changes in reinforcement level for hits from low, to medium, to high from left to right across the x-axis. Pearson correlation coefficients are given for each dog, and for the mean, and are significant ($p < .05$) unless shown (NS).

for all individual dogs and significantly so for two of the six dogs. All r values are shown in Figure 4. Figure 5 (on page 68) displays the data used to calculate the mean correlation and clearly shows a strong negative relationship between hit and false-alarm rate, in that, as hit rate increases, false-alarm rate decreases.

Weekly hit and false-alarm rates for each dog, and for the mean, were grouped according to reinforcement-rate condition (low, medium, and high). These data are shown in Figure 6 (on page 68). A one-way analysis of variance indicated that hit rates in the three groups differed significantly [$F(2, 15) = 5.34$, $p < .05$]. A Fisher's LSD post-hoc test⁹ showed that the medium and high reinforcement rates produced significantly higher hit rates than the low reinforcement rate condition, but that the medium and high conditions did not differ significantly from one another in terms of hit rate. No significant difference in false-alarm rates were found across the three reinforcement

conditions [$F(2, 15) = 0.89$, $p > .05$]. However, Figure 6 (on page 68) shows that false-alarm rate was lowest during the medium-reinforcement rate condition for four of the six dogs, and for the mean.

Discussion

Hit rate and false-alarm rate were overall significantly negatively correlated.

Thus, as hit rate increased, false alarms decreased. According to signal-detection theory, these negative correlations are to be expected if the distance between the noise peaks and the signal-plus-noise functions changed. In other words, the correlations between hit and false-alarm rate were caused either by changing discriminability between positive and negative filters, or by changing the dog's sensitivity to the odor, and not by changing response bias (decision criterion). Given that the filters' discriminability was not manipulated, the likely reason for the negative correlation between hit and false alarm rate was the dog's increasing sensitivity due to changes in the overall reinforcement rate for hits.

This result suggests that the experimental method's nature, reinforcing hits and not correct rejections, does not produce changes in the dog's response bias. In other words, greater reinforcer availability for hits did not cause a bias toward indicating. Instead, in the present experiment, low reinforcement rates for hits produced poorer performance on negative and positive filters, while medium and high reinforcement levels produced more accurate responses on both filter types. In the present experiment, performance peaked under the medium level of hit reinforcement. Increasing the reinforcement frequency beyond this medium level did not result in greater accuracy on positive or negative filters. One implication of this finding is that procedures to improve the REST system's accuracy should focus on increasing the animals' hit rates, and that any hit rate increase will be accompanied by a false-alarm rate decrease.

Manipulating reinforcement ratios is one way to alter an animal's response accuracy. Another way is through the experimental procedure itself. The current procedure was a "go/no-go" procedure, whereby animals indicated, by sitting, the presence of explosive odor but made no response to filters containing no explosive odor. Such a procedure producing a bias toward indicating, rather than ignoring, is possible because ignoring is not explicitly reinforced. Alternatively, due to the greater numbers of negative filters (between 94 percent and 99 percent of filters were negative), the dog's behavior could become biased toward ignoring because it is the most frequently-required response. An analysis of bias, using [$\log b = \frac{1}{2} \log (FA / Hits)(CR / Miss)$], showed




Figure 5: Mean hit rate as a function of false-alarm rate. A straight line has been fit to the data to illustrate the pattern represented by the datum points.

that the behavior of four of the six dogs was biased toward indicating, and this bias strength decreased as reinforcement for hits increased for all six dogs. The behavior of two dogs was biased toward ignoring, and this bias was unaffected by reinforcement-rate manipulations. Thus, the present procedure appeared to not produce consistent effects on response bias, nor did it produce bias in one direction over another. Instead, each dog tended to maintain a fairly reliable preference for either indicating or ignoring, and biases toward indicating were counter-intuitively reduced by increasing reinforcement availability for correct indications.

REST programs should include ongoing monitoring of response bias, so they can redress any imbalance. Manipulation of reinforcement rates can eliminate response bias more easily in procedures where responses to positive and negative filters are directly reinforced. In procedures where responses to only one type of filter are reinforced, such as in the present REST system, response bias may be eliminated by careful manipulation of the ratio between positive and negative filters. REST programs should seek to determine the optimum ratio for their procedure and animals, and maintain this ratio while continuing to monitor ongoing response bias.

Other factors which affect the overall accuracy of animals' responses concern the quality of the samples. Sampling can be optimized in terms of filter material, climatic condition, avoidance of contamination, and so on. Once collected, filters should be handled to minimize cross-contamination. By maintaining as clear a signal on the filter as possible, the animal is given the best chance to obtain high hit rates.

see endnotes page 82

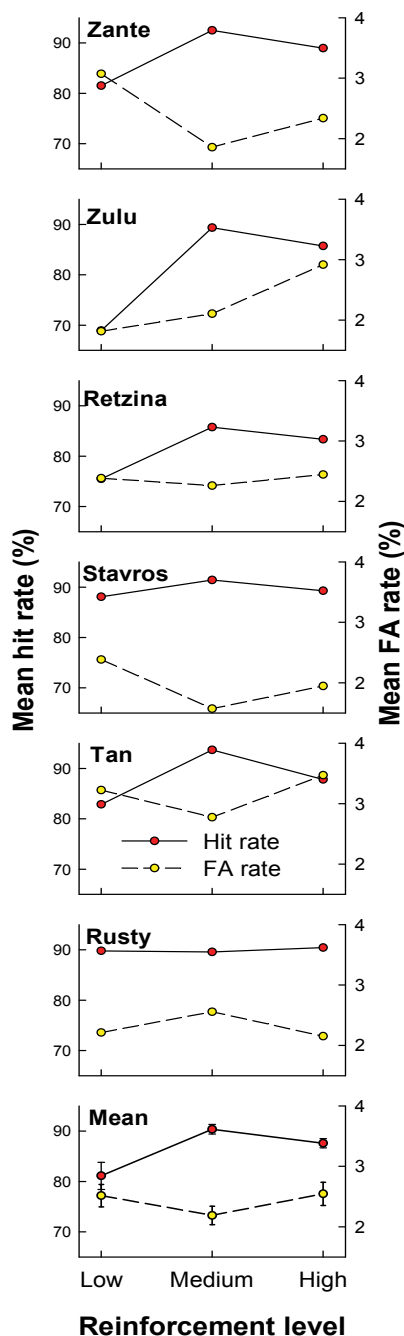


Figure 6: Mean hit (red circles) and false-alarm (yellow circles) rate for each dog and for the mean in each of the three reinforcement conditions (low, medium, and high).

Author note: The authors conducted this research while employed by the Geneva International Centre for Humanitarian Demining. We thank members of the REST team in Angola, especially Andolosi Sanjala and Felisberto Joao, Birgitte Lauritzen and Rune Fjellanger for their help. Norwegian People's Aid, and the government of Switzerland through a grant to GICHD, funded the research.



After completing a Ph.D. in psychology, **Rebecca J. Sargisson** was a Research Consultant at the Geneva International Centre for Humanitarian Demining from 2003 to 2006 working on many aspects of the use of dogs in demining. Sargisson is currently employed by the University of Waikato, New Zealand. She remains interested in dog research but is also researching issues related to children's play and playground design.

Rebecca J. Sargisson
Senior Tutor
University of Waikato
Private Bag 2105
Hamilton 3240 / New Zealand
Tel: +64 7 856 2289
Fax: +64 7 838 4300
E-mail: sargisson@waikato.ac.nz
Website: http://waikato.ac.nz



Ian G. McLean worked at the Geneva International Centre for Humanitarian Demining, conducting research on landmine-clearance systems, studying the environmental influences on demining and developing the Remote Explosive Scent Tracing system. McLean has taught environmental policy and wildlife management at the Universities of Otago and Waikato in New Zealand, and is currently raising his two children and consulting on environmental issues.

Ian G. McLean
Tel: +64 7 544 9703
E-mail: tawakix@hotmail.com