Bayes-optimal detection of TNT content by nuclear quadrupole resonance

Sven Aerts[1,*], Dirk Aerts[1,†], Franklin Schroeck[2,‡], and Jügen Sachs[3,♯]
[1] *Centre for interdisciplinairy studie CLEA, FUND Department of Mathematics, Vrije Universiteit Brussel, Triomflaan 2, 1050 Brussel, Belgium*
[2] *Department of Mathematics, John Greene Hall, 2360 S. Gaylord St., University of Denver, Denver, USA and*
[3] *Electronic Measurement Lab, Technische Universität Ilmenau, Germany,*
[*] *saerts@vub.ac.be,* [†] *diraerts@vub.ac.be ,* [‡] *fschroec@du.edu,* [♯] *juergen.sachs@tu-ilmenau.de*

We study the statistical performance and applicability of a simple quantum state discrimination technique for the analysis of data from nuclear quadrupole resonance experiments on a TNT sample. The target application is remote detection of anti-personnel landmines. We show that, even for data that allows the determination of only one time dependent component of the NQR subsystem, the use of the Bayes optimal detector leads to greatly improved ROC curves with respect to the popular demodulation technique, especially for spin echo signals with a low signal to noise ratio. The method can easily be extended to incorporate results from other sensing modalities and the incorporation of informationally complete measurements that estimate the full density matrix of the NQR subsystem.

PACS numbers:

INTRODUCTION

In this report we address the problem of optimally deciding whether a given closed volume contains TNT content by means of nuclear quadrupole resonance (NQR) measurements. Nuclear quadrupole resonance signals result from the relaxation of nuclear quadrupole momenta to their original thermal equilibrium position after an initial, high power RF pulse has been applied. The thermal equilibrium configuration of the nuclear spins is a function of the electromagnetic field in the vicinity of the quadrupole active nuclei. As a result, the NQR spectrum is very specific with respect to chemical compounds in the substance involved and can serve as a fingerprint to identify that substance. Because of its high potential value in remote explosive detection, there is renewed interest in NQR methods for landmine and UXO detection, as well as for securing high risk areas such as airports by non-intrusive means. Quite incidentally, NMR and NQR systems have recently also received great attention for their applicability in the fast growing field of quantum information [12]. In fact, state of the art quantum computers are currently based on the principles of NMR. Because the relaxation of the nuclear spin is quantized, one would conjecture that a quantum statistical analysis of the data is optimal [4], [6] and [10]. However, the macroscopic bulk size and consequently large number of spins necessary for an appreciable signal strength, as well as the far from absolute zero temperature of the sample in realistic conditions, call for a classical statistical approach. It has been shown that the principle of Bayes optimal observation [1] is effective in both the quantum and classical domain. If only one particular kind of measurement is performed, the mathematical analysis will be identical for the quantum and classical cases. We will briefly introduce both Bayes-optimal detection and the popular demodulation technique and show by means of ROC curves that Bayes-optimal detection offers a vast improvement over the latter, especially for a very low signal to noise ratio, as in the case of NQR based TNT detection.

LANDMINES AND NUCLEAR QUADRUPOLE RESONANCE

The detection of landmines turns out to be an extremely difficult task. Even though more and more landmines are made of plastic, the bulk of landmines is still detected using metal detectors. The reason for this, is that all landmines contain at least a small amount of metal content in the detonator and the metal detector gives a clear signal that is trusted by field workers. By increasing the sensitivity of the metal detector it is possible to reliably detect landmines. The problem is that the increased sensitivity will make the metal detector responsive to other metal objects that abound in postwar territory. The large false alarm rate that is accompanied by the increased sensitivity, results on average in 500 to 1000 objects to be wrongly classified as potential mines, for each real mine encountered. The overhead in time, energy and cost, not to mention the high rate of accidental detonation of real landmines as a result of this enormous high false alarm rate, has spurred the search for a better classification method of the detector signals. This classification of detected signals is made more difficult by the enormous variety of mines, soil parameters, vegetation and weather conditions. A possible solution involves the use of nuclear quadrupole resonance (NQR) techniques. A necessary condition for the use of NQR, is the presence of a substance with a nuclear electric quadrupole moment. An ideal candidate is the naturally stable nitrogen isotope ^{14}N, (with a nat-

ural abundance of 99.64 %) with nuclear spin 1 and corresponding nuclear quadrupole moment. All known explosives contain ^{14}N, so that, in principle, it is possible to detect any non-metallic mine by NQR [14]. The NQR spectrum for ^{14}N has transitions in the frequency range between 0 and 6 MHz, actual values depending mostly on the electric field gradient tensor, which is primarily determined by the charge distribution of the electrons that bind the nitrogen to the rest of the explosive. The resulting NQR signal is therefore highly dependent on the chemical structure of the sample, and delivers a potentially very reliable classification with an accompanying very low false alarm rate. Compared with other popular mine detection techniques such as the metal detector and the ground penetrating radar, NQR-based detector performance is not very sensitive with respect to weather conditions. Add to this that it is possible to construct a hand held NQR detector, and it seems that NQR is an ideal candidate for explosive detection. The main challenge for NQR techniques, is the inherently low energy content of the signal, resulting in a very low signal to noise ratio (SNR). To improve the SNR, many repetitions of the experiment are necessary. Rather than just measuring the free induction decay of a single excitation, one can set up an appropriate sequence of RF pulses, and measure the returned echo after each such pulse. In this way we obtain a larger data set from which inferences can be made. However, the rate at which the repetition is physically possible, is bound from below in a fundamental way by the physical parameters of the relaxation process. The nuclear relaxation is a result of two different mechanisms, called the spin-spin relaxation and the spin-lattice relaxation. The relaxation time that characterizes the spin-lattice relaxation, usually denoted T_1, determines the time necessary for the system to regain its original thermal equilibrium state, and gives a bound on how quickly a pulse sequence can be initiated after another. The spin-spin relaxation time, denoted T_2, is indicative of the decoherence as a result of spin-spin interactions and determines the length of the spin echo sequence. Spin-spin relaxation times are generally (much) shorter than spin-lattice relaxation times. In practice, we can apply a pulse sequence of length T_2, and repeat this pulse sequence every T_1. For most explosives, the relaxation times are short enough so that NQR detection becomes feasible. Unfortunately, about 60% of the landmines contain $\alpha-$trinitrotoluene (TNT), which has relaxation times that lead to prohibitively long detection times for efficient NQR detection within the operational limits of landmine detection. It is therefore projected that an NQR based landmine detector will probably serve mainly as a confirmation sensor, i.e. a detector that is employed to decrease the false alarm rate only after a metal detector or a ground penetrating radar system has detected a potential landmine. Whether used as a confirmation or as a primary detector, NQR detection efficiency for TNT will benefit from a reduction in the time necessary for reliable detection. Because one cannot shorten the relaxation parameters of TNT, much effort has gone into cleverly designing the emitted RF pulse and increasing the sensitivity of the receiver. Besides these efforts, it is worthwhile to pursue better signal analytic detection techniques.

QUANTUM OPERATIONS AND THE EVOLUTION OF THE NQR SIGNAL

It is not feasible to describe the entire quantum-physical state of the landmine, nor would this be interesting. What causes the NQR signal, is the change in the magnetization along the direction of the solenoid. In the case of ^{14}N, we are dealing with a spin-1 system so that the relevant quantum mechanical subspace is spanned by just three orthogonal vectors. A full determination of the state in this three-dimensional subspace could, in principle, lead to efficient (and provable optimal) strategies for detection and classification of the NQR signal. However, even the full characterization of the state in just this three dimensional subspace is not feasible with a single pulse sequence, as is the case for our data here. However, we do not necessarily need detailed knowledge of the whole state. The mathematical formalism required for optimal distinction between arbitrary quantum states can easily be simplified to accommodate our limited knowledge about the state. We will briefly show how quantum operations can serve as a framework to relate the measured quadrature components of the current in the coil to quantum state discrimination tools. In theoretical descriptions of NQR ([5], [9], and [11]), the state of the system is a classical statistical mixture of pure quantum states, described by a density operator ρ belonging to the class of linear, positive operators that sum to one when they act upon a complete set of eigenvectors. If we consider as system the landmine, its immediate surroundings, and the NQR detector, the detection system can be considered as closed and the dynamics of the total density operator ρ_{closed} is governed by the unitary evolution that solves the Schrödinger equation

$$\frac{d\rho_{closed}(t)}{dt} = -\frac{i}{\hbar}[\mathcal{H}, \rho_{closed}(0)] \tag{1}$$

Here $\rho_{closed}(0)$ is the initial density operator and $\mathcal{H} = \mathcal{H}_{rf} + \mathcal{H}_Q$, with $\mathcal{H}_Q$ is the nuclear quadrupole Hamiltonian, $\mathcal{H}_{rf}$ the Hamiltonian corresponding to the RF pulse and [,] is the commutator. The strength of the quadrupolar Hamiltonian depends mainly on the coupling between the electric field gradient (EFG) and the quadrupolar moment. The efficiency of the excitation by an RF field depends on the relative orientation between the incident radiation and the EFG principal axis frame. Because the EFG principal axis frame depends on

the molecular orientation, it is not possible to excite all quadrupole levels with the same efficiency in a powder crystalline sample. A calculation shows that the signal strength resulting from a crystalline powder is approximately only 43% the strength of a signal stemming from a single crystal with the same number of NQR active nuclei [9]. In absence of the RF pulse, a canonical ensemble of NQR spin-1 systems at temperature T, is described by a density operator

$$\rho = \frac{\exp(-\mathcal{H}_Q/kT)}{Tr\exp(-\mathcal{H}_Q/kT)]} = \frac{1}{Z}(1-\mathcal{H}_Q/kT) + O(\frac{1}{T^2}) \quad (2)$$

Here k is the Boltzmann constant, T the temperature in Kelvin, and Z the partition function, which acts as a normalization. The RF pulse perturbs the thermal equilibrium and it is the relaxation from this perturbed state to Eq..(2), according to Eq..(1), that is responsible for the NQR signal that we are interested in. For a pulse sequence, the Hamiltonian $\mathcal{H}_{rf}$ is $\mathcal{H}_{pulse}$ for a period of time, followed by absence of a pulse interaction for another period of time, after which, again $\mathcal{H}_{pulse}$ is switched on, and so on. The pulse will change the EFG, and hence the coupling strength of the quadrupole moment to the EFG. It is usual to approximate this as a series $\mathcal{H}_0$, $\mathcal{H}_1$, $\mathcal{H}_2$, $\mathcal{H}_3, \ldots$ describing the Hamiltonians at the time instances $t_0, t_1, t_{2,\ldots}$ The evolution Eq. (1) can then be formally solved for ρ to yield

$$\rho(t_0+t_1+\ldots) = e^{-i\mathcal{H}_n t_n}\ldots e^{-i\mathcal{H}_0 t_0}\rho(0)e^{i\mathcal{H}_0 t_0}\ldots e^{i\mathcal{H}_n t_n}$$

Because our data comes from the electron current in the coil, we need a way to connect the state of the mixture of the quadrupole active spins to this current. The coil used is a Faraday detector, and the electron current in the coil is the direct result of the load of the preamplifier connected to the coil and the change of the magnetic flux inside the coil. The expectation of the magnetization in the direction of the axis of symmetry of the solenoid (say, the z-axis), is obtained by tracing over the product of the state ρ (the mixture of quadrupole active spin-1 states) with the magnetization operator μ_z along that spatial axis:

$$\langle M_z \rangle = Tr(\mu_z \rho) \quad (3)$$

Such a tracing operation, is an example of a so-called *quantum operation*. A quantum operation offers the most general possible description of an evolution [12], and is defined as a mapping ε that transforms an initial stateρ_0 to a final state ρ

$$\rho = \varepsilon(\rho_0) \quad (4)$$

such that there exists a set $\mathcal{O}$, called *operation elements*,

$$\mathcal{O} = \{E_k : \sum_k E_k E_k^\dagger = I, \forall \rho : Tr(E_k\rho) \geq 0\}, \quad (5)$$

for which ε can be written as

$$\varepsilon(\rho_0) = \sum_k E_k \rho_0 E_k^\dagger \quad (6)$$

The operations, by definition (5), satisfy $\sum E_k E_k^\dagger = I$, and are hence *trace-preserving*. Important examples of operations that are trace-preserving, are projective measurements, unitary evolutions and partial tracing. If the quantum operation is a general description of a quantum measurement (or evolution), then to each outcome k, we associate one member E_k of the collection of measurement operators $\mathcal{O} = \{E_k, k = 1, 2, ...\}$ that act on the state space. If the state is ρ immediately before the measurement , then the probability that the outcome k occurs is

$$p(k|\rho) = Tr(E_k \rho E_k^\dagger) \quad (7)$$

and the state after the interaction, if k occurs, is

$$\rho_{fin} = \frac{E_k \rho E_k^\dagger}{Tr(E_k \rho E_k^\dagger)} \quad (8)$$

The two most common examples of quantum operations, are unitary transformations ($\varepsilon(\rho_0) = U\rho_0 U^\dagger$, U a unitary transformation) and von Neumann projective measurements ($\varepsilon_m(\rho_0) = P_m \rho_0 P_m^\dagger$, with P_m a projector on the subspace labelled m). Many more examples, such as in quantum computation, can be found in [12] and modern descriptions of quantum experiments, as in [3], [13]. In the latter, a set $\{M_k\}$ of positive operators satisfying $\sum M_k = I$ and $E_k = M_k^{1/2}$ is used.

Quantum operations are also a natural way to describe quantum noise and the evolution of an open system. The mathematical prescription of a quantum operation arises when one considers the system to be in interaction with an environment, and that together form a closed system, for which Eq. (1) applies. To see how this applies here, we denote the initial state of the system under investigation by ρ_{sys}, and the state of the environment (soil and interfering RF fields) as ρ_{env}, then the compound system can be written as a tensor product of those states: $\rho_{sys} \otimes \rho_{env}$. Following the standard rules of quantum mechanics, the expected mixture ρ is the partial trace over the degrees of the environment of the time evolved state of the closed system:

$$\rho = Tr_{env}(U(\rho_{sys} \otimes \rho_{env})U^\dagger) \quad (9)$$

It can be shown [12] that Eq. (9) is only slightly more general than Eq. (6), hence ρ can be described as resulting from a quantum operation acting on the system density matrix. Depending on whether the environment contains TNT or not, the examined system has a density matrix written as ρ_{sys}^{tnt}, or ρ_{sys}^1. We then expect either of two generic types of operation to have occurred:

$$\begin{aligned} \varepsilon'(\rho_{sys}^{tnt}) &= Tr_{env}(U(\rho_{sys}^{tnt} \otimes \rho_{env})U^\dagger) = \rho^0 \\ \varepsilon'(\rho_{sys}^{1}) &= Tr_{env}(U(\rho_{sys}^{1} \otimes \rho_{env})U^\dagger) = \rho^1 \end{aligned} \quad (10)$$

Here ρ^0 is the resulting mixture that produces the magnetization in the presence of TNT, and ρ^1 is the resulting mixture after the interaction, in absence of TNT. An optimal detection of TNT, hence entails optimally distinguishing the two quantum states ρ^0 and ρ^1. As mentioned above, we do not posses detailed knowledge of the states ρ^0 and ρ^1 in practice, but we have the quadrature components $V(t)$. We assume the quadrature components are induced by the magnetization in Eq.3 by means of another quantum operation acting on the unknown mixture ρ_{sys}:

$$V(t) = \varepsilon_{qc}(\rho_{sys}) \tag{11}$$

Quantum operations are closed under conjunction; two consecutive quantum operations can always be represented as a single quantum operation. What we need to distinguish in the laboratory, are then the situations represented by

$$\begin{aligned} V_0(t) &= \varepsilon_{qc}(\varepsilon'(\rho_{sys}^{tnt})) = \varepsilon(\rho_{sys}^{tnt}) \\ V_1(t) &= \varepsilon_{qc}(\varepsilon'(\rho_{sys}^{1})) = \varepsilon(\rho_{sys}^{1}) \end{aligned} \tag{12}$$

We will show shortly that quantum operations can only have the effect of reducing the trace distance, which in turn will increase the minimal Bayes risk associated with distinguishing the two situations. Hence Bayes optimal detection using the quadrature components induces some loss in detector performance in comparison with the same procedure applied to a reconstruction of the state ρ_{sys}; because we skip one quantum operation in Eq.12, this would lead to a lower Bayes-risk.

DETECTION SCHEMES

Bayes optimal observation of NQR data

In essence, Bayes-optimal detection deals with the optimal decision of a hypothesis from a set of mutually exclusive hypotheses. Consider the binary decision problem

$$\begin{aligned} H_0 &: \text{the signal indicates TNT presence} \\ H_1 &: \text{the signal indicates no TNT presence} \end{aligned}$$

If a given set of data is compatible only with one of the two hypotheses, the decision problem becomes trivial. However, in practice, the data generally supports both hypotheses, albeit with a different probability, and the decision task is consequently complicated by this fact. If we are given data x_i from a set of possible outcome results $X = \{x_1, x_2, \ldots, x_i, \ldots, x_n\}$, and the factual occurrence of x_i supports both hypotheses, we need to infer what the probability was of getting the result x_i as a result of either hypothesis being true. That is, we need some means to evaluate $p(x_i|H_0)$ and $p(x_i|H_1)$. Any additional (prior) information can be included under the label D and then we compare $p(x_i|H_0, D)$ and $p(x_i|H_1, D)$. Of course, what we are after, is the probability of H_0 or H_1 being true, on the condition that D holds and x_i was the outcome of the experiment. By the use of Bayes' theorem [8], we have

$$p(H_0|x_i, D) = p(H_0|D)\frac{p(x_i|H_0, D)}{p(x_i|D)} \tag{13}$$

$$p(H_1|x_i, D) = p(H_1|D)\frac{p(x_i|H_1, D)}{p(x_i|D)} \tag{14}$$

We eliminate the denominator by calculating the ratio of Eq. (13) and Eq. (14):

$$\frac{p(H_0|x_i, D)}{p(H_1|x_i, D)} = \frac{p(H_0|D)}{p(H_1|D)}\frac{p(x_i|H_0, D)}{p(x_i|H_1, D)} \tag{15}$$

In absence of any preference which of the two hypotheses is more likely than the other on the basis of the prior information alone, we set $\frac{p(H_0|D)}{p(H_1|D)} = 1$. In complete absence of any prior information, we omit dependence on D. The quantity of interest for optimally choosing between two alternative hypotheses, is the likelihood ratio (also called *the odds* in the binary case):

$$\Lambda_i = \frac{p(H_0|x_i)}{p(H_1|x_i)} \tag{16}$$

If the signal strength is low in comparison with the noise content of data, the probability $p(H_1|x_i)$ of no TNT being present given the data x_i, will generally be greater than $p(H_1|x_i)$. We call the detection apparatus *Bayes-optimal* iff the obtained outcome x_i is the outcome that maximizes the odds, Eq.(16), that the outcome given, pertains to the system under investigation rather than to noise in the detection system. It turns out that this is a model for quantum as well as classical observation [2]. Assuming our detector is Bayes-optimal, allows for an optimal detection strategy by reversing the logic of the detector. Of course, in practice we do not know in advance whether the physical detector satisfies the condition of Bayes-optimality, and the performance will depend on how well this condition will be met. In accordance with quantum mechanics, we assume the probability $p(H_0|x_i)$ (and $p(H_1|x_i)$) is a monotone function of the trace distance between the actually measured signal, and the ideal (averaged over many samples) signal obtained in the presence (absence) of TNT. The rationale for taking the trace distance, is that it arises naturally when one considers the Bayes risk in the state discrimination problem.

Trace distance and Bayes risk of distinguishing quantum states

If we are given two states ρ^0 and ρ^1 with *a priori* probabilities p_0 and $p_1 = 1 - p_0$, then, following Eq. (4), we

look for two operations elements $\mathcal{O}=\{E_0, E_1)$ such that $E_0 + E_1 = I$ and $E_0, E_1 \geq 0$ that minimize the *Bayes risk* or probability of error [4] :

$$R_{\mathcal{O}}(p_0) = p_0 Tr(\rho^0 E_1) + p_1 Tr(\rho^1 E_0) \tag{17}$$

rewriting Eq.(17) once with $E_1 = I - E_2$ and once with $E_2 = I - E_1$, adding and dividing, yields

$$R_{\mathcal{O}}(p_0) = \frac{1}{2}[1 - Tr[(p_0\rho^0 - p_1\rho^1)(E_0 - E_1)]$$

To proceed, we define the trace distance between $\rho^0$and ρ^1, as

$$D(\rho^0, \rho^1) = \frac{1}{2} Tr\sqrt{(\rho^0 - \rho^1)(\rho^0 - \rho^1)^\dagger} \tag{18}$$

An important property of the trace distance is that it is symmetric in its arguments, positive iff $\rho^0 \neq \rho^1$, zero iff $\rho^0 = \rho^1$, and satisfies the triangle inequality. In other words, it is a bona-fide distance measure on the set of density matrices. Another important property of the trace distance, is given by

$$D(\rho^0, \rho^1) = \max_{E_i \in \mathcal{O}} Tr(E_i(\rho^0 - \rho^1))$$

With this we can show [12] that the minimum value the Bayes risk $\min_{\mathcal{O}} R_{\mathcal{O}}(p_0)$ can attain does not depend on $\mathcal{O}$ and equals

$$\min_{\mathcal{O}} R_{\mathcal{O}}(p_0) = R_{Bayes}(p_0) \tag{19}$$

$$= \frac{1}{2} - \frac{1}{2} Tr[\sqrt{(p_0\rho^0 - p_1\rho^1)(p_0\rho^0 - p_1\rho^1)^\dagger} \tag{20}$$

In our treatment of the data, each specific sample could equally well contain TNT, or not, so that we have as prior probabilities $p_0 = p_1 = 1 - p_0 = 1/2$:

$$R_{Bayes}(p_0) = \frac{1}{2} - D(\rho^0, \rho^1) \tag{21}$$

We see the minimal Bayes risk is attained for two states that maximize the trace distance. Trace preserving quantum operations can be shown to cause a contraction in the space of density operators [12]. Because the trace distance is a true distance measure on the space of density operators, it can only decrease as a result of an arbitrary trace-preserving quantum operation ε:

$$D(\rho^0, \rho^1) \geq D(\varepsilon(\rho^0), \varepsilon(\rho^1)) \tag{22}$$

If the current in the coil is the result of Eq. (10), then being able to distinguish the currents reliably (i.e., the trace distance is greater than can be explained from fluctuations), indicates we have successfully distinguished the situations represented by H_0 and H_1.

The Demodulation technique

A popular method to establish the presence of a given substance in a NQR tested sample, is the use of the so-called demodulation technique. This method is particularly simple and consists of calculating an estimate $\sigma(\upsilon_n)$ of the power spectral density $S(\upsilon)$ of the signal $s(t_n)$, $n = 1, \ldots, 256$, by first fast Fourier transforming the signal and taking its modulus squared. Let us call $\upsilon_{\max}$ the frequency $\upsilon_{\max} = \arg(\max(S(\upsilon)))$ where one expects the spectral line with the highest intensity in presence of TNT. The value of the *estimated* power spectral density $\sigma(\upsilon_{\max})$ evaluated at the frequency $\upsilon_{\max}$, is then the test statistic for a treshold detector. If $\sigma(\upsilon_{\max})$ exceeds a given treshold, the presence of TNT is accepted, if not, it is rejected. The estimated $\sigma(\upsilon_n)$ will in general deviate from $S(\upsilon)$ at the precise values υ_n, but may be approximately regarded as an average over the interval $[\frac{\upsilon_n - \upsilon_{n-1}}{2}, \frac{\upsilon_{n+1} - \upsilon_n}{2}]$. To account for this, sometimes the average under $\sigma(\upsilon_n)$ over a few frequency bins is taken as a statistical test parameter. Whether this is useful or not depends, among other things, on the magnitude of the width of the spectral line with the highest intensity relative to the width of the frequency bins. Moreover, as NQR spectra are generally a function of the temperature of the sample, and because this parameter is difficult to estimate in demining applications within a range of 5 to 10 Kelvin, the value of $\upsilon_{\max}$ will depend on the temperature too. To make sure we do not miss the peak, one can then take the area over a region in the frequency domain where one expects the peak. This complication presents no real problem to the method employed and, however important to the actual demining problem, is not taken into account here (see, however, [7] and [11]). All our experimental samples are taken at the same temperature. As expected, we see little change in the efficiency of the method, whether we use $\sigma(\upsilon_{\max})$, or a sum of values $\sum \sigma(\upsilon_{\max})$ for a tiny region surrounding the relevant frequency bin. However, a considerable improvement is obtained when we allow for the demodulation technique to sample multiple peaks simultaneously. The results that we present here, employ the single peak value $\sigma(\upsilon_{\max})$ of the frequency bin containing the mean excitation frequency 841.5 kHz, as well as an improved demodulation algorithm exploiting knowledge of three resonance frequencies of TNT within a range of a few tens of kHz around the mean excitation frequency 841.5 kHz.

EXPERIMENTAL RESULTS

Set up and data acquisition

The data employed for our analysis was kindly provided by the NQR group of King's College, London, under supervision of Professor J.A.S. Smith. In the experi-

mental set up employed, a pure monoclinic TNT sample, typical of that found in an anti-personnel mine, is placed inside a solenoidal coil. The same coil is used for emission of the RF-pulse, as well as for the reception of the subsequent echo. The returned echo is routed through a hardware band-pass filter (approximately 50 kHz width) and subsequently sent to a Tecmag Libra spectrometer that splits the signal in two signals which are then mixed with two quadrature components, yielding a complex discrete time series. Because of the ideal laboratory conditions under which the signal has been obtained, the results will compare unrealistically optimistic with respect to those obtained under field conditions. In particular, the absence of RF interference and the use of a coil that contains the sample in its entirety, must be taken into account when attempting to compare the results of our analysis with those of data obtained under more realistic conditions. The emitted RF signals are pulsed, spin locked echo signals with a mean excitation frequency of 841.5 kHz. The mean and width of the excitation are such, that 4 spectral lines of TNT can be detected within the frequency range of the band pass filter. Because the same coil is used for the emission of the RF signal (which has a mean power of several kilo watts), as for the reception of the echo (which is extremely weak), the returned echo contains so-called antenna ringing effects. To cancel the effect of the antenna ringing, a phase cycling technique, popular in the more established field of NMR, is employed. The phase cycling technique requires forming an appropriate sum of four signals. The signals used for the analysis, are the sum of 5 such phase-cycled sums and hence consist of 20 repeated data acquisitions, averaged to improve signal to noise ratio. The sampling time is 5μS. and each set has 8192 data points, which consists of 32 sequential echo signals, each containing 256 data points. The pulse sequence is of the type

$$\pi - \tau - \pi - 2\tau - \pi - 2\tau - \pi - 2\tau - ...$$

Here π denotes the RF pulses and the 1280 μS. of data (256 times 5 μS.) for each echo signal is acquired during the 2τ periods between the pulses. All algorithms are programmed in MATLAB 7 on a 2,2 Ghz PC with 512 MB RAM. Both methods are fast: the determination of whether a given signal was obtained in the presence of TNT or not, requires in both cases a calculation time less than a second, more than one order of magnitude below the necessary data acquisition time.

Detector performance

The statistical assessment of detector performance is based on the sensitivity, specificity and ultimately on the functional relationship which exists between these two, as expressed in the receiver operating characteristic (ROC). The true positive rate, or sensitivity, is the probability that the detector indicates the presence of a mine, when there is indeed a mine present. The false positive rate is defined as the probability that the detector indicates the presence of a mine when there was no mine present. The specificity is then defined as 1−false positive rate. Increasing the sensitivity of the detector, lowers the specificity and vice versa. The overall performance is therefore expressed by means of a receiver operating characteristic (ROC) curve, which plots the sensitivity as a function of the false positive rate. By definition, all ROC curves have at least two points in common: the origin (sensitivity zero, specificity 100%, i.e. the detector never indicates the presence of a mine) and the point (1,1) (sensitivity 100%, specificity 0%, i.e. the detector indicates presence of a mine regardless of sample content). An ideal detector is one with a ROC curve that jumps to 100% immediately after the point (0,0) and stays there. Such a detector can reliably distinguish signals from noise and vice versa. A completely ignorant detector (i.e., a detector that indicates presence or non-presence in a completely arbitrary fashion) is then one for which the sensitivity always equals (1-specificity) and corresponds to a ROC curve that coincides with the line joining (0,0) with (1,1). The data used to calculate the ROC curves were obtained by using 100 data samples with TNT, and 100 data samples without TNT. Because of spin-lattice relaxation, each progressive echo contains less information. Consequently we expect the signal quality to decrease as a function of the echo number, a behavior we see reflected in the ROC curves and in the line intensities of the three most visible resonances as a function of the echo number, as depicted in figure 1.

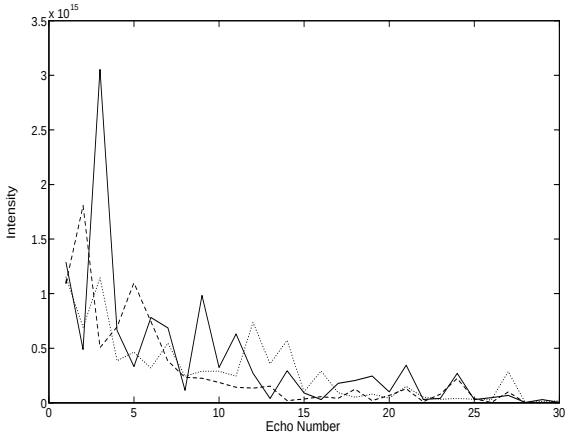

FIG. 1: The intensity of the three TNT quadrupole resonances within experimental reach, averaged over 100 signals, as a function of the echonumber. The decrease in intensity is approximately exponential. The intensity of the last few echoes is two orders of magnitude smaller than the first.

In figure 2 and figure **??**, we have depicted ROC curves

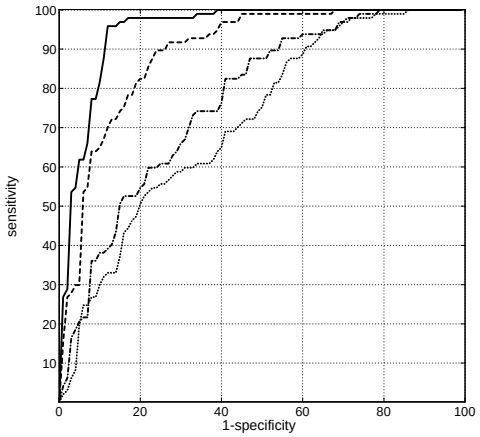

FIG. 2: ROC curves for the demodulation technique for a single peak. In decreasing order of performance (ever lower ROC curves), we plotted ROC curves for echoes 9 (solid), 13 (dashes), 17 (point-dash) and 21 (points) respectively. The first four echoes yield close to perfect detectors.

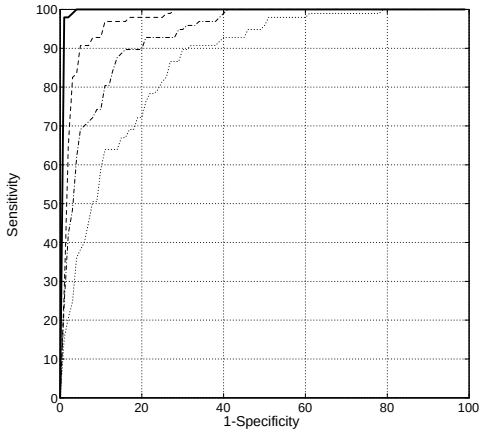

FIG. 3: ROC curves for the Bayes optimal detector. As in the previous graph, shown are ROC curves for echo numbers 9 (solid), 13 (dashes), 17 (point-dash) and 21 (points), in decreasing order of performance respectively. The first ten echoes yield close to ideal detectors.

for both detection methods. As is well-known, ROC curves can be "convexified" using mixed measurements, i.e., measurements that are linear combinations of measurements with a threshold value that corresponds to extremal points of the experimental ROC curve. This procedure results in an improved detector. Nevertheless, we have depicted all ROC curves calculated only for single threshold values; the convexified ROC curves can easily be visualized from the given experimental curves.The Bayes optimal detector uses the whole signal without noise reduction (except for the hardware band-pass filter), whereas the demodulation detector measures the intensity of the single peak with the highest intensity, which comes down to a very narrow boxcar band-pass filter. In this sense the two methods are opposites of one another. An intermediate detector can be obtained by using the three most distinct resonances in the NQR spectrum of TNT within the frequency band allowed by the band-pass filter. As can be seen from the ROC curves in figure 4, the performance of the three peak demodulation technique, although still lagging behind the Bayes optimal detector, is better than the single peak detector.The normalized area under a ROC curve can be taken as a crude measure of the overall performance of the detection scheme. The ideal case then corresponds to an area of one, the completely ignorant detector scores one half. Due to the high risks involved in the practice of demining, it is crucial no mine be missed, i.e., one wants to have a sensitivity as close to 100% as possible. As indicated before, the main problem with many of the current demining detectors, is that the necessity of having such a high sensitivity, will lower the specificity to such an extent, that many false alarms become unavoidable.The Bayes optimal detector allows to use any of the first 10 echo numbers to obtain a detector that is very close to ideal. In contrast, the demodulation detector only yields a close to ideal detector for the first 4 echoes. We see in figure 5 that the overall performance of the demodulation detector decreases much more rapidly than the Bayes optimal detector. As the first echo already yields a perfect detector for both methods, there seems no obvious incentive to improve the detection capabilities. However, the spin-lattice relaxation constrains the

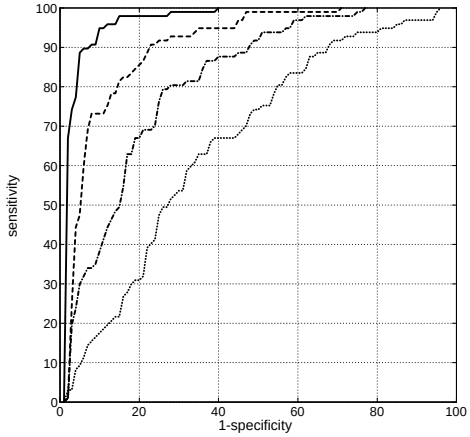

FIG. 4: An improved demodulation detector utilising the three most distinct peaks in the NQR spectrum. As before, depicted are echoes 9 (solid), 13 (dashes), 17 (point-dash) and 21 (points). One can see the detector is a considerable improvement over the single peak detector, but lags behinds the Bayes optimal detector.

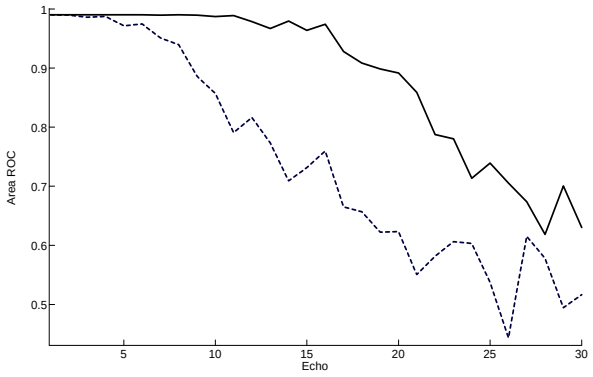

FIG. 5: A comparison between the overall performance of the demodulation (dotted curve) and Bayes optimal detection (solid curve) methods. Depicted is the area under the ROC curve as a function of the echo number. One can see how for both methods the performance decreases with increasing echo number. The Bayes optimal detection clearly outperforms the demodulation technique after echo number 4.

time between the spin-locked pulses to a minimum of 10 seconds. Hence the necessary data acquisition time for the each single data sample, is approximately 20 acquisitions*10 seconds = 200 seconds. In actual demining applications, the necessary acquisition time will further increase as a result of RF interference, other NQR active soil constituents such as piezoelectric ceramics, and the fact that only single sided (as opposed to the sample being *within* the coil, as is the case for our data), remote acquisition is possible. However, one can substantially decrease this acquisition time by combining the information in the different echoes. It is hence of vital importance to improve the detector performance for *all* the echoes in the pulse sequence. The proposed detector succeeds in doing just that.

CONCLUDING REMARKS

We examined the applicability of a simple Bayes-optimal quantum state discrimination technique to see if it is possible to improve the detection capabilities of remote TNT detection by NQR measurements. Although the experimental setup employed here is only able to give an estimate of a single the projection of the NQR state onto the axis of symmetry of the solenoid, the method delivers a very reliable detector. A comparison was made with the popular demodulation detector. Both methods are simple and fast, but our results indicate the Bayes-optimal scheme offers a considerable improvement over the demodulation approach. The difference in performance between the two methods becomes greater as the echo number increases. For the last few echoes, the advantage becomes less pronounced, which we attribute to the fact that the signal strength of the echo diminishes exponentially with increasing echo number, so that eventually both methods will fail to deliver for very weak echoes. Handling signals with a low SNR is important, as one expects a deterioration of the already low SNR inherent in NQR measurements in actual field tests. The proposed detector offers two distinct and important advantages for demining applications: increase of the specificity (without sacrificing the close to perfect sensitivity necessary for demining applications), and a decrease of the necessary detection time. It is possible to include data from primary detectors (such as a metal detector or ground penetrating radar) in the form of prior probabilities in Eq. 19, so that the NQR detector becomes a confirmation sensor. It would be interesting to combine different pulse sequences that allow for a two or three dimensional reconstruction of the full density matrix of the spin-1 NQR system, and see whether this leads to a better detector. Besides the ability to have a more faithful characterization of the density operator, varying the pulse scheme may offer other advantages. By tailoring the pulse sequences to enact on disjoint excitations in the frequency plane, one may, depending on the magnitude of the cross-relaxation between the modes, be able to improve the extraction of the information content by questioning different modes. If different pulse sequences are transmitted by different antenna, one can improve the fraction of nuclei participating in the NQR signal above the 43% limit for a single field orientation. As the state contains all information about the system, a detector based on the reconstructed density operator, yields an approximation to a truly optimal detector. It remains to be seen whether an implementation of such a detector offers practical improvements for demining applications.

ACKNOWLEDGEMENTS

This work was done as part of Flemish Fund for Scientific Research (FWO) research project G.0362.03N. We gratefully acknowledge J.A.S. Smith and M. Rowe for supplying us their NQR data and helpful feedback on our results. We thank Hichem Sahli and Luc van Kempen for fruitful discussions and feedback.

[1] Aerts S. (2005). The Born rule from a consistency requirement on hidden measurements in complex Hilbert space. Int. J. Theor. Phys. **44**, 7

[2] Aerts S. (2005). Quantum and classical probability from Bayes optimal observation, arXiv:

[3] Busch P., Grabowski M., Lahti P. J., Operational Quantum Physics, Springer, Lecture Notes in Physics,1995

[4] Helstrom C.W. (1976). Quantum detection and estimation theory. Academic, New York.
[5] Das T.P., Hahn E.L. (1958). Nuclear quadrupole resonance spectroscopy. Solid state physics. New York, Academic. Suppl. 1
[6] Holevo A.S., (1982). Probabilistic and Statistical aspects of quantum theory. North-Holland, Amsterdam
[7] Jakobsson A., Mossberg M., Rowe M.D., Smith J.A.S.(2005). Exploiting temperature dependency in the detection of NQR signals. IEEE Transactions on Signal Processing (in press)
[8] Jaynes E.T. (2003). Probability theory, the logic of science, Cambridge University Press
[9] Lee Y.K. (2002). Spin-1 nuclear quadrupole resonance theory with comparison to nuclear magnetic resonance. Concepts in Magnetic Resonance **14**, 3
[10] Malley D.M., Hornstein J., (1993). Quantum statistical inference. Statistical Science, **8**, 4, 433-457
[11] Mikhalsevitch V.T., Rudakov T.N. (2004), On the NQR detection of nitrogenated substances by multi-pulse sequences. Phys. stat. sol. (b), **241**, 2, 411-419
[12] Nielsen M.A., Chuang, I.L. (2000). Quantum computation and quantum information. Cambridge University Press
[13] Schroeck F.E., (1996). Quantum Mechanics on Phase space. Kluwer Academic Publishers, Dordrecht
[14] An exception is the PFM-1 landmine which contains a liquid explosive, which is outside the scope of current NQR techniques.