Resource-efficient adaptive Bayesian tracking of magnetic fields with a quantum sensor

K. Craigie, E. M. Gauger, Y. Altmann and C. Bonato

School of Engineering and Physical Sciences, SUPA, Heriot-Watt University, Edinburgh, EH14 4AS, UK

E-mail: klc31@hw.ac.uk, c.bonato@hw.ac.uk

Abstract. By addressing single electron spins through Ramsey experiments, nitrogen-vacancy centres can act as high-resolution sensors of magnetic field. In applications where the magnetic field may be changing rapidly, total sensing time is crucial and must be minimised. Bayesian estimation and adaptive experiment optimisation protocols work by computing the probability distribution of the magnetic eld based on measurement outcomes and, by computing acquisition settings for the next measurement. These protocols can speed up the sensing process by reducing the number of measurements required. However, the computations feeding into the next iteration measurement settings must be performed quickly enough to allow real-time updates. This paper addresses the issue of computational speed by implementing an approximated Bayesian estimation technique, where probability distributions are approximated by a superposition of Gaussian functions. Given that only three parameters are required to fully describe a Gaussian, we find that the magnetic field probability distribution can typically be described by fewer than ten numbers, achieving a reduction in the number of operations by factor 20 compared to existing approaches, allowing for faster processing.

Keywords: Nitrogen-vacancy centre, Quantum metrology, Magnetic field tracking, Bayesian filtering

1. Introduction

Control and measurement of individual electron spins, achieved in the last two decades, enables highly sensitive measurements of magnetic fields [1–3], with spatial resolution on the order of tens of nanometres [4]. This is typically achieved through a defect in diamond, the nitrogen-vacancy (NV) centre, which allows spin readout even at room temperature through Optically Detected Magnetic Resonance (ODMR) and long spin coherence time [5, 6].

NV centres can also be used as nanoscale sensors for temperature [7], strain [8,9] and electric eld [10, 11]. Their sensing capabilities have been applied to studying nanoscale magnetic phenomena in materials and biological processes [4,12]. For example, biological compatibility for nanodiamonds containing NV centres allows monitoring of nanoscale,

in-vivo processes [13–16]. Performing sensing faster, or tracking a changing physical quantity in real-time could give new insight into previously inaccessible timescales.

Recent work has demonstrated the power of Bayesian modeling and estimation, coupled with adaptive rules for experiment optimisation in order to speed up the magnetic field sensing process [17–21]. While several algorithms have been proposed and analysed [22, 23], only one experimental implementation so far has enabled fully online operation [24]. An important point is that real-time implementation requires fast computations, on the microsecond scale, to optimise the settings for the following measurements. Indeed, a slow computational turnaround can yield an adaptive sensing protocol whose performance is inferior to non-adaptive alternatives. Time-critical computation with minimal latency can be performed in parallel with a fast digital electronic system, such as a field-programmable gate array (FPGA). However, the number of (sequential) operations still needs to be kept as low as possible to keep the overall computational overhead real-time compatible.

Here we address this issue by adopting an approximate Bayesian estimation technique. In particular, at each point in time, we approximate the likelihood function and the posterior distribution of the parameter of interest as finite sums of Gaussian functions, i.e., Gaussian mixtures. Since Gaussian functions can be fully described by only three parameters (amplitude, location and width), this allows faster processing as the number of parameters propagated over time is small compared to what would be required if the distributions were discretised on a grid or approximated via particle filtering [25]. We find that our Bayesian approach can typically be performed using only one or two Gaussian functions, achieving a 20-fold reduction in the number of operations compared to previous non-approximate implementations [22, 23].

While the work detailed here focuses on quantum sensing with NV centres in diamond, the protocol we examine can be readily applied to any other single-qubit quantum sensor [26, 27].

2. Background

A magnetic field applied to the electron spin induces a Zeeman splitting of the energy levels, which can be measured by a Ramsey experiment [28]. In a Ramsey experiment, an equal spin superposition freely evolves under the applied magnetic field B, so that the spin eigenvalues acquire a relative phase, corresponding to a rotation at the Larmor frequency. The probability for outcome $\mu \in \{0; 1\}$ given a Larmor frequency f_B (corresponding to a magnetic field $B = f_B/\gamma$, where γ is the gyromagnetic ratio, is

$$P(\mu|f_B, \theta, \tau) = \frac{1 + e^{i\mu\pi} cos(2\pi\tau f_B + \theta)}{2} \,. \tag{1}$$

where τ is the sensing time and θ the rotation angle of the measurement basis (which is controlled by the phase of the second $\pi/2$ pulse in the Ramsey measurement sequence).

In Equation 1, $P(x|y,z)$ denotes the distribution of x, conditioned on the value of (y,z), ie, the distribution of x given y and z.

In this paper, we consider the problem of tracking a magnetic field (through f_B) that fluctuates on timescales longer than a single measurement time, but potentially shorter than a large number of repetitions of the measurement time. We assume the fluctuations to be described by a Wiener process yielding a sequence of Larmor frequency values that are generated as

$$f_B^{(t+\delta t)} = f_B^{(t)} + \kappa \mathrm{d}W^{(t)}\,. \tag{2}$$

The diffusion coefficient, κ, is a measure of the magnetic field rate of change and $\mathrm{d}W^{(t)}$ is an innitesimal Wiener increment during a time inverval δt. This process is simulated by discretising the time axis to intervals of length τ_{min} (minimum sensing time), and generating a normal distribution with variance τ_{min}. In this way, a changing Larmor frequency signal is generated that will act as the ground truth in tracking simulations. More generally for any small time interval δt, equation 2, leads to the following Gaussian random walk distribution

$$P\left(f_B^{(t+\delta t)}\middle| f_B^t, \kappa\right) \propto \exp\left[-\frac{\left(f_B^{(t+\delta t)} - f_B^{(t)}\right)^2}{2\delta t \kappa^2}\right], \tag{3}$$

which will be used in the tracking algorithm described in the next section.

2.1. Bayesian Estimation

Bayesian online tracking of f_B consists of approximating in real time of the probability distribution f_B, which can in turn allow us to optimise the experimental settings through θ and τ in equation 1. After the n-th measurement, the (posterior) distribution $P(f_B^{(t_n)}|\boldsymbol{\mu}^{(t_n)}, \theta_n, \tau_n)$, of $f_B^{(t_n)}$ is updated using Bayes' rule

$$P(f_B^{(t_n)}|\boldsymbol{\mu}^{(t_n)}, \theta_n, \tau_n) \propto P(\mu^{(t_n)}|f_B^{(t)}, \theta_n, \tau_n) P(f_B^{(t_n)}|\boldsymbol{\mu}^{(t_{n-1})}, \theta_n, \tau_n), \tag{4}$$

where $\boldsymbol{\mu}^{(t_n)} = \left\{\mu^{(t_0)}, \ldots, \mu^{(t_n)}\right\}$. In equation 4, $P(f_B^{(t_n)}|\boldsymbol{\mu}^{(t_{n-1})}, \theta_n, \tau_n)$ acts as a prior distribution and can be obtained via

$$\begin{aligned} &P(f_B^{(t_n)}|\boldsymbol{\mu}^{(t_{n-1})}, \theta_n, \tau_n) = \\ &\quad \int P\left(f_B^{(t)}\middle|\, f_B^{(t_{n-1})}, \kappa\right) P(f_B^{(t_{n-1})}|\boldsymbol{\mu}^{(t_{n-1})}, \theta_n, \tau_n) \mathrm{d}f_B^{(t_{n-1})}, \end{aligned} \tag{5}$$

with $P(f_B^{(t_{n-1})}|\boldsymbol{\mu}^{(t_{n-1})}, \theta_n, \tau_n)$ obtained via the Bayesian update after the $(n-1)$-th measurement. The update rules for (θ_n, τ_n), are described in more detail in Sections 3.5.2 and 3.5.1.

Figure 1 illustrates the main principle of Bayesian adaptive tracking algorithms. However, as discussed in the introduction, exact Bayesian inference based on equations 4 and 5 is not tractable because of the shape of the likelihood in equation 1, which makes the integral in equation 5 not computationally tractable efficiently. While it is theoretically possible to use particle filters [25], their computational complexity make

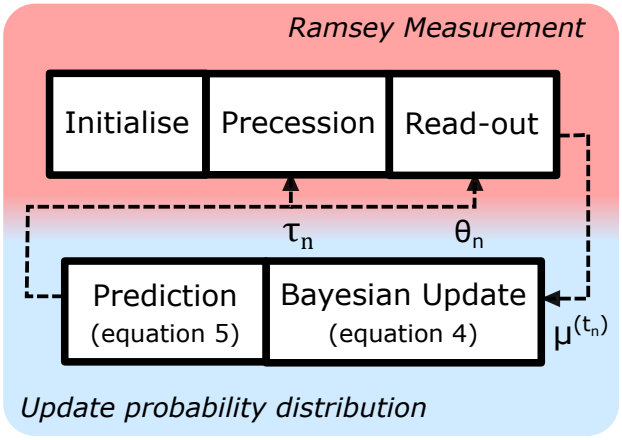

Figure 1. Graphical representation of adaptive protocol, showing flow of information. The measurement outcome, μ, from the Ramsey experiment is fed into the Bayesian update, followed by a prediction step. From this computation, the adaptive phase, θ_n, and sensing time, τ_n, are determined. These values are fed into the following Ramsey experiment.

them less attractive than approximate methods using Gaussian mixture approximations, as proposed here.

For completeness, a simplified example of the tracking protocol, similar to that described in [22], is provided in figure 2, where the adaptive phase and sensing time are chosen to completely eliminate one of the two peaks in panel (d).

3. Method

In this article, we aim to reduce computational time and memory requirements by approximating $P(f_B^{(t_n)}|\boldsymbol{\mu}^{(t_n)}, \theta_n, \tau_n)$ as a finite sum of Gaussian distributions. This is performed by approximating the likelihood $P(\mu^{(t_n)}|f_B^{(t_n)}, \theta_n, \tau_n)$ by a weighted sum of Gaussian functions (with respect to $f_B^{(t_n)}$). Figure 2a shows how the actual cosine-shaped likelihood function is approximated by two shifted Gaussian peaks. As is apparent, the approximation is relatively accurate at the top of the peaks, but is rather poor at the bottom due to the Gaussian tails. In the following, we address the dual question of how well a magnetic field can be tracked despite this simplification, and examine the associated advantage of doing so in terms of the algorithm's reduction in memory requirements and computation time.

3.1. Gaussian approximation of the likelihood function

We begin by approximating the initial likelihood in equation 1 as:

$$P(\mu|f_B, \theta_n, \tau_n) \approx \sum_{l=0}^{N_G} A e^{-\frac{(f_B - a_l)^2}{2\sigma_a^2}}. \tag{6}$$

This approximation can be computed by a Taylor expansion of equation 1 at $2\pi\tau f_B + \theta_n + \mu\pi \approx 2\pi l$ and re-writing as a Taylor expansion for an exponential.

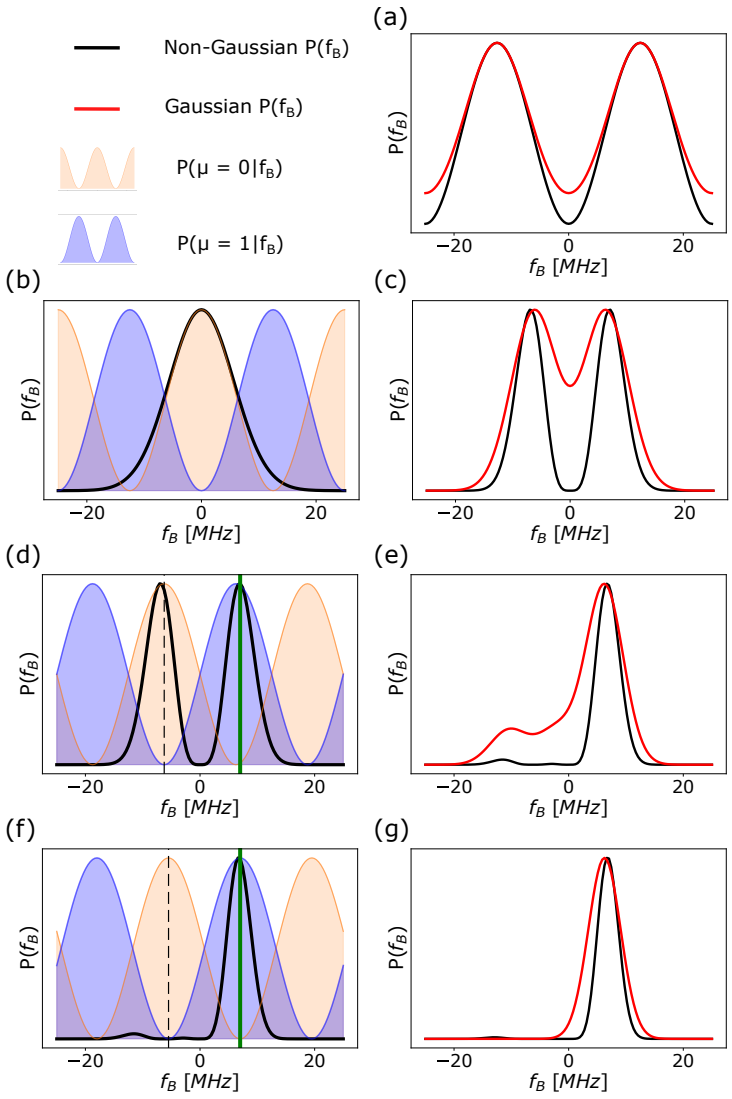

Figure 2. Tracking steps after a change in magnetic field. (a) likelihood function with Gaussian approximation, (b) prior and likelihood function for measurement n=1, (c) posterior for n=1, μ= 1, with Gaussian approximation, (d) prior and likelihood function for measurement n=2, (e) posterior for n=2, μ= 1, with Gaussian approximation, (f) prior and likelihood function for measurement n=1, (g) posterior for n=3, μ= 1, with Gaussian approximation.

$$\begin{aligned}\frac{1+\cos(2\pi\tau_n f_B+\theta_n+\mu\pi)}{2} &\approx 1-\frac{(2\pi\tau f_B+\theta_n+\mu\pi)^2}{4}\\ &\approx \exp^{-\left(\pi\tau_n f_B+\frac{\theta_n+\mu\pi}{2}\right)^2}\end{aligned} \tag{7}$$

For the initial likelihood and under this approximation, we let all Gaussians have the same amplitude $A = 1$ (the normalisation is handled) and width $\sigma_a = 1/(\sqrt{2}\pi\tau)$, meaning they only differ in their centres a_l as follows:

$$a_l = \frac{2\pi l+\pi\mu+\theta_n}{2\pi\tau}\,. \tag{8}$$

This choice ensures that the displacement between adjacent Gaussians coincides with the period of the oscillatory likelihood function and that the Gaussian peaks align with the local maxima of the cosine function.

The number of periods of the likelihood function that fit in the prior frequency range is equal to the sensing time coefficient (the sensing time as a fraction of the minimum sensing time, τ_0). This implies that the required number N_G of Gaussians for the approximation is given by $2^N + 1$, since this is the maximum sensing coefficient (plus one additional one to cater for edge peaks which are only partially visible).

3.2. Bayesian update

Once the likelihood is approximated by a finite sum of Gaussians, the Bayesian update in equation 4 becomes tractable and the posterior distribution reduced to a product of two mixtures of Gaussians (one arising from the prior distribution and one from the likelihood function). Defining A, a, σ_a and B, b, σ_b as, respectively, the amplitude, centre and standard deviation of two Gaussians, their product will be another Gaussian that is fully characterised by the parameters C, c, σ_c with

$$\sigma_c = \sqrt{\frac{\sigma_a^2\sigma_b^2}{\sigma_a^2+\sigma_b^2}}\,, \tag{9}$$

$$c = \frac{a\sigma_b^2 + b\sigma_a^2}{\sigma_a^2+\sigma_b^2}\,, \tag{10}$$

$$C = ABe^{\frac{(a\sigma_b^2+b\sigma_a^2)^2/(\sigma_a^2+\sigma_b^2)-(a^2\sigma_b^2+b^2\sigma_a^2)}{(2\sigma_a^2\sigma_b^2)}}\,. \tag{11}$$

Note that after the initial assignment, the amplitudes of different Gaussians will no longer generally be identical, requiring calculation of C above.

3.3. Prediction Step

As described in equation 5, the (predictive) probability distribution of $f_B^{(t_n)}$ at time $t_n = t_{n-1} + \delta t_n$ can be found as the convolution of the probability distribution at time t_{n-1} with a zero-mean Gaussian with variance $\delta t_n \kappa^2$. The time elapsed between measurements is given by $\delta t_n = \tau_n + t_{oh}$, where τ_n is the sensing time and t_{oh} is the overhead time. This overhead time is added to account for the time it takes to physically carry out measurements in a lab. Naturally, this value varies depending on the precise experimental equipment and set-up. The distributions before and after prediction are therefore:

$$P(f_B^{(t_{n-1})}|\boldsymbol{\mu}^{(t_{n-1})}, \theta_n, \tau_n) = \sum_l C_l e^{-\frac{\left(f_B^{(t_{n-1})}-c_l\right)^2}{2\sigma_l^2}}\,, \tag{12}$$

$$P(f_B^{(t_n)}|\boldsymbol{\mu}^{(t_{n-1})}, \theta_n, \tau_n) = \sum_l \frac{C_l\sigma_l}{\sqrt{\sigma_l^2+\kappa^2 dt}} e^{-\frac{\left(f_B^{(t_n)}-c_l\right)^2}{2(\sigma_l^2+\kappa^2\delta t_n)}}\,, \tag{13}$$

where $\sigma' = \sqrt{\sigma^2 + \kappa^2 \delta t}$ and $D = \frac{C\sigma}{\sigma'}$.

3.4. Pruning the set of Gaussians

It is important to note that, using mixtures of Gaussians to approximate the likelihood and the prior distribution (say with m and n terms, respectively), the number of Gaussians ($m \times n$) in the resulting posterior distribution keeps on increasing over time. To prevent this, we introduce an automatic pruning step based on amplitude thresholding. Any Gaussian with amplitude smaller than the predefined threshold is discarded. Additionally, we combine Gaussians that are very similar to each other, as defined by their Kullback-Leibler (KL) divergence [29]:

$$KL(g_1, g_2) = \log \frac{\sigma_2}{\sigma_1} + \frac{\sigma_1^2 + (a_1 - a_2)^2}{2\sigma_2^2} - \frac{1}{2}\,, \tag{14}$$

where g_1 and g_2 are the two Gaussians being compared and a_1, a_2 are their respective central positions. When the KL divergence for a pair of Gaussians is below a given threshold, the pair is merged into one single Gaussian whose mean and variance are obtained by averaging the parameters of the two original Gaussians and whose amplitude is the sum of the original amplitudes. We found that the KL divergence threshold $KL_{th} = 0.001$ for merging and the amplitude threshold $A_{th} = 0.04$ for pruning work well in practice within the parameter range we have studied.

The proposed tracking scheme might temporarily lose the track of f_B and this yields all Gaussians having small amplitudes. If no amplitude exceeds the threshold, all the previous amplitudes and positions are kept unchanged and the variances are doubled. By simply broadening the previous distribution when the tracking is lost, the protocol picks up the signal again after a few iterations.

3.5. Protocol Overview

The proposed protocol is divided into phases, namely the initial adaptive sensing, followed by the adaptive tracking phase. The sensing portion is required to obtain a starting Larmor frequency value, which can then be tracked. Experimentally, this could also be obtained from locating the initial Larmor frequency by observing the spin resonance signal while sweeping the spin drive frequency. During sensing, the measurement time is not chosen adaptively, but in a predetermined sequence to narrow in on the Larmor frequency. Repeating each sensing time is required to minimise errors in sensing (see line 4 of Algorithm 1). However, the goal of adaptive sensing is to use the fewest Ramsey measurements and so a balance must be achieved. The number of Ramsey measurements performed at each sensing time is determined by integers G and F, which have been selected in accordance with previous work on adaptive sensing [24]. Due to this previous work studying in detail the initial sensing process, as well as the Gaussian approximation implementation providing an advantage only for tracking, this phase is not of primary interest for the current work. Algorithm 1 gives an overview

of the initial sensing process, in which the posterior distribution of f_B typically evolves from a uniform distribution (prior to any measurements) to an almost unimodal density which can be well approximated by a single Gaussian distribution.

Algorithm 1 : Adaptive sensing overview. Variables: adaptive phase (θ_n), measurement outcome ($\mu^{(t_n)}$), sensing time (τ_n), minimum sensing time (τ_{min}), number of sensing times (N). The remaining variables, M_n, G and F, define the number of repetitions for each sensing time.

for n= 0 to N-1 **do**
  $\tau_n = 2^n \tau_{min}$
  choose θ_n
  $M_n = G + F(n-1)$
  for $m = 1$ *to* M_n : **do**
    $\mu^{(t_n)} = Ramsey(\theta = \theta_n, \tau = \tau_n)$
    $Bayesian\_update(\mu = \mu^{(t_n)}, \theta = \theta_n, \tau = \tau_n)$
  end for
end for

In contrast to the first phase, during the second phase of the protocol, the sensing time is chosen before each Ramsey measurement, based on the level of uncertainty of the current probability distribution of f_B. Algorithm 2 illustrates the sequence of steps involved in adaptive tracking. In both tracking and sensing, the Gaussian approximation only affects steps "choose θ_n," "Bayesian_update" and "Prediction". The other steps do not use $P(f_B^{(t)}|\boldsymbol{\mu}^{(t)}, \theta_n, \tau_n)$ and thus remain unaffected.

Algorithm 2 : Adaptive tracking overview. Variables: adaptive phase (θ_n), measurement outcome ($\mu^{(t_n)}$), sensing time (τ_n).

for total tracking time interval **do**
  choose τ_n
  choose θ_n
  $\mu^{(t_n)} = Ramsey(\theta = \theta_n, \tau = \tau_n)$
  $Bayesian\_update(\mu = \mu^{(t_n)}, \theta = \theta_n, \tau = \tau_n)$
  $Prediction(\kappa)$
end for

3.5.1. Adaptive Sensing Time In this work, to determine whether the optimum sensing time has been chosen, we use the figure of merit proposed in [22] and computed from an estimate of the standard deviation of the posterior distribution of f_B (as detailed in equation 21 of [22]). This figure of merit is associated with a threshold above which the sensing time is judged insufficiently accurate and reduced by a factor two for the next measurement. This procedure is repeated until the threshold condition is met. If the

figure of merit is lower than the threshold, the sensing time is instead increased by a factor two.

3.5.2. Adaptive Phase In order to maximise the information from each measurement, we adaptively set the angle of the measurement basis as [20]:

$$\theta_n = \frac{1}{2} arg\{p_{2t_n}\}\,, \tag{15}$$

where p_{2t_n} is the prior probability distribution in Fourier space, p_k for $k = 2t_n$, where t_n is the sensing time coefficient. In terms of Gaussians, p_{2t_n} is computed as:

$$p_{2t_n} = \sum_l \sqrt{2\pi} C_l \sigma_l e^{-2\pi^2 (2t_n)^2 \tau^2 \sigma_l^2 + l 2\pi (2t_n)\tau c_l} \tag{16}$$

4. Results

To test the performance of the Gaussian-approximation tracking protocol, we performed numerical simulations. We assume perfect spin readout fidelity and $T_2^* = 100$ μs. Other constant parameters were the amplitude pruning threshold $A_{th} = 0.04$, merging threshold $KL_{th} = 0.001$ as well as $G = 5$ and $F = 3$. The time interval over which the signal was tracked was set to 50 ms for direct comparisons and varied during statistical comparisons, so that each run would contain 1000 Ramsey measurements. Overhead time, the extra experimental time taken to perform a Ramsey measurement, and κ, the prediction coefficient, were varied to examine the robustness of the protocol.

The performance is assessed with the mean squared error (MSE)

$$\epsilon^2 = \frac{1}{T} \int_0^T |f_b - f_b^{est}|^2\,, \tag{17}$$

where f_{est} is the estimated frequency and f_B the true frequency. Figure 3 illustrates a successful tracking run, characterised by a typical MSE value of around 0.08 MHz/ms. In addition, to assess the improvement in terms of computational cost, we compared the average number of parameters used in the discretisation of the distribution of f_B. For the non-approximate implementation, this is the number of equally-space frequency points in the discretisation. For the Gaussian case, this is simply the number of Gaussian parameters i.e. 3× the number of Gaussians employed for approximating the probability distribution.

4.1. Discussion

As suggested by figure 2, without factoring computation time, the reliability of the tracking is generally degraded when using Gaussian tracking. In figure 4, panel a, 50 ms of the exact same magnetic field fluctuations were tracked 200 times using both methods. Neither are perfect but it is clear that the first histogram bin contains slightly fewer Gaussian runs (88.5%) than non-Gaussian runs (93%). Any run in this first bin has successfully tracked f_B without losing track of the Larmor frequency in a major way.

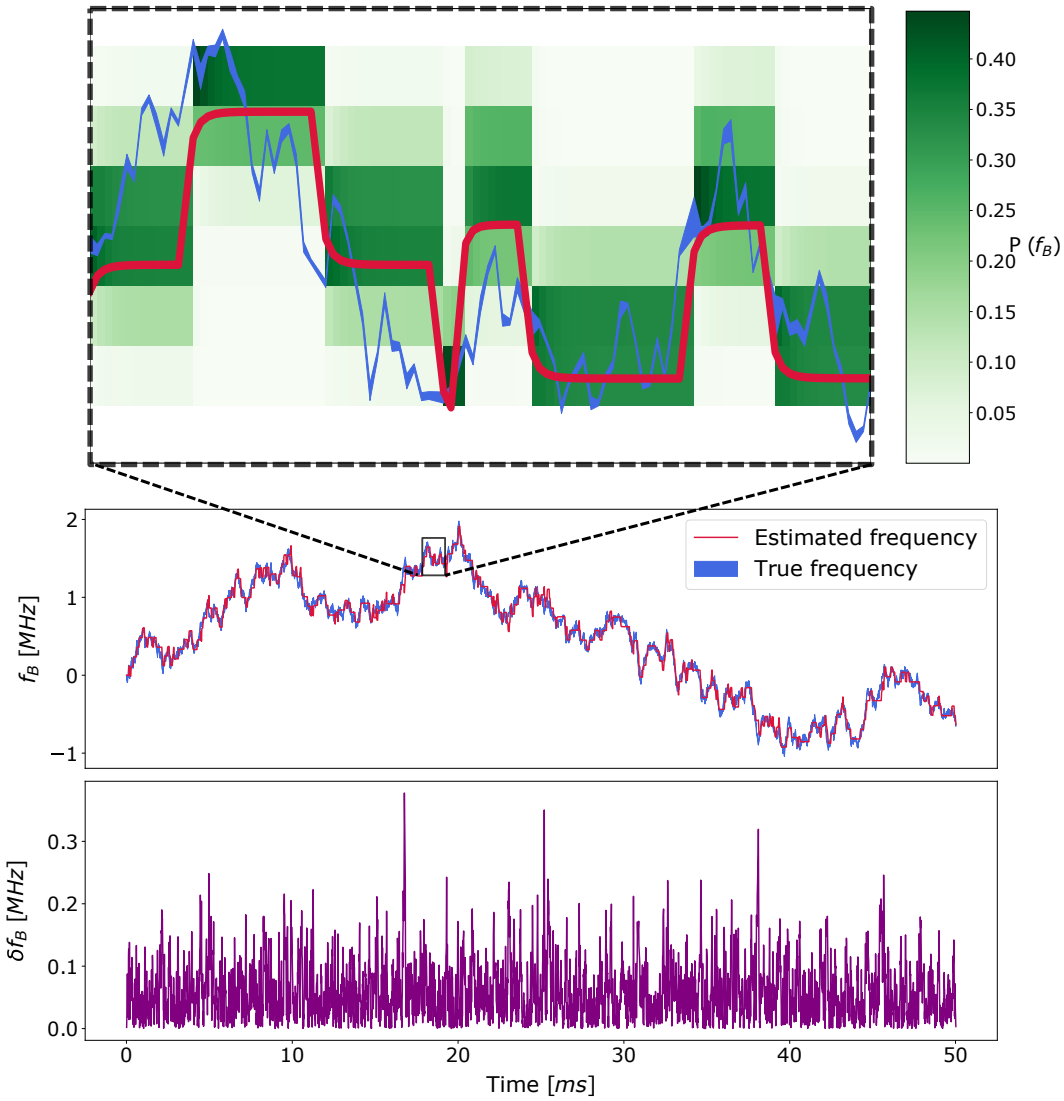

Figure 3. (a) example Gaussian tracking run, with an MSE of 0.078 MHz/ms, $\kappa = 10$ MHz $\mathrm{Hz}^{1/2}$, $t_{oh} = 10\mu s$. Bottom subplot shows the difference between the estimation and ground truth, green colour indicates the probability distribution calculated by the Bayesian update

We chose to vary overhead time and prediction coefficient, κ, as they had been used to benchmark previous comparisons of adaptive versus non-adaptive schemes [22]. In this way, we could test the robustness of the methods, as these are also not parameters we have control over in an experiment. From subplots 4b and c, it can be seen that the mean squared error tends to increase as these variables increase. We also see that the Gaussian method is overall more likely to break down at larger values of κ and overhead time, than the smaller values. This is because, at larger values of both these variables, we can fit in fewer measurements per change of magnetic field. We have less time to narrow in on the correct frequency, and this puts the Gaussian method at a

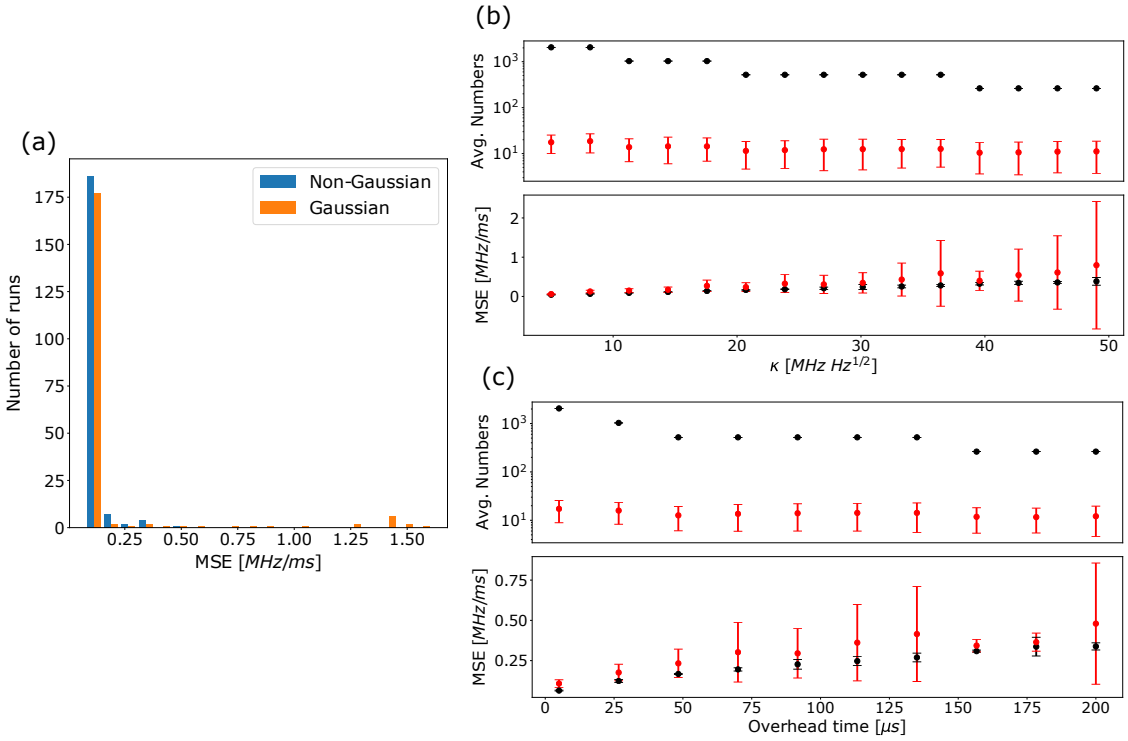

Figure 4. (a) direct comparison of 200 tracking runs on the same set of 'true' Larmor frequency values, for both Gaussian approximated and non-approximate algorithms. A single run is described in figure 3 (b) sweeping prediction coefficient, $\varkappa$, and performing a statistical comparison between the two methods. Each data point is 1000 tracking runs, overhead time fixed at $10\mu s$ (c) similarly sweeping overhead time, with κ set to 10 MHz $\text{Hz}^{1/2}$. Again, each data point is 1000 runs.

disadvantage, since we cannot factor in the potential computational time reduction in the simulation.

The average numbers is a reflection of the Bayesian update computation time. In the case of the original method, this is mostly in the thousands. Note that these numbers are not constant as the method discretises of the Larmor frequency probability distribution with a grid whose resolution is proportional to the number of sensing times N. Moreover, N varies depending on κ and t_{oh} since the maximum sensing time is optimised according to equation 14 of [22]. For the Gaussian method though, the average number of parameters is around eight or nine for this length of tracking run, suggesting a substantial reduction in computational complexity and thus a gain in computational time. The average number of parameters is dragged up by the initial sensing, which initially requires hundreds of Gaussian parameters to describe it. The tracking uses mostly three, sometimes six, and rarely nine or more parameters, indicating that a single Gaussian peak, with only occasionally a second or third, is largely sufficient and delivers adequate performance. This translated to a reduction in the number of operations required to track a changing magnetic field by a factor of 20. Therefore, one would expect this Gaussian-approximated computation to be 20 times faster than the

non-approximate method in practice.

5. Conclusion

In this paper, we have simulated and tracked a fluctuating magnetic field via Ramsey experiments on an NV centre. Ramsey measurements were optimised through an adaptive Bayesian update protocol, with Gaussian approximation of all probability distributions. A comparison of the Gaussian-approximated with the original protocol revealed that the approximation yielded potentially significant increases in computation speed, with the number of operations involved in calculating the Bayesian update and adaptive measurement inputs decreased by factor 20. As expected, the approximation performed slightly poorer at tracking than the non-approximated protocol. However, under the simulation parameters detailed here, the fail rate of the tracking increased only by around 5%.

This protocol could find applications in sensing settings where one needs to track a fluctuating signal with a statistics that is, at least approximately, known in advance. For example, one could use an NV centre in a nanodiamond to monitor temperature drift inside a living cell [15,16]. Measuring temperature is an integral part of studying energy metabolism [30] or developmental processes [31,32]. One other issue with nanodiamonds, is that while moving in a fluid medium, they rotate considerably, often very rapidly. Our protocol could be extended to track this rotation with minimal resource consumption.

Another possible application is in experiments with levitated nano-diamond, where our technique could be used for fast tracking of rotation. In these experiments, the nanodiamond containing NV centres is held in place translationally using ion traps [33,34], optical traps [35] or magnetic traps [36]. The librational (rotational) frequencies of trapped nanodiamond vary from 100s [37] of Hz to 1GHz [38]. For the lower frequencies, in which tracking resolutions of 1ms per data point are suitable, our method could be directly applied. However, the method could also be used in conjunction with ac magnetomety [39, 40], using spin-echo instead of Ramsey measurements, to achieve tracking of faster, periodic librations.

Tracking provides the information required for realigning the nanodiamond orientation, for whichever feedback mechanism the traps use. This feedback mechanism could potentially be 3D Helmholtz coils in the case of ion traps or optical traps. For a magnetic trap, it has been proposed [41] that the diamond orientation be confined with an electrode using the dielectric force on the non-spherical diamond. Though this method does not conventionally require feedback, tracking may nonetheless prove useful in testing the confinement.

Finally, the Gaussian-approximation described here could be applied to track a quantum signal, such as the magnetic field arising from a bath of nuclear spins surrounding a central electron spin. Previous theoretical work has shown that, by adaptively tracking the fluctuating nuclear magnetic field and narrowing its distribution through the back-action of the quantum measurement process, one can considerably

extended the coherence time of the central spin [42]. The protocol described here can reduce the computational complexity of this task, enabling faster and more precise tracking.

Acknowledgments

The authors would like to thank Eleanor Scerri for helpful discussions. This project is supported by the Engineering and Physical Sciences Research council through grant EP/S000550/1, and by a Weizmann-UK Joint Research Program grant. K.C. acknowledges studentship funding from EPSRC under grant no. EP/L015110/1. Y.A. is supported by the Royal Academy of Engineering under the Research Fellowship scheme RF201617/16/31.

References

[1] Degen C L, Reinhard F and Cappellaro P 2017 Rev. Mod. Phys. **89** 035002
[2] Rondin L, Tetienne J P, Hingant T, Roch J F, Maletinsky P and Jacques V 2014 Rep. Prog. Phys. **77** 056503
[3] Levine E V, Turner M J, Kehayias P, Hart C A, Langellier N, Trubko R, Glenn D R, Fu R R and Walswortho R L 2019 Nanophotonics **8** 1945
[4] Wrachtrup J and Finkler A 2016 J. Magn. Reson. **269** 225
[5] Balasubramanian G, Neumann P, Twitchen D, Markham M, Kolesov R, Mizuochi N, Isoya J, Achard J, Beck J, Tissler J, Jacques V, Hemmer P R, Jelezko F and Wrachtrup J 2009 Nature Materials **8** 383–387 ISSN 14764660
[6] Bar-Gill N, Pham L M, Jarmola A, Budker D and Walsworth R L 2013 Nature Communications **4** 1–6 ISSN 20411723
[7] Kucsko G, Maurer P C, Yao N Y, Kubo M, Noh H J, Lo P K, Park H and Lukin M D 2013 Nature **500** 54–58
[8] Trusheim M E and Englund D 2016 New J.Phys. **18** 123023
[9] Broadway D A, Johnson B C, Barson M S J, Lillie S E, Dontschuk N, McCloskey D J, Tsai A, Teraji T, Simpson D A, Stacey A, McCallum J C, Bradby J E, Doherty M W, Hollenberg L C L and Tetienne J P 2019 Nano Lett. **19, 7** 4543–4550
[10] Dolde F, Fedder H, Doherty M W, Nbauer T, Rempp F, Balasubramanian G, Wolf T, Reinhard F, Hollenberg L C L, Jelezko F and Wrachtrup J 2011 Nature Physics **7** 459463
[11] Michl J, Steiner J, Denisenko A, Bulau A, Zimmermann A, Nakamura K, Sumiya H, Onoda S, Neumann P, Isoya J and Wrachtrup J 2019 Nano Lett. 2019 **19** 4904–4910
[12] Schirhagl R, Chang K, Loretz M and Degen C L 2014 Annu. Rev. Phys. Chem. **65** 83
[13] Chipaux M, van der Laan K J, Hemelaar S R, Hasani M, Zheng T and Schirhagl R 2018 Small **14** 1704263
[14] McGuinness L P, Yan Y, Stacey A, Simpson D A, Hall L T, Maclaurin D, Prawer S, Mulvaney P, Wrachtrup J, Caruso F, Scholten R E and Hollenberg L C L 2011 Nature Nanotech **6** 358363
[15] Choi J, Zhou H, Landig R, Wu H Y, Yu X, Kucsko G, Maurer P, Needleman D, Samuel A D T, Park H and Lukin M D 2020 Proceedings of the National Academy of Sciences **117,26** 201922730
[16] Yukawa H, Fujiwara M, Kobayashi K, Kumon Y, Miyaji K, Nishimura Y, Oshimi K, Umehara Y, Teki Y, Iwasaki T, Hatano M, Hashimoto H and Baba Y 2020 Nanoscale Adv. **2** 1859–1868
[17] Lumino A, Polino E, Rab A S, Milani G, Spagnolo N, Wiebe N and Sciarrino F 2018 Phys Rev. Applied **10** 044033
[18] Zhang Y H and Yang W 2018 New J. Phys. **20** 093011

[19] Dinani H T, Berry D W, Gonzalez R, Maze J R and Bonato C 2019 Phys. Rev. B **99** 125413
[20] Cappellaro P 2012 Phys. Rev. A **85** 030301(R)
[21] Van Den Berg E 2020 Efficient Bayesian phase estimation using mixed priors (*Preprint* `2007.11629v1`)
[22] Bonato C and Berry D W 2017 Phys Rev A **95** 052348
[23] Santagati R, Gentile A A, Knauer S, Schmitt S, Paesani S, Granade C, Wiebe N, Osterkamp C, McGuinness L P, Wang J, Thompson M G, Rarity J G, Jelezko F and Laing A 2019 Phys. Rev. X **9** 021019
[24] Bonato C, Blok M S, Dinani H T, Berry D W, Markham M L, Twitchen D J and Hanson R 2016 Nature Nanotech **11** 247252
[25] Doucet A, Smith A, de Freitas N and Gordon N 2001 Sequential Monte Carlo Methods in Practice Information Science and Statistics (Springer New York) ISBN 9780387951461
[26] Kraus H, Soltamov V A, Fuchs F, Simin D, Sperlich A, Baranov P G, Astakhov G V and Dyakonov V 2014 Scientific Reports **4** ISSN 20452322
[27] Yan F F, Yi A L, Wang J F, Li Q, Yu P, Zhang J X, Gali A, Wang Y, Xu J S, Ou X, Li C F and Guo G C 2020 npj Quantum Information **6** 1–6 ISSN 20566387
[28] Vandersypen L M and Chuang I L 2004 Rev. Mod. Phys. **76** 1037–1069 ISSN 00346861
[29] Kullback S and Leibler R A 1951 Ann. Math. Statist. **22,1** 79–86
[30] Meyer C W, Ootsuka Y and Romanovsky A A 2017 Frontiers in Physiology **8** 520
[31] Begasse M L, Leaver M, Vazquez F, Grill S W and Hyman A A 2015 Cell Rep. **10** 647
[32] Chong J, Amourda C and Saunders T E 2018 J. R. Soc. Interface. **15** 20180304
[33] Conangla G P, Schell A W, Rica R A and Quidant R 2018 Nano Lett. **18** 3956–3961
[34] Kuhlicke A, Schell A W, Zoll J and Benson O 2014 Appl.Phys.Lett. **105** 073101
[35] Frangeskou A C, Rahman A T M A, Gines L, Manda S, Williams O A, Barker P F and Morley G W 2018 New J.Phys. **20** 043016
[36] Hsu J F, Ji P, Lewandowski C W and DUrso B 2016 Sci. Rep. **6** 30125
[37] Delord T, Huillery P, Nicolas L and Htet G 2020 Nature **580** 5659
[38] Reimann R, Doderer M, Hebestreit E, Diehl R, Frimmer M, Windey D, Tebbenjohanns F and Novotny L 2018 Physical Review Letters **121** 033602 ISSN 10797114
[39] Maze J R, Stanwix P L, Hodges J S, Hong S, Taylor J M, Cappellaro P, Jiang L, Dutt M V, Togan E, Zibrov A S, Yacoby A, Walsworth R L and Lukin M D 2008 Nature **455** 644–647 ISSN 14764687
[40] De Lange G, Ristè D, Dobrovitski V V and Hanson R 2011 Phys.Rev. Lett. **106** ISSN 080802
[41] Pedernales J S, Morley G W and Plenio M B 2019 Motional dynamical decoupling for matter-wave interferometry (*Preprint* `arXiv:1906.00835`)
[42] Scerri E, Gauger E M and Bonato C 2020 New J.Phys. **3,22** 035002