Extended Kalman smoothing of free spin precession signals for precise magnetic field determination

Lutz Mertenskötter[†1], Jasper Riebesehl[†2], Wilhelm Stannat[3], Wiebke Pohlandt[4],

Wolfgang Kilian[4]

[1]   Weierstrass Institute
      Mohrenstr. 39
      10117 Berlin, Germany
      E-Mail: lutz.mertenskoetter@wias-berlin.de

[2]   Technical University of Denmark (DTU)
      DK-2800 Kgs. Lyngby, Denmark
      E-Mail: janri@dtu.dk

[3]   Technische Universität Berlin (TU Berlin)
      Straße des 17. Juni 136
      10623 Berlin, Germany
      E-Mail: stannat@math.tu-berlin.de

[4]   Physikalisch-Technische Bundesanstalt (PTB)
      Abbestr. 2-12
      10587 Berlin, Germany
      E-Mail: wiebke.pohlandt@ptb.de
              wolfgang.kilian@ptb.de

[†]These authors contributed equally to this work.

Extended Kalman smoothing of free spin precession signals for precise magnetic field determination

Lutz Mertenskötter[†], Jasper Riebesehl[†], Wilhelm Stannat, Wiebke Pohlandt, Wolfgang Kilian

Abstract

We present a novel application of the Extended Kalman Smoother (EKS) for high-precision frequency estimation from free spin precession signals of polarized ^{3}He. Traditional approaches often rely on nonlinear least-squares fitting, which can suffer from limited robustness to signal decay and time-dependent frequency shifts. By contrast, our EKS-based method captures both amplitude and frequency variations with minimal tuning, adapting automatically to fluctuations via an expectation-maximization algorithm.

We benchmark the technique in extensive simulations that emulate realistic spin precession signals with exponentially decaying amplitudes and noisy frequency drifts. Compared to least-squares fits with fixed block lengths, EKS systematically reduces estimation errors, particularly when frequencies evolve or signal-to-noise ratios are moderate to high. We further validate these findings with experimental data from a free-precession decay ^{3}He magnetometer.

Our results indicate that EKS-based analysis can substantially improve precision in nuclear magnetic resonance-based magnetometry, where accurate frequency estimation underpins absolute field determinations. This versatile approach promises to enhance the stability and accuracy of future high-precision measurements.

1 Introduction

Spin precession magnetrometry Magnetometry, i.e. the determination of the static magnetic flux density, is an extremely broad field of research in which a wide variety of technical methods are used [1]. The potentially most precise technique which allows to directly trace back the unit of Tesla to the SI unite of the second is the measurement of the free spin precession frequency using either electronic spin resonance (ESR) or nuclear magnetic resonance (NMR). Given the fundamental relation of the Larmor resonance frequency f_{L} and magnetic field strength B_0

$$f_{\mathrm{L}} = \frac{\gamma}{2\pi} B_0, \tag{1}$$

only the substance-specific gyromagnetic ratio γ has to be known to deduce the external magnetic field strength from Larmor precession measurements. The gyromagnetic ratio γ'_p of protons has been known for decades with a relative uncertainty well below 10^{-6} and has more recently been determined with an uncertainty even below 10^{-8} [2]. Consequently, the classical proton free induction decay (FID) measurement – aside from its widespread use in chemical analysis and medical imaging – has naturally become a prominent method for high-precision magnetic field measurements [3, 4]. With recent advances in measuring the magnetic moment of ^{3}He$^+$ and calculation of the diamagnetic shielding, the gyromagnetic ratio γ'_h of gaseous ^{3}He is now known with even higher precision, achieving a relative uncertainty below 10^{-9} [5, 2]. Given these precisely known scaling factors for γ'_p and γ'_h,

the overall accuracy of magnetic field measurements now critically depends on the precision of the frequency determination from NMR signals. In this work we employ a new frequency tracking technique, based on frequency domain Kalman smoothing to improve this precision. Consequently, ^{3}He NMR allows for the most precise field determination in future when all systematic errors have to be accounted correctly [6]. These systematic errors are not further discussed in this work.

A wide range of analytical methods – and their variants – have been proposed for this task, including frequency-domain analysis using fast Fourier transformation (FFT) [7, 8], zero-crossing counting [9], singular value decomposition [10], and separable nonlinear least-squares analysis using the variable projection method [11]. More recently, machine learning-based approaches have also been explored [12, 13].

To our knowledge, the nonlinear least-squares method remains the benchmark for estimating frequencies in ultra-high-precision NMR signals in ^{3}He-^{129}Xe-co-magnetometry measurements [14, 15, 16]. Our frequency domain Kalman smoothing method is a promising candidate for a more precise estimator, as it effectively filters out the majority of the noise, and – unlike the least-squares method – is able to assimilate the entire time-series into one estimation. In the following chapters we give an in-depth comparison of this method to the least-squares method on a host of simulated data and real measurements.

The Kalman smoother [17] not only estimates the hidden state $\mathbf{x}_k$, in our case frequency $f(t)$ and amplitude $A(t)$, but also computes its covariance matrix $\mathbf{P}_k$. The recursive update of $\mathbf{P}_k$ provides a robust measure of the confidence in the state estimates, which is crucial in sensor applications, where the uncertainty is of paramount importance.

Kalman smoothers have found numerous applications in post-processing of sensor data, geophysical signal reconstruction, and time series forecasting. For instance, in biomedical applications, Kalman smoothers improve electroencephalography (EEG) [18]. In magnetometry, Kalman filters have been successfully used to significantly improve the resolution of atomic sensors [19].

2 Frequency estimation methods

2.1 The Kalman smoother

The Kalman smoother is a recursive algorithm designed to estimate the state of a dynamical system from noisy observations. It is formulated within a state-space framework, which separates the system dynamics from the measurement process. In the nonlinear case, the system dynamics are described by the *state equation*

$$\mathbf{x}_k = \phi(\mathbf{x}_{k-1}, t_{k-1}) + \mathbf{w}_k, \tag{2}$$

where $\mathbf{x}_k$ represents the state vector at time t_k, and $\phi(\mathbf{x}, t)$ is a (potentially nonlinear) function. The *process noise* $\mathbf{w}_k$ is white noise with covariance Q. The measurement process, is in turn described by the *measurement equation*

$$\mathbf{y}_k = h(\mathbf{x}_k, t_k) + \mathbf{v}_k, \tag{3}$$

where $h(\mathbf{x}, t)$ is a nonlinear function mapping the state to the observed measurements $\mathbf{y}_k$, and the *measurement noise* $\mathbf{v}_k$ is white noise with covariance R.

The Kalman Smoother then estimates the hidden state $\mathbf{x}_k$ at time t_k given all N measurements, i.e. it computes $p(\mathbf{x}_k|\mathbf{y}_{1:N})$. It does so, such that it minimizes the mean squared error of the estimated state vs. the true state. Strictly, this applies only when $\phi(\mathbf{x}, t)$ and $h(\mathbf{x}, t)$ are linear. Yet, the extended Kalman smoother (EKS) used here performs very well if the linearizations are as adequate as in our case. The Kalman Filter in turn estimates $\mathbf{x}_k$ at time t_k given only the measurements up to time t_k, i.e. it computes $p(\mathbf{x}_k|\mathbf{y}_{1:k})$, as needed in real-time applications.

2.2 The model

In this work we will take a different approach in modeling the atomic sensor, that strays from the traditional conceptualization in which $\phi(\mathbf{x}, t)$ is typically a discretization of a differential equation – 'the dynamics' – and $h(\mathbf{x}, t)$ is a function modeling the relationship of the hidden state to the measured quantity – 'the measurement'. Rather we focus on the statistics of the data, namely that (2) and (3) both have additive white noise to derive a more versatile model.

To extract the Larmor frequency of the spin precession signal, we adapted the approach of La Scala et al. [20], which applies an extended Kalman filter (EKF) to track a harmonic signal with low signal-to-noise ratio (SNR) and slowly varying frequency $f(t)$ as is typical in spin precession measurements. The main idea is to derive the analytic expression for the Fourier coefficients of the harmonic signal (4) and use this as the measurement equation (3) to compare the model prediction to the Fourier coefficients of the signal, rather than using the time-domain model (4) itself. To this end, the signal model is separated into non-overlapping blocks, each of duration $T_{\mathrm{bl}} = \Delta t N_{bl}$ with N_{bl} sample points (separated by the sampling period Δt) in which the frequency $f_0 + \delta f_k$ and amplitude A_k are assumed to be constant

$$z_k(t_n) = A_k \cos\{2\pi [f_0 + \delta f_k] t_n + \varphi_k\} + \eta_k(t_n), \tag{4}$$

with additive Gaussian noise $\eta(t)$ of variance σ_η^2. $\varphi_k = \Sigma_{i=0}^{k-1} 2\pi \delta f_k T_{bl} + \varphi_0$ then is the accumulated phase at the beginning of the $k-$th block and f_0 is the part of the frequency that fits integer periods into the block, i.e. $\cos(2\pi(f_0 + \delta f_k)T_{bl}) = \cos(2\pi \delta f_k T_{bl})$. The discrete Fourier transform of each block then reads

$$\mathcal{F}z_k(f_m) = \sum_{n=0}^{N_{bl}-1} z_k(t_n)e^{-i2\pi f_m t_n} = h_m(\mathbf{x_k}) + v_{m,k} \tag{5}$$

$$= \frac{A_k}{2}\left(e^{i\varphi_k}\frac{1 - e^{i2\pi \delta f_k T_{bl}}}{1 - e^{i2\pi(f_0 + \delta f_k - f_m)\Delta t}} + e^{-i\varphi_k}\frac{1 - e^{-i2\pi \delta f_k T_{bl}}}{1 - e^{-i2\pi(f_0 + \delta f_k + f_m)\Delta t}}\right) + v_{m,k}, \tag{6}$$

where , $f_m = \frac{m}{T_{bl}}$, $m = 1, ..., N_{\mathrm{bl}}$. Since the Fourier transform is linear, the Gaussian white noise $\eta_k(t)$ transforms into independent Gaussian noise components $v_{m,k} = \mathcal{F}\eta_k(f_m)$ in the Fourier coefficients, with variance of the independent real and imaginary parts

$$\sigma_v^2 = \frac{2}{N_{\mathrm{bl}}}\sigma_\eta^2 \tag{7}$$

that is significantly reduced compared to that of η_k. Thus, the $\mathcal{F}z_k(f_m)$ themselves form a time series with white Gaussian noise components, and are therefore compatible with the EKS framework.
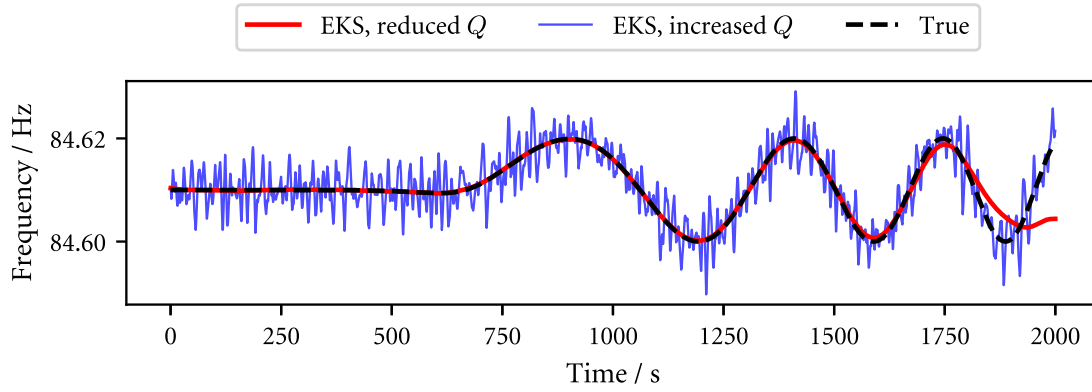
Figure 1: EKS tracking of a simulated signal with time-dependent frequency variation illustrating the variance/stiffness trade-off. The true frequency (black dashed line) is followed by the model with artificially large Q (blue) even when the frequency changes rapidly in time. Meanwhile, the model with artificially small Q (red) exhibits a maximum slew rate that is at some point exceeded by the rate of change of the frequency. Conversely, where the frequency is constant, the model with small Q has much smaller variance and outperforms the model with large Q.

The evolution of the hidden variables is modeled by a typical augmented Kalman filter approach,

$$\mathbf{x}_k = \begin{bmatrix} A_k \\ \Delta A_k \\ \varphi_k \\ \delta f_k \\ \Delta \delta f_k \end{bmatrix} = \begin{bmatrix} A_{k-1} + \Delta A_{k-1} \\ \Delta A_{k-1} \\ \varphi_{k-1} + 2\pi \delta f_{k-1} \\ \delta f_{k-1} + \Delta \delta f_{k-1} \\ \Delta \delta f_{k-1} \end{bmatrix} + \mathbf{w}_k, \tag{8}$$

with independent integrated random walks in the amplitude and frequency offset, while the phase is only an auxiliary variable that is computed as the numerical integral of the frequency. The process noise $\mathbf{w}_k$ is white noise with covariance

$$Q = \mathsf{diag}\left(\begin{bmatrix} 0 & Q_{\Delta A} & 0 & 0 & Q_{\Delta \delta f} \end{bmatrix} \right), \tag{9}$$

such that the variables A, δ_f, and φ that feed into the model through (6) are integrated and twice-integrated white noise, respectively. This choice of integrated model for the state evolution reduces the variance of the tracking, but also reduces its slew rate, as the integration amplifies slow components of the model while suppressing fast ones. The filter then compares the estimated state not to the signal directly in the time-domain, but rather to its Fourier coefficients (6), which serve as the measurement model

$$\mathbf{y}_k(\mathbf{x}_k) = \begin{bmatrix} h_{M-L}(\mathbf{x}_k) \\ \vdots \\ h_{M+L}(\mathbf{x}_k) \end{bmatrix} + \begin{bmatrix} v_{M-L,k} \\ \vdots \\ v_{M+L,k} \end{bmatrix}. \tag{10}$$

Rather than operating on all Fourier coefficients, the EKS only incorporates a range of frequencies around $f_0 = f_M$, controlled by the parameter L, which determines how many spectral bins contribute to the estimation. A larger L improves robustness to frequency drifts but introduces more noise and thus reduces precision, while a smaller L improves precision but reduces adaptability. By tuning L, this method enables reliable frequency tracking even in extreme noise conditions. Frequency drifts in DC magnetometry systems are typically very small, and it is sufficient to choose $L = 1$.

The static parameters of the model

$$\theta = \{Q, R, \delta f_0, A_0\}, \tag{11}$$

that is, the process noise covariance Q, the measurement noise covariance R, and the starting values δf_0 and A_0 are all optimized separately using an algorithm based on expectation maximization. Details on the algorithm can be found in the supplementary material. An important strength of this model is that (3) and (2) are entirely free of system-dependent parameters that would need to be known in order to perform the tracking. The only static parameters in the Kalman smoother itself are noise-covariances and starting values in (11), which are inherently present in any state-space filtering problem. Optimization of the process covariances Q is particularly vital to the performance of the EKS, as the tracking becomes too stiff when it is too small, while a Q that is too large results in a large variance of the estimates, as is illustrated in Fig. 1.

2.3 Least Square Analysis as Reference Method

To separate the nonlinear signal dependence on frequency f and phase offset φ_0 at $t = 0$ for a signal obeying $S(t) = A \sin(2\pi ft + \varphi_0)$ within the least squares framework, an equivalent model function, commonly referred to as the sine-cos-fit (SCF) is used

$$S(t) = A_s \cos(2\pi ft) + A_c \cos(2\pi ft) + C_0, \tag{12}$$

which allows to calculate the signal amplitude of the corresponding data set by $A = \sqrt{A_s^2 + A_c^2}$ and the phase at $t = 0$ by $\varphi_0 = \arctan(A_c/A_s)$ and accounts for a constant signal offset C_0. Fitting data to (12) permits the use of the variable projection (VP) method [11] wherein the linear terms are estimated separately from the nonlinear frequency dependence. The method alternates between a linear least-squares fit of the linear parameters A_s, A_c, and C_0, and a non-linear least-squares fit of the frequency f using the Levenberg-Marquardt algorithm. The uncertainties of the four parameters u_f, u_{Ac}, u_{As}, and u_{C0} are obtained by scaling the covariance matrix with the mean squared errors of the residuals. The covariance matrices for f and the other three parameters are obtained from the two separate fitting procedures of the VP, respectively. The start parameter for f for the fitting process, is determined beforehand by performing a FFT on the complete dataset and searching for the frequency bin with the highest peak in a given range. The start values for the parameters A_s, A_c, and C_0 are set to 1, since their estimation via linear least squares is largely insensitive to the choice of starting values.

As the function (12) does not take into account the exponential decay of the amplitude, the entire time domain data is divided into non-overlapping blocks of equal length. These blocks are chosen to be short enough that the signal amplitude can be assumed constant, yet long enough to sufficiently reduce the variance of the least-squares estimates. As the uncertainty (i.e., standard deviation) of the fitted parameters scales inversely with the square root of the number of sample points, assuming independent, homoscedastic Gaussian noise [21], in principle a longer block length should reduce the uncertainties. However, in our experimental data this is compromised by the signal amplitude decay and frequency drifts due to a drift in B_0.

3 Comparing the methods on simulated data

To quantitatively determine which of the methods can provide better frequency estimates in free precession decay (FPD) signals, we turn to a broad simulation study. The characteristic properties

of such a signal are its exponential amplitude decay and the time-dependence of the true frequency, which is in turn determined by the fluctuations of the true magnetic field strength. To represent this class of signals, we use the phenomenological model

$$y(t_k) = A_0 \, e^{-t_k/T_2^*} \sin\{2\pi f(t_k)t_k + \varphi_0\} + \eta(t_k),\tag{13}$$

where A_0 is the initial signal amplitude, T_2^* the effective time constant of the transverse spin-relaxation leading to loss of phase-coherence and thus to exponential signal decay, $f(t)$ the time-varying frequency, φ_0 an initial phase and finally, $\eta(t)$ subsumes all noise sources accumulated within the measurement data. The index $k = 0 \dots N$ indicates a discretely sampled signal with sampling rate f_s.

We generated simulated signals for different value combinations of these parameters to evaluate the methods on. We are mainly interested in investigating the influence of A_0 and the strength of frequency fluctuations, while T_2^* and the variance of η are fixed for a certain experimental setup.

The stochastic nature of field fluctuations is emulated by modeling the time dependence of the frequency with a random walk plus a constant offset f_c. An instance of a random walk is generated by cumulatively summing discrete random increments δf drawn from a Gaussian distribution with variance $\sigma_{\delta f}^2$. The magnitude of the frequency drift is quantified by the diffusion constant D of the random walk, with

$$D = \frac{\sigma_{\delta f}^2 f_s}{2}.\tag{14}$$

The second parameter to be varied is the initial amplitude A_0, giving an initial SNR

$$\mathrm{SNR}_0 = \frac{A_0^2}{2\sigma_\eta^2}.\tag{15}$$

For simplicity, $\eta(t)$ is white Gaussian noise with fixed variance σ_η^2. Real experimental signals, however, can also contain low-frequency magnetic field perturbations and additional noise signals with frequency components that do not originate from the spin precession, e.g. from power-line or setup vibrations. Because the Kalman smoother method effectively applies a band-pass filter these deviations from the model can be neglected unless they are close in frequency to the signal itself.

Due to the stochastic nature of the simulated signals, multiple repetitions were simulated per set of parameters and the ensemble means are reported. The fixed parameters of this study are shown in the supplementary material. They were selected to closely resemble those of the measured data. We generated an ensemble of signals according to (13) for a range of values of D and SNR_0. For each instance, both the SCF and the EKS methods were applied.

For the SCF, a broad range of block sizes were evaluated for each D and SNR_0 independently. Ultimately, the block size which results in the estimate with the lowest MSE was selected as the best estimate (supplementary material 2.4). Note that this is the best case for the SCF estimate, and this selection method is only available in simulation.

The application of the EKS includes the full parameter optimization routine using the EM algorithm. Here, we intentionally kept the initial parameters for the optimization fixed to emulate zero prior knowledge about the signal. This demonstrates the method's robustness as the optimization's convergence does not critically depend on the initial parameters. The only parameter which has to be adapted by simulation studies is $T_{bl,\mathrm{EKS}} = 4.5$ s providing valid uncertainty bounds for the EKS (supplementary material 2.3).
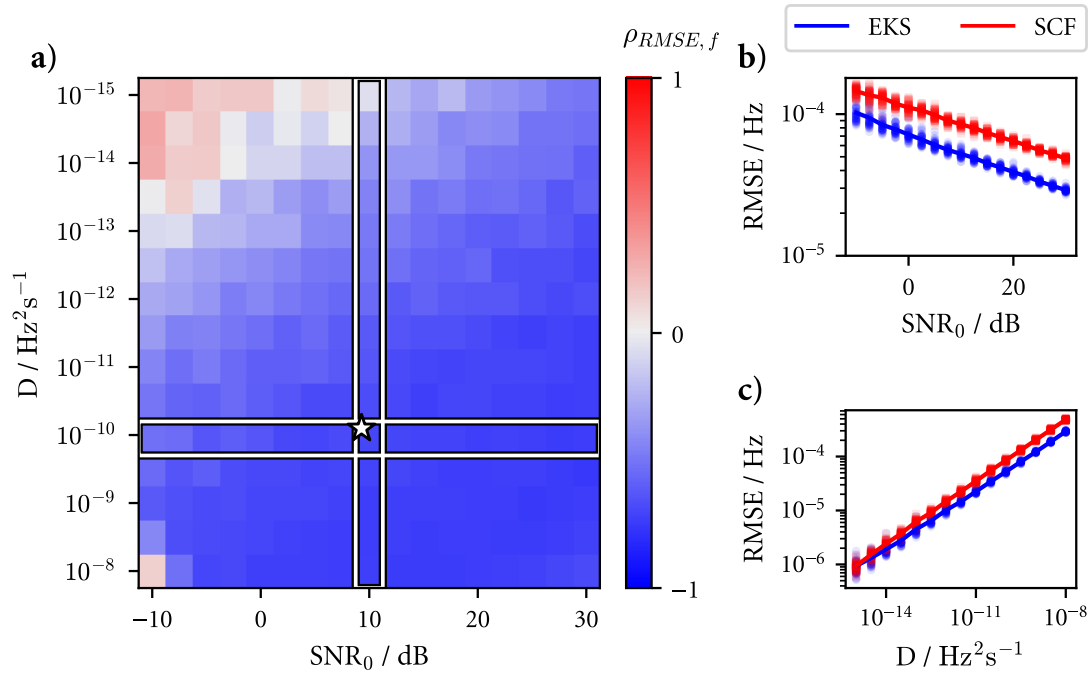
Figure 2: Quantitative comparison of the SCF against the EKS in a simulation study. **a)** $\log_2$ of the mean ratio of the frequency estimate errors. Blue indicates a better performance of the EKS over the SCF, red the opposite. The star marker indicates a rough estimate of our experimental conditions. **b)** Horizontal slice through a). The RMSEs of individual simulation runs are displayed as markers, the solid lines indicate the mean. **c)** displays a vertical slice.

3.1 Results

The results of this study are summarized in Fig. 2. The frequency estimation errors of both methods are compared by calculating the ratio of their root mean squared errors (RMSE). Using the expression

$$\rho_{RMSE,f} = \log_2 \frac{\text{RMSE}_{f,EKS}}{\text{RMSE}_{f,SCF}}, \tag{16}$$

a comparative measure is obtained. $\rho_{RMSE,f}$ is zero when the errors are equal, becomes negative for lower error of the EKS frequency estimate, and positive for the opposite case. For a large portion of the parameter space in Fig. 2 a), $\rho_{RMSE,f}$ is close to -1, indicating an improved frequency estimate of the EKS by a factor of 2. For low SNR and very low frequency drift, $\rho_{RMSE,f}$ becomes positive, indicating that the SCF estimate should be used in these regions. This seems consistent with the observation that the SCF operates closer to the Cramér–Rao lower bound (CRLB) for no frequency drift (supplementary material 2.5). Continuously reducing the frequency drift asymptotically approaches this regime.

We also compared the amplitude estimates, where the EKS significantly outperforms the SCF. This is to be expected, as the models don't consider the exponential signal decay in order to stay free of unknown parameters. For the longer SCF block lengths, this model mismatch leads to a stronger degradation of estimation performance in the time-dependent amplitude (supplementary material 2.2).
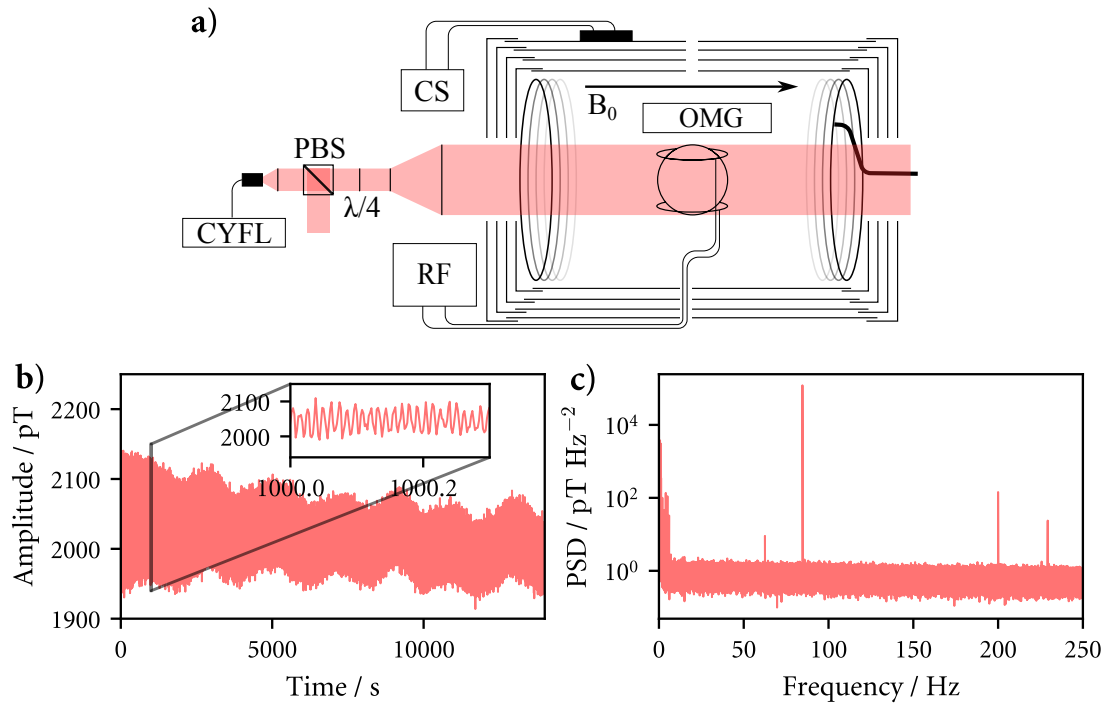
a)

b)

c)

Figure 3: **a)** Schematic of the setup as used for ^{3}He spin precession measurements. For MEOP, light from a laser (CYFL) via some optics (PBS and a $\lambda/4$) is shone on the metastable ^{3}He generated by a radio-frequency (RF) discharge. A current source (CS) attached to a coil inside the four-layer shield generates the B_0 field. The OMG sensor is used for readout. **b)/c)** Experimental gradiometer signal from the OMG sensor. **b)** Time domain of the signal. The inset shows the fast oscillations caused by the ^{3}He spin precession. **c)** Power spectral density of the signal.

4 Application to experimental data

Finally, we assess both methods on real experimental data. For this ^{3}He spin precession measurements were performed by a setup as sketched in Fig. 3 a). A glass cell of 3 cm diameter filled with 10 torr ^{3}He was positioned within a mu-metal shield which allows to generate a constant holding field B_0 in the μT-range and a pulsed B_1-field ($\perp B_0$) to induce the $\pi/2$ spin flip. Signal detection of the FPD was performed alike shown in [22], applying a commercial dual cell optical pumped rubidium magnetometer in gradiometer arrangement (OMG). To maintain sufficiently high signal metastability exchange optical pumping (MEOP) technique [23] was used to polarize nuclear spins of the ^{3}He atoms.

Due to the $T_2^\star$ time of about 50 minutes, signals over a time-span of nearly four hours could be measured as shown in Fig. 3 b). The inset shows a zoomed-in view of the sinusoidal signal stemming from the precessing ^{3}He magnetization which is generating a time-varying field difference seen by the two Rb cells. The ≈ 50 pT amplitude can be calculated to stem from $\approx 8\%$ ^{3}He polarization [24]. The large offset of ≈ 2 nT reflects the background field gradient of $\approx 1\,\mathrm{nT/cm}$ explaining the relatively short $T_2^\star$ time as compared to measurements obtained in a large shielded room [14] and common in small shielding environment [25]. The dominant peak at $\approx 84.6\,\mathrm{Hz}$ in the power spectrum in Fig. 3 c) stems from the FPD signal corresponding to a constant background field of $\approx 2.61\,\mathrm{\mu T}$.

Both analysis methods were applied directly to the OMG time-domain signal without any pre-processing. Again, $T_{bl,\mathrm{EKS}} = 4.5$ s was used for the EKS as determined in the simulation study. For the SCF, multiple block lengths were evaluated as no clear criterion is available to select the optimal block length in the RMSE sense purely from data. The simulation study indicated that the optimal block length for a
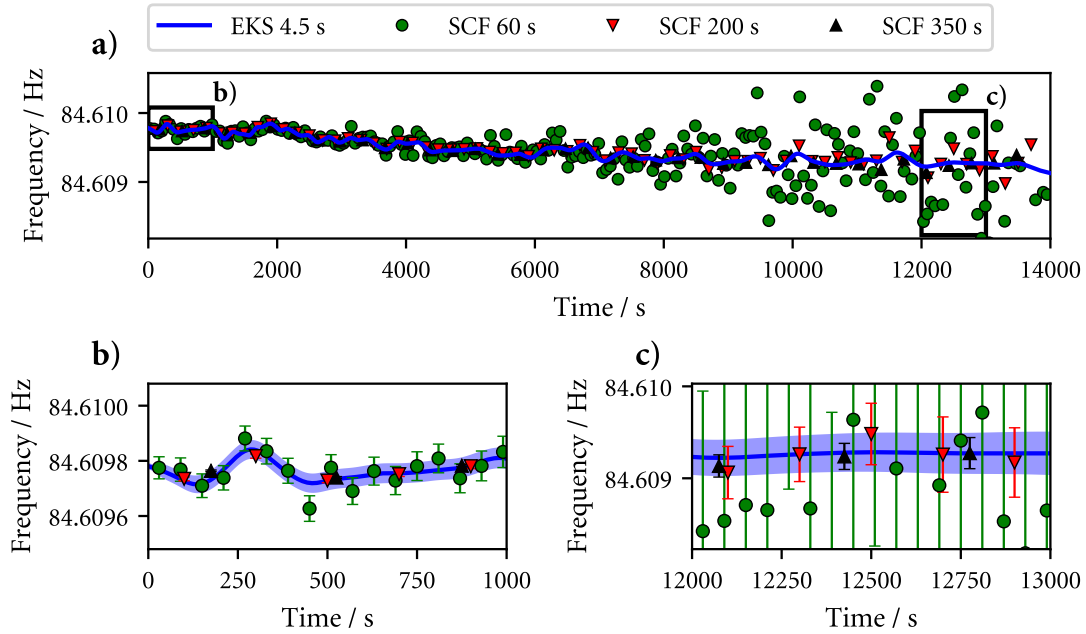
Figure 4: Frequency tracking on experimental data, comparison of the methods. Uncertainty intervals are two standard deviations. **b)** and **c)** are zoomed insets of **a)**.

signal with the given parameters should be close to $200\,\mathrm{s}$ (supplementary material 2.4).

The frequency estimates over time are shown in Fig. 4. In a), the estimates over the full time series are shown without error bars for visibility. The SCF estimates of any block lengths are scattered around the EKS estimate with increasing variance at larger times. This is caused by the exponential amplitude decay and subsequent decrease in SNR. Fig. 4 b) shows a zoom to the beginning of the measurement, where the SNR is high. Here, the effect of a reduced time resolution when selecting very large block lengths for the SCF becomes apparent. For the smaller block length of 60 s the tracking of fluctuation dynamics is possible, however, resulting in higher uncertainty margins, a typical bias-variance tradeoff. Note that the SCF with the largest block length applied here can not resolve the upwards swerve at $\approx 300\,\mathrm{s}$ while the uncertainty interval is much smaller than this variation, clearly seen with the shorter block lengths and the EKS. This hints at unreliable uncertainty estimates of the SCF for too large block lengths, however, with no analytical measure to judge on this.

In Fig. 4 c) the frequency estimates towards the end of the measurement with low SNR are shown. All results agree within their uncertainty intervals. Notably, the EKS uncertainty have a similar magnitude to SCF uncertainty intervals which use close to 2 orders of magnitude larger block length. An obvious improvement to the SCF method would be the introduction of an SNR-adaptive block length in future work. However, it is not clear which criterion should be used to determine said block length.

These results highlight the flexibility of the EKS, as it does not need to be tuned manually to perform best in a wide range of parameter regimes. All tuning is handled automatically by the EM algorithm.

5   Conclusion

In this work we have demonstrated the use of extended Kalman smoothing for precise frequency estimation in ^{3}He nuclear spin precession signals aimed to be used for absolute field magnetometry.

Through simulation studies, we have shown that the EKS outperforms the nonlinear least squares fit based state of the art method for a wide range of experimental conditions. In addition, we have demonstrated the robustness of the method by applying it to an experimental signal with significant perturbations. For accessibility, we have provided a robust and easy to use implementation of the method.

Data availability statement

The code that supports the findings of this study is available from the corresponding author upon reasonable request.

Supplementary Material

A Mathematical details on the extended Kalman smoother method

In this section, the equations for the Kalman smoothing approach presented in the main text are given for reproducibility. For derivations and more detail please refer to [26].

A.1 Extended Kalman filtering and smoothing equations

The EKS discussed in this work relies on a forward-backward recursion over the signal to be applied. In the first forward recursion, the extended Kalman filter is applied by using the filtering equations given as

$$
\begin{aligned}
\hat{\mathbf{x}}_k &= \phi(\mathbf{x}_{k-1}, t_k) \\
\hat{\mathbf{P}}_k &= J_\phi(\mathbf{x})\Big|_{\mathbf{x}_{k-1}} \mathbf{P}_{k-1} J_\phi(\mathbf{x})^T\Big|_{\mathbf{x}_{k-1}} + \mathbf{Q}_k \\
\mathbf{S}_k &= J_h(\mathbf{x})\Big|_{\hat{\mathbf{x}}_k} \hat{\mathbf{P}}_k J_h(\mathbf{x})^T\Big|_{\hat{\mathbf{x}}_k} + \mathbf{R}_k \\
\mathbf{K}_k &= \hat{\mathbf{P}}_k J_h(\mathbf{x})^T\Big|_{\hat{\mathbf{x}}_k} \mathbf{S}_k^{-1} \\
\mathbf{v}_k &= \mathbf{y}_k - h(\hat{\mathbf{x}}_k) \\
\mathbf{x}_k &= \hat{\mathbf{x}}_k + \mathbf{K}_k \mathbf{v}_k \\
\mathbf{P}_k &= \hat{\mathbf{P}}_k - \mathbf{K}_k \mathbf{S}_k \mathbf{K}_k^T
\end{aligned}
\tag{17}
$$

where $J_\phi(\mathbf{x})\Big|_{\mathbf{x}_k}$ indicates the Jacobian of a function ϕ with respect to $\mathbf{x}$, evaluated at $\mathbf{x}_k$. For $k = 0$ an initial state $\mathbf{x}_0$ and its covariance $\mathbf{P}_0$ are required. In our method, we keep $\mathbf{Q}_k$ and $\mathbf{R}_k$ constant over k.

In the second step, the EKS estimates are calculated using the smoothing equations:

$$
\begin{aligned}
\mathbf{G}_k &= \mathbf{P}_k \left[J_\phi(\mathbf{x})\Big|_{\mathbf{x}_k} \right]^T \hat{\mathbf{P}}_{k+1}^{-1} \\
\mathbf{x}_k^s &= \mathbf{x}_k + \mathbf{G}_k \left[\mathbf{x}_{k+1}^s - \hat{\mathbf{x}}_{k+1} \right] \\
\mathbf{P}_k^s &= \mathbf{P}_k + \mathbf{G}_k \left[\mathbf{P}_{k+1}^s - \hat{\mathbf{P}}_{k+1} \right] \mathbf{G}_k^T
\end{aligned}
\tag{18}
$$

This backwards recursion is initialized at $k = N$ with $\mathbf{x}_N^s = \mathbf{x}_N$ and $\mathbf{P}_N^s = \mathbf{P}_N$. The smoothing estimates $\mathbf{x}_k^s$ for $k = N \ldots 0$ contain the EKS frequency and amplitude estimates. The diagonal elements of the covariance matrices $\mathbf{P}_k^s$ contain the marginal uncertainties of the estimates.

A.2 Robust and fast parameter optimization with the Expectation-Maximization algorithm

The filtering and smoothing equations contain static parameters which need to be optimized for the method to perform optimally. These parameters are the process/measurement noise covariance matrices $\mathbf{Q}$ and $\mathbf{R}$ and the initial state and state covariance, $\mathbf{x}_0$ and $\mathbf{P}_0$.

A well documented algorithm for this purpose is the Expectation-Maximization (EM) algorithm [26, 27]. It iteratively maximizes the likelihood function and can be approximately solved using closed form analytic expressions. Using the auxiliary variables

$$
\begin{aligned}
\mathbf{C}_k &= \mathbf{P}_k^s \mathbf{G}_{k-1}^T + \mathbf{x}_k^s \left[\mathbf{x}_{k-1}^s \right]^T \\
\mathbf{\Sigma}_k &= \mathbf{P}_k^s + \mathbf{x}_k^s \left[\mathbf{x}_k^s \right]^T \\
\mathbf{\Lambda}^{(j+1)} &= \frac{1}{T} \sum_{k=1}^{N} \mathbf{\Sigma}_k - \mathbf{C}_k \left[J_\phi(\mathbf{x}) \Big|_{\mathbf{x}_k^s} \right]^T - J_\phi(\mathbf{x}) \Big|_{\mathbf{x}_k^s} \mathbf{C}_k^T - J_\phi(\mathbf{x}) \Big|_{\mathbf{x}_k^s} \mathbf{\Sigma}_{k-1} \left[J_\phi(\mathbf{x}) \Big|_{\mathbf{x}_k^s} \right]^T \quad (19) \\
\mathbf{\Omega}^{(j+1)} &= \frac{1}{T} \sum_{k=1}^{N} \mathbf{v}_k \mathbf{v}_k^T - J_h(\mathbf{x}) \Big|_{\mathbf{x}_k^s} \mathbf{P}_k^s \left[J_h(\mathbf{x}) \Big|_{\mathbf{x}_k^s} \right]^T
\end{aligned}
$$

the EM update equations are expressed as

$$
\begin{aligned}
\mathbf{R}^{(j+1)} &= (1 - \alpha_{\mathbf{R}}) \mathbf{\Omega}^{(j+1)} + \alpha_{\mathbf{R}} \mathbf{R}^{(j)} \\
\mathbf{Q}^{(j+1)} &= (1 - \alpha_{\mathbf{Q}}) \mathbf{\Lambda}^{(j+1)} + \alpha_{\mathbf{Q}} \mathbf{Q}^{(j)} \\
\mathbf{P}_0^{(j+1)} &= (1 - \alpha_{\mathbf{P}_0}) \left[\mathbf{P}_0^s + (\mathbf{x}_0^s - \mathbf{x}_0)(\mathbf{x}_0^s - \mathbf{x}_0)^T \right] + \alpha_{\mathbf{P}_0} \mathbf{P}_0^{(j)} \\
\mathbf{x}_0^{(j+1)} &= (1 - \alpha_{\mathbf{x}_0}) \mathbf{x}_0^s + \alpha_{\mathbf{x}_0} \mathbf{x}_0^{(j)}.
\end{aligned}
\quad (20)
$$

Here we have added exponential smoothing to the update equations in the form of a convex combination of the new and previous parameter estimate. In practice, this improves the convergence of all parameters since it allows individual tuning of the convergence speeds. Throughout this work, we have used $\alpha_{\mathbf{R}} = \alpha_{\mathbf{x}_0} = \alpha_{\mathbf{P}_0} = 0.8$ and $\alpha_{\mathbf{Q}} = 0$. $\mathbf{Q}$ is the slowest parameter to converge and this adjustment can help to equalize the parameter convergence speeds.

Due to numerical inaccuracies, it is possible that the covariance matrices loose the positive semi-definiteness property they are required to have. To ensure that $\mathbf{Q}^{(j+1)}$, $\mathbf{R}^{(j+1)}$ and $\mathbf{P}_0^{(j+1)}$ are positive semi-definite, we enforce this property by applying to function

$$
\mathbf{A} \to \mathrm{Re} \left(\sqrt{\mathbf{A}\mathbf{A}^T} \right). \quad (21)
$$

While other projections exists, this particular one has proved to be the most reliable.

A.3 Bisection-like algorithm for EM

Simply iterating over the EM update equations (20) does eventually lead to convergence. However, it often requires a set of well-selected initial parameters. In addition, convergence can be slow as the

optimum is approached, especially if the initial conditions are not well conditioned. To improve these the convergence speed and to reduce dependence on initial parameters, we add an optional bisection inspired algorithm on top of EM. The algorithm is defined in Alg. 1 in pseudo-code.

Algorithm 1: Bisection-like algorithm for EM to improve convergence speed of $\mathbf{Q}$.

Input : $(\boldsymbol{\beta}, \boldsymbol{\gamma}, \beta_{end}, \Xi, J, J_{init}, J_{max})$

1   perform J_{init} EM iterations using Eqs. (20);
2   $j \leftarrow J_{init}$;
3   **while** *any($\beta_i > \beta_{end}$) and $(j < J_{max})$* **do**
4   perform J EM iterations using Eqs. (20);
5      $j \leftarrow j + J$;
6      $\mathbf{B} \leftarrow \mathbb{1}$;
7   **for** $i = 1..dims(\mathbf{Q})$ **do**
8         $\Delta\mathbf{Q} \leftarrow \mathbf{Q}_{i,i}^{(j)} - \mathbf{Q}_{i,i}^{(j-J)}$;
9   **if** $\Delta\mathbf{Q} > 0$ **then**
10   **if** $\gamma_i \geq 0$ **then**
11             $\mathbf{B}_{i,i} \leftarrow \sqrt{\beta_i}$;
12   **else**
13             $\beta_i \leftarrow \beta_i^{\Xi}$;
14             $\mathbf{B}_{i,i} \leftarrow \sqrt{\beta_i}$;
15   **end**
16           $\gamma_i \leftarrow 1$;
17   **else**
18   **if** $\gamma_i \geq 0$ **then**
19             $\mathbf{B}_{i,i} \leftarrow \sqrt{1/\beta_i}$;
20   **else**
21             $\beta_i \leftarrow \beta_i^{\Xi}$;
22             $\mathbf{B}_{i,i} \leftarrow \sqrt{1/\beta_i}$;
23   **end**
24           $\gamma_i \leftarrow -1$;
25   **end**
26   **end**
27      $\mathbf{Q}^{(j)} \leftarrow \mathbf{B}\mathbf{Q}^{(j)}\mathbf{B}^T$
28 **end**

This algorithm applies J EM iterations and amplifies the convergence trend of the diagonal elements of $\mathbf{Q}$. If the matrix element $\mathbf{Q}_{i,i}$ has converged to higher values within J steps, it is multiplied with a factor $\sqrt{\beta_i}$ larger than 1. For a downwards trend, it is multiplied with $1/\sqrt{\beta_i} < 1$. β_i is initially large to cause large jumps in $\mathbf{Q}_{i,i}$. Initial large jumps ensure that a large portion of the parameter space is explored during the optimization. While $\mathbf{Q}_{i,i}$ is closing in on the true value through continuous application of the EM algorithm, β_i is shrunk closer to unity using $\beta_i \leftarrow \beta_i^{\Xi}$ with $0 < \Xi < 1$ whenever the sign of the previous trend amplification had the opposite sign γ_i. This causes smaller and smaller jumps to home in on the true value for $\mathbf{Q}_{i,i}$. This ensures that β_i monotonically shrinks. The process terminates when $\beta_i < \beta_{end}$ with β_{end} close to 1.

A.4 Initial values for EM parameter optimization

The modifications and tuning of the EM algorithm presented above lead to apparent independence of the convergence of the initial parameters. Hence we are able to use the same hyper- and initial parameters for all applications of the EKS in this work. The only exception are the initial values for amplitude (A_0^{est}) and frequency (δ_f^{est}) of the signal, which appear in $\mathbf{x}_0$. They are estimated using the same approach as for the nonlinear least squares method. All other initial parameters are:

Table 1: Hyper- and initial parameters for EKS parameter optimization

Parameter	Value
$Q^{(j=0)}$	$\mathrm{diag}(10^{-5}, 10^{-5}, 10^{-5}, 10^{-5}, 10^{-4})$
$R^{(j=0)}$	$\mathrm{diag}(10^2)$
$\mathbf{x}_0^{(j=0)}$	$\left[\sqrt{A_0^{est}/(4\pi)}, 1.0, \delta_f^{est}, 0, 0\right]^T$
$P_0^{(j=0)}$	$\mathrm{diag}(10^{-1}, 10^{-1}, 10^{-1}, 10^{-2}, 10^{-2})$
Ξ	0.75
γ_i	0
β_i	100
β_{end}	$[\beta_i]^{-2^6}$
J	20
J_{init}	200
J_{end}	$100 \cdot J$

B Further simulation study results

Here we present further results from the simulation study shown in the main text, but were left out due to space constraints.

B.1 Simulation Study Parameters

B.2 Amplitude RMSE comparison

Just as the ratio of the frequency estimation RMSE of both methods are compared in the main text, we can also compare the amplitude estimates of the methods. Analogously to the definition in the main text, we define the ratio of amplitude RMSE as

$$\rho_{RMSE,A} = \log_2 \frac{\mathrm{RMSE}_{A,EKS}}{\mathrm{RMSE}_{A,SCF}}. \tag{22}$$

Table 2: Fixed parameters in the simulation study.

Parameter	Value
N	$5.4\,\mathrm{Msamples}$
f_s	$500\,\mathrm{Hz}$
f_c	$84.06\,\mathrm{Hz}$
T_2^*	$3142\,\mathrm{second}$
σ_η	$10\,\mathrm{pT}$
φ_0	0
$\mathrm{N}_{repetitions}$	50
$\mathrm{N}_{bl,EKS}$	$4.5\,\mathrm{s}\cdot f_s$

This ratio, alongside with slices through the parameter space, are displayed in Fig. 5. The ratio is below 0 everywhere, indicating that the amplitude time series estimation of the EKS is significantly better that the SCF estimation. This is due to the selection of the block length of the SCF: The block length that produces the lowest frequency RMSE is selected, which is often very long compared to the time scales of the exponential decay of the amplitude. Hence, the time resolution is too low to resolve the exponential decay, leading to large RMSE.

While the SCF block length could be optimized to produce a significantly lower amplitude RMSE, this would result in a larger frequency RMSE. This tradeoff can be avoided using the EKS, as the same model parameters produce both good frequency and amplitude estimates.

B.3 Validity of EKS uncertainty intervals

As the EKS is a Bayesian method, it inherently produces uncertainty estimates for its frequency estimates. However, since the EKS can only solve the nonlinear state space model approximately via linearization, we checked whether the uncertainty intervals accurately reflect the frequency estimation error.
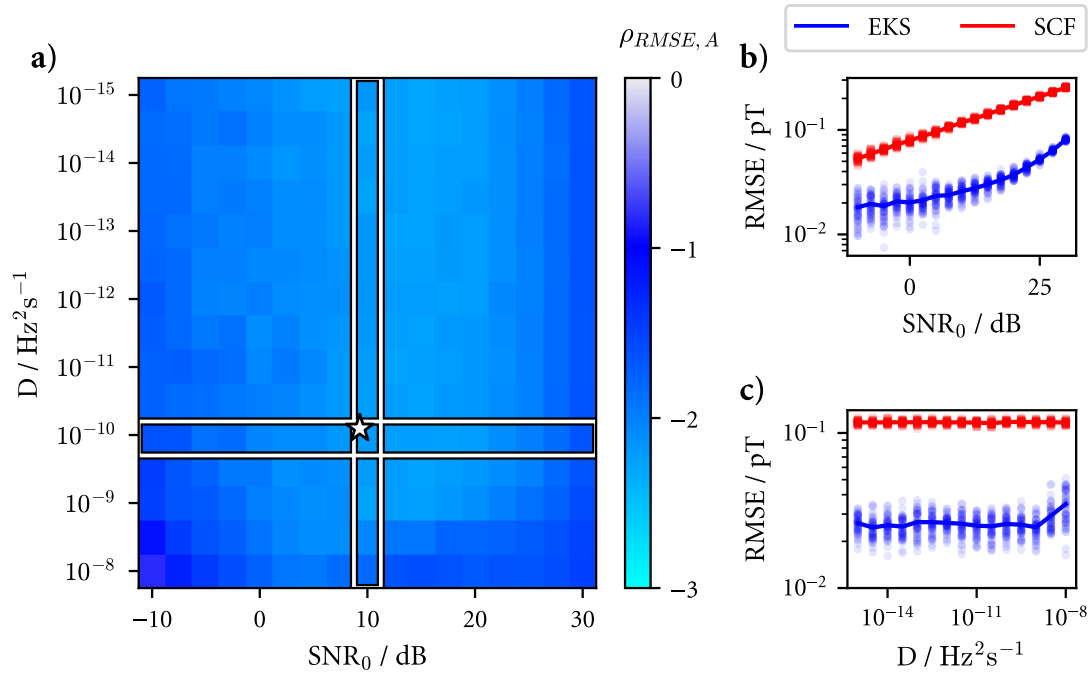
Figure 5: Quantitative comparison of the SinCos Fit against the EKS in a simulation study. **a)** $\log_2$ of the mean ratio of the amplitude estimate errors. Blue indicates a better performance of the EKS over the SCF, red the opposite. The star marker indicates a rough estimate of our experimental conditions. **b)** Horizontal slice through a). The RMSEs of individual simulation runs are displayed as markers, the solid lines indicate the mean. **c)** displays a vertical slice.
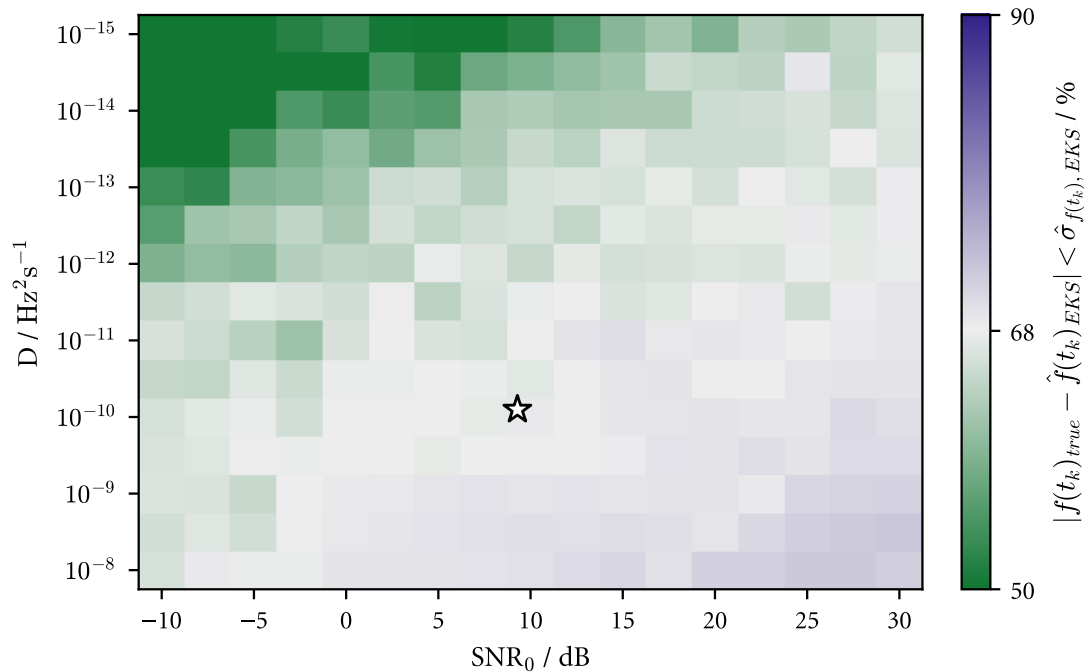

Figure 6: Percentage of true frequency samples $f(t_k)_{true}$ within the 1σ bounds of the EKS frequency estimate.

The EKS assumes Gaussian probability distributions of its estimates. For each signal sample, the EKS

produces a probability density function (PDF) that should represent the probability of the value of the true frequency. Therefore, we can simply count the number of true frequency samples that lie within the 1σ interval of the estimated Gaussian PDF. If the estimated PDF represents the estimation error well, about 68% of true frequency samples of the full time series should fall within the 1σ interval.

This count is visualized in Fig. 6. In a large region of the parameter space, the count is very close to 68% of samples within the 1σ interval, indicating that the uncertainty estimates are valid. This regime corresponds to the region in which the EKS outperforms the SCF in frequency estimation. For very low frequency fluctuation and low SNR, the count is closer to 50%, which indicates that the uncertainty is underestimated. This coincides with the parameter regime where the SCF performs better than the EKS.

Not shown here is that the uncertainty estimate depends on the selected block length of the EKS. With small block lengths, the uncertainty tends to be underestimated while with large block lengths, it is overestimated. Therefore, the selection of the correct block length is critical if the uncertainty estimates are of importance for the application. Currently, no selection criterion other than comparison to simulations exists.

B.4 Optimal SCF block length

As mentioned in the main text, the optimal block length of the SCF is selected based on the minimal frequency estimation RMSE. In Fig. 7, these optimal block lengths are shown.
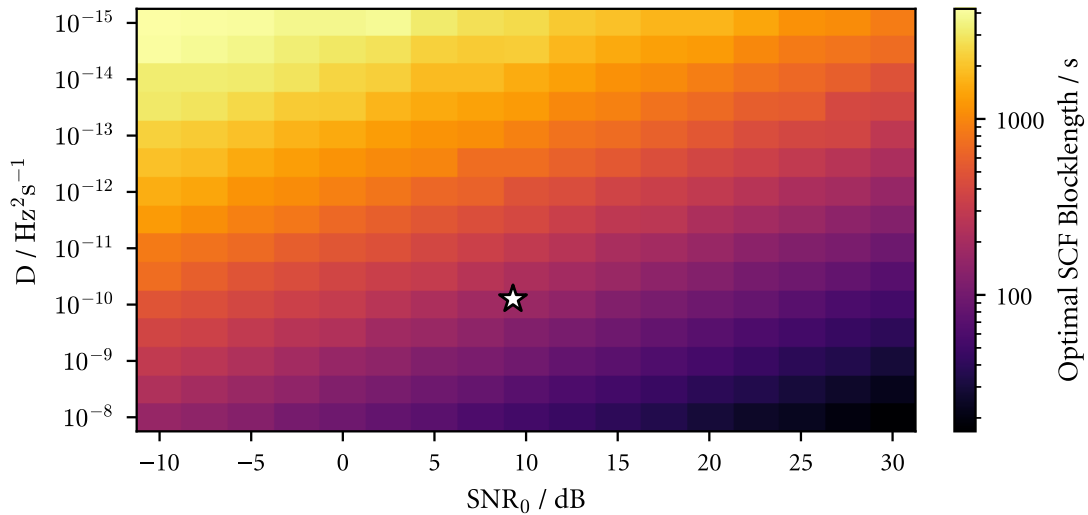

Figure 7: Optimal block length for SCF in the minimal frequency RMSE sense.

B.5 Cramér–Rao lower bound (CRLB) for zero drift

In the case of zero drift ($D = 0$) we can calculate the CRLB [14] and compare it to the performance of our estimators (see Fig. 8).

Figure 8: At zero drift the SCF performs better then the EKS, while both methods perform close to the theoretical variance limit of the CRLB.

References

[1] A. Grosz, M. J. Haji-Sheikh, S. C. Mukhopadhyay (Eds.), High Sensitivity Magnetometers, Vol. 19 of Smart Sensors, Measurement and Instrumentation, Springer, Cham., 2017. `doi:10.1007/978-3-319-34070-8`.

[2] E. Tiesinga, P. J. Mohr, D. B. Newell, B. N. Taylor, The 2022 codata recommended values of the fundamental physical constants (2022).
URL `https://physics.nist.gov/constants`

[3] X. Fei, V. Hughes, R. Prigl, Precision measurement of the magnetic field in terms of the free-proton NMR frequency, Nuclear Instruments and Methods in Physics Research Section A 394 (1997) 349–356. `doi:10.1016/s0168-9002(97)84161-7`.

[4] K. Weyand, Maintenance and dissemination of the magnetic field unit at PTB, IEEE Trans. Instrum. Meas. 50 (2001) 470–473. `doi:10.1109/19.918168`.

[5] A. Schneider, B. Sikora, S. Dickopf, M. Müller, N. S. Oreshkina, A. Rischka, I. A. Valuev, S. Ulmer, J. Walz, Z. Harman, C. H. Keitel, A. Mooser, K. Blaum, Direct measurement of the ^{3}He$^+$ magnetic moments, Nature 606 (2022) 878–883. `doi:10.1038/s41586-022-04761-7`.

[6] M. Farooq, T. Chupp, J. Grange, A. Tewsley-Booth, D. Flay, D. Kawall, N. Sachdeva, P. Winter, Absolute magnetometry with ^{3}He, Phys. Rev. Lett. 124 (2020) 223001. `doi:10.1103/PhysRevLett.124.223001`.

[7] K. Stamataki, V. Papadakis, M. A. Everest, S. Tzortzakis, B. Loppinet, T. Peter Rakitzis, Monitoring adsorption and sedimentation using evanescent-wave cavity ringdown ellipsometry, Applied Optics 52 (2013) 1086. `doi:10.1364/ao.52.001086`.

[8] A. Moriat, Sytem and method for estimating tones in an input signal. `arXiv:https://forums.ni.com/t5/Signal-Conditioning/frequency-of-sine-waveform-VI/m-p/2638895#M7428`.
URL `https://patents.justia.com/patent/6965068`

[9] D. Yurin, M. Makukov, D. Kuvatova, A. Gluchshenko, C. Omarov, Applying zero-crossing method for frequency map analysis of dynamical systems, Mon. Not. R. Astron Soc. 531 (2024) 1483–1495. `doi:10.1093/mnras/stae1216`.

[10] Y.-Y. Lin, P. Hodgkinson, M. Ernst, A. Pines, A novel detection–estimation scheme for noisy NMR signals: Applications to delayed acquisition data, J. Magn. Reson. 128 (1997) 30–41. `doi:10.1006/jmre.1997.1215`.

[11] G. Golub, V. Pereyra, Separable nonlinear least squares: the variable projection method and its applications, Inverse Problems 19 (2) (2003) R1. `doi:10.1088/0266-5611/19/2/201`.

[12] J. C. Visschers, D. Budker, L. Bougas, Rapid parameter estimation of discrete decaying signals using autoencoder networks, Mach. Learn.: Sci. Technol. 2 (2021) 045024. `doi:10.1088/2632-2153/ac1eea`.

[13] H. R. Almayyali, Z. M. Hussain, Deep learning versus spectral techniques for frequency estimation of single tones: Reduced complexity for software-defined radio and iot sensor communications, Sensors 21 (2021) 2729. `doi:10.3390/s21082729`.

[14] C. Gemmel, W. Heil, S. Karpuk, K. Lenz, C. Ludwig, Y. Sobolev, K. Tullney, M. Burghoff, W. Kilian, S. Knappe-Grüneberg, W. Müller, A. Schnabel, F. Seifert, L. Trahms, S. Baeßler, Ultra-sensitive magnetometry based on free precession of nuclear spins, Eur. Phys. J. D 57 (2010) 303–320. `doi:10.1140/epjd/e2010-00044-5`.

[15] K. Tullney, F. Allmendinger, M. Burghoff, W. Heil, S. Karpuk, W. Kilian, S. Knappe-Grüneberg, W. Müller, U. Schmidt, A. Schnabel, F. Seifert, Y. Sobolev, L. Trahms, Constraints on spin-dependent short-range interaction between nucleons, Physical Review Letters 111 (10) (2013) 100801. `doi:10.1103/PhysRevLett.111.100801`.

[16] N. Sachdeva, I. Fan, E. Babcock, M. Burghoff, T. Chupp, S. Degenkolb, P. Fierlinger, S. Haude, E. Kraegeloh, W. Kilian, S. Knappe-Grüneberg, F. Kuchler, T. Liu, M. Marino, J. Meinel, K. Rolfs, Z. Salhi, A. Schnabel, J. Singh, S. Stuiber, W. Terrano, L. Trahms, J. Voigt, New measurement of the permanent electric dipole moment of ^{129}Xe using ^{3}He comagnetometry and SQUID detection, Phys. Rev. Lett. 123 (2019) 143003. `doi:10.1103/PhysRevLett.123.143003`.

[17] H. E. Rauch, F. Tung, C. T. Striebel, Maximum likelihood estimates of linear dynamic systems, AIAA Journal 3 (8) (1965) 1445–1450. `doi:10.2514/3.3166`.

[18] M. Tarvainen, J. Hiltunen, P. Ranta-aho, P. Karjalainen, Estimation of nonstationary eeg with kalman smoother approach: an application to event-related synchronization (ERS), IEEE Transactions on Biomedical Engineering 51 (3) (2004) 516–524. `doi:10.1109/TBME.2003.821029`.

[19] R. Jiménez-Martínez, J. Kołodyński, C. Troullinou, V. G. Lucivero, J. Kong, M. W. Mitchell, Signal tracking beyond the time resolution of an atomic sensor by kalman filtering, Physical Review Letters 120 (4) (2018) 040503. `doi:10.1103/PhysRevLett.120.040503`.

[20] B. La Scala, R. Bitmead, B. Quinn, An extended Kalman filter frequency tracker for high-noise environments, IEEE Transactions on Signal Processing 44 (2) (1996) 431–434. `doi:10.1109/78.485940`.
URL `http://ieeexplore.ieee.org/document/485940/`

[21] S. M. Kay, Fundamentals of Statistical Signal Processing: Estimation Theory, Prentice Hall PTR, 1993.

[22] H.-C. Koch, G. Bison, Z. D. Grujić, W. Heil, M. Kasprzak, P. Knowles, A. Kraft, A. Pazgalev, A. Schnabel, J. Voigt, A. Weis, Design and performance of an absolute ^{3}He/Cs magnetometer, Eur. Phys. J. D 69 (2015) 202. `doi:10.1140/epjd/e2015-60018-7`.

[23] T. R. Gentile, P. J. Nacher, B. Saam, T. G. Walker, Optically polarized ^{3}He, RMP 89 (2017) 045004. `doi:10.1103/RevModPhys.89.045004`.

[24] H.-C. Koch, G. Bison, Z. D. Grujić, W. Heil, M. Kasprzak, P. Knowles, A. Kraft, A. Pazgalev, A. Schnabel, J. Voigt, A. Weis, Investigation of the intrinsic sensitivity of a ^{3}He/Cs magnetometer 69 (2015) 262. `doi:10.1140/epjd/e2015-60509-5`.

[25] E. Klinger, T. Liu, M. Padniuk, M. Engler, T. Kornack, S. Pustelny, D. F. J. Kimball, D. Budker, A. Wickenbrock, Optimization of nuclear polarization in an alkali-noble gas comagnetometer 19 (2023) 044092. `doi:10.1103/physrevapplied.19.044092`.

[26] S. Särkkä, Bayesian Filtering and Smoothing, 1st Edition, Cambridge University Press, 2013. `doi:10.1017/CBO9781139344203`.
URL `https://www.cambridge.org/core/product/identifier/9781139344203/type/book`

[27] D. Dreano, P. Tandeo, M. Pulido, B. Ait-El-Fquih, T. Chonavel, I. Hoteit, Estimating model-error covariances in nonlinear state-space models using Kalman smoothing and the expectation–maximization algorithm, Quarterly Journal of the Royal Meteorological Society 143 (705) (2017) 1877–1885. `doi:10.1002/qj.3048`.