Ridley and Rodgers *EPJ Quantum Technology* (2024) 11:79
https://doi.org/10.1140/epjqt/s40507-024-00292-4

EPJ Quantum Technology
a SpringerOpen Journal

RESEARCH Open Access

An investigation of errors in ellipse-fitting for cold-atom interferometers

Kevin Ridley[1*] and Anthony Rodgers[1]

*Correspondence:
kdridley@physics.org
[1]University of Birmingham, Birmingham, UK

Abstract

Ellipse fitting is a technique which is used to extract differential phase in cold-atom interferometers, particularly in situations where common-mode noise needs to be suppressed. We use numerical simulation to investigate errors in the ellipse fitting process; specifically, errors due to the presence of additive noise, linear drift in ellipse offset and amplitude, as well as an error that can arise from fringe normalisation. Errors are found to manifest in two ways: bias in the ellipse phase measurement and incomplete suppression of common mode phase noise. We quantify these errors for three different ellipse fitting algorithms and discuss the applicability of these results to future cold atom sensors.

Keywords: Ellipse fitting; Gravity gradiometry; Quantum sensors; Cold atoms

1 Introduction

Ellipse fitting is used as a method for extracting differential phase in cold-atom interferometers [1, 2]. An example is a gravity gradiometer, where simultaneous measurements are carried out with two spatially-separated cold-atom interferometers [3, 4]. Ellipse fitting has also been used in atom interferometer experiments that measure the gravitational constant [5, 6], the fine-structure constant [7] and test the weak equivalence principle [8]. It has also been proposed for use in future gravitational-wave detectors [9].

Ellipse fitting can be particularly useful in situations where two interferometers have common-mode phase noise which varies over more than a single 2π range, that is, more than one interference fringe. In the gravity gradiometer, the common-mode phase noise comes from common accelerations, which may arise from vibration of the gradiometer or acceleration of a moving platform on which the instrument is mounted. Referring to the common mode phase as θ, the suitably-normalised interferometer signals can be written as

$$x = a\sin\theta + b \tag{1}$$

$$y = c\sin(\theta + \phi) + d \tag{2}$$

Here ϕ is the differential phase, which contains a component proportional to the gravity gradient being measured.

The ellipse fitting method involves taking a number of separate measurements of x and y which have different phases θ, but the same phase difference ϕ. The different values of θ may arise naturally, due to vibration, for example, or they may be imposed deliberately. It can be shown that when the values of y are plotted against those of x the locus of points is an ellipse [10]. By fitting an ellipse to the data, values for the parameters in (1) and (2) can be obtained, including the desired phase ϕ. Note that, because it appears within a sine function, the phase ϕ can only be recovered over a range of π radians, not 2π. We will use the convention that ϕ lies between 0 and π radians. Imperfections in the sensor may result in data which is not precisely elliptical in form. The ellipse parameters a, b, c, d may change during acquisition of a single ellipse, for example. Noise arising during the detection of optical signals will manifest as additive, random noise, which can be mathematically equivalent to random variations in the offsets b and d. All of these effects may lead to errors in the extraction of the phase by the ellipse-fitting process. If the data do not lie on an ellipse, then the ellipse fitting algorithm will produce a 'best fit' ellipse, which will not necessarily give the true phase ϕ. In the present paper, we investigate a number of ways in which the data can be non-elliptical and quantify the resulting phase errors. We show that the ellipse-fitting process exhibits bias: the average measured phase can differ from the true phase. We also show that ellipse errors lead to imperfect suppression of common-mode phase noise and therefore 'break-through' of this noise into the measured phase.

As future cold-atom interferometers with higher sensitivities are developed, the influence of small ellipse fitting errors on the phase measurement will become more significant. The results of this paper can be used to quantify how stable the various ellipse parameters need to be in order to achieve a given phase accuracy. We concentrate on the gravity gradiometer application but note that the results may also be applicable to other systems, such as dual-species atom-interferometers [8, 11].

A gravity gradiometer consists of two cold-atom interferometers separated by a distance L. A difference of Δg in the gravitational accelerations experienced by the atom clouds in the two interferometers results in a phase difference of [12]

$$\phi_g = -2\Delta g k T^2 \tag{3}$$

Where $k = 2\pi/\lambda$, λ being the wavelength of the Raman laser, and $2T$ is the drop time. Here, the subscript g indicates that this is the gravitational component of the phase ϕ, as distinct from other components which may be present as well. Consider, as an example, an instrument with T=50 ms (corresponding to a drop distance slightly under 5 cm for an atom cloud starting at rest) and L=1 m. A Rubidium cold-atom interferometer would use a Raman wavelength of 780 nm. The gravity difference in equation (3) is equal to the gravity gradient times L, provided the gradient does not change significantly over the distance L. Thus, the Earth's mean gravitational gradient of $3.08\times10^{-6}\mathrm{s}^{-2}$ would produce a phase shift of 0.124 rad. In geophysical applications one might be interested in measuring gradients which are less than 1/100th of the Earth's gradient [13], which would give phase changes of order milliradians or less. Therefore, a requirement for milliradian (or better) precision in ellipse fitting is a realistic performance goal for this type of system.

2 Methods

Various methods for fitting ellipses can be found in the literature [8, 14–17]. One can distinguish between geometric fitting methods, which minimise the distance between the data points and the fitted ellipse (in the x-y plane), and algebraic methods, which minimise a certain function of x and y. Geometric methods are iterative and relatively slow compared to the non-iterative algebraic methods. However, geometric methods have the advantage of minimising the error between the fit and the data in a more satisfactory manner; for detailed discussions of the issues involved the reader is referred to references [1, 14–16].

Geometric methods are all aimed at obtaining the same fit, albeit via different routes (although we note that the minimum distance can be defined in different ways: an L1 norm or an L2 norm, for example). Different algebraic methods, however, may give different fits. Algebraic methods typically minimise a function of x and y which arises from the equation of a conic section. However, there is a mismatch between the conic section equation, which has six parameters, and equations (1) and (2), which have five parameters. This mismatch can be resolved by reducing the number of parameters in the conic section equation. This can be achieved by either making one of the parameters a constant, or by introducing some relationship between the parameters so that the effective number of parameters is reduced by one. Clearly, there are many possible ways of doing this parameter reduction. In the present paper we use two different methods, each of which has some theoretical justification. The first method we use, referred to as method A, is described by Gander et al [10]; they call it a Euclidian-invariant algorithm because it is dependent on the shape of the ellipse only and does not change if the ellipse is translated or rotated (there are two such methods in Gander et al, the one we use has the constraint $\lambda_1+\lambda_2=1$, in their notation). The second method, referred to as method B, is that introduced by Fitzgibbon et al [14]. This uses a Lagrange multiplier to force the fit to be elliptical, thus excluding solutions corresponding to other conic sections. This was the method used by Foster et al for their gravity gradiometer [1].

The geometric fit method, referred to as method G, is that described in reference 10. It starts with an algebraic fit (in our case using method A) and then uses an iterative Gauss-Newton method to find a fit which minimises the geometric distance. The computer code for this fitting method is based on that produced by Richard Brown [18]. In this method, two variables are introduced in order to determine when to stop iterating and return a result. First, there is a goodness-of-fit measure which the algorithm aims to achieve. For each iteration it compares the norm of the current ellipse parameters with the norm of the increments to those parameters; the measure used is the ratio of the latter to the former: if this is less than a specified tolerance value the iteration is stopped. Unless stated otherwise, a default value of 1×10^{-5} was used for this tolerance. The second variable is the maximum number of iterations used before stopping. This prevents the algorithm getting stuck in an endless loop. A default value of 200 iterations was used. If the maximum number of loops is reached, the algorithm returns the initial algebraic fit, along with a message to say that the geometric fit has failed.

Real data will never be precisely elliptical, if only because of inevitable noise on the measurements. We will investigate a number of different ways in which the data may depart from an ellipse. This is not intended to be exhaustive but is based on experience in operating cold-atom interferometers within our research group [4]. One feature that is common

to many of these effects is a random component to the data, which necessitates a statistical analysis. This is obviously the case when considering random noise in the optical detection process; however, it can also arise when taking into account the effects of common-mode vibration, i.e. the θ term in equations (1) and (2). If the system is subjected to vibration, at one or more frequencies, and that vibration moves the interference fringes by much more than a single period, then it is reasonable to approximate the effect by assuming the phase θ is a random variable which is uniformly distributed between 0 and 2π, and this is the assumption we use for our simulations. Each measurement has a random phase which is uncorrelated with previous phases; this is based on the assumption that the time period between measurements is relatively long compared to vibration frequencies and that the timing of the measurement is not synchronised with the vibration frequency. We use this random-phase model in all of the simulations.

We introduce the statistical element by using random number generation to produce different realisations of the underlying random variables. Every individual estimate of the phase ϕ, produced by an ellipse fit, represents a sample from a population of ellipse phases. By repeating many times, to generate many samples, we investigate the statistical properties of the phase. Specifically, we calculate the mean and standard deviation of the sample phases. For sufficiently large samples, these give accurate estimates of the population mean and the population standard deviation.

3 Results

3.1 Additive Gaussian noise

First, we consider the case where the measured quantities x and y are subject to additive noise with Gaussian statistics. An atom interferometer system will have a number of different noise sources. The most fundamental of these is usually considered to be atom shot noise (also referred to as quantum projection noise), which will be governed by Poisson statistics. Additional noise sources are photon shot noise, both in the atom fluorescence signal and stray background light, and photodiode detector noise. Photodiode detector noise is thermal in origin and governed by Gaussian statistics. It is well known that Poisson noise tends towards Gaussian statistics when the mean number (of photons or atoms) is high. The model in this paper was developed for the system described in reference 4, and atom shot noise is not the dominant noise source in this system. The dominant noise sources are detector noise, which is Gaussian, and photon shot noise on background light, which can be treated as Gaussian because the photon number is always large. If atom shot noise was the limiting noise source, and the interferometer fringe contrast was high, the assumption of Gaussian statistics would be invalid and there may be some differences compared to the results presented in Sect. 3.1. Further work would be needed to quantify such differences. Thus, we use additive Gaussian noise with zero-mean and a given standard deviation (note that any non-zero mean can be simply subsumed into the offsets b and d). For simplicity, we use the same standard deviation, σ_s, for the noise in both the x and the y measurements. Each ellipse is formed from N measurements of x and y with independent additive noise and random phases θ, the latter being independent and sampled from a uniform distribution between 0 and 2π. The ellipse parameters used in this section are, a=1, b=2, c=1, d=2. Note that the results are, in fact, independent of the values of the offsets b and d and depend only on the ratios of a and c to σ_s. An example of simulated ellipse data with N=100 is given in Fig. 1, this is for ϕ=1.0 rad and σ_s =0.05.

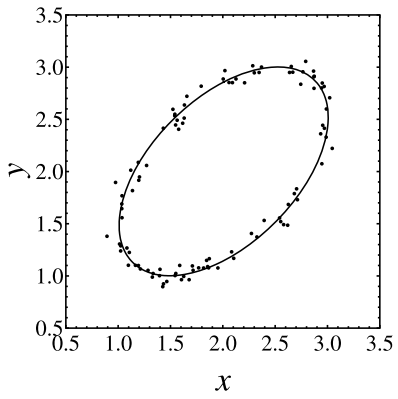

Figure 1 Simulated ellipse data with additive noise. The solid line is the ellipse curve without noise

In Fig. 1, the points are the $\{x, y\}$ values and the solid line shows the ellipse without noise. Note that, here, the degree of scatter in the points is independent of the position around the ellipse. Referring back to the earlier discussion around atom shot noise, if shot noise was the dominant noise source one would see more scatter around the top right hand side of the ellipse, because atom number is higher in this region. The extent to which the scatter varies between the lower left and upper right would depend on the interferometer contrast, i.e. how close to zero the atom number becomes in the lower left region. As mentioned previously, one might expect some difference in the results in such an atom-shot-noise limited regime, but this is beyond the scope of the present paper.

The value of ϕ determines how open the ellipse is: $\phi = \pi/2$ gives a circle, as ϕ approaches zero the ellipse collapses on to the line $y=x$, and as it approaches π the ellipse collapses to the line $y=-x$. In practice, one would aim to avoid values near 0 and π because ellipse fitting errors become greater as the locus of points approaches a straight line. For many of the results in this paper we use the value of ϕ=1.0 rad, this gives an open ellipse but one which is still clearly elliptical and not the special case of a circle. Results with ϕ =1.0 rad were used to investigate the influence of the number of points N to which the ellipse is fitted. Clearly, one would expect the accuracy of the estimate of ϕ, derived from the ellipse fit, to improve as N gets larger. Equally, one would expect poor ellipse fits for very small values of N. There are 5 free parameters in equations (1) and (2), so a minimum of N=5 points are required to determine them unambiguously. The estimate of ϕ can be written as the true value, plus some constant bias, plus a random noise component of zero mean,

$$\hat{\phi} = \phi + \phi_b + \phi_n. \tag{4}$$

The accuracy of the phase measurement depends on both the bias and the random noise component. The random noise can be characterised by its standard deviation. Since the bias is constant, the standard deviation of the phase estimates is equal to the standard deviation of the noise component ϕ_{n}. One would expect that for large N the standard deviation of the noise will reduce by the square root of N, as is usual for stationary random noise processes, that is, for large N

$$\sigma_{\hat{\phi}} \propto \mathrm{N}^{-\frac{1}{2}}. \tag{5}$$

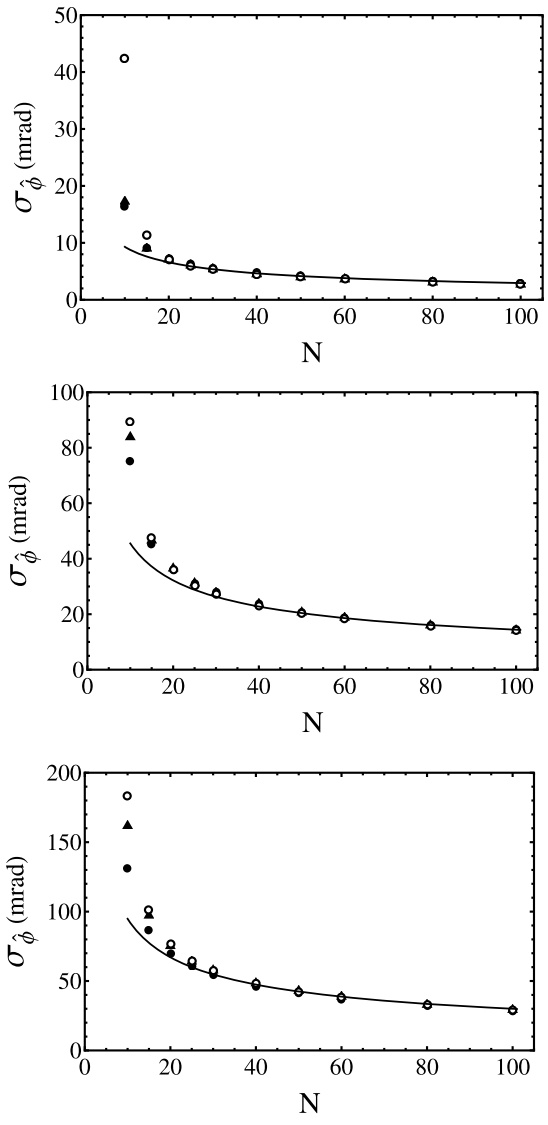

Figure 2 Ellipse fit as a function of the number of measurements N. Upper plot for σ_s =0.01; middle, σ_s =0.05, and lower, σ_s =0.1. Triangles: method A; solid circles: method B; open circles: method G

In Fig. 2, the standard deviation of the phase estimates is plotted as a function of N for values of N between 10 and 100 and three different amounts of additive noise.

Each point on the graphs is the standard deviation obtained from 50,000 ellipse fits. Note that the values plotted are in milliradians. The square-root dependency of equation (5) is shown by the solid lines, where the constant of proportionality has been determined by fitting to the four points with the largest N values. It can be seen that, for all three additive noise levels, there is a notable departure from the square-root dependence when N is less than around 30, and the departure becomes greater for smaller N. This implies that ellipses with fewer than 30 data points may give noticeably sub-optimal phase estimates. This is a relevant consideration when dealing with large numbers of x and y values. It may be that it is preferable to sub-divide the data into a series of short ellipses rather than fitting a single ellipse to all the data, for reasons to do with short-term stability, for

example. These results show that there will be no substantial reduction in overall errors if the individual ellipses have N greater than around 30 (the standard deviation when all the individual phase estimates are averaged will be the same as if a single ellipse were used); when, however, fewer points are used there will be excess phase noise and the standard deviation after averaging all the phase estimates will be worse than that obtained if a single ellipse was fitted to all the data. Note that the value of N=30 is chosen, somewhat arbitrarily, as the N value below which the departure from ideal behaviour becomes apparent: it is not an absolute cut-off below which ellipse fitting cannot be used. There are some differences in the performance of the different ellipse fitting algorithms when N is small but at the higher N values all three algorithms have very similar noise performance.

The errors associated with each point in Fig. 2 can be estimated as follows: the 95% confidence interval for the standard deviation is $2\sigma/\sqrt{2n}$, where n=50,000 is the number of phase estimates contributing to each point. Thus, when σ =100 mrad, for example, the 95% confidence interval is +/-0.63 mrad. For the numbers of ellipses used in these simulations, these values are generally too small to show as error bars on the graphs. Note, however, that this error analysis assumes that the phase estimates are Gaussian variables. For most of the data this is a good assumption, however, for N=10 it was found that the geometric fit, method G, departed from Gaussian behaviour owing to the presence of occasional 'outliers' with much larger errors. These occurred when the goodness-of-fit metric was small, but a good fit was not, in fact, achieved and was particularly significant for the case of σ_s =0.01, which is why the standard deviation for method G is so much greater than for the other methods in that case. Generally, method G gave greater fluctuations compared to the other methods when N was small.

The bias was also calculated for this data, although plots are not included here for reasons of space. It was found that, in all cases, the bias changes significantly with N for small N (i.e. N<30), but much less so when N is increased to values increased larger values, beyond 30.

The next set of results investigate changing the value of the phase ϕ. These are shown in Fig. 3 and used additive noise with a standard deviation of 0.05, and N=100 data points per ellipse.

The phase ϕ was varied from 0.17 rad to 2.97 rad in steps of 0.1 rad. The results for methods A and B have a continuous line interpolated through these points. Results for method G do not include points at ϕ =0.17 rad and ϕ =2.97 rad because the geometric fit was not always able to locate a global minimum when the ellipses were very narrow. The quantity plotted on the vertical axis of the upper graph in Fig. 3 is the phase error (bias), which is defined as the mean of the difference between the measured phase and the actual value of ϕ, that is, the quantity ϕ_b in (4). For each phase value, 100,000 samples were used to calculate the mean and standard deviation. The standard deviations are plotted in the lower graph in Fig. 3. Note that the vertical axis does not start from zero, this was in order to show the variations more clearly. In the results of Fig. 2, the different points, corresponding to the different ellipse fitting methods, cannot be distinguished when N=100, but in Fig. 3 it can be seen that there are some small differences. For values of ϕ near $\pi/2$, the standard deviation of the phases is slightly less with method B than with methods A and G. As the ellipse phase moves away from $\pi/2$, the values for method B increase slightly but those for method G decrease, becoming significantly lower before rising again. This means that a small improvement in performance, of order a few percent, could be obtained

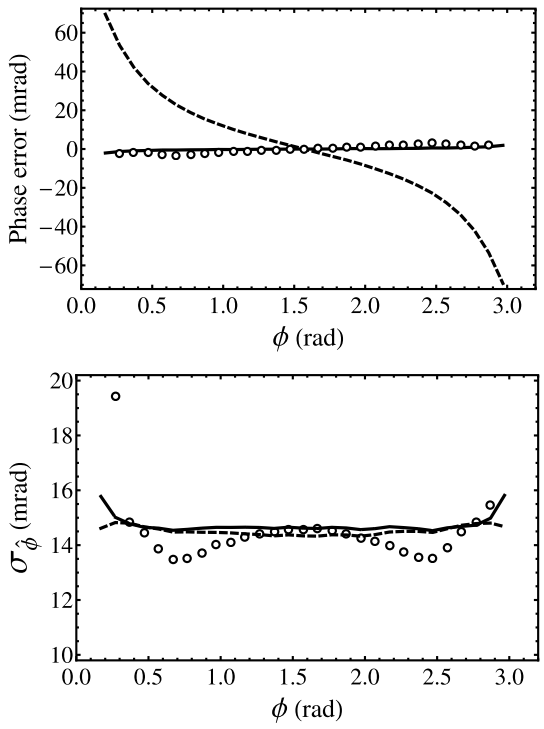

Figure 3 Ellipse fit with additive noise. Upper plot: phase error as a function of true phase (bias). Lower plot: standard deviations of the phase estimates. Dotted line method A, dashed line method B, circles method G

by using method G when the ellipse phase is near 0.7 rad. With method G, the final point at ϕ =0.27 rad is anomalously large owing to convergence issues, i.e. the geometric fit not always locating the global minimum. We note, also, that as the two extremes of 0 and π are approached, the standard deviation increases with method A but reduces slightly with method B.

The standard deviation can be used to estimate the accuracy of these results. The 95% confidence interval of the mean is $2\sigma/\sqrt{n}$, n being the number of samples. Thus, the errors in the upper graph in Fig. 3 are of order 0.1 mrad. As mentioned previously, the corresponding confidence intervals for the standard deviation (lower graph) are $2\sigma/\sqrt{2n}$, which gives errors of order 0.07 mrad.

All the ellipse fitting methods are found to be biased in the presence of additive noise. That is, the mean of the measured phase is not equal to the true phase. The only exception is when the phase is exactly equal to $\pi/2$, when the ellipse becomes a circle. It is clear from Fig. 3 that method B exhibits much more bias than methods A and G. Note that the A and G bias goes in the opposite direction to that of B, the slopes of the central linear region being 0.5 mrad/rad, -20 mrad/rad and 3 mrad/rad for methods A, B and G respectively. In order to investigate the bias as a function of the amount of noise, further results were calculated with a constant phase of ϕ=1.0 rad. These are given in Fig. 4, which used 50,000 ellipses per noise value. Method B is found to have the greatest bias and method A the lowest (upper graph).

All the methods show similar standard deviations (Fig. 4, lower graph), with some minor differences becoming apparent at higher noise levels. The initial increase in phase standard deviation is in direct proportion to the noise standard deviation, with some slight

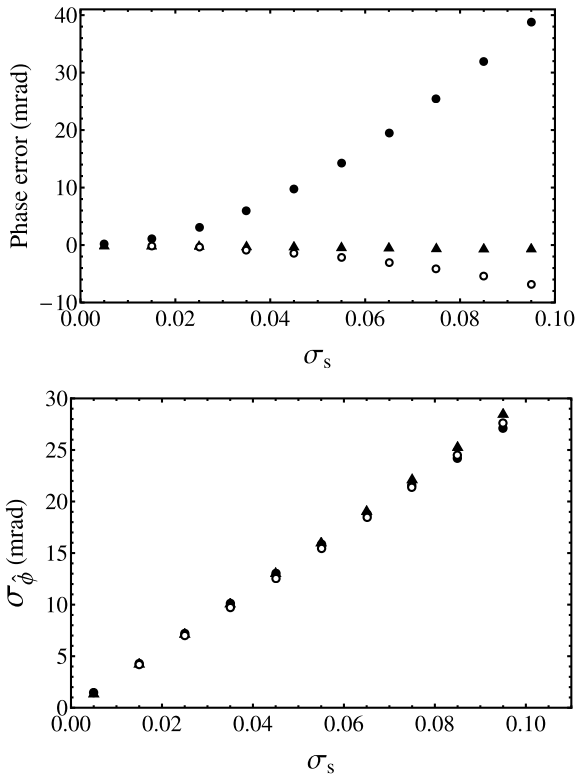

Figure 4 Ellipse Phase errors (upper) and standard deviation (lower) as a function of noise standard deviation for ϕ=1.0 rad. Triangles: method A; solid circles: method B; open circles: method G

Table 1 Results for least squares fits to data in Figs. 4, 9 and 15. Fits are either linear, of the form y=ax, for the phase standard deviation, or quadratic, of the form $y=bx^2$, for the phase bias. Here, the parameter x is indicated in the first column. In the second column y is the phase bias and in the third it is the phase standard deviation, both in radians. The values of the fit coefficients, b and a are given in these last two columns

Investigation and Method	Phase bias (b)	Phase standard deviation (a)
Additive noise (Fig. 4); x=σ_s		
Method A	-70.6	297
Method B	443	287
Method G	-752	286
Offset variation (Fig. 9); x=η		
Method A	-550	56.2
Method B	-355	56.4
Method G	-324	67.7
Amplitude variation (Fig. 15); x=μ		
Method A	-280	68.0
Method B	-0.97	67.8
Method G	-184	71.3

departure from straight line behaviour at higher noise levels. It was found that the points in the upper graph in Fig. 4 are well fitted by a simple quadratic function. The coefficients of those quadratic fits are given in Table 1, along with the coefficients of linear fits to the points in the lower graph.

3.2 Variation in ellipse offsets

In this sub-section and the following one we consider what happens when the ellipse amplitude and offset are not stable during the ellipse measurement. We assume that the variations in these parameters are sufficiently slow that only the first derivative is significant on the timescale of a single ellipse. That is, it is only the slope of the amplitude or offset variation that needs to be considered. Accordingly, equations (1) and (2) can be modified as follows:

$$x = a\,(1 + \mu_x t)\sin\theta + b + a\eta_x t \tag{6}$$

$$y = c\left(1 + \mu_y t\right)\sin\left(\theta + \phi\right) + d + c\eta_y t \tag{7}$$

where t is a dimensionless measurement time, which goes in equal steps from 0, when the first measurement is taken, to 1-1/N, when the final measurement is taken, there being N measurements in total. This means that the η values characterise the slope by the change in offset that occurs during the acquisition of one ellipse (the multiplication by a or c in (6) and (7) means that the slopes are relative to the ellipse amplitudes). In this section, we will consider the case when there are no amplitude variations, so the μ slopes are zero. Thus, there are two variable parameters, η_x and η_y, in addition to the variable ϕ. We do not include additive noise in this sub-section.

It is found that the variation of the ellipse parameters has two effects. First, it introduces bias: the measured phase differs from the true value of ϕ. Second, the common-mode variations in θ are no longer completely removed by the ellipse fitting. This second effect manifests as fluctuations in the measured phase. As in Sect. 3.1, our numerical simulations treat θ as a random variable which changes with each new pair of x and y values. The effect of this is to produce a random variation in the estimated phase $\hat{\phi}$, which we characterise, as before, by the standard deviation. One situation of practical interest is when $\eta_x = \eta_y$, so that the offset variations are the same in both interferometers. This could come about, for example, from intensity fluctuations in a laser beam which interacts with both atom clouds in the gradiometer. Examples of simulated ellipse data in this case are given in Fig. 5 for three different values of ϕ.

In order to better show the overall locus of points, 500 points are plotted. Note, however, that it is still 100 points that are used for all of the ellipse fitting results in this sub-section. The ellipse parameters a, b, c, d were the same as in the previous sub-section. The distortion of the ellipses in Fig. 5 can be interpreted as being due to the ellipse moving over time along a line with a slope of 45 degrees. Therefore, the greatest scatter in the locus of points occurs in regions where the curve of the ellipse is orthogonal to the 45 degree line. The corresponding plots with the offset slopes moving in orthogonal directions, that is $\eta_x = -\eta_y$, are related to those of Fig. 5 by the following transformation: replacing ϕ by $\pi - \phi$ and rotating the plot by 90 degrees. Thus, the ϕ =0.5 rad plot would transform to a rotated version of the ϕ =2.6 rad one (taking π -0.5≈2.6).

Figures 6, 7 and 8 show results for $\eta_x = \eta_y$ and $\eta_x = -\eta_y$, as well as some intermediate cases. These intermediate cases are constructed by keeping the sum of the squares of the two slopes constant and equal to 0.04 and varying the angle between them in steps of $\pi/8$; this can be thought of as using a vector in the $\{\eta_x, \eta_y\}$ space with a constant length of 0.2 but varying angle. The resulting offset variations are shown in the inset diagram in Fig. 5 (bottom right), which also provides a key for the differently coloured lines in the following

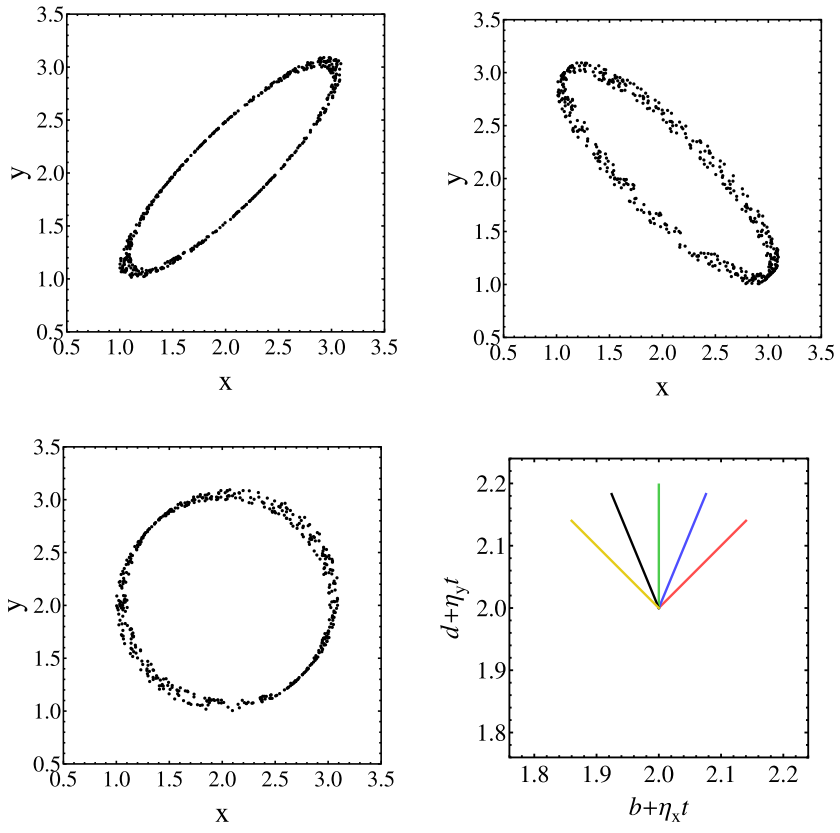

Figure 5 Three data sets with $\eta_x = \eta_y$ =0.1. Upper left ϕ =0.5 rad, upper right ϕ =2.6 rad, lower left ϕ =1.57 rad. The bottom right plot shows the values of η_x and η_y which will be used later in Figs. 6, 7 and 8

Figures. Note that, owing to the symmetry between x and y and between positive and negative directions of variation, results are only required for one quadrant of the η_x, η_y plane. Results for methods A, B and G are shown in Figs. 6, 7 and 8, respectively.

For Figs. 6 and 7 the phase ϕ was varied between 0.17 rad and 2.97 rad in steps of 0.1 rad. In Fig. 8 a step size of 0.1 rad was also used, but with a restricted range of phase values (between 0.67 and 2.37) because of the issues with convergence with method G for values near 0 and π rad, discussed earlier. For each phase value, 50,000 ellipses, each with 100 points, were used. In these figures, straight lines have been used to join the points.

Considering the random phase fluctuations first (lower plots in Figs. 6, 7 and 8), it can be seen that the standard deviations of the phase fluctuations are not symmetric about the central $\pi/2$ value, except in the case when the direction of offset motion is parallel to one of the axes (green lines). This can be easily understood by reference to the ellipse plots in Fig. 5. The difference in the appearance of the ellipse for $\phi < \pi/2$ (upper left plot in Fig. 5) compared to the one with $\phi > \pi/2$ (upper right) indicates how offset variations couple to phase variations. Phase variations correspond to opening and closing of the ellipse and, for this reason, one would expect to measure greater apparent phase variations in the upper right plot, where the scatter in the ellipse points is in the direction of opening and closing of the ellipse. This means that when offset variations in x and y are correlated, phase fluctuations will be smaller if the ellipse phase is small. Conversely, if the offset variations are anti-correlated the phase fluctuations will be smaller if the ellipse phase is large. If the

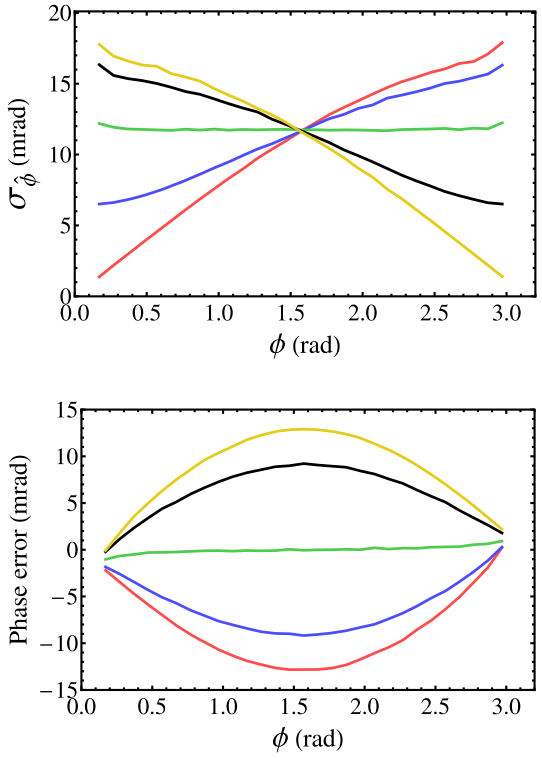

Figure 6 Ellipse fit with offset variation, method A. See the inset diagram in Fig. 5 (bottom right) for the key to the different lines

direction of offset motion is parallel to one of the axes there is little overall difference. This behaviour is seen with all the ellipse-fitting methods. There are some minor differences: with method A the standard deviation varies between 5.3 mrad and 15.7 mrad, with method B the variation is between 5.4 mrad and 14.4 mrad, and with method G the variation is between 7.8 mrad and 13.1 mrad. There are also some differences seen as ϕ is extended towards 0 and π, with results from method B showing a reduction which is not seen with method A.

The results for bias (upper plots) show some notable differences in the behaviours of the different ellipse fitting methods. One common factor is that the greatest bias occurs when the direction of the offset variation is either along the 45 degree line (in the x, y plane) or orthogonal to it (i.e. maximum correlation or maximum anti-correlation), and the minimum bias is when the offset direction is parallel to the x or y axis. With method A (Fig. 6) the greatest bias is when $\phi = \pi/2$ and the bias tends to zero as ϕ goes to 0 or π. With method B, however, the bias is less symmetrical and, in most cases, is low at one end of the 0 to π range and high at the other. The behaviour with method G is similar to method B but with opposite slope.

Figure 9 shows results for a fixed value of ϕ =1.0 rad and the magnitude of the offset slopes varied. These simulations used 100,000 ellipses for each slope value. Here, η_x and η_y were both equal (so the subscript on η has been dropped), which means the offset is moving along the 45 degree line, which corresponds to the red lines in Figs. 6 to 8. For this value of ϕ, method A has the greatest phase error and method G the least (upper graph). It is found that the dependence on η is well fitted by a line proportional to the square of η,

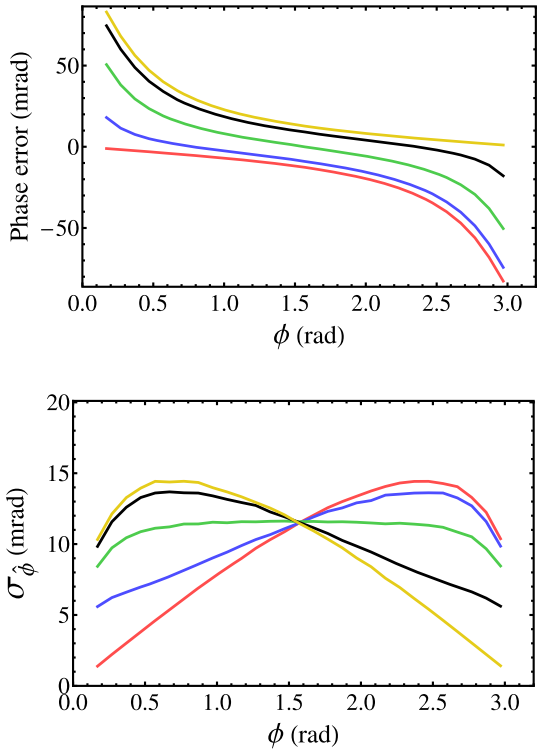

Figure 7 Ellipse fit with offset variation, method B. Lines as indicated in Fig. 5

that is, the bias increases as the second power of the offset slope; an example is shown by the solid line, which is fitted to the method A results. Coefficients of quadratic fits to all the points are given in Table 1, along with linear fits to the data in the lower plot.

The lower graph in Fig. 9 shows the standard deviation of the phase. For small values of η this has a linear dependence. For ϕ =1.0 rad, method A and method B give virtually identical standard deviations whereas method G gives somewhat larger values. All these results used ellipses with N=100 points. Simulations using different numbers of ellipse points, but the same η values, showed that the phase error (bias) is independent of N, even when N is as low as 10. This shows that it is the change in ellipse offset between the beginning and end of the ellipse measurement that determines the bias, rather than the number of ellipse points. The random phase fluctuations do, however, increase when there are fewer ellipse points, as one would expect. The behaviour of the phase fluctuations as a function of number of ellipse points N was found to be the same as that seen in the previous sub-section for additive noise: the phase fluctuation standard deviation is inversely proportional to the square root of N, as long as N is greater than around 30. So, we conclude that, if the data points are being produced at fixed intervals, the impact of ellipse offset variations can be minimised by dividing the data into relatively small groups, with N around 30, because this will minimise the overall effect of random phase fluctuations.

3.3 Variation in ellipse amplitudes

A similar approach to that of the previous sub-section was used to investigate the effect of variations in the ellipse amplitude, as parameterised by the slopes μ_x and μ_y in equations (6) and (7). This was done using fixed offsets (that is, η_x=η_y =0) and no additive noise. The

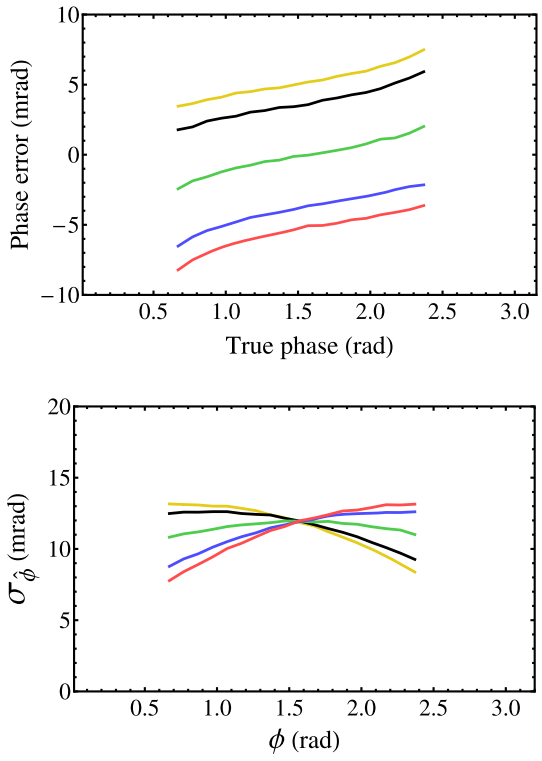

Figure 8 Ellipse fit with offset variation, method G. Lines as indicated in Fig. 5

same values of the ellipse parameters a, b, c, d were used. Some examples of y vs x scatter plots are given in Figs. 10 (for $\mu_x = \mu_y$) and 11 (for $\mu_x = -\mu_y$), which each plot having 1000 measurements with random phases θ. The situation here is more complicated than with the offset variation (Fig. 5), because the case of $\mu_x = \mu_y$ (Fig. 10) cannot be simply related to that $\mu_x = -\mu_y$ (Fig. 11) via symmetry.

Results of ellipse fitting with amplitude variation are shown in Figs. 12, 13 and 14 for methods A, B and G respectively. For each ϕ value, 50,000 ellipses were generated, each with N=100.

Considering the central region (that occupied by the data from method G), the overall range of bias values are fairly similar. With method A, the bias goes from -5.8 mrad to 5.7 mrad, with method B it goes from -3.4 mrad to 4.3 mrad, and with method G it goes from -6.0 mrad to 5.1 mrad. The phase standard deviations (lower plots) are similar for all three methods, with the greatest values occurring at $\phi = \pi/2$: method A varies from 5.8 mrad to 11.1 mrad, method B from 5.7 to 11.0, and method G from 6.1 mrad to 11.1 mrad.

The case of $\mu_x = \mu_y$ was used to investigate behaviour as a function of the amplitude slope, with $\mu_x = \mu_y = \mu$ and $\phi = 1$ rad. Results are given in Fig. 15; these used 400,000 realisations for methods A and B and 100,000 realisations for method G. In this case of equal x and y slopes, the phase error is very low with method B. The results with methods A and G have a similar functional dependence to that seen with the offset variation (Fig. 8) and can be well fitted by a line proportional to μ^2. Also, the standard deviation of the phase fluctuations is a linear function of μ when μ is small, with some minor differences in the

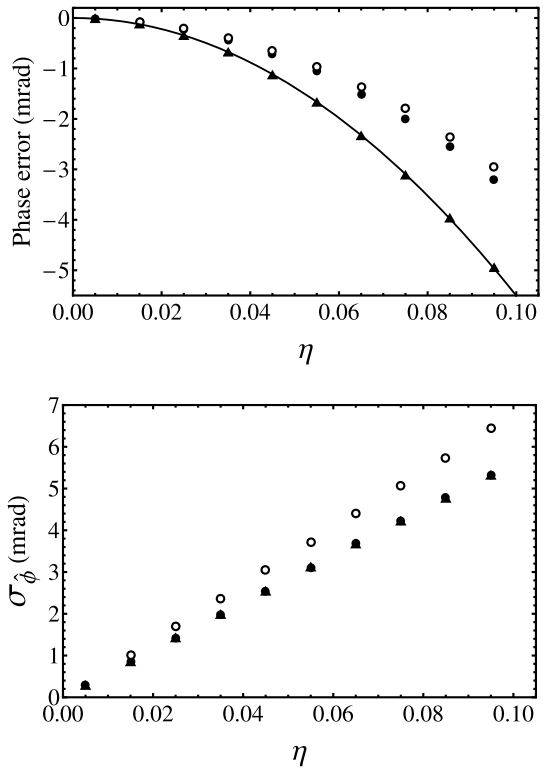

Figure 9 Upper plot, phase error as a function of offset slope. Lower plot, phase standard deviation. These are with slopes for x and y offsets, with ϕ =1.0 rad. Triangles: method A; solid circles: method B; open circles: method G

slope of the straight line depending on the fitting method. Data giving both sets of fitting coefficients can be found in Table 1.

Simulations were also carried out to investigate the effect of different values of N. As with the offset variations in the previous sub-section, the phase bias was found to be largely independent of N for fixed μ and the standard deviation reduces in inverse proportion to the square root of N when N is large.

3.4 Normalisation errors

The errors considered in this section are somewhat different to those in the previous section. They arise not because of time variation in the ellipse parameters but because the normalised fringes themselves are not sinusoidal and therefore when y is plotted against x the shape of the curve is not elliptical. This effect can occur during the detection of fluorescence from the atoms. The desired measurement is the fraction of atoms in the upper atomic level. This fraction can be determined by first measuring fluorescence from atoms in the upper level and then applying a laser pulse which transfers all atoms to the upper level and finally repeating the fluorescence measurement. Normalisation is carried out by taking the ratio of the two measurements to give the fractions x and y. However, the first measurement can perturb some of the atoms that are in the upper level so that they do not, later, contribute to the second measurement. This results in errors in the fractions x and y, errors which depend on the number of atoms originally in the upper level. Since the number of atoms in the upper level is directly proportional to $\sin(\theta)$ the number of atoms lost

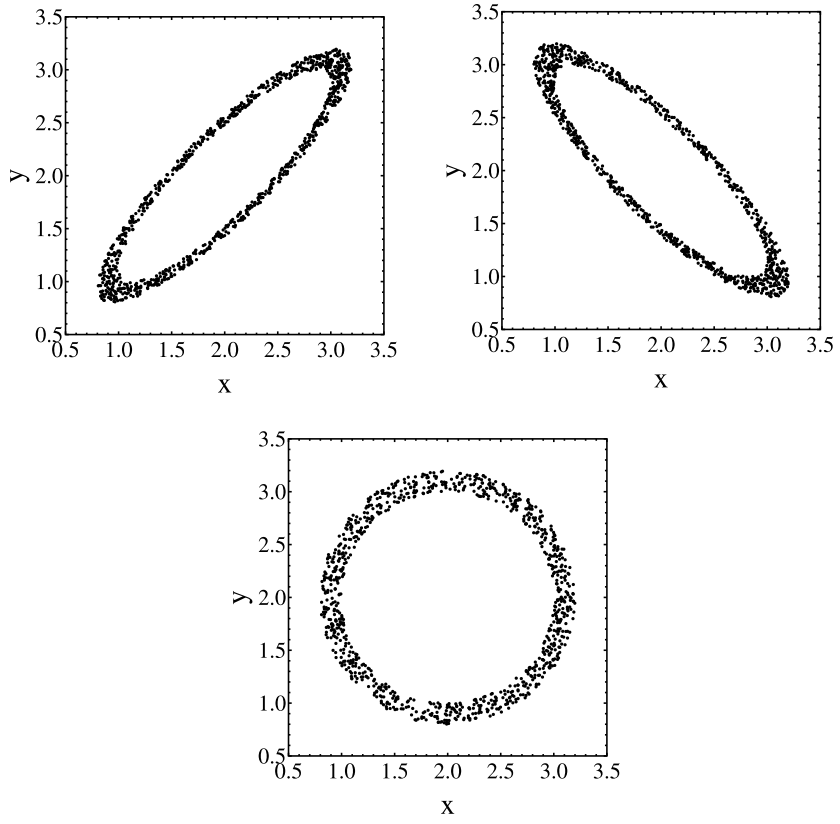

Figure 10 Three data sets with $\mu_x = \mu_y$ =0.2 and 1000 points. Upper left ϕ =0.5 rad, upper right ϕ =2.6 rad, lower ϕ =1.57 rad

from the second (normalisation) measurement is also proportional to $\sin(\theta)$. The effect on the measured fringe can be represented by introducing a negative term proportional to $\sin(\theta)$ in the denominator of the expression for x, as follows:

$$x = \frac{a \sin\theta + b}{1 - \varepsilon \sin\theta}, \tag{8}$$

and similarly for y

$$y = \frac{c \sin(\theta + \phi) + d}{1 - \gamma \sin(\theta + \phi)} \tag{9}$$

The deviations from the elliptical forms are characterised by new parameters ε and γ. When these parameters are non-zero, equations (8) and (9) can be written in a simpler form. For example, equation (8) can be written as

$$\frac{f}{1 - \varepsilon \sin\theta} + g, \tag{10}$$

where $g{=}a/\varepsilon$ and $f{=}b{+}a/\varepsilon$. The same transformation can be applied to equation (9). This new form represents a scaled and shifted version of a function which depends on ε and θ only (with ϕ additionally appearing in the expression for y). Since the ellipse fitting process accounts for offset and scaling factors, it follows that the phase angle derived from the

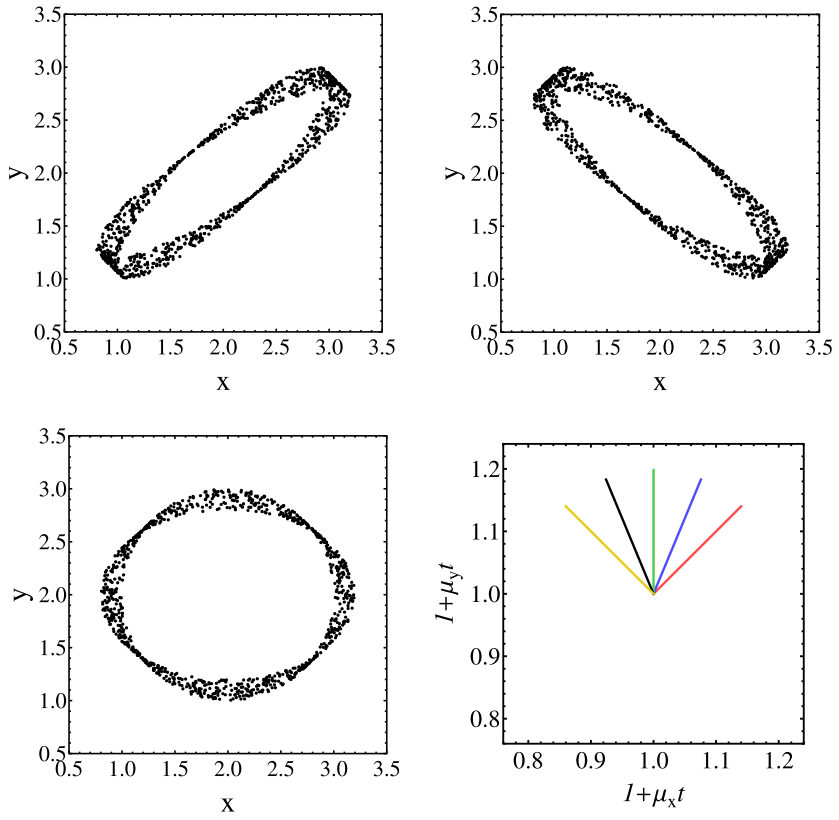

Figure 11 Three data sets with μ_x = 0.2, μ_y =-0.2 and 1000 points. Upper left ϕ =0.5 rad, upper right ϕ =2.6 rad, lower ϕ =1.57 rad. The bottom right diagram shows the values of μ_x and μ_y used in Figs. 12, 13 and 14

fitted ellipse will be a function of ε, γ and ϕ only, with f and g not affecting the phase. This simplifies the investigation of ellipse-fitting errors because they are independent of the actual values used for a, b, c and d. All the results presented here use a=1, b=0, c=1 and d=0. Some examples of distorted ellipses are shown in Fig. 16. These used values of ε=γ =0.1. These values are large compared to what one might see in practice and were chosen to clearly show the distortions introduced into the x-y plots.

The performance of the different fitting algorithms for ε=γ=0.1 are given in Fig. 17. The upper plot shows the bias as a function of ϕ. This used 100 points per ellipse and the same random statistics for θ that were used in the previous Sects. 50,000 estimates of ellipse phase were produced for each ϕ value. The lower plot shows the standard deviation of the phase estimates. It can be seen that the effects of the distortions get worse when the value of ϕ is larger. The effects with smaller values of ε and γ were also investigated (it was assumed that the values of the two parameters are equal, which makes sense for a sensor with a pair of identical interferometers). It was found that the standard deviation of the phase values decreased linearly and the bias decreased approximately quadratically, at least for the value of ϕ =1.57 rad, which was the only ϕ value investigated. For this value of ϕ the slope of the straight line fit to the standard deviation results was 75 mrad.

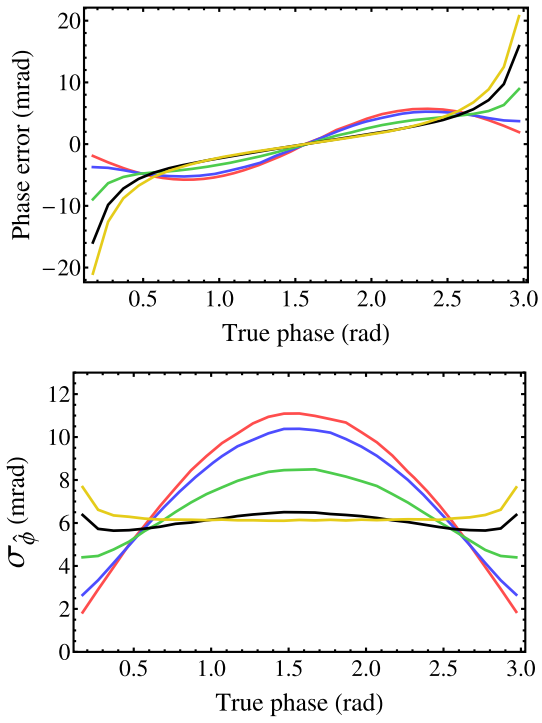

Figure 12 Ellipse fit with amplitude variation, method A. See the inset diagram in Fig. 11 (bottom right) for the key to the different lines

4 Discussion

We have investigated a number of ways in which ellipse fitting errors can manifest in an atom interferometer system and simulated these errors using three different ellipse-fitting methods. To the best of our knowledge, this is the first time that these specific errors have been investigated, apart from the case of bias in the presence of additive noise, which has been discussed in previous literature [17]. We note that Xu et al [16] investigated errors arising from non-common-mode phase fluctuations; however, this is a different source of error than those considered in the present paper.

We have considered two main consequences of ellipse fitting errors. The first is bias: the mean of the measured ellipse phase deviates from the true phase difference. The second is noise: either a direct manifestation of detection noise or arising from incomplete suppression of common mode noise in the phase θ. The implications of these errors will depend on what the cold-atom interferometer is being used for.

In some applications, a constant bias will not be problematic because the measurements are differential and the measured ellipse phase is always relative to a reference ellipse phase which is not greatly different in value. In a geophysical application, for example, a gravity gradiometer might be used to measure differences in the gravity gradient between different spatial locations, or a single location at different times. In this case, the resulting error will depend on the slope of the bias curve, because the differential nature of the measurement removes the first-order bias. As noted in the discussion following equation (3), for practical measurements the changes in the true phase will be very small, of order a few milliradians, which means that second order effects are not large. Considering the results in Fig. 3, for example, changes along the horizontal axis (the true phase) when carrying out

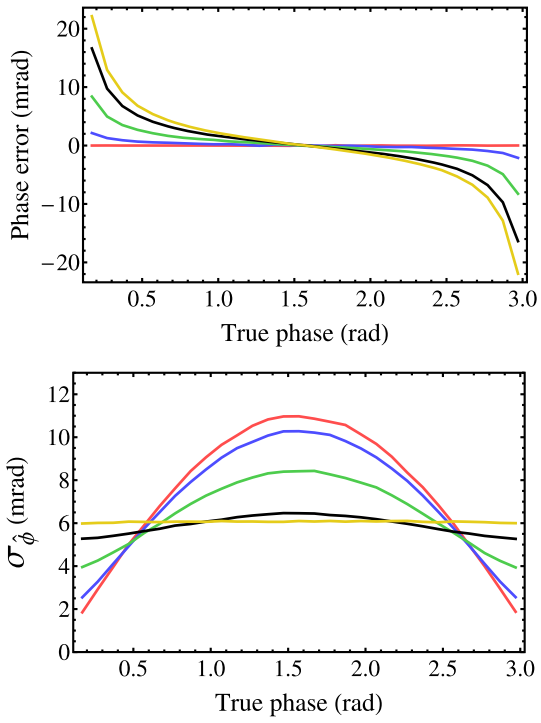

Figure 13 Ellipse fit with amplitude variation, method B. See the inset diagram in Fig. 11 (bottom right) for the key to the different lines

measurements of gravity at different locations or times may be of order 1 mrad. Even the ellipse fitting method with the greatest bias (method B) only has a slope of 20 mrad/rad, so a 1 mrad variation in the true phase will only produce a worst-case error of order 20μrad in a differential measurement: an error of one part in 50. If this level of bias is problematic then method A would be preferred, as this method has a bias that is more than an order of magnitude less than method B. This type of constant bias arises from the effect of additive noise (Sect. 3.1) and normalisation errors (Sect. 3.4), although in the latter situation the bias is only constant if the values of ε and γ are themselves stable over time. When considering ellipse offset variations (Sect. 3.2) and amplitude variations (Sect. 3.3), however, the bias cannot be assumed to be constant. In equations (4) and (5) it is assumed that the variations in ellipse parameters are sufficiently slow that they can be approximated as linear functions of time on the timescale of a single ellipse. However, on the timescale of multiple ellipses one would expect the slope to vary, from positive to negative and back again, as the underlying mechanism which affects the ellipse parameters varies. This will result in a bias which changes from ellipse to ellipse. Here, the significant performance-determining factor is how the change in the bias relates to the changing slopes; as shown, for example, in the upper plot in Fig. 9 for the offset. These plots can be used to impose a maximum allowable variation in offset or amplitude, given a requirement for phase accuracy (and thus gravity gradient accuracy). Clearly, to minimise the effects of changes that occur during acquisition of a single ellipse it is desirable to use as few points per ellipse as possible. However, using too few points leads to greater errors. We have shown that it is undesirable to use fewer than of order 25 to 30 points per ellipse for the ellipse-fitting methods investigated.

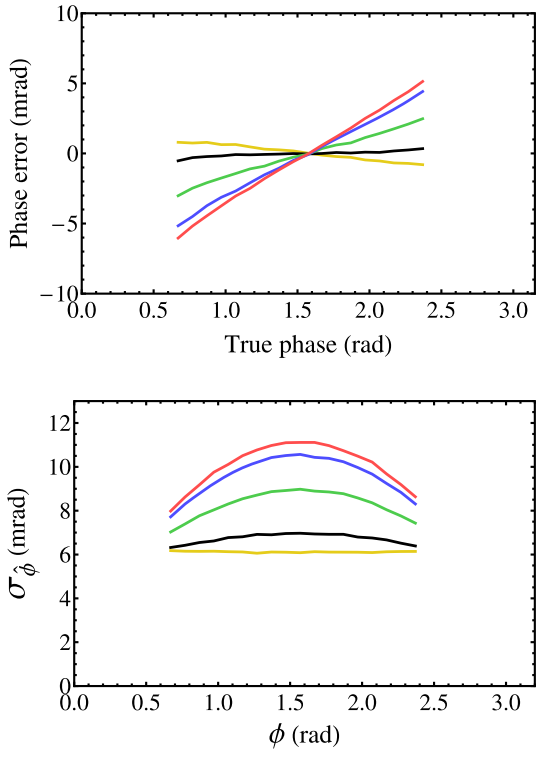

Figure 14 Ellipse fit with amplitude variation, method G. See the inset diagram in Fig. 11 (bottom right) for the key to the different lines

As far as phase fluctuations are concerned, these are expected in the presence of additive noise but their presence due to variation in ellipse offset or amplitude is a significant extra effect. This comes about because of the 'break-through' of common-mode noise, which is no longer entirely rejected by the ellipse-fitting procedure. For clarity, we have considered the effects of additive noise and ellipse parameter changes separately. In practice, they will occur together, since some level of additive noise is unavoidable. It is reasonable to assume that the resultant random ellipse-phase fluctuations will add in quadrature, since they arise from independent random variables. It is not as clear how the different biases will combine, and this could be a topic for further study.

In a cold-atom interferometer, there are various mechanisms by which the ellipse phase can be set to an arbitrary value (by changing the phase of the final Raman pulse, for example [4]), which means that there is scope to optimise the phase. It is clear that the ellipse phase should not be set close to 0 or π, as this produces narrow ellipses which have greater fitting errors. If getting the highest signal to noise ratio is of greatest importance, then there is a (small) benefit to be gained by using the geometric fitting method and setting the ellipse phase close to 0.7 rad or 2.4 rad (see Fig. 3). As far as optimising against offset or amplitude variations are concerned, it depends on whether the x and y variations are synchronised or not. In a gravity gradiometer, the two atom interferometers share a common pair of Raman laser beams. These beams are pulsed to apply the so-called $\pi/2$ and π interactions [12]. These pulses need to be of a specific intensity and duration. If the intensity varies then the operation of the interferometer is no longer optimum, and this will manifest as a reduction in amplitude and offset. Since the Raman pulses are com-

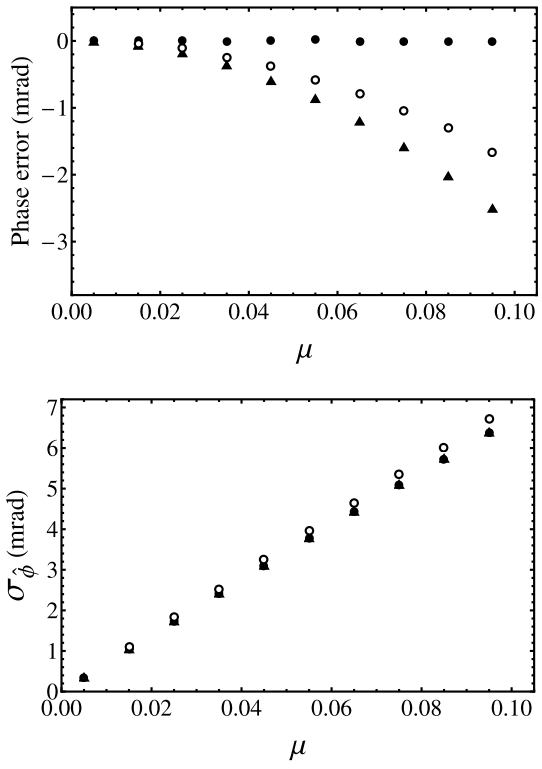

Figure 15 Upper plot, phase error as a function of amplitude slope. Lower plot, phase standard deviation. Equal slopes for *x* and *y* amplitudes., Triangles: method A; solid circles: method B; open circles: method G

mon to both interferometers of a gradiometer, the offset and amplitude variations will be correlated between x and y. This situation corresponds to the red lines in the figures in Sects. 3.2 and 3.3. In this case, it is clear that, with all of the ellipse fitting methods, the random phase fluctuations induced by offset variations are best suppressed by using a relatively small value of the ellipse phase, corresponding to a relatively narrow ellipse orientated about the x=y line. This is also the case for amplitude variations. With amplitude variations, the curves are symmetrical about $\pi/2$; however, there seems no good reason to choose a value near π because that would be worse in suppressing the effects of offset variation.

Stabilisation of ellipse amplitude and phase requires stabilisation of the Raman laser intensity. Referring, for example, to Fig. 9, which gives the increased phase noise resulting from offset variation, the maximum value of variation on the plot corresponds to a 10% offset change during the acquisition of the ellipse. This level of stability would indicate a rather poorly designed system, so is probably unlikely to be encountered in practice. A 1% variation seems more plausible for equipment that has to operate outside of a laboratory and is subject to large external temperature variations. Even at this lower level of variation, phase fluctuations of order 1 mrad are predicted by the results in Fig. 9, and these may not be negligible. Laser stabilisation requirements could be further investigated by combining ellipse modelling with a model of a cold-atom interferometer, in order to make a quantitative link between Raman intensity variations and ellipse offset and amplitude variations.

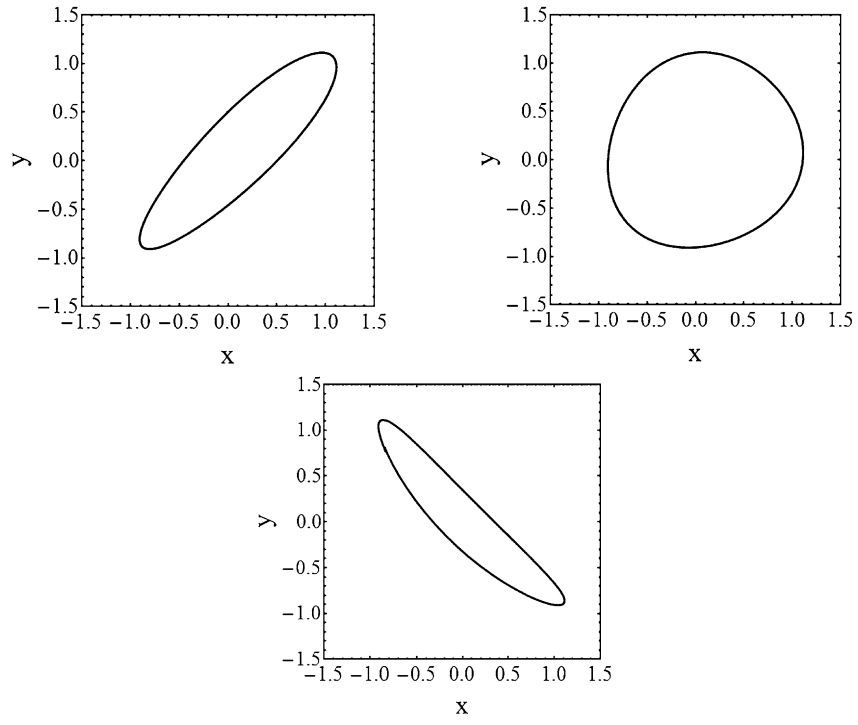

Figure 16 Distorted ellipses, with ε=γ=0.1. Upper left ϕ =0.5 rad, upper right ϕ =1.5 rad, lower ϕ =2.8 rad

The calculations presented here have considered variations that are slow compared to the time taken to acquire an ellipse. This implies slow variations in Raman intensity. If the intensity variations are faster than the time taken to acquire an ellipse, then the offset variations can be considered as an extra source of additive noise. To some extent, fast amplitude variations can be considered in the same way, but unlike the offset variations they will vary according to position around the ellipse, so further modelling would be required to fully understand their effect.

Finally, it should be noted that all of the results in this paper are based on the ellipse fitting method of extracting differential phase, but there are other methods which do not use ellipse fitting. In reference [19], a vibration-isolation platform was used to reduce the amount of common mode phase noise to the point where individual interferometer phases could be unambiguously extracted prior to being removed by phase differencing. In reference [20], an active self-calibration method was used to extract differential phase, and it was shown that this led to improved performance compared to ellipse fitting. In practice, the decision whether to use ellipse fitting or a differential phase extraction technique along the lines of those used in these references, will be determined by the specific application, and the environment in which the sensor is being deployed. An environment outside of a laboratory may impose size-weight and power limitations and be subject to high ambient vibration levels, both of which could make the relatively-simple ellipse fitting approach more attractive than methods that require additional vibration isolation or feed-back loops, even when those methods could, in principle, produce more accurate data.

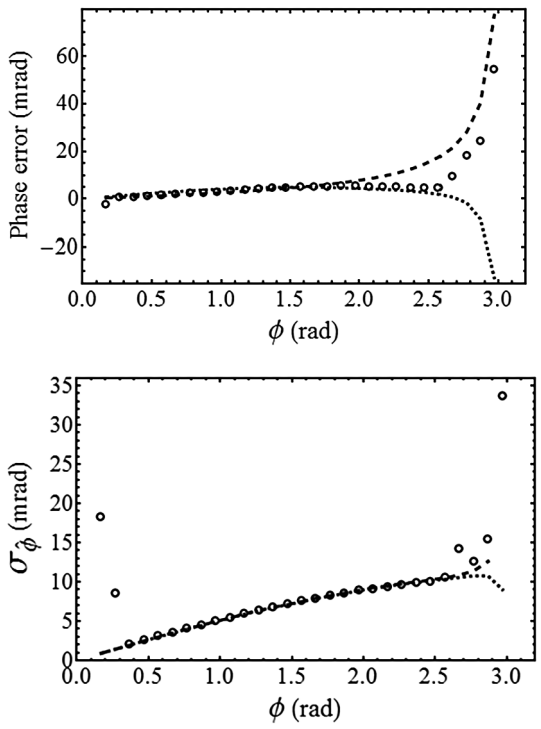

Figure 17 Distorted ellipse fits, with $\varepsilon=\gamma=0.1$. Dashed line, method A; dotted line, method B, circles, method G

5 Conclusions

We have used numerical simulation to investigate the effects of noise when an ellipse is fitted to a set of data points. This is of relevance to cold-atom interferometers, particularly when used in a differential mode to measure gravity gradients. Our results quantify how additive detection noise and unsuppressed common mode noise lead to errors in the measured interferometer phase difference: errors consisting of both bias and noise. An understanding of these effects is important for the design and operation of quantum sensors that use ellipse fittingto measure gravity gradients, particularly to meet the requirements for greater precision in future quantum sensor systems.

Author contributions

K.R. conceived the idea for this study and produced the results for ellipse-fitting methods A and B. A.R. produced the results for ellipse-fitting method G. Both authors contributed to the data analysis. K.R. wrote the first draft of the paper and generated all of the figures. A.R. contributed to subsequent drafts. All authors read and approved the final manuscript.

Funding

This work was funded by the Engineering and Physical Sciences Research Council (EP/T001046/1) as part of the UK National Quantum Technologies Programme, and the UK Defence Science and Technology Laboratory under the Gravity Imager project, contract number DSTLX-1000095040.

Data availability

The datasets used and/or analysed during the current study are available from the corresponding author on reasonable request.

Declarations

Competing interests

The authors declare no competing interests.

Received: 11 October 2023 Accepted: 6 November 2024 Published online: 22 November 2024

References

1. Foster GT, Fixler JB, McGuirk JM, Kasevich MA. Method of phase extraction between coupled atom interferometers using ellipse-specific fitting. Opt Lett. 2002;27:951–3.
2. Sorrentino F, Bodart Q, Cacciapuoti L, Lien Y-H, Prevedelli M, Rosi G, Salvi L, Tino GM. Sensitivity limits of a Raman atom interferometer as a gravity gradiometer. Phys Rev A. 2014;89:023607
3. Snadden MJ, McGuirk JM, Bouyer P, Haritos KG, Kasevich M. A measurement of the Earth's gravity gradient with an atom interferometer-based gravity gradiometer. Phys Rev Lett. 1998;81:971–4.
4. Stray B, Lamb A, Kaushik A, Vovrosh J, Rodgers A, Winch J, Hayati F, Boddice D, Stabrawa A, Niggebaum A, Langlois M, Lien Y-H, Lellouch S, Roshanmanesh S, Ridley K, de Villiers G, Brown G, Cross T, Tuckwell G, Faramarzi A, Metje N, Bongs K, Holynski M. Quantum sensing for gravitational cartography. Nature. 2020;602:590–4. https://doi.org/10.1038/s41586-021-04315-3.
5. Fixler JB, Foster GT, McGuirk JM, Atom KMA. Interferometer measurement of the Newtonian constant of gravity. Science. 2007;315:74–7. https://doi.org/10.1126/science.1135459.
6. Prevedelli M, Cacciapuoti L, Rosi G, Sorrentino F, Tino GM. Measuring the Newtonian constant of gravitation G with an atomic interferometer. Philos Trans R Soc A. 2014;372:20140030. https://doi.org/10.1098/rsta.2014.0030.
7. Parker RH, Chenghui Y, Weicheng Z, Estey B, Müller H. Measurement of the fine-structure constant as a test of the standard model. Science. 2018;360:191–5.
8. Zhou L, Long S, Tang B, Chen X, Gao F, Peng W, Duan W, Zhong J, Xiong Z, Wang J, Zhang Y, Zhan M. Test of equivalence principle at 10-8 level by a dual-species double-diffraction Raman atom interferometer. Phys Rev Lett. 2015;115:013004.
9. Graham PW, Hogan JM, Kasevich MA, Rajendran S. New method for gravitational wave detection with atomic sensors. Phys Rev Lett. 2013;110:171102
10. Gander W, Golub GH, Strebel R. Least-squares fitting of circles and ellipses. BIT Numer Math. 1994;34:558–78.
11. Kuhn CCN, McDonald GD, Hardman KS, Bennetts S, Everitt PJ. A Bose-condensed, simultaneous dual-species Mach-Zehnder atom interferometer. New J Phys. 2014;16:073035.
12. McGuirk JM, Foster GT, Fixler JB, Snadden MJ, Kasevich MA. Sensitive absolute-gravity gradiometry using atom interferometry. Phys Rev A. 2002;65:033608.
13. Hinze WJ, von Frese RRB, Saad AH. Gravity and magnetic exploration: principles, practices and applications. Cambridge: Cambridge University Press; 2013.
14. Fitzgibbon A, Pilu M, Fisher RB. Direct least square fitting of ellipses. IEEE Trans Pattern Anal Mach Intell. 1999;21:476–80.
15. Stockton JK, Wu X, Kasevich MA. Bayesian estimation of differential interferometer phase. Phys Rev A. 2007;76:033613
16. Xu A, Wang Z, Kong D, Fu Z, Lin Q. A new ellipse fitting method of the minimum differential-mode noise in the atom interference gravimeter. Chin Phys B. 2018;27:070203.
17. Szpak ZL, Chojnacki W, van den Hengel A. Guaranteed ellipse fitting with a confidence region and an uncertainty measure for centre, axes, and orientation. J Math Imaging Vis. 2015;52:173–99.
18. MATLAB Central File Exchange. Available online. https://www.mathworks.com/matlabcentral/fileexchange/15125-fitellipse-m (accessed on 5 Nov 2019).
19. McGuirk JM, Foster GT, Fixler JB, Snadden ML, Kasevich MA. Sensitive absolute-gravity gradiometry using atom interferometry. Rhys Rev A. 2002;65:033608
20. Chiow S, Williams J, Yu N. Noise reduction in differential phase extraction of dual atom interferometers using an active servo loop. Rhys Rev A. 2016;93:013602

Publisher's Note