<https://helda.helsinki.fi>

A geometric approach to non-linear correlations with intrinsic scatter

Pihajoki, Pauli

2017-12

Pihajoki , P 2017 , ' A geometric approach to non-linear correlations with intrinsic scatter ' ,
Monthly Notices of the Royal Astronomical Society , vol. 472 , no. 3 , pp. 3407-3424 . <https://doi.org/10.1093/mnras/s>

<http://hdl.handle.net/10138/228246>

<https://doi.org/10.1093/mnras/stx2179>

unspecified

publishedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

A geometric approach to non-linear correlations with intrinsic scatter

Pauli Pihajoki[★]

Department of Physics, University of Helsinki, Gustaf Hållströmin katu 2a, FI-00560 Helsinki, Finland

Accepted 2017 August 22. Received 2017 August 10; in original form 2017 April 13

ABSTRACT

We propose a new mathematical model for $n - k$ -dimensional non-linear correlations with intrinsic scatter in n -dimensional data. The model is based on Riemannian geometry and is naturally symmetric with respect to the measured variables and invariant under coordinate transformations. We combine the model with a Bayesian approach for estimating the parameters of the correlation relation and the intrinsic scatter. A side benefit of the approach is that censored and truncated data sets and independent, arbitrary measurement errors can be incorporated. We also derive analytic likelihoods for the typical astrophysical use case of linear relations in n -dimensional Euclidean space. We pay particular attention to the case of linear regression in two dimensions and compare our results to existing methods. Finally, we apply our methodology to the well-known $M_{\text{BH}} - \sigma$ correlation between the mass of a supermassive black hole in the centre of a galactic bulge and the corresponding bulge velocity dispersion. The main result of our analysis is that the most likely slope of this correlation is ~ 6 for the data sets used, rather than the values in the range of ~ 4 – 5 typically quoted in the literature for these data.

Key words: methods: data analysis – methods: statistical – galaxies: fundamental parameters.

1 INTRODUCTION

An important question in all the sciences is whether different measured observables seem to be connected by some form of mathematical relation, or whether they seem to be completely independent. If a relation is suspected, it then becomes important to estimate the type of relation connecting the measurements and to estimate its parameters.

Historically, these problems have been most often approached by the use of various correlation coefficients and linear and non-linear regression. Linear regression in particular has been the subject of lively debates and much research over the past two centuries or so, ever since first derived in the guise of least-squares regression by Gauss (or possibly Legendre, see Stigler 1981). The discussion has included such points of contention as how to treat the variables symmetrically – that is how to avoid dividing the variables into dependent and independent variables – and whether this is necessary (Pearson 1901; Boggs, Byrd & Schnabel 1987; Isobe et al. 1990; Feigelson & Babu 1992; Robotham & Obreschkow 2015). This issue is particularly important when investigating the hypothesis that the observables obey some mathematical relation. By definition, such a relation as a geometric structure (e.g. a line, a plane or some more complicated subset) is unique, even though it may have many equivalent formulations. As such, any method used for estimating such a relation should eventually yield the same structure even

if the variables are permuted or a different variable is chosen as the dependent variable. This requires that the method must involve all observed variables in an essentially symmetric way. However, this is not true for least-squares regression and many other similar approaches as well (see Section 3.3.1). This fact has been well known for a long time (see e.g. Pearson 1901; Isobe et al. 1990), but it appears somewhat underappreciated. Since the focus of this paper is on formulating and estimating relations between observables, we will pay close attention to this symmetry invariance throughout.

Another recurrent theme is the question of how to treat data subject to censoring (i.e. lower and upper limits), truncation (non-detections) and heteroscedastic, independent and correlated errors in general, and finally, how to incorporate intrinsic scatter, or intrinsic uncertainty in the regression hyperplane (see e.g. Kelly 2007; Hogg, Bovy & Lang 2010; Robotham & Obreschkow 2015; and the references therein). Intrinsic scatter can be intuitively understood to mean that the data probability distribution can be non-zero also in the neighbourhood of the subset defined by the relation; that is, *we also allow for data that may be somewhat off the relation* and thus not exactly satisfying the equation(s) defining the relation. For example, in the case of a line in two dimensions, intrinsic scatter would manifest as a data distribution that is not a linear delta function ridge, but a ‘fuzzy’ line instead. Despite this intuitive obviousness, intrinsic scatter has a history of various somewhat non-rigorous definitions, such as through an additional additive component in measurement errors. In this paper, we will present a more rigorous definition of intrinsic scatter that is also usable in non-Euclidean contexts.

[★] E-mail: pauli.pihajoki@iki.fi

Specific examples of papers addressing some of the problems above include Pearson (1901), who introduced a least-squares method to fit lines and hyperplanes based on minimizing residuals orthogonal to the regression plane (OR, orthogonal regression). Later, algorithms were developed to solve this problem for non-linear relations as well (see e.g. Boggs et al. 1987; Boggs et al. 1988; and the references therein). In Kelly (2007) and Hogg et al. (2010), problems related to outliers, truncation, censoring and intrinsic scatter are solved with a fully Bayesian approach, with some limitations. Namely, Kelly (2007) requires specifying a single dependent variable, and it is for this variable only that censoring is supported, while Hogg et al. (2010) only consider the two-dimensional case. Finally, Robotham & Obreschkow (2015) extends the results in Hogg et al. (2010) to n dimensions, for arbitrary n and presents analytic likelihoods for $n - 1$ -dimensional hyperplanes with intrinsic scatter (i.e. lines in a plane and planes in three-dimensional space), assuming data with Gaussian errors.

However, there still exist some remaining issues related to fitting non-linear relations to data that have been discussed to a much lesser extent. One of these is the notion of pre-existing geometry in the measured quantities, such as for angular quantities, or measurements of points on curved surfaces (however, see e.g. Pennec 2006; Calin & Udriste 2014). Another is the question of a proper characterization of intrinsic scatter for non-linear relations, and how to incorporate intrinsic scatter when the n -dimensional data is not well described by an $n - 1$ -dimensional subspace (i.e. of codimension one), but an $n - k$ -dimensional subspace (codimension k), for an arbitrary $k > 1$. A suitable resolution of these issues is of interest, since it would enable powerful hypothesis testing via finding the most likely value of k and distribution of intrinsic scatter for each proposed linear or non-linear relation simultaneously.

In this paper, we propose a solution to these issues by formulating the concept of non-linear relations with intrinsic scatter of general codimension k through Riemannian geometry. The novelty of our approach is in extending the idea of intrinsic scatter to curved spaces and correlations that are non-linear and have codimension greater than one, that is, not restricted to linear hyperplanes of dimension $n - 1$. The approach also accommodates arbitrary measurement errors, censoring and truncation. Furthermore, we give analytic results for the likelihood and posterior probability for linear $n - k$ -dimensional relations in Euclidean spaces. These are easy to implement in fitting codes, and useful for hypothesis testing by enabling quick determination of the most likely codimension of a potential correlation in n -dimensional data, along with the parameters of the intrinsic scatter.

In Section 2, we give definitions of correlation relations and intrinsic scatter distributions through the use of Riemannian submanifolds and an intrinsic coordinate system. The definitions allow the concept of intrinsic scatter distributions to be smoothly extended to non-linear relations and correlations of codimension greater than one. We also give a Bayesian solution to the problem of finding the most likely parameters for a given relation and intrinsic scatter, incorporating censoring, truncation and independent, general distributions of measurement errors. In Section 3, we apply our formulation to the special case of linear relations in Euclidean spaces. In particular, we extend the results in Robotham & Obreschkow (2015) to $n - k$ -dimensional subspaces with intrinsic scatter, which correspondingly extends their usefulness to a much wider range of applications. We also extensively discuss the implications of the formalism in the case of linear regression in two dimensions, with comparisons to well-established existing methods. Finally, in Section 4, we apply our methodology to the well-known $M_{\text{BH}} - \sigma$

relation between the mass M_{BH} of a supermassive black hole in a galactic bulge and the bulge velocity dispersion σ . Finally, we present a summary of the paper in Section 5.

We note that Section 2 is by nature somewhat technical, and relies heavily on concepts in Riemannian geometry. However, the salient points of the formulation we propose are illustrated in Figs 1–3. For the reader with immediate applications in mind, we suggest looking at the likelihood functions derived for the special cases of an $n - 1$ -dimensional linear relation (equation 19), $n - k$ -dimensional linear relation (equation 28) and the results for the two-dimensional case (equations 33 and B3–B6).

2 RELATIONS AND INTRINSIC SCATTER

2.1 Relations as submanifolds

The dual aim of this paper is first to give a geometric formulation of relations with intrinsic scatter that is intuitive and adapted to the geometry of the relation and the space of measurements. Secondly, to fit these relations to data, we aim to provide a method that is fully Bayesian and places no observable in a privileged position, i.e. there is no division into independent and dependent variables. In addition, we wish to accommodate truncated and censored data sets and independent and general distributions of measurement errors. To this end, we need to carefully define what we mean by a relation and its intrinsic scatter. We also need equal care in defining what we mean by being off the relation to be able to define intrinsic scatter in the first place.

In this paper, we will consider relations as submanifolds of the manifold of all possible combinations of measurement values, which we take to have a Riemannian geometry. Being off the relation will be related to the geodesic (shortest path) distance between a point in the measurement manifold, representing a single possible combination of measurements, and the relation submanifold. In this approach, the joint distribution of the measured observables is the primary object of interest. It is this joint distribution that we wish to model with the relation submanifold and the intrinsic scatter. Consequently, the distributions of the individual observables, obtained as marginal distributions, are of secondary interest only.

We start by defining a relation as a smooth map $f : M \rightarrow \mathbb{R}^k$, where M is an n -dimensional manifold with a Riemannian metric g_{ab} and $k \leq n$. The relation f is fully specified by the k component functions f_i , so that $f(m) = (f_1(m), \dots, f_k(m))$, for $m \in M$. Furthermore, the relation f may depend on n_p parameters $\theta \in \mathbb{R}^{n_p}$. In the following, we assume that we are always working with some local coordinate chart $\phi_x : M \rightarrow \mathbb{R}^n$ and at times use a shorthand $\mathbf{x}$ for the point $m \in M$ for which $\phi_x(m) = (x^1, \dots, x^n) = \mathbf{x}$. These coordinates represent the observable quantities (e.g. angle, charge, distance, velocity and mass) in the chosen units, or as scale-free logarithmic measurements. The manifold M then represents all possible combinations of measured values for these observables.

Given M and f , the level set

$$\begin{aligned} S &= \{m \in M | f(m) = \mathbf{0}\} \\ &= \bigcap_{i=1, \dots, k} S_i = \bigcap_{i=1, \dots, k} \{m \in M | f_i(m) = 0\} \end{aligned} \quad (1)$$

defines an $n - k$ -dimensional (or of codimension k) subset $S \subseteq M$.¹ We further require that S is a regular level set; that is, there are no

¹ A general level set would be $f(m) = \mathbf{c}$, $\mathbf{c} \in \mathbb{R}^k$, but we assume that the constant $\mathbf{c}$ has been subsumed in the definition of f .

Figure 1. A figure demonstrating the relationships between the measurement manifold M (shown in violet), the manifold S defined by the relation (brown), the exponential mapped normal space $\exp_p(N_p S)$ (orange) and the coordinate charts ϕ_x and ϕ_f in a case when $\dim(M) = 3$. Shown are cases where $\dim(S) = 2$ (left) and $\dim(S) = 1$ (right). The coordinates for a point p on S are either $\phi_x(p) = (x^1(p), x^2(p), x^3(p))$ or $\phi_f(p) = (\tau^1(p), \tau^2(p), 0)$ (left) or $\phi_f(p) = (\tau^1(p), 0, 0)$ (right). For m , $\phi_f(m) = (\tau^1(m), \tau^2(m), v^1(m))$ (left) or $\phi_f(m) = (\tau^1(m), v^1(m), v^2(m))$ (right). The origin and orientation of the components τ can be set arbitrarily on S , but the origin of the v is fixed on the point p .

Figure 2. An example of a one-dimensional relation submanifold S in a two-dimensional Euclidean manifold M . The coordinate transformation between ϕ_x and ϕ_f can be solved with equations (3).

$m \in M$ for which $f(m) = 0$ and the pushforward df_m fails to be surjective, so that S is additionally an embedded (or regular) submanifold of M , as is each S_i (see Lee 2013, for a complete discussion). S then represents the locus of our relation.

The most geometrically obvious way to define being off the relation is to relate it to the *geodesic distance* of a point from the submanifold S . To this end, it is useful to construct a new coordinate system defined with the help of the normal bundle

$$NS = \bigcup_{p \in S} N_p S = \bigcup_{p \in S} \{z \in T_p M \mid g(z, w) = 0, \forall w \in T_p S\} \quad (2)$$

of S , essentially consisting of all the tangent vectors of M that are orthogonal to S at each point p on S . We now define a coordinate system ϕ_f on M by defining $\phi_f(m) = (\tau(p), v(z)) \in \mathbb{R}^{n-k} \times \mathbb{R}^k$, where $m = \exp_p(z)$, $(p, z) \in NS$, τ is an arbitrary coordinate system on S and v is an orthonormal coordinate system on $N_p S$. Here, $\exp_p: TM \rightarrow M$ is the exponential mapping around point p . In these coordinates, the geodesic distance δ between m and the submanifold S is just $\delta = \sqrt{g(z, z)}$. Intuitively, the coordinates ϕ_f correspond to covering the original manifold M using normal coordinates around each point $p \in S$, or $M \subseteq \bigcup_{p \in S} \{p\} \times EN_p$, where we have used a shorthand $EN_p = \exp_p(N_p S)$. See the explanatory Fig. 1. The spaces EN_p contain all the points that can be reached from p by geodesics orthogonal to S at p . However, the subspaces EN_p do not necessarily fit together neatly. In general, ϕ_f^{-1} fails to be injective. This happens when two or more geodesics $\gamma_i(t) = \exp_{p_i}(t z_i)$, $i \in I \subset \mathbb{Z}_+$, intersect each other at some $m \in M$. Usually normal coordinates are

Figure 3. Samples from a normally distributed p_{int} (equation 6) for non-linear one-dimensional relations $f(x) = y - g(x) = 0$ in $\mathbb{R}^2$. The chosen relations are an exponential (left), a parabola (middle) and a logistic curve (right). The relations are shown in black, and the sampled points with red dots. The samples are given a two-dimensional zero-correlation Gaussian probability distribution with standard deviations σ_x and σ_y . The background heatmaps show the summed sample probability density. The salient feature in the images is the enhancement of the probability density towards the centre of curvature, most prominently seen in the middle and right images.

only defined up to these points, which define the boundary of the region where $\exp_p$ is a diffeomorphism. In this case, however, it is advantageous to let EN_p contain the points reachable from p by geodesics of unlimited length. In this case, ϕ_f may be multivalued, such that $\phi_f(m) = \{(\tau(p_i), \mathbf{v}(v_i))\}$, $i \in I$. This is discussed in the Sections 2.2 and 2.3. When M is complete and connected and S is closed, it is known that ϕ_f^{-1} is surjective (Wolf & Zierau 1996). In some pathological cases, ϕ_f^{-1} may fail to be surjective, such as when the relation submanifold S has sharp corners, but in this paper, we consider only cases where S is sufficiently smooth.

2.1.1 A concrete example

A prototypical astrophysical case is that of a relation between two real-valued observables. In this case, the manifold M is $\mathbb{R}^2$, the local coordinates are the identity map $\phi_x(x, y) = (x, y)$ and the relation is $f(x, y; \theta)$, parametrized by n_p parameters $\theta \in \mathbb{R}^{n_p}$. The normal spaces are $N_p S = \{\lambda \cdot (\nabla f)(p) | \lambda \in \mathbb{R}\}$ and we can identify $\exp_p(N_p S)$ with $N_p S$ itself. If we can put our relation f in the form $y - g(x) = 0$, we can find the coordinate transformation $\phi_f \circ \phi_x^{-1}$ as follows. If $m \in M$ is the point under consideration, and $p \in S$ is the point on S geodesically closest to m , with $\phi_x(m) = (m_x, m_y)$ and $\phi_x(p) = (p_x, g(p_x))$, then we can take for example $\tau(p_x, g(p_x)) = p_x$. With this definition, we can in principle solve v and τ as a function of m_x and m_y from

$$m_x = p_x + v \frac{-g'(p_x)}{\sqrt{1 + g'(p_x)^2}} \quad (3a)$$

$$m_y = g(p_x) + v \frac{1}{\sqrt{1 + g'(p_x)^2}}. \quad (3b)$$

See Fig. 2. Unfortunately, analytic solutions for these equations are not straightforward to derive except in the case where $g(x)$ is linear.

2.2 Intrinsic distributions

We now have a coordinate system ϕ_f with which we can define intrinsic scatter in a satisfying way. In fact, we can go beyond simple intrinsic scatter and define a general unnormalized probability

distribution $p_{\text{int}}(\tau, \mathbf{v})$ on M , adapted to the relation f through the coordinates ϕ_f . This amounts to setting up a k -dimensional probability distribution on $N_p S$ for each point $p \in S$ and using the exponential mapping to create an n -dimensional distribution on M . However, as mentioned above, ϕ_f is not injective when the geodesics normal to S intersect each other. In these cases, all the contributing points should be included, with corresponding probabilities added, so that for a point $m \in M$ given in the original coordinates, $\phi_x(m) = \mathbf{x}$, we have

$$p_{\text{int}}(\mathbf{x}) = \sum_{i \in I} p_{\text{int}}(\tau_i(\mathbf{x}), \mathbf{v}_i(\mathbf{x})), \quad (4)$$

where $I \subset \mathbb{Z}_+$ indexes all the points on S corresponding to m . This definition is necessary for cases such as the wrapped normal distribution on a circle S^1 .

The distribution p_{int} can represent e.g. a physical process producing objects with correlated values of some physical parameters, but with intrinsic scatter given by some probability distribution. An example would be the process producing the correlation between the central black hole mass and velocity dispersion in galaxies (see Section 4). The distributions of the possible physical measurements, as observed in the chart ϕ_x , are implicitly defined as the marginal distributions of the distribution p_{int} . These can be improper, as can the distribution p_{int} itself. That is, we may have $\int_M p_{\text{int}} \text{vol}_{\phi_x} = \infty$, where $\text{vol}_{\phi_x} = \sqrt{|g|} dx^1 \wedge \dots \wedge dx^n$ is the volume form on M in the chart ϕ_x , and $g = \det(g_{ab})$. A concrete example of an improper p_{int} is given below.

For applications, a useful distribution is the k -dimensional normal distribution, independent of τ , or the position on S , defined by

$$p_{\text{int}}(\mathbf{v}; \Sigma) = \frac{\exp[-\frac{1}{2} \mathbf{v}^T \Sigma^{-1} \mathbf{v}]}{\sqrt{(2\pi)^k \det(\Sigma)}}, \quad (5)$$

where Σ is the covariance matrix. Of particular importance is the case where $k = 1$, in which case we have a one-dimensional normal distribution depending only on v , for which

$$p_{\text{int}}(v; \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{v^2}{2\sigma^2}\right), \quad (6)$$

where the parameter σ , standard deviation, is now an accurate representation of what is typically called ‘intrinsic scatter’ in astrophysical literature.

It should be noted that the intrinsic scatter distributions (5) and (6) can yield a probability distribution that is not proper. An example is a line $y = \alpha + \beta x$ in flat two-dimensional space, for which, in the coordinate chart, we have (for more details, see Section 3.3)

$$p_{\text{int}}(x, y) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{(y - \alpha - \beta x)^2}{2\sigma^2(1 + \beta^2)} \right], \quad (7)$$

for which $\int_{\mathbb{R}^2} p_{\text{int}}(x, y) dx dy$ diverges. Similarly, in this case the marginal distributions for x and y ,

$$p_x(x) = \int_{\mathbb{R}} p_{\text{int}}(x, y) dy = \sqrt{1 + \beta^2} \quad (8)$$

$$p_y(y) = \int_{\mathbb{R}} p_{\text{int}}(x, y) dx = \sqrt{1 + \beta^2}/|\beta|, \quad (9)$$

are similarly improper. If necessary, improper distributions p_{int} arising this way can be made proper e.g. with a suitable truncation in the coordinates τ followed by normalization.

Fig. 3 depicts samples of the distribution (6) for three non-linear relations. From the figure, it is easy to appreciate the fact that intrinsic scatter can lead to distinct patterns of amplification and attenuation in the probability density near the regions where the underlying relation is strongly curved. As such, a misleading result would be obtained from a fit to such data, if the intrinsic scatter was not modelled properly, or at all.

Finally, we note that intrinsic scatter could also have been defined in two other ways, either along coordinate directions or the scalar geodesic distance to S . This seems to exhaust the possibilities, since in general the only directions we have available are the coordinate directions, or alternatively directions along the relation and orthogonal to the direction. In addition, only geodesic distance is a sensible natural choice in a Riemannian context. It should be emphasized that the definitions of intrinsic scatter using the intrinsic coordinates ϕ_f or the coordinate chart ϕ_x are in principle completely equivalent, and merely formulated in different coordinates. Both are more general than a definition based only on geodesic distance from the relation S .

A definition based on coordinate directions is often used at least in astronomical literature, where in the case of a line in two dimensions (x, y) , the intrinsic scatter is typically taken to lie in the direction of the y -coordinate (e.g. Press et al. 1992; Tremaine et al. 2002; Gültekin et al. 2009). This gives a relation that is ‘puffed up’ in the y -direction. This approach has the obvious problem of not being invariant with respect to a change of coordinates. However, it may be reasonable in some cases, where we are certain that whatever processes are producing the intrinsic scatter only operates strictly along a particular measured observable. Nevertheless, considering the often very correlated nature of astronomical observables, establishing this fact seems to be a difficult proposition. Consider for example the $M_{\text{BH}}-\sigma$ relation between the mass of a supermassive black hole and the central velocity dispersion of its host galaxy, where the intrinsic scatter is typically taken to lie entirely along the M_{BH} axis. As galaxies together with their SMBHs evolve, they move on the $M_{\text{BH}}-\sigma$ plane, but with different galaxies starting at (nearly) the same point, not necessarily ending up near each other later due to differences in composition, environment, merger history and so on. This produces intrinsic scatter. However, it seems very difficult to justify that these processes should operate only along the black hole mass axis. Rather, it seems that they might operate in any which direction indeterminately, but of these, the only one evident to us is *the direction away from the underlying relation*. Indeed, if the processes producing intrinsic scatter would only operate along the relation, the resulting data would have no scatter around the

relation at all. Finally, comparing the magnitude of intrinsic scatter between data sets becomes difficult, since the value will end up depending on the orientation of the line (see Section 3.3.1). In this sense, it seems reasonable to formulate intrinsic scatter problems in terms of the intrinsic coordinates ϕ_f , separately considering the directions orthogonal and parallel to the proposed relation.

As the last option, the intrinsic scatter could also have been defined as a probability density in the original measurement coordinate system ϕ_x , depending only on the scalar geodesic distance to S . This approach would give relations that are ‘puffed up’ symmetrically in all directions, with density falling off as a given function of the geodesic distance. Probability densities using this approach would not present the enhancements shown in Fig. 3. However, this approach is strictly less general than the one outlined above, and does not allow for e.g. wrapped distributions on spheres, toroids and so on, and cannot cope with directionally varying scatter for submanifolds S with codimension higher than one.

2.3 Fitting a relation with intrinsic scatter

In a typical case, we have obtained n_d measurements $\mathcal{D} = \{(m_1, h_1), \dots, (m_{n_d}, h_{n_d})\}$, where $m_i \in M$ are the measured values and $h_i : M \rightarrow \mathbb{R}$ are the probability distributions $p(\eta_i | m_i)$ for the true value η_i , given the measurement m_i , typically, representing the estimated measurement errors. The h_i are otherwise unspecified, and arbitrary distributions of measurement errors, including upper or lower limits (left or right censoring, respectively), can be naturally accommodated. For optimal results, the distributions h_i should be Bayesian posterior probability densities obtained in the measurement process, fully representing the available information. However, in some cases h_i are completely unknown, in which case they must be estimated, typically as Gaussians with unknown covariance matrices, which then must taken as additional nuisance parameters, in addition to the true values η . In addition to the measurements, we have a postulated relation f , parametrized by n_p parameters $\theta \in \mathbb{R}^{n_p}$, from which we derive coordinates $\phi_f(m; \theta)$, also parametrized by θ . We believe the data to fulfil the relation f , up to some intrinsic scatter p_{int} , parametrized by n_s parameters $\phi \in \mathbb{R}^{n_s}$. We would now like to find the most probable values of θ and ϕ , constrained by the data $\mathcal{D}$.

The Bayesian way to proceed is to note that for a single datum (m, h) , the conditional probability of the measurement factors as $p(m, h | \theta, \phi) = \int p(m, h | \eta) p_{\text{int}}(\eta | \theta, \phi) d\eta$, where η is the true value, drawn from the intrinsic distribution, p_{int} . As such, the true value η is taken as a nuisance parameter and integrated out. In addition, it may be that the measurement process is not fully sensitive across the range of possible values, leading to truncation, or data points that are not seen in the sample. If the probability that a given value η leads to a detection is given by $p_{\text{det}}(\eta)$, we have for the measurement probability that $p(m, h | \theta, \phi) = \int p(m, h | \eta) p_{\text{det}}(\eta) p_{\text{int}}(\eta | \theta, \phi) d\eta$ (see Appendix A for full derivation).

Now, taken as a function of the parameters, the probability $p(\mathcal{D} | \theta, \phi)$ defines the likelihood $\mathcal{L}$ of θ and ϕ . Assuming that the data $\mathcal{D}$ are independent, we get

$$\begin{aligned} \mathcal{L}(\theta, \phi | \mathcal{D}) &= \prod_{i=1}^{n_d} \int_M h_i(x(\tau, \mathbf{v}; \theta)) p_{\text{det}}(x(\tau, \mathbf{v}; \theta)) p_{\text{int}}(\tau, \mathbf{v}; \phi) \text{vol}_{\phi_f} \\ &= \prod_{i=1}^{n_d} \int_M h_i(x) p_{\text{det}}(x) \sum_{j \in I} p_{\text{int}}(\tau_j(x; \theta), \mathbf{v}_j(x; \theta); \phi) \text{vol}_{\phi_x}, \end{aligned} \quad (10)$$

where the integration can be done either in the coordinate chart ϕ_x or the intrinsic chart ϕ_f . The integration in ϕ_f automatically takes care of points $m \in M$ that have multiple representations (τ_j, ν_j) , but if the integration is done over ϕ_x , these have to be explicitly summed over.

If θ and φ have a joint prior distribution $\pi_{\theta, \varphi}$ (possibly uninformative), we can define the joint posterior probability distribution of θ and φ ,

$$p_{\text{post}}(\theta, \varphi) = \frac{\mathcal{L}(\theta, \varphi | \mathcal{D}) \pi_{\theta, \varphi}(\theta, \varphi)}{\iint \mathcal{L}(\theta', \varphi' | \mathcal{D}) \pi_{\theta, \varphi}(\theta', \varphi') d\theta' d\varphi'}. \quad (11)$$

The posterior distribution (equation 11) contains all the information of how the data constrains the parameters θ of the relation itself, as well as the parameters φ of the intrinsic scatter model, given our prior information $\pi_{\theta, \varphi}$. In some special cases, the posterior distribution can be computed analytically, but in general, it is necessary to employ numerical techniques, such as a suitable Markov Chain Monte Carlo (MCMC) method (see e.g. Gelman et al. 2013). Finally, point estimates such as the maximum (log-)likelihood estimate (MLE) and maximum a posteriori (MAP) estimate can be obtained using the likelihood via standard procedures. Full derivation of the likelihood (10) and the posterior distribution (11) following the style of Jaynes (2003) is found in Appendix A.

2.4 Some considerations

For the formulation presented above to make sense, a reasonable geometry must exist between the measured quantities. This is not a problem e.g. for measurements of a 3D angle, where we have the natural geometry of the sphere S^2 . In other cases, such as for joint measurements of galaxy luminosity and velocity dispersion it is reasonable to question whether a distance measured ‘across’ the measured quantities makes any sense. This observation has been discussed numerous times since Pearson in 1901 originally presented a least-squares method using orthogonal distances to the regression line (or plane), typically called orthogonal regression (OR; Pearson 1901). See for example Isobe et al. (1990) for a discussion.

The Riemannian approach presented in this paper makes it possible to specify the units and any possible geometry between the measured quantities beforehand. After the geometry is set, the metric character of the approach ensures that results are invariant with respect to all coordinate transformations, which in this case would typically represent scaling of the measured quantities (i.e. a change of units). Another natural possibility is to introduce the measurements in an entirely scale-free manner by using logarithms.²

A practical problem is the fact that the coordinate transformation between ϕ_x and ϕ_f may prove difficult or impossible to derive analytically. Finding the transformation and computing the posterior probabilities numerically necessitates computing numerous geodesic distances in a curved space. In general, this is highly computationally demanding, although efficient algorithms exist for some special cases (Crane, Weischedel & Wardetzky 2013).

Finally, we note that there are to be some consequences from and natural restrictions to the choice of parametrization and the prior distributions of the parameters. First, it is clear that an arbitrary

transformation of parameters $\theta = g(\theta')$ will result in the change of the prior probability measure so that $\pi(\theta)d\theta = \pi(g(\theta'))\det(J_g)d\theta'$, where J_g is the Jacobian of the transformation. This can be understood also through the fact that a probability distribution π defined in the parameter space also yields through the relation f and p_{int} a probability distribution on M . This latter distribution, once defined, should naturally be invariant under a change of parametrization, which is guaranteed by the transformation formula above.

Secondly, prior distributions are usually considered to be somewhat arbitrary and application-specific, often chosen for convenience to yield analytic results for the posterior probability. However, as noted above, this freedom can be restricted by what we know of how the data is distributed on M combined with the choice of a relation f and p_{int} . An important special case is if we have no prior knowledge of how the data are distributed, but only the possible values that it may take, through the choice of M . In this case, the maximum entropy distribution on M is uniform. This distribution places restrictions on the choice of $f(x; \theta)$, and after f is chosen, forces a particular choice of $\pi(\theta)$, which can then in this context be called *the* uninformative prior. The general argument is outside the scope of this paper, but see e.g. George & McCulloch (1993), Broemeling & Broemeling (2003) and Fraser et al. (2010) for historical and modern discussions on vague and uninformative priors from different points of view. However, a special case of an uninformative prior defined in the way described above is considered in detail in Section 3.3 and later used in Sections 3.3.1 and 4.

3 WORKED OUT EXAMPLES

The following sections contain analytical results for particular choices of the measurement manifold M and the relation f . In each case, we show how the framework presented above is used to obtain the likelihood, which can then be used for parameter estimation. The outline of the process, or a ‘practical how-to’, is as follows:

- (i) Determine the manifold M , along with the metric g . These are typically fixed by the nature of the measured quantities.
- (ii) Fix the functional form $f(x; \theta)$ of the relation to be fit, along with the parameters θ of the relation.
- (iii) Define the intrinsic coordinates (τ, ν) . A well-motivated choice for the coordinates τ along the relation is using geodesic distances from a given fixed origin. The coordinates ν are then defined up to a rotation of the ν -coordinate axes.
- (iv) Determine the coordinate transformation $\phi_x \circ \phi_f^{-1}$ from (τ, ν) to x or the inverse transformation $\phi_f \circ \phi_x^{-1}$. This may only be possible analytically in one direction only, or may not be possible analytically at all.
- (v) Choose a suitable intrinsic scatter model $p_{\text{int}}(\nu, \tau)$ as a function of the intrinsic coordinates.
- (vi) Evaluate the likelihood integral (equation 10). For analytical results, this typically needs to be done only for a single arbitrary data point, if all the data points have a similar error distribution model.
- (vii) Compute the likelihood over all data as a product of the likelihood of the individual likelihoods for each data point. Equivalently, choose a suitable prior and compute the posterior distribution (equation 11) over all the data.

We will refer to these steps in the following derivations so that the process can be more easily followed.

² Linear measurements x^i with a metric $g = \text{diag}(1, \dots, 1)$ are equivalent to logarithmic measurements $\log_{10} x^i$ and a metric with an exponential dependence on the measured values, $g = \log(10) \text{diag}(10^{x^1}, \dots, 10^{x^n})$.

3.1 Linear $n - 1$ -dimensional case

A highly useful special case is the case of normally distributed intrinsic scatter (equation 6) in a linear relation $f : \mathbb{R}^n \rightarrow \mathbb{R}$ between n observables, defining an $n - 1$ -dimensional (codimension 1) hyperplane S , with measurements drawn from an Euclidean space $M = \mathbb{R}^n$. For now, we assume that the metric g is given by the $n \times n$ identity matrix $\mathbf{I}$. These definitions complete step (i) from above. This special case represents well the typical astrophysical problem of modelling a linear relation between multiple physical variables that have no special geometry, such as mass, luminosity or velocity dispersion.

To address step (ii), we fix the parametrization. All the planes S can be parametrized with $\theta = \mathbf{p} \in \mathbb{R}^n$ so that

$$f(\mathbf{x}; \theta) = (\mathbf{x} - \mathbf{p})^T \mathbf{p} = \mathbf{x}^T \mathbf{p} - \|\mathbf{p}\|^2 = 0, \quad (12)$$

where $\mathbf{x} \in \mathbb{R}^n$. With this parametrization, the vector $\mathbf{p}$ corresponds to the point of S where it is closest to the origin, and as such $\mathbf{p}$ is also normal to the plane S . A different parametrization, more often used in the astrophysical literature, emphasizes one particular coordinate dimension, written as

$$x_1 = a_1 + \sum_{j=2}^n a_j x_j. \quad (13)$$

Converting between these two parametrizations is accomplished through

$$a_1 = \frac{\sum_{i=1}^n p_i^2}{p_1}, a_i = -\frac{p_i}{p_1}, i = 2, \dots, n \quad (14)$$

$$p_1 = \frac{a_1}{1 + \sum_{i=2}^n a_i^2}, p_i = -a_i p_0, i = 2, \dots, n. \quad (15)$$

We further assume that for each measurement $(\mathbf{m}, h)$, where now $\mathbf{m} \in \mathbb{R}^n$, the measurement error can be modelled with an n -dimensional normal distribution

$$h(\mathbf{x}; \mathbf{m}, \Sigma) = \frac{\exp \left[-\frac{1}{2} (\mathbf{x} - \mathbf{m})^T \Sigma^{-1} (\mathbf{x} - \mathbf{m}) \right]}{\sqrt{(2\pi)^n \det(\Sigma)}}, \quad (16)$$

where Σ is the measurement error covariance matrix. This assumption is necessary to obtain the analytic results below, but it is also well suited to published astrophysical data, for which the complete posterior distributions for each datum are typically not available. Upper limits can be incorporated by substituting one or more degrees of freedom x_i in equation (16) with suitable one-dimensional distributions representing the limit distributions.

To complete steps (iii) and (iv), we need to specify the intrinsic coordinates and the coordinate transformations. One method to construct the coordinates ϕ_f is to transform the natural coordinate frame of the measurement space M with a translation by $\mathbf{p}$ and a rotation $\mathbf{R} \in SO(n)$, which takes the n th basis vector $\hat{\mathbf{e}}_n$ to the unit vector $\hat{\mathbf{p}}$ (hereafter, we use a hat to signify a unit vector). The new components of a vector $\mathbf{x}$ are then $\phi_f(\mathbf{x}) = \mathbf{R}^T (\mathbf{x} - \mathbf{p}) = (\boldsymbol{\tau}, v) \in \mathbb{R}^{n-1} \times \mathbb{R}$. Defining $\mathbf{w} = (p_0, \dots, p_{n-1}, 0)$, $\cos \theta = \hat{\mathbf{p}}^T \hat{\mathbf{e}}_n$ and $\sin \theta = \hat{\mathbf{p}}^T \hat{\mathbf{w}}$, we can compute $\mathbf{R}$ with

$$\mathbf{R}(\mathbf{p}) = \mathbf{I} - \hat{\mathbf{w}} \hat{\mathbf{w}}^T - \hat{\mathbf{e}}_n \hat{\mathbf{e}}_n^T + (\hat{\mathbf{w}} \hat{\mathbf{e}}_n) \mathbf{R}_2(\theta) (\hat{\mathbf{w}} \hat{\mathbf{e}}_n)^T, \quad (17)$$

where

$$\mathbf{R}_2(\theta) = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \quad (18)$$

is the usual two-dimensional rotation matrix. The origin of the $\boldsymbol{\tau}$ coordinates is then at $\mathbf{p}$ (in chart ϕ_x), with the orientation of the

coordinate axes $\hat{\boldsymbol{\tau}}_i$ given by $\hat{\boldsymbol{\tau}}_i = \mathbf{R} \hat{\mathbf{e}}_i$, for $i = 1, \dots, n - 1$. The orthogonal intrinsic coordinate axis is $\hat{\mathbf{v}}_1 = \mathbf{R} \hat{\mathbf{e}}_n$.

Step (v) requires specification of the intrinsic scatter model p_{int} . We assume the normal intrinsic distribution (equation 6). With this definition, combined with the assumed measurement error distribution, and assuming no censoring, that is no upper or lower limits, we can complete step (vi). We evaluate the equation (10) for the likelihood of a single measurement data point. The result is

$$\mathcal{L}(\mathbf{p}, \sigma | \mathbf{m}, \Sigma) = \frac{1}{\sqrt{2\pi(\tilde{\Sigma}_{nn} + \sigma^2)}} \exp \left[\frac{-\tilde{m}_n^2}{2(\tilde{\Sigma}_{nn} + \sigma^2)} \right], \quad (19)$$

where

$$\tilde{\mathbf{m}} = \mathbf{R}^T (\mathbf{m} - \mathbf{p}) \quad (20)$$

$$\tilde{\Sigma} = \mathbf{R}^T \Sigma \mathbf{R}. \quad (21)$$

We can also write $\tilde{m}_n = \hat{\mathbf{p}}^T \mathbf{m} - \|\mathbf{p}\|$ and $\tilde{\Sigma}_{nn} = \hat{\mathbf{p}}^T \Sigma \hat{\mathbf{p}}$, which transforms equation (19) to the form used in Robotham & Obreschkow (2015).

We can also address the case where the measurements are given in coordinates where the metric is represented by a *constant* matrix $\mathbf{G} \neq \mathbf{I}$. This can be the case when the metric $\mathbf{G}$ encodes the choice of units for the variables, or when it has off-diagonal terms. A situation that would lead to a non-diagonal metric is, for example, a case where we are interested in two observables A_1 and A_2 , but can only measure $B_1 = A_1$ and $B_2 = A_1 + A_2$. The observables may be e.g. inflows of current or liquid from two independent sources, but we can only measure one source directly, and the sum of the flows somewhere downstream. The flows A_i can by assumption have any values, so the metric in coordinates A_i should be a diagonal product metric, whereas in coordinates B_i off-diagonal terms appear. For the case of a non-identity constant metric $\mathbf{G}$, the equation (19) is still valid, but we have

$$\tilde{\mathbf{m}} = \mathbf{R}^T \mathbf{W} \mathbf{P}^T (\mathbf{m} - \mathbf{p}) \quad (22)$$

$$\tilde{\Sigma} = \mathbf{R}^T \mathbf{W} \mathbf{P}^T \Sigma \mathbf{P} \mathbf{W}, \quad (23)$$

where $\mathbf{W}$ is a diagonal matrix of the square roots of the eigenvalues of $\mathbf{G}$, and $\mathbf{P} \in O(n)$, so that $\mathbf{W}^{-1} \mathbf{P}^T \mathbf{G} \mathbf{P} \mathbf{W}^{-1} = \mathbf{I}$. This is possible, since the metric is assumed to be Riemannian, in which case all the eigenvalues of $\mathbf{G}$ must be positive. Note that if the metric is not constant, the shortcut presented above does *not* work, and all the steps (i)–(vii) have to be followed using the metric explicitly.

3.2 Linear $n - k$ -dimensional case

The approach in the previous section can be readily generalized to the case of $k > 1$ simultaneous linear relations, defining an $n - k$ -dimensional (codimension k) affine subspace S_{n-k} . In this case, it is easiest to define the relation through defining the intrinsic coordinates first. As such, we go through steps (iii) and (iv) first. This can be accomplished by starting with the $n - 1$ -dimensional subspace S_{n-1} defined by the first relation, through the n parameters $\mathbf{p}_n \in \mathbb{R}^n$, as above. Now, the subspace S_{n-k} must lie in the intersection of S_{n-1} and an $n - 2$ -dimensional subspace S_{n-2} , which itself must also lie entirely in S_{n-1} . We may parametrize S_{n-2} within S_{n-1} with the $n - 1$ parameters $\mathbf{p}_{n-1} = (p_{n-1,1}, \dots, p_{n-1,n-1}, 0) \in \mathbb{R}^n$. This process is then continued until we have given p_{n-k+1} , parameterizing S_{n-k} in S_{n-k+1} . The number of parameters is

$(2nk - k^2 + k)/2$ in total. Each parameter vector $\mathbf{p}_i$ yields a rotation matrix $\mathbf{R}_i$ as in the section above, which together with $\mathbf{p}_i$ defines the coordinate transformation from the intrinsic coordinates ϕ_f on S_i to the intrinsic coordinates on S_{i-1} , with the understanding that $S_n = M$. Following this process through, we find that the coordinate transformation from ϕ_x to the intrinsic coordinates on S_{n-k} is given by

$$\begin{aligned} (\phi_f \circ \phi_x^{-1})(\mathbf{x}) &= \mathbf{R}_{n-k+1}^T \cdots \mathbf{R}_n^T (\mathbf{x} - \mathbf{p}_n) \\ &\quad - \mathbf{R}_{n-k+1}^T \cdots \mathbf{R}_{n-1}^T \mathbf{p}_{n-1} - \cdots \\ &\quad - \mathbf{R}_{n-k+1}^T \mathbf{p}_{n-k+1} \\ &= \mathbf{R}^T (\mathbf{x} - \mathbf{p}) \\ &= (\boldsymbol{\tau}, \mathbf{v}) \in \mathbb{R}^{n-k} \times \mathbb{R}^k, \end{aligned} \quad (24)$$

where

$$\mathbf{R}^T = \mathbf{R}_{n-k+1}^T \cdots \mathbf{R}_n^T \quad (25)$$

and

$$\mathbf{p} = \mathbf{p}_n + \mathbf{R}_n \mathbf{p}_{n-1} + \cdots + \mathbf{R}_n \cdots \mathbf{R}_{n-k+2} \mathbf{p}_{n-k+1}. \quad (26)$$

Again, the origin of the intrinsic coordinate axes is at $\mathbf{p}$. The orientation is similarly given by $\hat{\mathbf{t}}_i = \mathbf{R}\hat{\mathbf{e}}_i$ for $i = 1, \dots, n-k$ and $\hat{\mathbf{v}}_i = \mathbf{R}\hat{\mathbf{e}}_{n-k+i}$ for $i = 1, \dots, k$. This process defines the relation f in a roundabout way as the equation $\mathbf{v} = \mathbf{0}$, which in the chart ϕ_x is then given by the system of k equations

$$f(\mathbf{x}; \boldsymbol{\theta}) = [\mathbf{R}^T (\mathbf{x} - \mathbf{p})]_{n-k+1, \dots, n} = \mathbf{0}, \quad (27)$$

where $\boldsymbol{\theta} = (\mathbf{p}_n, \dots, \mathbf{p}_{n-k+1})$. This completes step (ii).

We again assume that measurement errors are normally distributed, given by equation (16). For step (v), we now assume that the intrinsic scatter is normally distributed in $\mathbf{v}$ with a covariance matrix $\boldsymbol{\Sigma}_{\text{int}}$, as in equation (5). With this definition, we can complete step (vi) and compute the likelihood given by a single measurement $\mathbf{m}$. The result is

$$\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\Sigma}_{\text{int}} | \mathbf{m}, \boldsymbol{\Sigma}) = \frac{\exp \left[-\frac{1}{2} \tilde{\mathbf{m}}_v^T (\tilde{\boldsymbol{\Sigma}}_{vv} + \boldsymbol{\Sigma}_{\text{int}})^{-1} \tilde{\mathbf{m}}_v \right]}{\sqrt{(2\pi)^k \det(\tilde{\boldsymbol{\Sigma}}_{vv} + \boldsymbol{\Sigma}_{\text{int}})}}, \quad (28)$$

where

$$\tilde{\mathbf{m}} = \mathbf{R}^T (\mathbf{x} - \mathbf{p}) = (\tilde{\mathbf{m}}_\tau, \tilde{\mathbf{m}}_v) \quad (29)$$

$$\tilde{\boldsymbol{\Sigma}} = \mathbf{R}^T \boldsymbol{\Sigma} \mathbf{R} = \begin{pmatrix} \tilde{\boldsymbol{\Sigma}}_{\tau\tau} & \tilde{\boldsymbol{\Sigma}}_{\tau v} \\ \tilde{\boldsymbol{\Sigma}}_{v\tau} & \tilde{\boldsymbol{\Sigma}}_{vv} \end{pmatrix}. \quad (30)$$

A possible constant non-Euclidean metric can be accommodated as in the previous section.

The likelihood, equation (28) is readily generalized for intrinsic distributions other than the multivariate normal distribution through mixture models (see e.g. the approach in Kelly 2007). In addition, the result in this section is useful for hypothesis testing in the sense of finding the most likely value of k for a given n -dimensional data set.

3.3 The line in two dimensions

It is useful to work out the two-dimensional special case of a line with intrinsic scatter in detail, considering the amount of literature focusing on symmetric fitting of linear relations with and without intrinsic scatter (e.g. Pearson 1901; Boggs et al. 1987; Isobe et al. 1990; Feigelson & Babu 1992; Robotham & Obreschkow 2015, and many others). As such, step (i) consists of

setting $M = \mathbb{R}^2$ with an Euclidean metric. In step (ii), for ease of comparison with existing methods, instead of the parametrization used above, we define the relation f and the line S with

$$f(x, y; \boldsymbol{\theta}) = y - \beta x - \alpha = 0, \quad (31)$$

where now $\boldsymbol{\theta} = (\alpha, \beta) \in \mathbb{R}^2$. The measurement error in the case of no censoring can now be represented with a two-dimensional normal distribution, given by

$$\begin{aligned} h(x, y; x_0, y_0, \sigma_x, \sigma_y, \rho) \\ = \frac{\exp \left\{ -\frac{1}{2(1-\rho^2)} \left[\frac{(x-x_0)^2}{\sigma_x^2} + \frac{(y-y_0)^2}{\sigma_y^2} - \frac{2\rho(x-x_0)(y-y_0)}{\sigma_x \sigma_y} \right] \right\}}{2\pi\sigma_x\sigma_y\sqrt{1-\rho^2}}, \end{aligned} \quad (32)$$

where x_0 and y_0 specify the measured values, σ_x, σ_y represent the measurement uncertainties in the x - and y -directions, and $\rho \in [-1, 1]$ specifies the correlation between the measurement errors. Upper and lower limits can be introduced as in Section 3.1. While the parametrization through α and β is convenient and intuitive, it is not manifestly symmetric with respect to the coordinates. This will be investigated further below.

Keeping the assumption that the internal scatter is normally distributed, via equation (6), we can use the results for steps (iii)–(vi) from Section 3.1. The likelihood of a single measurement in the case of no censoring is then obtained from equation (19), yielding

$$\mathcal{L}(\alpha, \beta, \sigma_x, \sigma_y, \rho) = \frac{1}{\sqrt{2\pi\tilde{\sigma}^2}} \exp \left(-\frac{v^2}{2\tilde{\sigma}^2} \right), \quad (33)$$

where now

$$v^2 = \frac{(y_0 - \alpha - \beta x_0)^2}{1 + \beta^2} \quad (34)$$

$$\tilde{\sigma}^2 = \frac{\beta^2 \sigma_x^2 + \sigma_y^2 - 2\beta\rho\sigma_x\sigma_y}{1 + \beta^2} + \sigma^2, \quad (35)$$

so that v^2 is the squared orthogonal distance from the relation, and $\tilde{\sigma}$ is an extended uncertainty incorporating both intrinsic scatter and the measurement errors. Analytic results for particular censored error distributions can also be found (see Appendix B).

We note that equation (35) has the correct asymptotic behaviour with respect to β , in the sense that when the regression line tends towards the vertical, or $\beta \rightarrow \infty$, we have $\tilde{\sigma} \rightarrow \sigma_x^2 + \sigma^2$, agreeing with intuitive result that all of the uncertainty should in this case be a combination of the horizontal (along x -axis) and intrinsic scatter. Likewise, when the regression line tends towards the horizontal, or $\beta \rightarrow 0$, we have $\tilde{\sigma} \rightarrow \sigma_y^2 + \sigma^2$, similarly agreeing with geometric intuition.

We now apply the discussion in Section 2.4 and derive a strictly unique uninformative prior probability density for the parameters $\boldsymbol{\theta} = (\alpha, \beta)$. This is appropriate for the case where no data has yet been collected, and we have no information on where on the plane the relation is and in what orientation. Thus, we need to find a prior density for $\boldsymbol{\theta}$ that yields a uniform distribution on M when integrated over. That is, the lines must cover the plane evenly. A direct approach seems to necessitate set-based analysis, but in this special case, we can apply existing results from the literature, derived via other means.

For example, demanding invariance under Euclidean coordinate transformations (simultaneous rotation and translation) yields an invariant prior probability measure $\pi(\alpha, \theta) d\alpha d\theta = d\alpha d\sin\theta$, where $\theta = \arctan\beta$ is the angle between the line and the x -axis. This is equivalent to $\pi(\alpha, \beta) = (1 + \beta^2)^{-3/2}$. The same result can be

obtained by demanding that the prior density be invariant under a switch of the coordinates, that is under $x \leftrightarrow y$, $\sigma_x \leftrightarrow \sigma_y$, $\beta \leftrightarrow 1/\beta$ and $\alpha \leftrightarrow -\alpha/\beta$. This result was apparently originally derived by E. T. Jaynes in 1976 (reprinted in Jaynes 1983).

Finally, for σ there is considerably more leeway in the literature with regards to the choice of an uninformative prior. A typical choice is the Jeffreys prior, $\pi(\sigma) \propto 1/\sigma$. The complete uninformative prior for this special case is then

$$\pi_{\phi, \theta}(\alpha, \beta, \sigma) = \frac{1}{\sigma(1 + \beta^2)^{3/2}}. \quad (36)$$

Equations (33) and (36) can now be combined to yield the un-normalized posterior distribution

$$p(\alpha, \beta, \sigma | \mathcal{D}) = \frac{1}{\sigma(1 + \beta^2)^{3/2}} \prod_{i=1}^{n_d} \mathcal{L}(\alpha, \beta, \sigma | x_i, y_i, \sigma_{x,i}, \sigma_{y,i}, \rho_i). \quad (37)$$

This gives a Bayesian solution to the problem of symmetric fitting of a linear relation with intrinsic scatter to two-dimensional data with heteroscedastic errors in both measured variables, including possible upper or lower limits and truncation. This result thus extends the earlier Bayesian results in Zellner (1971), Gull (1989), Jaynes (1991, unpublished) and Kelly (2007). It should be noted that while the likelihood, equation (33) is invariant under the switch of coordinates, the posterior distribution in itself is not, if the Jacobian of the transformation of parameters is not included, as discussed in Section 2.4. In practice, this means that in the usual case of $f = y - \alpha - \beta x$ we wish to find the maximum of

$$p(\alpha, \beta, \sigma | x, y, \sigma_x, \sigma_y) = p(\alpha, \beta, \sigma) \mathcal{L}(\alpha, \beta, \sigma | x, y, \sigma_x, \sigma_y). \quad (38)$$

In the inverse case, where $f = x - \alpha' - \beta'y$, and $\alpha' = -\alpha/\beta$, $\beta' = 1/\beta$, we should maximize

$$\begin{aligned} |\beta'^3| p(\alpha', \beta', \sigma) \mathcal{L}(\alpha', \beta', \sigma | y, x, \sigma_y, \sigma_x) \\ = p(\alpha, \beta, \sigma) \mathcal{L}(\alpha, \beta, \sigma | x, y, \sigma_x, \sigma_y), \end{aligned} \quad (39)$$

where β'^3 is the Jacobian of the transformation $(\alpha, \beta) \mapsto (\alpha', \beta')$. This works since for the priors we have $p(\alpha', \beta', \sigma) = |\beta'^3| p(\alpha, \beta, \sigma)$ and for the likelihood we have

$$\mathcal{L}(\alpha', \beta', \sigma | y, x, \sigma_y, \sigma_x) = \mathcal{L}(\alpha, \beta, \sigma | x, y, \sigma_x, \sigma_y) \quad (40)$$

as found above.

The result in this section can be easily extended to cases where the intrinsic scatter is not normally distributed or the measurement errors are not distributed with a bivariate normal distribution, by approximating the distributions with a weighted sum of Gaussians, as used e.g. in the LINMIX_ERR-method of Kelly (2007). However, the result is not applicable for situations more general than the linear case considered here, such as when the M itself is not trivially Euclidean, but includes e.g. angular or directional measurements or when the relation f is non-linear. In these cases, the likelihood function and consequently the posterior distribution may have to be evaluated numerically.

3.3.1 Comparison to some existing approaches

The equations (33)–(35) are a fundamentally symmetric way to describe the likelihood of parameters (α, β, σ) specifying a line with orthogonal intrinsic scatter. Several earlier works have incorporated intrinsic scatter as an error parameter, added in quadrature to the measurement errors along some specific measurement axis, as in e.g. the least likelihood method in Gültekin et al. (2009) and the FITEXY

method (Press et al. 1992) as modified in Tremaine et al. (2002). These methods introduce intrinsic scatter into the measurement errors of the dependent variable (for now taken to be y), yielding a total variance of the form

$$\bar{\sigma}^2 = \sigma_y^2 + \beta^2 \sigma_x^2 + \sigma_{\text{int},y}^2, \quad (41)$$

in the case of normally distributed intrinsic scatter in the y -direction, where $\sigma_{\text{int},y}^2$ is the variance of the intrinsic scatter and $\rho = 0$ is assumed. This is not equivalent to equation (35), and in particular, the missing normalization factor $1 + \beta^2$ leads to $\bar{\sigma} \rightarrow \infty$ as $\beta \rightarrow \infty$. We will now discuss how this difference arises.

First, assume that the underlying distribution of the data is a normal distribution orthogonal to a line $y = \alpha + \beta x$, i.e. given by equation (6) with $v = (y - \alpha - \beta x)/\sqrt{1 + \beta^2}$. In this case, the conditional distributions of the intrinsic scatter in the x - and y -directions are also normal, with variances

$$\sigma_{\text{int},x}^2 = \frac{1 + \beta^2}{\beta^2} \sigma^2, \quad (42)$$

$$\sigma_{\text{int},y}^2 = (1 + \beta^2) \sigma^2. \quad (43)$$

However, the marginal distributions of x and y are *not* normal. Now, the result in equation (41) can be obtained from the likelihood produced by a single measurement,

$$\begin{aligned} \mathcal{L}(\alpha, \beta | x_0, y_0, \sigma_x, \sigma_y, \sigma) \\ = \iint_{\mathbb{R}^2} h(x, y; x_0, y_0, \sigma_x, \sigma_y, \rho) p_{\text{int}}(v; \sigma) dx dy, \end{aligned} \quad (44)$$

where h is the two-dimensional Gaussian given by equation (32) and p_{int} is the normal distribution (equation 6). The assumptions yielding equation (41) are to take $v = y - \alpha - \beta x$, i.e. distance from the regression line *in the y -direction*, and to set $\sigma = \sigma_{\text{int},y}$. This yields a normal distribution, as in equation (33), but with

$$v^2 = v_a^2 = (y_0 - \alpha - \beta x_0)^2 \quad (45)$$

$$\bar{\sigma}^2 = \bar{\sigma}_a^2 = \beta^2 \sigma_x^2 + \sigma_y^2 - 2\beta\rho\sigma_x\sigma_y + \sigma_{\text{int},y}^2, \quad (46)$$

where the subscript a refers to asymmetric, which is a point we will discuss below. If we set $\rho = 0$, we have the likelihood used in Gültekin et al. (2009),

$$\mathcal{L} = \frac{1}{\sqrt{2\pi\bar{\sigma}^2}} \exp(-\chi^2) = \frac{\exp\left[-\frac{(y_0 - \alpha - \beta x_0)^2}{2(\sigma_y^2 + \beta^2\sigma_x^2 + \sigma_{\text{int},y}^2)}\right]}{\sqrt{2\pi(\sigma_y^2 + \beta^2\sigma_x^2 + \sigma_{\text{int},y}^2)}}, \quad (47)$$

as well as the χ^2 value used in Tremaine et al. (2002). Likewise, if we instead use $v = (y - \alpha - \beta x)/\sqrt{1 + \beta^2}$, i.e. orthogonal distance from the regression line, in the equation (44), the result is the set of equations (33)–(35). However, at this point we find that making the obvious substitution from equation (43) does *not* make the set of equations (33)–(35) equivalent to equation (33) combined with the substitutions (45) and (46). Furthermore, the χ^2 in equation (47) is invariant under the switch of independent and dependent coordinates, together with $\sigma_{\text{int},y} \leftrightarrow \sigma_{\text{int},x}$. However, the variance $\bar{\sigma}_a^2$ and consequently the likelihood are manifestly not invariant. The equation (47) will give smaller likelihoods than equation (33) for $\beta \gg 1$ or $\beta \sim 0$, depending on which variable is taken as the dependent one.

The fundamental reason for this state of affairs is the fact that if the intrinsic scatter is modelled as a proper probability distribution in any particular coordinate direction, the normalization of

the intrinsic scatter as a distribution orthogonal to the relation will change in normalization with change in β . This reflects the breaking of the rotational symmetry of the problem. For example, assume that the intrinsic scatter is modelled as a normal distribution in the y -direction, with variance $\sigma_{\text{int},y}^2$. In this case, the distribution of the data, being the intrinsic scatter in a direction orthogonal to the regression line, will need to have a normalization constant that goes down with increasing slope β . Indeed, as $\beta \rightarrow \infty$, the normalization constant will need to tend towards zero to keep the conditional probability distribution in the y -direction normalized to unity. This leads to a vanishing likelihood for a vertical line $\beta \rightarrow \infty$, which results from the fact that in this case $\tilde{\sigma}_a \rightarrow \infty$ and consequently the resulting likelihood $\mathcal{L}(\alpha, \beta \rightarrow \infty | \mathcal{D}) \rightarrow 0$. In mathematical terms, if we have a distribution of the data depending only on the orthogonal distance, $p_{\text{int}}(v)$, and we have demanded that in y -direction we should have a proper probability density function $p_{\text{int},y}(y)$, that is

$$\int_{-\infty}^{\infty} p_{\text{int},y}(y) dy = \int_{-\infty}^{\infty} p_{\text{int}}(v(x(y), y)) dy = 1, \quad (48)$$

then necessarily

$$\int_{-\infty}^{\infty} p_{\text{int}}(v) dv = \int_{-\infty}^{\infty} p_{\text{int}}(v(x(y), y)) \frac{dv}{dy} dy = \frac{1}{\sqrt{1 + \beta^2}}, \quad (49)$$

which goes to zero as $\beta \rightarrow \infty$. This is inconsistent, if we expect the distribution of the data, p_{int} , to have the same normalization no matter which way the regression line might point. Armed with this knowledge, if we now multiply the right side of equation (47) by $\sqrt{1 + \beta^2}$ and substitute $\sigma_{\text{int},y}^2$ from equation (43), we *do* arrive at the equations (33)–(35). Based on this analysis, the likelihood of equation (33) is recommended for the purpose of fitting a regression line with normally distributed intrinsic scatter to data with normally distributed measurement errors. This point is also raised in Robotham & Obreschkow (2015).

To investigate this conclusion numerically, we generated sets of simulated observations from linear relations with intrinsic scatter. The data were then fit using LINMIX_ERR, FITEXY,³ and by maximizing the posterior probability (equation 37). This last method is referred to as Geo-MAP (geometric maximum a posteriori) hereafter. The parameter estimates and 1σ uncertainties for the LINMIX_ERR method were taken using the median and standard deviation of the Markov chain outputs. For the FITEXY method, the estimate and uncertainty provided by the algorithm were used. For the Geo-MAP method, a bootstrapping method was used together with a numerical maximization of the posterior distribution to yield the parameter estimates and 1σ uncertainties. The simulated data sets were generated assuming a relation $y = \alpha + \beta x$, with a normally distributed orthogonal intrinsic scatter $\sigma = 0.1$ and normally distributed measurement error with $\sigma_x = \sigma_y = 0.1$ and $\rho = 0$. In total, six data sets were generated, with $n_d = 100$ data points each and with $\alpha \in \{0, 10\}$ and $\beta \in \{0.1, 1, 10\}$. The data sets represent a situation where the intrinsic scatter and measurement errors contribute equally to the observed scatter, and as such approximations assuming the relative smallness of either are maximally violated. The data were generated with a uniform distribution $U(-1, 1)$ along the relation, and then shifted first according to σ and then according to σ_x and σ_y . The data was then fit both in the forward direction, with $y = \alpha + \beta x$ as well as the inverse direction, with $x = \alpha' + \beta' y$. For the inverse direction,

Table 1. Slopes of linear relations fit with the MAP estimate of equation (37) and the LINMIX_ERR and FITEXY methods.

Slope of data method	Slope of fit $\pm 1\sigma$ uncertainties		
	0.1	1	10
$\alpha = 0$			
LINMIX_ERR fwd	0.09 ± 0.02	0.99 ± 0.04	1.67 ± 0.70
LINMIX_ERR inv	0.42 ± 0.12	1.08 ± 0.05	15.4 ± 5.99
FITEXY fwd	0.09 ± 0.02	0.99 ± 0.04	1.68 ± 0.39
FITEXY inv	0.41 ± 0.06	1.08 ± 0.04	15.5 ± 6.10
Geo-MAP fwd	0.10 ± 0.03	1.04 ± 0.05	11.8 ± 2.80
Geo-MAP inv	0.10 ± 0.03	1.04 ± 0.04	11.8 ± 2.79
$\alpha = 10$			
LINMIX_ERR fwd	0.04 ± 0.02	0.93 ± 0.04	1.79 ± 0.57
LINMIX_ERR inv	0.51 ± 0.27	1.01 ± 0.04	10.9 ± 3.31
FITEXY fwd	0.04 ± 0.02	0.93 ± 0.04	1.80 ± 0.35
FITEXY inv	0.52 ± 0.13	1.01 ± 0.04	10.9 ± 3.24
Geo-MAP fwd	0.04 ± 0.02	0.98 ± 0.04	8.66 ± 1.68
Geo-MAP inv	0.04 ± 0.02	0.98 ± 0.04	8.66 ± 1.68

Figure 4. Fits to simulated measurements with measurement errors $\sigma_x = \sigma_y = 0.1$, sampled from a linear relation $y = 10x$ with intrinsic scatter $\sigma = 0.1$. Solid lines show the fitted relations, and dashed lines indicate 1σ intrinsic scatter estimates.

the slope $\beta = 1/\beta'$ was computed after the fit. The slopes of the resulting fits are listed in Table 1. Fig. 4 illustrates the differences between the methods for the $\alpha = 0, \beta = 10$ case.

The results show that for values of $\beta \gg 1$, the inverse fits with the LINMIX_ERR and FITEXY methods agree better with the data, albeit with high uncertainties. Similarly, for $\beta \ll 1$, the forward fits give a better agreement with the data. This is in line with the observations above, showing that these methods give results tending towards lower values of β for $\beta \gg 1$ when used in the forward direction, and correspondingly towards higher values of β for $\beta \sim 0$, when used in the inverse direction. It should be noted that in addition to yielding estimates of β that are much too high or low in these cases, both LINMIX_ERR and FITEXY methods give corresponding uncertainty estimates that are small, so that the true value of β ends up as a multiple- σ outlier. Similarly as expected, the Geo-MAP method

³ We used the implementation available in the MPFIT package (Markwardt 2009) through the MPFITEXY wrapper routine (Williams, Bureau & Cappellari 2010).

gives nearly identical results for forward and inverse fits, up to the noise caused by the bootstrapping procedure. In addition, the true value of β is always contained within the 1σ bounds except for the case of $(\alpha, \beta) = (10, 0.1)$. Finally, we note that none of the methods appear invariant with respect to a shift with α .

We conclude that methods based on likelihood examining distances from the regression line in the direction of a particular coordinate may give estimates of slope that are far from the true value while simultaneously providing tight error bounds. In particular, if the distance is measured along the dependent variable, the resulting method will be biased towards $\beta = 0$. For methods, like the maximum-likelihood method in Gültekin et al. (2009), this effect remains even when the intrinsic scatter $\sigma \rightarrow 0$, in which case the distribution p_{int} becomes a delta-function ridge, and the problem reduces to ordinary least-squares (OLS) regression. This tendency towards lower slopes was reported in Park et al. (2012) for all the methods used in the paper, namely OLS, BCES (Akritas & Bershady 1996), FITEXY, LINMIX_ERR and the Gültekin et al. (2009) maximum-likelihood method. Our results also agree with Isobe et al. (1990), who studied different regression methods in the limit of vanishing measurement error.

3.4 Ellipses in Euclidean space

Next, we consider a case where the relation f is properly non-linear. With astrophysical applications in mind, we seek analytic results for ellipses embedded in a Euclidean space. A possible application could be, for example, to model the orbit of a stream of low-mass particles. As such, for step (i) we consider our measurements to be points in $M = \mathbb{R}^3$ with the Euclidean metric. For step (ii), we define the ellipse with

$$f(\mathbf{x}; \boldsymbol{\theta}) = \frac{\tilde{x}_1^2}{a^2} + \frac{\tilde{x}_2^2}{(1-e^2)a^2} - 1 = 0, \quad (50)$$

where $\boldsymbol{\theta} = (\mathbf{p}, a, e, i, \omega, \Omega)$. Here, the parameters $\boldsymbol{\theta}$ constitute $\mathbf{p} \in \mathbb{R}^3$, which specifies the position of the centre of the ellipse, and a, e, i, ω and Ω , which are the orbital elements, namely semimajor axis, eccentricity, inclination, argument of pericentre and longitude of the ascending node, respectively. The tilde signifies a coordinate transformation to a vector $\tilde{\mathbf{x}} = \mathbf{R}^T(\mathbf{x} - \mathbf{p})$, where the matrix $\mathbf{R}$ rotates the Euclidean coordinate triad $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \hat{\mathbf{e}}_3\}$ into a coordinate triad $\{\tilde{\mathbf{e}}_1, \tilde{\mathbf{e}}_2, \tilde{\mathbf{e}}_3\}$ aligned with the ellipse. That is, $\tilde{\mathbf{e}}_1$ points along the semimajor and $\tilde{\mathbf{e}}_2$ along the semiminor axis. The matrix $\mathbf{R}$ can be written explicitly as the product

$$\mathbf{R} = \begin{pmatrix} \cos \Omega & -\sin \Omega & 0 \\ \sin \Omega & \cos \Omega & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 & 1 \\ 0 & \cos i & -\sin i \\ 0 & \sin i & \cos i \end{pmatrix} \times \begin{pmatrix} \cos \omega & -\sin \omega & 0 \\ \sin \omega & \cos \omega & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (51)$$

In the ellipse-aligned frame, the ellipse can now be characterized also as the curve

$$\mathbf{g}(E) = a \cos E \tilde{\mathbf{e}}_1 + \sqrt{1-e^2} a \sin E \tilde{\mathbf{e}}_2, \quad (52)$$

where $E \in [0, 2\pi)$ is typically called the eccentric anomaly. To use this convenient form, we will work in the coordinate frame aligned with ellipse in the following.

We can now define the intrinsic coordinates [step (iii)] by defining τ to be the arc length of the ellipse, with $\tau = 0$ in the direction of

$\tilde{\mathbf{e}}_1$ and increasing towards $\tilde{\mathbf{e}}_2$. We will need the relation between the eccentric anomaly E and the arc length τ , given by

$$\tau(E) = \sqrt{1-e^2} a \mathcal{E}(E, -e^2(1-e^2)^{-1}), \quad (53)$$

where

$$\mathcal{E}(\phi, m) = \int_0^\phi \sqrt{1-m \sin^2 \theta} d\theta \quad (54)$$

is the incomplete elliptic integral of the second kind. The orthogonal coordinate axes can be defined so that $\hat{\mathbf{v}}_1$ is parallel to $\tilde{\mathbf{e}}_3$ and $\hat{\mathbf{v}}_2$ points orthogonally away from the ellipse, and thus can be given as

$$\begin{aligned} \hat{\mathbf{v}}_2 &\propto \frac{2\tilde{x}_1}{a^2} \tilde{\mathbf{e}}_1 + \frac{2\tilde{x}_2}{(1-e^2)a^2} \tilde{\mathbf{e}}_2 \\ &= \frac{2 \cos E}{a} \tilde{\mathbf{e}}_1 + \frac{2 \sin E}{\sqrt{1-e^2}a} \tilde{\mathbf{e}}_2, \end{aligned} \quad (55)$$

with normalization omitted for clarity.

For step (iv), we can now construct the coordinate transformation $\phi_x \circ \phi_f^{-1}$ to find

$$\begin{aligned} \tilde{\mathbf{x}}(\tau, v_1, v_2) \\ = \mathbf{g}(E(\tau)) + v_1 \tilde{\mathbf{e}}_3 + v_2 \left(\frac{2 \cos E}{a} \tilde{\mathbf{e}}_1 + \frac{2 \sin E}{\sqrt{1-e^2}a} \tilde{\mathbf{e}}_2 \right). \end{aligned} \quad (56)$$

However, the inverse transformation is already very difficult to find analytically, and further problems caused by the non-linearity will now also manifest.

At this point, step (v), we should specify the intrinsic scatter model. To mitigate analytic problems, we would again like to use the simplest normal distribution, so that

$$p_{\text{int}}(\tau, v_1, v_2; \sigma) = \frac{\exp \left[-\frac{1}{2\sigma^2} (v_1^2 + v_2^2) \right]}{2\pi\sigma^2}. \quad (57)$$

However, we find that even if we consider a similarly simplified measurement error distribution,

$$h(\tilde{\mathbf{x}}; \tilde{\mathbf{m}}, \sigma_m) = \frac{\exp \left[-\frac{1}{2\sigma_m^2} (\tilde{\mathbf{x}} - \tilde{\mathbf{m}})^T (\tilde{\mathbf{x}} - \tilde{\mathbf{m}}) \right]}{(2\pi\sigma_m^2)^{3/2}}, \quad (58)$$

it is very difficult to obtain a closed form result for the likelihood integral (equation 10). Neither is it obvious how to compute the form of p_{int} in the measurement coordinates ϕ_x , whether in the original or ellipse-aligned form.

To make progress, we have to make some assumptions. The main assumption we need is that the intrinsic scatter σ is small compared to the minimum radius of curvature of the ellipse, or $\sigma \ll (1-e^2)b$. If this is the case, we can approximate the intrinsic scatter distribution with a convolution of the ellipse and a normal distribution, which in the ellipse-aligned coordinate chart yields

$$\begin{aligned} p_{\text{int}}(\tilde{\mathbf{x}}) &\sim \int_{\mathbb{R}^3} \delta \left(\frac{\tilde{y}_1^2}{a^2} + \frac{\tilde{y}_2^2}{(1-e^2)a^2} - 1 \right) \\ &\times \frac{\exp \left[-\frac{1}{2\sigma^2} (\tilde{\mathbf{y}} - \tilde{\mathbf{x}})^T (\tilde{\mathbf{y}} - \tilde{\mathbf{x}}) \right]}{(2\pi\sigma^2)^{3/2}} d\tilde{\mathbf{y}}, \end{aligned} \quad (59)$$

where δ is the Dirac delta distribution. While this integral also resists evaluation in closed form, we can approximate it by a mixture of normal distributions, by distributing individual normal distributions on points equidistantly spaced in arc length. The number N of component distributions required to make the distribution smooth

depends on the intrinsic scatter σ . A Euclidean separation of less than σ is sufficient, and we can thus set

$$N = \left\lceil \frac{\tau(E = 2\pi)}{\sigma} \right\rceil, \quad (60)$$

where the brackets denote the ceiling function. We can now complete step (vi) and evaluate the likelihood for a single measurement, yielding

$$\begin{aligned} \mathcal{L}(\theta, \sigma | \tilde{\mathbf{m}}, \sigma_m) &= \int_{\mathbb{R}^3} \frac{\exp \left[-\frac{1}{2\sigma_m^2} (\tilde{\mathbf{x}} - \tilde{\mathbf{m}})^T (\tilde{\mathbf{x}} - \tilde{\mathbf{m}}) \right]}{(2\pi\sigma_m^2)^{3/2}} \\ &\times \frac{1}{N} \sum_{i=1}^N \frac{\exp \left[-\frac{1}{2\sigma^2} (\tilde{\mathbf{x}} - \mathbf{g}(E_i))^T (\tilde{\mathbf{x}} - \mathbf{g}(E_i)) \right]}{(2\pi\sigma^2)^{3/2}} d\tilde{\mathbf{x}} \\ &= \frac{1}{N} \sum_{i=1}^N \frac{\exp \left[-\frac{1}{2(\sigma_m^2 + \sigma^2)} (\mathbf{g}(E_i) - \tilde{\mathbf{m}})^T (\mathbf{g}(E_i) - \tilde{\mathbf{m}}) \right]}{(2\pi[\sigma_m^2 + \sigma^2])^{3/2}}, \end{aligned} \quad (61)$$

where the eccentric anomalies E_i need to be numerically solved from

$$\tau(E_i) = \tau(2\pi) \frac{i}{N}. \quad (62)$$

The equation (61) is then straightforward to evaluate numerically for each data point, completing step (vii).

3.5 Circles on a spherical surface

Finally, we consider an example of a relation defined in a non-Euclidean space. Specifically, we assume the measurements lie on $M = S^2$, the two-dimensional spherical shell. We use the standard spherical coordinate chart $(\theta, \phi) \in \mathbb{R}^2$, for which the metric is given by the matrix

$$\mathbf{G} = \begin{pmatrix} 1 & 0 \\ 0 & \sin^2 \theta \end{pmatrix}. \quad (63)$$

This choice is appropriate for positions on the plane of the sky, for example. This definition completes step (i). For step (ii), we choose a relation that defines a circle, by setting

$$f(\theta, \phi; \theta) = \tilde{\theta} - \alpha_c = 0, \quad (64)$$

where $\theta = (\theta_c, \phi_c, \alpha_c)$. Here, the parameters $(\theta_c, \phi_c) \in \mathbb{R}^2$ specify the centre of the circle, and α_c is equal to one-half of the angular size of the circle. The tilde denotes a coordinate transformation that takes the centre of the circle to the north pole, or $\tilde{\theta}_c = 0$. This transformation is most conveniently realized by first transforming the points (θ, ϕ) to Cartesian three vectors, using the rotation matrix constructed in Section 3.1, and then converting back to spherical coordinates. The analytic form of the transformation $(\theta, \phi) \mapsto (\tilde{\theta}, \tilde{\phi})$ can be given directly in terms of (θ, ϕ) and (θ_c, ϕ_c) , but the result is unwieldy and omitted.

For steps (iii) and (iv), we now define the intrinsic coordinates and the coordinate transformation between them and the tilde-transformed coordinates. A convenient choice is to set $\tilde{\phi} = \tau$, so that the τ coordinate increases with longitude as we go along the circle. We likewise set $\tilde{\theta} = \alpha_c + \nu$. This coordinate transformation is multivalued in the sense that a given point $(\tilde{\theta}, \tilde{\phi})$ corresponds to an infinite number of intrinsic coordinate pairs, which can be

written as

$$\begin{aligned} \phi_f \circ \phi_x^{-1}(\tilde{\theta}, \tilde{\phi}) &= \{(\tilde{\theta} - \alpha_c + n 2\pi, \tilde{\phi}) | n \in \mathbb{Z}\} \\ &\times \cup \{(\tilde{\theta} + \alpha_c + k 2\pi, (\tilde{\phi} + \pi)) | k \in \mathbb{Z}\}. \end{aligned} \quad (65)$$

The second set of intrinsic coordinate pairs above corresponds to geodesics originating from the point of the circle on the other side of the sphere, at $\tilde{\phi} + \pi$ (modulo 2π).

For the definition of the intrinsic scatter distribution in step (v), we again use the one-dimensional normal distribution (equation 6). At this point, it is advantageous to express the intrinsic scatter distribution in the original coordinate frame as well. Since the coordinate transformation is now multivalued, we must proceed as explained in Section 2.2 to get

$$\begin{aligned} p_{\text{int}}(\tilde{\theta}; \alpha_c, \sigma) &= \frac{1}{\sqrt{2\pi\sigma^2}} \sum_{n=-\infty}^{\infty} \left[e^{-\frac{(\tilde{\theta}-\alpha_c+n 2\pi)^2}{2\sigma^2}} + e^{-\frac{(\tilde{\theta}+\alpha_c+n 2\pi)^2}{2\sigma^2}} \right] \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} \left[e^{-\frac{(\tilde{\theta}-\alpha_c)^2}{2\sigma^2}} \vartheta \left(i \frac{\pi(\tilde{\theta} - \alpha_c)}{\sigma^2}, e^{-\frac{2\pi}{\sigma^2}} \right) \right. \\ &\quad \left. + e^{-\frac{(\tilde{\theta}+\alpha_c)^2}{2\sigma^2}} \vartheta \left(i \frac{\pi(\tilde{\theta} + \alpha_c)}{\sigma^2}, e^{-\frac{2\pi}{\sigma^2}} \right) \right], \end{aligned} \quad (66)$$

where ϑ is a Jacobi theta function corresponding to the definition

$$\vartheta(z, q) = \sum_{n=-\infty}^{\infty} q^{n^2} e^{2niz}, \quad (67)$$

with $q, z \in \mathbb{C}$, $|q| < 1$. Note that there is no dependence on $\tilde{\phi}$. From this form of the p_{int} distribution, it is easy to appreciate the close relation to the wrapped normal distribution defined on the circle, also defined through a Jacobi theta function. As such, this example shows how the wrapped distributions naturally arise from the intrinsic coordinate formalism.

To proceed further, we need to make some assumptions of the error distribution of the measurements, which we denote by pairs $(\tilde{\theta}_m, \tilde{\phi}_m)$. An attractive choice would be one of the spherical generalizations of the normal distribution, such as the Kent distribution or the Von Mises–Fisher distribution. However, these combined with the intrinsic scatter distribution (66) and the metric of S^2 do not seem to easily yield closed form results for the likelihood in step (vi). Instead, we have to assume that the measurement error is very small, and approximately given by a uniform distribution within a small cell $[\tilde{\theta}_m \pm \Delta_m/2, \tilde{\phi}_m \pm \Delta_m/(2 \sin \tilde{\theta}_m)] \subset S^2$, where Δ_m specifies the (small) magnitude of the error.

With the help of this admittedly severe assumption, we can proceed to step (vi) and compute the likelihood for a single measurement, finding

$$\begin{aligned} \mathcal{L}(\theta, \sigma | \tilde{\theta}_m, \tilde{\phi}_m, \Delta_m) &\sim p_{\text{int}}(\tilde{\theta}_m; \alpha_c, \sigma) + \frac{1}{12} \Delta_m^2 \cot(\tilde{\theta}_m) \frac{dp_{\text{int}}(\tilde{\theta}; \alpha_c, \sigma)}{d\tilde{\theta}} \Big|_{\tilde{\theta}=\tilde{\theta}_m} \end{aligned} \quad (68)$$

to second order in Δ_m . While the likelihood (68) is complicated by the presence of the theta functions and their derivatives, these can be efficiently evaluated numerically, and it is again straightforward to numerically compute the value of the likelihood for each data point to complete step (vii).

We can compare the result to some existing approaches, such as presented in Jupp & Kent (1987) and Fujiki & Akaho (2009). The former method is based on splines, and the latter is based on Euclidean approximation of small spherical distances and least-squares fitting. The method in Jupp & Kent (1987) can in principle

cope with arbitrarily large normally distributed measurement errors, but cannot incorporate intrinsic scatter, limits or truncation. In addition, since the result is always a spline, no parameter estimation for a predetermined curve is possible. The method in Fujiki & Akaho (2009) is suitable for parameter estimation, but otherwise shares the disadvantages of the Jupp & Kent (1987) method in addition to not incorporating measurement errors at all. In this sense, the geometric approach presented here compares favourably, although it presents mathematical difficulties if curves other than circles are to be fit.

3.6 What is learned from the examples

The examples clearly demonstrate that analytic results are only easy to obtain in the case where M is Euclidean and the relation to be fit is linear, or more exactly a geodesic. In the last two examples, simplifying assumptions are necessary to obtain a closed form for the likelihood. As such, while the linear Euclidean case of normally distributed orthogonal scatter is essentially solved here completely, much future work is required in the context on non-linear relations and non-Euclidean spaces.

In addition, the choice of parametrization of the relation to be fit has some important consequences. First, the chosen parametrization directly affects the form of the prior probability distribution, as discussed in Section 2.2. In addition, different parametrizations may not be numerically well-behaved everywhere. For example, the p -parametrization used in Sections 3.1 and 3.2 will likely exhibit numerical instability if any of the relation hyperplanes passes close to the coordinate origin, in which case $\|p\| \rightarrow 0$ and the orientation of the plane becomes indeterminate. The (α, β) -parametrization in Section 3.3 suffers from similar difficulties for nearly vertical lines, as then the values $\beta \rightarrow \pm\infty$ become degenerate.

4 THE $M_{\text{BH}}-\sigma$ RELATION

As a typical application, the methodology described above can be applied to the $M_{\text{BH}}-\sigma$ relation, typically written in the form

$$\log_{10} \left(\frac{M_{\text{BH}}}{M_{\odot}} \right) = \alpha + \beta \log_{10} \left(\frac{\sigma}{200 \text{ km s}^{-1}} \right). \quad (69)$$

The $M_{\text{BH}}-\sigma$ relation is an important correlation between the mass M_{BH} of the supermassive black hole in the centre of a galaxy and the velocity dispersion σ of the galactic bulge.⁴ Since the initial discovery of this correlation (Ferrarese & Merritt 2000; Gebhardt et al. 2000; but see also Magorrian et al. 1998), it has been re-established several times, using both larger data sets and different statistical procedures (e.g. Tremaine et al. 2002; Novak, Faber & Dekel 2006; Gültekin et al. 2009; Graham et al. 2011; McConnell et al. 2011; Beifiori et al. 2012; McConnell & Ma 2013; van den Bosch 2016; Saglia et al. 2016). The existence of the $M_{\text{BH}}-\sigma$ relation and analogous relations, such as the correlation $M_{\text{BH}}-M_{\text{bulge}}$ with the galaxy bulge mass, have also been confirmed in numerical simulations, both in galaxy merger simulations (e.g. Di Matteo, Springel & Hernquist 2005; Johansson, Naab & Burkert 2009a; Johansson, Burkert & Naab 2009b; Choi et al. 2014) and in cosmological simulations (e.g. Sijacki et al. 2007; Di Matteo et al. 2008; Booth & Schaye 2009; Sijacki et al. 2015). However, all of these studies have used statistical methods that treat one observable as independent and the other observable as dependent. Redoing the fit with the

independent observable as the dependent and vice versa produces a significantly different correlation slope, and in some cases also affects the estimated value of the intrinsic scatter (Park et al. 2012). Consequently, there is some controversy in the astronomical literature as to whether it is more suitable to fit M_{BH} as a function of σ (forward regression, using the definition in Park et al. 2012), or the other way around (inverse regression, respectively) in the presence of intrinsic scatter. For a review of the debate, see Graham (2016) and the references therein. That different slopes are produced by switching dependent and independent variables has been known for a long time (e.g. Pearson 1901) and amounts to asking two different questions: what is the most likely value for M_{BH} , given σ or vice versa. However, it is equally well known that if one is interested in the functional relation between the observables, then a symmetric method should be used (see e.g. Isobe et al. 1990).

The approach in Section 3.3 presents a symmetric Bayesian solution to fitting a linear relation between two observables, incorporating heteroscedastic errors, upper limits and intrinsic scatter. As such, we would expect it to yield a non-biased estimate of the intrinsic slope β of the relation (69), regardless of which way the relation is fit. To investigate this, the three data sets used in table 1 of Park et al. (2012) were analysed. The data are originally from Gültekin et al. (2009), McConnell et al. (2011) and Graham et al. (2011). In order to compare with the results in Park et al. (2012), the measurement errors were modelled as uncorrelated bivariate Gaussians in the logarithmic space, leaving out all upper limits. Following Park et al. (2012), the standard deviations were set equal to mean errors, i.e. $\sigma_{\log_{10} M_{\text{BH}}} = (\log_{10} M_{\text{BH,high}} - \log_{10} M_{\text{BH,low}})/2$ and similarly for the velocity dispersions. In addition, for the Graham et al. (2011) sample, which lacks velocity dispersion error data, the velocity dispersion errors were set to 10 per cent and then propagated to averaged logarithmic errors. Finally, a data set compiled from a multitude of sources used in van den Bosch (2016) was used. For this data set, the errors were used as given, and upper limits were also incorporated in the fit.

For each data set, the values for the same set of parameters as in Park et al. (2012) were computed: the intercept α , slope β and intrinsic scatter along the M_{BH} -axis $\sigma_{\text{int}, M_{\text{BH}}}$. The results were computed for forward regression (equation 69), as well as for the inverse regression, converting back to equivalent forward values in the end. The parameter values and uncertainties were estimated using the posterior distribution (equation 37) in two complementary ways.

The first set of estimates (Geo-MAP, hereafter) were derived by numerically maximizing the posterior distribution combined with a bootstrap resampling procedure. For each data set, a total of $n_d(\log n_d)^2$ bootstrap samples were constructed (Babu & Singh 1983; Feigelson & Babu 2012), where n_d is the number of data points. The parameters were then estimated using the median parameter values. Parameter uncertainties at 1σ level were estimated using median absolute deviations scaled to correspond to standard deviations for a normal distribution.

For the second set of estimates (Geo-MCMC, hereafter), the posterior distribution was sampled with the MCMC sampler `emcee` (Foreman-Mackey et al. 2013). Convergence during sampling was monitored with the potential scale reduction factor R (Gelman et al. 2013), and sampling was continued until $R < 1.01$ was achieved. The parameter estimates were obtained as the values corresponding to the sample with maximum posterior probability. The parameter uncertainties at 1σ level were then estimated by constructing credible regions containing 0.6827 of the posterior probability mass. This was done by starting from the maximum posterior

⁴ To avoid confusion, the intrinsic scatter will be denoted with σ_{int} in this section.

Table 2. $M_{\text{BH}}-\sigma$ relations $\log_{10}(M_{\text{BH}}/M_{\odot}) = \alpha + \beta \log_{10}(\sigma/200 \text{ km s}^{-1})$ derived using the data sets in Gültekin et al. (2009), McConnell et al. (2011), Graham et al. (2011) and van den Bosch (2016), using methods FITEXY, LINMIX_ERR and the Geometric method of this paper. For the Geometric method, results with both bootstrapped (Geo-MAP) and MCMC-sampled (Geo-MCMC) MAP estimate are shown. The results for the FITEXY and LINMIX_ERR methods for the first three data sets are from Park et al. (2012).

Method	Intercept α	Slope β	$\sigma_{\text{int}, M_{\text{BH}}}^a$	σ_{int}^a
Gültekin et al. (2009) data				
FITEXY fwd	8.19 ± 0.06	4.06 ± 0.32	0.39 ± 0.06	0.093
FITEXY inv	8.21 ± 0.07	5.35 ± 0.66	0.45 ± 0.09	0.083
LINMIX_ERR fwd	8.19 ± 0.07	4.04 ± 0.40	0.42 ± 0.05	0.101
LINMIX_ERR inv	8.21 ± 0.08	5.44 ± 0.56	0.49 ± 0.09	0.089
Geo-MAP fwd	8.22 ± 0.07	5.42 ± 0.72	0.43 ± 0.08	0.076 ± 0.008
Geo-MCMC fwd	$8.21^{+0.14}_{-0.14}$	$5.41^{+1.10}_{-0.79}$	$0.44^{+0.12}_{-0.09}$	$0.080^{+0.022}_{-0.016}$
McConnell et al. (2011) data				
FITEXY fwd	8.28 ± 0.06	5.07 ± 0.36	0.43 ± 0.05	0.083
FITEXY inv	8.32 ± 0.06	6.29 ± 0.49	0.47 ± 0.06	0.073
LINMIX_ERR fwd	8.27 ± 0.06	5.06 ± 0.36	0.44 ± 0.05	0.085
LINMIX_ERR inv	8.32 ± 0.07	6.31 ± 0.46	0.49 ± 0.07	0.077
Geo-MAP fwd	8.32 ± 0.05	6.26 ± 0.51	0.45 ± 0.06	0.070 ± 0.008
Geo-MCMC fwd	$8.32^{+0.13}_{-0.12}$	$6.31^{+0.90}_{-0.71}$	$0.46^{+0.11}_{-0.08}$	$0.071^{+0.017}_{-0.013}$
Graham et al. (2011) data				
FITEXY fwd	8.15 ± 0.05	5.08 ± 0.34	0.31 ± 0.04	0.060
FITEXY inv	8.16 ± 0.05	5.84 ± 0.42	0.33 ± 0.05	0.056
LINMIX_ERR fwd	8.15 ± 0.05	5.08 ± 0.36	0.31 ± 0.05	0.060
LINMIX_ERR inv	8.17 ± 0.06	5.85 ± 0.42	0.34 ± 0.06	0.057
Geo-MAP fwd	8.17 ± 0.04	6.00 ± 0.41	0.31 ± 0.05	0.051 ± 0.006
Geo-MCMC fwd	$8.17^{+0.11}_{-0.11}$	$5.98^{+0.87}_{-0.67}$	$0.32^{+0.11}_{-0.10}$	$0.052^{+0.018}_{-0.016}$
van den Bosch (2016) data				
LINMIX_ERR fwd ^b	8.32 ± 0.04	5.30 ± 0.22	0.49 ± 0.03	0.091
Geo-MAP fwd	8.42 ± 0.04	5.90 ± 0.26	0.55 ± 0.06	0.091 ± 0.008
Geo-MCMC fwd	$8.44^{+0.09}_{-0.09}$	$6.15^{+0.42}_{-0.40}$	$0.54^{+0.07}_{-0.06}$	$0.087^{+0.011}_{-0.009}$
van den Bosch (2016) data, no upper limits				
FITEXY fwd	8.35 ± 0.04	4.91 ± 0.23	0.48^c	0.096
FITEXY inv	8.42 ± 0.05	6.51 ± 0.30	0.55^c	0.084
LINMIX_ERR fwd	8.34 ± 0.04	4.92 ± 0.25	0.49 ± 0.03	0.098
LINMIX_ERR inv	8.43 ± 0.05	6.57 ± 0.32	0.56 ± 0.05	0.084
Geo-MAP fwd	8.43 ± 0.04	6.70 ± 0.40	0.54 ± 0.07	0.080 ± 0.008
Geo-MCMC fwd	$8.43^{+0.10}_{-0.09}$	$6.68^{+0.67}_{-0.55}$	$0.55^{+0.08}_{-0.06}$	$0.082^{+0.011}_{-0.010}$

Notes. ^aThe orthogonal intrinsic scatter σ_{int} has been converted to intrinsic scatter along the M_{BH} coordinate via the equation (43) for both Geo methods. Similarly, the scatter in M_{BH} -direction, $\sigma_{\text{int}, M_{\text{BH}}}$, has been converted to equivalent orthogonal intrinsic scatter via the same equation for the LINMIX_ERR and FITEXY methods.

^bThe LINMIX_ERR method only supports upper limits on the dependent variable. As such, inverse fits cannot be computed.

^cThe FITEXY method does not give error estimates for the intrinsic scatter.

probability sample, and descending in posterior probability until the limit was exceeded. The extent of the credible region in each parameter direction was then used to compute the upper and lower 1σ limits.

The results of this analysis are displayed in Table 2. Shown also are the results from Park et al. (2012) containing both forward and inverse regressions with the FITEXY (Press et al. 1992; Tremaine et al. 2002) and ‘Bayesian’ (i.e. LINMIX_ERR) methods. In addition, fits for the van den Bosch (2016) data set without upper limits were computed separately for all methods. A graphical representation of the data sets and the relations obtained with the different methods is shown in Fig. 5.

Table 2 indicates that the results computed using the geometrical approach presented in this paper are more consistent with the inverse fits done using the LINMIX_ERR and FITEXY algorithms. This is not surprising in light of the discussion in the previous section, since the slope of the $M_{\text{BH}}-\sigma$ relation is high and we ex-

pect the forward fits to be biased towards low slopes in this situation. The effects of this bias can also be seen from the best-fitting values of the orthogonal intrinsic scatter, σ_{int} , for which the forward fits yield consistently higher values than the inverse fits for LINMIX_ERR and FITEXY. It seems that in these methods the lower value for the slope is compensated by a higher estimate for the intrinsic scatter. Note that this behaviour cannot be appreciated by looking at the values of the scatter in M_{BH} -direction, $\sigma_{\text{int}, M_{\text{BH}}}$, since these are not truly intrinsic, but depend on β and indeed show opposite behaviour. In contrast, the geometric methods yield identical results (up to sampling noise) in forward and inverse directions, so only forward results are shown in Table 2. The geometric methods also give consistently smaller estimates for the orthogonal intrinsic scatter. The 1σ errors for all parameters are in general comparable between the methods, with the exception of the fully Bayesian MCMC approach, which consistently gives more conservative 1σ errors.

Figure 5. $M_{\text{BH}}-\sigma$ relation $\log_{10}(M_{\text{BH}}/M_{\odot}) = \alpha + \beta \log_{10}(\sigma/200 \text{ km s}^{-1})$ derived using the data sets in Gültekin et al. (2009) (top left), McConnell et al. (2011) (top right), Graham et al. (2011) (bottom left) and van den Bosch (2016) (bottom right). The data sets have been rendered as density maps, assuming two-dimensional Gaussian probability densities derived from the 1σ errors of the data points. The data points are shown with red dots. The upper limits in van den Bosch (2016) data are rendered as a product of one-dimensional Gaussian and the logarithmic y upper limit density (equation B6). The fitted relations are overplotted with lines. Intrinsic scatter 1σ limits are shown with shaded regions for the Geo-MAP and LINMIX_ERR forward fits for demonstration. The values of intrinsic scatter produced by the rest of the methods are comparable (see Table 2).

Finally, it seems that if there indeed is a fundamental approximately linear relation between the logarithms of the mass of a supermassive black hole mass and the velocity dispersion of its host galaxy, the slope of the relation is likely to be $\gtrsim 6$, at least based on the van den Bosch (2016) data, which is the most comprehensive data set used here. This is in contrast to the values around 4–5 often obtained in the literature (e.g. 4.8, Ferrarese & Merritt 2000; 3.75, Gebhardt et al. 2000; 4.02, Tremaine et al. 2002; 4.24, Gültekin et al. 2009; 5.64, McConnell & Ma 2013), including the value 5.35 derived in van den Bosch (2016). These values were all derived using a method or a variation of a method described in Section 3.3.1, fitting the $M_{\text{BH}}-\sigma$ relation in the ‘forward’ direction (M_{BH} as a function of σ), in which it has a high numerical value for the slope.

This then causes the fitted value of the slope to be biased towards lower values. However, as can be seen from Fig. 5, the numerically rather different estimates of the slope are not visually at all that obvious, with a change in slope of 5–6 corresponding only to an $\sim 1.8^\circ$ change in angle of the regression line with respect to a fixed direction, such as the x -axis.

The obtained slope also crucially depends on the observational and other biases the data may have. Indeed, it is evident from Table 2 that the choice of data set has an effect that is roughly comparable to the effect of the choice of method. It should also be noted that the fundamental relation, if there is one, may involve more than two physical quantities or their measurable proxies, as suggested in van den Bosch (2016). Fitting relations to a sampled projection

of this hyperplane would then in general yield a higher value for the intrinsic scatter and slope that is offset from the ‘true’ value. As such, there is an urgent need for more high-quality data in order to say much with any certainty regarding the slope of the $M_{\text{BH}}-\sigma$ relation (or the possible multivariable generalizations), its possible evolution with redshift, or whether it truly is linear across the entire range of black hole masses.

5 SUMMARY

We have presented a mathematical formalism for representing physical relations as submanifolds S of a Riemannian manifold of observables M . In this geometric approach, intrinsic scatter in the relation can be accommodated with probability distributions defined on the normal spaces of S . Given a probability distribution of data, and parametrizations of the relation and the intrinsic scatter distribution, the formalism then yields a Bayesian posterior probability for the parameters of the relation and the intrinsic scatter distribution (equation 11). The novelty of our formulation is that it fully accommodates arbitrary measurement errors, both left and right censored data (upper and lower limits, respectively), truncation (non-detections) and extends the concept of intrinsic scatter both to non-linear relations and relations that define a submanifold of codimension greater than one.

We have derived explicit analytic results for the likelihood and the posterior distribution first in the case where the postulated relation defines a linear $n - 1$ -dimensional hyperplane. We then extended this result to the case where the relation defines an $n - k$ -dimensional affine subspace, for an arbitrary k , a result we believe to be potentially highly useful for seeking out the most likely codimension of a correlation within a set of n -dimensional data. Finally, we have derived the likelihood and posterior distribution in the case of a line in two dimensions and discussed its implications at length. We also compared the results given by our method with two established methods widely used in astronomical literature, namely `FITEXY` and `LINMIX_ERR`. We demonstrated that our inherently symmetrical geometrical approach is preferable in situations where the data obeys a relation with a slope much larger or smaller than one, and measurement errors and intrinsic scatter are severe and equally important.

Finally, we used our method to fit the $M_{\text{BH}}-\sigma$ relation, between the mass M_{BH} of a supermassive black hole in a galactic bulge and the stellar velocity dispersion σ of the bulge, using several published data sets. We compared our results to the fits in the literature and find that our results support a slope of ~ 6 , clearly higher than the slopes $\sim 4-5$ derived in the literature. We note that this difference is mainly due to the methods used to derive the literature results. We show that if these methods are used ‘in reverse’ to fit σ as a function of M_{BH} and then inverting the slope, the results are in much better agreement with ours. This is due to the tendency of standard methods, such as `FITEXY` or `LINMIX_ERR`, which do not respect the geometric symmetry of the problem, to misestimate steep slopes in the presence of intrinsic scatter (see Section 3.3.1 for the discussion).

ACKNOWLEDGEMENTS

The author is most grateful to the anonymous referee for comments and ideas that have contributed to a much improved manuscript, as well as to Peter H. Johansson for extensive comments on the manuscript drafts. The numerical computations in this work have benefited from the PYTHON libraries NUMPY, SCIPY and SYMPY (Jones

et al. 2001; van der Walt, Colbert & Varoquaux 2011; Meurer et al. 2017). The figures have been rendered with the help of the PYTHON library Matplotlib (Hunter 2007). This work has made use of NASA’s Astrophysics Data System Bibliographic Services. The research for this publication was supported by the Academy of Finland grant no. 1274931.

REFERENCES

- Akritis M. G., Bershadsky M. A., 1996, *ApJ*, 470, 706
- Babu G. J., Singh K., 1983, *Ann. Stat.*, 11, 999
- Beifiori A., Courteau S., Corsini E. M., Zhu Y., 2012, *MNRAS*, 419, 2497
- Boggs P. T., Byrd R. H., Schnabel R. B., 1987, *SIAM J. Sci. Stat. Comput.*, 8, 1052
- Boggs P. T., Spiegelman C. H., Donaldson J. R., Schnabel R. B., 1988, *J. Econom.*, 38, 169
- Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
- Broemeling L., Broemeling A., 2003, *Biometrika*, 90, 728
- Calin O., Udriste C., 2014, *Geometric Modeling in Probability and Statistics. Mathematics and Statistics*. Springer International Publishing, Switzerland
- Choi E., Naab T., Ostriker J. P., Johansson P. H., Moster B. P., 2014, *MNRAS*, 442, 440
- Crane K., Weischedel C., Wardetzky M., 2013, *Geodesics in Heat: A New Approach to Computing Distance Based on Heat Flow*. *ACM Trans. Graph.*, 28, 4
- Di Matteo T., Springel V., Hernquist L., 2005, *Nature*, 433, 604
- Di Matteo T., Colberg J., Springel V., Hernquist L., Sijacki D., 2008, *ApJ*, 676, 33
- Feigelson E. D., Babu G. J., 1992, *ApJ*, 397, 55
- Feigelson E. D., Babu G. J., 2012, *Modern Statistical Methods for Astronomy: With R Applications*. Cambridge Univ. Press, Cambridge
- Ferrarese L., Merritt D., 2000, *ApJ*, 539, L9
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Fraser D. A. S., Reid N., Marras E., Yi G. Y., 2010, *J. R. Stat. Soc. Ser. B (Stat. Methodol.)*, 72, 631
- Fujiki J., Akaho S., 2009, in 2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops. Kyoto, Japan, p. 250
- Gebhardt K. et al., 2000, *ApJ*, 539, L13
- Gelman A., Carlin J., Stern H., Dunson D., Vehtari A., Rubin D., 2013, *Bayesian Data Analysis*, 3rd edn. Chapman & Hall, CRC Press, Boca Raton, FL
- George E. I., McCulloch R., 1993, *J. Stat. Plan. Inference*, 37, 169
- Graham A. W., 2016, *Galactic Bulges*, 418, 263
- Graham A. W., Onken C. A., Athanassoula E., Combes F., 2011, *MNRAS*, 412, 2211
- Gull S. F., 1989, *Maximum Entropy and Bayesian Methods*. Springer Science+Business Media, Dordrecht, p. 511
- Gültekin K. et al., 2009, *ApJ*, 698, 198
- Hogg D. W., Bovy J., Lang D., 2010, preprint ([arXiv:1008.4686](https://arxiv.org/abs/1008.4686))
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Isobe T., Feigelson E. D., Akritis M. G., Babu G. J., 1990, *ApJ*, 364, 104
- Jaynes E., 1983, in Rosenkrantz R. D., ed., *E. T. Jaynes: Papers on Probability, Statistics and Statistical Physics*, Vol. 158. Reidel, Dordrecht, p. 190
- Jaynes E. T., 1991, *Straight Line Fitting – A Bayesian Solution*, Unpublished, available online at: <http://bayes.wustl.edu/sfg/line.pdf>
- Jaynes E. T., 2003, *Probability Theory – The Logic of Science*. Cambridge Univ. Press, Cambridge
- Johansson P. H., Naab T., Burkert A., 2009a, *ApJ*, 690, 802
- Johansson P. H., Burkert A., Naab T., 2009b, *ApJ*, 707, L184
- Jones E. et al., 2001, *SciPy: Open Source Scientific Tools for Python*, <http://www.scipy.org/>
- Jupp P. E., Kent J. T., 1987, *J. R. Stat. Soc. Ser. C (Appl. Stat.)*, 36, 34
- Kelly B. C., 2007, *ApJ*, 665, 1489

- Lee J., 2013, *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics. Springer, New York
- Magorrian J. et al., 1998, *AJ*, 115, 2285
- Markwardt C. B., 2009, in Bohlender D. A., Durand D., Dowler P., eds, *ASP Conf. Ser. Vol. 411, Astronomical Data Analysis Software and Systems XVIII*. Astron. Soc. Pac., San Francisco, p. 251
- McConnell N. J., Ma C.-P., 2013, *ApJ*, 764, 184
- McConnell N. J., Ma C.-P., Gebhardt K., Wright S. A., Murphy J. D., Lauer T. R., Graham J. R., Richstone D. O., 2011, *Nature*, 480, 215
- Meurer A. et al., 2017, *PeerJ Comp. Sci.*, 3, e103
- Novak G. S., Faber S. M., Dekel A., 2006, *ApJ*, 637, 96
- Park D., Kelly B. C., Woo J.-H., Treu T., 2012, *ApJS*, 203, 6
- Pearson K., 1901, *Phil. Mag. Ser. 6*, 2, 559
- Pennec X., 2006, *J. Math. Imaging Vis.*, 25, 127
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, *Numerical Recipes in FORTRAN: The Art of Scientific Computing*, 2nd edn. Cambridge Univ. Press, Cambridge
- Robotham A. S. G., Obreschkow D., 2015, *PASA*, 32, e033
- Saglia R. P. et al., 2016, *ApJ*, 818, 47
- Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007, *MNRAS*, 380, 877
- Sijacki D., Vogelsberger M., Genel S., Springel V., Torrey P., Snyder G. F., Nelson D., Hernquist L., 2015, *MNRAS*, 452, 575
- Stigler S. M., 1981, *Ann. Statist.*, 9, 465
- Tremaine S. et al., 2002, *ApJ*, 574, 740
- van den Bosch R. C. E., 2016, *ApJ*, 831, 134
- van der Walt S., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22
- Williams M. J., Bureau M., Cappellari M., 2010, *MNRAS*, 409, 1330
- Wolf J. A., Zierau R., 1996, in Gindikin S., ed., *Progress in Nonlinear Differential Equations and Their Applications*, Vol. 20, Riemannian Exponential Maps and Decompositions of Reductive Lie Groups. Birkhäuser Boston, p. 349
- Zellner A., 1971, *An introduction to Bayesian Inference in Econometrics*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons, Inc., US

APPENDIX A: DERIVATION OF THE POSTERIOR DISTRIBUTION

We present a derivation of the likelihood (equation 10) and the posterior distribution (equation 11) by following the Bayesian style promoted in Jaynes (2003). We consider the following propositions:

- (i) η = ‘the true value of the observable is $\eta \in M$ ’
- (ii) X = ‘a value $x \in M$ was measured for the observable’
- (iii) V = ‘a detection was made’
- (iv) H = ‘the true value of the observable is drawn from the distribution p_{int} defined on a relation S , with parameters $\boldsymbol{\varphi}$ and $\boldsymbol{\theta}$, respectively’

In addition, we will use I to specify all the other relevant prior information. This includes the prior distribution $\pi_{\boldsymbol{\theta}, \boldsymbol{\varphi}}$ of $\boldsymbol{\theta}$ and $\boldsymbol{\varphi}$, the measurement error distribution $h(\eta; x)$ and the fact that the probability of a detection for a value η of the observable is given by $p_{\text{det}}(\eta)$.

At the outset, we then know the following probabilities:

$$P(\eta|HI) = p_{\text{int}}(\eta; \boldsymbol{\theta}, \boldsymbol{\varphi}) := p_{\text{int}}(\tau(\eta; \boldsymbol{\theta}), \nu(\eta; \boldsymbol{\theta}); \boldsymbol{\varphi}) \quad (\text{A1})$$

$$P(\eta|XI) = h(\eta; x) \quad (\text{A2})$$

$$P(V|\eta I) = p_{\text{det}}(\eta). \quad (\text{A3})$$

We will need the probability $P(X|\eta VI)$ of a measured value, given a true value and the fact that there is a detection. This is obtained with Bayes’ theorem:

$$P(X|\eta VI) = P(\eta|XVI) \frac{P(X|VI)}{P(\eta|VI)} = h(\eta; x), \quad (\text{A4})$$

as a function of x (which we write as $h(x; \eta)$ in the following), since the prior probabilities $P(X|VI)$ and $P(\eta|VI)$ must be uninformative and equal everywhere, since nothing in I tells us where the true and measured values are a priori.

We can now compute the probability of a measured value given the true value, yielding

$$\begin{aligned} P(X|\eta I) &= P(X(V + \bar{V})|\eta I) = P(XV|\eta I) + P(X\bar{V}|\eta I) \\ &= P(X|\eta VI)P(V|\eta I) + P(X|\eta \bar{V}I)P(\bar{V}|\eta I) \\ &= h(x; \eta)p_{\text{det}}(\eta) + 0 \cdot [1 - p_{\text{det}}(\eta)] \\ &= h(x; \eta)p_{\text{det}}(\eta), \end{aligned} \quad (\text{A5})$$

where $\bar{V}$ is the negation of V (i.e. there was no detection). Here, we used the fact that V and $\bar{V}$ form a complete ($P(V + \bar{V}|I) = 1$) and independent set of propositions, together with the identity $P(AB|C) = P(A|C)P(B|AC) = P(B|C)P(A|BC)$.

Using Bayes’ theorem again, we can now obtain the probability $P(H|XI)$ that the data obey the relation and are drawn from p_{int} , given the measured value. The theorem gives

$$P(H|XI) = P(X|HI) \frac{P(H|I)}{P(X|I)}. \quad (\text{A6})$$

Since the η also form a complete and independent set of propositions (the true value must be somewhere, and the possible positions are independent), we can write

$$\begin{aligned} P(X|HI) &= \int P(X\eta|HI)d\eta = \int P(X|\eta HI)P(\eta|HI)d\eta \\ &= \int h(x; \eta)p_{\text{det}}(\eta)p_{\text{int}}(\eta; \boldsymbol{\theta}, \boldsymbol{\varphi})d\eta \\ &= \mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\varphi}|x), \end{aligned} \quad (\text{A7})$$

which gives the likelihood (equation 11). The factor $P(H|I) = \pi_{\boldsymbol{\theta}, \boldsymbol{\varphi}}(\boldsymbol{\theta}, \boldsymbol{\varphi})$ is the parameter prior probability, and the denominator is a normalizing constant, given formally by

$$\begin{aligned} P(X|I) &= \int P(XH|I)dH = \int P(X|HI)P(H|I)dH \\ &= \iint h(x; \eta)p_{\text{det}}(\eta)p_{\text{int}}(\eta; \boldsymbol{\theta}, \boldsymbol{\varphi})\pi_{\boldsymbol{\theta}, \boldsymbol{\varphi}}(\boldsymbol{\theta}, \boldsymbol{\varphi})d\eta d\boldsymbol{\theta} d\boldsymbol{\varphi}. \end{aligned} \quad (\text{A8})$$

We then have the posterior probability (equation 11),

$$P(H|XI) = \frac{\int h(x; \eta)p_{\text{det}}(\eta)p_{\text{int}}(\eta; \boldsymbol{\theta}, \boldsymbol{\varphi})d\eta \pi_{\boldsymbol{\theta}, \boldsymbol{\varphi}}(\boldsymbol{\theta}, \boldsymbol{\varphi})}{\iint h(x; \eta)p_{\text{det}}(\eta)p_{\text{int}}(\eta; \boldsymbol{\theta}, \boldsymbol{\varphi})\pi_{\boldsymbol{\theta}, \boldsymbol{\varphi}}(\boldsymbol{\theta}, \boldsymbol{\varphi})d\eta d\boldsymbol{\theta} d\boldsymbol{\varphi}}. \quad (\text{A9})$$

APPENDIX B: AN EXAMPLE OF LIMIT DISTRIBUTIONS

Here, we present a simple example of how the framework presented in the paper works with upper limits, and how analytic results can be obtained for upper limits as well. This is desirable since typically the analytic expression for the likelihood is much faster to compute than performing a numerical integration to obtain the required value. As mentioned in Section 2.3, limits are taken into account by including them into the measurement error distribution. This can be done for the examples in Sections 3.1–3.3 by substituting one or more

degrees of freedom x_i in equations (16) and (32) with suitable one-dimensional distributions. Conceptually, the simplest possibility is the uniform distribution, so that

$$u(x_i; x_i^u) = \frac{\chi_{[0, x_i^u]}(x_i)}{x_i^u}, \quad (\text{B1})$$

where $\chi_{[a, b]}(x) = 1$ if $x \in [a, b]$ and 0 otherwise, and x_i^u is the limiting value. Other choices with less pronounced cut-offs are naturally also possible. If the measurements are already given in logarithmic units with base k , the form

$$u_{\log_k}(x_i; x_i^u) = u(k^{x_i}; k^{x_i^u}) k^{x_i} \log k \quad (\text{B2})$$

must be used instead. Lower limits can be introduced in an analogous manner.

Analytic results of reasonable complexity based on these limit distributions can be obtained for the example in Section 3.3. For the linear case, we have

$$\begin{aligned} \mathcal{L}_x(\alpha, \beta, \sigma | x^u, y_0, \sigma_y) \\ = \frac{1}{2x^u} \frac{\sqrt{1+\beta^2}}{\beta} \left\{ \operatorname{erf} \left(\frac{\alpha + x^u \beta - y_0}{\sqrt{2(\sigma_y^2 + (1+\beta^2)\sigma^2)}} \right) \right. \\ \left. - \operatorname{erf} \left(\frac{\alpha - y_0}{\sqrt{2(\sigma_y^2 + (1+\beta^2)\sigma^2)}} \right) \right\} \text{ and} \end{aligned} \quad (\text{B3})$$

$$\begin{aligned} \mathcal{L}_y(\alpha, \beta, \sigma | x_0, y^u, \sigma_x) \\ = \frac{1}{2y^u} \sqrt{1+\beta^2} \left\{ \operatorname{erf} \left(\frac{\alpha + x_0 \beta}{\sqrt{2(\beta^2 \sigma_x^2 + (1+\beta^2)\sigma^2)}} \right) \right. \\ \left. - \operatorname{erf} \left(\frac{\alpha + x_0 \beta - y^u}{\sqrt{2(\beta^2 \sigma_x^2 + (1+\beta^2)\sigma^2)}} \right) \right\} \end{aligned} \quad (\text{B4})$$

for the upper limits in x - and y -directions, respectively, where x^u and y^u are the corresponding limiting values. If the measurements have been given in logarithmic scale with base k , we have instead

$$\begin{aligned} \mathcal{L}_{\log_k x}(\alpha, \beta, \sigma | x^u, y_0, \sigma_y) \\ = \frac{\log k}{2x^u} \sqrt{\frac{1+\beta^2}{\beta^2}} k^{\frac{2(y_0 - \alpha)\beta + [\sigma_y^2 + (1+\beta^2)\sigma^2] \log k}{2\beta^2}} \\ \operatorname{erfc} \left(\frac{(y_0 - \alpha)\beta + [\sigma_y^2 + (1+\beta^2)\sigma^2] \log k - \beta^2 \log_k x^u}{\sqrt{2[\sigma_y^2 + (1+\beta^2)\sigma^2] |\beta|}} \right) \text{ and} \end{aligned} \quad (\text{B5})$$

$$\begin{aligned} \mathcal{L}_{\log_k y}(\alpha, \beta, \sigma | x_0, y^u, \sigma_x) \\ = \frac{\log k}{2y^u} \sqrt{1+\beta^2} k^{\alpha + x_0 \beta + \frac{1}{2}[\beta^2 \sigma_x^2 + (1+\beta^2)\sigma^2] \log k} \\ \operatorname{erfc} \left(\frac{\alpha + x_0 \beta + [\beta^2 \sigma_x^2 + (1+\beta^2)\sigma^2] \log k - \log_k y^u}{\sqrt{2[\beta^2 \sigma_x^2 + (1+\beta^2)\sigma^2]}} \right), \end{aligned} \quad (\text{B6})$$

where $\operatorname{erfc}(x) = 1 - \operatorname{erf}(x)$. Note that here the limiting values x^u and y^u are given in linear scale. For numerical applications, it should be noted that the arguments of the k -exponential and the erfc function may have large numerical values, and asymptotic expansions should be used when necessary.

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.