

PAPER • OPEN ACCESS

Mixed noise and posterior estimation with conditional deepGEM

To cite this article: Paul Hagemann *et al* 2024 *Mach. Learn.: Sci. Technol.* **5** 035001

View the [article online](#) for updates and enhancements.

You may also like

- [Gradients should stay on path: better estimators of the reverse- and forward KL divergence for normalizing flows](#)
Lorenz Vaitl, Kim A Nicoli, Shinichi Nakajima et al.
- [-deep Probabilistic Inference \(-DPI\): Efficient Uncertainty Quantification from Exoplanet Astrometry to Black Hole Feature Extraction](#)
He Sun, Katherine L. Bouman, Paul Tiede et al.
- [Deep Potential: Recovering the Gravitational Potential from a Snapshot of Phase Space](#)
Gregory M. Green, Yuan-Sen Ting, et al.



PAPER

Mixed noise and posterior estimation with conditional deepGEM

Paul Hagemann^{1,*} , Johannes Hertrich² , Maren Casfor³, Sebastian Heidenreich³ and Gabriele Steidl¹

¹ Institute of Mathematics, TU Berlin, Berlin, Germany

² Department of Computer Science, University College London, London, United Kingdom

³ Physikalisch Technische Bundesanstalt, Berlin, Germany

* Author to whom any correspondence should be addressed.

E-mail: hagemann@math.tu-berlin.de, j.hertrich@ucl.ac.uk, maren.casfor@ptb.de, Sebastian.Heidenreich@ptb.de and steidl@math.tu-berlin.de

Keywords: normalizing flow, inverse problem, posterior, expectation maximization

OPEN ACCESS

RECEIVED

19 February 2024

REVISED

29 May 2024

ACCEPTED FOR PUBLICATION

17 June 2024

PUBLISHED

2 July 2024

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Abstract

We develop an algorithm for jointly estimating the posterior and the noise parameters in Bayesian inverse problems, which is motivated by indirect measurements and applications from nanometrology with a mixed noise model. We propose to solve the problem by an expectation maximization (EM) algorithm. Based on the current noise parameters, we learn in the E-step a conditional normalizing flow that approximates the posterior. In the M-step, we propose to find the noise parameter updates again by an EM algorithm, which has analytical formulas. We compare the training of the conditional normalizing flow with the forward and reverse Kullback–Leibler divergence, and show that our model is able to incorporate information from many measurements, unlike previous approaches.

1. Introduction

In a variety of healthcare and other contemporary applications, the variables of primary interest are obtained through indirect measurements, such as in the case of magnetic resonance imaging and computed tomography. For some of these applications, the reliability of the results is of particular importance. The accuracy and trustworthiness of the outcomes obtained through indirect measurements are significantly influenced by two critical factors: the degree of uncertainty associated with the measuring instrument and the appropriateness of the (forward) model used for the reconstruction of the parameters of interest (measurand). In this paper, we consider Bayesian inversion to obtain the measurand from signals measured by the instrument and a noise model that mimics background noise coming from the instrument and the variation of the measurement, depending on the forward model. Within this framework, we developed an extension of the expectation maximization (EM) algorithm that is able to handle a Bayesian inversion with a measurement noise model. As a result, we obtain the posterior distribution for the parameters of interest (distribution of the measurand), which is a measure of the reliability of the measurement results. To demonstrate the applicability and effectiveness we apply the algorithm to two real examples in nanometrology, i.e. EUV Scatterometry. The key focus of the work is the development of a noise-adapted posterior sampler based on DeepGEM [24], which can incorporate information from several measurements simultaneously.

In this context we consider Bayesian inverse problems

$$Y_\theta = F(X) + \eta_\theta, \quad (1)$$

with a possibly nonlinear forward operator $F: \mathbb{R}^d \rightarrow \mathbb{R}^n$ and a random noise variable η_θ which depends on an unknown parameter θ and on $F(X)$. Note, Y describes the signals of the instrument whereas X are the parameters of interest. The posterior (parameter distribution) $P_{X|Y_\theta=y}$ for observations y will ultimately depend on these parameters θ , as the likelihood $P_{Y_\theta|X=x}$ depends on them. Therefore, we aim to estimate the parameter θ from observations $y_i \in \mathbb{R}^n$, $i = 1, \dots, N$, where N is possibly small. There exists plenty of literature on estimating the standard deviation σ within the Gaussian noise model $\eta_\theta = \eta_\sigma \sim \mathcal{N}(0, \sigma^2 I_n)$.

However, motivated by applications in nanometrology [30], we are interested in a mixture of additive and multiplicative Gaussian noise of the form

$$\eta_\theta = \eta_{(a,b)} = \eta_1 + \eta_2, \quad \eta_1 \sim \mathcal{N}(0, a^2 I_n), \quad \eta_2 \sim \mathcal{N}\left(0, b^2 \text{diag}\left(F(x)^2\right)\right), \quad (2)$$

where $F(x)^2 = (F_i(x)^2)_{i=1}^n$ and the identity in $\mathbb{R}^{n \times n}$ is given by I_n . For convenience we assume that the instrument noise and other sources can be described by the simple Ansatz made here. In general different noise models may appear in the applications. The noise model equation (2) was used in several previous studies in optics [29, 30, 32, 59] and analyzed in [19]. A similar noise model appears in analytical chemistry [58] and the study of gene expression arrays [57]. It belongs to the class of heteroskedastic noise models [22] and an algorithm for parameter estimation in a slightly different problem was proposed in [23]. Learning the noise model without any parametric form was done using NFs in [1].

The standard approach for parameter estimation is *maximum likelihood estimation*. That is, we choose θ as the minimizer of the negative log likelihood function

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N \log(p_\theta(y_i)),$$

where p_θ is the probability density function of Y_θ . However, in our case this function involves a high-dimensional integral of the form $\int_x p_{Y_\theta|X=x}(y) dP_X(x)$ which is intractable to compute. As a remedy, we exploit EM algorithms which were introduced in [16], see also [10] for an overview and [50, 51] for applications for parameter estimation in inverse problems. The basic idea is to iteratively compute the posterior distribution $P_{X|Y_\theta(r)=y_i}$ for a current estimate $\theta^{(r)}$ of θ and then updating this estimate to $\theta^{(r+1)}$ based on this posterior distribution. Here, the computation of $P_{X|Y_\theta(r)=y_i}$ is called E-step, while the update of θ is called M-step. Intuitively, this corresponds to the idea that the distribution of $F(X)$ is approximately concentrated on a lower-dimensional manifold and consequently the distance of the y_i to this manifold contains the information of the noise parameters θ .

Recently, Gao et al [24] proposed to solve the E-step by a normalizing flow (NF) [18] using the reverse/backward Kullback–Leibler (KL) divergence as loss function. For the M-step, they apply a stochastic gradient ascent algorithm. Note that in general the same procedure can be applied to estimate parameters of the forward operator instead of the noise level, see [41]. This approach has several drawbacks. The model needs to be retrained for every new observations y and cannot profit from many observations that follow the same error parameters. Furthermore, the reverse KL is known to be *mode-seeking*. That is, it tends to recover only one mode of multimodal distributions, which incorporates a significant approximation error, see the discussions in [26, 48, 67] for more details. Very related is also the JANA framework [55], where the authors propose to learn forward (likelihood), posterior and summary network together in an amortized (i.e. conditional manner). However, they optimize it not iteratively, since they do not treat it as an EM framework and do not discuss noise modelling. The main idea of this paper was presented by some of the authors in a one-page extended abstract in [56].

1.1. Contributions

First, we propose to use the conditional NFs [7, 66] in the E-step. This allows the incorporation of *several* measurements from the same error model and to solve the inverse problem for all measurements *simultaneously*. Fortunately, the forward KL [7] can be used as loss function for training the conditional NFs which makes the method mode covering. Second, we propose an inner EM algorithm for solving the M-step more efficiently. For our special noise model (2), we deduce analytic expressions for E- and M-steps of this inner algorithm. The performance of our approach will be demonstrated on two applications from nano-optics. In particular, we propose a conditional version of DeepGEM and benchmark it against forward conditional DeepGEM, where the reverse KL is replaced by a forward KL.

1.2. Organization

We start in section 2 by recalling the general EM algorithm. Then, in section 3, we construct the E-step and M-step for our application. That is, we show how conditional NFs can be incorporated into the E-step and describe how the M-step can be solved for our noise model with an ‘inner’ EM algorithm which steps can be given in a closed analytical form. Some of the technical computations are postponed to appendix A. We test our algorithms on two nano-optics problems, which is done in section 4. Finally, conclusions are drawn in section 5.

2. EM algorithm

In this section, we introduce the EM algorithm as a maximization-maximization algorithm of an evidence lower bound. A general introduction into the EM algorithm can be found, e.g. in [10].

Let $\{Y_\theta : \theta \in \Theta\}$ be a family of n -dimensional random variables having probability density functions p_θ , $\theta \in \Theta$. Given i.i.d. samples $y_1, \dots, y_N \in \mathbb{R}^n$ from Y_{θ^*} for some unknown θ^* , which we want to approximate by computing the maximum log-likelihood estimator

$$\hat{\theta} = \arg \max_{\theta \in \Theta} \mathcal{L}(\theta), \quad \mathcal{L}(\theta) := \frac{1}{N} \sum_{i=1}^N \log(p_\theta(y_i)).$$

In the literature, the term $\log(p_\theta(y))$ is also called *evidence* of y under θ . As in many applications it is hard to maximize \mathcal{L} , we introduce an absolute continuous d -dimensional auxiliary random variable X such that the joint density p_{X,Y_θ} exists and is easy to evaluate. Then, it holds by the law of total probability and Jensen's inequality, for any probability density function q on \mathbb{R}^d , that

$$\begin{aligned} \log(p_\theta(y)) &= \log\left(\int_{\mathbb{R}^d} p_{X,Y_\theta}(x,y) \, dx\right) = \log\left(\int_{\mathbb{R}^d} \frac{p_{X,Y_\theta}(x,y)}{q(x)} q(x) \, dx\right) \\ &\geq \int_{\mathbb{R}^d} \log\left(\frac{p_{X,Y_\theta}(x,y)}{q(x)}\right) q(x) \, dx =: \mathcal{F}(q, \theta|y). \end{aligned}$$

We call the random variable X the hidden variable and the expression $\mathcal{F}(q, \theta|y)$ the *evidence lower bound* (ELBO). Now, instead of maximizing the log-likelihood function directly, the EM algorithm is a maximization-maximization algorithm for the ELBO, i.e. starting with an initial estimate $\theta^{(0)}$, it consists of the following two steps:

$$\text{E-step: } q_i^{(r+1)} = \arg \max_{q \in \text{pdf}(d)} \mathcal{F}(q, \theta^{(r)}|y_i), \quad i = 1, \dots, N, \tag{3}$$

$$\text{M-step: } \theta^{(r+1)} = \arg \max_{\theta \in \Theta} \sum_{i=1}^N \mathcal{F}(q_i^{(r+1)}, \theta|y_i), \tag{4}$$

where pdf is the space of d -dimensional probability density functions.

The E-step (3) can be solved based on the following standard lemma which can be found, e.g. in [9, section 9.4]. For convenience, we provide the simple proof. Recall that the *Kullback–Leibler* (KL) *divergence* of two probability measures P, Q with densities p, q is defined by

$$\text{KL}(P, Q) := \int_{\mathbb{R}^d} p(x) \log \frac{p(x)}{q(x)} \, dx,$$

if P is absolutely continuous with respect to Q , and $\text{KL}(P, Q) = +\infty$ otherwise. Further, we use the convention $0 \log 0 = 0$.

Lemma 1. *Let $X \in \mathbb{R}^d$ be a absolute continuous random variable and let Q be an absolutely continuous measure on \mathbb{R}^d with probability density function q . Then it holds, for any $y \in \mathbb{R}^n$, that*

$$\log(p_\theta(y)) - \mathcal{F}(q, \theta|y) = \text{KL}(Q, P_{X|Y_\theta=y}).$$

Proof. By definition of the conditional distribution, we have

$$\begin{aligned} &\log(p_\theta(y)) - \mathcal{F}(q, \theta|y) \\ &= \log(p_\theta(y)) - \int_{\mathbb{R}^d} q(x) \log\left(\frac{p_{X|Y_\theta=y}(x) p_\theta(y)}{q(x)}\right) \, dx \\ &= \log(p_\theta(y)) - \int_{\mathbb{R}^d} q(x) \log(p_\theta(y)) \, dx + \int_{\mathbb{R}^d} q(x) \log\left(\frac{q(x)}{p_{X|Y_\theta=y}(x)}\right) \, dx \\ &= \log(p_\theta(y)) - \log(p_\theta(y)) + \text{KL}(Q, P_{X|Y_\theta=y}) = \text{KL}(Q, P_{X|Y_\theta=y}). \end{aligned}$$

□

As the KL divergence $KL(Q, P)$ is minimal if and only if $Q = P$, the lemma implies that the solutions q_i of the E-step (3) are given explicitly by

$$\text{E-step: } q_i^{(r+1)} = p_{X|Y_{\theta^{(r)}}=y_i}, \quad i = 1, \dots, N. \tag{5}$$

For solving the M-step (4), we decompose the ELBO as

$$\mathcal{F}(q, \theta|y) = \mathbb{E}_{x \sim q} [\log(p_{X, Y_{\theta}}(x, y))] - \mathbb{E}_{x \sim q} [\log(q(x))]. \tag{6}$$

Note, that only the first summand depends on θ . Using this decomposition and the explicit form (5) of the q_i in the M-step, we obtain the classical EM algorithm as proposed in [16]:

$$\begin{aligned} \text{M-step: } \quad \theta^{(r+1)} &= \arg \max_{\theta \in \Theta} \mathcal{Q}(\theta, \theta^{(r)}), \\ \text{E-step: } \quad \mathcal{Q}(\theta, \theta^{(r)}) &= \sum_{i=1}^N \mathbb{E}_{x \sim p_{X|Y_{\theta^{(r)}}=y_i}} [\log(p_{X, Y_{\theta}}(x, y_i))]. \end{aligned} \tag{7}$$

A convergence analysis of the EM algorithm based on KL proximal point algorithms was done in [13, 14]. In particular, we obtain the following convergence properties.

Proposition 2. *Let the sequence $(q^{(r)}, \theta^{(r)})$ be generated by the EM algorithm. Then, the following holds true.*

- (i) *The sequence of ELBO-values $(\mathcal{F}(q^{(r)}, \theta^{(r)}))_r$ is monotone increasing.*
- (ii) *The sequence of likelihood values $\mathcal{L}(\theta^{(r)})$ is monotone increasing.*

Remark 3 (generalized EM algorithms). Several papers propose so-called generalized EM algorithms [34, 42, 46, 52]. The key idea of these generalizations is to replace the maximization steps (3) and (4) by increase steps. More precisely, in each iteration the values $q_i^{(r+1)}$ and $\theta^{(r+1)}$ are chosen such that

$$\mathcal{F}(q_i^{(r+1)}, \theta^{(r)}|y) \geq \mathcal{F}(q_i^{(r)}, \theta^{(r)}|y), \quad \sum_{i=1}^N \mathcal{F}(q_i^{(r+1)}, \theta^{(r+1)}|y) \geq \sum_{i=1}^N \mathcal{F}(q_i^{(r+1)}, \theta^{(r)}|y).$$

Using such increase steps, generalized EM algorithms often achieve simpler and faster steps, than the original EM algorithm even though they might require more steps until convergence. By construction, part (i) of proposition 2 remains for generalized EM algorithms, while part (ii) is not longer proved for certain of these algorithms.

3. Parameter estimation in Bayesian Inverse Problems

Now we consider the inverse problem (1), where we assume that X has density p_X . Given N observations y_1, \dots, y_N of Y_{θ} , we aim to determine the parameter θ . We will derive an EM algorithm for this problem, where the hidden variable is given by the ground truth random variable X . In particular, we will deal with the noise model (2). Here the parameter $\theta = (a, b)$ can be updated in the M-step analytically.

3.1. E-Step: conditional NFs

As we have seen in (5), the E-step corresponds to finding the posterior densities

$$p_{X|Y_{\theta^{(r)}}=y_i} = \arg \min_q \mathcal{F}(q, \theta^{(r)}|y_i), \quad i = 1, \dots, N,$$

for given $\theta^{(r)}$. We propose to approximate these posteriors by conditional NFs. This extends the so-called DeepGEM from [24] to the conditional case. We will see that our approach brings the forward KL instead of the reverse KL into the play which has several advantages, see remark 4.

A conditional NF is a mapping $\mathcal{T}_{\phi}: \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ depending on some parameters ϕ such that $\mathcal{T}_{\phi}(y, \cdot)$ is invertible for any $y \in \mathbb{R}^n$. In this paper, \mathcal{T}_{ϕ} is a neural network. There were several architectures for NFs proposed in the literature. They include GLOW [39], real NVP [18], invertible ResNets [8, 12, 33] and autoregressive Flows [15, 20, 36, 53]. They were extended to the conditional setting in [7, 17, 27, 40, 66]. The parameters ϕ are learned such that

$$P_Z \circ T_{\phi}(y_i, \cdot)^{-1} = T_{\phi}(y_i, \cdot)_{\#} P_Z \approx P_{X|Y_{\theta^{(r)}}=y_i},$$

for all $i = 1, \dots, N$, where P_Z is some latent distribution, usually a standard Gaussian one. Once we have learned the conditional NF \mathcal{T}_ϕ for an appropriate parameter θ this provides us with a desired approximation of the posterior. In other words, given $y \in \mathbb{R}^n$, we can sample z from Z and produce a sample from $P_{X|Y=y}$ by $\mathcal{T}_\phi(y, z)$.

Now we could learn the conditional NF \mathcal{T}_ϕ by minimizing the loss function

$$\mathcal{J}_{\text{reverse}}(\phi) = \sum_{i=1}^N \mathcal{F} \left(p_{\mathcal{T}_\phi(y_i, \cdot) \# P_Z}, \theta^{(r)} | y_i \right) \propto \sum_{i=1}^N \text{KL} \left(\mathcal{T}_\phi(y_i, \cdot) \# P_Z, P_{X|Y_\theta=y_i} \right),$$

where the last relation follows from lemma 1 and ‘ \propto ’ indicates equality up to a constant. In literature, this loss function is known as *reverse* or *backward* KL loss function. Applying the change of variable formula for push-forward measures and Bayes’s formula this can be rewritten as

$$\begin{aligned} \mathcal{J}_{\text{reverse}}(\phi) &\propto \sum_{i=1}^N \text{KL} \left(P_Z, \mathcal{T}_\phi(y_i, \cdot)^{-1} P_{X|Y_\theta=y_i} \right) \\ &\propto - \sum_{i=1}^N \mathbb{E}_{z \sim P_Z} \left[\log \left(p_{Y_\theta^{(r)}|X=\mathcal{T}_\phi(y_i, z)}(y_i) \right) + \log(p_X(\mathcal{T}_\phi(y_i, z))) + \log(|\det(\nabla_z \mathcal{T}_\phi(y_i, z))|) \right], \end{aligned}$$

see [2, 3, 40] for a detailed explanation and applications. In order to evaluate these terms we have to be able to evaluate the prior density p_X as well as the conditional densities $p_{Y_\theta^{(r)}|X=x}(y)$, which contains the forward operator and the noise model for given parameters $\theta^{(r)}$. Unfortunately, it is known from the literature that the reverse KL is prone to mode collapse, see [48]. That is, in the case that $P_{X|Y_\theta=y_i}$ is multimodal, it tends to generate only samples from one of the modes.

As a remedy, we interchange the arguments in the KL divergence in $\mathcal{J}_{\text{reverse}}$ and replace the sum over the y_i by the expectation over $P_{Y_\theta^{(r)}}$. Then, we arrive at the so-called *forward* KL loss function

$$\begin{aligned} \mathcal{J}_{\text{forward}}(\phi) &= \mathbb{E}_{y \sim P_{Y_\theta^{(r)}}} \left[\text{KL} \left(P_{X|Y_\theta^{(r)}=y}, \mathcal{T}_\phi(y, \cdot) \# P_Z \right) \right] \\ &= \mathbb{E}_{y \sim P_{Y_\theta^{(r)}}} \left[\mathbb{E}_{x \sim P_{X|Y_\theta^{(r)}=y}} \left[\log \left(\frac{p_{X|Y_\theta^{(r)}=y}(x)}{p_{\mathcal{T}_\phi(y, \cdot) \# P_Z}(x)} \right) \right] \right] \\ &\propto - \mathbb{E}_{(x,y) \sim P_{X,Y_\theta^{(r)}}} \left[\log \left(p_Z \left(\mathcal{T}_\phi(y, \cdot)^{-1}(x) \right) \right) + \log |\det(\nabla \mathcal{T}_\phi(y, \cdot)^{-1}(x))| \right]. \end{aligned}$$

To compute these terms we need samples $(\tilde{x}_j, \tilde{y}_j), j = 1, \dots, \tilde{N}$ from the joint distribution $P_{X,Y_\theta^{(r)}}$. Note that such samples can be generated from just knowing the \tilde{x}_j by evaluating the forward operator and the noise model. In this setting, we do not need access to the prior density p_X or the conditional densities $p_{Y_\theta^{(r)}|X=x}$. The forward KL is more standard in (conditional) generative modelling [7, 18, 66] due to these properties and is also known as maximum likelihood training [67].

Remark 4 (forward versus Reverse KL). Note that in the case that \mathcal{T}_ϕ is an universal approximator, we have for both loss functions that the optimal parameters $\hat{\phi}$ fulfills $\mathcal{T}_{\hat{\phi}}(y_i, \cdot) \# P_Z = P_{X|Y_\theta^{(r)}=y_i}$. This is important, as we propose to replace the reverse KL in the E-step by the forward KL. Moreover, the assumptions for training and the approximation properties differ. For the reverse KL, we have to be able to evaluate the density p_X of the prior distribution, while the forward KL needs samples from P_X . In practice it depends on the problem which assumption is more realistic. On the other hand, the forward KL loss function is not that prone to mode collapse. The universality of conditional NFs has been discussed in [44].

3.2. M-Step: inner EM for mixed noise model

As described in (7), the M-step is given by

$$\theta^{(r+1)} = \arg \max_{\theta} \sum_{i=1}^N \mathbb{E}_{x \sim P_{X|Y_\theta^{(r)}=y_i}} [\log(p_{X,Y_\theta}(x, y_i))]. \tag{8}$$

Unfortunately, to the best of our knowledge, an analytic solution of (8) is not available. Therefore, we discretize the expectation in (8) by

$$\theta^{(r+1)} = \arg \max_{\theta} \sum_{i=1}^N \sum_{k=1}^M \log(p_{X,Y_\theta}(x_i^k, y_i)), \tag{9}$$

where the $x_i^k, k = 1, \dots, M$ are sampled from $P_{X|Y_{\theta(r)}=y_i}$. This can be solved by various iterative methods, e.g. by a stochastic gradient algorithm [38] as done in [24].

For our special noise model (2) with $\theta = (a, b)$, we propose to use again an EM algorithm, since both the E- and M-step of the ‘inner’ EM can be computed analytically, which will be shown in the following paragraphs. We use that for this noise model we have

$$p_{Y_{\theta}|X=x}(y) = \mathcal{N}(y|F(x), a^2 I_n + b^2 \text{diag}(F(x))).$$

For simplicity, we assume that we have only $M = 1$ samples. The case $M > 1$ can be reduced to this case by considering M copies of y_i . In our EM algorithm for (9) we use $V_{\theta} \sim \mathcal{N}(0, 1)$ as hidden variable, which corresponds to the ‘additive part’ of the noise.

3.2.1. Inner E-step

We have to compute the conditional distribution $P_{V_{\theta}|(X, Y_{\theta})=(x, y)}$. Using Bayes’ formula, we obtain

$$\begin{aligned} \log(p_{V_{\theta}|(X, Y_{\theta})=(x, y)}(v)) &\propto \log(p_{Y_{\theta}|V_{\theta}=v, X=x}(y)) + \log(p_{V_{\theta}|X=x}(v)) \\ &\propto \sum_{j=1}^n -\frac{(y_j - F_j(x) - v_j)^2}{2b^2 F_j(x)^2} - \frac{v_j^2}{2a^2} \\ &\propto \log\left(\mathcal{N}\left(v \left| \frac{a^2(y-F(x))}{a^2+b^2F(x)^2}, \text{diag}\left(\frac{a^2 b^2 F(x)^2}{a^2+b^2F(x)^2}\right)\right)\right), \end{aligned}$$

where the quotients in the last line are understood componentwise and ‘ \propto ’ indicates, that we have equality up to an additive constant independent of v . Consequently, the conditional distribution $P_{V_{\theta}|(X, Y_{\theta})=(x, y)}$ is given by

$$P_{V_{\theta}|(X, Y_{\theta})=(x, y)} = \mathcal{N}\left(\frac{a^2(y-F(x))}{a^2+b^2F(x)^2}, \text{diag}\left(\frac{a^2 b^2 F(x)^2}{a^2+b^2F(x)^2}\right)\right).$$

3.2.2. Inner M-step

We will just outline the final result. The quite technical proof is deferred to appendix A, where the M-step can be rewritten as

$$\theta^{(r+1)} = \underset{\theta=(a, b) \in \mathbb{R}_{\geq 0}^2}{\text{arg max}} A_1(b) + A_2(a),$$

where

$$A_1(b) = \frac{1}{2b^2} c_1^{(r)} - n + \log(b) + \text{const}, \quad A_2(a) = \frac{1}{2a^2} c_2^{(r)} - n \log(a) + \text{const},$$

with

$$\begin{aligned} c_1^{(r)} &= -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^n \frac{(y_{ij} - F_j(x_i))^2 (b^{(r)})^4 F_j(x_i)^2}{\left((a^{(r)})^2 + (b^{(r)} F_j(x_i))^2\right)^2} + \frac{(a^{(r)} b^{(r)})^2}{(a^{(r)})^2 + (b^{(r)} F_j(x_i))^2} \\ c_2^{(r)} &= -\frac{1}{N} \sum_{i=1}^N \left(\sum_{j=1}^n \left(\frac{(a^{(r)})^2 (y_{ij} - F_j(x_i))}{(a^{(r)})^2 + (b^{(r)})^2 F_j(x_i)^2} \right)^2 + \frac{(a^{(r)} b^{(r)} F_j(x_i))^2}{(a^{(r)})^2 + (b^{(r)} F_j(x_i))^2} \right), \end{aligned}$$

and $y_i = (y_{i1}, \dots, y_{in})$ and $F = (F_1, \dots, F_n): \mathbb{R}^d \rightarrow \mathbb{R}^n$. By setting the derivatives of A_1 and A_2 to zero, this is equivalent to

$$a^{(r+1)} = -\frac{c_2^{(r)}}{n}, \quad b^{(r+1)} = -\frac{c_1^{(r)}}{n},$$

which are the update rules we will use.

Algorithm 1. EM algorithm mixed noise estimation via CNFs.

Input: $y_1, \dots, y_N \in \mathbb{R}^n$, conditional normalizing flow \mathcal{T} and initial estimate $a^{(0)}, b^{(0)}$, number of x in total $K = 2000 \geq N$.
for $r = 0, 1, \dots, R$ **do**

E-Step:

for $p = 0, 1, \dots, P$ **do** update \mathcal{T} according to

$$\arg \min_{\phi} \mathbb{E}_{y \sim P_{Y_\theta}} \left[\text{KL} \left(p_{X|Y_\theta=y}, p_{\mathcal{T}_\phi(y, \cdot)} \# P_Z \right) \right] \quad (\text{reverse KL})$$

or

$$\arg \min_{\phi} \mathbb{E}_{y \sim P_{Y_\theta}} \left[\text{KL} \left(p_{\mathcal{T}_\phi(y, \cdot)} \# P_Z, p_{X|Y_\theta=y} \right) \right] \quad (\text{forward KL})$$

via a gradient descent method where $\theta = (a^r, b^r)$.

end for

Repeat $\tilde{y} = (y_1, \dots, y_1, \dots, y_N, \dots, y_N) \in \mathbb{R}^{Kn}$.

Sample $x_i \sim \mathcal{T}(\tilde{y}_i, z_i)$ for standard normal z_i .

M-Step:

for $l = 0, 1, \dots, L$ **do**

$$c_1^{(r,l)} = -\frac{1}{K} \sum_{i=1}^K \sum_{j=1}^n \frac{(\tilde{y}_{ij} - F_j(x_i))^2 (b^{(r)})^4 F_j(x_i)^2}{((a^{(r)})^2 + (b^{(r)} F_j(x_i))^2)^2} + \frac{(a^{(r)} b^{(r)})^2}{(a^{(r)})^2 + (b^{(r)} F_j(x_i))^2}$$

$$c_2^{(r,l)} = -\frac{1}{K} \sum_{i=1}^K \left(\sum_{j=1}^n \left(\frac{(a^{(r)})^2 (\tilde{y}_{ij} - F_j(x_i))}{(a^{(r)})^2 + (b^{(r)} F_j(x_i))^2} \right)^2 + \frac{(a^{(r)} b^{(r)} F_j(x_i))^2}{(a^{(r)})^2 + (b^{(r)} F_j(x_i))^2} \right).$$

$$\text{Update } (a^{(r,l)})^2 = -\frac{c_1^{(r,l)}}{n}, \quad (b^{(r,l)})^2 = -\frac{c_2^{(r,l)}}{n}$$

end for

Set $(a^{(r+1)})^2 = (a^{(r,L)})^2, \quad (b^{(r+1)})^2 = (b^{(r,L)})^2$

end for

3.3. Resulting algorithm

The summary of the two nested algorithms can be seen in algorithm 1. Here, both the E-steps and the M-steps are not run for 1 iteration, but several. In particular the analytical M-step is cheap, and therefore it is intuitive to make use of this. For the E-step we take usually 10 steps to perform posterior updates. The initialization of a, b are done in such a way that we approximate the posterior distribution ‘from above’. This is important so that the observed measurements are included in the distribution P_{Y_θ} , which is similar to the logdet schedule proposed in [63]. It indeed can be shown that making the logdet term larger corresponds to scaling the noise higher for additive Gaussian noise, which makes the estimated distributions broader and therefore prevents mode missing, although we still found that this does not solve the problem completely.

4. Experiments

We will benchmark our algorithm on two problems from nano-optics, the first one being low-dimensional and the second one harder and more recent. The first was introduced in [29, 30] and the second one is part of a current research project. The goal of this is to learn both a reasonable posterior reconstruction as well as the error parameters a, b jointly. To showcase the advantages of making the models conditional we also vary the number of measurements and hope that more measurements lead to better reconstructions.

Generally, we use the PyTorch framework [54] and use FrEia package for implementation of the conditional NFs [6]. The code is available on GitHub⁴. We train our models using the Adam optimizer [38] and fix some hyperparameter choices across the experiments. In particular, we only use a learning rate of 1×10^{-3} , $P = 10$, set $K = 2000$, $R = 5000$ in 1 and $L = 20$. The choice of K is in particular constant no matter how many measurements N are used. This allows us to compare whether the information of many measurements is beneficial for the estimation of a and b . However, these hyperparameters were not optimized in a grid search and therefore it is likely that one can improve the performance. We generate synthetic measurements via surrogate forward operators with known noise levels $a_{\text{true}}, b_{\text{true}}$, similar as in [3, section 3]. In both these experiments, we take these surrogate neural networks as forward operators. The extension to real world measurements and the relation to the true PDE inverse problem is left for future work. This allows us, given some noise parameters, to sample (x, y) data on the fly.

⁴ <https://github.com/PaulLyonel/ConditionalDeepGEM>.

Then we apply our proposed algorithms to learn a and b as well as the posterior reconstructions. Then we are able to compare the models and error parameters on two metrics. The metrics and models evaluated are summarized below.

4.1. Models evaluated

We will evaluate two EM-based models, one is the conditional version of the DeepGEM method [24] which we combined with our M-step. Note that amounts to using the reverse KL divergence in algorithm 1. However we propose to use the forward KL divergence which we call conditional forward DeepGEM. To see the general comparison of EM algorithms, we also implement a grid search over (a, b) and save the ‘best’ model of this. The grid search is still possible since we are searching over a two dimensional space, but becomes quickly infeasible for higher dimensional noise models. We will call this grid conditional NF and also evaluate its forward and reverse KL version.

4.2. Metrics

We are going to benchmark the models using the following two metrics.

- *Distance to true a and b* : we will consider synthetic data, where the a, b the observations were generated with, are known. This metric is given by

$$D((a, b), (a_{\text{true}}, b_{\text{true}})) = \frac{|a - a_{\text{true}}|}{a_{\text{true}}} + \frac{|b - b_{\text{true}}|}{b_{\text{true}}}.$$

However, this is a terrible metric as there can be other combinations of a and b which explain the observations equally well. However, we hope that for sufficient observations we will converge to the true a and b .

- *ELBO*: from lemma 1 we see that maximizing the ELBO \mathcal{F} leads to minimizing the KL distance to the true posteriors as well as maximizing the likelihoods of the observations under the estimated error parameters a, b . This is a good proxy, as both the likelihood of the observations as well as the KL distance to the posteriors are intractable in high dimensions.

By the above discussion, we also obtain a suitable *model selection* criterion. We train all the models for the same amount of steps, but we validate it after every EM-step according to the ELBO for the measurements. We then load the best model for every run and evaluate our metrics.

4.3. Scatterometry

For chip manufacturing the control of nanopattern in the lithography process is essential and non-destructive measurement methods with high throughput are desirable. In addition to standard scanning electron microscopy (low throughput, destructive) scatterometry is gaining importance. Scatterometry is a non-destructive optical measurement technique for assessing lithography’s periodic nanostructures’ critical dimensions [37]. In this measuring method, nanostructured periodic surfaces are illuminated with light and refraction patterns are detected. From these patterns geometry parameters are reconstructed by solving an inverse problem. According to equation (1) observations are given by the refraction patterns, the forward operator is determined by time-harmonic Maxwell’s equations and the noise is given by the instrument as well as the model error.

In the following, we consider two examples to demonstrate the performance of the developed algorithm for applications in nanometrology of chip production. The first example considers a typical photomask for extreme ultra violet light (EUV) and the second a line grating.

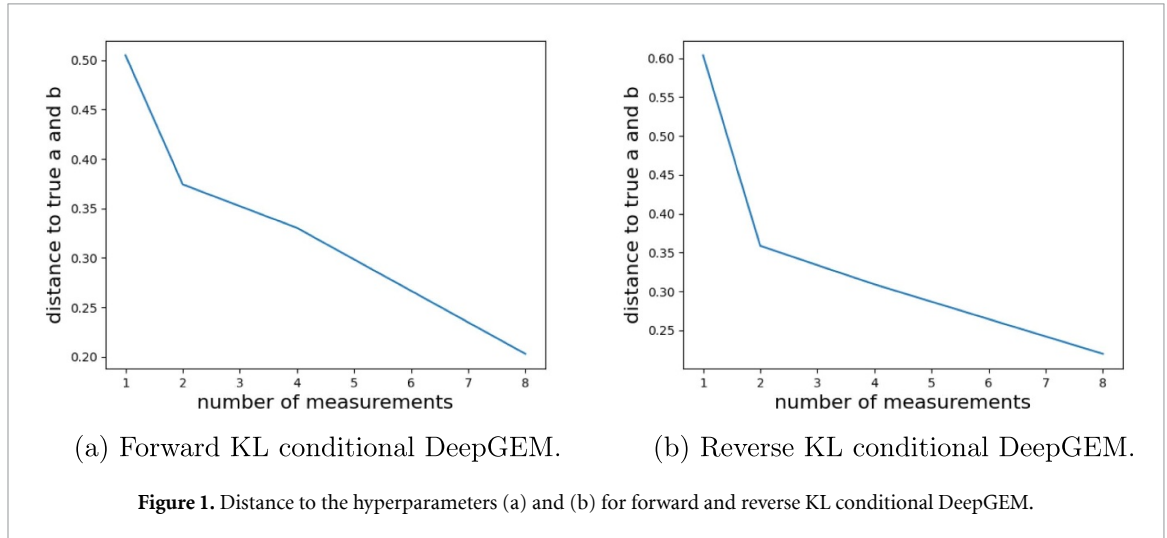
4.3.1. Photo mask

The EUV-photomask considered here consists of periodic absorber lines, capping layers and a multilayer stack functioning as a mirror for 13.4 nm wavelength waves (EUV range). Key geometry parameters include the line width, height and the angle of the sidewall (3 parameters). The refraction patterns comprise 23 intensities (maxima of the refractive orders) and the measurement/model noise is assumed to be distributed according to our mixed noise model.

The problem has x -dimension 3 and y -dimension 23 and therefore is well-suited for first experiments. Furthermore, by [30] it is known that the posterior is indeed multimodal. The prior is chosen uniformly in $[-1, 1]$ and its density is approximated like in [26] for the reverse KL. For the example we train both the conditional DeepGEM as well as the conditional forward DeepGEM using data from the finite element method (FEM) based forward model [30], which is approximated by a surrogate neural network. The forward DeepGEM is a bit quicker to train. The true a and b used to generate simulated signals of the

Table 1. Distance of estimated a and b to the true ones over 10 runs. The best values are indicated in bold.

Number measurements	1	2	4	8
Forward condDeepGEM	0.50	0.37	0.33	0.20
Reverse condDeepGEM	0.60	0.36	0.31	0.22
Forward gridCNF	0.65	0.41	0.44	0.32
Reverse gridCNF	0.53	0.34	0.31	0.42


**Table 2.** ELBO of the algorithms over 10 runs. Calculated for the measurements based on 2000 samples. The best values are indicated in bold.

Number measurements	1	2	4	8
Forward condDeepGEM	81.12	80.55	79.36	79.20
Reverse condDeepGEM	80.92	79.80	78.62	78.59
Forward gridCNF	77.54	76.89	76.33	76.61
Reverse gridCNF	78.99	77.55	76.32	76.50


instrument were set to 0.005 and 0.1 respectively. We benchmark now the four methods, the conditional DeepGem with forward and reverse KL as well as the grid conditional NFs (gridCNF) with forward and reverse KL. For the grids we chose an equispaced grid with 8 points for a and for b . For a this ranged from 0.001 to 0.03 and for b from 0.01 to 0.2. Concerning training time, the forward and reverse conditional DeepGEM were similar with 9 min per run. The grid versions took approximately 13 min to train, where we took 1200 optimizer steps per grid point. Generally, we can see that the grid methods get outperformed by our EM versions, although they take a longer time. From table 1 and figure 1 we can see that the forward KL and the reverse KL have both similar performance in terms of distance to the true a and b , where the forward KL seems to have a slight edge in the case of many measurements. However, in terms of ELBO, we observe in table 2 that the forward KL performs favorably. This is somewhat remarkable, as the reverse KL is the ELBO objective when ignoring the parameters independent of the flow. Considering posterior measures obtained from simulated measurements we realize that the reverse KL does not exactly reproduce the modes in some of the examples, see figure 2 whereas the forward KL performs quite well. The inability of the reverse KL to detect the correct modes of the posterior can indeed explain the better performance of the forward conditional DeepGEM. Both algorithms are improving with more measurements, see figures 1(a) and (b).

4.3.2. Line grating with oxid layer

The second example involves a periodic line grating consisting of a silicon bulk and an oxide layer on top. Similar samples were investigated e.g. in [43]. In addition to the geometry parameters, as used in the previous example, the optical constants (OC) of the materials are assumed to be not accurately known. In practice this is often the case if the material composition was changed due to oxidation and contamination of the sample. So for each material, there are two parameters for the complex refractive index (real and imaginary part) [31], which depend on the material density. Hence we change the OC by varying the densities of the material, i.e. silicon and silicon-oxide. This results in two parameters changing the OC and five parameter describing the geometry of the line grating. The refraction pattern are detected for a single



(a) Posterior reconstruction for forward KL conditional DeepGEM for one simulated measurement. (b) Posterior reconstruction for reverse KL conditional DeepGEM for one simulated measurement.



(c) Posterior reconstruction for forward KL conditional DeepGEM for another simulated measurement. (d) Posterior reconstruction for reverse KL conditional DeepGEM for another simulated measurement.

Figure 2. Posterior reconstructions for different measurements using forward/reverse conditional DeepGEM via one dimensional histograms on the diagonal and two dimensional on the offdiagonal. Ground truth x is depicted by the blue line.

wavelength of the incoming light beam under an angle of incidence of 30° , from the sample plane and a set of seven azimuth angles between 0° and 6° , the sample is rotated by in the plane. In sum we obtained 77 simulated intensities and hence end up with x -dimension seven and y -dimension 77. For simulations the forward model was solved with the software package JCMSuite⁵, based on the FEM which solves a boundary value problem following from the Maxwell's equations [21]. In order to get a strong response for the OC of the oxide layer we used a wavelength of 12.99 nm, right before the absorption edge [4, 31]. For this work we standardized the data from the forward simulation [11] on $[0, 1]$ and chose a uniform prior for the x -data.

Again, we plot two example posterior distributions calculated. The distribution shapes seen in 3 clearly reflect the sensitivity of the forward operator against the line height (parameter 0), the silicon oxide density (parameter 4) and non-sensitive against the layer roughness (parameter 6). The true a and b were set to 0.03 and 0.25 respectively. Again as in the first scatterometry example we can see in tables 3 and 4 that the forward

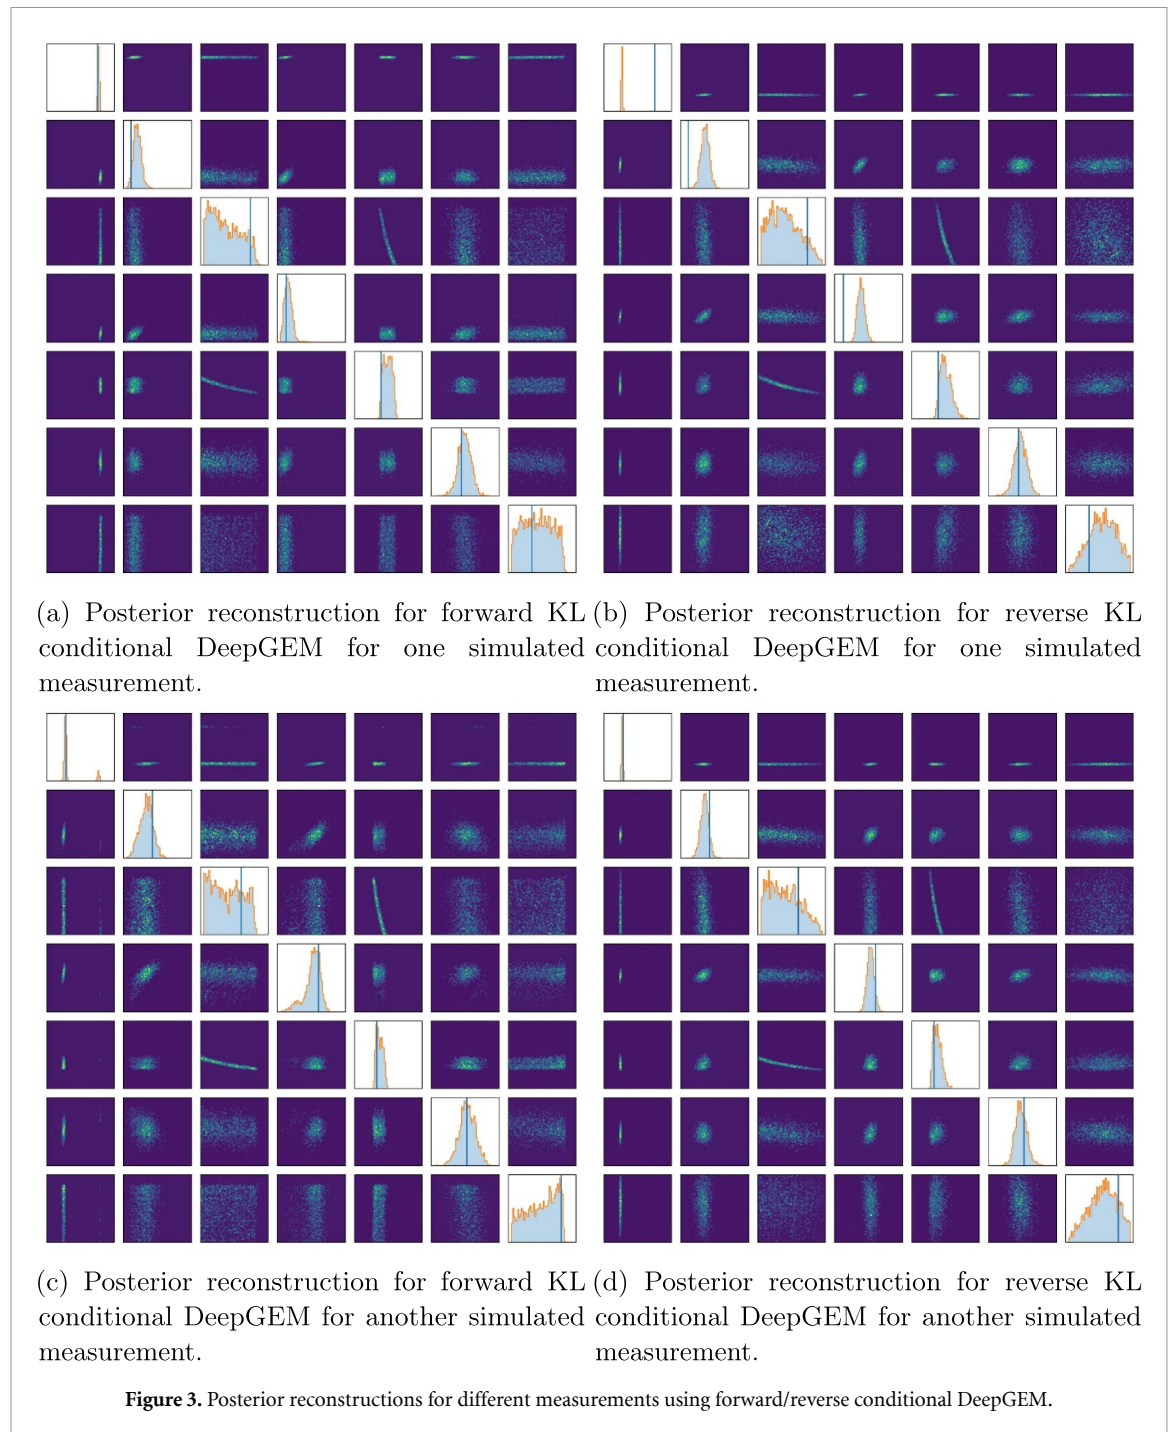
⁵ <https://jcmwave.com/>.

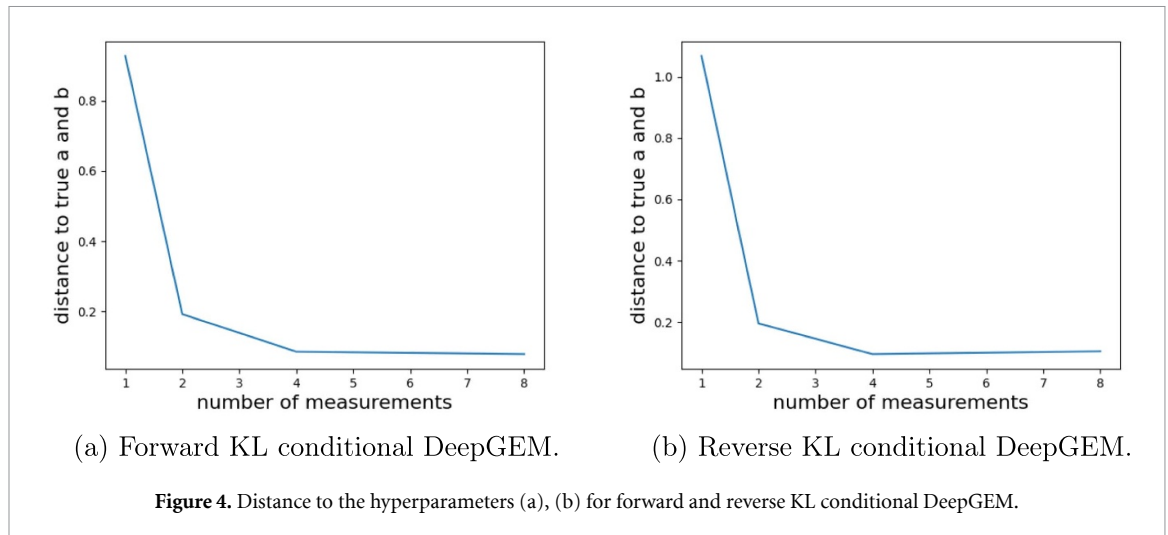
Table 3. Distance of estimated a and b to the true ones over 5 runs. The best values are indicated in bold.

Number measurements	1	2	4	8
Forward	0.93	0.19	0.09	0.08
Reverse	1.07	0.20	0.10	0.11

Table 4. ELBO of the algorithms over 5 runs. Calculated for the measurements based on 10 000 samples. The best values are indicated in bold.

Number measurements	1	2	4	8
Forward	195.9	196.0	190.5	188.1
Reverse	195.8	196.1	189.1	186.9





KL performs a bit better in distance to the true a and b as well as ELBO. Similarly, one can observe that the first x -component, the height can be multimodal, where the reverse KL can indeed miss the mode. This can be observed in figure 3. Similarly, the distance to the true a and b decreases by adding more simulated measurement values, which can be seen in figure 4.

5. Conclusions and limitations

We developed a nested EM algorithms, one for estimating the posterior distribution via a conditional NF and a second one to solve the M -step within the former EM algorithm to estimate the error model parameters. For the special kind of non-additive noise appearing in our applications we derived analytic formulas for the inner E - and M -steps. We showed advantages of using the forward KL for modelling multimodal distributions. The reverse KL often led to mode collapse. However, there has been a plethora of literature tackling this issue of the reverse KL, namely [5, 45, 47, 64]. It would be interesting to compare these approaches to the forward KL. Moreover, we could replace the conditional NF by other methods for posterior sampling like score-based diffusion models [35, 60], conditional GANs [49] or posterior MMD flows [25]. Furthermore, we chose synthetic a_{true} and b_{true} . One of the next steps is to test these approaches on real world measurements. Even if the novel algorithm was applied to two specific real world experiments, it may have an impact to a wide range of applications where indirect measurements are involved. The extension of the algorithm to other noise distributions than Gaussian is analogous. An advantage over standard approaches like Markov-Chain Monte Carlo methods is the fact that once the network has been trained, further similar measurements can be evaluated very quickly. This benefit opens the possibility of scatterometry and similar measurement techniques for real time applications, e.g. important for process control. In terms of *limitations*, it would be also interesting to test the algorithm on other inverse problems. Intuitively, we believe that the scatterometric inverse problem is particularly well-suited for these estimations since the observed $f(X)$ -data is living on a low-dimensional manifold in a nominal high-dimensional space. One can indeed easily think of an inverse problem, where recovering noise parameters is much harder if the observed data already lies in the full space.

Data availability statement

The code and data in the form of surrogate networks that support the findings of this study are openly available. This means, the surrogates and the code for running and reproducing the experiments is available under <https://github.com/PaulLyonel/ConditionalDeepGEM>. For the second photo mask problem we have training data from the forward simulation/simulation of the measurement under <https://doi.org/10.5281/zenodo.10580011>.

Acknowledgments

P H acknowledges support from the DFG within the SPP 2298 ‘Theoretical Foundations of Deep Learning’ (STE 571/17-1). J H acknowledges funding by the EPSRC programme grant ‘The Mathematics of Deep Learning’ with reference EP/V026259/1. M C and S H acknowledges the support of the EMPIR Project

22IND04-ATMOC. This Project (20IND04 ATMOC) has received funding from the EMPIR programme cofinanced by the Participating States and from the European Union’s Horizon 2020 research and innovation programme.

Appendix A. Derivation of the inner M-step

For the simplicity of the notation, we use the abbreviation

$$Q_i = P_{V_{\theta^{(r)}} | (X_i, Y_{\theta^{(r)}}) = (x_i, y_i)} = \mathcal{N} \left(\frac{(a^{(r)})^2 (y_i - F(x_i))}{(a^{(r)})^2 + (b^{(r)})^2 F(x_i)^2}, \text{diag} \left(\frac{(a^{(r)})^2 (b^{(r)})^2 F(x_i)^2}{(a^{(r)})^2 + (b^{(r)})^2 F(x_i)^2} \right) \right). \quad (\text{A.1})$$

Using the decomposition (6) of the ELBO, and noting that the second summand within (6) does not depend on the parameters $\theta = (a, b)$, we obtain that the optimization problem (4) is equivalent to

$$\begin{aligned} \theta^{(r+1)} &= \arg \max_{\theta \in \Theta} \sum_{i=1}^N \frac{1}{N} \mathbb{E}_{v \sim Q_i} [\log (p_{V_{\theta}, X, Y_{\theta}} (v, x_i, y_i))] \\ &= \arg \max_{\theta \in \Theta} \sum_{i=1}^N \frac{1}{N} \mathbb{E}_{v \sim Q_i} [\log (p_{V_{\theta}, Y_{\theta} | X=x_i} (v, y_i))] + \log (p_X (x_i)) \\ &= \arg \max_{\theta \in \Theta} \sum_{i=1}^N \frac{1}{N} \mathbb{E}_{v \sim Q_i} [\log (p_{V_{\theta}, Y_{\theta} | X=x_i} (v, y_i))]. \end{aligned} \quad (\text{A.2})$$

Now, the objective function reads as

$$\begin{aligned} &\sum_{i=1}^N \frac{1}{N} \mathbb{E}_{v \sim Q_i} [\log (p_{V_{\theta}, Y_{\theta} | X=x_i} (v, y_i))] \\ &= \sum_{i=1}^N \frac{1}{N} \mathbb{E}_{v \sim Q_i} [\log (p_{Y_{\theta} | V_{\theta}=v, X=x_i} (y_i)) + \log (p_{V_{\theta} | X=x_i} (v))]. \end{aligned}$$

As it holds by definition that $p_{Y_{\theta} | V_{\theta}=v, X=x_i} (y_i) = \mathcal{N}(y_i | F(x_i) + v, b^2 \text{diag}(F(x_i))^2)$ and $p_{V_{\theta} | X=x_i} (v) = p_{V_{\theta}} (v) = \mathcal{N}(v | 0, a^2 I_n)$, this is equal to

$$\sum_{i=1}^N \frac{1}{N} \mathbb{E}_{v \sim Q_i} [\log (\mathcal{N}(y_i | F(x_i) + v, b^2 \text{diag}(F(x_i))^2) + \log (\mathcal{N}(v | 0, a^2 I))) = A_1(b) + A_2(a),$$

where (leaving out constants with respect to a or b)

$$\begin{aligned} A_1(b) &= \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{v \sim Q_i} [\log (\mathcal{N}(y_i | F(x_i) + v, b^2 \text{diag}(F(x_i))^2))] \\ &= -\frac{1}{2Nb^2} \sum_{i=1}^N \sum_{j=1}^n \frac{1}{F_j(x_i)^2} \left((y_{ij} - F_j(x_i))^2 - 2 \mathbb{E}_{v \sim Q_i} [v_j] (y_{ij} - F_j(x_i)) + \mathbb{E}_{v \sim Q_i} [v_j^2] \right) - \frac{n}{2} \log(b^2), \end{aligned}$$

and

$$A_2(a) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{v \sim Q_i} [\log (\mathcal{N}(v | 0, a^2 I))] = -\frac{1}{2a^2 N} \sum_{i=1}^N \sum_{j=1}^n \mathbb{E}_{v \sim Q_i} [v_j^2] - \frac{n}{2} \log(a^2).$$

Now, by (A.1), the expressions $\mathbb{E}_{v \sim Q_i} [v_j]$ and $\mathbb{E}_{v \sim Q_i} [v_j^2]$ are the first and second moment of certain normal distributions, such that

$$\mathbb{E}_{v \sim Q_i} [v_j] = \frac{(a^{(r)})^2 (y_{ij} - F_j(x_i))}{(a^{(r)})^2 + (b^{(r)})^2 F_j(x_i)^2}$$

and

$$\mathbb{E}_{v \sim Q_i} [v_j^2] = \left(\frac{(a^{(r)})^2 (y_{ij} - F_j(x_i))}{(a^{(r)})^2 + (b^{(r)})^2 F_j(x_i)^2} \right)^2 + \frac{(a^{(r)})^2 (b^{(r)})^2 F_j(x_i)^2}{(a^{(r)})^2 + (b^{(r)})^2 F_j(x_i)^2}.$$

Putting everything together, we obtain that (A.2) is equivalent to

$$\theta^{(r+1)} = \underset{(a,b) \in \mathbb{R}_{\geq 0}^2}{\operatorname{argmax}} A_1(b) + A_2(a)$$

with

$$A_1(b) = \frac{1}{2b^2} c_1^{(r)} - n \log(b) + \text{const}, \quad A_2(a) = \frac{1}{2a^2} c_2^{(r)} - n \log(a) + \text{const},$$

where const denotes an unspecified constant independent of a and b . Further, the $c_i^{(r)}$ are given by

$$c_2^{(r)} = -\frac{1}{N} \sum_{i=1}^N \left(\sum_{j=1}^n \left(\frac{(a^{(r)})^2 (y_{ij} - F_j(x_i))}{(a^{(r)})^2 + (b^{(r)})^2 F_j(x_i)^2} \right)^2 + \frac{(a^{(r)} b^{(r)} F_j(x_i))^2}{(a^{(r)})^2 + (b^{(r)} F_j(x_i))^2} \right),$$

and

$$c_1^{(r)} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^n \frac{1}{F_j(x_i)^2} \left((y_{ij} - F_j(x_i))^2 - 2 \frac{(a^{(r)})^2 (y_{ij} - F_j(x_i))^2}{(a^{(r)})^2 + (b^{(r)})^2 F_j(x_i)^2} + \left(\frac{(a^{(r)})^2 (y_{ij} - F_j(x_i))}{(a^{(r)})^2 + (b^{(r)})^2 F_j(x_i)^2} \right)^2 + \frac{(a^{(r)} b^{(r)} F_j(x_i))^2}{(a^{(r)})^2 + (b^{(r)} F_j(x_i))^2} \right).$$

Bringing the first three terms onto one denominator, this can be simplified to

$$c_1^{(r)} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^n \frac{(y_{ij} - F_j(x_i))^2 (b^{(r)})^4 F_j(x_i)^2}{\left((a^{(r)})^2 + (b^{(r)} F_j(x_i))^2 \right)^2} + \frac{(a^{(r)} b^{(r)})^2}{(a^{(r)})^2 + (b^{(r)} F_j(x_i))^2}.$$

Note that by definition $c_i^{(r)}$, $i = 1, 2$ are non-positive. Thus, $a^{(r+1)}$ and $b^{(r+1)}$ are given by

$$a^{(r+1)} = \underset{a \geq 0}{\operatorname{argmax}} A_2(a), \quad b^{(r+1)} = \underset{b \geq 0}{\operatorname{argmax}} A_1(b).$$

By setting the derivatives of A_1 and A_2 to zero, this is equivalent to

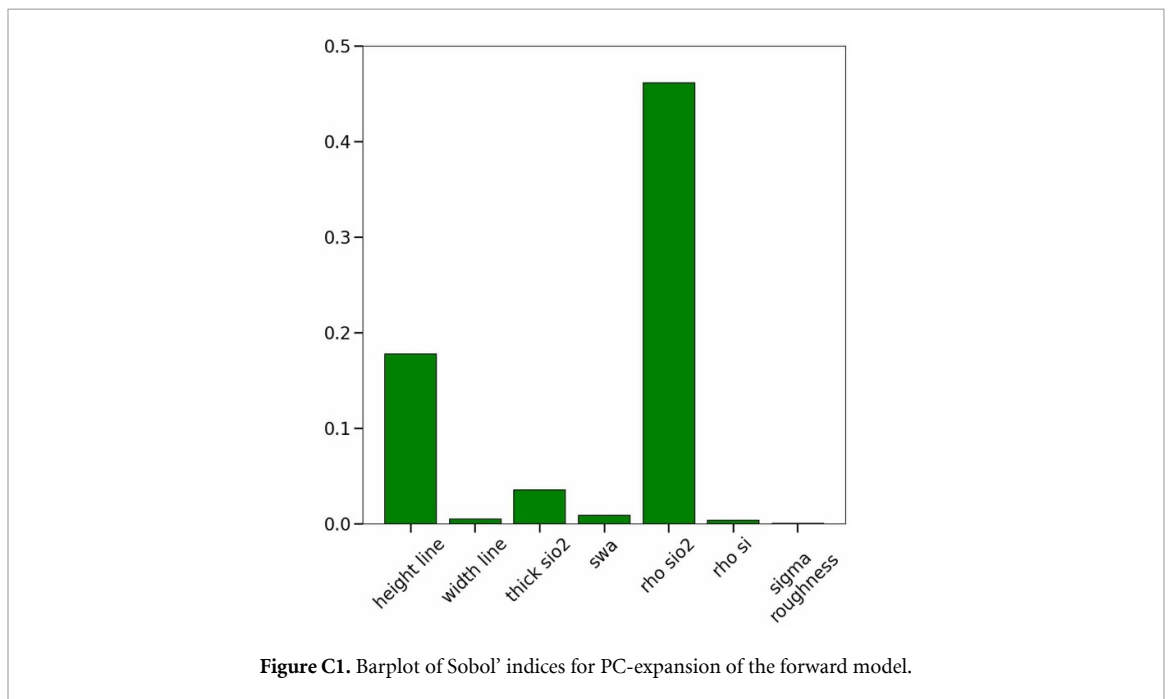
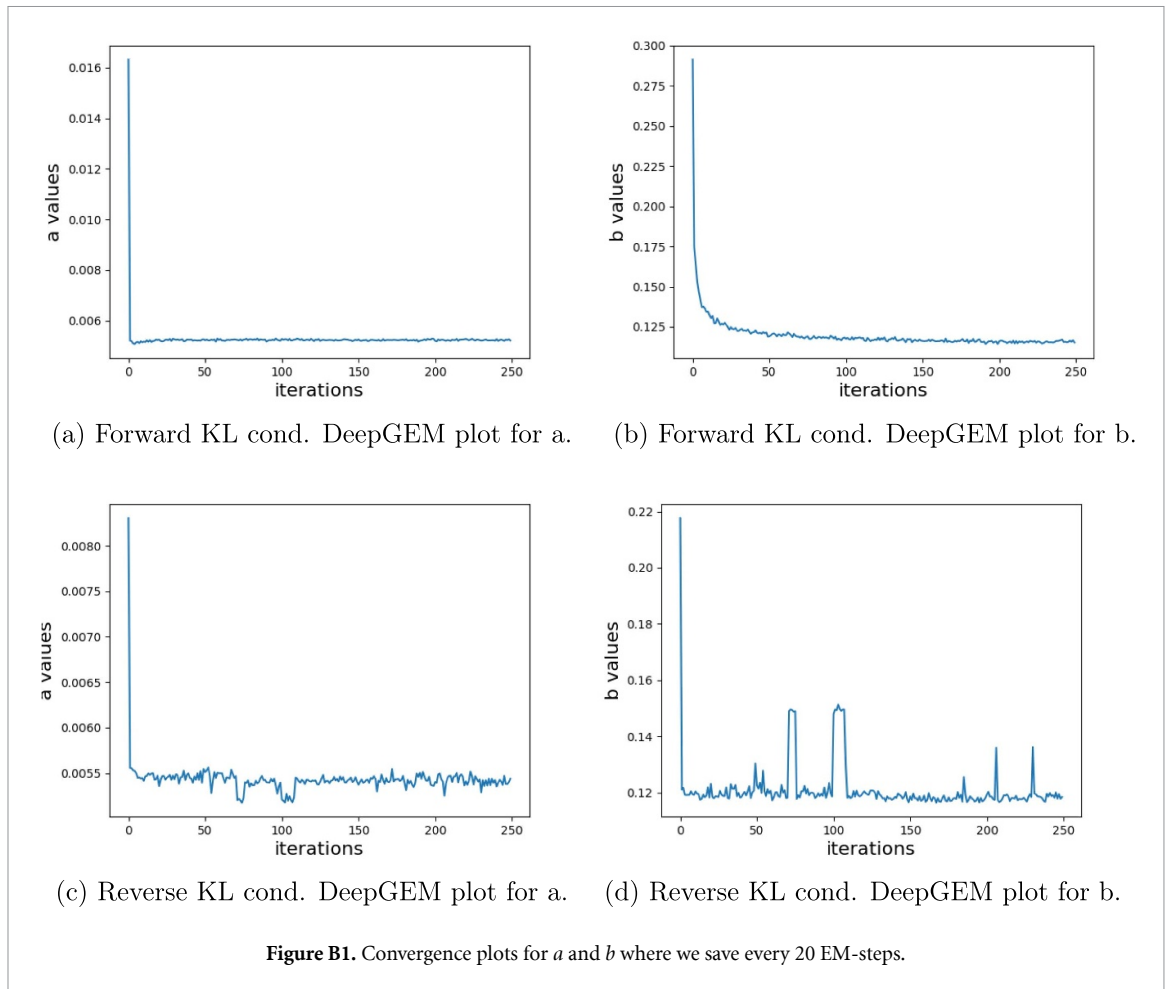
$$a^{(r+1)} = -\frac{c_2^{(r)}}{n}, \quad b^{(r+1)} = -\frac{c_1^{(r)}}{n}.$$

Appendix B. Convergence plots of a and b

Here we showcase the convergence of the parameters a and b in the first scatterometry photo mask example. For one particular run the convergence of a and b is shown, see figure B1. Note that both methods seem to converge to the true a and b values, however the reverse KL trains more unstably, which might be tackable with more careful hyperparameter selection or other stabilization techniques.


Appendix C. Sensitivity analysis of the forward model

To verify our results for the reconstruction of the line grating in 4.3.2, we make a sensitivity analysis of the corresponding forward model. For the sensitivity analysis we have to determine the Sobol' indices, which are coefficients of a decomposition of the forward model and describe the impact of each parameter combination on the forward model. Those indices are normalized to $[0, 1]$ and sum up to 1 [61, 62]. Making an approximation of the forward model in a polynomial basis using Polynomial Chaos (PC) [21, 65] makes it very easy to calculate the Sobol' indices. The indices in figure C1 come from a PC-approximation with a relative L_2 -error of about 0.076 and show the dependence on each single parameter. It is clearly seen, that the height of the grating line and the density of the oxide layer and hence the OC for silicon-oxide has a huge impact on the forward model. In general the sensitivity analysis fits very well to the reconstruction of the line grating, since in figure 3 the distributions for parameter 0 and 4 are very sharp defined, while those which are very broad distributed also show a low impact on the forward model. For the sensitivity analysis the source software tool PyThia was used [28].



ORCID iDs

Paul Hagemann  <https://orcid.org/0009-0006-9679-5654>

Johannes Hertrich  <https://orcid.org/0000-0003-4433-8604>

References

- [1] Abdelhamed A, Brubaker M A and Brown M S 2019 Noise flow: noise modeling with conditional normalizing flows *IEEE/CVF Int. Conf. on Computer Vision* pp 3165–73
- [2] Altekrüger F and Hertrich J 2023 WPPNets and WPPFlows: the power of wasserstein patch priors for superresolution *SIAM J. Imaging Sci.* **16** 1033–67
- [3] Andrle A, Farchmin N, Hagemann P, Heidenreich S, Soltwisch V and Steidl G 2021 Invertible neural networks versus MCMC for posterior reconstruction in grazing incidence x-ray fluorescence *Int. Conf. on Scale Space and Variational Methods in Computer Vision* (Springer) pp 528–39
- [4] Andrle A, Hönicke P, Vinson J, Quintanilha R, Saadeh Q, Heidenreich S, Scholze F and Soltwisch V 2021 The anisotropy in the optical constants of quartz crystals for soft x-rays *J. Appl. Crystallogr.* **54** 402–8
- [5] Arbel M, Matthews A G D G and Doucet A 2021 Annealed flow transport monte carlo *Proc. 38th Int. Conf. on Machine Learning (Proc. Machine Learning Research)*
- [6] Ardizzone L, Bungert T, Draxler F, Köthe U, Kruse J, Schmier R and Sorrenson P 2018–2022 Framework for easily invertible architectures (FrEIA)
- [7] Ardizzone L, Lüth C, Kruse J, Rother C and Köthe U 2019 Guided image generation with conditional invertible neural networks (arXiv:1907.02392)
- [8] Behrmann J, Grathwohl W, Chen R T, Duvenaud D and Jacobsen J-H 2019 Invertible residual networks *Int. Conf. on Machine Learning* pp 573–82
- [9] Bishop C M 2006 *Pattern Recognition and Machine Learning (Information Science and Statistics)* (Springer)
- [10] Bishop C M and Nasrabadi N M 2006 *Pattern Recognition and Machine Learning* (Springer)
- [11] Casfor M 2024 Forward simulation of scatterometry for a nanostructure (available at: <https://zenodo.org/doi/10.5281/zenodo.10580011>)
- [12] Chen R T Q, Behrmann J, Duvenaud D K and Jacobsen J-H 2019 Residual flows for invertible generative modeling *Advances in Neural Information Processing Systems* vol 32 (Curran Associates, Inc.)
- [13] Chrétien S and Hero A O 2000 Kullback proximal algorithms for maximum-likelihood estimation *IEEE Trans. Inf. Theory* **46** 1800–10
- [14] Chrétien S and Hero A O 2008 On EM algorithms and their proximal generalizations *ESAIM: Probabil. Stat.* **12** 308–26
- [15] De Cao N, Titov I and Aziz W 2019 Block neural autoregressive flow (arXiv:1904.04676)
- [16] Dempster A P, Laird N M and Rubin D B 1977 Maximum likelihood from incomplete data via the EM algorithm *J. R. Stat. Soc. B* **39** 1–38
- [17] Denker A, Schmidt M, Leuschner J and Maass P 2021 Conditional invertible neural networks for medical imaging *J. Imaging* **7** 243
- [18] Dinh L, Sohl-Dickstein J and Bengio S 2017 Density estimation using real NVP *Int. Conf. on Learning Representations*
- [19] Dunlop M M 2019 Multiplicative noise in bayesian inverse problems: well-posedness and consistency of map estimators (arXiv:1910.14632)
- [20] Durkan C, Bekasov A, Murray I and Papamakarios G 2019 Neural spline flows *Advances in Neural Information Processing Systems*
- [21] Farchmin N, Hammerschmidt M, Schneider P-I, Wurm M, Bodermann B, Bär M and Heidenreich S 2019 An efficient approach to global sensitivity analysis and parameter estimation for line gratings *Modeling Aspects in Optical Metrology VII* p 15
- [22] Foi A 2009 Clipped noisy images: heteroskedastic modeling and practical denoising *Signal Process.* **89** 2609–29
- [23] Foulley J-L 1997 ECM approaches to heteroskedastic mixed models with constant variance ratios *Genet. Sel. Evol.* **29** 297–318
- [24] Gao A F, Castillo J C, Yue Y, Ross Z E and Bouman K L 2021 DeepGEM: generalized expectation-maximization for blind inversion *Advances in Neural Information Processing Systems* vol 35
- [25] Hagemann P, Hertrich J, Altekrüger F, Beinert R, Chemseddine J and Steidl G 2024 Posterior sampling based on gradient flows of the MMD with negative distance kernel *ICLR 2024* (arXiv:2310.03054)
- [26] Hagemann P, Hertrich J and Steidl G 2022 Stochastic normalizing flows for inverse problems: a Markov chains viewpoint *SIAM/ASA J. Uncertain. Quantification* **10** 1162–90
- [27] Hagemann P, Hertrich J and Steidl G 2023 *Generalized Normalizing Flows via Markov Chains* (Cambridge University Press)
- [28] Hagemann N and Heidenreich S 2023 Pythia: a Python package for uncertainty quantification based on non-intrusive polynomial chaos expansions *J. Open Source Softw.* **8** 5489
- [29] Heidenreich S, Gross H and Bär M 2015 Bayesian approach to the statistical inverse problem of scatterometry: comparison of three surrogate models *Int. J. Uncertain. Quantification* **5** 51–26
- [30] Heidenreich S, Gross H and Bär M 2018 Bayesian approach to determine critical dimensions from scatterometric measurements *Metrologia* **55** S201
- [31] Henke B, Gullikson E and Davis J 1993 x-ray interactions: photoabsorption, scattering, transmission and reflection at $e = 50$ –30, 000 eV, $z = 1$ –92 *At. Data Nucl. Data Tables* **54** 181–342
- [32] Herrero A F, Pflüger M, Puls J, Scholze F and Soltwisch V 2021 Uncertainties in the reconstruction of nanostructures in euv scatterometry and grazing incidence small-angle x-ray scattering *Opt. Express* **29** 35580–91
- [33] Hertrich J 2023 Proximal residual flows for Bayesian inverse problems *Int. Conf. on Scale Space and Variational Methods in Computer Vision* (Springer) pp 210–22
- [34] Hirschberger F, Forster D and Lucke J 2021 A variational EM acceleration for efficient clustering at very large scales *IEEE Trans. on Pattern Analysis and Machine Intelligence*
- [35] Ho J, Jain A and Abbeel P 2020 Denoising diffusion probabilistic models *Advances in Neural Information Processing Systems* vol 33 pp 6840–51
- [36] Huang C-W, Krueger D, Lacoste A and Courville A 2018 Neural autoregressive flows *Int. Conf. on Machine Learning* pp 2078–87
- [37] Huang H-T and Terry F L 2004 Spectroscopic ellipsometry and reflectometry from gratings (scatterometry) for critical dimension measurement and in situ, real-time process monitoring *Thin Solid Films* **455** 828–36
- [38] Kingma D P and Ba J 2017 Adam: a method for stochastic optimization (arXiv:1412.6980)
- [39] Kingma D P and Dhariwal P 2018 Glow: generative flow with invertible 1×1 convolutions (arXiv:1807.03039)
- [40] Kruse J, Detommaso G, Scheichl R and Köthe U 2020 HINT: hierarchical invertible neural transport for density estimation and Bayesian inference (arXiv:1905.10687)
- [41] Laroche C, Almansa A and Coupété E 2024 Fast diffusion EM: a diffusion model for blind inverse problems with application to deconvolution *Proc. IEEE/CVF Winter Conf. on Applications of Computer Vision (WACV)* pp 5271–81

- [42] Liu C and Rubin D B 1994 The ECME algorithm: a simple extension of EM and ECM with faster monotone convergence *Biometrika* **81** 633–48
- [43] Lohr L M, Ciesielski R, Glabisch S, Schröder S, Brose S and Soltwisch V 2023 Nanoscale grating characterization using euclidean scatterometry and soft x-ray scattering with plasma and synchrotron radiation *Appl. Opt.* **62** 117–32
- [44] Lyu J, Chen Z, Feng C, Cun W, Zhu S, Geng Y, Xu Z and Yongwei C 2022 Para-CFlows: C^k -universal diffeomorphism approximators as superior neural surrogates *Advances in Neural Information Processing Systems* ed A H Oh, A Agarwal, D Belgrave and K Cho
- [45] Máté B and Fleuret F 2023 Learning interpolations between boltzmann densities *Trans. Mach. Learn. Res.* (available at: <https://openreview.net/forum?id=TH6YrEcbth>)
- [46] Meng X-L and Rubin D B 1993 Maximum likelihood estimation via the ECM algorithm: A general framework *Biometrika* **80** 267–78
- [47] Midgley L I, Stimper V, Simm G N C, Schölkopf B and Hernández-Lobato J M 2023 Flow annealed importance sampling bootstrap *11th Int. Conf. on Learning Representations*
- [48] Minka T 2005 Divergence measures and message passing *Technical Report* (Microsoft Research)
- [49] Mirza M and Osindero S 2014 Conditional generative adversarial nets (arXiv:1411.1784)
- [50] Mohammad-Djafari A 1993 On the estimation of hyperparameters in bayesian approach of solving inverse problems *1993 IEEE Int. Conf. on Acoustics, Speech and Signal Processing* vol 5 pp 495–8
- [51] Nan Y, Quan Y and Ji H 2020 Variational-EM-based deep learning for noise-blind image deblurring *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*
- [52] Neal R M and Hinton G E 1998 A view of the EM algorithm that justifies incremental, sparse and other variants *Learning in Graphical Models* (Springer) pp 355–68
- [53] Papamakarios G, Pavlakou T and Murray I 2017 Masked autoregressive flow for density estimation *Advances in Neural Information Processing Systems* pp 2338–47
- [54] Paszke A et al 2019 Pytorch: an imperative style, high-performance deep learning library *Advances in Neural Information Processing Systems* vol 32, ed H Wallach, H Larochelle, A Beygelzimer, F d'Alché-Buc, E Fox and R Garnett (Curran Associates, Inc.)
- [55] Radev S T, Schmitt M, Pratz V, Picchini U, Köthe U and Bürkner P-C 2023 Jana: jointly amortized neural approximation of complex Bayesian models *Proc. 39th Conf. on Uncertainty in Artificial Intelligence (Proc. of Machine Learning Research* vol 216), ed R J Evans and I Shpitser (PMLR) pp 1695–706
- [56] Ramos A, Furtado C, Colaço A, Arteiro A, Furtado A, Horas C, Lopes I, Carvalho R and Pereira S 2023 *Proc. 7th ECCOMAS Young Investigators Conf. (ECCOMAS YIC 2023)* (Zenodo)
- [57] Rocke D M and Durbin B 2001 A model for measurement error for gene expression arrays *J. Comput. Biol.* **8** 557–69
- [58] Rocke D M and Lorenzato S 1995 A two-component model for measurement error in analytical chemistry *Technometrics* **37** 176–84
- [59] Saadeh Q et al 2021 Time-frequency analysis assisted determination of ruthenium optical constants in the sub-EUV spectral range 8 nm–23.75 nm *Opt. Express* **29** 40993–1013
- [60] Singh I R, Denker A, Barbano R, Kereta Ž, Jin B, Thielemans K, Maass P and Arridge S 2023 Score-based generative models for PET image reconstruction (arXiv:2308.14190)
- [61] Sobol I 1993 Sensitivity estimates for nonlinear mathematical models *MMCE* **1** 407–14
- [62] Sobol I 2001 Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates *Math. Comput. Simul.* **55** 271–80
- [63] Sun H and Bouman K L 2021 Deep probabilistic imaging: uncertainty quantification and multi-modal solution characterization for computational imaging *Proc. AAAI Conf. on Artificial Intelligence* vol 35 pp 2628–37
- [64] Vaitl L, Nicoli K A, Nakajima S and Kessel P 2022 Gradients should stay on path: better estimators of the reverse- and forward kl divergence for normalizing flows *Mach. Learn.: Sci. Technol.* **3** 045006
- [65] Wiener N 1938 The homogeneous chaos *Am. J. Math.* **60** 897
- [66] Winkler C, Worrall D, Hoogeboom E and Welling M 2023 Learning likelihoods with conditional normalizing flows (arXiv:1912.00042)
- [67] Wu H, Köhler J and Noé F 2020 Stochastic normalizing flows *Advances in Neural Information Processing Systems* vol 33 pp 5933–44