



Citation for published version:

Dorta, G, Vicente, S, Agapito, L, Campbell, N & Simpson, I 2018, 'Structured Uncertainty Prediction Networks' Paper presented at IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2018, 18/06/18 - 22/06/18, .

Publication date:
2018

Document Version
Publisher's PDF, also known as Version of record

[Link to publication](#)

University of Bath

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Structured Uncertainty Prediction Networks

Garoe Dorta^{1,2} Sara Vicente² Lourdes Agapito³ Neill D.F. Campbell¹ Ivor Simpson²
¹University of Bath ²Anthropics Technology Ltd. ³University College London

¹{g.dorta.perez,n.campbell}@bath.ac.uk ²{sara,ivor}@anthropics.com ³l.agapito@cs.ucl.ac.uk

Abstract

This paper is the first work to propose a network to predict a structured uncertainty distribution for a synthesized image. Previous approaches have been mostly limited to predicting diagonal covariance matrices [15]. Our novel model learns to predict a full Gaussian covariance matrix for each reconstruction, which permits efficient sampling and likelihood evaluation.

We demonstrate that our model can accurately reconstruct ground truth correlated residual distributions for synthetic datasets and generate plausible high frequency samples for real face images. We also illustrate the use of these predicted covariances for structure preserving image denoising.

1. Introduction

Deep probabilistic generative models have recently become the most popular tool to synthesize novel data and reconstruct unseen examples from target image distributions. At their heart lies a density estimation problem, which for reconstruction models is commonly solved using a factorized Gaussian likelihood. While attractive for its simplicity, this factorized Gaussian assumption comes at the cost of overly-smoothed predictions, as shown in Fig. 1d. In contrast, this paper introduces the first attempt to train a deep neural network to predict full structured covariance matrices to model the residual distributions of unseen image reconstructions. We postulate that the residuals are highly structured and reflect limitations in model capacity – we therefore propose to estimate the reconstruction uncertainty using a structured Gaussian model, to capture pixel-wise correlations which will in turn improve the sampled reconstructions, as shown in Fig. 1c and 1f.

Learning in generative models involves some form of density estimation, where the model parameters are fitted to generate samples from the target image distribution. This paper concentrates on methods, such as the Variational Autoencoder (VAE) [15, 23], where an explicit representation of the data density is given by the model, however the re-

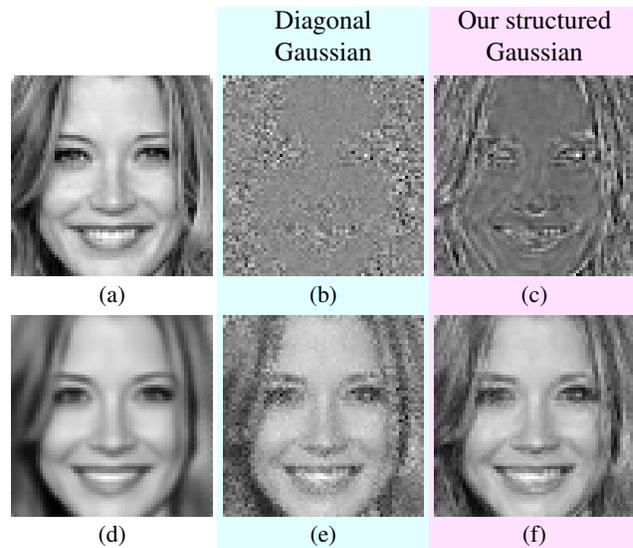


Figure 1: Given an input image (a), a simple diagonal Gaussian likelihood model learns a smooth reconstruction (d), with an unstructured residual sample (b). When (b) is added to (d) it generates the unrealistic image (e), demonstrating a failure to capture residual structure. In contrast, we learn a structured residual model, with residual samples like (c) that, when added to (d), generate a plausible and realistic image (f).

sults could in principle be extended to implicit approaches in future work. In VAE models, a mapping is learned from a latent representation to image space. These models commonly impose some conditions on the latent space, in the form of a prior, and the nature of the residual distribution.

It is common practice to use a Gaussian likelihood, as this provides a simple formulation for maximum likelihood estimation. However, it is well known that samples from factorized Gaussian distributions tend to be overly-smooth. This effect is pronounced in common simplifications of the likelihood, such as the mean squared error, which assumes the errors at all pixels are i.i.d (independent and identically distributed) or a diagonal covariance [15, 4], which allows some local estimation of noise level but maintains the strong and flawed assumption that pixels in the residual are uncor-

related, as shown in Fig. 1b and 1c. These common choices of likelihood mean that if one were to draw a sample from any of these models including the noise term, white noise would be added to the reconstruction, which is unlikely to ever appear realistic, as shown in Fig. 1e. This emphasizes the incoherence of these simplifications, and this is addressed in this paper.

This work proposes to overcome the problems of factorized likelihoods by training a deep neural network to predict a more complex residual distribution for samples drawn from a trained probabilistic generative model. Specifically, we model the residual distribution as a structured Gaussian distribution. This work demonstrates that a network to predict this distribution can be tractably learned through maximum likelihood estimation of residual images; these predictions generalize well to previously unseen examples. Samples from this model are plausible, coherent, and can capture high-frequency details that may have been missed when training the original model with a factorized likelihood. We demonstrate the efficacy of this approach by: estimating ground truth covariances from synthetic data, adding coherent residual samples to face reconstructions from a VAE and include a further motivating example application for denoising face images.

2. Related work

Statistical quantification of uncertainty has been an area of interest for a long time. Many traditional statistical estimation models provide some measure of uncertainty on the inferred parameters of a fitted model. An additional source of uncertainty arises from the distribution of the model residual, which is often modeled as a spherical or diagonal Gaussian distribution. A common issue with these models is the false assumption of independence between pixels in the reconstruction residual image. Previous work on modeling correlated Gaussian noise is limited, it has been used for small data scenarios [19], for temporally correlated noise models [24] and in Gaussian processes [22].

The most recent related work on uncertainty prediction for deep generative models has been the prediction of heteroscedastic Gaussian noise for encoder/decoder models [13]. This approach is similar to the variational autoencoder with a diagonal covariance likelihood [15], but can be applied when the input and generated output are different, for example in semantic segmentation tasks. Interestingly, the maps of predicted variance in [13] correspond well to high frequency structured image features, similarly to the VAE with a diagonal noise model (see Fig. 10). These are structured regions that the model consistently struggles to accurately predict, which is a further encouragement for our work.

This paper proposes an approach to predict a structured uncertainty distribution for generated images from trained

networks. This method is applicable to models that reconstruct images without embellishing the prediction with details that were not present in the original input.

Generative adversarial networks (GANs) [8] are an implicit density estimation method widely used for generating novel images with a great deal of success. Samples from GANs have been shown to contain fine details and can be created at very high resolutions [12]. Although, they are not designed to provide reconstructions, some methods have been proposed that enable this [17, 21, 6]. Reconstructions from GAN models contain realistic high frequency details, however the generated images usually do not resemble well the inputs. This might be caused by mode-dropping [7], where parts of the image distribution are not well modeled. Despite recent work [2] addressing this issue, it remains an open problem. Furthermore, the residual distribution for reconstruction using GAN models is likely to be highly complex and therefore we do not use a GAN model as a starting point in this work.

Explicit density methods for learning generative models use either a tractable [20] or approximate [23] density to model the image distribution. These models allow maximization of the likelihood of the set of training observations directly through reconstruction. These approaches generate data that is more appropriate for the task this paper addresses and offer potential likelihood models to learn the uncertainty of prediction.

To train a network to predict the uncertainty of a reconstruction model, we need to choose a reconstruction likelihood that is efficient to calculate and allows structured prediction. PixelCNN [20] and derived work [9] provide an autoregressive sampling model, where the likelihood of a pixel is modeled by a multinomial distribution conditioned on the previously generated pixels. Although these models are capable of producing images with details, the generation process is computationally expensive. Approximate density models, such as the Variational Autoencoder (VAE) [15, 23], use a variational approximation of the marginal likelihood of the data under a diagonal Gaussian likelihood. These models are very efficient to reconstruct from, and the Gaussian noise model has an efficient likelihood that can be extended to include off-diagonal terms to allow correlated noise prediction. Therefore, the focus of this work is on explicit generative models with approximate likelihoods.

3. Methodology

Many generative models, with a (diagonal) Gaussian likelihood, define the conditional probability of the target image, \mathbf{x} , given some latent variable, \mathbf{z} , as:

$$p_{\theta}(\mathbf{x} | \mathbf{z}) = \mathcal{N}(\boldsymbol{\mu}(\mathbf{z}), \boldsymbol{\sigma}(\mathbf{z})^2 \mathbf{I}), \quad (1)$$

where \mathbf{x} is the target image flattened to a column vector, θ are the model parameters, and the mean $\boldsymbol{\mu}(\mathbf{z})$ and variance $\boldsymbol{\sigma}(\mathbf{z})^2$ are (non-linear) functions of the latent variables. This is equivalent to the forward model:

$$\mathbf{x} = \boldsymbol{\mu}(\mathbf{z}) + \boldsymbol{\epsilon}(\mathbf{z}), \quad (2)$$

where $\boldsymbol{\epsilon}(\mathbf{z}) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\sigma}(\mathbf{z})^2 \mathbf{I})$ is commonly considered as unstructured noise inherent in the data. However, reconstruction errors (or residuals) are often caused by data deficiencies, limitations in the model capacity and suboptimal model parameter estimation, while the residual itself is generally highly structured, as shown in Fig. 1. The noise term is being used mostly to explain failures in the model rather than just additive noise in the input. The novelty of this work is to encode coherent information about uncertainty in the reconstruction in $\boldsymbol{\epsilon}(\mathbf{z})$.

For the majority of models, $\boldsymbol{\sigma}^2$ is never used in practice when sampling reconstructions; it only adds white noise that would rarely improve the residual error, as shown in Fig. 1. Although in some cases $\boldsymbol{\sigma}^2$ is inferred from the data, in many cases it is assumed that $\sigma_i^2 = \sigma^2$, a user defined constant, which simplifies the log likelihood from Eq. 1 to a sum of squared errors scaled by $1/\sigma^2$.

In contrast, this paper extends the noise model to use a multivariate Gaussian likelihood with a full covariance matrix

$$p_{\theta}(\mathbf{x} | \mathbf{z}) = \mathcal{N}(\boldsymbol{\mu}(\mathbf{z}), \boldsymbol{\Sigma}_{\psi}(\mathbf{z})), \quad (3)$$

where $\boldsymbol{\Sigma}_{\psi}(\mathbf{z})$ is a (non-linear) function parametrized by ψ ; this is equivalent to $\boldsymbol{\epsilon}(\mathbf{z}) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\psi}(\mathbf{z}))$. The covariance matrix captures the correlations between pixels to allow sampling of structured residuals as demonstrated in our experiments.

A maximum likelihood approach is used to train the covariance prediction. We optimize the log-likelihood with respect to ψ keeping the generative model parameters θ fixed:

$$\arg \min_{\psi} \log(|\boldsymbol{\Sigma}_{\psi}(\mathbf{z})|) + (\mathbf{x} - \boldsymbol{\mu}(\mathbf{z}))^T (\boldsymbol{\Sigma}_{\psi}(\mathbf{z}))^{-1} (\mathbf{x} - \boldsymbol{\mu}(\mathbf{z})). \quad (4)$$

To simplify notation, in subsequent sections $\boldsymbol{\Sigma}$ and $\boldsymbol{\mu}$ are used to denote $\boldsymbol{\Sigma}_{\psi}(\mathbf{z})$ and $\boldsymbol{\mu}(\mathbf{z})$ respectively.

4. Covariance estimation

A deep neural network is used to estimate the covariance matrix $\boldsymbol{\Sigma}$ from the latent vector \mathbf{z} , henceforth referred to as the covariance network. This learning task is challenging on two fronts. It is ill-posed as no ground truth residual covariances are available for real data. Moreover, for each training example, a full covariance matrix must be estimated from a single target image. By definition, the covariance matrix is symmetric and positive definite. Therefore, for an image \mathbf{x}

with n pixels, the matrix contains $(n^2 - n)/2 + n$ unique parameters.

For any covariance estimation method there are three aspects to consider: (i) how difficult is it to sample from this covariance? (ii) how difficult is it to compute the terms in Eq. 4? and (iii) how difficult is it to impose symmetry and positive definiteness?

The need to draw samples arises from the fact that the covariance captures structured information about reconstruction uncertainty of the generative model. This means that drawing a sample from $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ and adding that to the reconstructed output may produce a result that is more representative of the target image, as shown in Fig. 1.

Given a decomposition of the covariance matrix: $\boldsymbol{\Sigma} = \mathbf{M}\mathbf{M}^T$, samples can be drawn from $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ as $\mathbf{x} = \mathbf{M}\mathbf{u} + \boldsymbol{\mu}$, where $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is a vector of standard Gaussian samples.

If $\boldsymbol{\Sigma}$ is the direct output of the covariance network, it needs to be inverted to calculate the negative log-likelihood in Eq. 4. Hence, it is more practical to estimate the precision matrix $\boldsymbol{\Lambda} = \boldsymbol{\Sigma}^{-1}$ as this term appears directly in the log likelihood, and the log determinant term can be equivalently computed as $\log(|\boldsymbol{\Sigma}|) = -\log(|\boldsymbol{\Lambda}|)$.

Cholesky Decomposition We represent the precision matrix via its Cholesky decomposition, $\boldsymbol{\Lambda} = \mathbf{L}\mathbf{L}^T$, where \mathbf{L} is a lower triangular matrix, and the covariance network only explicitly estimates \mathbf{L} .

Using the Cholesky decomposition, it is trivial to evaluate both terms in the negative log likelihood. The reconstruction error is $\mathbf{y}^T \mathbf{y}$, where $\mathbf{y} = \mathbf{L}^T(\mathbf{x} - \boldsymbol{\mu})$. The log determinant is $\log(|\boldsymbol{\Sigma}|) = -2 \sum_i \log(l_{ii})$, where l_{ii} is the i^{th} element in the diagonal of \mathbf{L} .

Sampling from $\boldsymbol{\Sigma}$ involves solving the triangular system of equations $\mathbf{L}^T \mathbf{y} = \mathbf{u}$ with backwards substitution, which requires $O(n^2)$ operations.

By construction, the estimated precision matrix $\boldsymbol{\Lambda}$ is symmetric. To ensure that it is also positive-definite it is sufficient to constrain the diagonal entries of \mathbf{L} to be strictly positive, *e.g.* by having the network estimate $\log(l_{ii})$ element-wise.

However, estimating this matrix directly is only a feasible solution for datasets with small dimensionality, as the number of parameters to be estimated increases quadratically with the number of pixels in \mathbf{x} .

Sparse Cholesky Decomposition To scale to larger resolution images we can reduce complexity by imposing a fixed sparsity pattern in the matrix \mathbf{L} , and only estimate the non-zero values of the matrix via the covariance network.

The sparsity pattern depends on the type of data being modeled. For image data, we propose that l_{ij} is only non-zero if $i \geq j$ and i and j are neighbours in the image plane, where pixels i and j are neighbours if a patch of size f

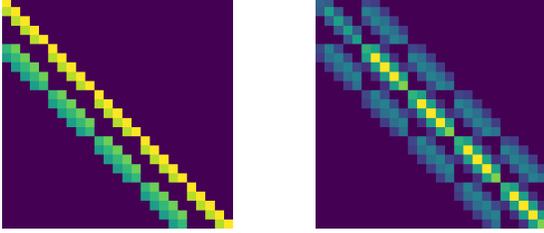


Figure 2: Left, an example of the sparsity patterns in the band-diagonal lower-triangular matrices \mathbf{L} , that are estimated by our model. Right, the precision matrix $\mathbf{\Lambda} = \mathbf{L}\mathbf{L}^T$.

centred at i contains j . This reduces the maximum number of non-zero elements in the matrix \mathbf{L} to $n((f^2 - 1)/2 + 1)$, where $f \ll n$.

The resulting sparse matrix \mathbf{L} is both lower-triangular and band-diagonal as shown in Fig. 2. This leads to a precision matrix $\mathbf{\Lambda}$ with a similar sparsity pattern with additional bands.

With a sparsity pattern of this form, our uncertainty model for each image can be interpreted as a Gaussian Random Field on its residual. A zero value in the precision matrix for pixels i and j implies that they are conditionally independent given all the other pixels. Similar Markov properties have been extensively used to model images [3].

With this representation the terms in the negative log likelihood can be evaluated efficiently, without the need to build the full dense matrix. Similarly, sampling can be performed by solving a sparse system of equations. Moreover, this approach is amenable to parallelization on the GPU, as each patch can be evaluated independently.

5. Results

We evaluate our model on two custom synthetic datasets to demonstrate the capability of our model to accurately describe known residual distributions. We also demonstrate our model on gray-scale cropped face images from the CelebA [18] dataset for sampling high frequency details to improve reconstructions. Finally, we show some examples of image denoising that takes advantage of the predicted covariance to better preserve structure. Results of our model evaluated on the CIFAR10 [16] dataset can be found in the supplemental material. All our models are implemented in Tensorflow [1] and they are trained on a single Titan X GPU using the Adam [14] optimizer. Unless otherwise stated, the input data for all the experiments is normalized in $[-1, 1]$.

5.1. Synthetic datasets

The goal of the two synthetic experiments is to evaluate the feasibility of training a covariance network to accurately estimate the residual distribution, where the true mean and covariance matrix are available for validation purposes.

Since the goal here is to evaluate the covariance prediction network, we simplify these experiments by bypassing

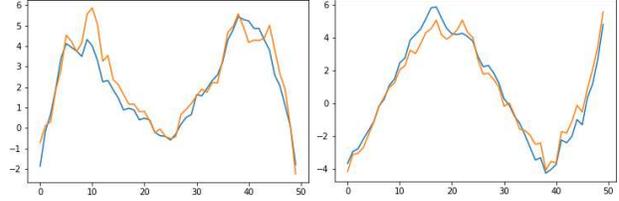


Figure 3: Reconstructions with samples from Σ . Each image corresponds to a different spline on the test dataset. The ground-truth spline is depicted in blue, and the estimated one in orange.

the use of a generative model, and directly predicting Σ from μ . This means that the input to our covariance prediction network is μ .

Each dataset contains 35,000 training examples and 1,000 test examples. In all the experiments, the test examples are reconstructed by drawing a sample from the estimated covariance and adding it to the true mean μ .

Both datasets are constructed by generating a set of μ , and then adding a random sample of correlated noise to them: $\mathbf{x} \sim \mathcal{N}(\mu, \Sigma)$, where Σ is a function of μ . Therefore, the added noise is dependent on the structure in the image, which imitates the situation on real data. This also reflects the main assumption of our model: that there is sufficient information in the latent variables \mathbf{z} (or in μ for the synthetic experiments) to estimate the residual distribution Σ .

We emphasize that despite the true covariance matrices being known for these synthetic experiments, we do not use them at train time. We train the prediction network using the objective in Eq. 4, which makes no use of the true covariance and mimics the situation with real datasets.

5.1.1 Splines

The first synthetic dataset is composed of one dimensional signals, with 50 points per example. Each spline is comprised of a low frequency component and a correlated high-frequency one. The high-frequency component is produced by a unique covariance matrix per example, that is generated by a deterministic function that takes as input the low-frequency signal. For more details please refer to the supplemental material.

The covariance prediction network is a multi-layer perceptron (MLP) with two layers of 100 units with relu activations and batch normalization [11], and a final layer with 1275 units. The final layer directly outputs the lower triangular part of the matrix \mathbf{L} , which does not use our proposed sparsity approach. The model is trained with a learning rate of $1e-4$ for 200 epochs.

Reconstructions for this dataset are obtained by adding a sample from Σ to μ . We show results for reconstructions in Fig. 3, where the uncertainty model is able to add a plausible

	$-\log p(\mathbf{x} \Sigma(\boldsymbol{\mu}))$	KL	$\ \Sigma(\boldsymbol{\mu}) - \Sigma_{gt}\ _2$
Ground truth	18.65 ± 0.75	-	-
Diagonal model	42.64 ± 0.83	72.06 ± 0.30	3.81 ± 1.21
Ours	21.12 ± 0.79	3.56 ± 0.18	1.26 ± 0.74

Table 1: Quantitative comparison of reconstructions on the splines dataset. KL denotes the KL divergence $D_{KL}(\mathcal{N}(\mathbf{0}, \Sigma(\boldsymbol{\mu})) || \mathcal{N}(\mathbf{0}, \Sigma_{gt}))$, where Σ_{gt} is the ground truth covariance matrix, and $\Sigma(\boldsymbol{\mu})$ is the estimated covariance. Our model obtains significant improvements under all metrics over a diagonal covariance model.

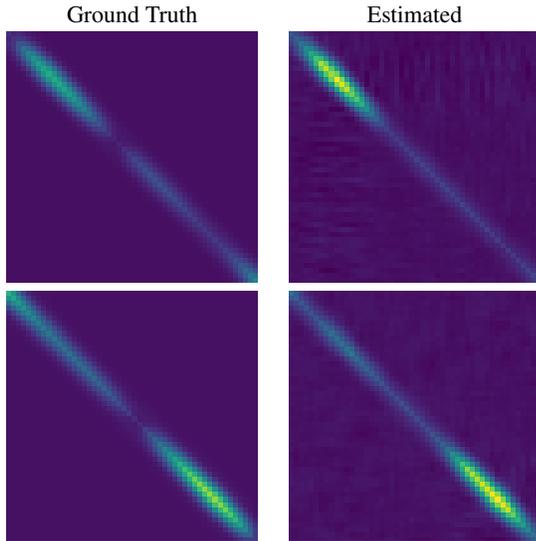


Figure 4: Estimated covariance matrices for the spline dataset. Each row corresponds to a different spline example on the test set. Our model is able to learn the variations in the structured residual distributions.

high-frequency component to the input $\boldsymbol{\mu}$.

Quantitative results are presented in Table 1. We compare with a covariance network that only estimates a diagonal covariance matrix. As the ground-truth covariances contain off-diagonal structure, the diagonal model is bound to fail in representing it. Our model achieves a negative log likelihood similar to the one evaluated using the real covariance matrices.

As we have access to the ground-truth covariance matrices for each test example, we can directly compare the estimated covariance with the ground-truth ones. We show qualitative results in Fig. 4. Note how the model is able to recover most of the off-diagonal values in the covariance.

5.1.2 Ellipses

The second synthetic dataset was built to evaluate the covariance prediction network for images and to highlight the

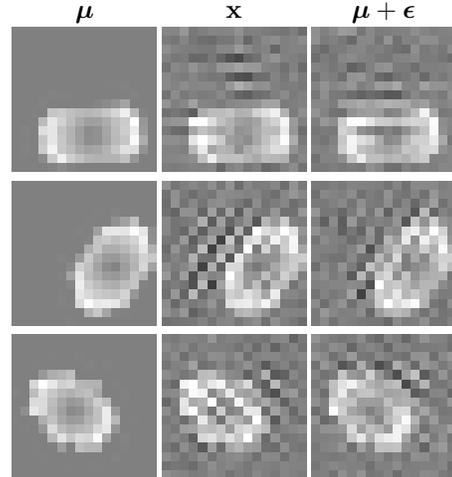


Figure 5: Reconstructions from our model, left column: input $\boldsymbol{\mu}$, middle column: original \mathbf{x} , and right column: reconstruction using a sample from the learned residual distribution, where $\epsilon \sim \mathcal{N}(\mathbf{0}, \Sigma)$.

	$-\log p(\mathbf{x} \Sigma(\boldsymbol{\mu}))$	KL	$\ \Sigma(\boldsymbol{\mu}) - \Sigma_{gt}\ _2$
Ground truth	-286 ± 2.8	-	-
Diagonal model	-149 ± 3.3	707 ± 7.2	1.80 ± 0.21
Ours	-259 ± 2.7	113 ± 2.6	1.06 ± 0.34

Table 2: Quantitative comparison on the ellipses dataset (see Table 1 for a description of the metrics). Our model is able to better model the real covariance matrices with its more complex uncertainty distribution.

limitations of estimating a dense matrix \mathbf{L} .

We generate a dataset of synthetic gray-scale 16×16 images. For each example, the mean image contains an ellipse with random width, height, position and rotation angle. The prototype covariance matrix for this dataset generates lines and is rotated by the same random rotation angle that was used for the ellipse, thus generating random lines that are aligned with the ellipse.

For this dataset, estimating directly a dense Cholesky matrix \mathbf{L} requires 32,896 values per image. Even at this limited image size we were unable to train a dense prediction model. Instead, we use the sparse prediction model with a neighborhood of size 5×5 . The model was trained for 200 epochs with a learning rate of $1e-3$,

Reconstructions of the test set are shown in Fig. 5, where we show results of taking a sample from Σ which is added to $\boldsymbol{\mu}$. The covariance prediction network is successful in mapping the uncertainty distribution from the mean $\boldsymbol{\mu}$. The samples from the covariance exhibit high-frequency detail that matches the true residual. A quantitative comparison with a diagonal Gaussian model is presented in Table 2. Our

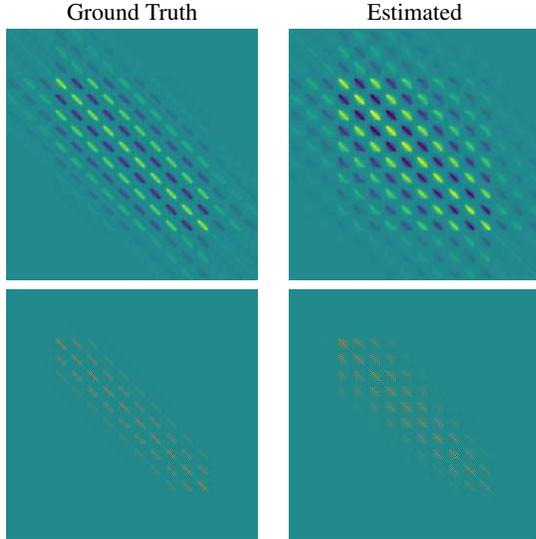


Figure 6: Covariance matrices estimated by our model, where each row corresponds to a different ellipse example on the test set. Much of the structure of the real covariance matrices is recovered.

model achieves a negative log likelihood similar to the one evaluated using the real covariance matrices.

Examples of ground-truth and estimated covariances are shown in Fig. 6 and illustrate the accuracy of the covariance prediction. The covariance structure for this model is more complex than the one for the splines, and yet the model is still able to recover it effectively with our sparse estimation of the Cholesky matrix.

5.2. CelebA

We now show results of employing the covariance prediction network on a real images dataset: CelebA [18]. The aligned and cropped version of the dataset is used, where a further cropping and resizing to 64×64 is performed and the images are converted from RGB to gray-scale. The dataset consists of 202,599 images of faces, which we split into 182,637 for training and 19,962 for testing as recommended by the authors.

We train both an Autoencoder (AE) [5] and a VAE on this dataset using the architecture in [21]. The Autoencoder is trained with a mean squared error, and it serves to show the performance of our model when using as input a \mathbf{z} from an uncontrolled latent space. We then trained two covariance prediction networks, one for each model, to estimate the residual uncertainty from the latent variables, \mathbf{z} . These networks have an initial block that is similar to the decoder of the VAE and a second block with four convolutional layers. The output of the networks is the sparse Cholesky decomposition of the precision matrix with a neighborhood of size 7×7 pixels. The sparsity imposed on the matrix means that instead of estimating 8,390,656 values as out-

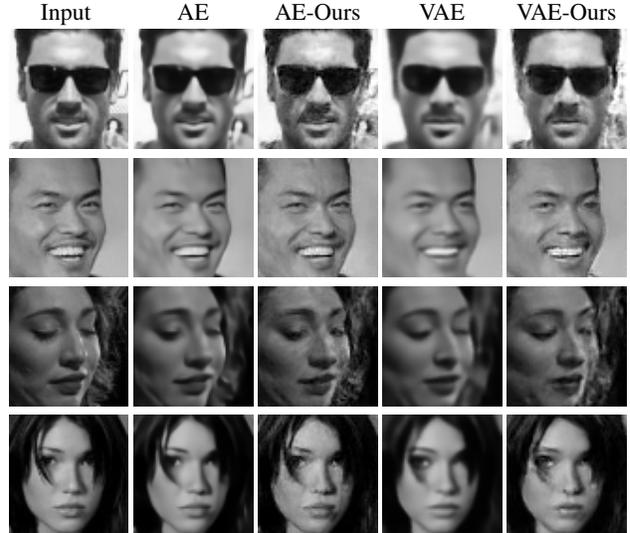


Figure 7: Comparison of image reconstructions for the different models. The AE and VAE both generate over-smoothed images. For both the AE and VAE, our model adds plausible high-frequencies from a single sample drawn from the predicted uncertainty distribution.

Model	NLL	$-\log p(\mathbf{x} \mathbf{z})$
AE [5]	-	-
VAE [15]	-5378 ± 931	-6079 ± 936
Ours-AE	-	-8242 ± 433
Ours-VAE	-7753 ± 1323	-8386 ± 1339

Table 3: Quantitative comparison of density estimation error. NLL denotes the upper bound of the marginal negative log likelihood, lower is better. The residuals are structured, thus the diagonal model produces poor estimations. Our estimation method is able to significantly improve over the AE and VAE simplified noise model.

put, the network only needs to estimate 102,400. The covariance networks are trained with a learning rate of $1e-3$ for 50 epochs. Additional implementation details are given in the supplemental material.

Example reconstructions for the VAE and the AE models are shown in Fig. 7. Reconstructions from both methods suffer from the over-smoothing effects of using a simplified Gaussian likelihood and in both cases high frequency details are lost. By adding a random sample of the predicted residual distribution our method is able to recover plausible high-frequency details, resulting in more realistic looking images. The added detail corresponds to important face features, which is lost by both autoencoder models, like teeth and hair.

A quantitative comparison of using either the predicted covariance or a diagonal covariance for calculating the negative log-likelihood of the reconstructions μ is presented in



Figure 8: Samples from a β -VAE with residual samples from our model. The covariance network predicts plausible structured residuals for the synthesized images.



Figure 9: Samples drawn with our model while interpolating on the latent space, from the top-left input to the bottom-right one. Using a fixed noise vector \mathbf{u} , samples are drawn from our model as $\mathbf{x} = \boldsymbol{\mu} + \mathbf{M}\mathbf{u}$, where the covariance $\boldsymbol{\Sigma} = \mathbf{M}\mathbf{M}^T$. The residuals are smoothly interpolated for the different images, which suggests that the estimated covariance matrices also vary smoothly.

Table 3. The reported values for VAE-based methods follow the protocol described in [4], where the upper bound of the marginal negative log likelihood is evaluated by numerically integrating over \mathbf{z} with 500 samples per image. For VAE and our model using as base a VAE, the $-\log p(\mathbf{x} | \mathbf{z})$ term is evaluated imposing zero variance in \mathbf{z} .

Images generated by decoding samples from the prior distribution on the latent space of a β -VAE with added residuals from our model are shown in Fig. 8. We found that the VAE with a diagonal Gaussian likelihood overfitted the reconstruction error, thus neglecting the KL term for the prior on the latent space, which in turn produces low quality samples. Instead, we trained our structured uncertainty network on a β -VAE [10] with $\beta = 5$. This corresponds to increasing the weight of the KL term on a VAE, which is known to improve sample quality. Our model is able to generate structured residuals of similar quality as those whose \mathbf{z} was created by encoding an image. The covariance network is still able to learn meaningful structured uncertainty for the VAE model, as shown in the supplemental material.

To evaluate the generalization of the model to different regions in the latent space, we show in Fig. 9 the result of

Model	MSE
DAE	$5.13\text{e-}3 \pm 2.52\text{e-}3$
Ours	$2.99\text{e-}3 \pm 7.98\text{e-}4$

Table 4: Quantitative comparison for denoising in terms of mean squared error (MSE) with respect to the noise-free input. Our model is on average able to produce better results than an Autoencoder trained for denoising.

interpolating between an image and its x-flipped mirror image. The generated images are plausible and the sampled residuals are consistent across the interpolated images.

To further highlight the differences between the diagonal and predicted covariance noise models, the variance maps for both are shown in Fig. 10. The diagonal model must explain all the errors with variance, while our model is able to explain some with correlations. The effect of this is evident when sampling from the estimated $\boldsymbol{\Sigma}$.

5.3. Denoising Example

One possible application of the covariance prediction network is image denoising, as shown in Fig. 11. We hypothesize that our predicted covariance matrices will only span the space of valid face residuals, thus projecting a noisy residual in that space will remove the noise. Here, the noisy image is reconstructed using a VAE, that was *not* trained with noisy data. The difference between the noisy input and the reconstruction is computed, and projected onto $\hat{\boldsymbol{\Sigma}}$, which is created by taking the 1000 eigenvectors of $\boldsymbol{\Sigma}$ with the largest eigenvalues. The projected difference is added to the VAE reconstruction to produce the final denoised image. A comparison is shown with an Autoencoder trained explicitly for denoising with the same architecture as the VAE. Note how the use of the structure covariance model is able to filter the noise to generate plausible structured high-frequency details, while the denoising Autoencoder fails to recover those details.

Quantitative results for this experiment are shown in Table 4, where MSE is reported for the first 2000 images in the test set. Our model achieves significantly lower error than an Autoencoder trained specifically for this task.

6. Conclusions

In this paper, we have demonstrated what we believe is the first attempt to train a deep neural network to predict structured covariance matrices for the residual image distribution of unseen image reconstructions. Our results show that ground truth covariances can be learned for toy data and good samples generated for celebA data. A further motivating experiment is also shown for denoising face images.

An interesting question is whether we can train a generative model and covariance network in tandem to produce better reconstructions. We postulate that a better residual

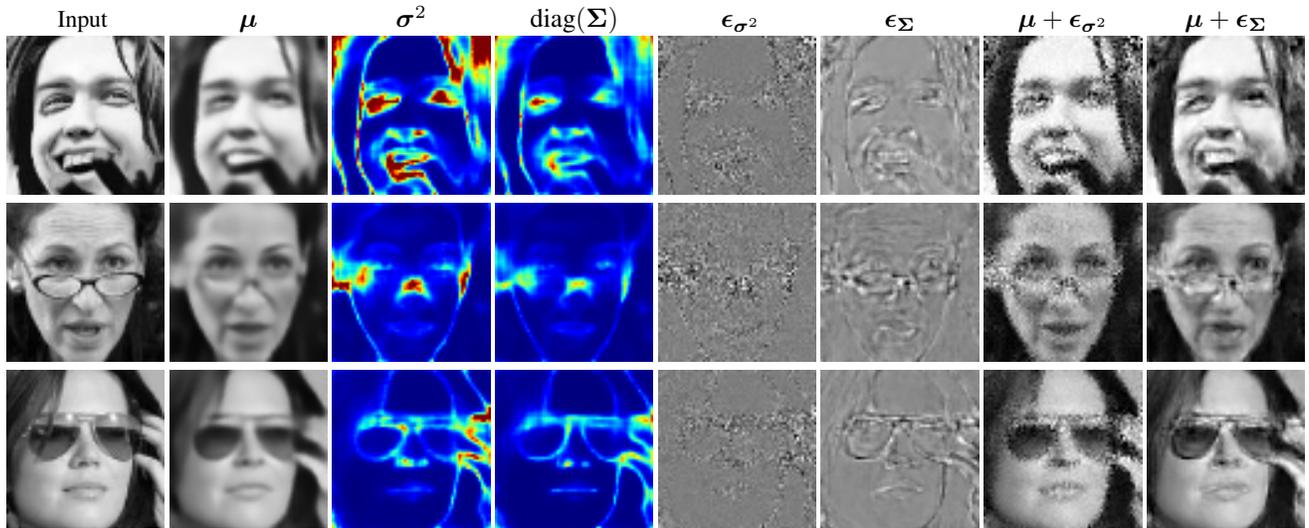


Figure 10: Variance maps for different inputs, where μ and σ^2 are predicted by a diagonal covariance VAE, and $\text{diag}(\Sigma)$ is the diagonal of our estimated covariance matrix. Residual predictions are sampled as $\epsilon_{\sigma^2} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ for the VAE, and $\epsilon_{\Sigma} \sim \mathcal{N}(\mathbf{0}, \Sigma)$ for our model. The diagonal noise estimation model mistakenly identifies teeth or skin wrinkles as variance, whereas the covariance model properly identifies them as regions with high covariance, yet low variance.

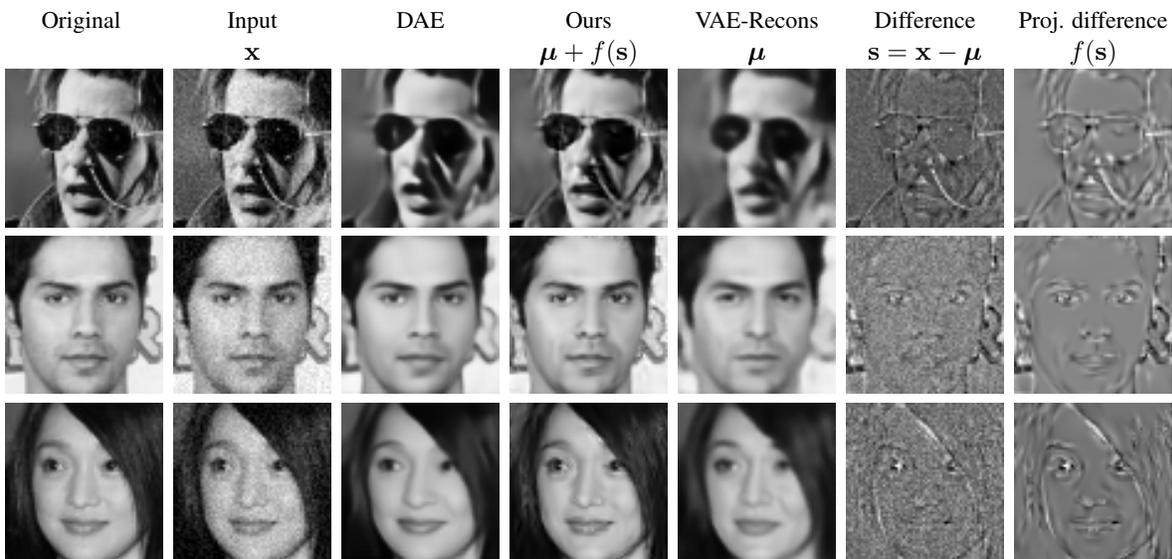


Figure 11: Denoising experiment, left column: original image without noise, second column: image with added noise, third column: denoising autoencoder (DAE) result, fourth column: our result, fifth column: VAE reconstruction from the noisy input, sixth column: difference between the VAE reconstruction and the noisy input, right column: the difference projected on $\tilde{\Sigma}$, the matrix constructed with 1000 eigenvectors of Σ . Our result is the sum of the projected difference and the VAE reconstruction. Our model is able to recover fine details that are lost with the DAE approach.

model would improve the reconstructions, as it may consider plausible but different realizations of high frequency image features as likely. However, there may be some effort required to keep the two networks consistent and additional investigations are needed into priors on the predicted covariances.

Another questions is what is the best input data to use for

the covariance network. Is it best to learn directly from \mathbf{z} or μ or some other pre-learned function of either of these?

Finally, in this work we have only examined the residual image distribution for reconstruction models trained with a Gaussian likelihood. An interesting avenue to explore would be a structured pixel uncertainty distribution for other reconstruction error metrics, such as perceptual loss.

Acknowledgements

This work has been supported by the EPSRC CDE (EP/L016540/1) and CAMERA (EP/M023281/1) grants.

References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, and et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. 4
- [2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein gan. *CoRR*, 2017. 2
- [3] C. M. Bishop. *Pattern recognition and machine learning*. springer, 2006. 4
- [4] Y. Burda, R. Grosse, and R. Salakhutdinov. Importance weighted autoencoders. *ICLR*, 2016. 1, 7
- [5] G. W. Cottrell, P. Munro, and D. Zipser. Learning internal representations from gray-scale images: An example of extensional programming. In *Conference of the Cognitive Science Society*, 1987. 6
- [6] J. Donahue, P. Krähenbühl, and T. Darrell. Adversarial feature learning. *ICLR*, 2017. 2
- [7] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. 2
- [8] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014. 2
- [9] I. Gulrajani, K. Kumar, F. Ahmed, A. A. Taïga, F. Visin, D. Vazquez, and A. Courville. PixelVAE: A Latent Variable Model for Natural Images. In *ICLR*, 2017. 2
- [10] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017. 7
- [11] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456, 2015. 4
- [12] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 2
- [13] A. Kendall and Y. Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems (NIPS)*, 2017. 2
- [14] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. 4
- [15] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *ICLR*, 2014. 1, 2, 6
- [16] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009. 4
- [17] A. B. L. Larsen, S. K. Sønderby, and O. Winther. Autoencoding beyond pixels using a learned similarity metric. In *ICML*, volume 48, pages 1558–1566. JMLR, 2016. 2
- [18] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *ICCV*, December 2015. 4, 6
- [19] C. L. Nikias and R. Pan. Time delay estimation in unknown gaussian spatially correlated noise. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 36(11):1706–1714, 1988. 2
- [20] A. Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. *ICML*, 2016. 2
- [21] Y. Pu, W. Wang, R. Henao, L. Chen, Z. Gan, C. Li, and L. Carin. Adversarial symmetric variational autoencoder. *NIPS*, 2017. 2, 6
- [22] C. E. Rasmussen and C. K. Williams. *Gaussian processes for machine learning*, volume 1. MIT press Cambridge, 2006. 2
- [23] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *ICML*, 2014. 1, 2
- [24] M. W. Woolrich, B. D. Ripley, M. Brady, and S. M. Smith. Temporal autocorrelation in univariate linear modeling of fmri data. *NeuroImage*, 14(6):1370 – 1386, 2001. 2