Uncertainty Quantification in CNN-Based Surface Prediction Using Shape Priors

Conference Paper**Author(s):**

Tóthová, Katarína ; Parisot, Sarah; Lee, Matthew C.H.; Puyol-Antón, Esther; Koch, Lisa M.; King, Andrew P.; Konukoglu, Ender; Pollefeys, Marc

Publication date:

2018

Permanent link:

<https://doi.org/10.3929/ethz-b-000310153>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

Lecture Notes in Computer Science 11167, https://doi.org/10.1007/978-3-030-04747-4_28

Uncertainty Quantification in CNN-Based Surface Prediction Using Shape Priors

Katarína Tóthová¹, Sarah Parisot², Matthew C. H. Lee³, Esther Puyol-Antón⁴,
Lisa M. Koch¹, Andrew P. King⁴, Ender Konukoglu¹, and Marc Pollefeys^{1,5}

¹ ETH Zurich, Switzerland

² Aimbrain Ltd.

³ HeartFlow Inc.

⁴ King's College London, United Kingdom

⁵ Microsoft

Abstract. Surface reconstruction is a vital tool in a wide range of areas of medical image analysis and clinical research. Despite the fact that many methods have proposed solutions to the reconstruction problem, most, due to their deterministic nature, do not directly address the issue of quantifying uncertainty associated with their predictions. We remedy this by proposing a novel probabilistic deep learning approach capable of simultaneous surface reconstruction and associated uncertainty prediction. The method incorporates prior shape information in the form of a principal component analysis (PCA) model. Experiments using the UK Biobank data show that our probabilistic approach outperforms an analogous deterministic PCA-based method in the task of 2D organ delineation and quantifies uncertainty by formulating distributions over predicted surface vertex positions.

Keywords: Surface reconstruction · Uncertainty quantification · Deep learning · Shape prior.

1 Introduction

Reconstruction of organ surfaces and segmentation of their bodies are amongst the most important tasks in medical image analysis. High-quality organ surface models are often sought after in disciplines such as cardiac or neuro-imaging, and provide a powerful tool in diagnosis, surgical planning, disease tracking, longitudinal studies and interpretation of functional data [21, 22, 1].

Traditional approaches to parametric surface modelling rely on evolving deformable shapes according to predefined forces [8, 12, 11, 23] or use atlas registration [22, 21, 10, 2]. Recent advances in machine learning have showed the possibility of training a deep neural network in an end-to-end manner: from images to parametrised shapes [19]. In their work, Milletari et al. [19] devised a convolutional neural network (CNN) to directly predict coordinates of the organ surface mesh from the imaging data. To build in prior shape knowledge, the

network contains an explicit principal component analysis (PCA) layer and predictions are made as linear combinations of the modes of variation determined from the training data.

Despite the advances in organ surface modelling and associated segmentation, ways of estimating the precision of the prediction are still sparse. Yet, it is of vital importance for interpreting medical data to be able to not only measure the accuracy of the result as a deterministic sample but to quantify the uncertainty associated with it as well. We can distinguish between two types of uncertainty [14]: *aleatoric* uncertainty inherent to the data, modelled by probability distribution over model outputs, and *epistemic* uncertainty accounting for uncertainty in the model parameters, which typically decreases with increasing data size.

In the context of medical imaging, the need for addressing aleatoric uncertainty stems from the nature of the data. Medical imaging data often suffers from high levels of noise, coarse resolution and imaging artifacts - all the factors conspiring towards heightened need for quantification of uncertainty about the produced results. In image segmentation, this has been approached by means of segmentation sampling [17, 6], where several plausible segmentations are gathered to estimate the variability of the output. While [6] uses MCMC to sample segmentations from an estimated posterior distribution where likelihood and prior functions are defined, [17] introduces Gaussian processes to sample from the posterior directly, knowing only its mean and covariance. On the other hand, uncertainties inherent to the prediction models can be captured by means of distributions over model parameters. Bayesian neural networks [7, 18, 20, 9], where one puts priors over the model parameters instead of using deterministic values, have been employed to this end. Extensions combining both aleatoric and epistemic uncertainty into one model have been proposed in [14, 16].

In our work, we build on a PCA-based method of surface reconstruction [19] and propose a probabilistic approach to integrate aleatoric uncertainty quantification within the model. We formulate the problem as a conditional probability estimation incorporating shape information in the form of a PCA model.

Hence, our approach addresses three main objectives:

- Direct probabilistic surface mesh prediction from imaging data. The proposed method improves upon a deterministic direct coordinate prediction by up to 12% on the UK BioBank dataset [25] as measured by DICE.
- Use of PCA-based shape priors to predict sensible shapes only. Here, our probabilistic approach improves upon the deterministic PCA-based method by up to 10.7% in terms of DICE score.
- Novel aleatoric uncertainty quantification formulation by means of assessing the posterior of the predicted surface.

2 Prediction Model

We formulate the surface prediction problem via a probabilistic model that utilises principles from probabilistic PCA [5]. Our goal is to build a model that

goes beyond deterministic prediction and allows sampling 2D delineations of the organ surface based on the corresponding MRI data. To this end, we express the prediction as a probability conditioned on the image, $p(y|x)$, where x refers to the MRI image and y to the parametrised surface, i.e. set of surface vertices. We model the conditional probability $p(y|x)$ with a latent variable model

$$p(y|x) = \int p(y|z, x)p(z|x)dz, \quad (1)$$

where z is the set of latent variables.

The PCA aspect of the model lies within the definition of $p(y|z, x)$

$$p(y|z, x) = \mathcal{N}(y|US^{\frac{1}{2}}z + \mu + s(x), \sigma^2 I), \quad (2)$$

Here, U , μ and S are the principal component matrix (principal components are columns of the matrix), mean and diagonal covariance matrix respectively, all precomputed using the surfaces in the training set, and s is a global spatial shift that depends on the image x . In this formulation, the latent variable z can be interpreted as the PCA weights corresponding to the surface y . Variance σ^2 refers to the noise level in the data.

The conditioning to the image is modelled with $p(z|x)$ and we use a deep CNN architecture for this purpose. Specifically, we express

$$p(z|x) = \mathcal{N}(z|\mu(x), \Sigma(x)). \quad (3)$$

Here, a deep network takes the image x as input and predicts $\mu(x)$ and $\Sigma(x)$ simultaneously.

The last component of the proposed model is the prior for the latent variables $p(z)$. The probabilistic PCA model assumes a unit Gaussian distribution for the PCA weights, i.e. $p(z) \sim \mathcal{N}(0, I)$. In our model we assume the same, which becomes important when training the model.

Training During training we optimise the parameters of $\mu(x)$, $\Sigma(x)$ and $s(x)$ using a training set. The optimisation objective consists of two terms: the first one aims to maximise the conditional probability $p(y|x)$ while the latter regularises the loss by minimising the Kullback-Leibler divergence (KLD) between the assumed prior distribution of z and the observed one in the training set, i.e. $\int p(z|x)p(x)dx \approx \sum_x p(z|x)$.

Direct maximisation of Equation 1 requires marginalisation of the latent variable. The marginal distribution is also Gaussian with analytical mean and variance that allow direct computation and therefore, optimisation of $\ln p(y|x)$. However, this requires inverting a not-necessarily-diagonal covariance matrix of the size $(\text{num of vertices})^2$ at each iteration, which can be infeasible due to the size and potential numerical instabilities, as we empirically observed. Therefore, instead of directly maximising $\ln p(y|x)$, we use Jensen's inequality to derive a lower bound as follows

$$\ln p(y|x) \geq \mathbb{E}_{z|x} [\ln p(y|z)] \cong \frac{1}{L} \sum_{l=1}^L \ln p(y|z_l), \quad (4)$$

where z_l is sampled from $p(z|x)$ defined in (3).

Maximisation of the lower bound in Equation 4 will not necessarily satisfy the prior in the probabilistic PCA model. To address this, we use a regularisation term that aims to align the observed latent variable distribution with its prior. In order to satisfy the PCA model

$$p(z) \cong \sum p(z|x_n), \quad x_n \sim p(x). \quad (5)$$

We use the KLD as a measure of deviation from this criteria and minimise

$$\text{KLD} \left(\sum_{n=1}^N p(z|x_n), p(z) \right) \quad (6)$$

Using (4) and (6), our full model can then be trained by solving the following minimisation problem:

$$\min_{\theta} \left\{ \lambda \text{KLD} \left(\sum_{n=1}^N p(z|x_n), p(z) \right) - \sum_{n=1}^N \frac{1}{L} \sum_{l=1}^L \ln p(y|z_l) \right\}, \quad (7)$$

where $z_l \sim \mathcal{N}(z|\mu(x_n), \Sigma(x_n))$ and λ is a regularisation parameter. For further details on the derivations please refer to Appendix A.

Inference As we mentioned previously, in the proposed model $p(y|x)$ is a Gaussian distribution. For a given test image, we perform prediction by computing the mean and the covariance matrix of this distribution.

Given (2) and (3) we can write

$$\mathbb{E}(y|x) = US^{\frac{1}{2}}\mu(x) + \mu + s, \quad (8)$$

$$\text{var}(y|x) = \sigma^2 I + US^{\frac{1}{2}}\Sigma(x)(US^{\frac{1}{2}})^T. \quad (9)$$

Note that the first term of the predicted variance, i.e. $\sigma^2 I$, represents a fixed noise level in the data. Full derivation can be found in Appendix B.

3 Experiments and Results

We have tested the proposed method on a task of delineation of myocardium boundaries in cardiac MRI using imaging volumes from UK BioBank [25]. The 2D surface reconstruction consisting of the prediction of 50 vertices was evaluated on small (60×60 crops around the heart) and full (200×200 crops) field of view (FOV) images with the following characteristics:

- Small FOV: isotropic pixel of size 1.8 mm; 572 training, 160 testing and 195 validation examples
- Full FOV: isotropic pixel of size 1.8269 mm; 1532 training, 455 testing and 499 validation examples

Table 1. Vertex coordinate regression results. Comparison of direct vertex prediction (Direct Vertex), deterministic PCA (detPCA) and our proposed probabilistic PCA-based approach (probPCA) with varying number of principal components. Segmentations used for computation of DICE were obtained by flooding the corresponding delineations. RMSE directly compares the predicted vertex coordinates to the reference ones.

Small FOV	DICE	RMSE	Full FOV	DICE	RMSE
Direct Vertex	0.86 ± 0.06	2.21 ± 1.04	Direct Vertex	0.75 ± 0.14	4.20 ± 2.82
detPCA 12	0.87 ± 0.07	2.42 ± 1.27	detPCA 12	0.78 ± 0.13	4.20 ± 2.65
detPCA 8	0.84 ± 0.07	2.51 ± 1.26	detPCA 8	0.75 ± 0.14	4.48 ± 2.69
<i>Ours:</i>			<i>Ours:</i>		
probPCA 12	0.88 ± 0.09	1.91 ± 1.05	probPCA 12	0.84 ± 0.10	2.59 ± 2.37
probPCA 8	0.88 ± 0.09	1.80 ± 0.97	probPCA 8	0.84 ± 0.11	2.60 ± 2.38

We used one imaging slice per original MRI volume. Active contours [13] delineations consisting of 50 connected (corresponding throughout the subjects) vertices were obtained from reference segmentations extracted from the UK BioBank dataset. These were prepared automatically using expert-segmentations and combination of learning and registration methods described in [3, 24].

The deep network architecture used in our model is analogous to the main branch described in [19] with 9 convolutional, 3 pooling and one dense layer (CL9P3DL1). Convolutional layers were followed by ReLU activations. Training was done by minimising loss (7) using RMS-Prop optimiser with a constant learning rate 10^{-6} for batches of size 5. Noise level in the data σ^2 and regularisation parameter λ were empirically set to $\sigma^2 = 5 \times 10^{-2}$ and $\lambda = 10^5$.

Table 1 shows evaluation of segmentations obtained from myocardium delineation in terms of DICE score measuring region overlap and Root Mean Square Error (RMSE) comparing distances between the corresponding predicted and reference vertices. We used two different methods as baselines for comparison. Firstly, a network with the same architecture as the one employed in our model (CL9P3DL1) was used to directly predict the coordinates of the surface vertices (Direct Vertex). Secondly, we utilised the deterministic PCA approach (detPCA) based on [19] again following the same architecture (without the spatial transformer refinement). For our method, the mean prediction $\mathbb{E}(y|x)$ served for computation of DICE scores and RMSE.

Working on the small FOV dataset, we exemplify the qualitative results of our probabilistic method in Figure 1. Here, $p(y|x)$ distributions of predicted vertices are illustrated by plotting their variance as 30% confidence ellipses along the predicted mean delineation. We only plot the 30% confidence intervals and show results for small FOV for visualization purposes. Notice how direction and size of the ellipses vary along the perimeter of the shape - the larger the variance the bigger the uncertainty over the position of the vertex that can be sampled from this distribution.

Figure 2 looks into the distribution of a specific vertex position in more detail. It essentially shows three types of results. On the left, we have high region

Fig. 1. Uncertainty quantification in small FOV myocardium delineations using 12 principal components at data noise level 5×10^{-2} : 30% confidence ellipses for points along the mean prediction. Colours correspond to the direction of the major axes. Only half of the points were plotted for clarity.

Fig. 2. Uncertainty quantification in small FOV myocardium delineation using 12 principal components at data noise level 5×10^{-2} : reference (red), predicted mean(cyan). Posterior distribution of specific points from predictions above (reference point and its mean prediction marked with a dot). Ellipses correspond to 30, 95 and 99.9% confidence regions. Observe how the reference point lies within the 99.9% area. Note we used corresponding vertices throughout the subjects.

overlap with high DICE (0.93) score, and relatively small variance; however large RMSE (3.23), which can lead to problems if one was to use the retrieved surface for e.g. registration purposes. The middle figure is a failure case with high variance, low DICE(0.66) and high RMSE (6.12). And finally on the right is a result with high DICE (0.92) and low RMSE (0.83), but higher uncertainty than in the first case. Several conclusions can be drawn from this. Firstly, failure to predict the correct delineation leads to heightened uncertainty about the prediction. Secondly, what may seem as a good solution in terms of overlap (and hence DICE) may not necessarily be ideal in terms of RMSE. While delineations may align, the corresponding vertices may not. Lastly, even if the corresponding points do not align, the uncertainty over their position may still be relatively low (compare images on the left and right in Figure 2).

Finally, Figure 3 illustrates results of the sampling process on full and small FOV images. Bottom row images here correspond to the cases outlined in Figure 2. Observe how the samples become more variable as the uncertainty grows through subjects.

4 Conclusion

In this paper, we presented a novel probabilistic deep learning approach for simultaneous surface reconstruction and aleatoric uncertainty prediction. Inspired by the works on deterministic shape models [19] and probabilistic PCA [5], our surface reconstruction method incorporates prior shape knowledge via a linear PCA model. Experiments using the UK Biobank data have shown that our probabilistic approach outperforms an analogous deterministic PCA-based method. In contrast to deterministic approaches, which provide a single surface, our method yields a distribution of positions for every vertex on the surface.

Fig. 3. Myocardium delineation: reference (red), predicted mean(cyan), sampled delineation (blue). We predicted positions of 50 vertices on the outline of myocardium using 12 principal components at the noise level of 5×10^{-2} . Top: Model trained on full FOV images. Bottom: Model trained on small FOV images.

This way, we can not only sample numerous predictions from one model, but ascertain the vital uncertainty pertinent to each prediction.

Providing uncertainty estimations is essential in the medical domain, and in particular for surgery planning where precision and accuracy are of utmost importance. While the proposed method is capable of producing surface predictions of healthy 2D cardiac data, we intend to extend it to more challenging scenarios in the future. The aim is to generalise to 3D setting and other types of organs as well. Extensions to organs of more variable shapes may require adaptation of the shape model. Finally, transfer learning or domain adaptation techniques could be investigated to apply the proposed approach to datasets of substantially smaller size.

Acknowledgements

This research has been conducted using the UK Biobank Resource under Application Number 17806.

References

1. Bai, W., et al.: Automated cardiovascular magnetic resonance image analysis with fully convolutional networks. arXiv preprint, arXiv:1710.09289 (2017)
2. Bai, W., et al.: A Probabilistic Patch-Based Label Fusion Model for Multi-Atlas Segmentation With Registration Refinement: Application to Cardiac MR Images. *IEEE Transactions on Medical Imaging* **32**(7), 1302–1315 (2013)
3. Bai, W., et al.: Semi-supervised learning for network-based cardiac MR image segmentation. In: *Med. Image Comput. Comput. Assist. Interv (MICCAI)*, 253–260 (2017)
4. Baumgartner, C.F., Koch, L.M., Pollefeys, M., Konukoglu, E.: An Exploration of 2D and 3D Deep Learning Techniques for Cardiac MR Image Segmentation. In: Pop M. et al. (eds.) *Statistical Atlases and Computational Models of the Heart. ACDC and MMWHS Challenges. Lecture Notes in Computer Science 10663*, Springer, (2018)
5. Bishop, C.M.: *Pattern Recognition and Machine Learning*. Springer, Singapore (2006)
6. Chang, J., Fisher III, J.W.: Efficient MCMC Sampling with Implicit Shape Representations. In: 2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2081–2088. IEEE (2011)
7. Denker, J., LeCunn, Y.: Transforming Neural-Net Output Levels to Probability Distributions. In: *Advances in Neural Information Processing Systems 3*, Citeseer (1991)
8. Fischl, B.: Freesurfer. *Neuroimage* 62, 774–781 (2012)
9. Gal, Y.: *Uncertainty in Deep Learning*. University of Cambridge, Cambridge (2016)
10. Garcia-Barnes, J., et al.: A normalized framework for the design of feature spaces assessing the left ventricular function. *Med. Imaging IEEE Trans.* **29**(3), 733–745 (2010)
11. Han, X., et al.: CRUISE: Cortical reconstruction using implicit surface evolution. *Neuroimage* 23, 997–1012 (2004)
12. Huo, Y. et al.: Consistent Cortical Reconstruction and Multi-Atlas Brain Segmentation. *NeuroImage* 138, 197–210 (2016)
13. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. *International Journal of Computer Vision* **1**(4), 321–331 (1988)
14. Kendall, A., Gal, Y.: What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision. In: *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, pp. 5574–5584. Curran Associates, (2017)
15. Kingma, D.P., Welling, M.: Auto-Encoding Variational Bayes. *ICLR*, (2014)
16. Kwon, Y., et al.: Uncertainty Quantification Using Bayesian Neural Networks in Classification: Application to Ischemic Stroke Lesion Segmentation. In: *Medical Imaging with Deep Learning (MIDL 2018)*, preprint, (2018)
17. Lê, M., Unkelbach, J., Ayache, N., Delingette, H.: Sampling Image Segmentations for Uncertainty Quantification. *Medical Image Analysis* 34, 42–51 (2016)
18. MacKay, D.J.C: A Practical Bayesian Framework for Backpropagation Networks. *Neural Computation* **4**(3), 448–472 (1992)
19. Milletari, F., Rothberg, A., Jia, J., Sofka, M.: Integrating Statistical Prior Knowledge into Convolutional Neural Networks. In: Descoteaux, M., et al. (eds.) *Medical Image Computing and Computer Assisted Intervention (MICCAI 2017)*. *Lecture Notes in Computer Science*, 10433, 161–168, (2017)
20. Neal, R.M.: *Bayesian Learning for Neural Networks*. University of Toronto, Toronto (1995)

21. Peressutti, D., et al.: A framework for combining a motion atlas with non-motion information to learn clinically useful biomarkers: Application to cardiac resynchronisation therapy response prediction. *Medical Image Analysis* 35, Pages 669–684 (2017)
22. Puyol-Antón, E. et al.: A multimodal spatiotemporal cardiac motion atlas from MR and ultrasound data. *Medical Image Analysis* 40, 96–110 (2017)
23. Schuh, A., et al.: A Deformable Model for the Reconstruction of the Neonatal Cortex. In: *IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, IEEE (2017)
24. Sinclair, M., Bai, W., Puyol-Anton, E., et al.: Fully automated segmentation-based respiratory motion correction of multiplanar cardiac magnetic resonance images for large-scale datasets. In: Descoteaux, M., et al. (eds.) *Medical Image Computing and Computer Assisted Intervention (MICCAI 2017)*. *Lecture Notes in Computer Science*, (2017)
25. UK BioBank Homepage, <https://www.ukbiobank.ac.uk/about-biobank-uk>. Last accessed 24 June 2018

Appendix A

4.1 Posterior

From (2)

$$p(y|z, x) = \mathcal{N}(y | \underbrace{US^{\frac{1}{2}}z + \mu + s(x)}_{:=A}, \underbrace{\sigma^2 I}_{:=\Sigma}), \quad (10)$$

and Jensen’s inequality we have

$$\ln p(y|x) \geq \int \ln[p(y|z)]p(z|x)dz = \mathbb{E}_{z|x} [\ln p(y|z)] \quad (11)$$

$$\cong \frac{1}{L} \sum_{l=1}^L \ln p(y|z_l), \quad (12)$$

where z_l is sampled from $p(z|x) = \mathcal{N}(z|\mu(x), \Sigma(x))$ and $\mu(x)$, $\Sigma(x)$ are provided by the network. In detail this translates to

$$\frac{1}{L} \sum_{l=1}^L \ln p(y|z_l) = \frac{1}{L} \sum_{l=1}^L -\frac{1}{2} \{r \ln(2\pi) + \ln |\Sigma| + (y - A)^T \Sigma^{-1} (y - A)\}, \quad (13)$$

r is the dimensionality of the vector y . In practice this value is summed over a batch of input vectors x_n . Note that to sample from $p(z|x)$ we employ the so-called ”reparametrisation trick” [15], where we first sample $\epsilon \sim \mathcal{N}(0, I)$ and then compute $z_l = \mu(x) + \Sigma^{1/2}(x) * \epsilon$.

4.2 Regularisation

Considering

$$p(z) = \int p(z|x)p(x)dx \quad (14)$$

$$\cong \sum p(z|x_n), \quad x_n \sim p(x), \quad (15)$$

the unitary Gaussian prior on z and the fact that

$$\text{KLD}(p||q) = \mathbb{E}_p \left[\ln \frac{p(z)}{q(z)} \right] \quad (16)$$

$$\cong \sum_l (\ln p(z_l) - \ln q(z_l)), \quad z_l \sim p \quad (17)$$

we write

$$\text{KLD} \left(\sum_{n=1}^N p(z|x_n), p(z) \right) \cong \frac{1}{L} \sum_l^L \left\{ \underbrace{\ln \left[\sum_{n=1}^N p(z_l|x_n) \right]}_{:=LNP} - \underbrace{\ln[\mathcal{N}(0, I)]}_{:=LNQ} \right\} \quad (18)$$

where $\mu_n = \mu(x_n)$ and $\Sigma_n = \Sigma(x_n)$. Furthermore,

$$LNP = \ln \left[\sum_n \frac{1}{\sqrt{((2\pi)^s |\Sigma_n|)}} e^{-\frac{1}{2}(z_l - \mu_n)^T \Sigma_n^{-1} (z_l - \mu_n)} \right] \quad (19)$$

and

$$LNQ = - \left[-\frac{s}{2} \ln(2\pi) - \frac{z^T z}{2} \right] = \frac{1}{2} [s \ln(2\pi) + z^T z], \quad (20)$$

with s equal to the dimensionality of the latent space.

Appendix B

We obtain the posterior distribution of our vertex coordinates as follows. Given

$$p(y|z, x) = \mathcal{N}(y|US^{\frac{1}{2}}z + \mu + s(x), \sigma^2 I), \quad (21)$$

and

$$p(z|x) = p(z|\mu(x), \Sigma(x)), \quad (22)$$

we can write

$$\begin{aligned} \mathbb{E}(y|x) &= \mathbb{E}(\mathbb{E}(y|z)|x) \\ &= \mathbb{E}(US^{\frac{1}{2}}z + \mu + s|x) \\ &= \int (US^{\frac{1}{2}}z + \mu + s)p(z|x)dz \\ &= US^{\frac{1}{2}}\mathbb{E}(z|x) + \mu + s \\ &= US^{\frac{1}{2}}\mu(x) + \mu + s. \end{aligned}$$

As for the variance,

$$\begin{aligned}
\text{var}(y|x) &= \mathbb{E}(\text{var}(y|z)|x) + \text{var}(\mathbb{E}(y|z)|x) \\
&= \sigma^2 I + \text{var}(US^{\frac{1}{2}}z + \mu + s|x) \\
&= \sigma^2 I + \text{var}(US^{\frac{1}{2}}z|x) \\
&= \sigma^2 I + US^{\frac{1}{2}}\text{var}(z|x)(US^{\frac{1}{2}})^T \\
&= \sigma^2 I + US^{\frac{1}{2}}\Sigma(x)(US^{\frac{1}{2}})^T.
\end{aligned}$$