

Supervised Uncertainty Quantification for Segmentation with Multiple Annotations

Shi Hu¹, Daniel Worrall¹, Stefan Knecht¹, Bas Veeling¹, Henkjan Huisman²,
and Max Welling¹

¹ University of Amsterdam

² Radboud University Medical Center

Abstract. The accurate estimation of predictive uncertainty carries importance in medical scenarios such as lung node segmentation. Unfortunately, most existing works on predictive uncertainty do not return *calibrated* uncertainty estimates, which could be used in practice. In this work we exploit multi-grader annotation variability as a source of ‘groundtruth’ aleatoric uncertainty, which can be treated as a target in a supervised learning problem. We combine this groundtruth uncertainty with a Probabilistic U-Net and test on the LIDC-IDRI lung nodule CT dataset and MICCAI2012 prostate MRI dataset. We find that we are able to improve predictive uncertainty estimates. We also find that we can improve sample accuracy and sample diversity. In real-world applications, our method could inform doctors about the confidence of the segmentation results.

Keywords: Uncertainty · Image segmentation · Deep learning.

1 Introduction

In recent years, deep learning has propelled the state of the art in segmentation in medical imaging [11,10,6,22,16]. However, previous works tend to focus on maximizing accuracy, ignoring predictive uncertainty. Modeling uncertainty at the per-pixel level is as important as accuracy, especially in medical scenarios, since it informs clinicians about the trustworthiness of a model’s outputs [21,13].

Typically, there are two main types of uncertainty one cares about, *aleatoric* and *epistemic* [15]. Aleatoric uncertainty is a measure of the intrinsic, irreducible noise found in data, usually associated with the data acquisition process. Epistemic uncertainty is our uncertainty over the true values of a model’s parameters, which arises from the finite size of training sets. With increasing training set size, epistemic uncertainty tends asymptotically to zero [8]. In practice, these two sources of uncertainty are difficult to quantify. Typically epistemic uncertainty is very hard to evaluate since one would need access to the groundtruth model to measure it, but it is possible to form a meaningful estimate of aleatoric uncertainty since we do have access to groundtruth data. Consider a training set of N images $\{x_i\}_{i=1}^N$. If for the i^{th} image we are able to acquire D grader

segmentations $\{y_i^1, \dots, y_i^D\}$, then we define aleatoric uncertainty to be the per-pixel variance among these segmentations $\mathbb{V}_{p(\mathcal{D})}[y_i] = \frac{1}{D} \sum_{j=1}^D (y_i^j - \bar{y}_i)^2$, where $\bar{y}_i = \frac{1}{D} \sum_{j=1}^D y_i^j$. Datasets containing multiple annotations exist in the literature, such as [1,2,7,18], and it is surprising that to date the authors cannot find examples of intergrader variability being exploited.

In this paper, we build a segmentation model based on the Probabilistic U-Net [16], exploiting intergrader variability as a target for aleatoric uncertainty. With this model one can draw diverse segmentations from its output and gain quantitative, calibrated aleatoric uncertainty estimates. We further add a source of epistemic uncertainty, which the model previously did not have. To view these two uncertainties we deploy an uncertainty decomposition in the output-space based on the law of total variance. We find improved predictive performance as well as better aleatoric uncertainty estimation over previous works, while also achieving higher sample diversity, which we did not explicitly design in.

2 Background And Related Works

Below we provide an overview of predictive uncertainty for deep learning.

Predictive Uncertainty Consider a training set $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ with inputs x_i and target segmentations y_i , and a neural network with parameters/hidden variables θ . We can think of the neural network as a conditional distribution $p(y|x, \theta)$. Given test image x_* , the posterior predictive distribution [19] is $p(y_*|x_*, \mathcal{D}) = \int p(y_*|x_*, \theta) p(\theta|\mathcal{D}) d\theta$, where $p(\theta|\mathcal{D})$ is a posterior distribution over θ given the training data. This quantity is intractable to find [4], so it is typically approximated by some $q_\lambda(\theta)$ from a tractable family of distributions, where λ is called the *variational parameters*. Typically the approximation is fitted by minimizing the reverse KL-divergence $\text{KL}[q_\lambda(\theta) \parallel p(\theta|\mathcal{D})]$. This is intractable, since it contains the intractable posterior term, but can be rearranged into the ELBO *evidence lower-bound* [4]:

$$\min_{\lambda} \text{KL}[q_\lambda(\theta) \parallel p(\theta|\mathcal{D})] = \max_{\lambda} \mathbb{E}_{q_\lambda(\theta)} [\log p(\mathcal{D}|\theta)] - \text{KL}[q_\lambda(\theta) \parallel p(\theta)] \quad (1)$$

where $p(\theta)$ is a prior on θ and $p(\mathcal{D}|\theta) = \prod_{i=1}^N p(y_i|x_i, \theta)$. The *predictive uncertainty* is the variance of the posterior predictive distribution.

Aleatoric and Epistemic Uncertainty The predictive uncertainty can be decomposed into two parts. By the law of total variance, we can write predictive variances as a sum of these two independent components:

$$\underbrace{\mathbb{V}_{p(y|x)}[y]}_{\text{predictive uncertainty}} = \underbrace{\mathbb{E}_{q_\lambda(\theta)} [\mathbb{V}_{p(y|\theta, x)}[y]]}_{\text{aleatoric uncertainty}} + \underbrace{\mathbb{V}_{q_\lambda(\theta)} [\mathbb{E}_{p(y|\theta, x)}[y]]}_{\text{epistemic uncertainty}}, \quad (2)$$

where we have used the notation \mathbb{E} and \mathbb{V} for the expectation and variance operator. We have labeled the two right-hand terms as *aleatoric* and *epistemic* uncertainty. The aleatoric term measures the average of the output variance

$\mathbb{V}_{p(y|\theta,x)}[y]$, under all settings of the variables θ . If $q_\lambda(\theta)$ were a delta peak, we would expect this term not to vanish and thus is it associated with aleatoric (data) uncertainty [21]. The epistemic term measures fluctuations in the mean prediction. These fluctuations exist because of uncertainty in the approximate posterior $q_\lambda(\theta)$. If $q_\lambda(\theta)$ were a delta peak, then this term would vanish to zero, and thus we associate it with epistemic (model) uncertainty [21,13].

Current techniques for estimating aleatoric and epistemic uncertainty follow similar line. In Tanno *et al.* [21] the authors treat MRI superresolution as a regression problem. They build a CNN directly outputting $\mathbb{E}_{p(y|\theta,x)}[y]$ and $\mathbb{V}_{p(y|\theta,x)}[y]$. They model epistemic uncertainty using variational dropout [14]. Bragman *et al.* [5] build on this technique, applying it to radiotherapy-treatment planning and multi-task learning. Concurrent to [21] Kendall and Gal proposed a similar method using Monte Carlo (MC) instead of variational dropout [9]. They also proposed a method which would work for classification, where they predict a mean and variance in the logit-space just before a sigmoid. Jungo *et al.* [12] estimate epistemic uncertainty in the context of postoperative brain tumor cavity segmentation using MC dropout [9]. In [3] Ayhan and Berens treat the data augmentation process as part of the approximate posterior $q_\lambda(\theta)$. They claim this is aleatoric uncertainty, but from their method it appears they really compute epistemic uncertainty. None of these works quantitatively evaluates the quality of the epistemic and aleatoric uncertainties. In this work, we show that the aleatoric uncertainty can indeed be measured.

The Probabilistic U-Net In the Probabilistic U-Net [16], the approximate posterior distribution is given the form $q_\lambda(z|x, y)$, where we have set $\theta = z$. The hidden variables are thus activations $z|x, y$ dependent on the training data. A (conditional) prior over z is given by a *prior network* $p_\lambda(z|x)$. To train this setup, the authors employ a variant of the ELBO with a β -weight on the KL-penalty

$$\max_{\lambda} \frac{1}{N} \sum_{i=1}^N (\mathbb{E}_{q_\lambda(z|x_i, y_i)} [\log p_\lambda(y_i|x_i, z)] - \beta \cdot \text{KL}[q_\lambda(z|x_i, y_i) \parallel p_\lambda(z|x_i)]) . \quad (3)$$

Again, λ represents the variational parameters to be optimized. Since at test time we do not have access to y , we use the prior network and Monte Carlo sample in $p(y_*|x_*, \mathcal{D}) = \int p_\lambda(y_*|x_*, z)p_\lambda(z|x_*)dz$. The specific form of the likelihood $p(y|x, z)$ can be found in the original paper [16]. This method is known to produce very diverse samples, from which we could estimate aleatoric uncertainty. In this paper, we endow the Probabilistic U-Net with a mechanism to estimate epistemic uncertainty and extend this method yet further, such that the aleatoric uncertainty estimates are automatically calibrated to the training set.

3 Method

We improve upon the Probabilistic U-Net model with two innovations. First, the original framework does not contain a mechanism to measure epistemic uncertainty. This can be included by adding variational dropout [14] after the

last convolution layer in the U-Net. This corresponds to setting $\theta = (z, w)$ and $q_\lambda(\theta) = q_\lambda(w)q_\lambda(z|x, y)$, where w are CNN weights. The objective defined in Eq. 3 then changes to:

$$\begin{aligned} \mathcal{L}_{\text{vd}}(\lambda) = & -\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{q_\lambda(w)q_\lambda(z|x_i, y_i)} [\log p_\lambda(y_i|x_i, z, w)] + \\ & \beta \cdot \frac{1}{N} \sum_{i=1}^N \text{KL}[q_\lambda(z|x_i, y_i) \parallel p_\lambda(z|x_i)] + \frac{1}{N} \text{KL}[q_\lambda(w) \parallel p(w)]. \end{aligned} \quad (4)$$

Notice that as N becomes very large the relative weight of the last KL term reduces, so the prior on w is ignored [8]. The objective is maximized when $q_\lambda(w)$ is a delta peak on the maximum likelihood parameters, corresponding to zero model uncertainty. For our second innovation, we use intergrader variability $\mathbb{V}_{p(\mathcal{D})}[y]$ as a training target for the predicted aleatoric uncertainty $\mathbb{E}_{q_\lambda(\theta)}[\mathbb{V}_{p(y|\theta, x)}[y]]$. We found directly minimizing the L_1 or L_2 distance between the two does not work well. Instead, since for binary variables the mean and variance are tied, we match the means $p_g = \mathbb{E}_{p(\mathcal{D})}[y]$ and $p_m = \mathbb{E}_{q_\lambda(\theta)}[\mathbb{E}_{p(y|\theta, x)}[y]]$ using a cross-entropy loss. This term is not part of the ELBO, so we are free to sample from the prior network since this is used at test-time. Introducing scaling coefficient γ our final training objective becomes

$$\mathcal{L}(\lambda) = \mathcal{L}_{\text{vd}}(\lambda) - \gamma \cdot \frac{1}{N} \sum_{i=1}^N [p_g \log p_m + (1 - p_g) \log(1 - p_m)]. \quad (5)$$

4 Experiments

Datasets and Implementation Details We use two datasets where images have different but plausible annotations. First, the LIDC-IDRI dataset with 4 lesion annotations per image [1,2,7]. This dataset contains 1,018 lung CT scans from 1,010 lung patients with manual lesion annotations. We use the LIDC Matlab Toolbox [17] to process the slices and annotations with dimension 512×512 , then center the lesions and crop the patches of size 128×128 . This results in 15,096 image patches in total. We do not change the in-plane resolution. Second, the MICCAI2012 dataset with 3 prostate peripheral zone annotations per image [18]. This dataset contains 48 prostate MRI images and each image has multiple slides. We discard images that have fewer than 3 annotations, which leaves 44 images in total. For each image, the original dimension of a slide is 320×320 , and we crop the central patch of size 128×128 . This results in 614 image patches in total. The patches are treated independently and we feed each 2D patch to a model. Since this is a very small dataset, we use elastic transformation [20] to augment the dataset to prevent overfitting.

For each dataset, we split the train/validation/test sets with ratio 70%/15%/15%. Different than [16], we put all annotations of an image in the same mini-batch. Table 1 shows some hyperparameters used for each dataset.

Hyperparameters	LIDC	MICCAI2012
# epochs	800	1000
Mini-batch size	32	12
β	1	100
γ	100	100
Data augmentation?	None	Elastic transformation
Adam learning rate	1e-6	1e-4

Table 1: Hyperparameters details.

Method	LIDC	MICCAI2012
Kohl et al. [16]	0.346 ± 0.038	0.382 ± 0.017
Kendall & Gal [13]	0.553 ± 0.010	0.571 ± 0.028
Ours	0.267 ± 0.012	0.373 ± 0.021

Table 2: D_{GED}^2 comparison (lower is better). Each result is computed over five random seeds.

The ones not presented in this table are similar to [16]. For ease of comparison, all experiments on the same dataset use the same hyperparameters. Lastly, we run all experiments on NVIDIA TITAN Xp GPUs.

Sample Accuracy and Diversity Figure 1 compares our generative results with Kendall and Gal model. In each plot, the first row is a patch in test and its true annotations. The second row is the samples generated by the Kendall and Gal model [13]. The third row is our samples. The annotations for most images exhibit some variability as they come from different graders. In general, we observe that the samples from our model are able to cover different modality in the annotations, whereas in Kendall and Gal there is limited diversity, and thus cannot cover all the variations in the true annotations.

Quantitatively, we evaluate the generative results using the generalized energy distance (or D_{GED}^2) metric [16]: $2\mathbb{E}[d(S, Y)] - \mathbb{E}[d(S, S')] - \mathbb{E}[d(Y, Y')]$, where d is the complement of the Intersection over Union (IoU): $d(a, b) = 1 - \text{IoU}(a, b)$. S and S' are independent samples from a model, and Y and Y' are independent samples from the graders. Thus, the first term measures the expected difference between the samples and annotations, the second among the samples themselves and the third among the annotations themselves. In other words, this metric evaluates both the accuracy and diversity of the samples. Table 2 compares the D_{GED}^2 scores on the LIDC and MICCAI2012 datasets. Since our model improves upon the Probabilistic U-Net model, we also present their numerical results for a reference³. Their generative results do not look very differently from ours, so we omit them in Figure 1. For each model in Table 2, we generate 50 samples for evaluation. The table shows our model achieves better D_{GED}^2 on both datasets.

³ We use the PyTorch implementation for the Probabilistic U-Net model from <https://github.com/stefanknegt/Probabilistic-Unet-Pytorch>.

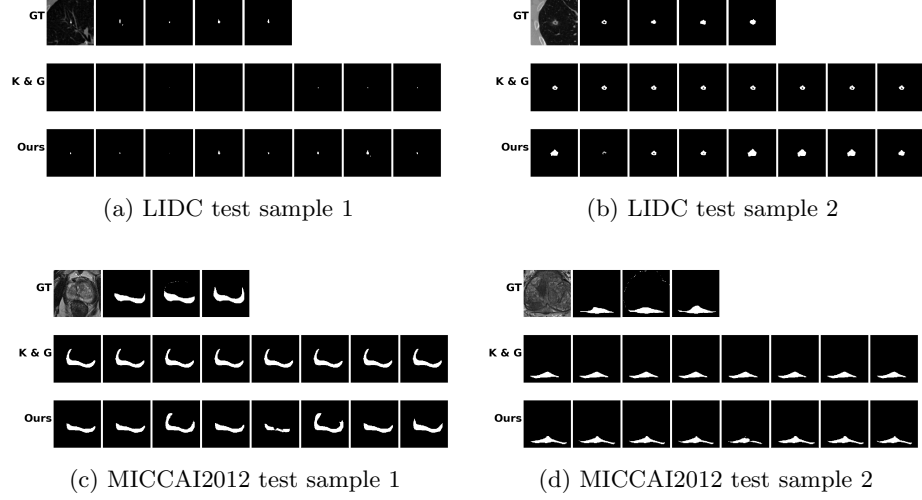


Fig. 1: Samples comparison (zoom in on detail).

Uncertainty Decomposition Figure 2 shows the aleatoric and epistemic uncertainty decomposition results. For each plot, the first row shows the results from Kendall and Gal [13] and the second row shows ours. To make the scales of these plots comparable, we set an upper threshold on the intensity values. Any value larger than the threshold will be treated as the threshold. For the true and predicted data uncertainty plots, we use the same threshold as we want to visually compare their similarity. In contrast, there is no label for the epistemic uncertainty, so we use a scale that fits well with most intensity pixels in a plot.

In general, we observe that Kendall and Gal makes plausible predictions on the shape of the data uncertainty, but tends to underestimate its scale, whereas we are relatively close to the ground truth in terms of both the shape and scale. Furthermore, the former tends to have high model uncertainty, especially at the image borders, whereas ours are usually around the point of interest. Although we do not know the true appearance of the model uncertainty, it should be high on the objects that do not occur often in the training set. In both test sets, the new objects usually appear around the center rather than at the borders. Therefore, we argue our epistemic uncertainty prediction is more sensible.

Quantitatively, Table 3 compares the data uncertainty prediction performance of the two models. As mentioned, the scale of the data uncertainty predictions from Kendall and Gal tends to be smaller than the ground truth. We want to establish a fair image similarity comparison that takes into account of this fact. Thus, for the true and predicted data uncertainty map a and b , we measure their similarity using the normalized cross-correlation $\frac{1}{n\sigma_a\sigma_b} \sum_{x,y} (a_{x,y} - \mu_a) \cdot (b_{x,y} - \mu_b)$, where n is the total number of pixels in an uncertainty map, and μ and σ are the mean and the standard deviation of a uncertainty map. Since we normalize the scales of the uncertainty maps a and

Method	LIDC	MICCAI2012
Kendall & Gal [13]	0.597 ± 0.006	0.299 ± 0.011
Ours	0.669 ± 0.011	0.345 ± 0.005

Table 3: Data uncertainty prediction comparison using the normalized cross-correlation (higher is better). Each result is computed over three random seeds.

b , the output value represents their intrinsic similarity. In Table 3, we report the average normalized cross-correlation score over all test images for each dataset. Our model achieves higher data uncertainty correlations in both cases.

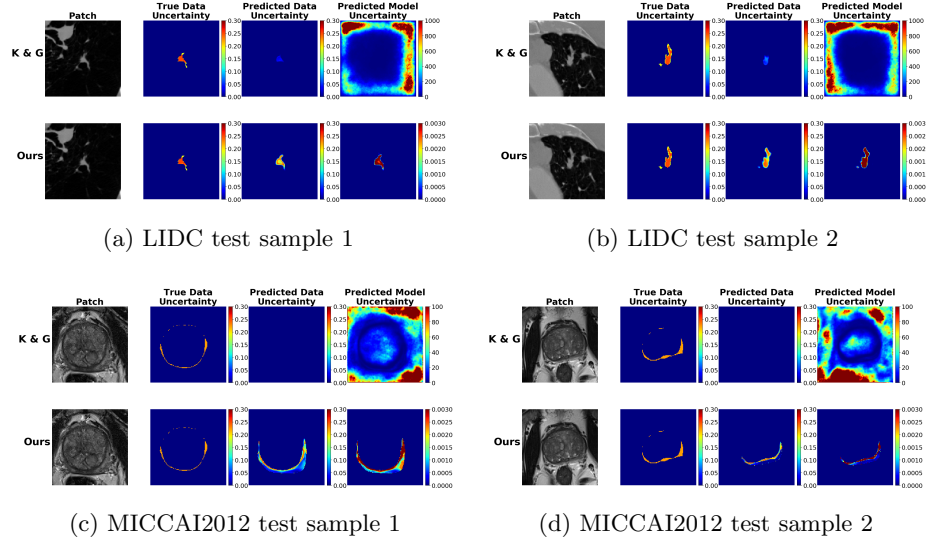


Fig. 2: Uncertainty quantification comparison.

5 Conclusions

In this work we designed a model for segmentation based on the Probabilistic U-Net [16] which outputs two kinds of quantifiable uncertainty, aleatoric (data) uncertainty and epistemic (model) uncertainty. We leveraged intergrader variability as a target for calibrated aleatoric uncertainty, which, as far as we know, related works have surprisingly not used. We showcased our model on the LIDC-IDRI lung nodule CT dataset [1,2,7] and MICCAI2012 prostate MRI dataset [18], demonstrating that we could improve predictive uncertainty estimates. We also found that we could improve sample accuracy and sample diversity. In real-world applications, our method could inform doctors if they can trust the segmentation

results to decide in real-time whether to continue a treatment path or not. As future work, we would like to improve the quality of the epistemic uncertainty.

Acknowledgements

We thank Dimitrios Mavroeidis for helpful discussions and Arsenii Ashukha for the variational dropout code. This research was supported by NWO Perspective Grants DLMedia and EDL, as well as the in-cash and in-kind contributions by Philips.

References

1. Armato, S.G., McLennan, G., et al., L.B.: The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Med Phys.* (2011)
2. Armato, S.G., McLennan, G., Bidaut, L., McNitt-Gray, M.F., Meyer, C.R., Reeves, A.P., Clarke, L.P.: Data from lidc-idri. The Cancer Imaging Archive (2015)
3. Ayhan, M.S., Berens, P.: Test-time data augmentation for estimation of heteroscedastic aleatoric uncertainty in deep neural networks. *MIDL* (2018)
4. Blei, D.M., Kucukelbir, A., McAuliffe, J.D.: Variational inference: A review for statisticians. *CoRR* **abs/1601.00670** (2016)
5. Bragman, F.J., Tanno, R., Eaton-Rosen, Z., L, W., Hawkes, D.J., Ourselin, S., Alexander, D.C., McClelland, J.R., Cardoso, M.J.: Quality control in radiotherapy-treatment planning using multi-task learning and uncertainty estimation. *MIDL* (2018)
6. Causey, J., Zhang, J., Ma, S., Jiang, B., Qualls, J., Polite, D.G., Prior, F.W., Zhang, S., Huang, X.: Highly accurate model for prediction of lung nodule malignancy with CT scans. *CoRR* **abs/1802.01756** (2018)
7. Clark, K., Vendt, B., Smith, K., Freymann, J., Kirby, J., Koppel, P., Moore, S., Phillips, S., Maffitt, D., Pringle, M., Tarbox, L., Prior, F.: The cancer imaging archive (tcia): Maintaining and operating a public information repository. *Journal of Digital Imaging* (2013)
8. Gal, Y.: Uncertainty in deep learning. PhD thesis, University of Cambridge (2016)
9. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. *ICML* (2016)
10. Gruetzemacher, R., Gupta, A., Paradise, D.B.: 3d deep learning for detecting pulmonary nodules in CT scans. *JAMIA* (2018)
11. Gu, Y., Lu, X., Yang, L., Zhang, B., Yu, D., Zhao, Y., Gao, L., Wu, L., Zhou, T.: Automatic lung nodule detection using a 3d deep convolutional neural network combined with a multi-scale prediction strategy in chest cts. *Comp. in Bio. and Med.* (2018)
12. Jungo, A., Meier, R., Ermis, E., Herrmann, E., Reyes, M.: Uncertainty-driven sanity check: Application to postoperative brain tumor cavity segmentation. *MIDL* (2018)
13. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? *NIPS* (2017)
14. Kingma, D.P., Salimans, T., Welling, M.: Variational dropout and the local reparameterization trick. *NIPS* (2015)

15. Kiureghian, A.D., Ditlevsen, O.: Aleatory or epistemic? does it matter? *Structural Safety* (2009)
16. Kohl, S.A., Romera-Paredes, B., Meyer, C., De Fauw, J., Ledsam, J.R., Maier-Hein, K.H., Eslami, S., Rezende, D.J., Ronneberger, O.: A probabilistic u-net for segmentation of ambiguous images. *NIPS* (2018)
17. Lampert, T.A., Stumpf, A., Gancarski, P.: An empirical study of expert agreement and ground truth estimation. *IEEE Transactions on Image Processing* (2016)
18. Litjens, G., Debats, O., van de Ven, W., Karssemeijer, N., Huisman, H.: A pattern recognition approach to zonal segmentation of the prostate on mri. *MICCAI* (2012)
19. MacKay, D.J.C.: Bayesian interpolation. *Neural Computation* **4**(3), 415–447 (1992)
20. Simard, P.Y., Steinkraus, D., Platt, J.C.: Best practices for convolutional neural networks applied to visual document analysis. *International Conference on Document Analysis and Recognition* (2003)
21. Tanno, R., Worrall, D.E., Ghosh, A., Kaden, E., Sotiropoulos, S.N., Criminisi, A., Alexander, D.C.: Bayesian image quality transfer with cnns: Exploring uncertainty in dmri super-resolution. In: *MICCAI* (2017)
22. Wang, S., Zhou, M., Liu, Z., Liu, Z., Gu, D., Zang, Y., Dong, D., Gevaert, O., Tian, J.: Central focused convolutional neural networks: Developing a data-driven model for lung nodule segmentation. *Medical Image Analysis* (2017)