JULY 12–16 2020 VIRTUAL
JOINT AAPM | COMP MEETING
EASTERN TIME [GMT-4]

Towards safer artificial intelligence-based treatment planning: adding uncertainty estimation to volumetric dose prediction using an approximate Bayesian method on deep neural networks

Dan Nguyen, Anjali Balagpal, Azar Barkousaraie, Rafe McBeth, Steve Jiang
Medical Artificial Intelligence and Automation (MAIA) Laboratory,
Department of Radiation Oncology,
UT Southwestern Medical Center, Dallas, Texas
Dan.Nguyen@UTSouthwestern.edu

MAIA Laboratory
UT Southwestern Medical Center
Radiation Oncology

INTRODUCTION

The field of radiation therapy is ever-changing, with technologies and distributions of patient populations evolving over time. One of the biggest advancements in recent years is the development of artificial intelligence technologies, and their application into radiation therapy. However, as we move forward, concerns regarding safety in using the model on patients has risen greatly. In this study we show how to add in an uncertainty estimation to a deep learning model's prediction—a way for the model to say "I don't know" when given data that is unlike that it has seen before—and we apply this method to dose prediction.

MONTE CARLO DROPOUT AS A BAYESIAN APPROXIMATION

Training phase

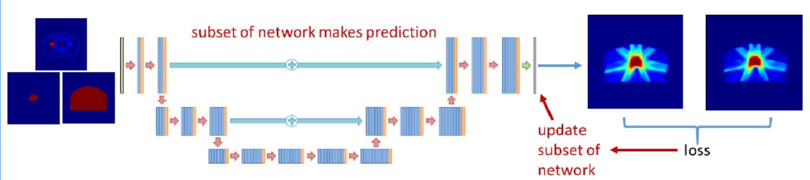

Inference phase

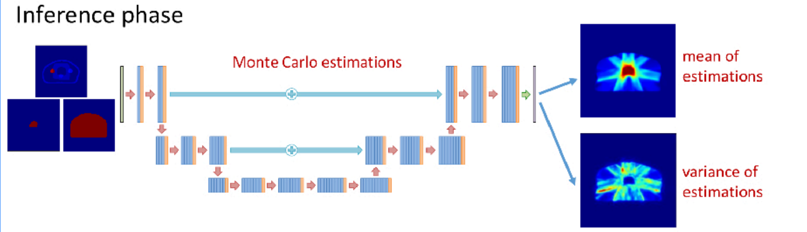

The prediction estimation and variance are defined as

Estimation:

$$\mathbb{E}_{q(y^*|x^*)}(y^*) \approx \frac{1}{T}\sum_{t=1}^{T}\hat{y}^*(x^*, W_1^t, \cdots, W_L^t)$$

Variance:

$$Var_{q(y^*|x^*)}(y^*) \approx \frac{1}{T}\sum_{t=1}^{T}\hat{y}^*(x^*, W_1^t, \cdots, W_L^t)^T\, \hat{y}^*(x^*, W_1^t, \cdots, W_L^t) - \mathbb{E}_{q(y^*|x^*)}(y^*)^T\mathbb{E}_{q(y^*|x^*)}(y^*)$$

where $\hat{y}^*(x^*, W_1^t, \cdots, W_L^t)$ is the trained model's prediction given an input x and a t set of weights $W_1^t, \cdots, W_L^t$.

DEEP LEARNING ARCHITECTURE

U-net Architecture

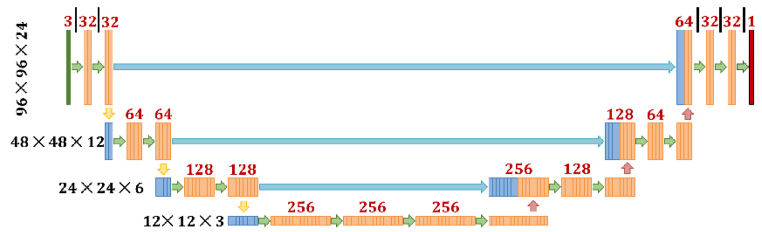

Legend

- 3 × 3 × 3 Padded Convolution, Activation, Group Norm, Dropout
- 2 × 2 × 2 Upsample, 3 × 3 × 3 Padded Convolution, Activation, Group Norm, Dropout
- Concatenate
- 2 × 2 × 2 Maxpool

DATA AND TRAINING

- 70 Prostate CTs with their segmentation
  - 54 training, 6 validation, 10 test patients
  - 96 x 96 x 24 array
  - 5 mm x 5 mm x 5 mm voxel size
  - 1200 Pareto optimal plans per patient
  - 84000 plans total
- Training the network
  - Mean squared error (MSE)
  - Dropout set to 0.125 throughout the network
- Scaling network uncertainty to the voxel error
  - Due to fact that uncertainty is not directly correlated to voxel error (e.g. uncertainty → 0 as dropout → 0)
  - We define a scaling factor, m, such that
    - $m * uncertainty \geq 95\%\ of\ validation\ data$

RESULTS

Scaling Uncertainty to Voxel Error

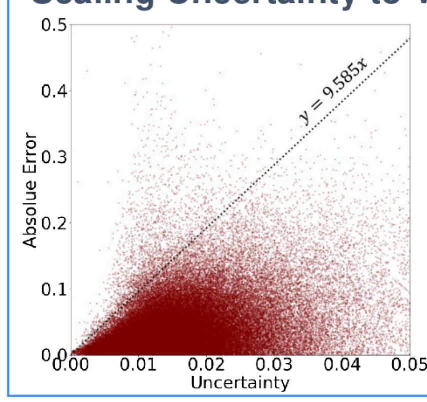

Figure 1: Plot of voxel-wise absolute error $|true - predicted|$ versus the uncertainty. The uncertainty is based off of the standard deviation of 50 Monte Carlo model predictions, and its value is dependent on the amount of Dropout used in the model (uncertainty → 0 as dropout → 0). We correlate our uncertainty to our error, by finding the slope, m, of a line $y = mx$, such that 95% of the validation data points are below the uncertainty. We solve $m = 9.585$. On the test data, we find that 97.49% of the error is below the dotted line. This can be used to characterized the lower and upper bounds of our prediction, shown later in Figure 4. High uncertainty does not necessarily equate to high error, but it means it is more *likely* to have high error.

Prediction and Uncertainty

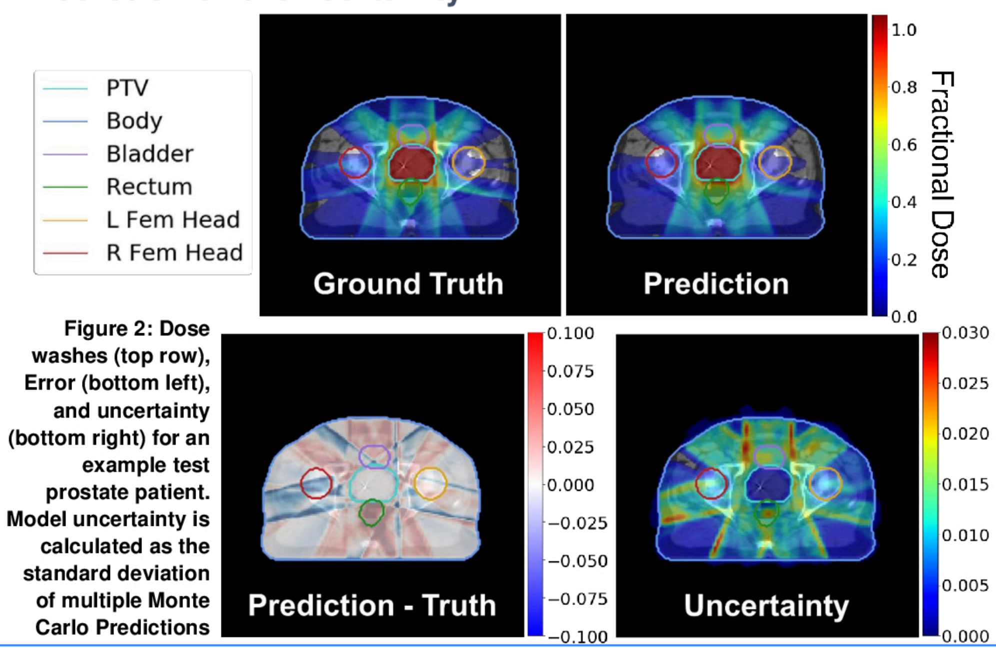

Figure 2: Dose washes (top row), Error (bottom left), and uncertainty (bottom right) for an example test prostate patient. Model uncertainty is calculated as the standard deviation of multiple Monte Carlo Predictions

DVH Uncertainty

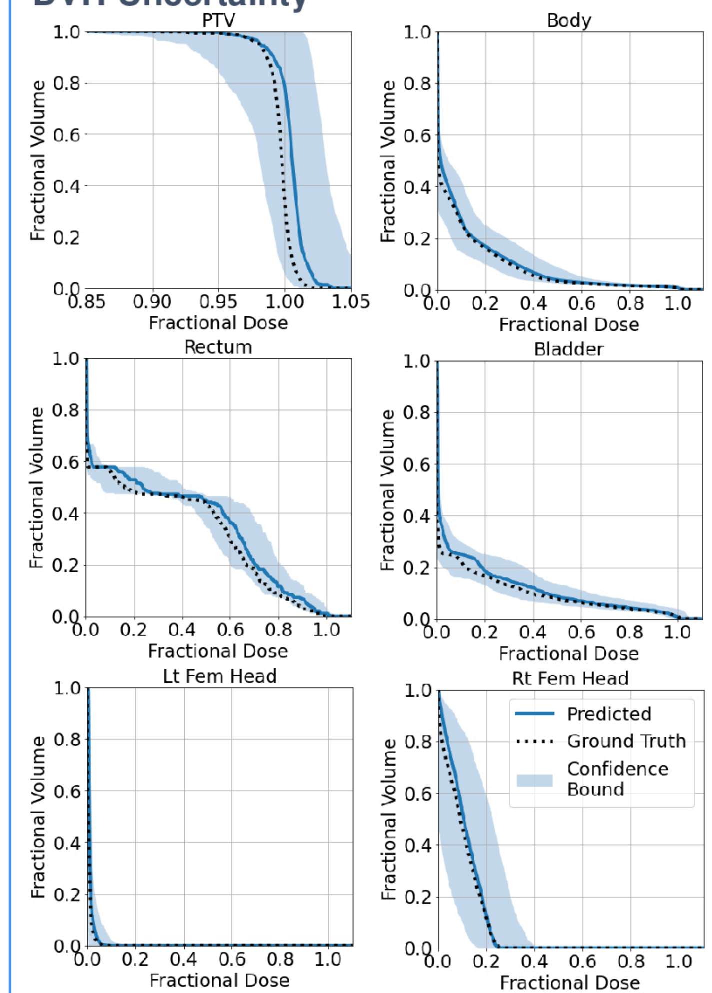

Figure 3: Comparison of predicted (solid blue) vs ground truth (dotted black) for an example test patient. Confidence bounds are calculated using $prediction \pm m * uncertainty$, where $m = 9.585$.

CONCLUSIONS

- We show a method that can be used so that the uncertainty of deep learning models can be obtained.
- We characterize a curve to relate the uncertainty to the error in the model.
- This work can be used to greatly improve the safety of model implementation in a clinical setting.
- This uncertainty represents the precision in the data as well as the model's own learning/generalization error

REFERENCES

[1] Gal, Y., & Ghahramani, Z. (2016, June). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In international conference on machine learning(pp. 1050-1059).

[2] Balagopal, A., Nguyen, D., Morgan, H., Weng, Y., Dohopolski, M., Lin, M. H., ... & Hannan, R. (2020). A deep learning-based framework for segmenting invisible clinical target volumes with estimated uncertainties for post-operative prostate cancer radiotherapy. arXiv preprint arXiv:2004.13294.