Title: THE USE OF HIERARCHICAL MODELING IN PROSTATE STEREOTACTIC BODY RADIATION THERAPY

Background: Prostate SBRT demands high geometric precision due to large fractional doses and steep dose gradients. Conventional PTV margins are typically derived from population-level estimates and may not reflect individual patient-specific motion. This thesis introduces a Bayesian hierarchical framework to simulate intrafraction motion and evaluate its dosimetric consequences, with emphasis on anatomically defined posterior margins used in clinical practice at OHSU.

Methods: A cohort of 205 patients treated with five-fraction prostate SBRT from 2020 to 2024 was retrospectively analyzed. Mid-treatment shifts were extracted from CBCT records and used to represent unintentional motion during treatment. These data trained a multivariate Student-T model with partial pooling, capturing population-level trends while retaining individual variability. Posterior predictive sampling generated synthetic shifts, which were applied to patient dose grids using rigid transformations. Cumulative dose distributions were recalculated and evaluated using standard DVH metrics.

Results: Simulated shifts reproduced the central tendency, spread, and directional correlations of recorded patient data. Failure rates under clinical margins closely matched those in patient records, and correlation structures between axes were preserved. Case studies confirmed directional fidelity and dose displacement, particularly in high-motion scenarios. Population-level analysis showed close agreement in DVH deviations across prostate, bladder, and rectum. For instance, mean anterior shifts were nearly identical (–0.151 mm patient vs. –0.119 mm simulated), and anterior–superior correlations were similar (r=0.422r = 0.422r=0.422 vs. r=0.357r = 0.357r=0.357, R2=0.178R^2 = 0.178R2=0.178 vs. 0.127). Most DVH differences were small and statistically insignificant.

Conclusion: This modeling framework offers a flexible approach to reconstructing patient-specific motion, even with limited empirical data. By leveraging partial pooling, it balances individual detail with population-informed inference. Its ability to replicate real-world distributions, failure rates, and dose perturbations underscores its validity as a simulation engine. Importantly, this approach enables probabilistic “what-if” testing without requiring additional imaging or patient intervention—offering a principled method for stress-testing clinical practices. In doing so, it advances the clinical conversation around personalization in radiation therapy and highlights the role of statistical modeling in improving treatment safety and precision.