A GPU-Based Multi-Criteria Optimization Algorithm for LDR Brachytherapy

Olivier Charette[1,2], Sylviane Aubin[2], Luc Beaulieu[1,2]

olivier.charette.4@ulaval.ca

[1] Département de Physique, de Génie Physique et d'Optique et Centre de Recherche sur le Cancer, Université Laval, Québec, Canada

[2] Département de Radio-Oncologie et Axe Oncologie du CRCHU de Québec, CHU de Québec, Université Laval, Québec, Canada

CONTEXT

Low-dose-rate (LDR) prostate brachytherapy uses permanent radioactive seed implants to treat the tumor. Treatment planning is a multi-criteria optimization (MCO) problem that balances target coverage against dose to the urethra, rectum, and bladder.

Conventional inverse-planning methods generate a single plan based on predefined dose objectives and priorities.

- Result is sensitive to these choices;
- Users need to adjust the optimization parameters;
- Process repeated until a valid plan is obtained.

This iterative tuning is time-consuming, user-dependent, and does not guarantee that the most clinically suitable plan is found.

Building on the original gMCO framework[1], gMCO-seed generates up to 1000 plans in one minute using GPU-based optimization,

➢ rapid exploration of clinical trade-offs while reducing planning time and user dependence.

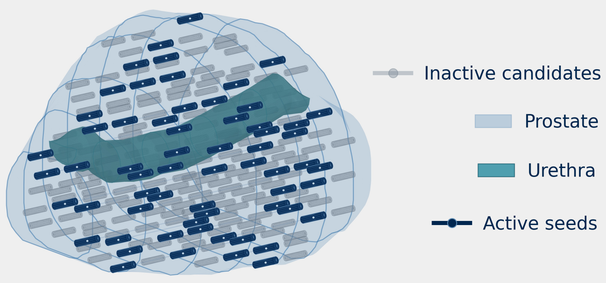

FIG. 1 The prostate LDR planning problem. The optimizer chooses which seed positions to turn on.

OBJECTIVES

1. Develop and benchmark gMCO-seed, a GPU multi-criteria optimizer for LDR prostate brachytherapy.

2. Raise the yield of clinically acceptable plans with patient-specific cost functions.

METHODS

Optimizer. gMCO-seed: GPU simulated annealing (gSA), TG-43 dose[2], I-125 SelectSeed, NVIDIA RTX 5090 GPU.

Cohort. 200 retrospective cases: 3000 random hidden-weight draws per case at each of three target dose-max objectives (259 / 230 / 216 Gy), locating every case's optimal weight center and spread.

Model. Hidden weights are a composition

- Statistics use isometric log-ratio (ilr) coordinates: ilr(w) = f(geometry) + g(dose), ridge regression[3] on 31 geometry features.
- Scored by 10-fold grouped cross-validation.

Validation. 60 cases re-optimized from scratch, 1000 plans/case, each using a center predicted out-of-fold by a model that never saw that patient.

Yield. Percentage of generated plans that meet all institutional objectives (see Table 1).

Table 1 Pre-implant institutional objectives

Metric	Objective	Metric	Objective
Target V100	**> 98%**	Urethra D5	**< 205 Gy**
Target V150	**62-68%**	Urethra V150	**0%**
Target V200	**30-35%**	DIL V150	**> 95%**
Target D90	**180-190 Gy**		

RESULTS & DISCUSSION

1 PERFORMANCE BENCHMARK

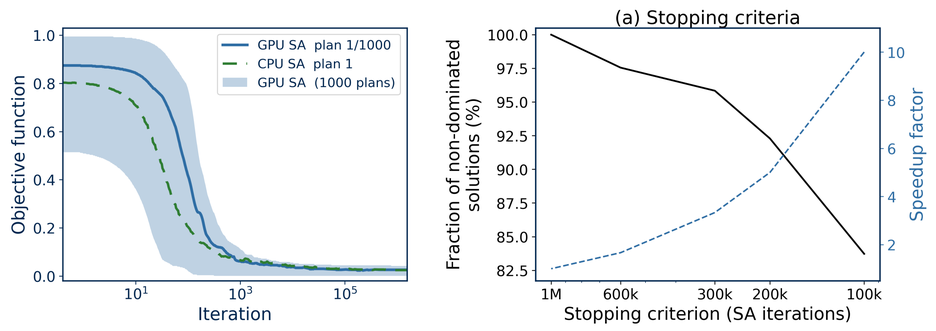

FIG. 2 Objective-function convergence for 1000 GPU plans versus a single CPU plan (left). Tradeoff between Pareto optimality and speedup factor of the optimization time (right). Mean time to generate 1000 plans (300k iterations/plan) : (62 ± 5) s.

2 OBJECTIVE FUNCTION AND CLASS SOLUTIONS

$$f_{ij}(d_{ij}) = \begin{cases} w_j^{\min}\left(D_j^{\min} - d_{ij}\right) & d_{ij} < D_j^{\min} \\ 0 & D_j^{\min} \le d_{ij} \le D_j^{\max} \\ w_j^{\max}\left(d_{ij} - D_j^{\max}\right) & d_{ij} > D_j^{\max} \end{cases}$$

$$F = \sum_{j=1}^{N_O} \frac{1}{N_{\mathrm{pnt},j}} \sum_{i=1}^{N_{\mathrm{pnt},j}} f_{ij}(d_{ij})$$

Notation

d_{ij}	point dose
$D_j^{\min}$, $D_j^{\max}$	dose-window limits
$w_j^{\min}$, $w_j^{\max}$	tunable penalty slopes
$N_{\mathrm{pnt},j}$	points in organ
N_O	optimized organs
f_{ij}	point penalty
F	plan cost

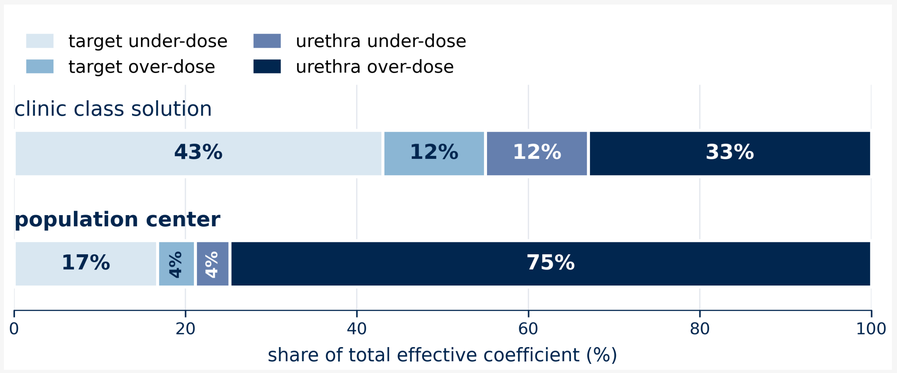

FIG. 3 Point-wise dose penalty and plan cost F (top). Effective coefficients of the two class solutions (bottom): the population center is urethra-dominated.

At our clinic, the current cost function targets the wrong region. A urethra-heavy cost function better controls overdoses and directs plans into the clinical objective box. Lowering the target dose maximum from 259 Gy (180% Rx) to 216 Gy (150% Rx) further improves overdose control and plan yield.

3 PATIENT-SPECIFIC SAMPLING

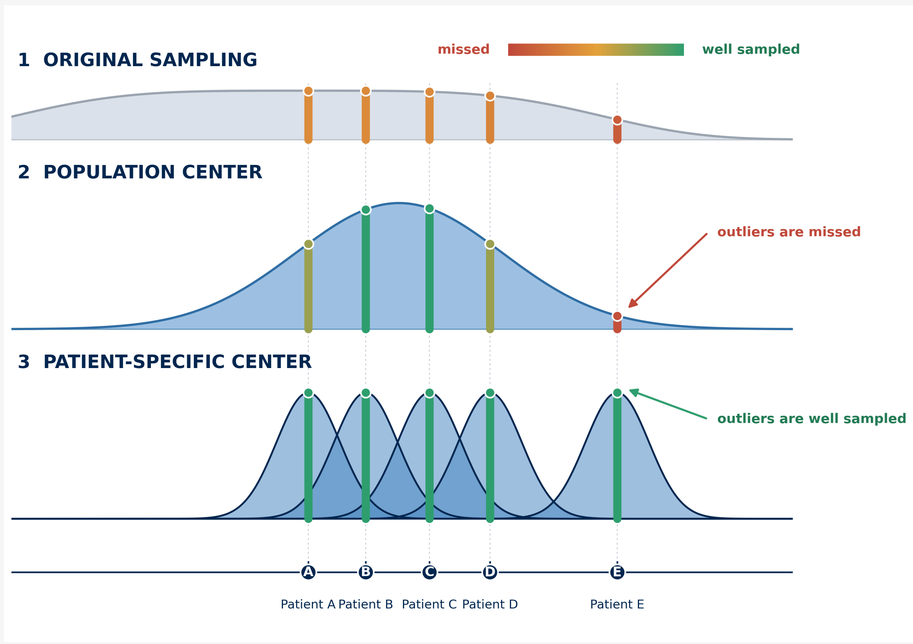

FIG. 4. Effect of tunable penalty-slope $w_j^{\min}$, $w_j^{\max}$ sampling on plan-yield robustness. Patient-specific centering ensures consistent sampling across the cohort, including for outlier patients.

4 EFFECT OF THE SAMPLING WIDTH

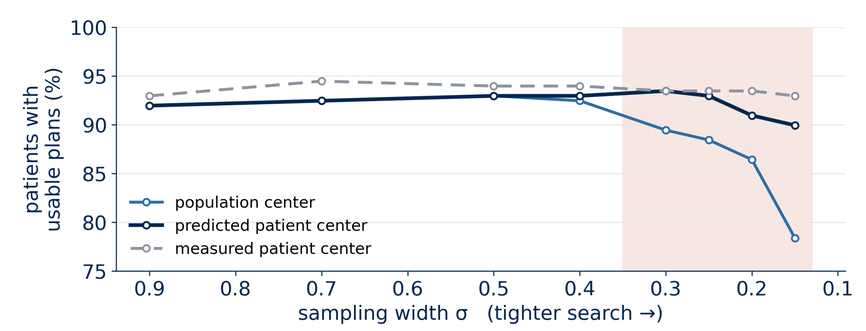

FIG. 5 Fraction of patients with usable plans versus sampling width σ. Predicted centering holds 90% where population centering falls to 78%.

5 PREDICTING THE RIGHT CLASS SOLUTION

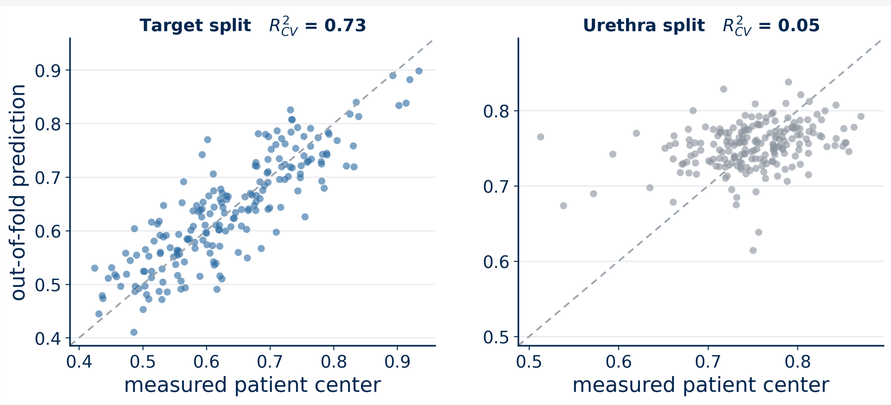

FIG. 6 Out-of-fold predicted versus measured weight center. Geometry captures much of the variation in target split (R^2 = 0.73) but not the urethra split (R^2 = 0.05).

6 YIELD OF ACCEPTABLE PLANS

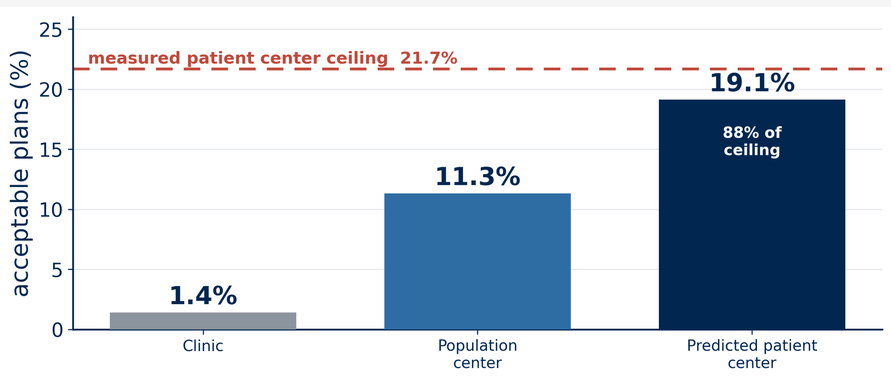

FIG. 7 Yield of clinically acceptable plans on the 60 held-out cohort. Predicted per-patient centering reaches 19.1%, compared to only 1.4% when using the original sampling and starting from the clinical class solution.

CONCLUSIONS

1. gMCO-seed generates 1000 plans in 62 s
2. The clinic's weight distribution is not adapted for MCO. Re-centering the class solution for the population and changing the sampling distribution raises yield from 1.4% to 11.3%.
3. Prostate geometry predicts the target weight split. A predicted-centered class solution reaches a yield of 19.1% : 88% of the 21.7% obtained by optimizing each patient first.

Next steps

➢ Evaluate optimization behavior under strict clinical constraints

➢ Optimization on Monte Carlo-based dose grids instead of TG-43

REFERENCES

1. **Bélanger C, et al.** *Phys Med Biol.* 2019;64:105005. doi:10.1088/1361-6560/ab1817.
2. **Rivard MJ, et al.** *Med Phys.* 2004;31:633–674. doi:10.1118/1.1646040.
3. **Pedregosa F, et al.** *J Mach Learn Res.* 2011;12:2825–2830.