AbstractID: 9105 Title: How many plans are required to represent Pareto surfaces in multi-objective IMRT planning?

Purpose:
To demonstrate that the number of IMRT plans needed to fully represent the trade-off space in multi-objective optimization is small (<50).
Method and Materials:
For clinical IMRT cases, we compute a database of Pareto optimal plans using established methods. A treatment plan is Pareto optimal if there does not exist another plan which is better in at least one objective and at least as good in every objective. We analyze Pareto optimal plan databases by three methods. The first method examines the points as they are added sequentially, and for each newly added point we determine how "new" that point is (i.e. how far that point is from the convex hull of the previously generated points). Principal component analysis (PCA) on the set of beamlet solutions is used to understand the shape of the Pareto surface in beamlet space. Correlations between objectives are found by computing Spearman correlation coefficients on each pair of objectives.
Results:
We compute Pareto surfaces from 3 to 10 dimensions for pancreas, prostate, and brain cases. For all cases, we find that after about 40 plans, subsequently generated plans are very close to what can be achieved by averaging the previous plans. Therefore, allowing for averaging of plans, which is possible if the plans have not been sequenced yet, the entire Pareto surface can be well approximated by a small number of plans.
Conclusion:
Since the shape of the feasible set is relatively simple, as demonstrated by PCA, the Pareto surface is simple to describe, regardless of the number or types of objectives used. Once plans that span this space are computed, further plans do not add to the diversity of solutions available. Based on the analyses, we offer a simplified strategy for finding a small set of spanning Pareto optimal plans.