The Gap between Visualization Research and Visualization Software (VisGap) (2020)
C. Gillmann, M. Krone, G. Reina, T. Wischgoll (Editors)

Lessons Learnt from Developing Visual Analytics Applications for Adaptive Prostate Cancer Radiotherapy

R.G. Raidou[1], K. Furmanova[1,2], N. Grossmann[1], O. Casares-Magaz[2], V. Moiseenko[3], J.P. Einck[3], M.E. Gröller[1], L.P. Muren[2]

[1]TU Wien, Austria, [2]Department of Medical Physics, Aarhus University Hospital, Denmark
[3]Department of Radiation Medicine and Applied Sciences, UC San Diego, United States

Abstract

In radiotherapy (RT), changes in patient anatomy throughout the treatment period might lead to deviations between planned and delivered dose, resulting in inadequate tumor coverage and/or overradiation of healthy tissues. Adapting the treatment to account for anatomical changes is anticipated to enable higher precision and less toxicity to healthy tissues. Corresponding tools for the in-depth exploration and analysis of available clinical cohort data were not available before our work. In this paper, we discuss our on-going process of introducing visual analytics to the domain of adaptive RT for prostate cancer. This has been done through the design of three visual analytics applications, built for clinical researchers working on the deployment of robust RT treatment strategies. We focus on describing our iterative design process, and we discuss the lessons learnt from our fruitful collaboration with clinical domain experts and industry, interested in integrating our prototypes into their workflow.

CCS Concepts
• ***Human-centered computing*** → ***Visual analytics;*** • ***Applied computing*** → ***Life and medical sciences;***

1. Introduction

Radiotherapy (RT) is a common therapeutic approach for prostate cancer [DJFB05]. It delivers high radiation doses to tumors, while preserving the adjacent healthy tissues and minimizing radiation-induced toxicity for the patient. Determining the adequate intensity and distribution of the dose is not trivial [SRM*19], and the so-called treatment plan needs to be optimised in dedicated software prior to treatment delivery. This lengthy process considers the locations and characteristics of the tumors and of the healthy surrounding organs. Then, the prescribed dose is administered to the patient—not all at once, but in multiple fractions over several weeks, to allow the recovery of healthy tissue [WL15].

During the treatment period, the involved pelvic organs (e.g., bladder and rectum) may exhibit high day-to-day anatomical variations, due to differences in urinary or rectal filling, due to weight loss or due to other tissue and anatomical characteristics [CMMH*17a, CMB*18, MSD03, CvHvdK*11]. Additionally, the pelvis anatomy presents unique variations in every human, which can be naturally varying across individuals [RDCO*17]. The standard treatment procedure is to generate one initial treatment plan per patient and to use it as a basis for all subsequent sessions [SRM*19, WL15], as real-time adaptation is not possible due to time constraints. Only alignment corrections are made before dose administration, prioritizing tumor coverage. For the rest, only considerable variations, such as empty instead of full bladder, triggers changes to the initial treatment plan. Recent clinical research suggests that anatomical variability can lead to increased radiation doses being delivered to healthy organs, such as the bladder or the rectum [CMMH*17a,MLK*07,CMB*18]. The necessity for a more drastic adjustment of the target volume in prostate cancer therapy on a per-treatment basis has been highlighted by several recent works [VYM*10,CvHvdK*11,RDCO*17].

Prostate cancer research starts looking into adaptive treatment approaches [THLM*13] that will take into account the shape variability and movement of the pelvic organs of the patient through treatment. Within this research, there is a hypothesis that certain subgroups of patients, e.g., patients with highly variable organ shapes, might be at a higher risk of manifesting side-effects, affecting their quality of life. To this end, retrospective cohort studies are conducted [CMMH*17a,MLK*07,CMB*18] and are expected to give information about potential subgroups of the population that face advanced risks of toxicity. However, no previous clinical studies have been able to relate organ variability and potential toxicity, mainly due to a lack of exploratory means.

Many visualization methods have resulted in a large collection of prototypes, visual designs, techniques and development kits for RT—without being adopted in clinical practice. Schlachter et al. [SRM*19] documented that older works in dose plan review and evaluation incorporating 2D representations [HST87, LFP*90, MYKO91, IFP95] are well-established in clinical routine. Yet,

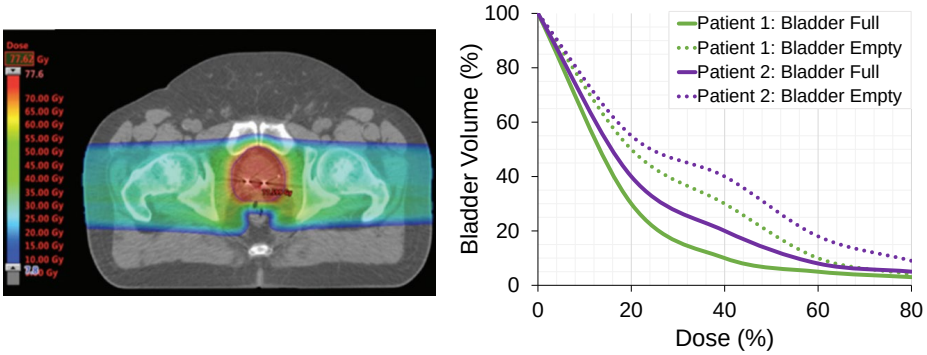

Figure 1: Left: *Anatomical 2D view on a radiation therapy plan of one patient.* Right: *Dose Volume Histogram (DVH) of the bladders of two patients for two treatment regimes (empty and full bladders).*

newer 3D approaches [GDF*00, FC16, GCC*16] have remained at a developmental stage, while visual analytics is still at its infancy within RT applications. Also, adaptive approaches incorporating anatomical feature analysis are also at a developmental stage—with few exceptions [WRH*08].

In this work, we showcase a successful example of bridging the gap between visualization and RT, by developing applications to improve RT in prostate cancer. Our process started with an application for the investigation of the impact of the shape variation of an individual organ (bladder) on the dose distribution and toxicity risk during the course of RT [RCMA*18]. Subsequently, the domain experts conducted a clinical study, using our application and additional statistical analysis extensions, to obtain the first clinical indications for bladder toxicity risk with regard to organ shape [CMRP*19]. The next step was to create an extension for exploring the variability of several segmented pelvic organs in multiple patients [GCMM*19]. This was followed by an additional clinical study [FRG*20], which was further extended to include also functionality for dose variability information exploration and analysis in the entire pelvis [FGM*on]. Our work is supporting clinical researchers in demonstrating the significance of dose plan adaptation to anatomical changes, and can be considered as a first positive step towards the analysis of variability in multi-organ cohorts and its inclusion in decision making.

The contribution of this paper is the documentation of the lessons learnt through the entire co-design process with the RT domain experts, which can be applicable to other visualization applications. We focus on the requirements and the design strategy for developing successful visual analytics applications for adaptive RT. For more details on the respective applications, the reader is referred to the original papers.

2. Learning the Domain and Defining the Problem

Current Workflow—In clinical practice, the evaluation of a treatment plan is currently done in two ways [SRM*19], as shown in Figure 1. First, anatomical 2D/3D views (left) depict how the dose affects the tumor and its surrounding organs for a given timepoint during the treatment period [NDSM*19]. Assessing the robustness of treatment strategies is often done in retrospective cohort studies, and anatomical views do not allow for an easy exploration of multiple patients at the same time. Second, dose volume histograms (DVHs, right) are plots to represent the amount of radiation that is administered to the volume of a specific organ [SRM*19]. DVHs scale well for a large number of patients and enable the quick identification of organs at risk of toxicity, but they do not provide a link to patient anatomy. In practice, they show how much of an organ might have been affected by harmful radiation levels, but they do not show the specific anatomical parts thereof.

Available Dataset—We had data from 29 anonymized patients, available from a cohort study conducted by UC San Diego, USA, and Aarhus University Hospital, Denmark. This dataset includes an initial 3D dose plan designed prior to treatment and 12 subsequent treatment sessions for each patient (13 fractions in total for each patient). The first five are from the five daily sessions of the first week, while the subsequent datasets were evenly sampled from the following treatment weeks [CMMH*17b]. The initial treatment plan was calculated for patients with an empty rectum and full bladder. At each session of their treatment, the patients were instructed to have roughly the same organ fillings. Before each treatment, a Cone Beam Computed Tomography (CBCT) acquisition was done for patient alignment using rigid translations. For each of these sessions, we were provided with pelvic organ segmentations in the form of contour lines, obtained by manual delineations conducted by the same radiologist. For all patients, the prostate, bladder and rectum delineations are available, as well as information about toxicity manifestation for each patient. A schematic depiction of the data is shown in the upper part of Figure 2.

Tasks Analysis—The main goal of clinical researchers working on adaptive radiotherapy treatment is to *demonstrate the significance of dose plan adaptation to changes in the anatomy of the patient throughout the course of the treatment*. To obtain more robust results, it is also necessary to account for the natural anatomical variability between individual patients. Therefore, clinical studies are conducted with regard to retrospective cohort data. This work started with a focus on a single organ—the bladder [RCMA*18]. Four clinical researchers from two different clinical institutions, working on the design of robust treatment strategies, determined the tasks for this analysis:

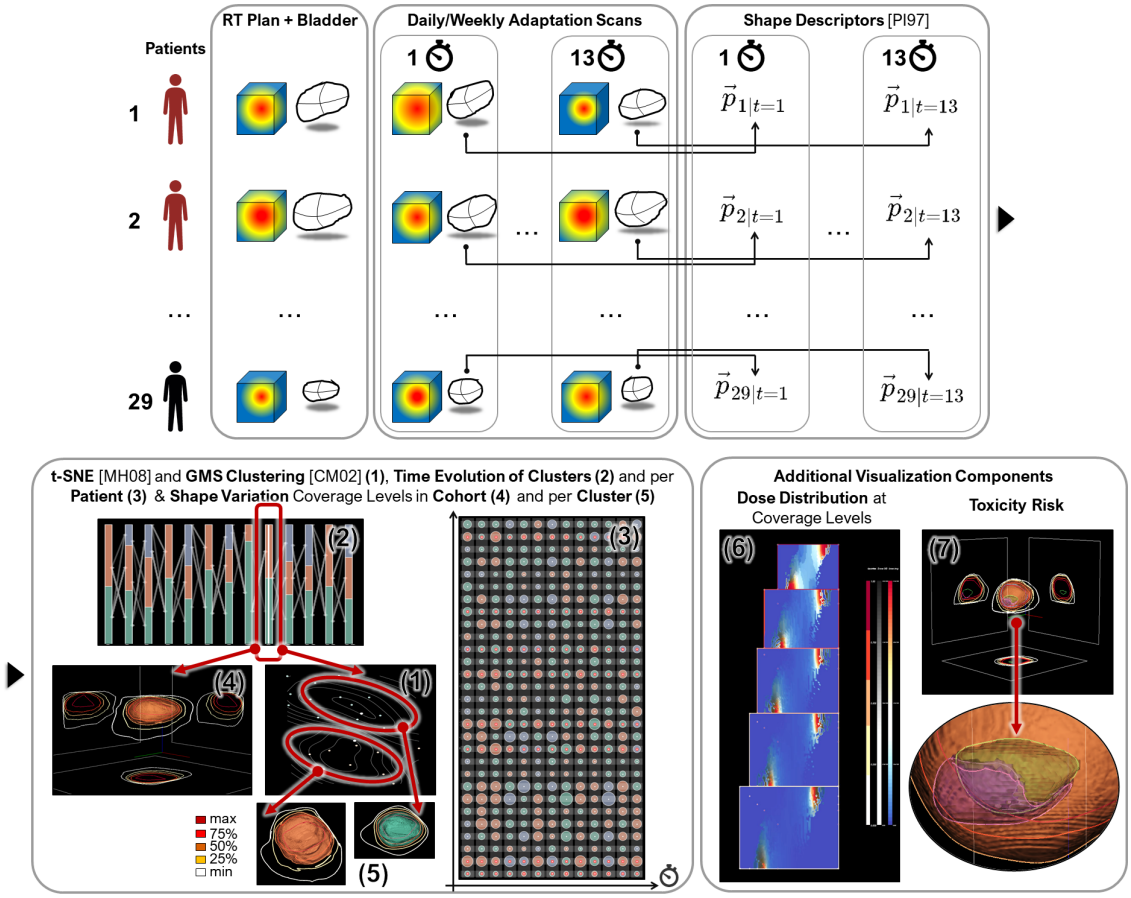

Figure 2: *The functionality of the Bladder Runner [RCMA*18], for individual patient and cohort exploration and analysis.*

(T1) Shape Variation Quantification and Exploration, i.e., how much does the organ shape vary through treatment time?
(T2) Dose Distribution Exploration and Analysis, i.e., how much does the administered RT dose vary through treatment time?
(T3) Toxicity Risk Exploration and Correlation to Anatomy, i.e., how much does the administered RT dose deviate from the initial plan through treatment time, with respect to the organ shape?

Once this functionality had been met, the clinical researchers wanted to extend the application to multiple organs [GCMM*19, FGM*on]. With this extension, their analysis would account for the variability of segmented pelvic organs in multiple patients, across the entire radiation therapy treatment process. Here, the extended tasks were:

(ET1) Global Pelvic Organ Variation Exploration, i.e., how much do the pelvic organs shapes vary through treatment time in the entire cohort of patients?
(ET2) Local Anatomical Pelvic Organ Variation Exploration, i.e., which anatomical parts of the pelvic organs vary more through treatment time?
(ET3) Dose Distribution Exploration and Analysis w.r.t. Anatomy, i.e., what is the impact of the dose distribution with respect to anatomy variability?

3. Visual Analytics in Adaptive Radiotherapy

We hereby guide the reader through a summary of our entire iterative process of developing visual analytics applications for adaptive RT. We refer the reader to the original papers of the applications, for more details on the design and implementation.

3.1. Developing the *Bladder Runner*

The *Bladder Runner* has been designed, with regard to the tasks and research questions mentioned in the Section 2. The application has been built as a stand-alone application in Python, using the Visualization Toolkit and PyQt, together with other standard libraries, such as sklearn, scipy, numpy and matplotlib. The general functionality of the application is presented in Figure 2. More details about the design choices within the *Bladder Runner* are provided in the original paper [RCMA*18].

For task **(T1)** (**Shape Variation Quantification and Exploration**), we extract a number of primitive shape descriptors [PI97], to build a high-dimensional feature vector for each bladder at each

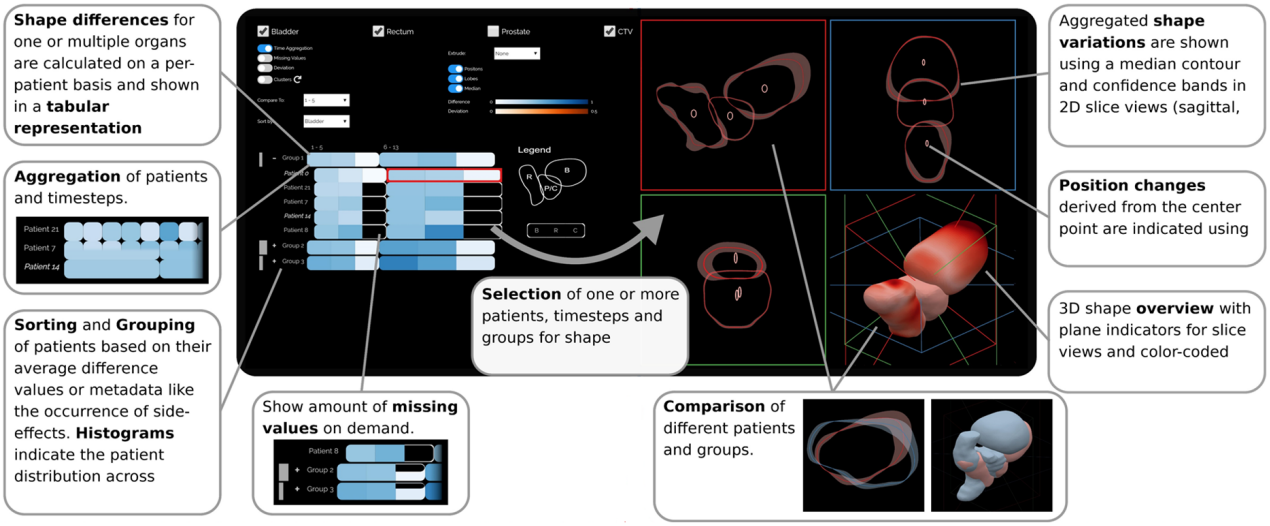

Figure 3: *The functionality of the Pelvis Runner [GCMM*19], for individual patient and cohort exploration and analysis.*

time point. Dimensionality reduction (t-SNE) [MH08] and clustering (GMS) [CM02] are employed to extract meaningful subgroups of bladders with similar shape characteristics (Figure 2, (1)). Additional representations are used to convey the time evolution of the clusters in a *stacked bar chart* configuration, and the evolution of each patient in a *contingency matrix* (Figure 2, (2) and (3), respectively). To investigate the bladder shapes and their variations in their anatomical space, we use the concept of *coverage levels*. This is similar to contour boxplots [WMK13] and denotes the likelihood that a voxel is within the bladder. This can be done for individual patients (Figure 2, (4)) and for clusters (Figure 2, (5)).

For task **(T2)** (**Dose Distribution Exploration and Analysis**), additional visualization components show the respective dose distributions through the treatment process and the calculated toxicity risk. To avoid occlusion, unfoldings of coverage levels on the plane are employed. On these, the average dose and standard deviation through the RT treatment process are color-encoded [HB03]. Then, we stack the five coverage levels on the so-called *matryoshka representation* (Figure 2, (6)). For the comparative visualization of more than one patients or clusters, a simple *juxtaposition* is employed [KCK17], given the small number of clusters.

Task **(T3)** (**Toxicity Risk Exploration and Correlation to Anatomy**) is addressed by providing an additional *rendering* of the over-radiated and under-radiated areas of the bladder. This is employed to show with silhouettes parts that have not received an adequate amount of dose (Figure 2, (7)), according to user-determined, acceptable limits of radiation dose.

Evaluation—To assess the value of the *Bladder Runner*, an initial informal evaluation was performed with three medical physicists from two clinical institutions [RCMA*18]. We conducted a real-life usage scenario, which could not be achieved with the previous means of exploration, i.e., DVHs. Clinical researchers were able to document, for the first time, the existence of a correlation between bladder shapes and risk of radiation-induced side-effects. The evaluation participants commented that "*patient information of the spatial dose delivered will allow a more accurate analysis of the administered dose*" and that "*this information will potentially widen our knowledge about patients more prone to develop toxicity*".

3.2. Clinical Case-Control Study with the *Bladder Runner*

Using the *Bladder Runner*, Casares et al. [CMRP*19] evaluated whether parameterized bladder shape descriptors can distinguish between patients exhibiting and patients not exhibiting toxicity. This study was conducted as a matched case-control in a cohort of 258 patients. A previous conventional analysis of the same data [CMMH*17a] had not shown differences between the above mentioned patient groups. Using the *Bladder Runner*, an analysis of the data was conducted, and patient bladders were clustered based on their shape descriptors. Two clusters with distinct shape characteristics included 85% of the patients, and the remaining 15% were considered outliers in their shape behavior. Then, ANOVA tests for each descriptor and each cluster were performed to find statistically significant shape descriptors. A repeated measurements model was fitted at each cluster to evaluate within-cluster trends for patients with and without toxicity. The study concluded that higher toxicity risk correlates with convex and round shapes for small bladders, and with concave and elliptical shapes for large bladders. For the first time, a clinical study confirmed—with the help of the *Bladder Runner*—that *bladder shapes can be used for toxicity prediction*.

3.3. Developing the *Pelvis Runner*

The previously discussed *Bladder Runner* demonstrated its clinical usefulness with a single focus on bladder toxicity. The *Pelvis Runner* has been designed, with regard to the tasks and research questions mentioned in the Section 2, for multiple pelvic organ exploration and analysis. To ensure scalability and interactivity within the *Pelvis Runner*, it was designed as a server-client application. A webserver in conjunction with MATLAB performs the computationally expensive processing operations, and the client browser

application is responsible for the visualizations, using three.js and D3.js. The functionality of the application is shown in Figure 3. More details about the design choices within the *Pelvis Runner* are provided in the original paper [GCMM*19].

Our approach starts with *data processing* to transform the registered data from given organ contours to volumetric data. Then, we use *linearization* [WFG*19] to create a data representation that facilitates statistical analysis and comparison within the cohort without losing patient and time correspondences. This is used to create a *low dimensional embedding* of the entire cohort, using Principal Component Analysis (PCA) for within organ class separation and t-SNE [MH08] for across organ class separation.

For task **(ET1)** (**Global Pelvic Organ Variation Exploration**), an *abstracted cohort representation* is created based on the outcome of the low dimensional embedding of the previous step (Figure 3, left). We use a flexible and interactive tabular representation to support the representation of multiple organs at the same time, hierarchical clustering for cohort partitioning, and additional trustworthiness information. Standard *F+C*, sorting and filtering, visual aggregations, e.g., based on timepoints, and additional partitioning based on patient metadata, e.g., toxicity data, are integrated.

For task **(ET2)** (**Local Anatomical Pelvic Organ Variation Exploration**), we need to convey *anatomical information* on demand (Figure 3, right). For this, the three standard anatomical 2D planes (sagittal, coronal and axial) and a 3D view are employed. By sampling the embedding space for the median and the standard deviations, we reconstruct the shape variations and we show them with a representation similar to contour boxplots [WMK13]. In the 2D and 3D views, superposition is employed to show multiple subjects at the same time, and additional glyphs are employed to show positional variance. Standard interaction, e.g., zooming, panning, slicing and linking, is integrated.

Evaluation—Four usage scenarios were conducted by two medical physicists [GCMM*19]. The domain experts commented that the application provides a flexible and systematic way to explore the data. Although this was not intended functionality, they commented that "*the tool offers a way of identifying the setup uncertainty of the entire treatment*", as it allows an overview of the motion, i.e., uncertainty, of the prostate. They expect that the application could give "*indications of patients that will fail or that may develop toxicity at the beginning of the treatment*", allowing them to adapt the employed strategy.

3.4. Clinical Study with the *Pelvis Runner*

Using the *Pelvis Runner*, Furmanova et al. [FRG*20] analyzed organ shape variations with respect to the mean shape, as computed from the first *N* treatment points. The study analyzed how the predictibility of organ variability changed with an increasing number of treatment points, used to generate the mean shape. When the variability of the entire treatment time was compared to the first time point and to the mean shape of the first five timepoints, with the latter, the variability decreased on average by 28% per patient for the bladder, 20% for the rectum, and 27% for the prostate. For most patients, most of the organ motion could be captured and predicted within the first four timepoints of treatment. Thus, using multiple pre-treatment scans captured over the first four consecutive days could be enough to identify patients with higher organ shape variability.

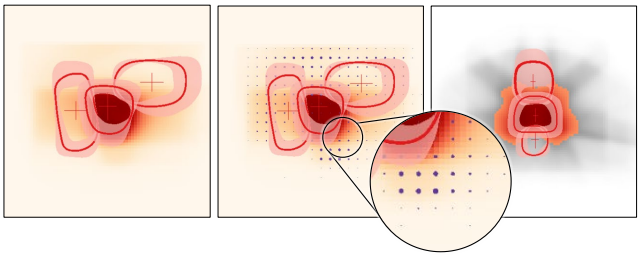

Figure 4: *Anatomical views for assessing the impact of organ variability to dose distribution in VAPOR [FGM*on].*

3.5. Developing *VAPOR*

The *Pelvis Runner* did not include functionality for assessing the impact of organ variability to dose distribution and toxicity risk **(ET3)**. The domain experts anticipated that incorporating this functionality into the tool could create a "*a promising and useful decision-making tool for radiation oncologists*" [GCMM*19]. To this end, we extended the *Pelvis Runner* into *VAPOR* [FGM*on].

The concept behind the development of *VAPOR* is that not all regions of the pelvic organs are equally important. The most critical regions are those where anatomical variability is high and the radiation dose is also high. To set this constraint, the domain experts can guide the global anatomical variability exploration and analysis by restricting the RT dose, with the use of a *user-selected threshold* (Figure 4, right). As the application supports the incorporation of retrospective toxicity information, it is possible to relate toxicity with anatomical variability and the locations of high dose.

For a more localized view on regions of interest, we compute the distribution of the administered RT dose, i.e., the average and the standard deviation. We show the average dose as a background sequential white-to-red (low-to-high dose) color map [HB03] in the 2D anatomical planes, as shown in Figure 4. The standard deviation is mapped on the area of superimposed circular glyphs [BKC*13]. To preserve anatomical context, *F+C* is employed [BCS96].

Evaluation—The functionality of *VAPOR* was showcased with four usage scenarios conducted by two domain experts. A systematic evaluation with the intended users and a clinical study with a larger cohort would be required in the future.

4. Lessons Learnt

Several lessons were learnt throughout the process of developing visual analytics applications for adaptive RT, step-by-step. Although the presented applications have very specific goals and objectives, the findings and knowledge obtained during the design process—which is actually a co-design process together with the domain experts—can be useful and applicable to other domains.

Building applications, which are "here-to-stay", requires following a long *iterative design process*, as summarized in Figure 5. This requires several steps:

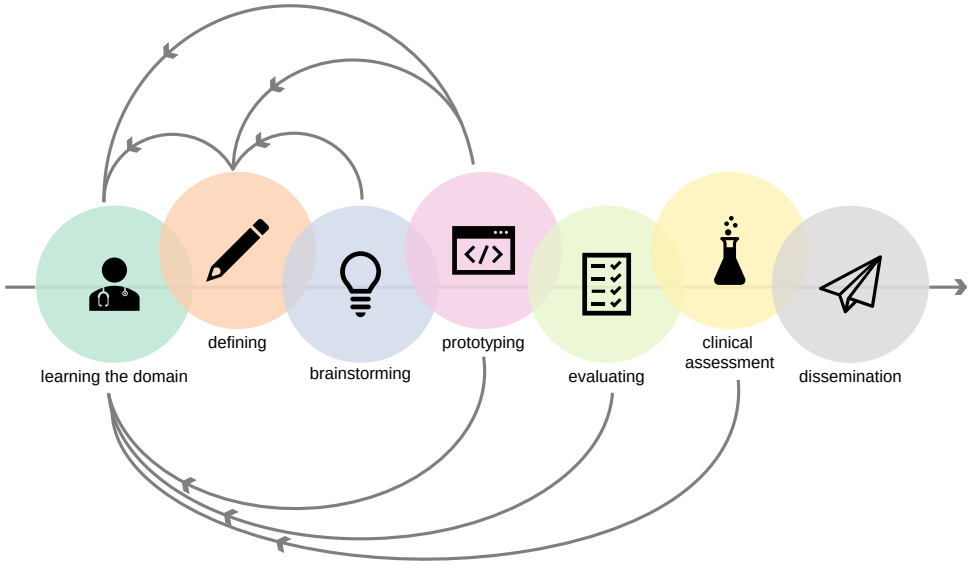

Figure 5: *The iterative process used during the design of our visual analytics applications for adaptive RT.*

- learning about the domain,
- defining together with (a small subset of) experts the problem,
- brainstorming together with (a small subset of) experts about possible solutions,
- providing an initial design (prototype),
- potentially discussing and refining,
- evaluating the initial outcome (with as many experts as possible),
- assessing within the domain experts' setting (clinical study),
- re-iterating to incorporate a problem of higher complexity (more patients, and/or more organs),
- and disseminating the result of one's work.

This is exactly the process that we followed in our work. This collaboration had already started with other projects, many years before the *Bladder Runner*. By the time we started working on the *Bladder Runner* project, the visualization researchers of this project already had a good understanding of the field of RT and the domain experts had a good overview of the limitless possibilities that visualization offers. The whole project started slowly by investigating solutions for the tasks of the *Bladder Runner*, brainstorming together with the domain experts. After several close interactions and iterations with one domain expert, who was actively involved in the design phase, the first prototype was ready for evaluation by more domain experts. The *inclusion of at least one domain expert in the design phase* is crucial for success, in our opinion. The evaluation was followed by a clinical study, from which several interesting points have been raised—resulting in the development of the *Pelvis Runner* and, later, *VAPOR*. For these two applications, we followed the same iterative design process.

For the visualization researchers involved in this project, obtaining a very *thorough knowledge of the data and of the already existing workflow* that is employed in the clinical research setting was fundamental. In our case, the effort revolved initially around reading textbooks from the field of RT, then spending significant time within the environment of the domain experts to observe their workflow, and mainly targeting contributions both in the visualization and the RT research field. By putting additional effort to obtain deeper knowledge about the application domain and to build a more meaningful and trustful relationship with the collaborating experts, we increase our chances to develop actually useful and adoptable solutions. We experienced that it was particularly beneficial to *spend significant time in the laboratories of the domain experts*, where we could experience how they work and we could clarify all necessary aspects, before starting the development.

For our applications, it has been crucial to conduct a *thorough tasks and requirements analysis* together with the collaborating domain experts, such as those presented in Section 2, or in the respective sections of the initial papers of the presented applications. This has been the most time-consuming step of the process, in all three applications. The tasks analysis might be refined during the iterative process of developing increasingly sophisticated applications, but it should be triggered by the experts to include all "must-have" features in the final product.

Although exploring and experimenting with different kind of solutions is an enjoyable and creative process, developing applications that have higher chances of being adopted in clinical practice requires more rigorous and systematic approaches. In our case, it was fruitful to have a continuous and open dialogue between visualization and domain experts about ideas and possibilities to include in the prototypes. Such open processes, though, require building a *high level of trust among partners*, which requires time and long-lasting, close collaborations—for example, in our case, the collaboration started already in 2011. Visualization researches were able to ask freely "*would X be a useful and desired feature?*" and domain experts were able to ask "*how could Y be included?*", generating new interesting concepts and novel approaches. Yet, not all investigated strategies made it into the final applications, while others were drastically simplified. The dose variability is a good example of a simplification. In the *Bladder Runner*, it included the step of

unfolding the bladders on the plane, which would be unfeasible to adapt and extend for the entire pelvis in *VAPOR*.

During the prototyping phase, we also realised that *creating and maintaining reusable visualization* software is not easy, while adopting visualization frameworks developed in other research groups is not always possible—even in the form of a plug-in. Although a more efficient use of resources would be required (also to achieve higher impact), prototyping for showcasing our ideas is also a common "way-to-go". Making our software open-source (unless restricted, for some reason) is also crucial, as it has a higher chance of practical uptake. In our case, the main findings and knowledge obtained during the software development was to account early enough for extensible and scalable solutions. This was not done in the *Bladder Runner*, and resulted in a complete change of environment for the *Pelvis Runner* to account for the higher complexity and dimensionality of the data. *VAPOR*, though, was built upon the *Pelvis Runner*.

In this project, we did not look into certification options, as the retrospective clinical studies could be performed with the existing prototypes. Yet, clinical applications are required to undergo *strong certification processes*, as they affect patient treatment and decision making. There are several benchmark standards for the development of medical software worldwide (e.g., IEC or ISO). A good practice towards certification would include standardization and benchmarking, as well as the inclusion of ethical restrictions and privacy requirements, such as data anonymization and GDPR.

An interesting aspect of this project was having *strong support from an industrial partner* that has interest in the developed applications, and/or is willing to fund the research, and/or is willing to integrate the developed solutions into robust frameworks for clinical practice. Teaming up with a company to develop research results into a product, significantly increases the chances of uptake for the prototypes, and of using them as a basis to build robust, solid and usable applications for clinical practice and research.

Evaluation is not always an easy task, and we relied on common best practices from the field of visualization [PRI18]. However, we found out that a particularly interesting and insightful practice of evaluating additionally our work is through an assessment within the domain experts' setting. Both for the *Bladder Runner* and the *Pelvis Runner*, clinical studies were performed employing our visualization tools. This was an intensive and detailed test to determine which features in our tools are helpful for the RT experts and to derive information about their usability. Each time, this step provided us with valuable insight for the following extension.

Last but not least, we highly recommend to think in advance about *dissemination methods of one's work*. Dissemination is a cornerstone of research, enabling us to transfer knowledge to fellow researchers, collaborators, our students, and the general public. In this case, dissemination has been necessary to convey the content of our work at all stages of development: (1) to our collaborators in a manner that is helpful for them to conduct further studies with the use of our applications and (2) to the general clinical research community, through the clinical studies of our collaborators. Dissemination should be seen as an *important step for further brainstorming* with our existing collaborators, and as a way of generating new ideas for future directions, also with other clinical partners.

5. Conclusions

This work is an example of a fruitful collaboration between the visualization and the medical domain. Specific visual analytics applications have been designed based on the actual needs of domain experts, evaluated and, finally, engaged in edge-cutting clinical research. This has also been possible with adequate support from industrial partners, who have shown interest in funding, refining and incorporating our prototypes into their systems. The visualizations designed for this specific application could be possibly adapted to other domains of model-based RT and extended to other RT applications, e.g., for training new personnel. Another interesting future direction is to use the deployed applications as the starting point for building and analyzing risk prediction models for incoming patients. The developed visual analytics applications offer new perspectives to clinical researchers working in the field of RT-induced toxicity, for designing better targeted treatment for the patients by the incorporation of anatomical variability into adaptive RT.

Acknowledgement

This work was supported by Varian Medical Systems of Palo Alto, California, USA in the frame of a research project entitled "A machine learning centered visualization system for model-based decision making in image-guided and adaptive radiotherapy of cancer" (principal investigator L. P. Muren).

References

[BCS96] BUJA A., COOK D., SWAYNE D. F.: Interactive high-dimensional data visualization. *Journal of computational and graphical statistics 5*, 1 (1996), 78–99. 5

[BKC*13] BORGO R., KEHRER J., CHUNG D. H., MAGUIRE E., LARAMEE R. S., HAUSER H., WARD M., CHEN M.: Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics (STARs)* (2013), pp. 39–63. 5

[CM02] COMANICIU D., MEER P.: Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 5 (2002), 603–619. 4

[CMB*18] CASARES-MAGAZ O., BÜLOW S., , PETTERSSON N., MOISEENKO V., M. T., EINCK J., HOPPER A., KNOPP R., MUREN L.: OC-0183: A case-control study of the relations between planned vs actually delivered rectal dose surface maps. *Radiotherapy and Oncology 127* (2018), S97–S98. 1

[CMMH*17a] CASARES-MAGAZ O., MOISEENKO V., HOPPER A., PETTERSSON N., THOR M., KNOPP R., DEASY J., MUREN L., EINCK J.: OC-0489: Variation in bladder volume and associated spatial dose metrics in prostate and pelvic radiotherapy. *Radiotherapy and Oncology 123* (05 2017), S260. 1, 4

[CMMH*17b] CASARES-MAGAZ O., MOISEENKO V., HOPPER A., PETTERSSON N. J., THOR M., KNOPP R., DEASY J. O., MUREN L. P., EINCK J.: Associations between volume changes and spatial dose metrics for the urinary bladder during local versus pelvic irradiation for prostate cancer. *Acta Oncologica 56*, 6 (2017), 884–890. 2

[CMRP*19] CASARES-MAGAZ O., RAIDOU R., PETTERSSON N., MOISEENKO V., EINCK J., HOPPER A., KNOPP R., MUREN L.: PO-0962: Bladder changes during first week of RT for prostate cancer determine the risk of urinary toxicity. *European SocieTy for Radiation & Oncology (ESTRO) 38* (2019). 2, 4

[CvHvdK*11] CHAI X., VAN HERK M., VAN DE KAMER J. B., HULSHOF M. C., REMEIJER P., LOTZ H. T., BEL A.: Finite element based bladder modeling for image-guided radiotherapy of bladder cancer. *Medical physics 38*, 1 (2011), 142–150. 1

[DJFB05] DELANEY G., JACOB S., FEATHERSTONE C., BARTON M.: The role of radiotherapy in cancer treatment. *Cancer 104*, 6 (2005), 1129–1137. 1

[FC16] FONSECA T. C. F., CAMPOS T. P. R.: SOFT-RT: Software for IMRT simulations based on MCNPx code. *Applied Radiation and Isotopes 117* (2016), 111–117. 2

[FGM*on] FURMANOVA K., GROSSMANN N., MUREN L., CASARES-MAGAZ O., MOISEENKO V., EINCK J., GRÖLLER M., RAIDOU R.: VAPOR: Visual Analytics for the Exploration of Pelvic Organ Variability in Radiotherapy—*Under submission*. 2, 3, 5

[FRG*20] FURMANOVA K., RAIDOU R., GROSSMANN N., CASARES-MAGAZ O., MOISEENKO V., EINCK J., MUREN L.: Using multiple planning scans to predict organ shape variability during RT for prostate cancer. *To Appear in European SocieTy for Radiation & Oncology (ESTRO) 39* (2020). 2, 5

[GCC*16] GEAR J. I., CUMMINGS C., CRAIG A. J., DIVOLI A., LONG C. D. C., TAPNER M., FLUX G. D.: Abdo-Man: A 3D-printed anthropomorphic phantom for validating quantitative SIRT. *EJNMMI Physics 3*, 1 (2016), 17. 2

[GCMM*19] GROSSMANN N., CASARES-MAGAZ O., MOISEENKO V., MUREN L. P., GRÖLLER M. E., RAIDOU R. G.: Pelvis Runner: Visualizing Pelvic Organ Variability in a Cohort of Radiotherapy Patients. *EG Workshop on Visual Computing for Biology and Medicine (EG VCBM)* (2019). 2, 3, 4, 5

[GDF*00] GAMBARINI G., DANESI U., FORONI R., MAURI M., PIROLA L., BIRATTARI C.: Prompt imaging of absorbed dose in tissue-equivalent gel-phantoms and new toolkit for 3D data visualization. In *IEEE Nuclear Science Symposium (NSS/MIC)* (2000), vol. 3, pp. 19/52–19/55 vol.3. 2

[HB03] HARROWER M., BREWER C. A.: Colorbrewer. org: an online tool for selecting colour schemes for maps. *The Cartographic Journal 40*, 1 (2003), 27–37. 4, 5

[HST87] HAHN P., SHALEV S., THERRIEN P.: Colour visualization as an aid to the comparison of treatment plans for prostatic carcinoma. *Acta Oncologica 26*, 4 (1987), 313–315. 1

[IFP95] INTERRANTE V., FUCHS H., PIZER S.: Enhancing transparent skin surfaces with ridge and valley lines. In *Proceedings Visualization '95* (1995), pp. 52–59. 1

[KCK17] KIM K., CARLIS J. V., KEEFE D. F.: Comparison techniques utilized in spatial 3d and 4d data visualizations: A survey and future directions. *Computers & Graphics 67* (2017), 138–147. 4

[LFP*90] LEVOY M., FUCHS H., PIZER S. M., ROSENMAN J., CHANEY E. L., SHEROUSE G. W., INTERRANTE V., KIEL J.: Volume rendering in radiation treatment planning. In *Proceedings of the Conference on Visualization in Biomedical Computing* (1990), pp. 4–10. 1

[MH08] MAATEN L. V. D., HINTON G.: Visualizing data using t-SNE. *Journal of Machine Learning Research 9*, Nov (2008), 2579–2605. 4, 5

[MLK*07] MOISEENKO V., LIU M., KRISTENSEN S., GELOWITZ G., BERTHELET E.: Effect of bladder filling on doses to prostate and organs at risk: a treatment planning study. *Journal of Applied Clinical Medical Physics 8*, 1 (2007), 55–68. 1

[MSD03] MUREN L. P., SMAALAND R., DAHL O.: Organ motion, set-up variation and treatment margins in radical radiotherapy of urinary bladder cancer. *Radiotherapy and Oncology 69*, 3 (2003), 291–304. 1

[MYKO91] MIYAZAWA T., YOSHIDA R., KIMURA M., OTSUKI T.: Visualization of 3D medical images for radiotherapy planning. In *IEEE Nuclear Science Symposium (NSS/MIC)* (1991), pp. 1553–1557 vol.3. 1

[NDSM*19] NEJAD-DAVARANI S. P., SEVAK P., MONCION M., GARBARINO K., WEISS S., KIM J., SCHULTZ L., ELSHAIKH M. A., RENISCH S., GLIDE-HURST C.: Geometric and dosimetric impact of anatomical changes for mr-only radiation therapy for the prostate. *Journal of applied clinical medical physics* (2019). 2

[PI97] PEURA M., IIVARINEN J.: Efficiency of simple shape descriptors. *Aspects of visual form* (1997), 443–451. 3

[PRI18] PREIM B., ROPINSKI T., ISENBERG P.: A Critical Analysis of the Evaluation Practice in Medical Visualization. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2018), The Eurographics Association, pp. 45–56. 7

[RCMA*18] RAIDOU R., CASARES-MAGAZ O., AMIRKHANOV A., MOISEENKO V., MUREN L., EINCK J., VILANOVA A., GRÖLLER M.: Bladder Runner: Visual Analytics for the Exploration of RT-Induced Bladder Toxicity in a Cohort Study. *Computer Graphics Forum 37*, 3 (2018), 205–216. 2, 3, 4

[RDCO*17] RIOS R., DE CREVOISIER R., OSPINA J. D., COMMANDEUR F., LAFOND C., SIMON A., HAIGRON P., ESPINOSA J., ACOSTA O.: Population model of bladder motion and deformation based on dominant eigenmodes and mixed-effects models in prostate cancer radiotherapy. *Medical image analysis 38* (2017), 133–149. 1

[SRM*19] SCHLACHTER M., RAIDOU R., MUREN L., PREIM B., PUTORA P., BÜHLER K.: State-of-the-art report: Visual computing in radiation therapy planning. In *Computer Graphics Forum* (2019), vol. 38, pp. 753–779. 1, 2

[THLM*13] THARIAT J., HANNOUN-LEVI J.-M., MYINT A. S., VUONG T., GÉRARD J.-P.: Past, present, and future of radiotherapy for the benefit of patients. *Nature reviews Clinical oncology 10*, 1 (2013), 52. 1

[VYM*10] VISWANATHAN A. N., YORKE E. D., MARKS L. B., EIFEL P. J., SHIPLEY W. U.: Radiation dose–volume effects of the urinary bladder. *International Journal of Radiation Oncology* Biology* Physics 76*, 3 (2010), S116–S122. 1

[WFG*19] WEISSENBÖCK J., FRÖHLER B., GRÖLLER E., KASTNER J., HEINZL C.: Dynamic volume lines: Visual comparison of 3d volumes through space-filling curves. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (2019), 1040–1049. 5

[WL15] WASHINGTON C. M., LEAVER D. T.: *Principles and practice of radiation therapy*. Elsevier Health Sciences, 2015. 1

[WMK13] WHITAKER R. T., MIRZARGAR M., KIRBY R. M.: Contour boxplots: A method for characterizing uncertainty in feature sets from simulation ensembles. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2713–2722. 4, 5

[WRH*08] WRIGHT P., REDPATH A. T., HØYER M., GRAU C., MUREN L. P.: The normal tissue sparing potential of adaptive strategies in radiotherapy of bladder cancer. *Acta Oncologica 47*, 7 (2008), 1382–1389. 2