

Visual Analytics for the Exploration of Tumor Tissue Characterization

R.G. Raidou¹, U.A. van der Heide², C.V. Dinh², G. Ghobadi², J.F. Kallehauge³, M. Breeuwer^{4,1}, A. Vilanova⁵

¹Eindhoven University of Technology, The Netherlands

²Department of Radiotherapy, the Netherlands Cancer Institute, Antoni van Leeuwenhoek Hospital, The Netherlands

³Department of Medical Physics, Aarhus University Hospital, Denmark

⁴Philips Healthcare Best, The Netherlands

⁵Delft University of Technology, The Netherlands

Abstract

Tumors are heterogeneous tissues consisting of multiple regions with distinct characteristics. Characterization of these intra-tumor regions can improve patient diagnosis and enable a better targeted treatment. Ideally, tissue characterization could be performed non-invasively, using medical imaging data, to derive per voxel a number of features, indicative of tissue properties. However, the high dimensionality and complexity of this imaging-derived feature space is prohibiting for easy exploration and analysis - especially when clinical researchers require to associate observations from the feature space to other reference data, e.g., features derived from histopathological data. Currently, the exploratory approach used in clinical research consists of juxtaposing these data, visually comparing them and mentally reconstructing their relationships. This is a time consuming and tedious process, from which it is difficult to obtain the required insight. We propose a visual tool for: (1) easy exploration and visual analysis of the feature space of imaging-derived tissue characteristics and (2) knowledge discovery and hypothesis generation and confirmation, with respect to reference data used in clinical research. We employ, as central view, a 2D embedding of the imaging-derived features. Multiple linked interactive views provide functionality for the exploration and analysis of the local structure of the feature space, enabling linking to patient anatomy and clinical reference data. We performed an initial evaluation with ten clinical researchers. All participants agreed that, unlike current practice, the proposed visual tool enables them to identify, explore and analyze heterogeneous intra-tumor regions and particularly, to generate and confirm hypotheses, with respect to clinical reference data.

Categories and Subject Descriptors (according to ACM CCS): I.3.8 [Computer Graphics]: Applications—Applications; J.3 [Computer Applications]: Life and Medical Sciences—Life and Medical Sciences

1. Introduction

Radiotherapy aims at irradiating tumors with a high dose, while sparing surrounding healthy tissues. However, tumors are heterogeneous tissues, enclosing multiple regions with distinct characteristics. Lately, it has been hypothesized that incorporating tissue characteristics information into radiotherapy planning can play an important role in tumor diagnosis and in designing even more effective treatment strategies, i.e., better targeted planning where distinct intra-tumor tissues are irradiated with tailored radiation doses.

Currently, intra-tumor tissue heterogeneity is investigated by studying histopathological data acquired from biopsies.

To substitute histopathology, clinical researchers need to associate histopathological findings, such as aggressiveness or resistance of a part of the tumor, to *features derived from imaging data*. For example, Dynamic Contrast Enhanced (DCE) and Diffusion Weighted (DW) Magnetic Resonance Imaging (MRI) are employed in tumor analysis to derive, using mathematical models, per-voxel features indicative of tissue characteristics. In addition to histopathology, clinical researchers often employ in their analysis *supplementary clinical data* that they use as reference, such as maps that predict high-risk tumor zones or maps depicting tumor control probability. These data also need to be explored, in association with the imaging-derived features. Moreover,

the derivation of tumor tissue characteristics from imaging data is based on -more or less- complex mathematical models [SB13]. Different modeling approaches make different assumptions, resulting in features with different distributions. A priori, it is not known whether different alternatives in modeling present significant differences or which option leads to better results. Additionally, all modeling approaches introduce model-fitting inaccuracy in the derived features. Exploring *variability* and incorporating *inaccuracy* induced by different modeling approaches, with respect to anatomical or clinical reference data, can have significant impact on the final clinical decision and treatment.

However, the high dimensionality and complexity of the imaging-derived feature space, along with model-induced variability and inaccuracy, is prohibiting for easy exploration. The structure of this feature space is also not fully comprehended and, in the state-of-the-art clinical research workflow, there is no easy and insightful way to link observations from histopathology or clinical reference to features derived from imaging.

In this paper, we introduce a visual analysis tool for tumor tissue characterization. We mainly focus on knowledge discovery and hypothesis generation and confirmation, by linking imaging-derived tissue characteristics to anatomical and clinical reference data. In our approach, we employ as central view, a 2D embedding to map the feature space of imaging-derived tissue characteristics into a 2D information space. Multiple linked interactive views provide functionality for the exploration and analysis of the local structure of the feature space, enabling the user to retrieve information on distinct intra-tumor regions and to associate observations between the feature space and clinical reference data.

The *contribution* of this paper is the design and implementation of a visual tool for the exploration of tumor tissue characterization. To the best of our knowledge, there is no other tool to serve this purpose. The proposed visual tool incorporates the following components: (C1) It supports the identification and exploration of intra-tumor regions with distinct imaging-derived tissue characteristics. (C2) It facilitates the exploration and analysis of the feature space structure, in relation to patient anatomy. (C3) It enables the association of observations from clinical reference information to the feature space, and vice versa. (C4) It allows the exploration and analysis of model-induced variability and inaccuracy of the feature space.

2. Clinical Background

Several imaging modalities are used in tumor diagnosis and analysis. In this paper, we will focus on prostate and cervical tumor, which are investigated using MRI data, such as DCE and DW images, but our approach can be easily generalized to include other modalities, like Computed Tomography (CT) or Positron Emission Tomography (PET). DCE-MRI data are 4D data, i.e., 3D volume+time data, which depict

the absorption and washout of a contrast agent (CA) in tissue over time. This technique is based on the idea that tumorous and healthy tissues have different CA uptake properties. Tumors tend to develop new, disorganized and permeable vessels, which have thinner and weaker walls [TTP*10]. Thus, they absorb and wash out CAs faster than healthy tissue. A quantitative way of measuring tissue properties from DCE-MRI data is to use one of the established pharmacokinetic (PK) models [SB13]. These models are employed to derive per voxel an output set of features, called PK parameters, which describe the distribution of the CA inside the tissue and are indicative of tissue characteristics [SB13]. A priori, it is not known which is the most suitable modeling approach and whether there are differences between them. The inaccuracy of each approach, i.e., the discrepancy between derived and expected feature values, also needs to be incorporated in the analysis. DW is another MRI method that produces in vivo images of biological tissues based on the random Brownian motion of water molecules within a voxel [SFHB08]. It is a powerful technique, used to identify high cellular tissue like tumors, where the diffusion of water is restricted primarily by cell membrane boundaries. This restriction can be quantitatively assessed using the apparent diffusion coefficient (ADC), as a measure of the magnitude of diffusion.

Additionally, clinical researchers require to associate findings in the imaging-derived feature space to observations from reference data and vice versa. This data might be anatomical images, such as T2-weighted MRI scans; or derived clinical data, e.g., from linear regression models predicting high-risk tumor zones [GBM*12]; or histopathological data that are obtained invasively from patient biopsies and can reveal information on tumor aggressiveness or resistance [EAJA*05]. For example, in histopathological images of prostate tumors, Gleason Scores (GS) are assigned to delineated tumor tissue foci based on their microscopic appearance. High GS tumors are more aggressive and have a worse prognosis [EAJA*05]. To substitute invasive histopathology, information such as the GS need to be associated with non-invasive, imaging-derived features.

In current clinical research, the exploration and analysis of the feature space of imaging-derived tissue characteristics is conducted using a simplistic slice-based technique. The values of each imaging-derived feature are encoded with a colormap [KHL*14]. Then, these so-called parameter maps are juxtaposed, manually inspected slice-by-slice and mentally compared and analyzed. The cross-sectional analysis of the imaging-derived features with clinical reference data is also done in a similar way. This approach has high memory demands, as it requires to mentally reconstruct relationships and patterns in the data and it is sub-optimal in insight.

Clinical researchers require an exploratory tool that allows them to perform the following *tasks*: (T1) Identify and explore intra-tumor regions with distinct imaging-derived tumor tissue characteristics. (T2) Analyze and understand

the structure of each distinct intra-tumor region and, then, compare these intra-tumor regions. (T3) Discover relations between the feature space of imaging-derived tumor tissue characteristics and clinical reference data. Also, generate and confirm hypotheses that can associate these two. (T4) Identify, explore and analyze the impact of variability or inaccuracy, induced during the derivation of tumor tissue characteristics from imaging data.

3. Related Work

Several systems have been proposed for the exploration and visual analysis of feature spaces derived from medical imaging. The most similar to our field of application were proposed by Preim et al. [POM*09], Glasser et al. [GTP10], Fang et al. [FMHC07] and Raidou et al. [RvdHvH*14]. However, these papers are mainly focusing on the exploration and analysis of perfusion parameters or on the distinction of tumors from healthy tissues. None of these applications addresses intra-tumor tissue characterization based on imaging-derived features, or cross-analysis with clinical reference and histopathological data.

Our approach is centered around a dimensionality reduced view, i.e., a 2D embedding, that preserves local structure in the feature space and that allows visual analysis of clusters and their intrinsic feature characteristics. For this reason, we reviewed also the literature related to these two topics.

Dimensionality reduction - Traditionally employed approaches for 2D projection of feature spaces include the Grand Tour [Asi85], XmdvTool [War94], WEAVE [GRW*00] and SimVis [ODH*07, Dol07]. These frameworks visualize high-dimensional datasets through projections, combined with a number of linked techniques for the visualization of the underlying feature space. However, they all support 2D projections of the multidimensional data by the selection of two variables from the feature space, which does not ensure preservation of the local structure. None of these systems supports linked brushing, or selection from a spatial view - the anatomical or clinical reference space, in our case - to an abstract view. Cluster analysis visualization is also not possible.

Frameworks to overcome some of these limitations were proposed by Blaas et al. [BBP07], Steenwijk et al. [SMB*10], Jeong et al. [JZF*09], Choo et al. [CLKP10], Ingram et al. [IMI*10] and Poco et al. [PEPM12]. They all integrate projection techniques with multiple information visualization views and bi-directionally linked scientific visualization views, to enhance data exploration. Nevertheless, the previously mentioned work is not fully applicable to our case. Even if all applications enable to identify clusters and to explore the respective feature space of the data, none of them supports incorporation of variability and inaccuracy in their analysis. Especially, they do not enable to cross-analyze the feature space with clinical reference data

and histopathology, for knowledge discovery and hypothesis generation and confirmation.

Visual analysis of clusters - We are mainly interested in visualizing clusters that are present in the multidimensional feature space of the data, as identified by the user, but also in visualizing the feature characteristics of each cluster and their relationship with each other. There are various methods for visual analysis of clusters, of which the most relevant to our work are H-BLOB [SBG00], Narcissus [HDWB95], Small-Worlds [vHvW04], but also the frameworks proposed by Seo and Shneiderman [SS02], Linsen et al. [LVLRO9], Glasser et al. [GLP14] and Turkay et al. [TPRH11]. However, in some of these frameworks, the visual complexity of the visualizations makes them unsuitable for cluster visualization of large data sets. Also, in all previously mentioned work no information is provided about the structure of a cluster and other metrics, such as the separation of two clusters.

To sum up, there are several different systems that are able to manage feature spaces similar to ours; others that employ projections and multiple views in their approaches; and others that facilitate visual analysis of clusters. However, to the best of our knowledge, there is no visual tool incorporating all the required functionality for our field of application.

4. Visual Analytics for Tumor Tissue Characterization

After data acquisition, tumor tissue characteristics are derived from medical images. However, the dimensionality and complexity of this imaging-derived feature space is prohibitive for visual exploration and data-driven analysis. The proposed visual tool aims at satisfying the specific exploratory needs of clinical researchers, working on tumor tissue characterization. Therefore, our design choices are oriented to fulfill their requirements, as described in Section 2. Our visual tool consists of three mutually linked components (Figure 1). A 2D embedding of the feature space forms the central view in our visual tool. This view is complemented by an anatomical view of a clinical reference space and multiple linked interactive views for the visual analysis of clusters.

The first requirement is the *identification and exploration of distinct intra-tumor regions with similar tissue characteristics (T1)*. An option could be to employ a scatterplot matrix to visualize all imaging-derived features. This would allow only pairwise comparison of the features. However, the relations between features in the different intra-tumor regions are expected to be more complex. To incorporate information from the high dimensional space in one simple view that enables visual exploration of all features, dimensionality reduction is required to map the high dimensional imaging-derived feature space (N-D) to a lower dimensional space (2D). This is performed as part of a pre-processing step and, therefore, different dimensionality reduction techniques could be used. In our approach, we choose to employ t-Distributed Stochastic Neighbor Embedding (t-SNE) [VdMH08]. t-SNE has the ability to map data

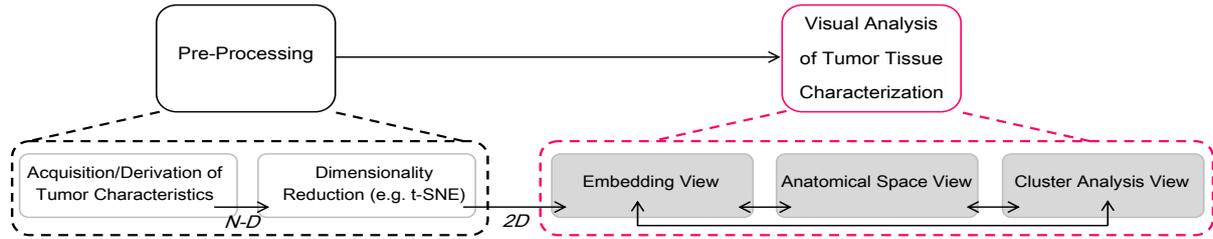


Figure 1: General profile of the proposed visual tool. The core of the tool (pink) consists of three linked components (grey).

from a high dimensional feature space into a 2D embedding space, preserving the local structure of the initial N-D space. In this way, the *embedding* can be represented in a simple 2D scatterplot. The dimensions of the resulting embedding do not have a direct association to the features: each high dimensional data point is embedded in an abstract 2D space, in such way that the resulting 2D data points plotted nearby represent N-D data points with similar values in the high dimensional feature space. In our case, 2D data points nearby in the embedding represent volumetric positions with similar imaging-derived tissue characteristics, while voxels with dissimilar tissue characteristics are plotted apart. We remark that the voxel location is not used in the embedding - only the imaging-derived features.

To provide context from the embedding map to the anatomical space, we apply a *2D colormap*, from a vast collection of colormaps [Bre94], to the points of the embedding, based on the position that they have in the 2D scatterplot. Then, this 2D colormap is propagated in the form of an overlay on the anatomical images (Figure 2). The 2D embedding map also supports interactive selection of regions of points, i.e., visual clusters (Figure 2). However, in some cases, visual clusters in the 2D embedding are not well-defined, or overlapping points may give a misleading representation of the space. For these reasons, we introduce a dual-view in the scatterplot, to depict also the point *density* (Figure 2). Hereby, the exploration can be guided by the density peaks, which might reflect the location of densely populated visual clusters. To intuitively depict high density regions versus low density, we employ the heated body colormap [Bre94]. When one visual cluster is selected, a color is assigned to it

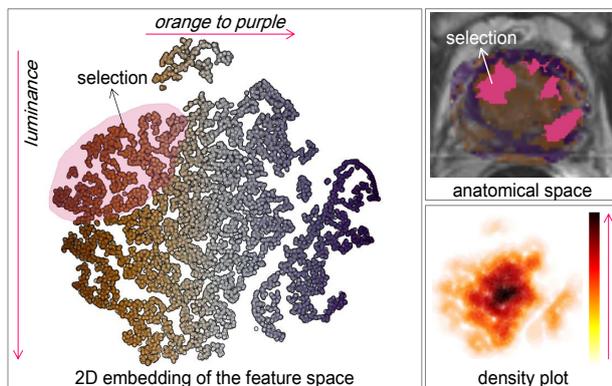


Figure 2: Visualizations used for the identification and exploration of distinct intra-tumor regions with similar tissue characteristics (*T1*).

and this color is used coherently in all views to represent it. To visualize the selections, we did not employ transparency in the points for F+C, because the visualization of the whole space is equally important, throughout the whole exploratory process. We use color, instead, to enable multiple selections. The association between the two spaces, i.e., *linking*, provides also an evaluation of the localization of the visual clusters selected in the embedding space.

The embedding provides a view on a reduced abstract 2D space. The dimensions of the embedding do not have a direct relation to the imaging-derived features. Yet, clinical researchers need to be able to link back to this feature space and *analyze the underlying high dimensional feature structure of the identified intra-tumor regions (T2)*. This part is achieved with several linked interactive views, which hold complementary data information and are interactively updated when the user selects one or more visual clusters in the embedding. Two interesting aspects of the underlying data are the *distributions* of the features of the selected regions and the *pairwise correlations* among these features, which we visualize respectively with boxplots and a simplified color-coded scatterplot matrix (SPLOM) [CL87]. These representations are intuitive and well-known to the intended clinical users (Figure 3). For a simplified correlation view, we abstract the SPLOM, by calculating and visualizing directly Pearson’s correlation (Pearson’s ρ), instead of all the points. We color-code the calculated correlation value to the divergent red-to-blue colormap [Bre94], to show the whole range from $\rho=-1$ (red) to $\rho=+1$ (blue)(Figure 3). To identify

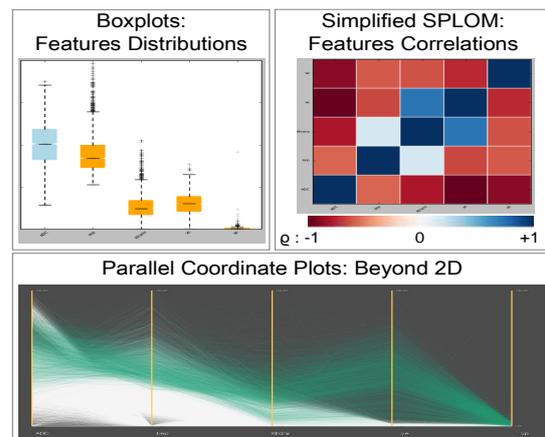


Figure 3: Visualizations used to analyze the underlying structure of the features of the each intra-tumor region (*T2*).

relations that go *beyond two dimensions* of the selected visual clusters, we additionally use parallel coordinate plots (PCPs) [Ins85] (Figure 3). PCPs are versatile in representing multiple dimensions in a single view and are commonly used to easily detect patterns, trends and outliers. In our design, PCPs can be either straight or curved polylines, rendered with low opacity for clutter reduction and improved readability. They also support bi-directional brushing and linking to the embedding, to establish connection with the high dimensional space. As a common way to link observations from different windows, colors are used coherently to denote in all views the same visual cluster.

For the easy comparison and assessment of the feature characteristics of distinct intra-tumor regions, we employ an additional view on the selected visual clusters. In this view, we provide information on the validity of each visual cluster, visualizing three commonly used internal validity measures: cohesion, separation and the average silhouette coefficient [TSK05]. Cohesion (WSS) is a measure of how closely objects are related within a visual cluster and is measured by the within sum of square distances to the mean feature vector: $WSS = \sum_{x \in C} (x - \mathbf{m})^2$, where C is the selected visual cluster, \mathbf{m} the mean feature vector of the visual cluster and \mathbf{x} a feature vector element of the visual cluster [TSK05]. Separation (BSS) reflects how distinct or well-separated a visual cluster is from another, using the between visual clusters sum of square distances: $BSS = \sum_i |C_i| (\mathbf{m} - \mathbf{m}_i)^2$, where C_i are the selected visual clusters, \mathbf{m}_i their respective mean feature vectors, $|C_i|$ the size of visual clusters and \mathbf{m} is the overall mean feature vector [TSK05]. The average silhouette coefficient (s) combines the notion of cohesion and separation: $s = \frac{BSS - WSS}{\max(BSS, WSS)}$. It ranges between 0 and 1, but it is usually interpreted in an ordinal way, i.e., 0-0.25 bad-defined visual cluster, 0.26-0.5 weak, 0.51-0.75 reasonable and 0.76-1 excellent [TSK05]. An initial option for the visualization of these indices would be a table, which would, however, need sequential analysis and be time-consuming, especially for the comparison of more than two clusters. For this reason, we abstract these numbers to glyph attributes in a 2D

view (Figure 4, left). First, we abstract each visual cluster to a sphere, which is an intuitive encoding for a group notion. The color of each sphere is used for the easy distinction of multiple visual clusters and is the same as the color used in all other views, for each visual cluster. *Cohesion* is mapped to the area and opacity of each sphere (Figure 4, bottom left): small and opaque spheres depict firm and coherent visual clusters, while large and transparent spheres depict incoherent, cloudy-like visual clusters. Transparency is needed to avoid occlusion and to emphasize the coherent visual clusters. Hereby, to preserve context despite the combined use of size and transparency, we force a minimum limit in both visual encodings. For the *separation* between two visual clusters, we considered two alternatives: encoding its value to the distance of the spheres or using an additional glyph in-between the spheres. The first option results in a cluttered view, where small spheres were often included or even hidden by the larger ones. Thus, we position among each pair of spheres a double-ended arrow glyph, whose thickness encodes the separation of the two visual clusters. Thin arrows depict well-separated visual clusters giving the illusion of distance, while thick arrows depict badly-separated visual clusters (Figure 4, bottom left). We also force a minimum limit in the arrow thickness. The choice of the double-ended arrow allows the incorporation of the *silhouette coefficient* in the visualization: each end belongs to one visual cluster of the pair and the color encodes the value of the coefficient, using a luminance colorscale (Figure 4, bottom left).

For a more detailed comparison of the visual clusters in the feature space, it is also necessary to show the most prominent features that differentiate them. At this point, we perform *Linear Discriminant Analysis (LDA)* [Bis06] between each pair of visual clusters, to calculate the vector that maximizes the linear separation between the means of these clusters, while minimizes the variance within each cluster. For each pair of clusters, we also obtain the separation histograms resulting from the projection of the high dimensional feature space of the visual clusters on the separation vector. We initially overlay the separation vector as text over the separation arrows of the visual cluster validity representation (Figure 4, upper left). A more intuitive and faster-to-perceive choice is to visualize it in a separate view as a stacked bar (Figure 4, right), to show which feature or combination of features contributes more to the separation. As we also want to show how good this separation is for each visual cluster pair, we use a matrix-like configuration to additionally show the distribution and the pairwise cluster projected histograms. This view can help identify whether and which multiple features contribute to the separation of clusters, as depicted in the right side of Figure 4. The colors in the histograms are used for visual cluster distinction and are consistent to the colors employed in all views, while the colors in the stacked bar are used for features distinction.

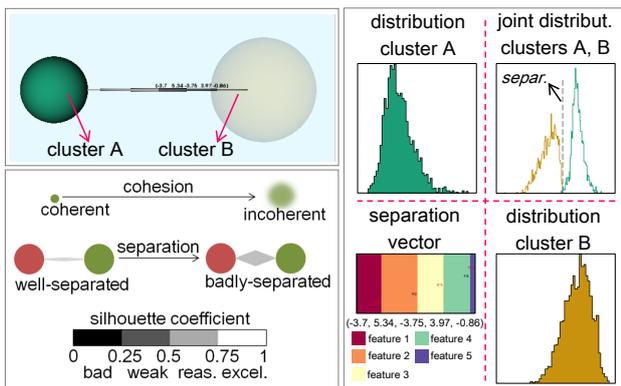


Figure 4: Cluster analysis view for comparison and assessment of two distinct intra-tumor regions (T2).

The association of the feature space with the anatomical or clinical reference space and vice versa (T3) is one

of the most important aspects for clinical users. Without the clinical reference context, users would not be able to generate and confirm hypotheses. For anatomical inspection, the visual tool provides functionality to slice through any kind of volumetric imaging data, and a linked 3D view of the organ where the tumor is located (Figure 5). To enable simultaneous inspection of clinical reference data, we need another 2D slice-based view (Figure 5). The anatomical and clinical reference views are bi-directionally linked to the embedding space, i.e., selections in one space are reflected in the other. Linking the feature space to the anatomical or reference space is performed as described in (T1). Linking the clinical reference space to the feature space is possible in two ways: by *brushing* or *color-encoding*. First, specific regions of interest can be interactively brushed in the reference space and linked to the other views (Figure 5, up). Second, the entire discretized regions of the reference data can be reflected on the embedding, with the use of a qualitative colormap [Bre94] to match each distinct region of the reference to the respective embedding points (Figure 5, bottom).

Finally, *easy exploration of the effect of modeling-induced variability and incorporation of measurements of modeling-induced inaccuracy* (T4) are also required. For *variability*, the visual tool enables side-by-side visualization of multiple linked 2D embedding maps, i.e., one per feature space. Different models result in different feature spaces and embeddings. Therefore, direct comparison of embedding spaces based on positions is not meaningful. To preserve context across multiple embedding maps and to link the points of the embedding map to their volumetric position in the anatomical space, we decided to use a simple visualization. For example, visualizations which would require to trace lines across maps, would be too complex and clutter the view. In our design, one map is used as reference. Then, the position-based 2D colormap discussed in (T1), is used to retain the information of the 2D position of each point across the multiple embedding spaces (Figure 6). All interactions in one map or in the anatomy are reflected on all maps, to strengthen the link between the different spaces. For *inaccuracy*, clinical researchers are interested either in exploring regions with

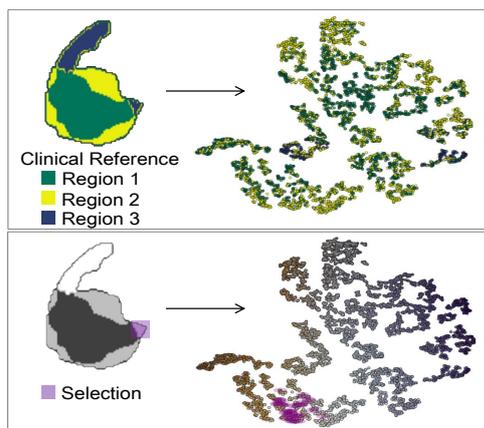


Figure 5: Linking the anatomical/clinical reference to the feature space (T3). The reverse linking is shown in Figure 2.

low accuracy to explain why inaccuracy occurs, or in restricting their analysis in regions with high accuracy. To simplify this, we employ a double slider that allows to selectively visualize accuracy ranges in the embedding. To enhance the visualization, points of the embedding map within the selected accuracy range are visualized as firm, opaque and with well-defined edges, if they have high accuracy, while lower accuracy points are transparent and with blurry edges (Figure 6).

The design of the visual tool at this point requires a large number of windows, each one of which is needed to show a specific complementary data aspect. The interface cannot be simplified drastically, given all the tasks that we want to accommodate. Yet, we enable users to selectively manipulate the profile of the tool and to selectively show the most useful representations for their specific exploratory tasks. We implemented the visual tool in Python as a DeVIDE module [BP08], employing the Visualization Toolkit (VTK), matplotlib and scikit-learn.

5. Evaluation

In order to assess the value of our visual tool, we performed an evaluation, inspired by the paper of Lam et al. [LBI*12]. The evaluation was performed with ten domain experts from three different institutions, three women and seven men. The group of participants included two research physicists, three medical physicists, four biomedical engineers and one computer scientist, who works on the automatic classification of tumor tissue. All participants have normal vision, five with and five without glasses, and nobody is colorblind. They ranked their computer expertise as intermediate to high and all have a high expertise in tumor tissue characterization. Four of the participants were already familiar with the visual tool, as they were actively involved in its design.

Before the evaluation, we gave an introduction to the visual tool, where we explained basic notions and main components. We simulated the visual environment for the exploration and analysis of the high dimensional feature space of tumor tissue characteristics, as the clinical researchers would do in real-life cases. At this point, the visual tool was initially operated by the first author, while the clinical researchers observed the demonstration of the individual components. Nevertheless, they could also operate the visual tool themselves anytime to understand better the functionality.

The first part of the evaluation included two case stud-

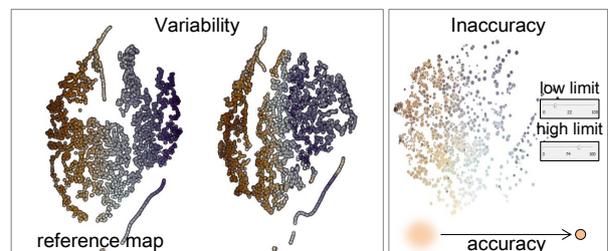


Figure 6: Visualizations for the easy exploration of the effect of variability and inaccuracy (T4).

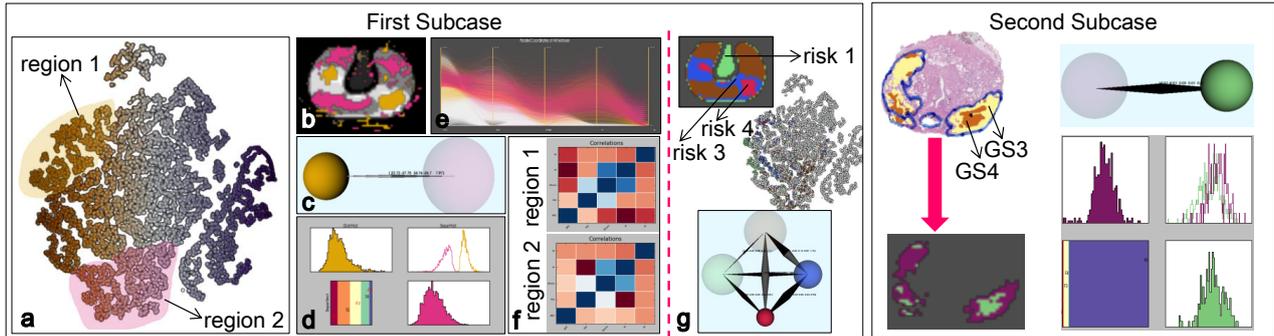


Figure 7: Case study of a prostate tumor patient.

ies, which required hands-on exploration of the data, aiming at analyzing the actual meaning of the insight provided by the visual tool, as it would occur in real clinical research settings. Each of the four tasks of Section 2 was performed with the "thinking-out-loud" method, as the clinical researchers explained and reasoned on the findings in the data. In the second part, the participants filled a questionnaire.

5.1. Evaluation: Case studies

In this section, we present part of the analysis performed by the evaluation participants, during the case studies.

Case study: Prostate tumor - In this case, data from a patient with advanced prostate tumor were employed. DCE-MRI and DW data were acquired. From these, five pharmacokinetic parameters maps (K^{trans} , k_{ep} , v_e , v_p) [SB13] and the ADC map with a b-value of 1000, were respectively derived. Afterwards, t-SNE was applied to obtain a 2D embedding of the six-dimensional feature space. This case consists of two subcases, where two different clinical reference data were used: (1) a linear regression model [GBM* 12] predicting the prostate risk zones and (2) Gleason Scores (GS) retrieved from histopathological data that reflect the aggressiveness of the tumor foci based on their microscopic appearance. To obtain the latter, the patient was scanned and then, the prostate was resected and histopathological slices were prepared and registered to the imaging data.

In the first subcase (Figure 7, left), two distant, well-separated regions of points were initially identified in the embedding map for further exploration (Figure 7-a). Linking to the clinical reference data shows that region 1 corresponds to part of the high risk zone of the prostate, while region 2 corresponds to part of the low risk zone (Figure 7-b). Using the cluster analysis view (Figure 7-c), the abstraction of region 1 shows that it is more coherent than region 2, but also that the two regions are well separated from each other. The joint histogram of the two regions also shows this separation (Figure 7-d, upper right), which is mainly due to a combined effect of three parameters (k_{ep} , K^{trans} and v_e), as it results from the vector of linear separation (Figure 7-d, lower left). This is also confirmed by the different patterns and relationships between the imaging-derived features of the two

regions, as shown by the parallel coordinates (Figure 7-e) and the correlation matrix view (Figure 7-f). The evaluation participants commented that the findings of this part of the analysis correspond to their theoretical knowledge [SB13].

When the analysis is done by color-coding the points in the projection maps based on the four discrete risk zones (Figure 7-g), the low risk cluster 1 is partially separated from the rest, but it is incoherent. Also, the two highest risk clusters 3 and 4 are reasonably coherent, despite some dispersion in the map, but not well-separated from each other. The employed risk prediction statistical model is built based on the same feature space as the one employed in the current analysis. Thus, the results from the analysis should be matching the model, but they do not. This can be an indication that the model is not able to detect potential subclusters in the high risk zones. To cross-check that this is not caused by inaccuracies during the derivation procedure of the features from the imaging data, high values of the goodness of fit, i.e., model accuracy, were selected.

In the second subcase (Figure 7, right), the analysis was done using the GS retrieved microscopically from the histopathology. In this case, the analysis was performed by going from the reference data to the embedding space. There are two regions in the histological data: GS3 and GS4, where higher GS corresponds to more aggressive tumor. We color-code the points of the embedding map with respect to the GS region that they correspond. The points of the two GS regions did not correspond to well-defined, separated visual clusters in the embedding space, which is an indication that the specific feature space is not able to reflect the GS system. However, the visual tool could be used to detect additional imaging-derived features that reflect GSs in a better way.

Case study: Cervical tumor - In this case, data from a patient with advanced cervix tumor were employed. Three different modeling approaches were used to derive pharmacokinetic parameters from DCE-MRI data. The employed pharmacokinetic models were the Tofts model (TM), the Extended Tofts model (ETM) and the 2-Compartment Exchange Model (2CXM) [SB13]. The first model results in two features (K^{trans} and v_e); the second, in three (K^{trans} , v_e and v_p); and the third in five (K^{trans} , v_e , v_p , F_p and

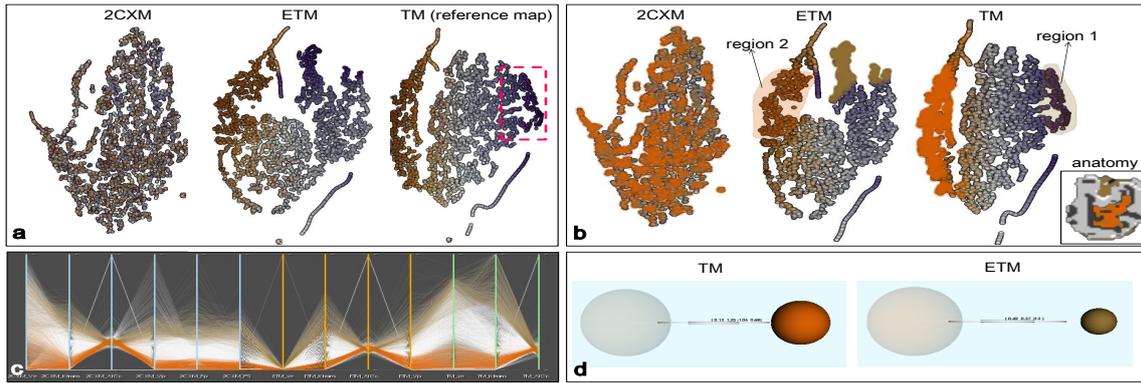


Figure 8: Case study of a cervical tumor patient.

PS) [SB13]. The Akaike information criterion (AIC_c) that relates to the quality of fit of each model, i.e., low values of AIC_c mean high relative quality, is included as an extra dimension in each feature space. For each one of the derived feature spaces, we compute a 2D embedding map. As clinical reference, we used segmented data that depict the three regions of the tumor: the periphery, the center and the in-between region. In literature [KHL*14], it has been hypothesized that different models perform better in different tumor regions and our visual tool was used for validation.

To begin with, the three embedding maps need to be compared. By colorcoding the points based on the TM, it appears that the 2CXM behaves differently than the other two models (Figure 8-a). For example, the region denoted in the red box is consistent in the ETM and TM, despite the fact that the cluster position is different across maps (Figure 8-a). However, in the 2CXM map, this visual cluster does not exist, as the points are spread. This is a confirmation to the theory that the 2CXM in some cases is unnecessarily complex and, thus, produces inaccurate results. Then, two regions were selected (Figure 8-b). From the anatomy, based on the theoretical hypothesis [KHL*14], region 1 represents the areas where the simplest model, the TM, should fit better. On the other hand, region 2 represents areas where the ETM model fits better. For region 1, this is supported by the ETM embedding map: the corresponding visual cluster presents a split (Figure 8-b). Similar observations can be drawn for the points of region 2 that also form a well-defined region in the ETM map, but present a bigger spread in the TM map. These observations are also validated by the parallel coordinates, visualizing the values of the features (Figure 8-c). To compare the underlying structure of these two regions, we use the cluster analysis view (Figure 8-d), from which we identify that the visual clusters are well-separated from each other in both ETM and TM maps, but the coherence is better in the map of better fit, i.e., region 1 in TM and region 2 in ETM.

5.2. Evaluation: Interviews

After the case studies, the participants were asked to complete a questionnaire. First, we asked four questions, related

to the four main tasks of Section 2. Each question required an open answer, but also grading using Likert scales (1-5) for the perceived effectiveness, perceived efficiency and perceived satisfaction. In order not to compromise the results, we separate in our statistical analysis the four people involved in the design and the other six who were not involved. It resulted that the two groups had in most cases comparable results. The measured variables received high scores, as summarized in Table 5.1 and in Figure 9. The scores of the variability and inaccuracy incorporation tasks were slightly lower than the other three, as some of the participants noted that they would like to use the inaccuracy component more, before giving a conclusive answer. One participant graded the efficiency of the second task lower, because he stated that the interpretation of the cluster analysis view takes time when the user is not familiar with the employed representations. The others considered that the representations adopted in the multiple views of the visual tool were intuitive.

The participants were also asked to compare the visual tool to what they are currently using and to evaluate the visual tool, as a whole, by discussing its suitability, strengths, limitations and missing features. In comparison to the currently employed approach, they commented: "It is a much more elegant approach than what I am currently using. It is very intuitive and versatile.", "I can learn more about the data and discover more about it, than with the current approach.", "I think the tool can help us to explore the feature space better.", "The information provided by this tool is very interesting and once collected for a wider population, it can be used to train a model based on more relevant features that provide a better separation of tissues."

All participants agreed that the visual tool is overall understandable and easy to learn. They also agreed that the suitability of the visual tool was mainly for data exploration, knowledge discovery and hypothesis validation or generation. Suitability for decision making might come as a result of the previous, after better familiarization with the visual tool. The strong features of the visual tool - and also what the evaluators liked more about it - are the multiple interactive linked views on the data, the link between anatomy

Table 1: Evaluation results for each of the perceived variables (effectiveness: *E*, efficiency: *e* and satisfaction: *s*) for each one of the tasks (*T1*, *T2*, *T3*, *T4*) of Section 2, as performed with the use of the proposed visual tool. We show the mean grade (μ), the standard deviation (σ) and the 95% Confidence Interval (0.95 CI), separately for the group involved in the design and the group not involved in the design. With bold, we denote cases where the second group graded the visual tool lower than the first.

	T1-E	T1-e	T1-s	T2-E	T2-e	T2-s	T3-E	T3-e	T3-s	T4-E	T4-e	T4-s
All $\mu \pm \sigma$	4.60±0.52	4.10±0.57	4.40±0.52	4.30±0.48	4.30±0.95	4.60±0.52	4.40±0.52	4.20±0.42	4.50±0.53	4.10±0.99	4.10±0.88	4.00±0.82
0.95 CI	[4.28, 4.92]	[3.75, 4.45]	[4.08, 4.72]	[4.00, 4.60]	[3.71, 4.89]	[4.28, 4.92]	[4.08, 4.72]	[3.94, 4.72]	[4.17, 4.46]	[3.48, 4.83]	[3.56, 4.64]	[3.49, 4.51]
Inv. $\mu \pm \sigma$	4.50±0.58	4.50±0.58	4.25±0.50	4.00±0.00	4.50±0.58	4.25±0.50	4.75±0.50	4.25±0.50	4.25±0.50	4.50±1.00	4.25±0.96	4.00±1.15
Not Inv. $\mu \pm \sigma$	4.67±0.52	3.83±0.41	4.50±0.55	4.50±0.55	4.17±1.17	4.83±0.41	4.17±0.41	4.17±0.41	4.67±0.52	3.83±0.98	4.00±0.89	4.00±0.63

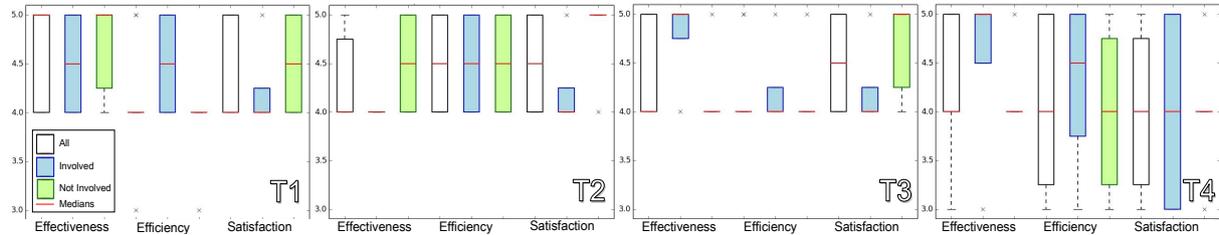


Figure 9: Schematic representation of the evaluation results of Table 5.1, for each one of the tasks of Section 2.

and feature space and the incorporation of inaccuracy in the analysis, even if the latter received a lower mean grade (Table 5.1). According to the evaluation participants, a missing feature or limitation of the visual tool is that it currently does not provide functionality for viewing simultaneously the PK parameter maps in a supplementary slice-based view, which is still important for clinical users, due to familiarity. Finally, some of the users commented that the areas of interest do not always correspond with descriptive density variation areas in the 2D embedding. Therefore, it is not always obvious how to define the selection boundaries in the embedding.

6. Conclusions and Future Work

The current exploratory approach for tumor tissue characterization based on imaging is time consuming, making it difficult to obtain the required insight. We propose a visual tool that enables clinical researchers to perform easy exploration and visual analysis of the feature space of imaging-derived tissue characteristics, and to discover new knowledge, with respect to reference data used in clinical research. We employ, as central view, a 2D embedding of the feature space, linked to multiple interactive views. These views provide information concerning the structure of the feature space, e.g., visual clusters and relations to anatomical and clinical reference information. We performed an initial evaluation with ten clinical researchers, who confirmed the usefulness of the visual tool in their analysis, as it opens new possibilities in the exploration of the feature space and provides access to new insight in the data. We illustrated this with two case studies performed during the evaluation. A direction for future work includes the extension of the application to allow meaningful follow-up or inter-patient analysis. It would also be interesting to extend the visual tool or to generalize its functionality for other applications, where also a higher number of features is involved. Finally, the incorporation of embedding precision information in the analysis

would also lead to more reliable observations. The proposed visual tool is a promising basis for clinical researchers to identify, explore and analyze heterogeneous intra-tumor regions and particularly, to generate and confirm hypotheses, with respect to clinical reference.

7. Acknowledgements

This project receives funding from the European Commission Seventh Framework Programme (Call: FP7-ICT-2011-9, activity ICT-9-5.2 - Virtual Physiological Human) and is part of the project "DR THERAPAT-Digital Radiation Therapy Patient". The authors would like to thank all the evaluation participants from the Department of Radiotherapy of Netherlands Cancer Institute (NL), the Department of Medical Physics of Aarhus University Hospital (DK) and the Department of Radiation Oncology of MAASTRO (NL).

References

- [Asi85] ASIMOV D.: The Grand Tour: a tool for viewing multidimensional data. *SIAM Journal on Scientific and Statistical Computing* 6, 1 (1985), 128–143. 3
- [BBP07] BLAAS J., BOTHA C. P., POST F. H.: Interactive visualization of multi-field medical data using linked physical and feature-space views. In *EuroVis* (2007), pp. 123–130. 3
- [Bis06] BISHOP C. M.: *Pattern recognition and machine learning*. New York: Springer., 2006. 5
- [BP08] BOTHA C. P., POST F. H.: Hybrid scheduling in the DeVIDE dataflow visualisation environment. In *SimVis* (2008), pp. 309–322. 6
- [Bre94] BREWER C. A.: Guidelines for use of the perceptual dimensions of color for mapping and visualization. In *IS&T/SPIE 1994 International Symposium on Electronic Imaging: Science and Technology* (1994), International Society for Optics and Photonics, pp. 54–63. 4, 6
- [CL87] CARR D. B., LITTLEFIELD R. J.: Scatterplot matrix techniques for large N. *Journal of the American Statistical Association* 82, 398 (1987), 424–436. 4

- [CLKP10] CHOO J., LEE H., KIHM J., PARK H.: iVisClassifier: An interactive visual analytics system for classification based on supervised dimension reduction. In *Visual Analytics Science and Technology (VAST), 2010 IEEE Symposium on* (2010), IEEE, pp. 27–34. 3
- [Dol07] DOLEISCH H.: SimVis: Interactive visual analysis of large and time-dependent 3D simulation data. In *Proceedings of the 39th conference on Winter simulation: 40 years! The best is yet to come* (2007), IEEE Press, pp. 712–720. 3
- [EAJA*05] EPSTEIN J. I., ALLSBROOK JR W. C., AMIN M. B., EGEVAD L. L., COMMITTEE I. G., ET AL.: The 2005 International Society of Urological Pathology (ISUP) consensus conference on Gleason grading of prostatic carcinoma. *The American journal of surgical pathology* 29, 9 (2005), 1228–1242. 2
- [FMHC07] FANG Z., MÖLLER T., HAMARNEH G., CELLER A.: Visualization and exploration of time-varying medical image data sets. In *Proceedings of Graphics Interface 2007* (2007), GI '07, ACM, pp. 281–288. 3
- [GBM*12] GROENENDAAL G., BORREN A., MOMAN M. R., MONNINKHOF E., VAN DIEST P. J., PHILIPPENS M. E., VAN VULPEN M., VAN DER HEIDE U. A.: Pathologic validation of a model based on diffusion-weighted imaging and dynamic contrast-enhanced magnetic resonance imaging for tumor delineation in the prostate peripheral zone. *International Journal of Radiation Oncology* Biology* Physics* 82, 3 (2012), e537–e544. 2, 7
- [GLP14] GLASSER S., LAWONN K., PREIM B.: Visualization of 3D Cluster Results for Medical Tomographic Image Data. In *In Proc. of Conference on Computer Graphics Theory and Applications (VISGRAPP/GRAPP)* (2014), pp. 169–176. 3
- [GTP10] GLASSER S., PREIM U., TÖNNIES K., PREIM B.: A visual analytics approach to diagnosis of breast DCE-MRI data. *Computers & Graphics* 34, 5 (2010), 602–611. 3
- [GRW*00] GRESH D. L., ROGOWITZ B. E., WINSLOW R. L., SCOLLAN D. F., YUNG C. K.: WEAVE: A system for visually linking 3-D and statistical visualizations, applied to cardiac simulation and measurement data. In *In Proceedings Visualization 2000* (2000), IEEE Computer Society Press, pp. 489–492. 3
- [HDWB95] HENDLEY R. J., DREW N. S., WOOD A. M., BEALE R.: Case study. Narcissus: visualising information. In *Information Visualization, 1995. Proceedings.* (1995), IEEE, pp. 90–96. 3
- [IMI*10] INGRAM S., MUNZNER T., IRVINE V., TORY M., BERGNER S., MOLLER T.: Dimstiller: Workflows for dimensional analysis and reduction. In *Visual Analytics Science and Technology (VAST), 2010 IEEE Symposium on* (2010), IEEE, pp. 3–10. 3
- [Ins85] INSELBERG A.: The plane with parallel coordinates. *The Visual Computer* 1, 2 (1985), 69–91. 5
- [JZF*09] JEONG D. H., ZIEMKIEWICZ C., FISHER B., RIBARSKY W., CHANG R.: iPCA: An Interactive System for PCA-based Visual Analytics. In *Computer Graphics Forum* (2009), vol. 28, Wiley Online Library, pp. 767–774. 3
- [KHL*14] KALLEHAUGE J., HAACK S., LINDEGAARD J., FOKDAL L., MOHAMED S., TANDERUP K., PEDERSEN E., NIELSEN T.: Tracer kinetic model selection for dynamic contrast-enhanced MRI of locally advanced cervical cancer. *Radiotherapy and Oncology* 111 (2014), 174. 2, 8
- [LBI*12] LAM H., BERTINI E., ISENBERG P., PLAISANT C., CARPENDALE S.: Empirical studies in information visualization: Seven scenarios. *IEEE Transactions on Visualization and Computer Graphics* 18, 9 (2012), 1520–1536. 6
- [LVLR09] LINSEN L., VAN LONG T., ROSENTHAL P.: Linking multidimensional feature space cluster visualization to multifield surface extraction. *Computer Graphics and Applications, IEEE* 29, 3 (2009), 85–89. 3
- [ODH*07] OELTZE S., DOLEISCH H., HAUSER H., MUIGG P., PREIM B.: Interactive visual analysis of perfusion data. *Visualization and Computer Graphics, IEEE Transactions on* 13, 6 (2007), 1392–1399. 3
- [PEPM12] POCO J., ELER D. M., PAULOVICH F. V., MINGHIM R.: Employing 2D projections for fast visual exploration of large fiber tracking data. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 1075–1084. 3
- [POM*09] PREIM B., OELTZE S., MLEJNEK M., GROLLER E., HENNEMUTH A., BEHRENS S.: Survey of the visual exploration and analysis of perfusion data. *Visualization and Computer Graphics, IEEE Transactions on* 15, 2 (2009), 205–220. 3
- [RvdHvH*14] RAIDOU R. G., VAN DER HEIDE U., VAN HOUTD P., BREEUWER M., VILANOVA A.: The iCoCoN: Integration of Cobweb Charts with Parallel Coordinates for Visual Analysis of DCE-MRI Modeling Variations. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2014), Ivan Viola K. B., Ropinski T., (Eds.), Eurographics, Eurographics Association, pp. 11–20. 3
- [SB13] SOURBRON S. P., BUCKLEY D. L.: Classic models for dynamic contrast-enhanced MRI. *NMR in biomedicine* 26, 8 (2013), 1004–27. 2, 7, 8
- [SBG00] SPRENGER T. C., BRUNELLA R., GROSS M. H.: H-BLOB: a hierarchical visual clustering method using implicit surfaces. In *Proceedings of the conference on Visualization'00* (2000), IEEE Computer Society Press, pp. 61–68. 3
- [SFHB08] SOMFORD D. M., FÜTTERER J. J., HAMBROCK T., BARENTSZ J. O.: Diffusion and perfusion MR imaging of the prostate. *Magnetic resonance imaging clinics of North America* 16, 4 (2008), 685–695. 2
- [SMB*10] STEENWIJK M. D., MILLES J., BUCHEM M., REIBER J., BOTHA C. P.: Integrated visual analysis for heterogeneous datasets in cohort studies. In *IEEE VisWeek Workshop on Visual Analytics in Health Care* (2010). 3
- [SS02] SEO J., SHNEIDERMAN B.: Interactively exploring hierarchical clustering results. *Computer* 35, 7 (2002), 80–86. 3
- [TPRH11] TURKAY C., PARULEK J., REUTER N., HAUSER H.: Interactive visual analysis of temporal cluster structures. In *Computer Graphics Forum* (2011), vol. 30, Wiley Online Library, pp. 711–720. 3
- [TSK05] TAN P.-N., STEINBACH M., KUMAR V.: *Introduction to Data Mining*. Addison Wesley, 2005. 5
- [TTP*10] TURKBEBY B., THOMASSON D., PANG Y., BERNARDO M., CHOYKE P. L.: The role of Dynamic Contrast-Enhanced MRI in cancer diagnosis and treatment. *Diagnostic and interventional radiology*. 16, 3 (2010), 186–92. 2
- [VdMH08] VAN DER MAATEN L., HINTON G.: Visualizing high-dimensional data using t-SNE. *Journal of Machine Learning Research* 9, 85 (2008), 2579–2605. 3
- [vHvW04] VAN HAM F., VAN WIJK J. J.: Interactive visualization of small world graphs. In *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on* (2004), IEEE, pp. 199–206. 3
- [War94] WARD M. O.: XmdvTool: Integrating Multiple Methods for Visualizing Multivariate Data. In *Proceedings of the Conference on Visualization '94* (Los Alamitos, CA, USA, 1994), VIS '94, IEEE Computer Society Press, pp. 326–333. 3