Eurographics Workshop on Visual Computing for Biology and Medicine (2023)
C. Hansen, J. Procter, and R. G. Raidou (Editors)

Visual Analytics to Support Treatment Decisions in Late-Stage Melanoma Patients

C. Pereira[1], U. Niemann[2], A. Braun[3], M. Mengoni[3], T. Tüting[3], B. Preim[1], M. Meuschke[1]

[1]Department of Simulation and Graphics, University of Magdeburg, Germany
[2]Department of Computer Science, University of Magdeburg, Germany
[3]University Dermatology Clinic, University Hospital Magdeburg, Germany

Abstract

We present a visual analytics system to support treatment decisions in late-stage Melanoma patients. With the aim of improving patient outcomes, personalized treatment decisions based on individual characteristics and medical histories are crucial. The research focuses on the design and development of a visual analytics system tailored specifically for tumor boards, where multidisciplinary teams collaborate to make informed decisions. By leveraging a comprehensive database containing treatment and tumor stage progression information from over 1100 patients, the system provides healthcare professionals with a holistic overview and facilitates the analysis of individual cases as well as comparisons between multiple patients. The distinction between tumor board preparation systems and systems used during discussions is emphasized to ensure user-centric design and usability. Through the use of visual analytics techniques, complex relationships between treatment outcomes, temporal features, and patient-specific factors are explored, enabling clinicians to identify patterns and trends that may impact treatment decisions. The findings of this research contribute to the growing field of visual analytics in healthcare and have the potential to enhance treatment decision-making and patient care in late-stage cancer scenarios.

CCS Concepts

• ***Human-centered computing*** → *Visualization;* • ***Computing methodologies*** → *Computer graphics;*

1. Introduction

Late-stage cancers have advanced local tumors, local and often remote metastasis. Due to rapid tumor therapy advances, even these advanced diseases can be treated with immune therapy or new chemotherapy. However, considering the patient's entire history makes treatment decisions at this stage difficult. As a specific example, we deal with melanoma (skin cancer)—a medical condition that affects about 150.000 persons worldwide[†].

Late-stage melanoma patients face complex treatment decisions requiring a thorough understanding of their long medical histories and temporal features. To make informed decisions, physicians need a good overview of patient databases with treatment regimens and tumor stage progression. Visual analytics leverages interactive and visual techniques to analyze and interpret complex patient data to aid treatment decisions.

A crucial part of this research is the creation of a tumor board interface, which acts as a collaborative platform for multidisciplinary teams engaged in cancer treatment decision-making. The objective is to design a personalized tumor board interface that considers patient-specific characteristics to provide tailored treatment recommendations. By differentiating between tumor board preparation systems and systems used during tumor board discussions, the emphasis is on a user-centric design catering to the preferences and needs of physicians. The developed interface enables in-depth analysis of individual patient cases while facilitating comparisons among multiple patients. Utilizing visual analytics techniques, clinicians can uncover intricate relationships between treatment outcomes, tumor characteristics, and patient-specific factors. This capability aids in identifying hidden patterns, trends, and potential treatment options within large and heterogeneous datasets.

In this paper, we present a visual analytics system to support treatment decisions in tumor board meetings on the example of late-stage melanoma patients. We discuss the importance of long patient histories and temporal features in understanding disease progression and treatment response. Furthermore, we highlight the significance of a comprehensive tumor board interface that supports personalized treatment recommendations based on patient-individual characteristics. Our research contributes to the growing field of visual analytics in healthcare by bridging the gap between data analysis and clinical decision-making, ultimately aiming to improve outcomes for late-stage cancer patients.

† https://www.wcrf.org/cancer-trends/skin-cancer-statistics/

DOI: 10.2312/vcbm.20231207

2. Related Work

Visual analytics provide clinicians with powerful tools to integrate and analyze diverse patient data, such as clinical records, histopathology images, genetic profiles, and treatment options. In this section, we review existing literature on visual analytics applications to analyze electronic health records (EHR), focusing on studies that address decision support in tumor boards. Moreover, we summarize approaches to compute the similarity between patient records.

Visualization of EHR. EHR contains longitudinal patient data, encompassing health status, concerns, therapies, and outcomes. Rind et al. [RWA*13] provide a comprehensive survey about advanced temporal visualization techniques to provide valuable insights. *LifeLines* [PMR*96] introduced horizontal bars representing event time periods and locations on a time axis. The zooming feature accommodates dense data. *TimeLine* [BAK07] consolidates, restructures, and displays EHR from multiple medical databases in a LifeLines-like manner. For supporting interactive exploration of EHR tools like *Lifelines2* [WWPS10], *Temporal summaries* [WPS*09], and *Outflow* [WG12] were proposed. Temporal summaries use a histogram-like chart to display the distributional trends of the occurrence rate of events over time. Outflow [WG12] utilized a customized user chart to examine disease progression trajectories with multiple events, their order, and their effects.

Our tool was influenced by LifeLines [PMR*96], which employed event and time glyphs. To enable comparisons, we adopted vertically aligned visualizations similar to those found in LifeLines2 [WWPS10]. However, unlike LifeLines2, our tool does not rely on known reference events for alignment. Defining events would be too restrictive and contradictory to our tool's goal of facilitating comparisons across various aspects of patients' clinical data. Static patient features, such as age and gender, hold equal importance alongside temporal data. Following visualization strategies of Oeser et al. [OGD*18] and Steinhauer et al. [SHB*20], we incorporated designated sections to accommodate static data.

Tumor board visual aid systems can be categorized into preparation systems and systems for information sharing during the tumor board. The 'Oncoflow' system [KNKMB12] manages data during the tumor board preparation, while Roche's 'NAVIFY' app [KCG18] streamlines administrative tasks for tumor board preparation. Several applications [KCG18, TDR*22, WKE*22] generate automated meeting presentations. Cypko et al. [CWS*17] proposed a TNM staging decision support system, and Macchia et al. [MFP*22] developed a case prioritization system.

However, these systems are not designed to support the tumor board itself, where easy access to information in a demanding environment is crucial. Therefore, visual tools such as 'CareVis' [AM06] and 'CareCruiser' [GAK*11] were developed to show how prior treatments affect patient health, displaying key parameters and therapy progression over time. Oeser et al. [OGD*18] created a head and neck tumor board interface. A map of information (MOI), presented key patient characteristics, recent treatment, and diagnostics. However, this system does not provide a view of disease progression in relation to previous treatments.

Later, Steinhauer et al. [SHB*20] introduced a dermatological tumor board MOI. Inspired by LifeLines, patient characteristics, treatments, and status are depicted by timelines and glyphs. However, the use of multiple visual encodings, such as location, border, shape, and color, results in a high cognitive load and complicates interpretation. Changes in TNM staging for assessing cancer malignancy [AGE*17], and ECOG performance status, measuring a patient's functional level [OCT*82], are shown by horizontal bars. Precise values are only available in the comparative view for two points in time. The analysis lacks a comprehensive consideration of TNM and ECOG, and stage durations must be deduced.

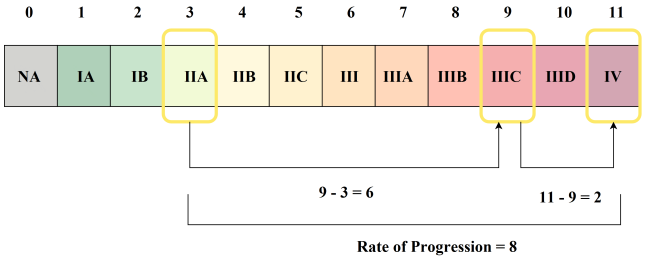

Figure 1: *Calculation of the progression rate based on the clinical stage progression of a patient from IIA → IIIC → IV.*

Based on the work by Steinhauer et al. [SHB*20], we introduce a tumor board visualization tool that effectively encodes patient data essential for tumor board discussions. Our tool addresses previous visual impediments through appropriate visualizations.

Patient similarity computation aids decision-making by identifying disease subtypes or patient groups [LCG*15, CZY*07]. Furthermore, patient similarity can be leveraged in clinical outcome prediction models, e.g., to predict ICU patient mortality [LMD15] or to compute patient survival probability [MMB*19].

Maximizing the effectiveness of cancer treatment relies on considering the location and spread of the disease. Luciani et al. [LWE*20] introduced a method to incorporate spatial information when comparing patients within a cohort. Srabanti et al. [STA*22] identified differences in demographic information, disease characteristics, treatments, and outcomes among cohorts. Floricel et al. [FNB*22] proposed THALIS, a tool that utilizes association rule and factor analysis to predict and explain the longitudinal development of patient symptoms in a cohort.

Depending on the goal, the most similar patients or a patient group with similar features are relevant. Clustering is used to create patient groups [PMSB18], but similarity metrics are more suited for finding the most similar patients. Our similarity computation retrieves patients most similar to the patient of interest.

Euclidean distance [WS15], Minkowski distance [LEYK16], Cosine distance [TWW18], and Mahalanobis distance [VMP*18] were used to measure patient similarity. Dynamic time wrapping (DTW) [YJF98], longest common sequence (LCS), and Edit distance [CN04] are used to determine sequence similarity for temporal event-sequence data.

There are several similarity metrics tailored to specific data types [WS15, SLG*16, LDC*22]. These metrics calculate feature-level similarities and combine them into a patient-level mea-

sure. Moreover, feature weighting techniques have been employed [WHL*19, LDC*22]. Drawing inspiration from Wang et al. [WHL*19] and Liu et al. [LDC*22], we create custom similarity metrics that integrate feature-level metrics into a single measure.

3. Requirement Analysis

Our requirement-gathering process followed a three-stage approach, inspired by Young [You02], to meet the needs of a tumor board meeting, with dermatology serving as a representative use case.

In the first stage, extensive discussions were held with dermatologists to gain a deep understanding of the melanoma tumor board meeting and to identify the core requirements for our system. Next, storyboards were utilized as a visual tool to illustrate the system's capabilities and interfaces. These storyboards were presented to dermatologists and other stakeholders to gather feedback and refine the requirements. Finally, a prototype was developed to provide a tangible demonstration of the system's interactions. This prototype underwent a thorough iterative process, incorporating feedback from dermatologists and other stakeholders. By iterating on the prototype, we ensured that the requirements were comprehensive.

While dermatology served as the specific use case for this paper, the derived requirements likely hold broader relevance for tumor board meetings across different medical specialties. In summary, the following requirements were identified:

R1 General system requirements. Fundamental requirements for the system to be used in the context of a tumor board:

- R1.1 Web-based application to avoid software or hardware dependence in clinical settings.
- R1.2 Ability to incorporate access control.
- R1.3 Search with patient overview information to find the patient of interest (POI).
- R1.4 Easy-to-use navigation methods.

R2 Visualization of patient history. The requirements regarding the visualization of single patient history:

- R2.1 Information required for introducing the patient should be easily accessible.
- R2.2 Effective visualization of temporal patient characteristics and treatments.
- R2.3 Patient condition changes must be easily identifiable.
- R2.4 Frequently needed patient information, such as clinical stage and treatments, should be easily accessible.
- R2.5 Additional supplementary information, like the exact duration of a stage, should be available on demand.

R3 Patient similarity computation. The requirements that the calculation of patient similarity must meet:

- R3.1 Retrieve patients with characteristics similar to the POI.
- R3.2 Enable users to customize patient similarity definition by changing the importance of characteristics.

R4 Patient similarity visualization. The requirements for visualization of similar patients:

- R4.1 Easily interpretable visualization that facilitates identification of patients most similar to the POI.
- R4.2 Effective visualization to facilitate characteristics of similar patients with those of the POI.
- R4.3 Summarized characteristics of POI along with a selected similar patient that can be extracted for import in other systems.
- R4.4 Easy-to-use similarity customization mechanism.
- R4.5 A visualization that allows comparison of the clinical course of the POI with a similar patient.

4. Data Acquisition and Pre-Processing

Our dermatologists provided access to 2402 melanoma patient records, ensuring their confidentiality and privacy. Patient data includes demographic information, clinical history, diagnostic procedures, treatment modalities, and follow-up details. These variables were carefully selected to capture relevant information related to melanoma diagnosis, prognosis, and treatment outcomes. Moreover, these variables align with the data commonly collected for cancer patients , indicating their relevance beyond melanoma cases.

On diagnosis, the patient's demographic data, specifically age, and gender, were documented alongside the clinical and TNM stage. The clinical stage indicates melanoma's depth of infiltration into the skin and its spread [AGE*17]. For effective treatment, it is also crucial to know the aggressiveness of the tumor. This variable is derived from the clinical stages progression of the patient. The clinical stages were assigned chronological numerical values. Based on the clinical stage the patient passed through, dermatologists calculated the progression rate as the aggregation of the difference between successive stages. Fig. 1 illustrates the calculation for the progression rate of a patient who advanced from stage IIA to stage IIIC and then stage IV. The transition from stage IIA to stage IIIC is assigned a value of 6, whereas the transition from stage IIIC to stage IV is allocated a value of 2. The patient's progression rate is 8.

The TNM staging from the AJCC is a standardized classification system for assessing the malignancy of cancer, both during the initial diagnosis and after surgical intervention [AGE*17]. TNM staging describes tumor thickness (T), the presence of cancerous cells near the lymph nodes (N), and metastatic spread (M). The location of the main and metastatic tumors, sentinel, and treatments administered on diagnosis were also recorded. Subsequently, upon each visit of a patient to the clinic, a new record was created, which documented the patient's clinical and TNM stages, prescribed therapies, sentinel, and location of the metastasis at the particular time point. During the metastatic stage, histological data from biopsies are also recorded.

While most variables apply to a wide range of cancer patients, sun exposure specifically pertains to melanoma cases. Sun exposure is a significant risk factor for melanoma development. It is derived from the primary location of the tumor. The inclusion of sun exposure allows for a focused investigation into the relationship between sun exposure patterns and melanoma incidence, aiding in the understanding of this particular cancer type. In summary, the following information is recorded for each patient:

- General details: admittance number, diagnosis date, death date
- Demographic information: sex, age at diagnosis
- Clinical stages: NA, IA, IB, IIA, IIB, IIC, III, IIIA, IIIB, IIIC, IIID, IV

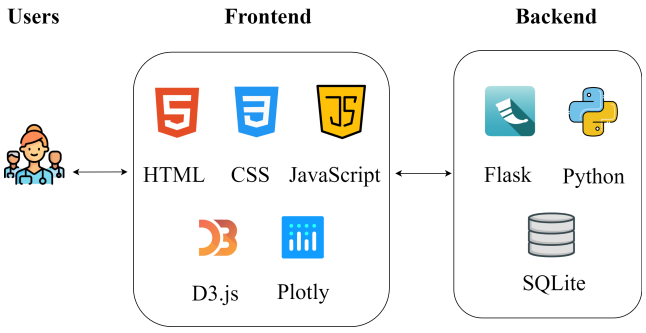

Figure 2: *Architecture of our melanoma tumor board system.*

- TNM stages: TNM stage's value is a combination of T, N, and M. T: Tis, T0, T1, T2, T3, T4; N: N0, N1, N2, N3; M: M0, M1
- Location information: primary tumor, secondary Tumor
- Metastasis data: date, organs of metastasis
- Histological data: Location, BRAF (an oncogene whose mutation is indicative of melanoma metastasis [LMN*05]), Clark (level of skin invasion of the melanoma [CFBM69]), thickness, actinic (indicative of sun damage and precursors to skin carcinomas [GSC*05]), ulceration, sentinel, availability of excision at the clinic.
- Therapy information: excision, targeted therapy, immunotherapy, chemotherapy, others
- Derived variables: sun exposure, rate of progression

Data pre-processing. 1177 of the 2402 anonymized patients had visited the clinic more than once; they constitute the final dataset. Tumor board deliberations only take recurring patients into account, so single-occurrence patients were excluded from the dataset. In order to maintain consistency, diverse representations of null values, including 'k.A', 'null', and 'NA', were standardized to null. Leading and trailing white spaces from textual data were trimmed.

5. Interactive Tumor Board Visualization

This section describes our system in detail, starting with an overview of its architecture and general user interface (Sec. 5.1). Then, the two major components are introduced, the visualization of the tumor information of the POI (Sec. 5.2), followed by the computation and visualization of similar patient records (Sec. 5.3 and Sec. 5.4).

5.1. System Architecture and General User Interface

The melanoma tumor board system architecture is shown in Fig. 2. The user accesses the system via the front end, which utilizes HTML and CSS to organize and style the web pages. JavaScript libraries, namely Plot.js [Inc15] and D3.js [BOH11], are used to create interactive visualizations. The web application was built using *Flask*, a lightweight and easy-to-use web framework in Python, making integration of other Python libraries easier (meet R1.1). The anonymized melanoma patient data is stored in an SQLite database.

The system contains sensitive patient data. Thus, user registration and authentication pages were incorporated (meet R1.2). After authentication, the user is redirected to the *Patient Overview* page,

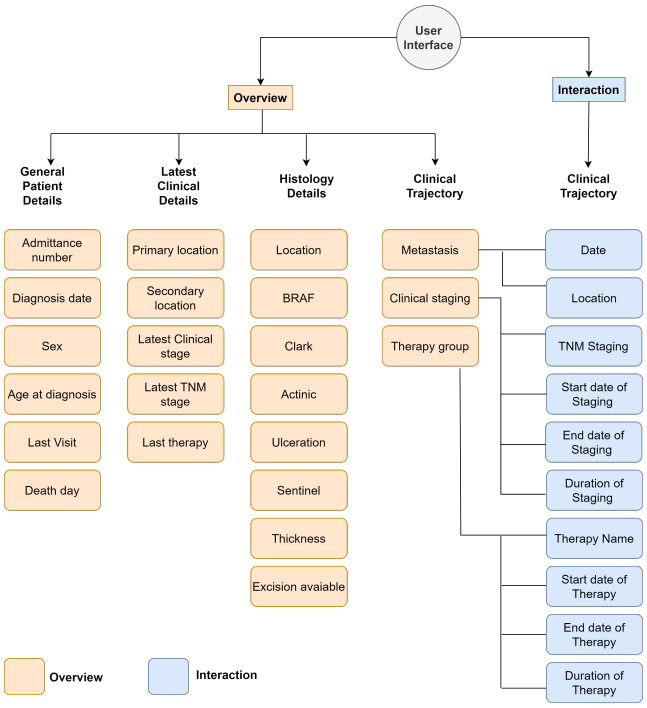

Figure 3: *The MOI is split into* Overview*: directly accessible information upon loading, and* Interaction*: on-demand information.*

which displays patients and their characteristics in a tabular format. This view enables users to search for the POI (meet R1.3). A search bar can be used to retrieve patients with a particular feature value. Moreover, records can be sorted according to their feature values. The upper right corner of the page contains links to pages that can be navigated from the current page (meet R1.4). When the POI is selected, the *Tumor Board* page for that patient is displayed.

5.2. Tumor Board Visualization for the POI

This section introduces the information architecture, structured as a map, for the development of the tumor board visualization specifically designed for the POI. Furthermore, it offers a comprehensive explanation of the visualization techniques employed in this view.

Map of Information (MOI). To present the patient's medical history, we created a *dermatologist-informed MOI* based on the information architecture model from Oeser et al. [OGD*18], see Fig. 3. The MOI is divided into two parts: the *Overview* and the *Interaction*. The overview provides essential patient information that is always accessible upon loading. The interaction part contains additional information that can be displayed on demand.

POI-centered visualization. The tumor board view for a patient is organized based on the MOI. For an effective clinical overview, the system integrates visualizations and textual content, see Fig. 4.

The information needed to introduce a patient is textually represented and organized into three categories (meet R2.1). The first category 'General Patient Details' lists the patient's admission number, diagnosis date, gender, age at diagnosis, last visit date, and

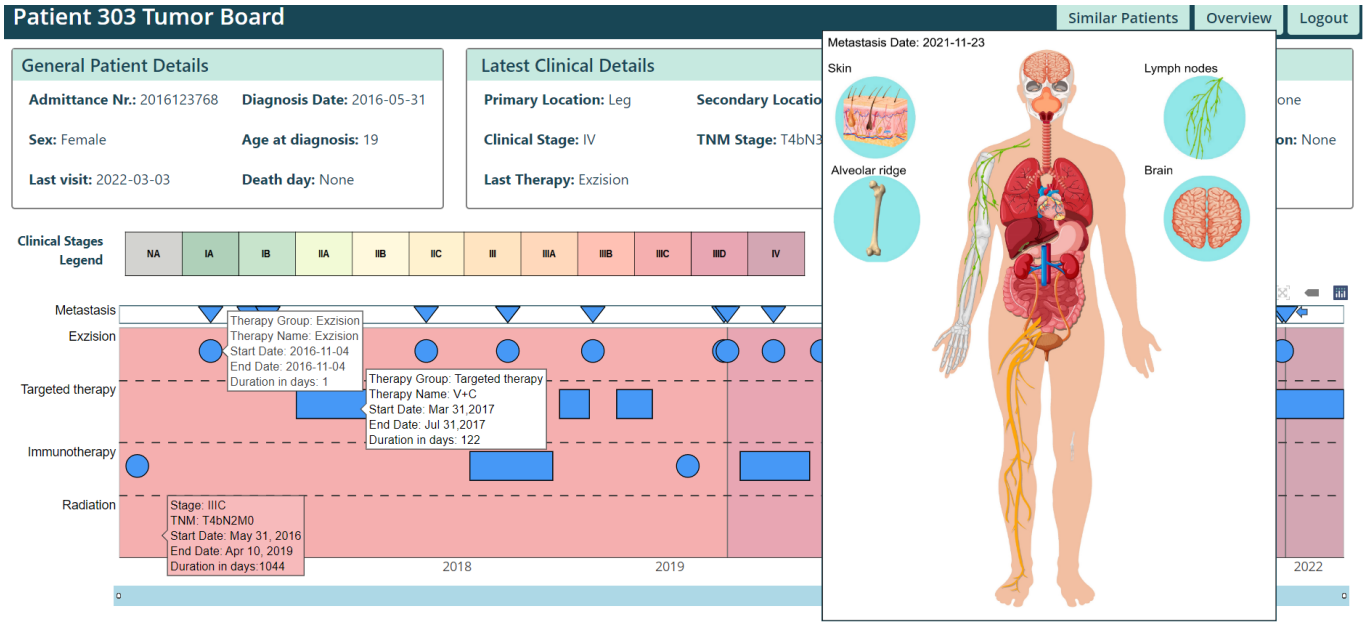

Figure 4: *Visualization of the tumor information for the POI based on the MOI. Information required for patient introduction is displayed in the textual format. Interactive visualization is used to depict the patient's clinical trajectory.*

death date. The patient's primary and metastatic locations, clinical and TNM stages, and last treatment are summarized under 'Latest Clinical Details'. The third category, 'Histology Details', lists histological characteristics such as Location, BRAF, Clark, Thickness, Actinic, Ulceration, Sentinel, and Excision available.

The patient's clinical trajectory is depicted by a timeline visualization. The temporal dimension spans from the date of diagnosis to the last visit. The range slider enables the selection of time periods of interest (meet R2.2). The administered therapies are along the y-axis, demarcated by horizontal dashed lines. The rectangles depict stages and stage transitions. A new rectangle is created each time there is an alteration in either the TNM or clinical stage (meet R2.3). Leveraging the cognitive linkage for green and red colors [MS20], a chromatic scale diverging from green to red, created using color brewer palettes [HB03], encodes the clinical stages (meet R2.4). On hover, the clinical and TNM stage, the start and end dates of that stage, and the stage duration are shown (meet R2.5).

A circular glyph denotes a therapy administered to a patient on a single day. Rectangular glyphs are used for therapies administered over a period of time. The glyphs are saturated in blue to make them prominent in the visualization (meet R2.4). On hovering glyphs, the therapy group and name, along with the start date, end date, and therapy duration, are displayed (meet R2.5). All occurrences of metastasis are depicted as triangles pointing downward. On hover, the date of metastasis, along with their locations, are shown. Supplementary functionalities, such as zoom, pan, and autoscale, have been incorporated to support data analysis.

5.3. Patient Similarity Computation

The definition of patient similarity is contingent upon specific criteria, necessitating the inclusion of domain knowledge of dermatologists [Hah14]. The features of sex, age at diagnosis, primary location, sun exposure, sentinel, clinical stages, TNM stages, organs of metastasis, and administered therapies were considered relevant for defining the similarity between melanoma patients. Similar to [WHL*19] and [LDC*22], we employed custom similarity metrics by computing feature-level similarity and combined them to get a patient-level similarity measure. We incorporated a weighting mechanism to assign significance to the features. Additionally, the Heterogeneous Euclidean Overlap Metric (HEOM) is considered as it is capable of accommodating both numerical and categorical values [WM97].

Custom similarity metrics. The feature-level metrics were chosen based on the data type of the feature. Let P_1 and P_2 be the patients under consideration, f_i be the feature for which the similarity is to be computed and $v^1_{f_i}$ and $v^2_{f_i}$ be the feature values of patient P_1 and P_2 for feature f_i, respectively.

For the numerical features, namely age at diagnosis and rate of progression, the similarity between P_1 and P_2 is calculated by a normalized difference (ND) (Eq. 1). Feature values are normalized between 0 and 1 to avoid the impact of magnitude.

$$S_{f_i}(P_1, P_2) = 1 - |v^1_{f_i} - v^2_{f_i}| \tag{1}$$

For the static categorical features, namely primary location, sun exposure, and sentinel, an equality check (EC) is used for the simi-

larity computation. If f_i has the same value for both patients, they are similar; otherwise, they are dissimilar (Eq. 2).

$$S_{f_i}(P_1, P_2) = \begin{cases} 1 & \text{if } v^1_{f_i} = v^2_{f_i} \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

The values of temporal categorical features are comma-separated sequences in the order of their occurrence. Temporal categorical features are clinical stage, TNM stage, treatments, and metastasis. LCS and the Jaccard index are utilized as alternatives to compute the similarity between patients for these features. The LCS between P_1 and P_2 for f_i is given by Eq. 3, where the resulting value is normalized between 0 and 1 for similarity computation.

$$S_{f_i}(P_1, P_2) = \frac{LCS(v^1_{f_i}, v^2_{f_i})}{max(len(v^1_{f_i}), len(v^1_{f_i}))} \tag{3}$$

With the Jaccard index, the value of f_i is treated as a quantity. Here, the similarity between P_1 and P_2 for f_i is calculated as the size of the intersection of $v^1_{f_i}$ and $v^2_{f_i}$ divided by their union (Eq. 4).

$$S_{f_i}(P_1, P_2) = \frac{\left| v^1_{f_i} \cap v^2_{f_i} \right|}{\left| v^1_{f_i} \cup v^2_{f_i} \right|} \tag{4}$$

These feature-level metrics were combined to create unified patient-level similarity metrics. The inclusion of weights allowed for the assignment of significance to features. The unified patient level similarity between P_1 and P_2 is computed as the weighted sum of the feature level similarities as in Eq. 5.

$$US(P_1, P_2) = \sum_{i=1}^{n} w_i \cdot S_{f_i}(P_1, P_2) \tag{5}$$

Finally, we define two custom similarity metrics (CM1 & CM2). Both metrics used Equality check and Numerical difference for static categorical and numerical features, respectively. For the temporal categorical features, CM1 & CM2 employed LCS and Jaccard, respectively.

Heterogenous Euclidean Overlap Metric (HEOM). The HEOM is considered since it handles mixed data types. By default, all features are assigned equal importance. The HEOM is updated to incorporate feature weighting as in Eq. 6 to customize similarity by providing feature importance.

$$HEOM(P_1, P_2) = \sqrt{\sum_{i=1}^{n} (w_i \cdot \delta(v^1_{f_i}, v^2_{f_i}))^2} \tag{6}$$

where

Metric	HEOM	CM 1	CM 2
Default weights	0.48	0.64	0.68
Modified weights	0.42	0.56	0.56

Table 1: *For default weights, CM2 returned the most patients similar to the POI, present in the expected results. CM1 and CM2 performed comparably in the modified weights scenario. HEOM had low precision with default weights and weighted features. Based on these findings, CM2 was implemented in the tumor board system.*

$$\delta(v^1_{f_i}, v^2_{f_i}) = \begin{cases} 1 & \text{if } v^1_{f_i} \text{ or } v^2_{f_i} \text{ is missing} \\ \text{olap}(v^1_{f_i}, v^2_{f_i}) & \text{if } f_i \text{ is nominal} \\ \text{diff}(v^1_{f_i}, v^2_{f_i}) & \text{if } f_i \text{ is continuous} \end{cases} \tag{7}$$

with

$$\text{olap}(v^1_{f_i}, v^2_{f_i}) = \begin{cases} 0 & \text{if } v^1_{f_i} = v^2_{f_i} \\ 1 & \text{otherwise} \end{cases} \tag{8}$$

and

$$\text{diff}(v^1_{f_i}, v^2_{f_i}) = v^1_{f_i} - v^2_{f_i} \tag{9}$$

Evaluation of similarity metrics. Obtaining ground truth data to evaluate the results of similarity metrics, with and without feature weightings, proved to be impractical. To address this limitation, we conducted an evaluation of our similarity computation using a limited set of manually curated ground truth data for five patients. We compared the results obtained with default feature weights to those obtained with altered feature weights, taking into account the concept of similarity discussed with dermatologists. Precision was used to assess the effectiveness of the similarity metrics. Precision measures the extent to which the expected patients are present in the returned results. Ideally, evaluating the similarity metrics based on the order of the results would be preferred. However, objectively ranking the expected results is challenging. Therefore, the similarity metric focuses on evaluating the presence of the expected results in the returned results rather than their specific position within the results. Table 1 summarizes default and modified feature weight similarity metrics. For default weights, CM2 returned the most patients similar to the POI, present in the expected results. CM1 and CM2 performed comparably in the modified weights scenario. HEOM had low precision with default weights and weighted features. Based on these results, the melanoma tumor board system implemented the CM 2 to identify patients similar to the POI (meet R3.1 & R3.2).

5.4. Patient Similarity Visualization

Our system visualizes the similarity computation results for the ten most similar patients to the POI. It comprises two tabs, the *Patient Similarity Overview* and the *Patient Similarity Comparison*.

Patient Similarity Overview. This tab compares characteristics of similar patients to the POI and allows customization of similarity by assigning feature weights. The tab is partitioned into four parts.

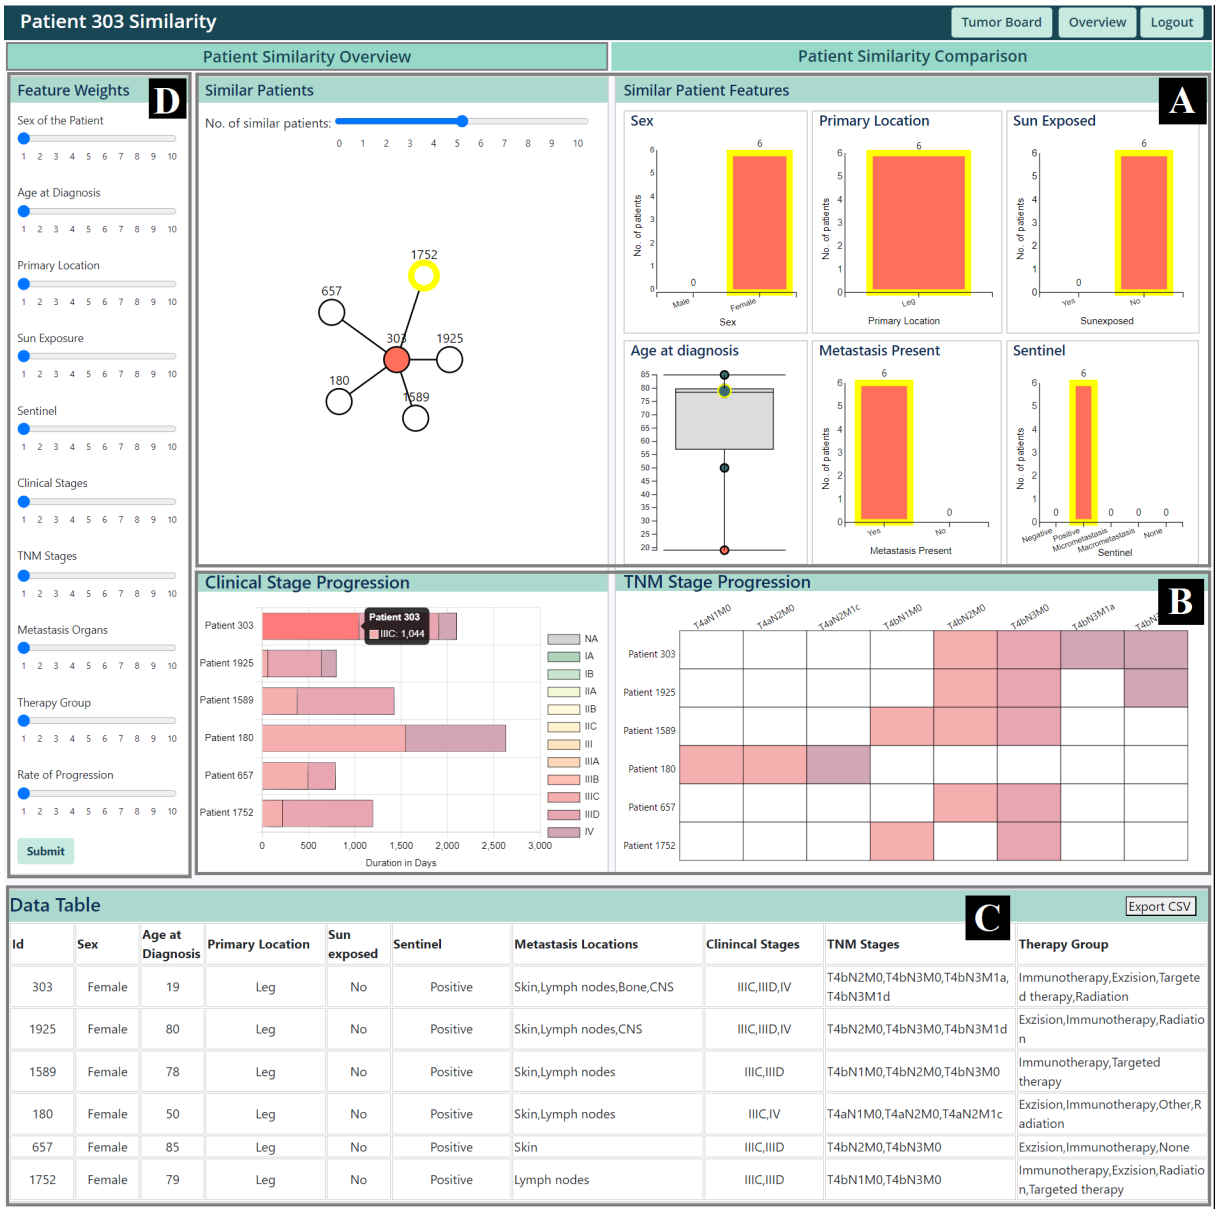

Id	Sex	Age at Diagnosis	Primary Location	Sun exposed	Sentinel	Metastasis Locations	Clinincal Stages	TNM Stages	Therapy Group
303	Female	19	Leg	No	Positive	Skin,Lymph nodes,Bone,CNS	IIIC,IIID,IV	T4bN2M0,T4bN3M0,T4bN3M1a, T4bN3M1d	Immunotherapy,Exzision,Targeted therapy,Radiation
1925	Female	80	Leg	No	Positive	Skin,Lymph nodes,CNS	IIIC,IIID,IV	T4bN2M0,T4bN3M0,T4bN3M1d	Exzision,Immunotherapy,Radiation
1589	Female	78	Leg	No	Positive	Skin,Lymph nodes	IIIC,IIID	T4bN1M0,T4bN2M0,T4bN3M0	Immunotherapy,Targeted therapy
180	Female	50	Leg	No	Positive	Skin,Lymph nodes	IIIC,IV	T4aN1M0,T4aN2M0,T4aN2M1c	Exzision,Immunotherapy,Other,Radiation
657	Female	85	Leg	No	Positive	Skin	IIIC,IIID	T4bN2M0,T4bN3M0	Exzision,Immunotherapy,None
1752	Female	79	Leg	No	Positive	Lymph nodes	IIIC,IIID	T4bN1M0,T4bN3M0	Immunotherapy,Exzision,Radiation,Targeted therapy

Figure 5: *Patient Similarity Overview that enables comparison of the POI to similar patients with respect to their features. The view comprises four parts. A) Similar patient overview and feature comparison, B) Comparison of the Clinical and TNM Stages, C) Data Table for overview and data export, and D) Feature weighting for customization of similarity definition.*

Part A illustrates the patients similar to the POI by means of a graph (see Fig. 5), as they provide a natural representation of pairwise similarities between patients (meet R4.1). Nodes are utilized to represent patients, while the equiangular radial edges differ in length, encoding the extent of similarity. By default, five patients with the highest degree of similarity are displayed; however, up to ten most similar patients can be viewed. The node of the POI is colored salmon pink, while the nodes of similar patients are uncolored.

Bar charts are the most effective and accurate way to display categorical data [SED19]. They are also preferred due to their high familiarity, which aids interpretation [QR22]. Hence, bar charts are utilized for the features of sex, location, sun exposure, metastasis, and sentinel. Box plots represent distributions and hence were used to depict the age distribution. Consistent with the color scheme in the graph, the features of the POI were colored salmon pink. Thus, each categorical feature's value for the POI can be quickly assessed. A highlight function was incorporated to determine the feature values of similar patients by clicking on the node as seen in Fig. 5 (meet R4.2).

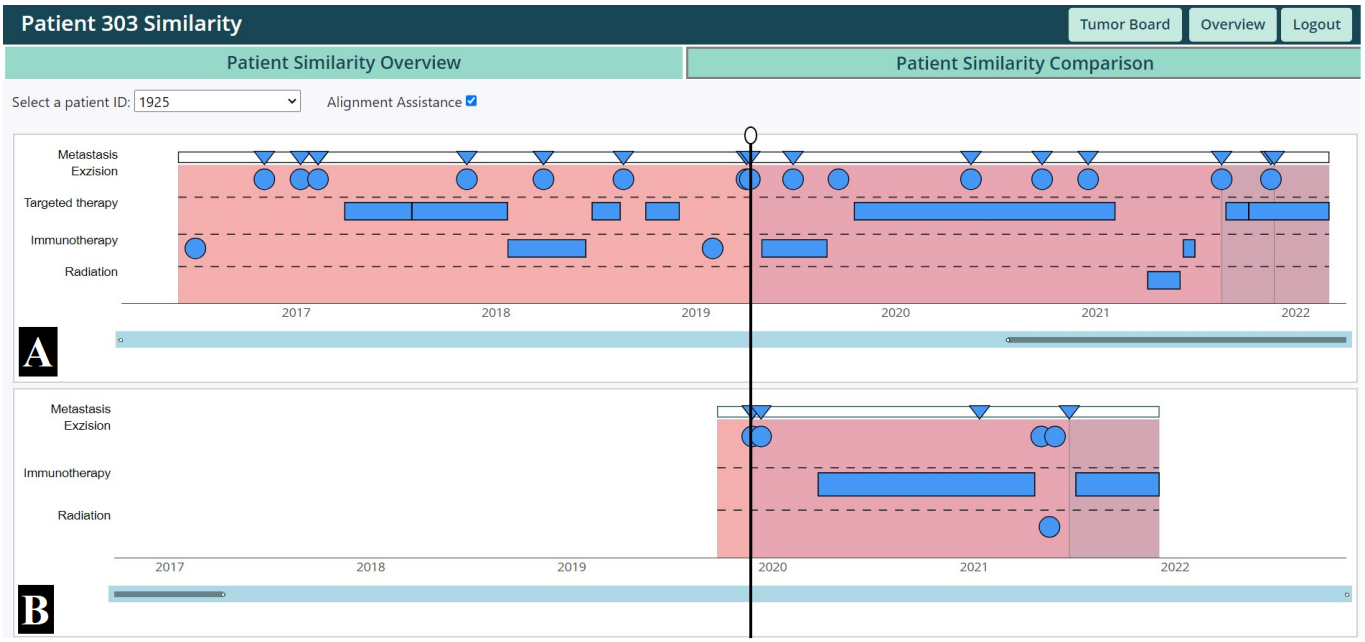

Figure 6: *Patient Similarity Comparison view to enable comparison of the clinical trajectory of the POI with a similar patient. It comprises two parts A) Clinical trajectory of the POI and B) Clinical trajectory of the selected similar patient.*

Part B enables to a comparison of clinical and TNM stages between similar patients and the POI. Howorko et al. [HBD*18] showed the effectiveness of stacked bar charts for comparing single attributes. Thus, horizontally stacked bar graphs are used to depict the duration of clinical stages, with hovering providing stage durations. A matrix visualization is employed to comparatively display TNM stages, effectively illustrating relationships along both horizontal and vertical axes. By placing these comparative visualizations adjacent to each other, it facilitates the examination of clinical stage durations alongside changes in TNM stages across different patients.

Part C features a summary data table for the POIs and similar patients' features. Tables are valuable for organizing information, fostering trust, and promoting understanding [BCT22]. To facilitate further analysis or integration with other systems (meet R4.3), the data can be downloaded in CSV format.

By default, the similarity calculation weights all features equally. The user can adjust these weights using the sliders for each feature in part D (meet R4.4). By clicking on *Submit*, the similarity is recomputed and the *Patient Similarity Overview* and the *Patient Similarity Comparison* tab are updated.

Patient Similarity Comparison. This tab allows us to compare the clinical trajectory of a patient similar to the POI (meet R4.5), see Fig. 6. On load, the clinical trajectory for the POI, as described in Sec. 5.2 is displayed. A drop-down menu showcasing similar patients, sorted by their similarity, is also provided. Upon selecting a similar patient from the menu, their clinical trajectory is shown below that of the POI. The vertical layout of the tab is designed to facilitate the comparison. The visualizations are initially generated along the temporal axis, representing the duration of each patient's treatment by default. Automated alignment of the patient's clinical trajectories would restrict the exploration and comparison to only limited aspects of the patient's clinical trajectories. There is a need for flexibility to enable the alignment of multiple aspects to make inferences. To support this, the system facilitates manual alignment of the clinical trajectories. The temporal parameters of the visualization can be adjusted to focus on specific time frames, such as considering only the initial year of treatment for both patients.

To facilitate the alignment of elements for comparison, a visual alignment aid has been incorporated. It can be activated by selecting the "Alignment Aid" checkbox. It displays a vertical line that spans both visual representations. The user can click and drag the line to any desired position to align the visualizations accordingly.

6. Evaluation

We performed a qualitative user study with five dermatologists (three females and two males) who regularly join tumor board meetings to evaluate our system's usefulness, usability, and acceptance by the physicians. Three participants had five or more years of experience, while two had less than three. One participant led tumor board discussions. Four participants had experience introducing patients to tumor boards. Two of the participants are co-authors of this work.

To address data privacy and security concerns associated with patient data, the system was not hosted externally. Instead, it was locally launched on the development machine for evaluation purposes. To gather user feedback, individual surveys were conducted using Google Forms. This approach ensured the protection of patient data while allowing for user input and assessment of the system.

Study design. The user study was conducted in four stages.

1. Our tumor board system was explained to each participant.
2. Participants were allowed to interact with the system in order to familiarize themselves with it and address any questions.
3. The melanoma tumor board system's usability and comprehension were assessed through user tasks and closed-ended questions on a five-point Likert scale (strongly disagree, disagree, neutral, agree, strongly agree) based on a selected patient case.
4. In an informal interview participants were asked about the system's pros, cons, and potential uses. Inquiries were also made about system enhancements.

The usability of the POI-centered tumor board visualization was assessed using the seven usability principles: Information coding, spatial organization, minimal actions, flexibility, recognition rather than recall, removal of the extraneous, and data set reduction from Forsell et al. [FJ10]. The proposed visualization aims to facilitate collaborative decision-making within the tumor board setting. In order to evaluate its effectiveness in this regard, the assessment includes queries derived from Lam et al.'s information visual scenarios [LBI*12]. These scenarios encompass various aspects such as understanding environments and work practices, visual data analysis and reasoning, collaborative data analysis, communication through visualization, and user experience. Similar questions were used by Steinhauer et al. [SHB*20] to evaluate their oncological tumor board. So we can compare both systems.

Concerning the usability of the patient similarity tabs, the emphasis is on the ease of information access and learning, so we take into account the usability principles of minimal actions, information coding, and spatial organization [FJ10]. The purpose of patient similarity tabs is to enable users to investigate similar patients' data and gain insights; thus, queries pertinent to the visual data analysis and reasoning scenario [LBI*12] are emphasized. Additionally, questions to validate the clinical usefulness of our system are included.

Results and Discussion. In the following, we discuss the most interesting results from the user study. The detailed results including diagrams can be found in the additional material (AM).

POI-centered visualization. In terms of the POI-centered tumor board visualization, all participants correctly answered six out of the nine task-based questions (Sec. 2.1 of the AM). For the remaining tasks, the majority of participants provided correct responses. However, their answers deviated from the expected response, but they were partially correct for two specific questions (Sec. 2.1, Quest. 10 & 13 of AM). This discrepancy may be attributed to differences in interpretation. Furthermore, there were two instances where participants inaccurately determined the duration of the clinical stage (Sec. 2.1 Quest. 12 of the AM). This could be attributed to the manual requirement for users to calculate the duration when the TNM stage changes while the clinical stage remains constant. As this increases cognitive load, the tooltip information should be improved to incorporate the TNM and clinical stages duration separately.

The majority of the participants agreed that the POI-centered tumor board visualization abides by the usability principles and meets the goals of the scenario-based questions regarding the tumor board visualization (Sec. 2.2 of the AM). The tumor board visualization for the POI meets also the R2 requirements. All participants unanimously agreed that the POI-centered tumor board visualization has clinical relevance that is better than the current setup of verbal patient case introduction (Sec. 2.4 Quest. 38 & 40 of the AM). It has the potential to be used in a tumor board meeting and aid in tumor board preparation (Sec. 2.4 Quest. 39 of the AM) The POI-centered tumor board visualization outperforms the tumor board system by Steinhauer et al. [SHB*20] in usability, suitability for the tumor board context, and in clinical relevance (Sec. 2.4 of the AM).

Patient similarity overview tab. Three of the seven task-based questions regarding this tab were answered correctly by all participants (Sec. 3.1 of the AM). Of the questions answered incorrectly, three of them are related to the interactions between the graph and the similar patient's feature part (Sec. 3.1 Quest. 35, 36 & 37 of the AM). In the similar patient's feature part, salmon pink denotes the POI; thus, the viewer may perceive the other bar as the value of the patient dissimilar to the POI overall, while in reality, it is dissimilar only with respect to the specific feature.

All participants agreed that the patient similarity overview tab complied with the usability principles of 'minimal actions' and 'information coding' (Sec. 3.2 Quest. 41, 42 & 43 of the AM). The qualitative survey results show that the system's similarity computation returns patients similar to the POI (Sec. 3.2 Quest. 45 of the AM) and allows intuitive adjustment of feature weighting that provides results that meet user expectations (Sec. 3.2 Quest. 43 & 47 of the AM), fulfilling requirements R3 and R4. Even though select users responded to tasks pertaining to the comprehension of the comparative visualization incorrectly, qualitative data shows that users generally consider the system as having a shallow learning curve (Sec. 3.2 Quest. 42 of the AM). Users unanimously agreed on the clinical relevance of the patient similarity overview tab for the exploration of patients similar to POI with respect to characteristics (Sec. 3.3 of the AM).

Patient similarity comparison tab. Of the four task-based questions pertaining to this tab, one was answered correctly by all participants (Sec. 4.1 Quest. 48 of the AM). The remaining tasks received correct responses from the majority of the participants (Sec. 4.1 Quest. 49-51 of the AM). Two of these tasks that received an incorrect response involved a change of the temporal axis (Sec. 4.1 Quest. 49 & 51 of the AM). This could be because the use of the temporal axis scroll and date modifications might require time for familiarization. However, this feature should be improved to enable changing and alignment of the temporal dimension with ease.

The usability evaluation showed that the patient similarity comparison tab requires minimum interaction for information retrieval (Sec. 4.2 Quest. 52 of the AM) and has a short learning curve (Sec. 4.2 Quest. 53 of the AM). Except for one indifferent respondent, most participants felt that the structure of the patient similarity comparison tab made it simple to find the required information (Sec. 4.2 Quest. 54 of the AM). All participants agreed that the patient similarity comparison tab allows clinical trajectory comparison of comparable patients to the POI (Sec. 4.2 Quest. 55 of the AM). Except for one participant, they agreed that the patient similarity comparison tab could help decision-making by comparing different aspects of the patient's clinical trajectory (Sec. 4.2 Quest. 56 of the AM). The one critical assessment arose from the current limitations of the system, as it lacks important information for decision-making due to limited data availability, such as molecular characteristics, blood indicators, genes, and more.

Generalizability of the System. Basically, our proposed system visualizes numerous patient information acquired for numerous tumor diseases. Variables specifically related to melanoma include sun exposure and progression rate. Therefore, our system could be adapted to other tumor diseases, such as breast cancer or prostate cancer, with little effort. Especially for the tumor board visualizations of POI, the participating dermatologists agreed that the proposed visualization is transferable to other tumor boards (Sec 2.2 Quest.21 of the AM). Patient similarity computation and comparative visualizations to explore characteristics of interest and support tumor boards in treatment recommendations can be incorporated.

7. Conclusion & Future Work

We proposed a visual analytics system to aid the decision-making process in tumor board meetings by presenting an overview of crucial clinical data and the progression of the patient. This addresses the main issue of tumor board meetings, where participants have to memorize the clinical information presented verbally during patient introductions. Our system is based on the dermatologist's MOI and customized for the tumor board meeting. In addition, to enable leveraging insights from existing data, a patient similarity view was incorporated that identifies patients most similar to the POI and provides appropriate visualizations for exploring and comparing their features. The qualitative study conducted with five dermatologists confirmed the clinical relevance of the proposed system and also gave insights into aspects of the system that can be enhanced.

The approach shows the value of patient similarity in tumor board decision-making, but the unavailability of ground truth data limits generalizability. The curation of ground truth data would facilitate the enhancement of the patient similarity computation. The present evaluation was lengthy, thus only key aspects of the similarity component were considered. We plan to conduct a comprehensive assessment of the patient similarity component in the future to further evaluate its impact and usability.

Therapy depends on the patient's lifestyle, physical condition, comorbidity, and other factors. Hematological markers like S100 and LDH help make decisions about tumors. Integrating all relevant data for decision-making into a unified system would improve its comprehensiveness and optimize workflow for physicians who would otherwise need to navigate multiple systems to obtain all relevant data. The addition of a worklist feature would allow physicians to curate a list of patients for tumor board discussion. Moreover, we are planning a comparative analysis of tumor board meetings and outcomes with and without our system to assess the impact of our system on clinical routine. Finally, we want to integrate more guidance and narrative techniques [GMF*21, MGS*22] to further support the data exploration by clinicians as well as to communicate risk factors and treatment outcomes to patients.

Acknowledgments

We thank all the clinicians of the University Dermatology Clinic Magdeburg, who participated in our user study.

References

[AGE*17] Amin M. B., Greene F. L., Edge S. B., Compton C. C., Gershenwald J. E., Brookland R. K., Meyer L., Gress D. M., Byrd D. R., Winchester D. P.: The eighth edition AJCC cancer staging manual: Continuing to build a bridge from a population-based to a more "personalized" approach to cancer staging. *CA Cancer J. Clin. 67*, 2 (2017), 93–99. 2, 3

[AM06] Aigner W., Miksch S.: Carevis: Integrated visualization of computerized protocols and temporal patient data. *Artif. Intell. Med. 37*, 3 (2006), 203–218. 2

[BAK07] Bui A. A. T., Aberle D. R., Kangarloo H.: TimeLine: Visualizing Integrated Patient Records. *IEEE Transactions on Information Technology in Biomedicine 11*, 4 (2007), 462–473. 2

[BCT22] Bartram L., Correll M., Tory M.: Untidy Data: The Unreasonable Effectiveness of Tables. *IEEE Trans. Vis. Comput. Graph. 28*, 1 (2022), 686–696. 8

[BOH11] Bostock M., Ogievetsky V., Heer J.: D³ data-driven documents. *IEEE Trans. Vis. Comput. Graph. 17*, 12 (2011), 2301–2309. 4

[CFBM69] Clark Wallace H. J., From L., Bernardino E. A., Mihm M. C.: The histogenesis and biologic behavior of primary human malignant melanomas of the skin1. *Cancer Research 29*, 3 (03 1969), 705–727. 4

[CN04] Chen L., Ng R.: On the marriage of lp-norms and edit distance. In *Proc. of Int. Conf. on Very Large Data Bases* (2004), pp. 792–803. 2

[CWS*17] Cypko M., Wojdziak J., Stoehr M., Kirchner B., Preim B., Dietz A., Lemke H. U., Oeltze-Jafra S.: Visual verification of cancer staging for therapy decision support. *Comput. Graph. Forum 36*, 3 (2017), 109–120. 2

[CZY*07] Chen S., Zhou S., Yin F.-F., Marks L. B., Das S. K.: Using patient data similarities to predict radiation pneumonitis via a self-organizing map. *Phys. Med. Biol. 53*, 1 (2007), 203–216. 2

[FJ10] Forsell C., Johansson J.: An heuristic set for evaluation in information visualization. *Proc. of Advanced Visual Interfaces* (2010). 9

[FNB*22] Floricel C., Nipu N., Biggs M., Wentzel A., Canahuate G., Van Dijk L., Mohamed A., Fuller C., Marai G.: THALIS: Human-Machine Analysis of Longitudinal Symptoms in Cancer Therapy. *IEEE Trans. Vis. Comput. Graph. 28*, 1 (2022), 151–161. 2

[GAK*11] Gschwandtner T., Aigner W., Kaiser K., Miksch S., Seyfang A.: CareCruiser: Exploring and visualizing plans, events, and effects interactively. In *Proc. of IEEE Pacific Visualization Symposium* (2011), pp. 43–50. 2

[GMF*21] Garrison L., Meuschke M., Fairman J., Smit N. N., Preim B., Bruckner S.: An Exploration of Practice and Preferences for the Visual Communication of Biomedical Processes. In *Proc. of EG Workshop on VCBM* (2021). 10

[GSC*05] Gandini S., Sera F., Cattaruzza M. S., Pasquini P., Zanetti R., Masini C., Boyle P., Melchi C. F.: Meta-analysis of risk factors for cutaneous melanoma: Iii. family history, actinic damage, and phenotypic factors. *European Journal of Cancer 41*, 14 (2005), 2040–2059. 4

[Hah14] Hahn U.: Similarity. *WIREs Cognitive Science 5*, 3 (2014), 271–280. 5

[HB03] Harrower M., Brewer C. A.: ColorBrewer.org: An Online Tool for Selecting Colour Schemes for Maps. *J. Cartogr. 40*, 1 (2003), 27–37. 5

[HBD*18] Howorko L., Boedianto J. M., Daniel B., et al.: The efficacy of stacked bar charts in supporting single-attribute and overall-attribute comparisons. *Vis. Inform. 2*, 3 (2018), 155–165. 8

[Inc15] Inc. P. T.: Collaborative data science, 2015. URL: https://plot.ly. 4

[KCG18] KRUPINSKI E. A., COMAS M., GALLEGO L. G.: A New Software Platform to Improve Multidisciplinary Tumor Board Workflows and User Satisfaction: A Pilot Study. *J. Pathol. Inform. 9* (2018). 2

[KNKMB12] KRISHNANKUTTY B., NAVEEN KUMAR B., MOODAHADU L., BELLARY S.: Data Management in Clinical Research: An overview. *Indian J. Pharmacol. 44*, 2 (2012), 168. 2

[LBI*12] LAM H., BERTINI E., ISENBERG P., PLAISANT C., CARPENDALE S.: Empirical studies in information visualization: Seven scenarios. *IEEE Trans. Vis. Comput. Graph. 18*, 9 (2012), 1520–1536. 9

[LCG*15] LI L., CHENG W.-Y., GLICKSBERG B. S., GOTTESMAN O., TAMLER R., CHEN R., BOTTINGER E. P., DUDLEY J. T.: Identification of type 2 diabetes subgroups through topological analysis of patient similarity. *Science Translational Medicine 7*, 311 (2015). 2

[LDC*22] LIU H., DAI H., CHEN J., XU J., TAO Y., LIN H.: Interactive similar patient retrieval for a visual summary of patient outcomes. *J. Vis. 26*, 3 (2022), 577–592. 2, 3, 5

[LEYK16] LABELLAPANSA A., EFENDI A., YULIANTI A., KADIR E. A.: Lambda value analysis on weighted minkowski distance model in cbr of schizophrenia type diagnosis. In *Proc. of Int. Conf on Information and Communication Technology* (2016), pp. 1–4. 2

[LMD15] LEE J., MASLOVE D. M., DUBIN J. A.: Personalized mortality prediction driven by electronic medical data and a patient similarity metric. *PLOS ONE 10*, 5 (2015). 2

[LMN*05] LIBRA M., MALAPONTE G., NAVOLANIC P. M., GANGEMI P., BEVELACQUA V., PROIETTI L., BRUNI B., STIVALA F., MAZZARINO M. C., TRAVALI S., MCCUBREY J. A.: Analysis of braf mutation in primary and metastatic melanoma. *Cell Cycle 4*, 10 (2005), 1382–1384. 4

[LWE*20] LUCIANI T., WENTZEL A., ELGOHARI B., ELHALAWANI H., MOHAMED A., CANAHUATE G., VOCK D., FULLER C., MARAI G.: A spatial neighborhood methodology for computing and analyzing lymph node carcinoma similarity in precision medicine. *J. Biomed. Inform. 112* (2020), 100067. 2

[MFP*22] MACCHIA G., FERRANDINA G., PATARNELLO S., AUTORINO R., MASCIOCCHI C., PISAPIA V., CALVANI C., IACOMINI C., CESARIO A., BOLDRINI L., ET AL.: Multidisciplinary tumor board smart virtual assistant in locally advanced cervical cancer: A proof of concept. *Front. Oncol. 11* (2022). 2

[MGS*22] MEUSCHKE M., GARRISON L. A., SMIT N. N., BACH B., MITTENENTZWEI S., WEISS V., BRUCKNER S., LAWONN K., PREIM B.: Narrative medical visualization to communicate disease data. *Computers & Graphics 107* (2022), 144–157. 10

[MMB*19] MARAI G. E., MA C., BURKS A. T., PELLOLIO F., CANAHUATE G., VOCK D. M., MOHAMED A. S. R., FULLER C. D.: Precision Risk Analysis of Cancer Therapy with Interactive Nomograms and Survival Plots. *IEEE Trans. Vis. Comput. Graph. 25*, 4 (2019), 1732–1745. 2

[MS20] M. SCHMIDT A. A. T.: *Learner and User Experience Research: An Introduction for the Field of Learning Design & Technology:*. Color Theory in Experience Design. EdTech Books, 2020. 5

[OCT*82] OKEN M. M., CREECH R. H., TORMEY D. C., HORTON J., DAVIS T. E., MCFADDEN E. T., CARBONE P. P.: Toxicity and response criteria of the eastern cooperative oncology group. *Am. J. Clin. Oncol. 5*, 6 (1982), 649–656. 2

[OGD*18] OESER A., GAEBEL J., DIETZ A., WIEGAND S., OELTZE-JAFRA S.: Information architecture for a patient-specific dashboard in head and neck tumor boards. *Int. J. Comput. Assist. Radiol. Surg. 13*, 8 (2018), 1283–1290. 2, 4

[PMR*96] PLAISANT C., MILASH B., ROSE A., WIDOFF S., SHNEIDERMAN B.: Lifelines: visualizing personal histories. In *Proc. of the SIGCHI Conf. on Human factors in Computing Systems* (1996), pp. 221–227. 2

[PMSB18] PARIMBELLI E., MARINI S., SACCHI L., BELLAZZI R.: Patient similarity for precision medicine: A systematic review. *J. Biomed. Inform. 83* (2018), 87–96. 2

[QR22] QUADRI G. J., ROSEN P.: A Survey of Perception-Based Visualization Studies by Task. *IEEE Trans. Vis. Comput. Graph. 28*, 12 (2022), 5026–5048. 7

[RWA*13] RIND A., WANG T. D., AIGNER W., MIKSCH S., WONGSUPHASAWAT K., PLAISANT C., SHNEIDERMAN B., ET AL.: Interactive information visualization to explore and query electronic health records. *Found. Trends Hum.-Comput. Interact. 5*, 3 (2013), 207–298. 2

[SED19] SAKET B., ENDERT A., DEMIRALP: Task-Based Effectiveness of Basic Visualizations. *IEEE Trans. Vis. Comput. Graph. 25*, 7 (2019), 2505–2512. 7

[SHB*20] STEINHAUER N., HÖRBRUGGER M., BRAUN A. D., TÜTING T., OELTZE-JAFRA S., MÜLLER J.: Comprehensive Visualization of Longitudinal Patient Data for the Dermatological Oncological Tumor Board. *Comput. Graph. Forum* (2020). 2, 9

[SLG*16] SUN L., LIU C., GUO C., XIONG H., XIE Y.: Data-driven automatic treatment regimen development and recommendation. In *Proc. of ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining* (2016), pp. 1865–1874. 2

[STA*22] SRABANTI S., TRAN M., ACHIM V., FULLER D., CANAHUATE G., MIRANDA F., MARAI G.: A Tale of Two Centers: Visual Exploration of Health Disparities in Cancer Care. In *Proc. of IEEE Pacific Visualization* (2022), pp. 101–110. 2

[TDR*22] TAMBORERO D., DIENSTMANN R., RACHID M. H., BOEKEL J., LOPEZ-FERNANDEZ A., JONSSON M., RAZZAK A., BRAÑA I., DE PETRIS L., YACHNIN J., ET AL.: The molecular tumor board portal supports clinical decisions and automated reporting for precision oncology. *Nature Cancer 3*, 2 (2022), 251–261. 2

[TWW18] TASHKANDI A., WIESE I., WIESE L.: Efficient in-database patient similarity analysis for personalized medical decision support systems. *Big Data Res. 13* (2018), 52–64. 2

[VMP*18] VITALI F., MARINI S., PALA D., DEMARTINI A., MONTOLI S., ZAMBELLI A., BELLAZZI R.: Patient similarity by joint matrix tri factorization to identify subgroups in acute myeloid leukemia. *JAMIA Open 1*, 1 (2018), 75–86. 2

[WG12] WONGSUPHASAWAT K., GOTZ D.: Exploring flow, factors, and outcomes of temporal event sequences with the outflow visualization. *IEEE Trans. Vis. Comput. Graph. 18*, 12 (2012), 2659–2668. 2

[WHL*19] WANG N., HUANG Y., LIU H., FEI X., WEI L., ZHAO X., CHEN H.: Measurement and application of patient similarity in personalized predictive modeling based on electronic medical records. *Biomed. Eng. Online 18*, 1 (2019). 3, 5

[WKE*22] WILLIAMS C. B., KARAGOZ K., ELSEY R., MEISSNER T., HIGASHI M., SCHADT E., JUN T., WANG X., CHEN R., ZHOU X., GUIN S., LEE J. H., OH W. K.: Qopi clinical informatics: A digital platform to enable real-time quality reporting, clinical decision support, and rapid learning. *J. Clin. Oncol. 40*, 16 (2022). 2

[WM97] WILSON D. R., MARTINEZ T. R.: Improved heterogeneous distance functions. *J. Artif. Intell. Res. 6* (1997), 1–34. 5

[WPS*09] WANG T. D., PLAISANT C., SHNEIDERMAN B., SPRING N., ROSEMAN D., MARCHAND G., MUKHERJEE V., SMITH M.: Temporal summaries: Supporting temporal categorical searching, aggregation, and comparison. *IEEE Trans. Vis. Comput. Graph. 15*, 6 (2009), 1049–1056. 2

[WS15] WANG F., SUN J.: Psf: A unified patient similarity evaluation framework through metric learning with weak supervision. *J. Biomed. Health Inform. 19*, 3 (2015), 1053–1060. 2

[WWPS10] WANG T. D., WONGSUPHASAWAT K., PLAISANT C., SHNEIDERMAN B.: Visual information seeking in multiple electronic health records. In *Proc. of International Health Informatics Symposium* (2010). 2

[YJF98] YI B.-K., JAGADISH H. V., FALOUTSOS C.: Efficient retrieval of similar time sequences under time warping. In *Proc. of Int. Conf. on Data Engineering* (1998), IEEE, pp. 201–208. 2

[You02] YOUNG R. R.: Recommended requirements gathering practices. *CrossTalk 15*, 4 (2002), 9–12. 3