Visual Analytics for Multivariate Sorting of Sport Event Data

D. H. S. Chung, P. A. Legg, M. L. Parry, I. W. Griffiths, R. Bown, R. S. Laramee, and M. Chen

Abstract— A critical job coaches and sport analysts are tasked with is the planning of key match videos for analytical coaching sessions. Each session may focus on a diverse range of topics, such as the strengths and weaknesses of a game. This needs to be tailored further based on a player's tactical position or skill. Hence, the criteria for sorting video is dynamic. This motivates a sorting criteria beyond individual attributes of a multi-dimensional data set. We propose a knowledge-assisted, event ranking framework to interactively model implicit sorting as formal parameters that can be used to perform multivariate sorting. We incorporate knowledge in the form of a user's event ranking which we formalize using regression analysis. Depending on the ranking criteria, the resulting function can be customized to many forms such as importance, or other performance metrics. We use visual analytics to depict the set of sortable attributes and weights determined by the model. Visual feedback helps the user comprehend the function, and aids in choosing the most appropriate model. We find that this approach significantly increases the usability of multivariate sorting and allows domain experts to incorporate their knowledge and expertise into the analysis. This work is undertaken in conjunction with a national rugby team. To demonstrate the effectiveness of our sorting system, we present a use case scenario in rugby event analysis, where coaches and analysts need to re-organize match videos in order to study and evaluate team and player performance.

1 Introduction

Event sorting is a fundamental task in visual analytics. This task becomes challenging when sorting involves several data dimensions, and the way in which each dimension influences the sorting is not well defined. Such a sorting task is commonplace in practical visual analytics, where one often encounters ad hoc request for organizing data in to some kind of order without precise specification of the relevant *sort keys* and a *sorting function*. Although some analytical methods such as multidimensional scaling (MDS) [6] or principle component analysis (PCA) [16] may help in some applications (e.g., [14]), they focus on the discovery of the most influential attributes in the data, rather than the discovery of a *sorting function* for an ad hoc requirement of a sorting task. This work addresses this challenge in the context of sports event analysis by using a knowledge-assisted visual analytics process.

We notice that when given an ad hoc requirement of event organization, a user normally knows how the sorting outcome should look like, without knowing explicitly about the sort keys and a sorting function with a visualization. In a knowledge framework [4], we can summarize the situations as follows:

- Users have *tacit* knowledge about sorting a set of events, but do not have the *formal* knowledge as to a sorting function. They may have partial knowledge about sort keys as they typically speculate a set of attributes that may influence the sorting.
- Although users can organize a given set of events in an 'accurate' manner using their tacit knowledge (because they define the expected sorting outcome), this does not scale up to a large number of events. It is generally easy for users to place a few most representative events (e.g., success, neutral, failure) into order. The task becomes inefficient when the number of events increases significantly, and ineffective (i.e., less 'accurate') for events with a similar principle criterion (e.g., how successful), but a diverse set of conditions (left or right, earlier or later, different players involved, etc.).
- On the other hand, the system does not have any a priori knowledge about the expected sorting outcome, since the sorting requirement is not predefined. Of course, it does not have the formal knowledge about a sorting function either. If the system has a sorting function, it can perform event sorting in a more scalable and consistent manner.

We thereby developed a visual analytics system that enables users to provide the system with some of their *tacit knowledge* by selecting a small set of events (typically 3-7), and placing them in an order as an example for the system. The users may also provide their *partial knowledge* about possible attributes (e.g., data dimensions) that should be considered. This partial knowledge is not essential, but can reduce the amount of computation significantly. The system uses an analytical method to convert the *tacit* and *partial knowledge* to some *formal knowledge* in the form of a potential sorting function and a measure of influence of different sort keys. The system then provide users with a visualization of the sorted results in relation to the potential sorting function and the weights of different sort keys. The former is shown in a glyph-based sorting canvas, and the latter in a parallel coordinates plot. Users can interactively refine the sorting results and the weights of different sort keys, or re-activate the knowledge discovery process by refining their initial specification of the example set or the speculated data dimensions. Satisfactory results can normally be obtained within a few iterations, and users can produce a sorted set of events (i.e., video clips) for supporting further analytical tasks such as compiling various statistical indicators in relation to the sorted events, and analyzing video clips in a coaching session. Our contributions are:

- We introduce a novel visual analytic approach to sorting multidimensional events by converting users' tacit and partial knowledge to formal knowledge.
- We develop a system that supports such a process iteratively through a close integration of interaction, analysis and visualization.
- We demonstrate the efficiency and effectiveness of visual analytics for multivariate sorting through a real-world application, and we evaluate our work objectively with a user consultation.

The remainder of the paper is organized as follows: In Section 2, we provide a brief overview of related work. Section 3 gives some background on rugby and outlines our motivation. The pipeline of the visual analytic system is described in Section 4. In Section 5

- *D. H. S. Chung, P. A. Legg and M. L. Parry are with the Department of Computer Science and the College of Engineering, Swansea University, UK.*
- *I. W. Griffiths, is with the College of Engineering, Swansea University, UK.*
- *R. Bown, is the head performance analyst at the Welsh Rugby Union.*
- *R. S. Laramee, is with the Department of Computer Science, Swansea University.*
- *M. Chen, is with Oxford e-Research Centre, Oxford University, UK.*

Manuscript received 31 March 2013; accepted 1 August 2013; posted online 13 October 2013; mailed on 27 September 2013.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

we describe the method for converting tacit and partial knowledge to an explicit sorting function in our knowledge-assisted, event ranking framework. Section 6 details the visual mappings and interaction of our visual analysis system. We present the system in Section 7. In Section 8, we evaluate and discuss the limitations of our work, and draw our concluding remarks in Section 9.

2 RELATED WORK

Sorting is the computational process of rearranging a sequence of items in order or categorizing entities with similar properties [18]. Our sorting concept closely follows the method of card sorting [26], a user-centered design that allows a user to decide how to categorize a set of items into groups or structures they are familiar with. This approach has been effectively used for classifying symbols in cartography [25], organizing online course sites [7], and clustering multivariate glyphs [2, 17], where the glyphs are treated as cards metaphorically. The fundamental difference here is that we define a sorting function that expresses different levels of knowledge to order events with an implicit sorting criteria.

The incorporation of knowledge in visual analytics is still an early research topic. Wang *et al.* [29] propose a framework to support communication between both domain and individual knowledge structures. Lipford *et al.* [20] explore the use of visualization to help recall a user's reasoning. Mistelbauer *et al.* [21] introduce Smart Super Views for analysis of different data sets in medical visualization. More work has been studied on using knowledge to tune parameters in statistical models. Heimerl *et al.* [10] present a user-study comparing three methods for training an interactive classifier for use in text search and filtering. Hoferlin *et al.* [11] describe the use of interactive learning in visual analytics to allow users to adjust and understand complex classifier models through visualization. Other approaches include using a distance function [3], semantic interaction [8] that captures the analytical reasoning of users, and multiple views for cross-filtering analysis [31]. To the best of our knowledge, this is the first work of its kind introducing visual analytics to sort events by converting a user's tacit and partial knowledge to formal parameters.

Sports Visualization is an area that is gaining interest within the community. Parry *et al.* [23] propose a framework for hierarchical event selection for video storyboard visualization. The authors demonstrate their work in snooker. Jin and Banks make use of treemaps for visualizing scoring results and match statistics in tennis [15]. Moore *et al.* [22] look at the potential of using visualization techniques for spatial temporal analysis of rugby data. Legg *et al.* [19] conducted a design study to show the effective use of glyph-based visualization within sports performance analysis. Pileggi *et al.* [24] introduce SnapShot: a system that enhances hockey analytics through visualization.

3 UNDERSTANDING THE PROBLEM

In modern sports, especially in high-level teams, coaches and analysts are experiencing a deluge of data due to the paid induction of various digital technologies for supporting match analysis and training. This work was carried out with the Wales National Rugby team, which uses videos extensively for analyzing performance indicators. The analysts in the team are often asked to organize a set of events (i.e., video clips) into an order based on some ad hoc specification, such as how successful a strategy is in some conditions. It is not difficult to observe that there would be numerous variations of requirements, and it is not feasible for a system to predetermine a set of sorting functions that could address the majority of the requirements. Hence, the system has to provide support in a flexible manner by enabling analysts to discover a set of *sort keys* and a *sorting function* efficiently and effectively to meet an ad hoc requirement of event analysis and organization. To fully consider the challenges involved with rugby event analysis and how a visual analytic system for sorting can be of significant advantage, we provide a background to the game in this section.

3.1 Rugby Union

Rugby Union is a popular team sport which consists of two teams (of 15 players) who advance an oval ball across a rectangular field (up to 144m long by 70m wide) with two H-shaped goal posts at either end. The game is played primarily by carrying the oval ball from one end of the pitch to the other. Points can be scored in several ways: A *try*, which involves grounding the ball in the opposition goal area, or through kicking the ball between the H-Shaped post from a *conversion*, *penalty kick* or *drop goal*. Each match is played in two 40-minute halves, where the objective is to score more points than the opponent.

3.2 Rugby Event Analysis

Analysts and sporting coaches heavily rely on using notational data [12] for player and team analysis. Notational analysis consists of "*tagging*" video footage with key events and semantic notations from which key performance indicators can be derived. One resulting output is a set of video clips that capture moments in a game of when a team receives and loses/changes possession. In rugby, such events are known as a *phase ball event*, and involves smaller *phases* that describe the period of play. A single match consists of a collection of phase ball events where each event may lead to a scoring outcome. These events are often enriched further with additional data sources (e.g., nutrition data, player data, and ball tracking) that increases the complexity of the analysis. The event descriptors (or attributes) we use are:

- **start event** — the type of event in which play is started.
- **gain** — the distance gained towards the goal area.
- **territory start position** — the starting position of ball when a team receives possession in relation to the the goal area.
- **time** — the starting time of the event.
- **tortuosity** — the tortuosity of the ball path.
- **number of phases** — a count of the phases.

When sorting events to meet some specific requirement, for example, by importance, selecting and finding the events is challenging. On average, a match typically contains around 104 events (52 per team). Existing notational analysis systems (e.g., SportsCode) does allow selecting clips easily, but it is time consuming to select clips according to some criteria (e.g., by time) and combining them together into a form in which the events can be compared. We find this approach does not scale very well when organizing clips from multiple matches.

Analysts, players, and even fans can often judge or rank a small set of events relatively well based on intuition and experience of the game without specific knowledge on a sorting function. The discovery of such a sorting function would enable users to order a larger set of video clips in a more effective and efficient manner. Our work aims to address this problem by modelling a user's tacit ordering requirement based on the underlying data attributes.

3.3 Tacit knowledge vs Formal knowledge

The incorporation of human knowledge in visual analytics describes the process of transferring *knowledge* into some explicit form (e.g., a function) [29]. For the purpose of this paper, we define the following:

Tacit knowledge. User's input of event ranking. The know-how cannot easily or explicitly be transferred to another user.

Partial knowledge. User's input of possible attributes that may affect the ranking. This narrows the search space.

Knowledge discovery. Using regression analysis for determining sortable attributes, how the attributes are combined into a ranking function, and evaluating the accuracy of the function.

Formal knowledge. The function (including the weights on various attributes) is formal knowledge as it can be explicitly written down and transferred to others.

Knowledge externalization. Visualizing the sortable attributes and their weights and their impact on sorting various events.

Knowledge application. Sorting video clips to be watched in an analytical coaching session.

In the next section, we use the above definitions to describe the pipeline to our visual analytic framework.

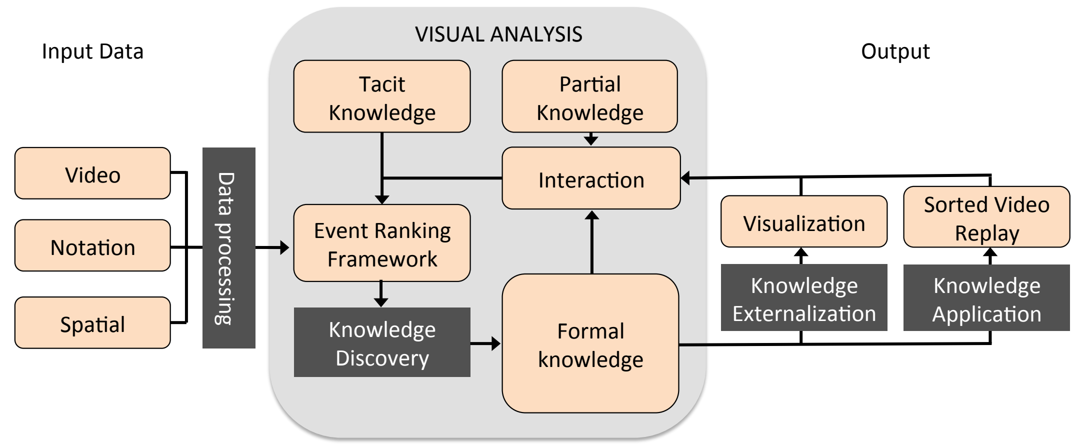

Fig. 1. A graphical pipeline illustrating our visual analytic system for multivariate sorting of rugby event data. It consists of four steps: processing the input rugby event data, a knowledge discovery process to derive a set of sortable attributes combined into a function (formal knowledge), knowledge externalization to determine the function's impact on sorting various events in a visual form, and finally, using formal knowledge to interactively sort and replay match videos (knowledge application).

4 Overview

Fig. 1 shows the process to our visual analytic sorting system. It consists of four key steps: *processing* of rugby event data, *knowledge discovery*, *knowledge externalization* and *knowledge application*.

The first step involves processing and integrating each data source into an event data structure in order to extract the underlying sorting attributes. We combine three types of data, namely: video, notation and spatial ball tracking data. Rugby event data is a mixed data set containing both quantitative attributes (e.g., gain, time, ball tortuosity) and categorical attributes (e.g., start event). In addition, these events may be associated with other types of information such as a video clip.

In the next step (detailed in Section 5), we use a knowledge-assisted, event ranking framework to determine a set of sort keys and a ranking function that expresses the user's ordering requirement. Tacit knowledge is stored in the form of the user's implicit ranking of events. We use regression analysis to develop a model that represents the tacit ordering as formal parameters. As the ordering specification may be based on prior-knowledge, the ranking function can be customized to many forms such as importance, or other types of performance metrics. We refer to this as *knowledge discovery*. The goal is to formally describe the user's ranking criteria as a function which can be explicitly written and transferred to other users. Optionally, the user may refine the model using partial knowledge by adjusting weights on various sort keys that affects their ranking.

In the third step, we visualize the set of sortable attributes and their weights in order to provide visual feedback to the user. Since there are many possible types of analytical algorithms, we use visual analysis as a method for choosing and optimizing the model (see Section 5.2). The visualization which is to be detailed in Section 6.1, enables analysts to understand how the data is sorted and its impact on sorting various events. We refer to this process as *knowledge externalization*. This allows the domain expert to incorporate their knowledge and experience into the analysis.

In the final step, we demonstrate an example of *knowledge application* through sorting of video events. We focus on two aspects: 1) using glyph-based visualization as an interface to interactively select events that need to be sorted, and 2) the ability to playback match videos based on the sorting criteria (see Section 6.3).

5 Knowledge-Assisted Event Ranking

The knowledge-assisted, event ranking framework involves defining a relationship between the user's sorting outcome and the set of sort keys (i.e., data dimensions). Let $e_1, e_2, \ldots, e_n$ be events and $e_{i,j}$ be its j-th attribute value. We can model the user's tacit ordering as $\mathbf{Y} = \mathbf{E}\beta$, where $\mathbf{E}$ is an $n \times m$ matrix, and $\beta_j \in \mathbb{R}$ are the weights or importance of each sort key. The goal is to estimate the weights β such that the

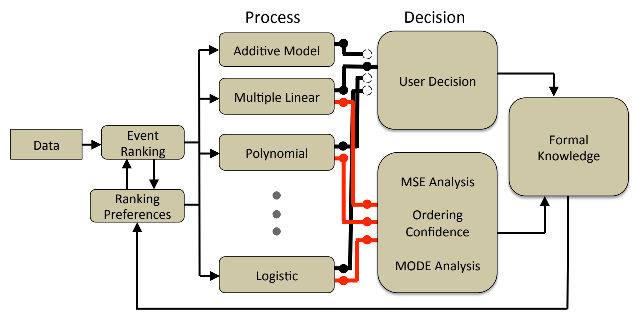

Fig. 2. A knowledge-assisted event ranking framework. The model takes a subset of events placed in order of the sorting outcome as input. We use regression analysis to discover a ranking function that represent the user's order of events. Visual feedback is provided to validate the model when the model is not sufficiently trained (black path), else automatically evaluated using model comparison metrics (red path).

user's ranking of events y_i is preserved. Typically, a user may speculate these weights during the ordering process. Since the requirement can be dynamic, it is often difficult for a user to quantify precisely the influence each sort key has on their ordering. This becomes significantly more difficult when the complex parameter space is large.

One effective approach for predicting such weights and a set of potential sort keys is through multiple regression analysis. Regression analysis is a statistical process for estimating the relationship between a dependant variable and one or more independent variables [1]. In this work, we deploy three main modelling techniques: multiple linear regression, polynomial regression, and logistic regression. Fig. 2 provides an overview of this process. The system relies on a user making an initial specification by choosing a subset of events $n \leq N$, and placing them in order $e_{s_1} < e_{s_2} < \ldots < e_{s_n}$ as training input for the model. We then approximate the sort key weights using a least squares fitting [1] which generalizes to:

$$\hat{\beta} = (\mathbf{E}^\mathbf{T}\mathbf{E})^{-1}\mathbf{E}^\mathbf{T}\mathbf{Y} \quad (1)$$

It is clear in Eq. 1 that a solution to β exists as long as the matrix $\mathbf{E}^\mathbf{T}\mathbf{E}$ is invertible. Hence, any constant attribute must be removed from the model. For example, this may occur if a user chooses to rank a set of events with similar ordinal or categorical values (e.g., ordering a set of *scrum* events based on successfulness).

The least square solution typically relies on an *over-determined* system of equations $\mathbf{E}$ (i.e., for $n > m$). Conversely, $\mathbf{E}$ is *under-determined* if the user does not specify enough events for the model to learn. Generally, such a system may have infinitely many or no solutions. We can pick one of these solutions by finding the smallest one such that

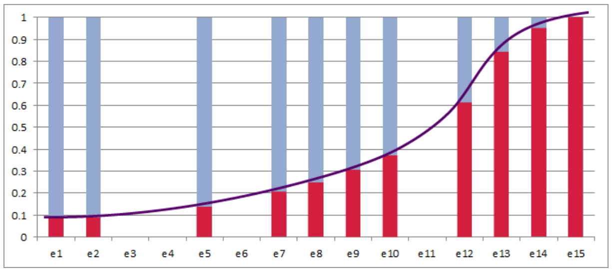

Fig. 3. Applying a Gaussian function (purple curve) as a weighting parameter to the predicted event ordering confidence τ. The blue values depict events that are ordered successfully. The adjusted values are shown in red. This emphasizes the importance of ordering higher ranked events, and de-emphasizes the ordering of lower ranked events.

$\hat{\beta}$ is minimized subject to the constraint $\mathbf{Y} = \mathbf{E}\beta$. This is solved using the method of Lagrange multipliers:

$$\hat{\beta} = \mathbf{E}^{\mathbf{T}}(\mathbf{E}\mathbf{E}^{\mathbf{T}})^{-1}\mathbf{Y} \tag{2}$$

Depending on the system in which the sort key weights β are estimated, we propose two approaches to validate the accuracy of the model's prediction as shown in Fig. 2. The under-determined case (black path) is evaluated based on visual feedback of the model. This approach has been effectively shown in other visual systems [11, 20] for validating model parameters. User's can refine the model parameters further, and re-assess whether the derived formal knowledge is representative of their ordering requirement (see Section 5.3). In the second approach (red path), we apply a series of comparison metrics to automatically select the most appropriate model which we will detail in the next section.

5.1 Regression Evaluation

As part of the knowledge discovery step, we adopt a set of comparison metrics to validate the quality of the formal sorting function derived by the model. For the purpose of this section, we differentiate between the event ranking *value* (i.e., the value in which the model estimates) and the predicted *rank* outcome (e.g., 1st, 2nd, 3rd, etc.) of an event. Many different criterions have been proposed (e.g., RMSE, RSE and MAE) in statistical modelling. We utilize the Mean Squared Error (MSE) [1, 27] as one metric for validating the model. MSE is the most common comparison test which describes the difference between the predicted event ranking value and the actual value determined by the user. This is described by the following:

$$MSE = \frac{1}{n - dof - 1}\sum_{i=1}^{n}(\hat{y}_i - y_i)^2 \tag{3}$$

where n is the number of events and dof is the degrees of freedom. In order to compare different regression models by MSE, the unit of error must be the same. This is important in predictive analysis when comparing bounded (e.g., logistic regression) versus unbounded models (e.g., linear regression). We address this by scaling the user's expected event ranking values y_i to the co-domain of the bounded function. It is easy to observe that the sorting outcome is maintained when MSE = 0. However, this may not be the case with MSE > 0, since the criterion does not describe how the predicted rank measures with respect to the user's ordering. Hence, to accurately determine the best model we introduce two ranking comparison metrics: an ordering confidence τ, and a Mean Ordering Distance Error (MODE).

Since the goal is to preserve the ordering $e_{s_1} < e_{s_2} < \ldots < e_{s_n}$, we introduce a confidence metric τ that measures the accuracy of the predicted order with respect to the user's actual order. Let $f : E \mapsto \mathbb{R}$ be the derived sorting function and $\phi : \mathbb{R}^2 \mapsto \{0,1\}$ be a binary function that returns 1 if the predicted order $f(e_{s_i}) < f(e_{s_{i+1}})$ for all $i = 1, \ldots, n-1$ is maintained. We derive the ordering confidence as:

Algorithm deriveOptimalRegressionModel(*RegressionModels reg*[])

Model m_1, m_2
$getModelByMinMSE(reg, m_1, m_2)$
$\delta_{MSE} = \frac{\min(MSE(m_1), MSE(m_2))}{\max(MSE(m_1), MSE(m_2))}$
$\delta_{\tau} = \frac{\min(\tau(m_1), \tau(m_2))}{\max(\tau(m_1), \tau(m_2))}$
$\delta_{MODE} = \frac{\min((MODE(m_1), MODE(m_2))}{\max(MODE(m_1), MODE(m_2))}$
if ($\delta_{MSE} < T_{MSE}$) **then**
    return $\min_{MSE}(m_1, m_2)$
else
    if ($\delta_{\tau} < T_{\tau}$) **then**
        return $\max_{\tau}(m_1, m_2)$
    end if
else
    if ($\delta_{MODE} < T_{MODE}$) **then**
        return $\min_{MODE}(m_1, m_2)$
    end if
end if

Fig. 4. Algorithm for automatically choosing the optimal model.

$$\tau = \frac{1}{n-1}\sum_{i=1}^{n-1}\phi(f(e_i), f(e_{i+1})) \tag{4}$$

In many event summarization tasks such as video storyboarding, we find that the importance of selecting higher level events will have a significant influence on the selection and ordering of lower level events [23]. Hence, the organization of higher ranked events (e.g., e_{n-2}, e_{n-1} and e_n) can often be established more easily and with greater confidence by the user. These events create a benchmark criteria in which a user can compare and rank subsequent events. Therefore, the predicted accuracy of such events are given a higher weighting in the model validation. We incorporate this by modulating the ordering confidence using a Gaussian function $G(x)$ where $x = (n-1) - i$. The parameter σ in $G(x)$ is pre-defined, and we set $\sigma = 2$ as default in our system. Fig. 3 illustrates this process. We can see that the error in the predicted order of lower ranked events e_3, e_4, e_6 and e_{11} have significantly less impact to the overall confidence of the model after moderation. As a result, the model would now be considered as a viable option to the user.

The third comparison metric we use is Mean Ordering Distance Error (MODE) which is the average difference between an event's actual rank and its predicted rank. Let i be an event's actual rank defined by the user and t_i be its predicted rank. We define MODE as:

$$MODE = \frac{1}{n}\sum_{i=1}^{n}||i - t_i|| \tag{5}$$

such that $||i - t_i|| = 0$ when an order is preserved. The MODE describes the similarity between the order of events predicted by the model, and the user's actual order.

5.2 Model Selection

Our visual analysis system incorporates three regression techniques to model a user's tacit knowledge. Each technique may discover a different set of sort keys and weights that potentially influences their ordering of events. Fig. 5 illustrates one example where each model has predicted a different set of sort key weights as shown by the colored parallel axes (see Section 6.1 for details). When sufficient training data (i.e., ordering of events) is learned by the system, we automatically choose the optimal model by evaluating the comparison metrics described in Section 5.1. Let $m_1, m_2 \in F = \{F_{linear}, F_{polynomial}, F_{logistic}\}$ be two regression models with least MSE. We determine the optimal model using the algorithm outlined in Fig. 4. The quality of the model is computed based on a series of if-conditions that compares the ratio $\delta_{MSE}, \delta_{\tau}$ and δ_{MODE} against the pre-defined threshold values $T_{MSE}, T_{\tau}, T_{MODE} \in [0,1]$. By default, we set a threshold of $T = 0.7$

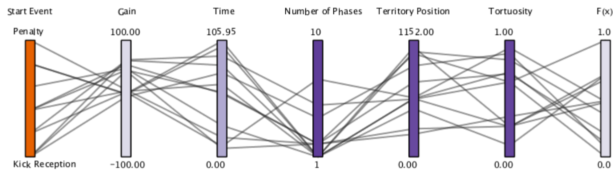

(a) Linear Regression

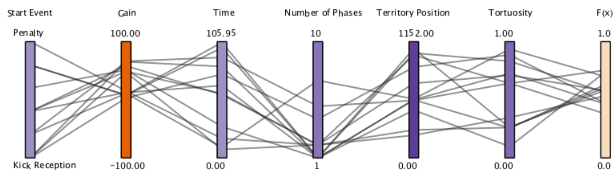

(b) Polynomial Regression

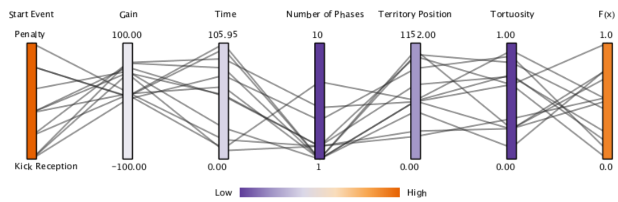

(c) Logistic Regression

Fig. 5. Visual comparison of the predicted sort key weights using (a) Linear, (b) Polynomial and (c) Logistic regression in parallel coordinates. The resulting sort key weights $||w_j\beta_j||$ of the model $F(X)$ is mapped to color respectively. In this example, $w_j = 1$. The sort function axis $F(X)$ is color mapped based on its event ordering confidence τ.

for each metric to produce the work in this paper. However, this can be customized according to the user's preference. The selected model will be a sort function that is most representative of the conditions specified by the user's ordering requirement.

5.3 User Interaction

When organizing a set of events to meet some requirement, users can often make intuitive or educated guesses on specific sort keys that may or may not affect their ranking criteria. We liken this to *partial knowledge*. To facilitate this in our visual analytic system, we allow the user to interactively tune the model parameters by applying additional weightings $w_j \in [0,1]$ to the sort key weights β such that:

$$y_i = F(\mathbf{w}, \beta, e_i) \tag{6}$$

The user can refine the sorting results effectively by de-emphasizing specific attribute axes and analyze new sorting strategies and how they impact the predicted order of events. Optionally, users can choose to remove a sort key completely ($w_j = 0$). We incorporate this function into our system as a series of interactive sliders which the user can adjust (see Fig. 8 for example).

6 Visual Mapping and Interaction

We have described the process of converting tacit and partial knowledge into formal parameters of a ranking function. To illustrate the associated model (i.e., knowledge externalization), we focus on three important aspects: 1) visualization of sort keys and their various weights, 2) informing the quality of formal knowledge through visualization, and 3) the interaction and visualization of sorted rugby events.

6.1 Visual Mapping

In order to visually convey the model parameters of the ranking function, we adopt the use of parallel coordinates which is proven to be effective in multivariate analysis [13]. Each attribute dimension is plotted as vertical axes and the events are drawn as polylines. To illustrate the contribution of each attribute, we color map the axes according to the magnitude of the weighted model parameters $||w_j\beta_j||$ for $j > 0$ (see Fig. 5). We use a diverging color scheme chosen from

Fig. 6. Screenshot of the video playback of sorted events. We integrate four different broadcasting feeds that correspond to the event.

Color Brewer [9] to emphasize the attributes that are least and most influential. The ranking function $F(X)$ is plotted as an additional axis and color mapped to depict the ordering confidence τ from the model. This enables the user to assess the quality of the formal knowledge in a visual manner. We find that parallel coordinates offers a holistic externalization of the formal knowledge which allows the user to comprehend model parameters and its impact on sorting events. This supports the user in making an informed decision on choosing the most appropriate model through visual feedback. For the example shown in Fig. 5, a logistic model would be the preferred option since $F(X)$ is colored orange.

6.2 Interaction and Glyph-based Visualization

Glyph-based visualization is an effective tool for representing multivariate data [30]. Glyphs are graphical entities that convey one or more data values using visual features such as size, shape and color. We take advantage of the recent work by Legg *et al.* [19], who demonstrate the usability of glyphs in rugby. We position the glyphs along two primary axes (see Fig. 7(c)). Although interactive multivariate sorting is the focus of this work, we are careful not to confuse the end-user with an unfamiliar visual design. To facilitate this, we adopt their glyph [19] to encode our event properties (see Section 3.2) as shown in Fig. 7(d). The glyphs highlighted in purple within the glyph-based canvas indicate events that resulted to a point scored. Other visual design choices (e.g., Chernoff Faces [5] and Star Glyphs [28]) may be used depending on its application context. Due to the inherent occlusion of using large glyphs [30], we provide interactive sliders that enables the user to adjust the length of the sorting axes. This can significantly reduce the amount of visual clutter. The user can then select the events (i.e., the glyphs) and import them into a ranking table. The table view provides an interface where the user can specify the event ordering by drag-and-drop. We found glyphs to be an intuitive mechanism for selecting and ranking events. This is due to similarity to our card metaphor.

6.3 Sorted Event Replay

Sporting analysts often rely on making semantic observations that can only be gained through studying video in order to determine the importance of an event and its event ordering. To support the transfer of tacit knowledge, our visual analysis system facilitates the inspection of key events by brushing the sorted results within the parallel coordinate or glyph-based view. Since the data are associated with single or multiple video clips, we incorporate a video playback user-option for viewing the sorted events (see Fig. 6). The playback of ordered video clips enable users to choose, view, and rank the events in a much more effective manner than the results of a typical search query.

7 Visual Analysis System for Sorting

Fig. 7 presents our visual analysis system analyzing a rugby match. The system contains four main views: (a) the parallel coordinate view for depicting the ranking function based on the the example ordering specification shown in table view (b). The table interface allows the user to configure, or modify their event ranking. Subsequently, we

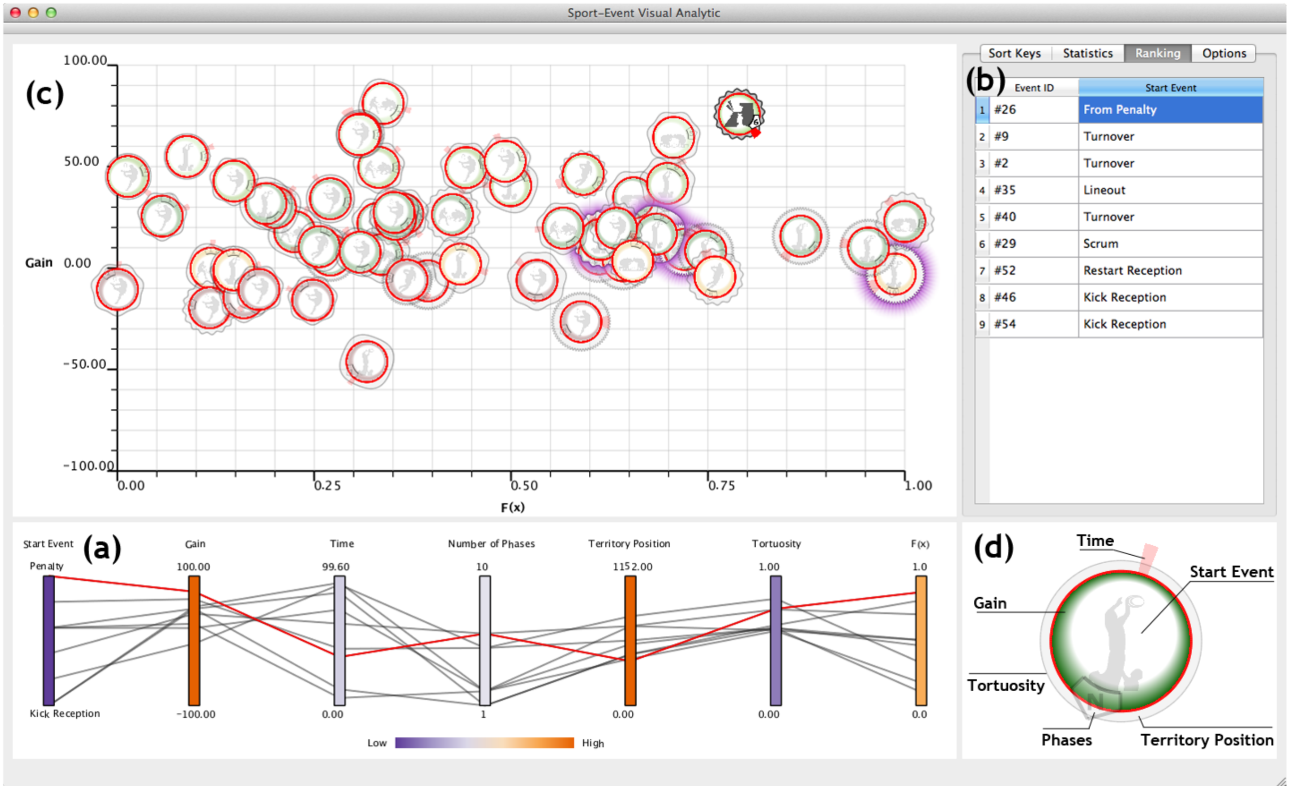

Fig. 7. A visual analysis sorting system used for multivariate sorting of rugby event data. It contains four main views: (a) is the parallel coordinate view of the ranking function. This allows the user to see the composition of the ranking function, and the accuracy in which the ordering of the event subset shown in table (b) is maintained. The user can adjust or refine the event ranking within the table view. (c) displays the sorted results using glyph-based visualization and (d) is an glyph-based interface for selecting the primary axes in (c).

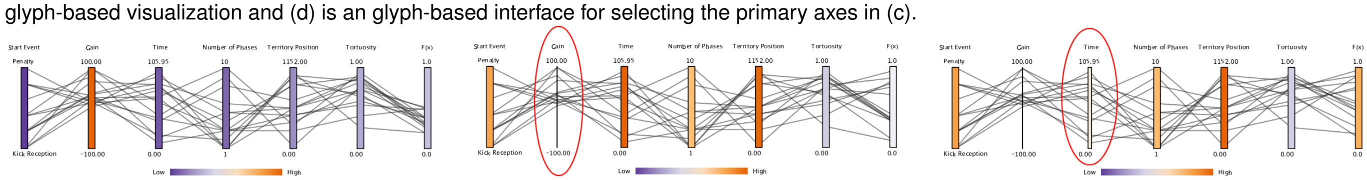

Fig. 8. Refining the knowledge discovery process (from left to right) using partial knowledge. The user's partial knowledge w_j is mapped to axis width. The initial model (left) starts with a low ordering confidence indicated by the blue $F(X)$ axis. After several iterations of refinement (centre and right), one solution is obtained that accurately preserves the user's ordering since $F(X)$ is colored orange. The adjusted partial knowledge (highlighted with a red circle) correspond to the values (centre) $w_{gain} = 0$ and (right) $w_{time} = 0.50$ respectively.

update the model according to the user's ranking and preference. The resulting model is then conveyed by the parallel coordinate view. In (c), we depict the sorted events using glyphs. We control the primary sorting axes within this view by clicking on the corresponding glyph component in the graphical interface (d). In order to sort the glyphs using formal knowledge, a drop down option menu is available for users to export the current ranking function to one of the X or Y axes. For this example, the X axis corresponds to the ranking function $F(X)$. This allows the user to explore how events in the match are reorganized according to the their derived sorting criteria.

There are many strategies towards using the system. The most generalized approach starts by choosing the glyphs (or events) of interest in the glyph-based visualization which the user wishes to order. Different layout methods can be used by changing the primary axes to enhance the searchability of such glyphs. These are imported into the ranking table view where the user can specify the event rank in a top-to-bottom (best-to-worst) order. This defines the event ranking which is parsed as a parameter to our knowledge-assisted, event ranking framework. The user can then visually assess the quality of the ranking function in the parallel coordinate view. In some cases, this may require several iterations refining the attribute weights (see Fig. 8) in order to achieve an optimal solution.

Each of the views are linked such that the user can interactively explore and brush the data in an intuitive manner. Selecting events in the table view will highlight the polyline red within the parallel coordinates, and render the corresponding glyph in focus (see Fig. 7). This enables the user to closely inspect and understand the relationship between the events, the event ordering, and the derived sorting function. Similarly, we provide brushing in both the parallel coordinate view and glyph-based visualization. The focus+context interaction offers a convenient method for analysts to choose and play back the sorted video clips associated with each event for further analysis.

8 Evaluation

To evaluate the work of this paper, we organized a consultation session consisting of 5 participants (3 computer scientists and 2 sport scientists). Each participant had reasonable knowledge of both rugby and visualization.

Method. The question we propose is, *"What are the most important positive outcomes?"*. In this example, importance is the tacit knowledge we are trying to formalize, and positive outcome are the measurable parameters. A positive outcome in rugby is considered when a team gains a tactical advantage over their opponent. For this study, such events are indicated when a team scores, and when a team is rewarded a penalty or free kick. The same rugby match was used throughout the study. During each session, the users were presented with several tasks outlined in Fig. 9. With each task, we provide the

Task	(optional) meta-answer	Result											
1. Identify and rank 5 events from best-to-worst	(a)	3		1					5	4	2		
	(a)		4	5		2		3	1				
	(b)		2					4	3	1		5	
	(a)		4					1	5	2			3
	(b)			1					5	3		4	2
2. Identify and rank 10 events from best-to-worst	(a)	4	8	2		1		7	10	9	3	6	5
	(b)	10	7	8		5	4	6	1	2	9	3	
	(b)	2	7		10	1	4	9	3	5		6	8
	(b)	7	8	2	3			4	10	5	1	9	6
	(b)	2		3	6	4	8		10	5	1	9	7
3. Identify a set of attributes that may affect the ranking	(a)	Gain (high), Tortuosity (low), Number of Phases (low)											
	(b)	(Tortuosity + Number of Phases), (Gain + Territory Position)											
	(b)	Tortuosity, Number of Phases, Start Event											
	(a)	Gain, Start Event, Number of Phases											
	(a)	Gain, Number of Phases											
4. Formulate a ranking based on the set of attributes	(c)	N/A											
	(c)	N/A											
	(c)	N/A											
	(a)	Combination of high gain, low tortuosity and a weighted start event (e.g., turnover is more important than scrum)											
	(b)	Sequences containing high gain or high number of phases from various start events											

Fig. 9. Table showing the consultation session results for sorting rugby events. Each sub-row within the four primary tasks correspond to five participants along with their optional meta-answer (see Section 8 for details). For task 1 and 2, each cell in the result column indicate 12 possible events. These are labelled and color-mapped from worst-to-best with 1-5 and 1-10 respectively according to the user's ranking criteria.

additional optional meta-answers that describe the following: (a) I am reasonably confident about my answer, (b) I am unsure about my answer and (c) I do not know how to do this. The user is required to give an answer for only (a) and (b). This would help us analyze the confidence of the participant's response. To select the events for tasks 1 and 2, participants used a basic system that consists of the three views shown in Fig. 7(b), (c), and (d) along with video playback. We consider this as our benchmark for analyzing and sorting rugby events.

8.1 Results

In task 1, we asked the users to identify and rank five events based on importance. The participants generally felt confident with their choices. We anticipated this bias to have some affect in the following task for ranking ten events. Instead, we noticed that users became less confident of their answers. Looking at the event choices closely, there were only two instances of which more than one participant ranked equivalently. These events were ranked 10 (best), and rank 5 respectively. The evidence here supports our hypothesis and application of a gaussian-moderated ordering confidence τ. On average, task 1 took longer by 5.1 minutes, since users had to study all 12 possible video events for the first time. The most time-consuming process came from the random approach towards selecting events. Upon completion of task 1, users could formulate some implicit strategy for identifying more important events. For example, one observed method was to choose events based on highest gain. This made the second task much faster and simpler to do. Even so, we notice a diverse variation in event ranking as illustrated in Fig. 9. We conclude from both tasks that each participant had different interpretations of importance.

Feedback from the third task proved to be very interesting. A number of participants understood their sorting strategy used in the previous tasks did not easily generalize to importance (e.g., the most important event is not determined solely by the largest gain). All users recognized the multitude and combination of factors that potentially affected their importance criteria. However, the participants could speculate several attributes that are most significant. Tortuosity and number of phases are two attributes that were often coupled together. The amount of gain is also considered influential. It shows that there are some tacit similarities when quantifying importance. Furthermore, we note that none of the participants chose *time* as a sortable attribute, which supports the use of partial knowledge for narrowing the search space. The final task involved formalizing their ranking function. Whilst most participants attempted to write their strategy, they agreed that this is too complex to do. The feedback from the study suggest the need for a tool to discover such a function. Participants believed that such a tool would be powerful and useful towards sorting rugby event data in a more effective manner. After the consultation, we demonstrated our software to each participant. The initial feedback was positive, and they could see the benefits of a visual analysis system to help understand, refine, and incorporate their knowledge within the analysis.

8.2 Discussion

Among several modeling techniques, we chose three types of regression methods to formalize a user's tacit and partial knowledge. A fundamental step in many machine learning methods involves training the model. Based on the study in Section 8, we demonstrate training off a relatively small sample (5-10 events) which makes it difficult to judge how accurate and robust the derived formal knowledge is. Clearly, a more extensive training and validation process of multiple matches would yield a better sorting function. With enough training, the formal knowledge can then be stored externally as a template metric.

We have demonstrated our system for sorting rugby events. The framework can easily be applied to other sports (e.g., football, basketball and tennis) since both regression analysis and parallel coordinates are generalizable. For our glyph-based visualization, a multivariate glyph such as Star Glyphs [28] can be used instead. However, one limitation of such designs is that a greater learning process may be required in comparison to domain specific glyphs that are semantically richer [19]. Extensions to higher dimensionality is also scalable, giving more (and possibly more accurate) solutions to the user's sorting criteria. A potential issue here is the increase in computational cost can affect the interactivity of the system for very large dimensions.

9 Conclusion

We proposed a knowledge-assisted, event ranking framework for interactive multivariate sorting of sport event data. Users provide tacit knowledge into the system by selecting and ranking a subset of events as input for the model. We use regression analysis to discover a set of influential sort keys, and a formal sorting function that reorganizes events based on the user's ordering requirement. This allows users to sort events by an ad hoc criteria such as importance. We have found that our visual analytic approach significantly enhances the usability of multivariate sorting and demonstrate its usefulness in rugby event analysis. As future work, we would like to investigate how our system performs over existing software, and to validate the accuracy of the derived formal knowledge when organizing different matches.

References

[1] F. S. Acton. *Analysis of Straight-line data*. Wiley, 1966.

[2] I. Borg and T. Staufenbiel. Performance of snow flakes, suns, and factorial suns in the graphical representation of multivariate data. *Multivariate Behavioral Research*, 27(1):43–55, 1992.

[3] E. Brown, L. Jingjing, C. Brodley, and R. Chang. Dis-function: Learning distance functions interactively. In *IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 83–92, 2012.

[4] R. Chang, C. Ziemkiewicz, R. Pyzh, J. Kielman, and W. Ribarsky. Learning-based evaluation of visual analytic systems. In *Proceedings of the 3rd BELIV'10 Workshop: BEyond time and errors: novel evaluation methods for Information Visualization*, BELIV '10, pages 29–34, 2010.

[5] H. Chernoff. Using faces to represent points in *k*-dimensional space graphically. *Journal of the American Statistical Association*, 68:361–368, 1973.

[6] T. Cox and M. Cox. *Multidimensional Scaling*. Chapman and Hall, 2001.

[7] A. Doubleday. Use of card sorting for online course site organization within an integrated science curriculum. *Journal of Usability Studies*, 8(2):41–54, 2013.

[8] A. Endert, P. Fiaux, and C. North. Semantic interaction for sensemaking: Inferring analytical reasoning for model steering. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2879–2888, 2012.

[9] M. Harrower and C. A. Brewer. *ColorBrewer.org: An Online Tool for Selecting Colour Schemes for Maps, in The Map Reader: Theories of*

Mapping Practice and Cartographic Representation. Wiley-Blackwell, 2011.

[10] F. Heimerl, S. Koch, H. Bosch, and T. Ertl. Visual classifier training for text document retrieval. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2839–2848, 2012.

[11] B. Hoferlin, R. Netzel, M. Hoferlin, D. Weiskopf, and G. Heidemann. Inter-active learning of ad-hoc classifiers for video visual analytics. In *IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 23–32, 2012.

[12] M. D. Hughes and I. M. Franks. *Notational analysis of sport*. London: E. & F.N. Spon., 1997.

[13] A. Inselberg. The plane with parallel coordinates. *The Visual Computer*, 1:69–91, 1985.

[14] D. H. Jeong, C. Ziemkiewicz, B. Fisher, W. Ribarsky, and R. Chang. ipca: an interactive system for pca-based visual analytics. In *IEEE VGTC conference on Visualization*, pages 767–774, 2009.

[15] L. Jin and D. C. Banks. Visualizing a tennis match. In *Proceedings of the 1996 IEEE Symposium on Information Visualization (INFOVIS '96)*, INFOVIS '96, pages 108–, Washington, DC, USA, 1996. IEEE Computer Society.

[16] I. T. Jolliffe. *Principal Component Analysis*. Springer, second edition, 2002.

[17] A. Klippel, F. Hardisty, R. Li, and C. Weaver. Colour-enhanced star plot glyphs: Can salient shape characteristics be overcome? *Cartographica: The International Journal for Geographic Information and Geovisualization*, 44(3):217–231, 2009.

[18] D. E. Knuth. *The Art of Computer Programming, Vol. 3: Sorting and Searching, Second Edition*. Addison-Wesley, Reading, Mass., 1998.

[19] P. A. Legg, D. H. S. Chung, M. L. Parry, M. W. Jones, R. Long, I. W. Griffiths, and M. Chen. Matchpad: Interactive glyph-based visualization for real-time sports performance analysis. *Computer Graphics Forum*, 31(3pt4):1255–1264, 2012.

[20] H. Lipford, F. Stukes, D. Wenwen, M. Hawkins, and R. Chang. Helping users recall their reasoning process. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)*, pages 187–194, 2010.

[21] G. Mistelbauer, A. Kochl, H. Bouzari, S. Bruckner, R. Schernthaner, M. Sramek, I. Baclija, and M. E. Groller. Smart super views: A knowledge-assisted interface for medical visualization. In *IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 163–172, 2012.

[22] A. Moore, P. Whigham, C. Aldridge, A. Holt, and K. Hodge. Rugby: (a) union of space and time. In P. A. Whigham, editor, *Proceedings: Thirteenth Annual Colloquium of the Spatial Information Research Centre*, pages 183–194, 2001.

[23] M. L. Parry, P. A. Legg, D. H. S. Chung, I. W. Griffiths, and M. Chen. Hierarchical event selection for video storyboards with a case study on snooker video visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):1747–1756, 2011.

[24] H. Pileggi, C. D. Stolper, J. M. Boyle, and J. T. Stasko. Snapshot: Visualization to propel ice hockey analytics. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2819–2828, 2012.

[25] R. E. Roth, B. G. Finch, J. I. Blanford, A. Klippel, A. C. Robinson, and A. M. MacEachren. Card sorting for cartographic research and practice. *Cartography and Geographic Information ScienceCartography and Geographic Information Science*, 38(2):89–99, 2011.

[26] G. Rugg and P. McGeorge. The sorting techniques: a tutorial paper on card sorts, picture sorts and item sorts. *Expert Systems*, 14(2):80–93, 1997.

[27] A. Sen and M. Srivastava. *Regression Analysis: Theory, Methods, and Applications*. Springer, 1990.

[28] J. Siegel, E. Farrell, R. Goldwyn, and H. Friedman. The surgical implication of physiologic patterns in myocardial infarction shock. *Surgery*, 72:27–35, 1972.

[29] X. Wang, D. H. Jeong, W. Dou, S.-W. Lee, W. Ribarsky, and R. Chang. Knowledge assisted visualization: Defining and applying knowledge conversion processes to a visual analytics system. *Computers and Graphics*, 33(5):616–623, 2009.

[30] M. O. Ward. A taxonomy of glyph placement strategies for multidimensional data visualization. *Information Visualization*, 1(3-4), 2002.

[31] C. Weaver. Multidimensional visual analysis using cross-filtered views. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)*, pages 163–170, 2008.