# Visualize It-Wise! An Iteration-Wise Computational Framework for Real-Time Visual Analytics

Jaegul Choo\* School of Computational Science and Engineering Changhyun Lee<sup>†</sup> School of Electrical and Computer Engineering Haesun Park\* School of Computational Science and Engineering

Georgia Institute of Technology

**Abstract**—Abstract Visual analytics has been gaining increasing interest due to its fascinating characteristic that leverages both humans' visual perception and the power of computing. Although various computational methods are being proposed, they do not properly support visual analytics. One of the biggest obstacles towards their real-time visual analytic integration is their high computational complexity. As a way to tackle this problem, this paper presents an iteration-wise computational framework, motivated by the fact that most advanced computational methods work by refining the solution iteratively. By visually delivering the results for each iteration to users, the proposed framework enables users to quickly acquire the information that the computational method provides as well as the ability to interact with them in real time. We show the benefits of the proposed framework by using various dimension reduction and clustering methods.

Index Terms-Radiosity, global illumination, constant time

# **1** INTRODUCTION

The innate ability of humans to quickly perceive insight through visual analysis and decision processes has been a key factor in the growth of visual analytic research [12, 17]. One of the most significant efforts made by visual analytics researchers is the integration of various computational methods from data mining and machine learning areas with visual analytics so that users can benefit from intelligent meaningful information generated by these techniques.

For example, dimension reduction and clustering methods have been commonly used in high-dimensional data visual analytics [4, 16]. More recently, latent Dirichlet allocation (LDA) [3], a popular method for document topic modeling, has been adopted in a wide variety of visual analytics systems for document analysis [19, 7, 13].

However, a critical hurdle in the integration of computational methods into visual analytics is the significant amount of computational time required by these methods. As computational methods become more advanced and capable, they usually run much slower, making it almost impossible to use and interact with them smoothly in real-time visual analytic scenarios.

Due to these reasons, even though numerous computational methods are currently being developed and some, e.g., *t*-distributed stochastic neighbor embedding (t-SNE) [18], even claim their suitability directly in visualization applications, the state-of-the-art systems in visual analytics are not currently keeping up with the advancements in computational methods. Consequently, people still adopt only a few basic computational methods, such as principal component analysis (PCA) [10] and multidimensional scaling (MDS) [5] for dimension reduction, hierarchical clustering and *k*-means [2] for clustering, etc., in many real-world analysis tasks.

We claim that, by leveraging the various characteristics which arise from the integration of computational methods into visual analytics,

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. a significant amount of running time can be removed from the computational methods. Thus, we propose a novel framework called an iteration-wise integration of computational methods for real-time visual analytics.

Our proposed framework exploits the fact that most computational methods are iterative methods. An iterative method refers to a mathematical procedure that iteratively refines a sequence of approximate solutions towards the final converged solutions. In contrast, a direct method solves a given problem in a finite number of steps using some closed-form expressions of the final solution. However, problems solvable with a direct method are limited, while many modern computational methods, which can handle complicated problems, are mostly iterative methods.

Another aspect of iterative computational methods that we exploit is that, throughout the iterations, major refinement of the solution typically occurs in early iterations while only minor changes occur in the later iterations. Considering that humans can quickly perceive the overall structure/trend in data, a result obtained from an early iteration may already provide enough information.

However, apart from well-principled stopping criteria studied in most computational methods, there is no straightforward method to determine the point at which to terminate the iteration without sacrificing a noticeable amount of information when compared to the final solution. Instead, our proposed approach is to *visualize the intermediate results from iterations when they become available*.

This idea takes advantage of the useful characteristic of many iterative computational methods that the intermediate result at every iteration maintains the same form as the final output. For instance, in the case of clustering, many clustering algorithms including k-means give the cluster memberships or coefficients of the entire set of data items at each iteration. The only difference between the results from each iteration is the preciseness of the result.

Therefore, by adopting the iteration-wise integration framework, users would be able to acquire an overview of the results much more quickly than in typical cases where they have to wait for the entire iterations to finish. Furthermore, the proposed framework makes user interactions with computation methods much more responsive since these interactions can be reflected immediately in later iterations. Users would not need to wait for a full run of the set of iterations if they were to make a particular change to the computational methods

<sup>\*</sup>e-mail: {jaegul.choo, hpark}@cc.gatech.edu

<sup>&</sup>lt;sup>†</sup>e-mail: {clee407@gatech.edu

such as parameter changes, data filtering, etc.

Motivated by these ideas, this paper describes the iteration-wise integration framework of computational methods (Section 3) and present the customization of well-known computational methods fitted to the proposed framework (Section 4). Using the presented computational methods, we show the quantitative iteration-wise analyses of computational methods and use cases where we applied the proposed framework in several well-known visual analytics systems (Section 5). Finally, we conclude the paper (Section 6).

# 2 RELATED WORK

#### **3** ITERATION-WISE INTEGRATION FRAMEWORK

First, we discuss an overall flow of the iteration-wise integration framework of computational methods in visual analytics (Fig. 1(a)), by highlighting its differences from the standard (non-iteration-wise) approach (Fig. 1(b)).

Let us begin with a general procedure when an iterative computational method is integrated in visual analytics. As shown in Fig. 1(a), input data, which are usually represented as multidimensional vectors, are given to the computational module along with its required parameter values. The computational module pre-processes the data, if necessary, and runs through iterations, which are usually divided into multiple sub-routines, until it converges. Upon convergence, the output goes through a post-processing step.

The final output of the computational module is then passed to the visualization module, which encodes it in a visual space and finally delivers its visualization to users. For example, the output of a dimension reduction method, e.g., PCA, can map data items onto the coordinates of the screen space, and the output of clustering can be used to color-code each group of data clusters.

Users can then better explore the visually represented data with the help of the information provided by the computational method and often interact with computational methods by adjusting their input data as well as their parameters. These interactions trigger another run of the computational method. For example, given the cluster summary for a set of text documents, if a user finds an interesting cluster, the user may perform another iteration of clustering on the particular subset to obtain more details about the chosen subset. On the other hand, users might want to adjust the number of clusters, which is usually a userspecified parameter in clustering methods, to find the best clustering result for the data.

In most of the described visualizations and interactions, the standard framework generally treats the computational module as a black box, which the visualization module has no control over, depicted by a gray rectangle in Fig. 1(a). In other words, once the computational module has been initiated, visual analytic systems must wait for it to finish its iterations before it outputs the visualization to users.

On the contrary, the iteration-wise integration framework takes the results of intermediate iterations out of the computational module and delivers them to the visualization module whenever they are available. More specifically, as highlighted with the red horizontal line in Fig. 1(b), the result from each iteration is always passed to the post-processing step, the output of which, in turn, reaches all the way to the visualization module, regardless of whether it has converged or not. Consequently, these intermediate results are visualized to users much more quickly than having to wait for the converged solutions.

In addition, the iteration-wise integration framework enables the above-discussed interactions to be instantly reflected by directly interacting with the process for each iteration of the computational module, as highlighted with the red vertical line in Fig. 1(b). For instance, given the result of a particular iteration, one could exclude certain data items from the following iterations, which accelerate the later iterations due to the reduced data size. Furthermore, users could change the number of clusters while a clustering method is running, which immediately affects the following iterations.

# 3.1 Issues and Solutions

#### **Computational Overhead and Multi-threading**

Computational overheads are one of the issues that can be potentially

introduced by this framework. As can be seen in the red-lined stacked rectangle blocks in Fig. 1(b), visual analytics systems have to process the output for each iteration repetitively while the standard approach needs to process only the final output once. These additional computations could undermine the effectiveness of the proposed framework. Let us suppose that a particular computational method, which requires 50 iterations to converge, converges in a minute. If the proposed framework runs only 4-5 iterations within the same amount of time, then users might prefer the standard approach instead of being able to check the intermediate results since the results from such early iterations may not be satisfactory.

However, we claim that this issue can be easily overcome by applying a multi-threaded approach to the proposed framework. As shown by the blue ellipses in Fig. 1(b), the entire process can be separated into two concurrent processes/threads. The first thread shown to the left is responsible only for the sub-routines inside the iteration while the second thread on the right handles actions from the post-processing block to the visualization block. These two threads communicate with each other via a message queue, as shown by the blue rectangle on top in Fig. 1(b), where the outputs for each iteration for post-processing are to be stored.

Modern computers are usually equipped with at least two or more cores on the CPU. These two threads can be executed virtually in parallel, which hardly slow down the computational methods compared to the standard approach. Although not included in this paper, for the computational methods we customized, we compared the total computing time between the iteration-wise and the standard frameworks, but with multi-threading implemented, there were essentially no differences in their running times.

Even in this multi-threading framework, the following case may still be problematic. Suppose the second thread involves more intensive computations than the first thread because, for example, the postprocessing block takes more time than the processes at each iteration. As a result, the second thread would act as a bottleneck in the overall flow of the proposed framework, resulting in the message queue increasing. One way to handle this issue is to store the results of each iteration periodically rather than storing every one of them in the first thread. Alternatively, the second thread could take the most recent iteration-wise results and discard the remaining older ones from the message queue. Under this situation, the visualization of the intermediate results may be somewhat discontinuous, but users would always be given the most recent result, which should be the most accurate solution up to the current iteration.

Finally, the other overhead comes from copying results from each iteration to the message queue, which results in a memory write operation. In the standard approach, these results for each iteration are usually written to the same memory space over iterations since the results from previous iterations do not need to be maintained. However, memory write operations are generally very fast. Furthermore, the outputs from each iteration of computational methods take up a much smaller memory compared to input data. For example, even if the data is a very high-dimensional, say, in the hundreds of thousands of dimensions, such as is the case in text data, the dimension reduction outputs would only be two-dimensional representations assuming they are visualized in a 2D space. Since the amount of additional computational time and memory that is required by our approach is minimal, we do not see memory overheads being a critical issue.

# Visual Inconsistency and User Control

The second issue in the proposed framework is the visual inconsistency, which occurs during visualization updates, due to dynamic results changing each iteration. The most severe case occurs when the visualization changes too frequently. Although the amount of change generally diminishes as the iterations proceed, frequently changing visualizations may prevent users from obtaining a consistent picture of the data.

To address these issues with visual inconsistencies, we've come up with several possible controls. The first most basic option would be a stop and resume control which would stop and resume updates of the visualization. Secondly, a time period control would manage the



(a) The standard (non-iteration-wise) framework



(b) The iteration-wise framework

Fig. 1: An overall diagram of the iteration-wise integration framework (b) in contrast to the standard (non-iteration-wise) one (a). In the standard framework (a), a computional method is treated as a black box, as depicted by a gray rectangle. On the other hand, the iteration-wise framework (b) breaks down the computational method at its iteration level, allowing it to be visualized at each iteration while taking into account any user iteractions. The blue line separates the overall procedure into two separate threads with their message queue, as shown in the blue rectangle, to remove potential computational overheads.

length of the visualization. Additionally, we could pair this time controller with two choices - the option to visualize the most up-to-date result or to visualize the result of the next item in the queue, which would provide the user with smoother visual transitions. Similar to the 'stop/resume' interaction, since our approach maintains each of the intermediate results, we could simply expand the controls to also add both the 'play backwards' and 'jump to...' options. These interactions would help users understand the overall trajectory of the results through each iteration. Through the use of these controls, it is very possible that the user may uncover an interesting insight into the data at a particular iteration or a series of iterations.

# 3.2 Related Work

In fact, there have been existing approaches which follow the proposed framework. For example, the force-directed graph layout or spring layout [8], e.g., the 'graphview' implemented in the Prefuse toolkit, <sup>1</sup> visualizes data immediately after the first iteration of the layout algorithm. Furthermore, the Prefuse toolkit provides various interactions, such as moving data points and changing parameters, e.g., 'Spring-Coefficient,' etc., that provide feedback and update the visualization during its iterations such as moving data points and changing parameters, e.g., 'Spring-Coefficient,' etc.

The reason for such a tight iteration-wise integration of the forcedirected graph layout is partly because its purpose is directly for visualization and the output is straightforwardly mapped to a visual space without involving much post-processing.

On the other hand, in many cases, computational methods are not entirely intended for visualization. Therefore, non-trivial preprocessing/post-processing steps may have to be involved for their adoption to visual analytics. In addition, advanced computational methods are often hard to understand. These numerous obstacles distract people from the fact that most computational methods have indeed an iterative nature similar to the force-directed graph layout, and they can be easily dismantled and tightly coupled with visual analytics at their iteration level. The rest of the paper will focus on how we can achieve this claim for various computational methods.

#### 4 CUSTOMIZED METHODS IN VISUAL ANALYTICS SYSTEMS

In this section, following the proposed framework, we present several customized computational methods in visual analytics systems. To begin with, we have chosen three visual analytics systems, FodavaTestbed,<sup>2</sup> Jigsaw,<sup>3</sup> and iVisClustering [13], which involve computational methods.<sup>4</sup>

FodavaTestbed is a visual analytics system for high-dimensional data, where users can apply various dimension reduction and clustering methods for exploratory analysis. Among various methods supported, we have chosen three dimension reduction methods, <u>1. MDS</u>, <u>2. PCA</u>, and <u>3. t-SNE</u>. Jigsaw is a well-known system for document analysis, and we have chosen <u>4. k-means</u>, which is used to provide a summary in terms of a compact set of clusters. Finally, iVisClustering is an interactive document clustering system which uses <u>5. LDA</u>, a popular topic modeling method.

In the following, we describe how each method is customized along with the additional interactions we implemented in the proposed framework.

# 4.1 Principal Component Analysis (PCA)

PCA [11] is a well-known dimension reduction method that captures the maximal variance in the data via a linear projection. PCA is mainly based on the method called eigendecomposition, the algorithms of which are categorized into two different methods, the QR algorithm and the Lanczos algorithm [9].

Basically, the Lanczos algorithm approximates a given data matrix by a much smaller one in the Krylov subspace [9], the dimension of

<sup>&</sup>lt;sup>2</sup>http://fodava.gatech.edu/fodava-testbed-software

<sup>&</sup>lt;sup>3</sup>http://www.cc.gatech.edu/gvu/ii/jigsaw/

<sup>&</sup>lt;sup>4</sup>We obtained the code from the original authors of the systems.

<sup>&</sup>lt;sup>1</sup>http://prefuse.org/

which iteratively expands, and efficiently solves the eigendecomposition on the latter matrix. Due to the nature that this matrix wellapproximates the largest eigenvectors of the original one, the Lanczos algorithm performs much faster than the QR algorithm in visual analytics in which only a few dimensions are needed.

We customize the Lanczos-based PCA implementation of FodavaTestbed so that the results for each iteration are dynamically visualized.

# 4.2 Multidimensional Scaling (MDS)

MDS [5] is a traditional dimension reduction method that attempts to preserve given distances/relationships of data items in a lowerdimensional space. Given the ideal distance  $\delta_{ij}$  between  $x_i$  and  $x_j$ , MDS solves

$$\min_{x_1,\dots,x_n} \sum_{1 \le i < j \le n} \left( d_{ij} - \delta_{ij} \right)^2, \tag{1}$$

where  $d_{ij}$  is the distance between the reduced dimensional vectors  $x_i$ and  $x_j$ . A Euclidean distance  $||x_i - x_j||_2$  is usually used for  $d_{ij}$ . Solving Eq. (1) iteratively refines  $x_i$ 's based on various optimization techniques [6]. We customize MDS in FodavaTestbed by extracting the  $x_i$ 's at each iteration from the MDS implementation.

## Interaction / User Intervention

Additionally, while the results for each iteration of MDS are visualized in a scatter plot, we support the interaction capability that enables users to move the data points by mouse via drag-and-drop, similar to the Prefuse force-directed layout. Then, during the MDS iterations, their new positions in the screen space are translated back to the MDS output coordinates,  $x_i$ 's. The changes in  $x_i$ 's at a particular iteration then affect the following iterations by generating different  $d_{ij}$ 's. In terms of how MDS behaves due to these changes, we provide two different capabilities: 'soft' vs. 'hard' placement. The soft placement continues iterations without any changes in MDS behaviors. It is equivalent to restarting MDS with the intermediate result at the particular iteration as the initial values for  $x_i$ 's.

The hard placement capability fixes the values of  $x_i$ 's for points moved by the user. This can be easily achieved by skipping the update step of these  $x_i$ 's in the following iterations. Note that, however, even though their values do not change, other data points are still influenced by these fixed points, and in this sense, our approach is a semi-supervised MDS that reflects user interventions.

When using the semi-supervised MDS, an important advantage of the proposed framework is that users can immediately check the effects of these interactions via the iteration-wise visualization. Our modifications in FodavaTestbed support both types of interactions.

## 4.3 t-Distributed Stochastic Neighbor Embedding (t-SNE)

t-SNE [18] is a relatively new dimension reduction method. It interprets pairwise distances as probabilities both in high-dimensional and lower-dimensional spaces and tries to minimize their Kullback–Leibler divergence, a distance measure between probability distributions. Unlike the previous methods discussed, it focuses on preserving neighborhood relationships instead of global ones, and it has shown its outstanding capabilities in visualizations.<sup>5</sup>

Although we skip the detailed formulations due to the page limit, the algorithm works iteratively by refining the lower-dimensional coordinates based on a gradient descent framework. In practice, however, t-SNE does not provide a clear stopping criterion, and thus it typically iterates several hundred times by default for any data set, which usually takes a significant amount of time. We customize the t-SNE in FodavaTestbed in a similar manner to the way we altered MDS.

#### Interaction / User Intervention

Likewise, we provide both the soft and hard placement interactions for t-SNE, as discussed in MDS. Although the algorithm details are different, the overall iterative procedure turns out to be quite similar

<sup>5</sup>http://homepage.tudelft.nl/19j49/t-SNE.html

to MDS. Thus, for the soft placement, we restart t-SNE with the intermediate results immediately during iterations. For the hard one, we skip the update step for data items moved by the user in the following iterations while they still influence other points in the t-SNE iterations. Therefore, our altered method can be viewed as a semi-supervised t-SNE.

## 4.4 k-means

k-means, which is a widely-used clustering method, performs the following steps iteratively: 1. computing the centroid of each cluster by averaging the data vectors in the corresponding cluster and 2. updating the cluster assignment of each data item based on its closest cluster centroid. The iteration terminates when there are no cluster membership changes.

Although Jigsaw provides a cluster view based on k-means, it currently visualizes only the pre-computed results since k-means is usually very slow to converge. We customize it so that users can run k-means in real time and the intermediate cluster memberships are dynamically visualized.

#### Interaction / User Intervention

Additionally, we add several interaction capabilities in the proposed framework. One is to split/merge clusters during iterations. On a split/merge interaction, similar to the soft placement in MDS and t-SNE, *k*-means restarts with the intermediate cluster memberships that reflect split/merged clusters, involving dynamic changes in a *k*-means parameter which represents the number of clusters.

Similar to the hard placement in MDS and t-SNE, another capability we provide is the option to fix the cluster assignments of the data in a particular cluster. To accomplish this, we skip the updating step of the cluster assignment for these data in the following iterations. However, they still contribute to the centroid computing step. A similar semi-supervised way of *k*-means was previously proposed [1], but our framework significantly accelerates such interactions with *k*-means.

# 4.5 Latent Dirichlet Allocation (LDA)

LDA [3] is a popular topic modeling method for documents based on a generative probabilistic model. Given a number of topics, it gives two outputs: the term-wise distribution of each topic and the topic-wise distribution of each document. The iterations of LDA basically update these two outputs alternately. From a clustering viewpoint, the former corresponds to a centroid vector of each topic cluster, and the latter to a soft-clustering coefficient. By taking the topic index that has the maximum value in the latter, a document is clustered to a particular topic.

iVisClustering uses one of the fastest LDA libraries called Mallet [14], which implements LDA based on a Gibbs sampling [15]. Although this sampling-based approach does not guarantee a convergence, it is being widely used because of its simplicity and robustness against overfitting, compared to the variational approximation method proposed in [3]. Due to this convergence issue, LDA usually iterates a pre-defined number of iterations and usually requires a significant amount of time. We customize the Mallet library so that it can give the outputs from each iteration to iVisClustering, allowing iVisClustering to dynamically update its visualization.

#### Interaction / User Intervention

In addition to the original iVisClustering interaction capabilities being available during iterations, we also add several interactions with LDA in iVisClustering, similar to those in Jigsaw: splitting/merging clusters and fixing the cluster assignments of particular data items during iterations. The customization of LDA for such interactions is similar to *k*-means, and thus we skip the details due to the page limit.

# **5** EXPERIMENTS

In this section, we present behaviors within each iteration for computational methods as well as their interactive aspects in the proposed framework.



(a) PCA criteria values and output (b) The scatter plot at the first iteration (c) The scatter plot at the third iteration(d) The scatter plot at the 80th iteration changes



(e) MDS criteria values and output (f) The scatter plot at the 10th iteration (g) The scatter plot at the 30th iteration(h) The scatter plot at the 156th (conchanges verged) iteration

Fig. 2: The behavior for each iteration of PCA and MDS and their visualization snapshots. In (a) and (e), the red lines represent the criteria values of each method, the lower-dimensional variance in PCA and the stress value, i.e., Eq. (1), in MDS. The blue line is the Euclidean distances of the lower-dimensional outputs between the current and the previous iterations, and the black line is those between the current and the final iterations. In (e), the black and the blue lines almost coincide. In PCA, 1,420 facial image data representing pixel values in 2,048 dimensions have been used, and in MDS, 500 handwritten digit data representing pen traces in 16 dimensions have been used.



Fig. 4: A point-moving interaction example using t-SNE. The two overlapping clusters, 'I' and 'o,' are separated due to a user interaction of moving apart a few points from each cluster. 1,558 spoken letter data represented in 618 dimensions have been used.



Fig. 3: Computing times of the example in Fig. 2.

#### 5.1 Iteration-wise Behaviors and Visualization

Fig. 2 shows the behaviors of each iteration for the customized PCA and MDS along with their computing times shown in Fig. 3. In PCA, the criteria value, i.e., the lower-dimensional variance, as well as the lower-dimensional outputs (Fig. 2(a)) converge within a few iterations, indicating that only a few iterations of the Lanczos algorithm

are needed in visual analytics applications (Figs. 2(b)-(d)). Nonetheless, each iteration takes roughly the same amount of computation time except for the first iteration which includes the pre-processing time. (Fig. 3(a)). Instead of having to perform a fairly large number of iterations, as most PCA algorithms do, the iteration-wise visualization enables users to obtain an equivalent visualization much quickly. A similar argument applies to MDS as well. Although its convergence is relatively slow compared to PCA (Fig. 2(e)), we obtain the results similar to the final one ahieved at the 10th iteration (Figs. 2(f)-(h)). We do not present the quantitative analyses for t-SNE, but we found tendencies similar to MDS, and we will focus on its interactive aspects in the following section.

In clustering, the behavior of each iteration of k-means is presented in Fig. 5 as well as their snapshots in Jigsaw in Fig. 7. In Jigsaw, in order to best assist users in easily identifying the location and the number of changes that occur while the visualization is dynamically updated, we draw arrows to represent where a particular data item has moved relative to the previous iteration, and color-code each data item to represent which cluster index it previously belonged to. As shown in Fig. 7 and in the redlines in Fig. 5, although significant change occurs in early iterations (Fig. 7(a)), they diminishes quickly, as seen in the



(a) The second-iteration result

(b) The sixth-iteration result



(c) The converged (25th-iteration) result after fixing clusters

(d) The converged (26th-iteration) result without any interaction

Fig. 7: The results of the iteration-wise integration of k-means in Jigsaw. At the sixth iteration, the interaction of fixing the yellow-colored clusters is made (b). The final result with and without this interaction is shown in (c) and (d), respectively. The NSF-awarded abstract data have been used.



(a) The third-iteration result (b) The fourth-iteration result after(c) The final (15th-iteration) result with(d) The final (7th-iteration) result with-split/merge interactions split/merge interactions out split/merge interactions

Fig. 8: An example of split/merge interactions. The yellow and green ones in (a) are merged to the same-colored ones, respectively, in (b), and the white one in (a) is split to the-same colored ones in (b). Webpages about autism have been used as input data.

sixth iteration (Fig. 7(b)), which is not much different from the final result (Fig. 7(d)). However, the time each iteration takes is almost the same (Fig. 5(d)).

Finally, LDA, which is a sampling-based approach, shows a significantly different behavior from the previous methods (Fig. 6). Although cluster membership changes between iterations generally decrease and the solution narrows to the final solutions (Fig. 6(a)), cluster memberships change significantly even after many iterations, in this case after 1,200 iterations. In iVisClustering, we could see the top keywords of each topic become somewhat stable after several hundreds of iterations (Fig. 6(a)), but the randomness of the sampling-based algorithms might make it harder to give consistent visualizations when compared to deterministic methods in the iteration-wise framework.

# 5.2 Interaction Examples

Basically, in all three systems, we provide basic interactions that control the visualization for each iteration, as discussed in Section 3.1. In the following, we show several use cases of the interactions discussed in Section 4.

#### 5.2.1 Moving data points in t-SNE

Fig. 4 shows an interesting interaction which involves moving a data point in t-SNE. Given some overlapping clusters in a particular visualization generated by t-SNE (Fig. 4(a)), users place several points from different clusters far apart (Fig. 4(b)), and then t-SNE reflects these changes in the following iterations, resulting in separation of most points in two clusters from each other (Figs. 4(c)(d)). This simple, yet powerful example clearly illustrates the advantage of providing users with the ability to interact with computational methods in our framework in real-time visual analytics.

## 5.2.2 Fixing cluster assignments in *k*-means

For our *k*-means method, we provide users with another interaction that allows them to fix cluster assignments for particular data items at a certain iteration. This interaction becomes especially useful when users feel that particular clusters are adequate and want to prevent them from changing much. In addition, fixing some clusters that are already stable in early iterations can accelerate the later iterations by excluding them from the cluster assignment step.

Fig. 5 shows the effects of such interactions. First, we start with the same example shown in Fig. 7, but we fix the clustering assignments of



(a) The third-iteration result

## (b) The 300th-iteration result

(c) The 700th-iteration result

Fig. 9: An example of filtering documents whose cluster memberships are unclear. This interaction is done in the 300th iteration, and the topics become clearer in the later iterations. 20 newsgroups data have been used.



Fig. 5: The behaviors for each iteration of k-means with and without the interaction made in Fig. 7(b). In (a), the decreasing lines are the cluster membership changes between the current and the previous iterations while the increasing ones are the correct cluster memberships with respect to the final solutions without the interaction. The black vertical line represents the iteration point of the interaction made.



Fig. 6: The iteration-wise behaviors of LDA. In (a), the blue line represents cluster membership changes between the current and the previous iterations while the red line represents the correct cluster memberships with respect to the final solutions.

the cluster highlighted in yellow rectangles, which amounts to 44% of the total data items, at the sixth iteration (Fig. 7(b)). Once this interaction is performed, the computing times for the following iterations of k-means drops significantly (Fig. 5(b)). However, only less than 10% of the final cluster memberships differ from the final results without this interaction, as shown in the increasing red line in Fig. 5(a). The final outputs of the cluster view in Jigsaw of the two cases can also be compared in Figs. 7(c) and (d), both of which are similar in terms of cluster sizes as well as keyword descriptions.

# 5.2.3 Split/merge clusters in k-means

Our customization of k-means enables users to merge multiple small or semantically related clusters or split large or unclear clusters. Fig. 8 shows its example in Jigsaw. In the third iteration, we merge yellow and green clusters and split a white cluster (Fig. 8(a)). The resulting is shown in Fig. 8(b). We obtain a much more balanced set of clusters (Fig. 8(c)) compared to the final result in which no splitting/merging was performed (Fig. 8(d)). Furthermore, after analyzing the documents in two split clusters, we found that one of the clusters primarily contained documents about the causes of autism while the other about the symptoms, as seen in the keyword summary in Fig. 8(c). Without the interaction, one will notice in Fig. 8(d) that these clusters are not easily separated.

# 5.2.4 Filtering noisy documents to improve topics in LDA

The ability to filter noisy documents has been an appealing interaction for LDA in iVisClustering. To be specific, given parallel coordinate representations of topic-wise distributions of documents, users can interactively filter out documents that are not strongly related to a single topic, i.e., documents that have a very small maximum value in the topic-wise distribution. By removing them and re-running LDA, iVis-Cluster generally obtains significantly clearer topics. In the iterationwise framework, we performed this interaction near the 300th iteration (Fig. 9(b)), which is an early iteration when compared to the total number of iterations performed by LDA. However, such an interaction successfully generates clearer topics (Fig. 9(c)) over the standard approach where users have to wait for the algorithm to finish its full iterations in order to perform the same interaction.

# 6 CONCLUSIONS AND FUTURE WORK

We have presented the iteration-wise integration framework of computational methods for real-time visual analytics. One of its apparent advantages is its ability to present users with intermediate results during the interactions, which could reveal a significant amount of information immediately in visual analytics. Another important advantage is that it indeed opens up the possibility of many new micro-level interaction capabilities, which in the past have been considered to be too inefficient, and allows real-time control over computational methods in visual analytics. For example, the interactions we proposed in this paper are relatively simple, which do not involve any major algorithmic modifications. Our framework makes them significantly useful by enabling users to perform these interactions easily and efficiently.

However, the advantage of our framework can be limited when the changes between iterations remain nontrivial, resulting in inconsistent visualizations. We have seen similar limitations when using LDA in our framework due to the random nature of the used LDA algorithm. We plan to tackle this problem more actively by, for example, post-processing the results or even imposing additional constraints in computational methods so that the results from the following iterations do not change much from the current ones.

#### REFERENCES

- S. Basu, A. Banerjee, and R. J. Mooney. Semi-supervised clustering by seeding. In *Proceedings of the Nineteenth International Conference on Machine Learning*, ICML '02, pages 27–34, San Francisco, CA, USA, 2002. Morgan Kaufmann Publishers Inc.
- [2] C. Bishop. Pattern recognition and machine learning. Springer, 2006.
- [3] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, March 2003.
- [4] A. Buja, D. Cook, and D. Swayne. Interactive high-dimensional data visualization. *Journal of Computational and Graphical Statistics*, 5(1):78– 99, 1996.
- [5] T. F. Cox and M. A. A. Cox. *Multidimensional Scaling*. Chapman & Hall/CRC, London, 2000.
- [6] J. De Leeuw. Applications of convex analysis to multidimensional scaling. *Recent developments in statistics*, pages 133–146, 1977.

- [7] W. Dou, X. Wang, R. Chang, and W. Ribarsky. Paralleltopics: A probabilistic approach to exploring document collections. In *Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on*, pages 231 –240, oct. 2011.
- [8] T. M. J. Fruchterman and E. M. Reingold. Graph drawing by forcedirected placement. *Software: Practice and Experience*, 21(11):1129– 1164, 1991.
- [9] G. H. Golub and C. F. van Loan. *Matrix Computations, third edition*. Johns Hopkins University Press, Baltimore, 1996.
- [10] I. T. Jolliffe. *Principal Component Analysis*. Springer-Verlag, New York, 1986.
- [11] I. T. Jolliffe. Principal component analysis. Springer, 2002.
- [12] D. Keim. Information visualization and visual data mining. Visualization and Computer Graphics, IEEE Transactions on, 8(1):1–8, jan/mar 2002.
- [13] H. Lee, J. Kihm, J. Choo, J. Stasko, and H. Park. iVisClustering: An interactive visual document clustering via topic modeling. *Computer Graphics Forum*, 31(3pt3):1155–1164, 2012.
- [14] A. K. McCallum. Mallet: A machine learning for language toolkit. http://mallet.cs.umass.edu, 2002.
- [15] I. Porteous, D. Newman, A. Ihler, A. Asuncion, P. Smyth, and M. Welling. Fast collapsed gibbs sampling for latent dirichlet allocation. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '08, pages 569–577, New York, NY, USA, 2008. ACM.
- [16] J. Seo and B. Shneiderman. Interactively exploring hierarchical clustering results [gene identification]. *Computer*, 35(7):80 –86, jul 2002.
- [17] J. Thomas and K. Cook. Illuminating the path: The research and development agenda for visual analytics, volume 54. IEEE, 2005.
- [18] L. van der Maaten and G. Hinton. Visualizing data using t-SNE. Journal of Machine Learning Research, 9:2579–2605, 2008.
- [19] F. Wei, S. Liu, Y. Song, S. Pan, M. X. Zhou, W. Qian, L. Shi, L. Tan, and Q. Zhang. Tiara: a visual exploratory text analytic system. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '10, pages 153–162, New York, NY, USA, 2010. ACM.