Visual Analysis of Large Dental Imaging Data in Caries Research

Guangchen Ruan*
School of Informatics and Computing
Indiana University

Hui Zhang†
Pervasive Technology Institute
Indiana University

ABSTRACT

With dental imaging data acquired at unprecedented speed and resolution, traditional serial image processing and single-node storage need to be re-examined in a "BigData" context. Furthermore, most previous dental computing has focused on the actual imaging acquisition and image analysis tools, while much less research has focused on enabling caries assessment via visual analysis of large dental imaging data. In this paper we present DENVIS, an end-to-end solution for cariologists to manage, mine, visualize, and analyze large dental imaging data for investigative carious lesion studies. DENVIS consists of two main parts: data driven image analysis modules triggered by imaging data acquisition that exploit parallel MapReduce tasks and ingest visualization archive into a distributed NoSQL store, and user driven modules that allow investigative analysis at run time. DENVIS has seen early use by our collaborators in oral health research, where our system has been used to pose and answer domain-specific questions for quantitative assessment of dynamic carious lesion activities.

Keywords: dental computing, visual knowledge discovery, MapReduce

Index Terms: I.6.8 [SIMULATION AND MODELING]: Types of Simulation—Parallel; H.5.2 [INFORMATION INTERFACES AND PRESENTATION]: User Interfaces—Graphical user interfaces

1 INTRODUCTION

Detecting dental caries at the earliest stage and assessing the dynamic activities of carious lesions has become one of the most active research areas in cariology. By quantitatively monitoring carious lesions over time and correlating variables of interest, cariologists desire a deep understanding upon how carious lesions develop, progress, and can be treated. To support this goal, emerging technologies for diagnosis of dental caries (e.g., microfocus computed tomography (μCT) and Cone Beam CT (CBCT)) have been developed that enable cariologists to acquire dental imaging data at unprecedented quantity and quality [1]. For example, many research efforts have focused on actual imaging data acquisition, image analysis tools, or 3D reconstructions over computed tomography slice stacks. While cariologists now can collect massive amounts of high-resolution imaging data, they still lack efficient and intuitive means to investigatively explore data at scale and pose domain-specific questions. The ever increasing size of imaging data has suggested our re-examination of image analysis and storage in a "Big Data" context. Furthermore, when proceeding to the next step we have an even more challenging task: to allow cariologists to perform investigative caries analysis over a large number of segmentations and derived dental structures.

Our task in this paper is to present such a visual analysis tool for cariologists to conduct investigative carious lesion studies over large collections of dental imaging datasets. Our major contribution is on the interactivity and effective integration of techniques from MapReduce-based fast carious lesion assessment, data-driven distributed storage for visualization archive, and template-assisted visual computing interfaces that are well-suited for various carious lesion assessments.

*e-mail:gruan@indiana.edu
†e-mail:huizhang@iu.edu (corresponding author)

IEEE Symposium on Large Data Analysis and Visualization 2014
October 9–10, Paris, France
978-1-4799-5215-1/14/$31.00 ©2014 IEEE

2 MOTIVATION

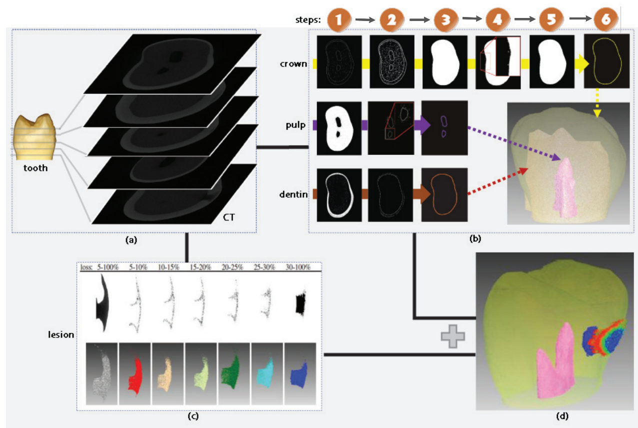

Figure 1: (a) Dental CT. (b) Examples of image analysis and derived structures. The first row: a series of image segmentation techniques used to extract tooth crown — ① binary gradient mask $\mapsto$ ② dilated gradient mask $\mapsto$ ③ binary image with filled holes $\mapsto$ ④ cleared border image $\mapsto$ ⑤ segmented image with object smoothened $\mapsto$ ⑥ extracted dental crown contour. The second row: extracting dental pulp — ① binary image with thresholding $\mapsto$ ② extracted interior holes $\mapsto$ ③ extracted dental pulp contour with removal of small objects. The third row: extracting dentin surface — ① binary image with thresholding $\mapsto$ ② extracted interior hole $\mapsto$ ③ extracted dentin surface with removal of small objects. (c) Carious lesion segmentations at various threshold levels: corresponding to each threshold level, pixels with the right amount of gray-scale loss are considered to be part of the carious lesions. (d) Visual analysis for carious lesion assessment. When addressing various threshold levels used in the carious lesion detection, volumetric models are assigned a surface color keyed to their levels of mineral content loss.

Assessing caries activity has traditionally been limited to evaluating patients *live* in the chair and by subjective clinical methods (i.e., visual or tactile inspection). By combining new imaging technologies with image analysis and visual computing methods, which accurately detect lesions over cross-sectional images, we can provide visual analysis tools that can in principle improve on carious lesion assessment. This improvement is possible because, for example, dental scans (e.g., μCT) can provide low-level details over each single slice, however, it is nearly impossible for cariologists to perform useful studies by manually checking a large number of images slice by slice. A visual analysis tool, where image-based "micro"-level information is turned into "macro"-level volumetric model, can greatly reduce the necessary effort for cariologists to adopt computational methods in carious lesion assessment. We can therefore use the same methods to help us understand and analyze the much more complicated cases with multiple specimens and variables of interest (which arise naturally in longitudinal experiments and mineral content distribution studies.) If geometry-based assessment is added to describe the quantitative measures of our derived dental variables (e.g., the area or volumetric size of a lesion, or the distance between a lesion and the dentin), one can pose domain-specific questions without necessarily having to use visual and tactile inspections at all; when used in combination with statistical analysis in the format of curve/bar charts, the tool can enable cariologists to base their reporting on derived information from a large number of dental images by providing only the most relevant methods in a clear way.

3D Example. To create a visual representation of a human tooth, we typically obtain dental CT scans out of each specimen (see e.g., Fig. 1(a)), extract contours that represent various types of dental structures, and construct a standard 3D computer graphics representation of the result (see e.g., the extraction and reconstruction of dental crown, pulp and dentin surface in Fig. 1(b)). In this way the 3D information can be generated for various dental structures. We may then, in effect, depend on the second order effects of our practical experience with 3D structures to reconstruct a 3D tooth structure in our mind. However, the 3D models for carious lesions that we are considering can be very interesting: on one hand, they may progress or reverse over time; on the other hand, cariologists often report their findings by checking carious lesions corresponding to various threshold levels. Fig. 1(c) shows an array of carious lesion segmentations corresponding to mineral content levels between 5% loss and 30% loss (by 5% increment), with different surface color keyed to the level being studied. The contours and regions of our interest, the logical series of segmentation steps, and some of the problems they induce are as follows:

- *Dental Crown* — Fully reconstructing dental crown is useful when lesions' 3-dimensional locations need to be analyzed and can provide an accurate way to align and compare lesions on 2D surface images with those on 3D reconstructed images. Dental crown is well visible with great contrast from background, and can be generally detected with the *edge* and *sobel* operators. The contour segmentation starts with an initially segmented dental crown contained in a binary mask (step ①), with lines of high contrast in the image. These lines do not quite delineate the outline of the dental crown. Compared to the original image, we see gaps in the lines surrounding the object in the gradient mask. These linear gaps can be removed if the *sobel* image is dilated using linear structuring elements. To fulfill this, we create vertical structuring elements followed by horizontal structuring elements. The dilated gradient mask shows the contour of dental crown quite nicely, except holes in the interior (step ②). The segmentation can be improved by filling the holes (step ③), removing connected objects on border (step ④), and smoothening (step ⑤) respectively (see the key steps in the 1^{st} row of Fig. 1(b) with step ⑥ representing the extracted tooth crown contour).

- *Dental Pulp* — In the 2^{nd} row of Fig. 1(b) are the major steps to extract dental pulp, which is a variable of interest in many measurements, e.g., the minimum distance from the dentin to lesions in vivo. We observe that dental crown, pulp and the background image have observable gray-scale contrast and through thresholding we can transform the gray-scale image into a binary one to facilitate contour segmentation. The typical procedure starts with the binary image after thresholding (step ①). Next we extract the interior holes (i.e. tooth pulp) within the tooth crown (step ②). Remained residues are considered components with no reasonable pixel sizes and are further removed. The resultant dental pulp segmentation is shown as step ③ in the 2^{nd} row of Fig. 1(b).

- *Dentin Surface* — Dentin surface is yet another variable of interest and accurate construction of dentin surface can quantitatively reveal how lesions progress (e.g., from staying in the enamel to extending into the dentin). The $3rd$ row of Fig. 1(b) illustrates the processes to extract the dentin surface in a non-interrupted scenario, which is similar to dental pulp segmentation.

- *Carious Lesion* — Leveraging the above segmentations, carious lesions can be fully segmented. Lesions in the enamel typically correspond to pixels with gray-scale 5% less than the surrounding healthy enamel, and those extending into dentin 5% less than surrounding healthy dentin. By thresholding pixels in the enamel and in the dentin respectively, we can extract carious lesions that are in the enamel and in the dentin.

- *Putting All Together* — Finally, as suggested in Fig. 1(d), we can explore a more complex 3D representation that reveals all quantitative information and relative spatial relationship of dental concepts and structures: this is the very first step that we shall present in this paper – visual computing methods that can be used to fully assess carious lesions and their activities.

Data Challenge. The visual computing methods in principle are sufficient to allow us to explore carious lesions in fully constructed 3-dimensional tooth structures. In practice, to navigate from observation to findings and to the most appropriate therapeutic solution in a situation with multiple variables and phenomena of interest, cariologists need to draw on scientific knowledge from very large dental imaging data. This naturally arise in one such longitudinal study shown in Fig. 2(a), where each of a collection of 95 tooth specimens were adopted in a multi-phase dental study using a remineralization[1]/demineralization[2] model (e.g., 3-day demineralization in the 1^{st} phase, 6-day remineralization in the 2^{nd} phase, then 4-day remineralization in the 3^{rd} phase, and finally 6-day remineralization in the 4^{th} phase). Furthermore, dental imaging data in oral health research are often collected from multiple modalities which adds to the visualization and analysis difficulty. The example dataset used to illustrate our work was collected from three imaging systems for various purposes of lesion assessment:

- *μCT slices* — acquired over four phases in the longitudinal study, using SkyScan® Microtomography system with a 2.25 μm spatial resolution; around $4,000$ CT scans are obtained from each tooth specimen at each phase.

- *Quantitative Light-Induced Fluorescence (QLF) surface images* — recorded from a QLF camera at a speed of $1 sec\ per$ image; 16 images were obtained from each tooth specimen at each phase.

- *Infrared thermal surface images* — recorded from an infrared camera at a speed of $\frac{1}{113} sec\ per$ image; 400 images were obtained from each tooth specimen at each phase.

How to perform investigative analysis and effectively discover new knowledge from rich and large dental imaging data poses a compelling challenge. Assume that there are N specimens assessed in a longitudinal study over P time points, and K properties can be extracted from each of M acquired dental images, we will have $N \times P \times K \times M$ temporal variables. With the advance of image acquisition techniques and more fine-grained assessment of carious lesions, we anticipate N, P, K and M to increase continuously. As a concrete example, we currently acquire $4,000$ μCT images from a single tooth specimen at each time point of the longitudinal study. Using a desktop computer with an Intel Pentium-4 3.20 GHz processor, it takes approximately 3.7 seconds to fully process a single image, i.e. extracting all structures of interest, and 4.1 hours to process the whole image stack (i.e. $4,000$ images). Moreover, the storage cost is nontrivial and it takes roughly 1.2 GB (i.e.

[1]dental treatments to help restore the mineral content of tooth tissues

[2]a process to reduce the content of mineral substances in tooth issues

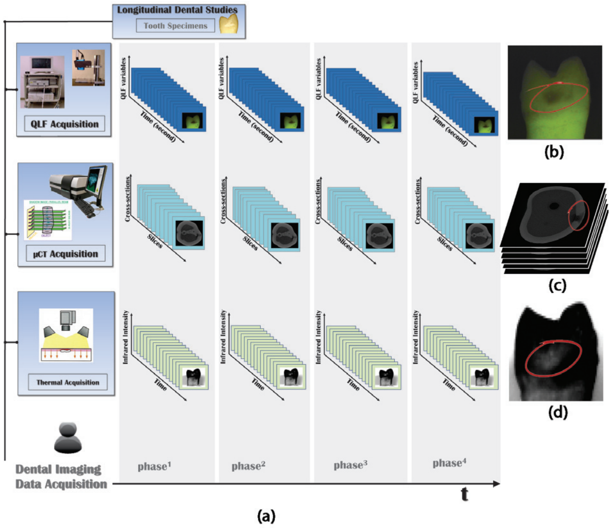

Figure 2: One typical experimental setting used in our longitudinal caries lesion studies. (a) Three types of dental imaging data were acquired during a 4-phase dental experiment. (b) Quantitative light-induced fluorescence (*QLF*) images, (c) μCT slices, and (d) Infrared thermal images. μCT technology enables very detailed images of tooth structures to be taken in slices, *QLF* and infrared thermal images indicate dental caries on single surface images. The cross-sections of demineralized enamel on μCT slices show an observable gray-scale difference from that of healthy enamels. Similarly, carious lesions appear dark on *QLF* and white on infrared thermal images when viewed. This is based on the principle that a demineralized tissue limits the penetration of light due to excessive scattering of photons entering the lesion with consequent limitation to the change of a photon being absorbed and fluorescence or infrared remitted.

300 KB $*$ 4,000) to store segmented images for a single specimen. Hence, distributed processing and storage become a requirement rather than an option.

Visual Analysis of Large Dental Imaging Data. Although a handful of dental computing and computational caries assessment efforts exist([6, 11, 12, 15]), they do not offer explicit guidelines on how to link user-driven interactive visual analysis of large data volumes with data-driven scalable image processing. Moreover, caries assessment tools should allow one to perform investigative analysis and pose domain-specific questions, which has not been a focus in previous efforts. Based on our analysis of these challenges, we propose a set of design principles that can guide how to implement a visual analysis tool that can transform large dental imaging data to knowledge discovery:

1. *Data-driven large-scale processing* — Given the scale of the data and the time complexity of image segmentation, scalable image processing needs to be applied when extracting various structures of interest from raw datasets. Furthermore, cariologists' online analytics are mostly *read-only* rendering and visualization operations applied to extracted structures, thus the large-scale processing can be conducted offline, triggered by data acquisition. In addition, cariologists are domain scientists rather than experts in distributed systems, therefore we would like to design the scalable image processing workflow to be automated and totally transparent to cariologists. Derived structures are made available for analytics once processed and will populate a data list widget in client user interfaces.

2. *User-driven investigative analysis* — To understand the carious process and to create effective caries treatments, cariologists desire an iterative knowledge discovery process. The user interface should provision clean exploratory interfaces for cariologists to analyze and correlate carious lesion phenomena. For example, cariologists often want to perform lesion analyses through easy creation of various assessments, develop hypotheses from assessment reports in both visual and statistical formats, pose domain-specific questions by issuing visual queries and etc.

3. *Loosely coupled server-client architecture* — To make our system easily extendable, modules should interface with each other in a loosely-coupled fashion and through well defined interfaces. We like to think of our design intuitively as existing in a *server-and-client* structure. For instance, the service agreement between the data store and visual-analysis server, and that between the server and client interface, are all defined as *RESTful* web services. This gives us flexibility on module update, e.g., switching from a relational database to a NoSQL store.

3 System Overview

Our system consists of two main parts (see Fig. 3): data driven scalable image processing is offline and triggered by imaging data acquisition, while investigative analysis module is online and user-driven at run time. All services provisioned by modules are exposed in the form of *RESTful* web services. We implement a visual analysis tool (DenVis) in a client/server architecture for cariologists to perform various carious lesion assessments.

Within the scalable image processing and data store module, μCT imaging data are processed once acquired with parallel MapReduce [4] tasks (see ①) and all resultant data, e.g., segmented images and derived structures (see ②) are ingested into Mongo DB [9], a distributed NoSQL store (see ③). *QLF* and infrared thermal images do not need further processing and are directly ingested into MongoDB. Upon completion of processing, data are made available to user-driven module for online analysis (see ④).

At run time, the user initializes data requests by interacting with DenVis client, which allows one to conduct various analysis, e.g., template-based assessment, interactive visual analysis, longitudinal evaluation, visual queries and etc. To avoid the performance bottleneck brought by a single server, DenVis server employs multiple hosts to distribute the workload evenly across machines through Domain Name System (DNS) and this load-balancing procedure is totally transparent to both DenVis server and client. The typical interaction flow between end users and DenVis is as follows: step ① — the user opens DenVis client, which upon start first queries DenVis server through the *listdata* web service to list available datasets, DenVis server in turn asks MongoDB for metadata of datasets and returns results to DenVis client, which in turn lists available datasets in a *data list* widget; step ② — the user selects a dataset for analysis and issues a request to DenVis server, e.g., finding the distance between the carious lesion and the dental pulp; step ③ — DNS (not shown in Fig. 3) directs the request to one of the server hosts that share the same service name of DenVis web services. The chosen server host reads geometries of the carious lesion and the dental pulp from mongoDB, calculates the distance between the two and generates the visualization and statistics; and step ④ — the server host sends the visualization/statistics back to the user and results are displayed on DenVis client.

4 Implementation Models and Methods

We now turn to our main objective, which is to leverage MapReduce for large-scale dental image processing and volumetric model generation, and then how these derived structures together with various types of surface images are explored and analyzed in investigative carious lesion assessment using DenVis software. Our fundamental techniques are based on a wide variety of prior art, including the use of MapReduce for simplifying data processing on clusters [4], the computational analysis methods for caries detection on dental CT and surface images (e.g., QLF) [1], and variants on visual query and template-assisted interfaces in information visualization systems [2, 7].

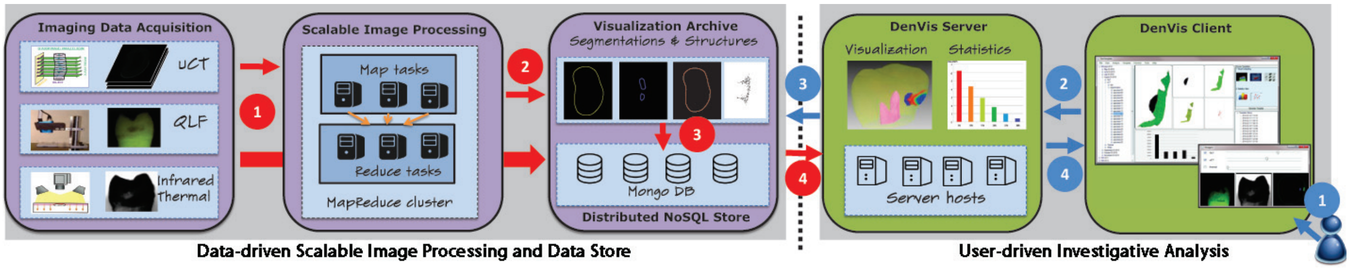

Figure 3: System overview. DenVis consists of two main parts: data driven or offline large-scale image processing modules (left, purple) are triggered by imaging data acquisition, while visualization/user-driven or online investigative analysis modules (right, green) are active at run time. Dental imaging data are processed in parallel with MapReduce and all generated data (e.g., segmented images and derived structures) are stored in Mongo DB, a distributed NoSQL store. At run time, the user initializes data requests by interacting with DenVis client, which enables the user to conduct template-based assessment, interactive visual analysis, longitudinal evaluation, visual queries, etc. (cf. Fig. 5). Data flow is denoted as numbered red circles while interaction flow as numbered blue circles.

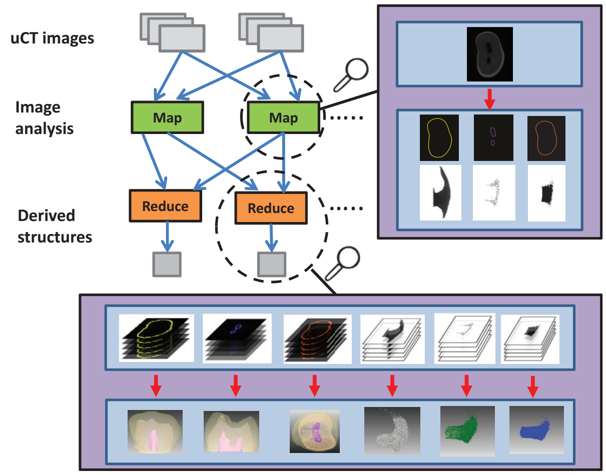

Figure 4: The MapReduce task that extracts structures of interest, e.g., dental crown and pulp, dentin surface, and carious lesions of different mineral lost levels. In Map phase, for each μCT image, image segmentations are performed to extract contours of dental crown and pulp, and dentin surface, and thresholding is performed to identify carious lesions of different mineral loss levels (e.g. by 5% interval). In Reduce phase, contours and lesion areas belonging to the same image stack are "stacked" together to produce geometries. We note that extracted contours/lesion areas need to be "stacked" following the correct order as how they were sliced in the original specimens. We modify V_i as a two-element tuple *(seqNum, segmentedVariable)* where *seqNum* specifies the order of the μCT image within its image stack and can be determined from the image filename. By this means, we can enforce the correct order in the reduce phase.

4.1 Scalable Image Processing and Geometry Generation

In the data driven modules, dental imaging data are pushed to a MapReduce cluster for immediate processing once acquired. The basic idea is that segmentation algorithms operate on each image of a specimen image stack independently, and a strucuture/geometry derived from a single image stack is independent of other stacks. Hence, the whole process can easily fit into MapReduce's loosely-coupled parallel framework, which centers around two user-defined functions that represent a problem domain:

- $Map(D_i) \rightarrow \text{list}(K_i, V_i)$
- $Reduce(K_i, \text{list}(V_i)) \rightarrow \text{list}(V_f)$

In simple terms, *Map* is a function which, given an input data value D_i, produces a list of an arbitrary number of key/value pairs. *Reduce* is a function which, given a single key and a list of associated values, produces a list of final output values V_f. Given a *Map* and a *Reduce* function and an input dataset D, applying MapReduce to the entire dataset follows a general pattern:

- Divide the dataset into individual data values D_i.
- Apply $Map(D_i)$ to each value, producing many lists of key/value pairs $list(K_i, V_i)$.
- Gather all data produced by *Map* operations and group them by key K_i, producing lists of associated values $list(V_i)$.
- Apply $Reduce(K_i, list(V_i))$ to each key K_i and associated list of values $list(V_i)$.

A lesion assessment based on analyzing μCT images can therefore be expressed with the MapReduce programming model. Table 1 describes the mapping of our data types in a longitudinal lesion assessment to those parameters in MapReduce's model. We use *structureID*, a 2-element tuple of form *(imageStackID, structureTypeID)*, to identify a structure of interest in the longitudinal study. *imageStackID* recognizes which specimen this extraction belongs to, and *structureTypeID* recognizes which dental structure (e.g., crown, pulp, dentin, or lesion) this extraction represents. Now that we know *Map* is followed by a grouping/shuffling phase to group intermediate results by key, we can leverage this process to generate a set of distinct structures of interest, e.g., the geometry of the tooth crown of a specimen at a specified phase in a longitudinal study. If *Map* generates STRUCTUREIDs as intermediate keys K_i, the grouping phase that follows can assemble all values associated with each distinct structure as an input to *Reduce*. Likewise, since *Reduce* produces its final output by aggregating all values V_i associated with a key K_i, we can choose V_i to represent a contour/lesion area segmented from a single μCT image which is then "stacked" in *Reduce* to produce a final geometry for each structure.

Table 1: Data types in MapReduce tasks.

D	Whole μCT imaging data
D_i	A single μCT image
K_i	STRUCTUREID
V_i	Segmented contours/lesion regions from a μCT image
V_f	(STRUCTUREID, geometry) pair

Creating assessment on a μCT image dataset using a MapReduce framework would result in following steps and Fig. 4 shows an overview of this process:

- Apply our *Map* function to an image, producing key/value pairs of structureIDs and segmented variables.
- Gather all intermediate key/value pairs produced by *Map* and group by key. Since we use structureIDs as keys, the result of this step will be a each distinct structure paired with a list of segmented variables (i.e. contours or lesion areas belonging to the same image stack).

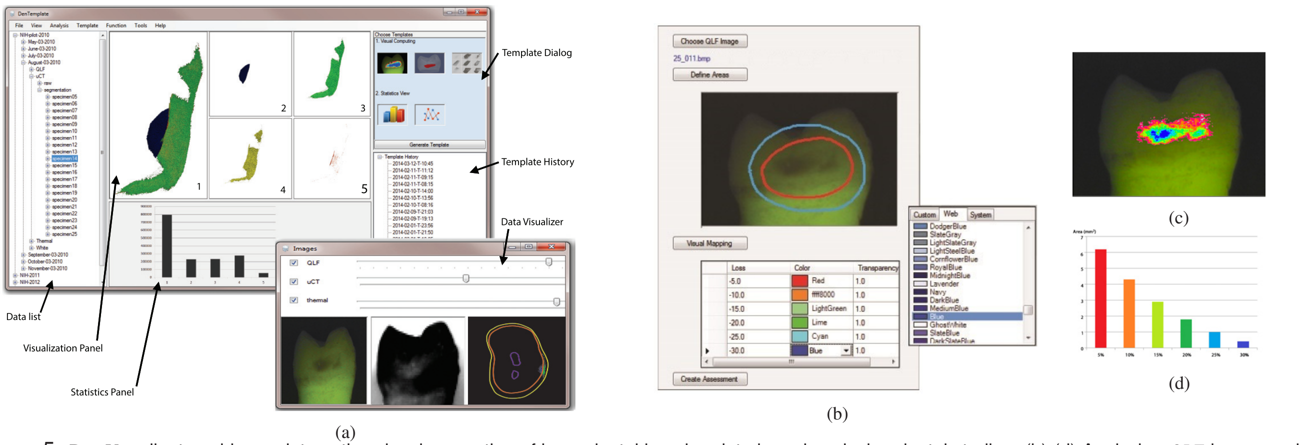

Figure 5: DENVIS client enables an interactive visual computing of large dental imaging data in carious lesion dental studies. (b)-(d) Analyzing *QLF* images using DENVIS client's image computing template in three simple steps: ① load data source $\mapsto$ ② define two patches $\mapsto$ ③ specify view mapping. (c)(d) The generated visual and statistical views as assessment results.

- Apply *Reduce* to each grouped result. Reduce will stack segmented variables of each structure in correct order to produce ordered pairs containing each structureID and its associated geometry.

Since MapReduce excels at processing a small number of large files instead of a large number of small files. In the implementation we pack the μCT images into a small number of large binary files before actually being processed. This preprocessing step is performed in parallel as well through a MapReduce task. We note that only μCT images need to be processed; *QLF* and infrared thermal images are directly ingested into MongoDB.

4.2 Visualization Archive

The large-scale imaging data not only pose challenges to computation but also to storage. We employ MongoDB, a distributed NoSQL data store to host segmented μCT images, *QLF* and infrared thermal images, and derived structures. Compared to traditional relational databases, NoSQL stores provide flexible schema design and better fault tolerance and performance [3]. For the concerned scalability, we can add storage nodes to MongoDB to accommodate more imaging data without stopping the online service. We wrap MongoDB's low-level data access APIs as *RESTful* web services so that data can be ingested into and read from it easily.

4.3 Quantitative and Interactive Analyses Using DENVIS

The user-driven investigative analysis module is based on a client/server architecture. As shown in Fig. 5, DENVIS client has two major display components in the frontend: an imaging data viewer and a visual computing dashboard. The data viewer is a digital media player that allows users to access raw images and their segmentations. The visual computing dashboard is the main tool that allows one to create caries lesion assessment and explore the resultant visual and statistical representations. When users adjust the queries, sensitivities, and specificities for an active assessment, the two display windows will be synchronized. There are three display areas in our visual computing dashboard. Datasets processed by data-driven modules are displayed in a tree widget in the upper left corner and are organized in a hierarchy following the temporal order in the longitudinal study. Users can select which ones they will import into the computing environment, create and combine various assessments using the visual templates provided in the upper right corner. Occupying the central area we have a visualization panel at the top and a statistics viewer at the bottom. The statistics viewer is the place to display statistical results, often in the format of bar or line charts. The visualization panel generates image or volumetric representations that correspond to the desired assessments being performed. Within the visualization panel, multiple views can be created if needed, e.g., when generating an array of 3D images as the representation of a longitudinal assessment. In this way one can generate various visual analyses to ask domain-specific questions, such as the mineral distributions of a tooth specimen, or the 3D-time representations of dynamic carious lesion activities. DENVIS client also provides various interactive functions to allow easy queries and measurements of areas, volumes, and distances among variables of interest, which we will discuss more later. Upon receiving DENVIS client's requests, DENVIS server retrieves corresponding data from MongoDB, performs necessary computation for visualization and statistics, and finally sends results back to the client.

4.3.1 Carious Lesion Assessment By Templates

Although cariologists may have clear domain-specific questions regarding their data, they can be novices when it comes to visualization and visual analysis, and naive as to how machine computation can support their investigative analysis needs. In investigative analysis, cariologists need to explore multiple visual data representations from diverse data sources and apply iterative analysis specifications in established assessment procedures. We feel interactive visual templates are very suitable for such scenarios. A visual template can be implemented to incorporate the sequence of steps to form valid caries lesion assessments with continuous visual feedback as the final outcome of users' choices.

***QLF* Analysis by Image Assessment Templates.** *QLF* is a dental diagnostic tool for in-vivo and in-vitro quantitative assessment of dental caries lesions, dental plaque, bacteria activity, calculus, staining, and tooth whitening. Once a *QLF* image is captured, the next stage is to analyze the lesions and produce a quantitative assessment of the demineralization status of the tooth. The basic idea is to manually use a "patch" to define an area of *healthy* enamel and identify pixels corresponding to a threshold of fluorescence loss. These pixels appear dark when being viewed, and they form the lesion of interest in carious lesion studies. Now we can begin to see how to exploit an interactive visual template in *QLF* assessment to facilitate this visual computing process. The logical series of steps in using image computing template are as follows:

- *Interactively create two patches to define healthy and lesion enamels.* The interactive process involves using two user-sketched patches to define areas of healthy enamel around the lesion of interest (see e.g., the blue patch for healthy enamel and the red for lesion in Fig. 5(b)).

- *Reconstruct surface image.* DENVIS server uses the pixel values of the healthy enamel to reconstruct the surface of the tooth and then subtracts those pixels which are considered to be the lesion.

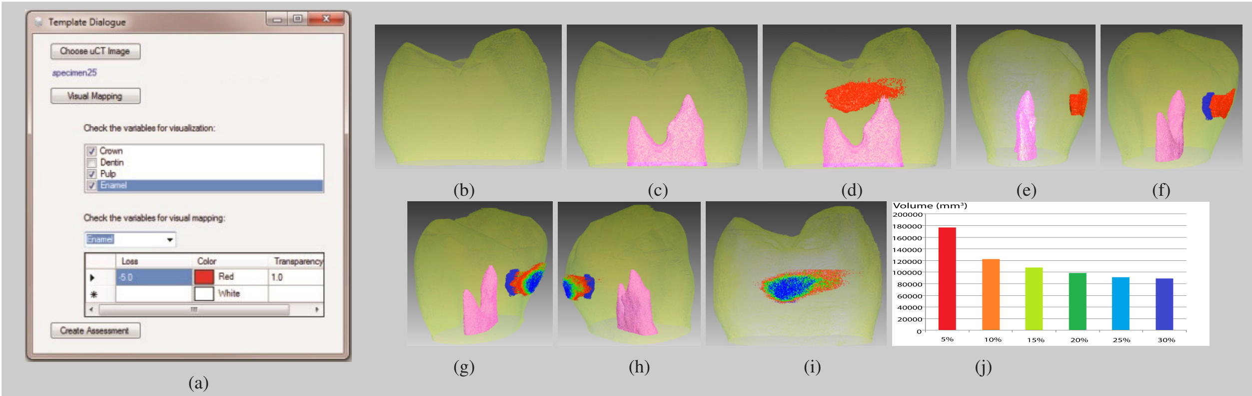

Figure 6: Typical interaction flow for generating multi-threshold volumetric assessments. (a) The volumetric assessment template allows users to choose data source, select static visual structure, and define visual mapping for a variety of threshold levels. (b)-(e) The user enables visual structures in an order of dental crown, dental pulp, and a reconstructed lesion with a 5% gray-scale loss. (f)-(i) More threshold levels are activated to show the mineral distributions in the lesion. (j) Barplot representation of the mineral distribution corresponding to threshold levels between 5% loss and 30% loss by 5% increment.

- *Generate quantitative analysis using threshold levels.* This is controlled by a series of thresholds of fluorescence loss specified by the user for assessments of various purposes. For instance, when the threshold is generally set to 5%, this means that all pixels with a loss of fluorescence greater than 5% of the average healthy value will be considered to be part of the lesion. With the visual template, one can specify a color map where each individual threshold level is keyed to a unique color representation.

- *Output visual and statistical results.* Once the pixels have been assigned “healthy enamel” or “lesion”, DENVIS server then outputs the visual representation by overwriting the pixels in the colors chosen by the user in the visual mapping step (see e.g., Fig. 5(c)) and sends the result back to client. Area measurements of the lesions are calculated in mm^2, corresponding to the whole sequence of threshold levels. See Fig. 5(d) for an example of generated statistical view using the widely-used barplot, which indicates a decrement of lesion areas when corresponding to greater fluorescence loss in the image.

Using image computing template, one can load in a *QLF* image of a tooth that has been captured, and generate a quantitative visual analysis of the demineralization status with just a few mouse clicks.

Exploit Volumetric Assessment Templates. Oftentimes cariologists need to go beyond a tooth surface image, and look into the tooth structures. This can be done with volumetric assessment template. DENVIS client has been designed and developed under the assumption that cariologists may have no previous knowledge of what visualization methods and representations are possible for different types of segmentations. The volumetric assessment templates aim to make the sequence of steps fairly simple to follow. Volumetric assessment is based on computing a large number of images and incorporating iterative analysis specifications. The process can be greatly facilitated with a visual template. We use Fig. 6 to walk through the major steps and screen images to show the interaction flow for performing volumetric assessments. The user first loads the data sources on the template interface (cf. Fig. 6(a)), and selects the static structures (e.g., tooth crown, pulp, dentin surface, etc.) to be visualized from a list (cf. Figs. 6(b)-(d)). In visual mapping control, the user assigns each volumetric model a surface color keyed to its threshold level of mineral content loss. Reconstructed volumetric models that correspond to a variety of threshold levels are added to the 3D scene (cf. Figs. 6(e)-(i)). The visualization (cf. Fig. 6(i)) and statistical results (cf. Fig. 6(j)) can provide a more accurate model for us to understand the results and phenomena generated from the *QLF* assessment.

Composing Complex Assessment with Multi-View Templates. There are many scenarios where we desire to create assessments by applying the same templates across multiple data sources and hope to examine all the results at a time. Structuring and bringing all the needed assessment results to cariologists have become highly challenging in such scenarios. To make easy creation of such batch-style processes, we complement our visual computing environment with a multi-view template. The multi-view template provides no new technical functionalities, but allows one to link one or more assessment templates with multiple data sources. Fig. 7(a) gives an example case scenario where such multi-view visual analyses are preferred. The assessment is concerned with a longitudinal dental study to examine the dynamic carious lesion activities in a 5-phase demineralization/remineralization experiment. A volumetric assessment template addressing three threshold levels is iteratively applied to 5 data sources representing time-resolved imaging data. The multi-view representation provides a 3D-time series that intuitively depicts how tooth minerals with three different levels of loss are evolving over time.

4.3.2 Carious Lesion Assessment by Visual Analysis

During the life cycle of DENVIS client, user needs may evolve and various types of assessment results may need to be linked for correlation analysis. Especially given novice users’ tendency for iterative analysis specifications and interactive visualization tasks, DENVIS client has several features for user friendly customization and query.

Volumetric Assessment with Selected Slices. At times we need to compare sub-regions across a volumetric model with their original image datasets or segmentations. With DENVIS, one can drag and drop a range of slices from data list panel to the visualization panel to force a volumetric assessment with selected slices. As shown in Fig. 7(b), this allows one to “clip” a volume and connect its sub-regions with the imaging datasets.

Linking Image and Volumetric Assessments. As an additional interface element, DENVIS automatically saves the successfully applied assessments in the format of a list of drag-and-drop templates. One value of doing this is to allow users’ exploration of alternatives without risking losing existing experimentations; all actions are reversible by tracing back to one of the previous assessments recorded (cf. Fig. 5). Another value of doing this is to allow users to link visualizations generated from different types of assessment templates. For example, one can drag-and-drop a *QLF* analysis result to link with the volumetric assessment corresponding to the same specimen (Fig. 7(c)).

Quantitative Assessment with Visual Queries. Many quantitative assessment tasks involve the relationship among mul-

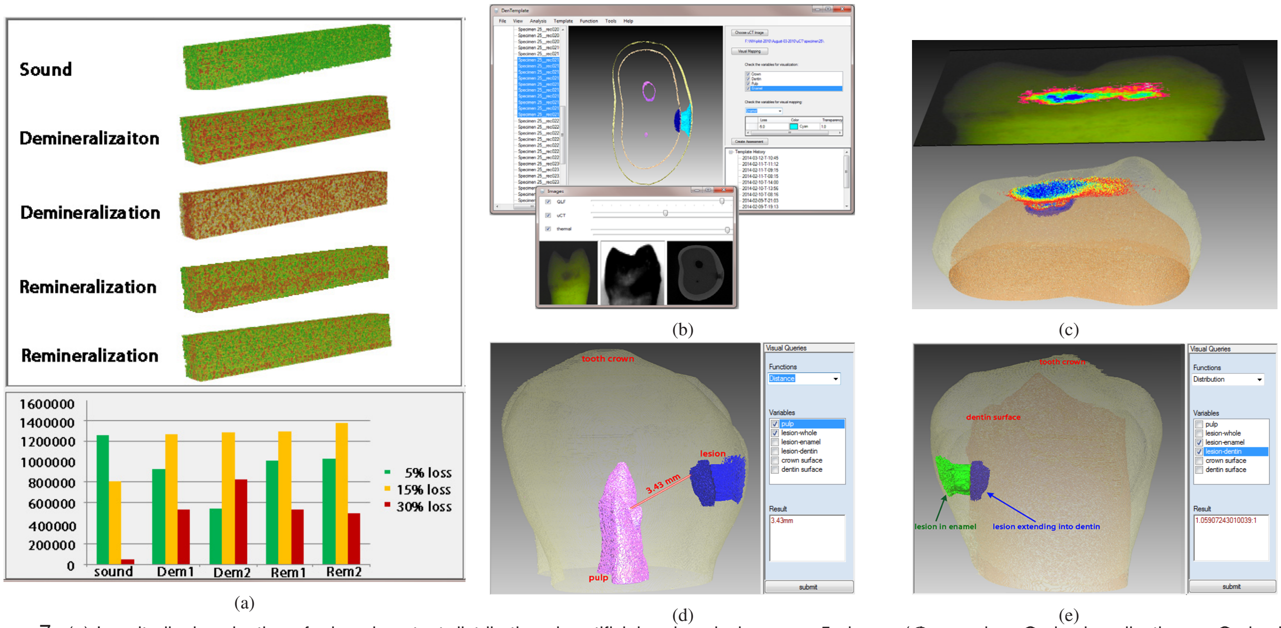

Figure 7: (a) Longitudinal evaluation of mineral content distributions in artificial carious lesions over 5 phases (① sound ↦ ② demineralization ↦ ③ demineralization ↦ ④ remineralization ↦ ⑤ remineralization.) Volumes corresponding to 5% mineral content loss are rendered in green, 15% in yellow, and 30% in red. Visualization indicates that tooth issues lost most mineral content at phase 3 (most pixels in red), with two deimineralization treatments in a row. After two remineralization treatments, the tooth structure at phase 5 has been restored to a status comparable to phase 1. (b) Customizing a clip of volumetric reconstruction by dropping only 10 slices of images. (c) Drag and drop a *QLF* image assessment and a μCT volumetric assessment into an integrated view. (d)(e) Quantitative assessment by visual queries. (d) Finding the distance between the carious lesion and the dental pulp. (e) Finding the ratio of carious lesions in enamel and extending into dentin.

tiple variables. Therefore it is natural to provide users with a variable-driven query interface to complement the visualizations. For instance, Fig. 7(d) shows the active query interface for one to calculate the distance between the reconstructed carious lesion and the dental pulp. To estimate the degree of encroachment of caries on the pulp, the minimum distance between the lesion and the dental pulp has to be measured. This information can help to identify cases with the potential for maintaining pulp vitality and monitor treatment success of indirect pulp capping, the procedure of covering the nearly exposed pulp with calcium hydroxide to stimulate formation of secondary dentin. X-ray imaging has been routinely used in clinical practice for diagnosis of carious lesions and estimation of their extension. However, the technique only provides two-dimensional projections limiting the information regarding location and size of the lesion and its distance to the pulp. Fig. 7(e) is the typical screen image for one to query the volumetric ratio between lesions in enamel (highlighted in green) and those extending into dentin (in blue).

5 Evaluation

Parallel Efficiency. Though image processing is conducted offline, we still want this process to be efficient so that acquired data can be made immediately available to the frontend for analysis. Here we focus on the parallel efficiency of our MapReduce-based approach. We use XSEDE resource Stampede cluster [13] as the experimental platform, where each compute node has two 8-core 2.7 GHz Intel Xeon E5-2680 processors and 32GB DDR3 memory. We configure 8 map slots and 4 reduce slots on each Hadoop node. The experiments are conducted under two settings: (1) increase cluster size while fixing the dataset size, which reveals the processing *speedup* by utilizing more computing resources; and (2) increase dataset size while fixing the cluster size, which reveals the *classical algorithm complexity*, i.e., the relationship between the increment of data size and that of corresponding processing time (e.g. linearly or exponentially). Fig. 8(a) shows the speedup test, with data size fixed to be 40,000 (80-stack * 500-image per stack) μCT images. We obtain superlinear speedup when the cluster is scaled from overly utilized 8 nodes to 16 nodes, due to more map/reduce tasks running in parallel. The speedup begins to deteriorate when cluster size goes beyond 32 nodes, due to low cluster utilization and increased communication cost. Fig. 8(b) shows the algorithm complexity test by varying the dataset size (i.e. from 10,000 to 50,000 images), with Hadoop cluster fixed to be 32 nodes. We observe that the cluster has serial linear asymptotic behavior.

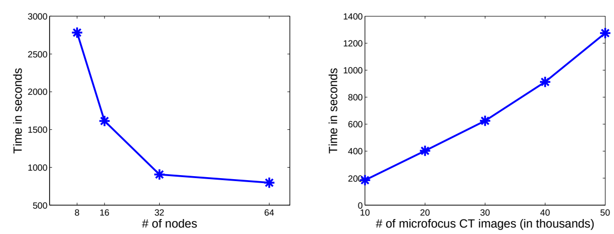

(a) (b)

Figure 8: Evaluation of parallel efficiency. (a) Speedup test by varying MapReduce cluster size with # of μCT images fixed to be 40,000. (b) Algorithm complexity test by varying # of μCT images with MapReduce cluster size fixed to be 32-node.

Response Time. Interactive analytics has stringent requirement on the response time to ensure good usability. Table 2 shows the average transfer size and response time of various data types for a single specimen at one phase. Response time measures the turnaround time between client issuing the request and receiving the response. The DenVis client, server and MongoDB in this test co-locate within the same local area network (LAN) with gigabit Ethernet connection. By transforming image-based "micro"-level information into "macro"-level volumetric model, we can clearly see from Table 2 that the data size has been reduced greatly and hence much faster response time is achieved.

User Evaluation. The DenVis has seen early use in dental domain-specific problems. We note that there are two specific case scenarios where our collaborators have used our system to perform carious lesion studies and produce visualization and analysis results for their domain publications. The first scenario is concerned

Table 2: Average transfer size and response time of various data types for a single specimen at one phase. Derived 3D geometries are dental crown and pulp, and dentin surface for the specimen (cf. Fig. 1(b)), lesions corresponding to a variety of mineral content loss levels for the mineral distribution study (cf. Fig. 1(c)), and specimen and mineral distribution combined for the longitudinal study (cf. Fig. 7(a)), respectively.

Data type		Data transfer size (MB)	Response time (second)
Derived 3D geometries	Single specimen	2.9	0.031
	Mineral Distribution study	4.2	0.038
	Longitudinal study	7.1	0.065
Raw imaging data	μCT (4,000 slices)	1,172	13.8
	QLF (16 images)	14.3	0.18
	Infrared thermal (400 images)	10.9	0.11

with quantitative assessment of dynamic carious lesion activities in a longitudinal 5-phase demineralization/remineralization dental study. For longitudinal evaluation of mineral content changes, the μCT images of sound, demineralized and remineralized enamel and that of the phantom are acquired. DENVIS is used to generate visual analysis to study the dynamic carious lesion activities in this longitudinal study (see one such analysis output in Fig. 7(a)).

The second study uses our system to visualize and correlate *QLF* and μCT images of white-spot Lesions. The objective of this study is whether multiple severity levels (thresholds) of white-spot lesion determined by a fluorescence technique would be corresponding to those of μCT images. Visualization is performed by 5% increment from 70% to 95% of threshold levels using DENVIS. Visual comparisons with μCT images are performed in two ways: 1) with original fluorescence images; and 2) with multiple threshold fluorescence images. The study observed that the shape of lesions in original fluorescence images corresponds well with μCT images. Within the limitations of this study, multiple severity levels (thresholds) of white-spot lesion determined by fluorescence technique correspond to those of μCT images. Fig. 7(c) is a typical screen image to show how *QLF* and μCT images of white-spot lesions are being compared and correlated by cariologists.

6 RELATED WORK AND OUR FUTURE DIRECTIONS

The idea of using imaging technologies as adjunct to clinical visual or tactile examinations for caries diagnosis has greatly facilitated oral health-care in dentistry [5]. Most previous researches focus on the actual image acquisition and subsequent three-dimensional dental visualization. For example, Farman's work reviewed the historic of digital imaging in dentistry to outline the fundamental issues related to digital imaging modalities [1]. Gehleitner illustrated the typical appearance of dental related diseases of the jaws with dental CT, and suggested where dental CT can serve as an addition to conventional imaging methods in dental radiology [6]. Tymofiyeva's study assessed the feasibility of MRI of three-dimensional visualization and quantification of carious lesions. Other representative efforts include a variety of ways of semi-automatic or fully automatic analysis of dental imaging data (see e.g., Van [14], Pretty [10], and Magne [8]), and to exploit high performance computing to accelerate the automatic analysis (see e.g., Zhang [15], Ruan [11], and Smelyanskiy [12]). However, the ability for researchers to locate, analyze, and use large, complex, and diverse dental imaging data is still limited for reasons related to access to relevant software and tools, expertise, and other factors. Indeed, much less research focuses on the next step — on how to enable efficient and intuitive cariological analysis of dental datasets of larger size and more complexity. Cariologists now have huge collections of high-quality high-definition dental images of various types and their segmentations, what is really needed is an efficient means for cariologists to base their knowledge discovery on exploratory visual analysis of large dental segmentation sets, posing and answering many high-level domain-specific questions. In this paper we have introduced DENVIS, an interactive visual computing environment that allows us to freely "wander" around dental imaging data managed in a tree widget, and make discoveries by checking visualizations and analyses generated from intensive longitudinal data.

There are several directions that we want to explore in the future: (1) implement the cache and prefetch mechanism in client side to further reduce the response time. For instance, we can cache the geometries, statistical and visual query results. Moreover, based on the current images examined by the user, we can foresee and prefetch images to be used subsequently; (2) we note that segmented images can benefit from an over 99% compression ratio since only contours are left. Hence we intend to examine the tradeoff between the storage savings and the extra computation cost for compression/decompression; (3) we can refine the processing logic used in MapReduce, for example, by transferring the identified regions of interest instead of the whole segmented images from Map tasks to Reduce tasks; and (4) reimplementation of DENVIS as a web-based science gateway portal.

REFERENCES

[1] B. T. Amaechi. Emerging technologies for diagnosis of dental caries: The road so far. *Journal of Applied Physics*, 105(10):102047–102047–9, May 2009.

[2] J. Beyer, A. Al-Awami, N. Kasthuri, J. W. Lichtman, H. Pfister, and M. Hadwiger. Connectomeexplorer: Query-guided visual analysis of large volumetric neuroscience data. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2868–2877, Dec. 2013.

[3] R. Cattell. Scalable SQL and NoSQL data stores. In *ACM SIGMOD Record*, volume 39, pages 12–27, 2010.

[4] J. Dean and S. Ghemawat. MapReduce: Simplified data processing on large clusters. In *6th Symposium on Operating System Design and Implementation (OSDI'04)*, volume 37, CA, USA, December 2004.

[5] A. Farman. Fundamentals of image acquisition and processing in the digital era. *Orthodontics & craniofacial research*, 6(s1):17–22, 2003.

[6] A. Gahleitner, G. Watzek, and H. Imhof. Dental ct: imaging technique, anatomy, and pathologic conditions of the jaws. *European radiology*, 13(2):366–376, 2003.

[7] L. Grammel, M. Tory, and M.-A. Storey. How information visualization novices construct visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):943–952, Nov. 2010.

[8] P. Magne. Efficient 3d finite element analysis of dental restorative procedures using micro-ct data. *Dental Materials*, 23(5):539–548, 2007.

[9] mongoDB. `https://www.mongodb.org/`.

[10] I. Pretty, W. Edgar, and S. Higham. The erosive potential of commercially available mouthrinses on enamel as measured by quantitative light-induced fluorescence (qlf). *Journal of dentistry*, 31(5):313–319, 2003.

[11] G. Ruan, H. Zhang, and B. Plale. Exploiting mapreduce and data compression for data-intensive applications. In *Proceedings of the Conference on Extreme Science and Engineering Discovery Environment: Gateway to Discovery*, page 38. ACM, 2013.

[12] M. Smelyanskiy, D. Holmes, J. Chhugani, A. Larson, D. M. Carmean, D. Hanson, P. Dubey, K. Augustine, D. Kim, A. Kyker, et al. Mapping high-fidelity volume rendering for medical imaging to cpu, gpu and many-core architectures. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6):1563–1570, 2009.

[13] Stampede computing cluster. `https://www.tacc.utexas.edu/stampede/`.

[14] M. Van Der Veen and E. De Josselin de Jong. Application of quantitative light-induced fluorescence for assessing early caries lesions. 2004.

[15] H. Zhang, H. Li, M. J. Boyles, R. Henschel, E. K. Kohara, and M. Ando. Exploiting hpc resources for the 3d-time series analysis of caries lesion activity. In *Proceedings of the 1st Conference of the Extreme Science and Engineering Discovery Environment: Bridging from the eXtreme to the Campus and Beyond*, XSEDE '12, pages 19:1–19:8, New York, NY, USA, 2012. ACM.