Medical Image Analysis 33 (2016) 98–101

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Editorial

Large-Scale medical image analytics: Recent methodologies, applications and Future directions

Shaoting Zhang[a], Dimitris Metaxas[b,*]

[a] Department of Computer Science, University of North Carolina at Charlotte, Charlotte, NC, 28223, USA
[b] Department of Computer Science, Rutgers University, Piscataway, NJ, 08854, USA

ARTICLE INFO

Article history:
Received 1 April 2016
Revised 6 June 2016
Accepted 13 June 2016
Available online 28 June 2016

Keywords:
Medical image analytics
Large-scale
Segmentation
Image retrieval
Visual analytics

ABSTRACT

Despite the ever-increasing amount and complexity of annotated medical image data, the development of large-scale medical image analysis algorithms has not kept pace with the need for methods that bridge the semantic gap between images and diagnoses. The goal of this position paper is to discuss and explore innovative and large-scale data science techniques in medical image analytics, which will benefit clinical decision-making and facilitate efficient medical data management. Particularly, we advocate that the scale of image retrieval systems should be significantly increased at which interactive systems can be effective for knowledge discovery in potentially large databases of medical images. For clinical relevance, such systems should return results in real-time, incorporate expert feedback, and be able to cope with the size, quality, and variety of the medical images and their associated metadata for a particular domain. The design, development, and testing of the such framework can significantly impact interactive mining in medical image databases that are growing rapidly in size and complexity and enable novel methods of analysis at much larger scales in an efficient, integrated fashion.

1. Introduction

An important goal in medical image analytics is transforming raw images into a quantifiable symbolic form for indexing, reasoning and analysis. This task, of interest to medical imaging research, medical practice and the healthcare industry, is challenging because medical image content quantification is complex and not a solved problem. More importantly, few methods are able to analyze large-scale medical image databases (e.g., hundreds of thousands of images or more) in real-time, let alone to incorporate user-provided criteria or domain knowledge interactively. These are important requirements for medical doctors to analyze and use these databases. These drawbacks limit the effectiveness of current approaches in medical imaging research and industry applications to analyze the ever-growing number of medical images stored digitally.

Given the fact that the development of large-scale medical image analysis algorithms has lagged greatly behind the increasing quality (and complexity) of medical images and the imaging modalities themselves, there is an urgent need to develop innovative and integrated frameworks enabling robust and timely medical imaging and analysis, disease detection and characterization, and search in relevant databases. Recent advances in web-scale image analytics and multimodal databases have paved the way for large-scale, data-driven methods for robust detection and modeling, fine-grained disease classification, and semantic segmentation. In medical image analytics, the ever-increasing amount of medical images provides a foundation for novel semantic analysis methods. Transforming these raw medical images into a quantifiable, symbolic form will facilitate indexing and retrieval and potentially lead to new avenues of knowledge discovery and decision support (Fang et al., 2016). Therefore, it is now feasible to advance large-scale medical data analytics and information retrieval such that interactive systems can be effective for knowledge discovery in large databases of medical images.

This research direction needs a strong multidisciplinary component that involves a nexus of ideas from machine learning, image analysis, modeling, information retrieval and visual analytics. Conceptually, a potential solution can be considered consisting of three inter-related modules: (1) a robust parsing, modeling and segmentation module that provides automatic delineation and measurement of both healthy and abnormal cases, enabling effective extraction and analysis of information within specific regions, (2) a scalable learning-based image query module that retrieves instances among large databases in real-time, with morphological profiles most relevant and consistent to a query image for decision

* Corresponding author.
E-mail address: dnm@cs.rutgers.edu (D. Metaxas).

http://dx.doi.org/10.1016/j.media.2016.06.010
1361-8415/

support, and (3) an intelligent visualization and interaction module that integrates a user's input using learning methods, for improved accuracy and performance of the system. In this paper, motivated by our own research methods and results, we elaborate on each of these modules, and also discuss potential future directions.

2. Robust segmentation and modeling

An essential first step in medical image analytics is to locate and segment regions-of-interest (e.g., organs, cells) in images before extracting important features. Image segmentation in images is a key step in almost all medical image analysis methods. Despite the intense research in this area, automated and robust medical image segmentation is still an unsolved problem, due to many factors, such as imaging texture non-uniformity inside and outside organs, complex organ shape, limitations of imaging methods, scanner variability, imaging artifacts and image noise. To address these issues we have pioneered the use of deformable models as methods for simultaneous segmentation and modeling of medical imaging methods (Metaxas, 1997). Deformable models are curves or surfaces that move under the influence of internal smoothness and external image forces. There are two major classes of deformable models: parametric (explicit) models and geometric (implicit) models. Physics-based parametric models represent deformable curves and surfaces in their parametric form and deform the model by minimizing an energy function (Kass et al., 1988). Geometric deformable models represent curves and surfaces implicitly as level sets of a higher dimensional scalar function (Chan and Vese, 2001; Osher and Fedkiw, 2006). Despite the different formulations and implementations of these models, many of them rely on primarily edge information from image gradient to derive external image forces to drive a shape-based model. Thus, they are sensitive to image noise and artifacts.

To address these limitations, we developed a new class of deformable models, named "Metamorphs" (Huang and Metaxas, 2008). This method naturally integrates both shape and texture information and unifies the boundary and interior region information into a variational framework. The novelty of this approach is that the model is a deformable disk, where the probability density function of the interior texture also changes while both the interior texture and the boundary information modify the shape/volume of the deformable model. This type of method is suitable for medical image segmentation, where both boundary and texture information are important, while the texture is non stationary. When available, shape priors (Cootes et al., 1995; Heimann and Meinzer, 2009) have proved to be an effective strategy to ensure the robustness of deformable models, since they ensure that the segmented result follows the shape patterns learned from a training database. Inspired by compressed sensing, we proposed sparse representation-based shape priors for robust segmentation (Zhang et al., 2012), based on the observation that a small subset of the existing data is necessary to represent the statistical variation of shapes. Therefore, this extension leads to less sensitivity to weak or misleading image appearance cues. Instead of assuming any particular parametric model of shape statistics, our method incorporates shape priors on-the-fly through Sparse Shape Composition to handle outliers, preserve shape details, or model complex shape variations in a unified framework. We employ two types sparsity: (1) given a large shape repository for an organ, a specific shape instance of the same organ can be approximated by the composition of a sparse set of instances in the shape repository; and (2) gross errors from local appearance cues might exist, but these errors are sparse in spatial space. By incorporating these two sparsity priors into deformable models, the model becomes robust to gross errors, and it can preserve shape details even if they are not statistically significant in the training repository. This shape composition method benefits both the model initialization and refinement. Together with our methods of deformable segmentation, we are able to delineate the contours accurately.

In order to model topologically complex organs we also introduced the incorporation of topological information into deformable models. For example, we have used them to obtain accurate segmentation of the papillary muscles and the trabeculae of the heart's left ventricle from high resolution CT images (Gao et al., 2013). This is critical in understanding the anatomical function and geometric properties of the heart and the formation of clots in case of pathologies, such as the cardiac hypertrophy. This novel coupling of deformable models with persistent homology ensures the accurate segmentation, modeling and analysis of the cardiac trabeculation (Fig. 1).

3. Efficient retrieval in large-Scale medical databases

Once segmentation results are acquired, we can extract image features and use them based on machine learning techniques for computer aided diagnosis (CAD) and decision support[1]. Traditional classification-based methods may not be efficient or effective enough to discover information in large-scale databases. In recent years, researchers have become increasingly interested in content-based image retrieval (CBIR) for medical image analysis. Compared to traditional classification methods, which directly compute the likelihood of the diagnostic result, large-scale CBIR approaches open a new avenue for mining knowledge from large data and supporting clinical diagnosis. Therefore, a new research area has emerged that requires the development of large scale analytic methods that are capable of efficiently analyzing large data sets. Specifically, it is necessary to design a scalable learning-based query system that allows users (e.g. doctors, medical school students) to search large-scale medical image databases at different levels in real-time. Users can query based on either the low-level region information, the mid-level tissue/object information, or the high-level annotation information, using different features. An example of such an approach for histopathological image analysis is shown in Fig. 2.

Regarding the workflow in Fig. 2, given cell detection and segmentation results (Xing and Yang, 2016), image data can be represented by various "signatures" or features, such as image histograms, local texture/shape features and bag-of-words, when comparing the similarity among different clinical cases. However, such features lie in a high-dimensional space, and there is a very large number of features in large databases. To address these challenges, we have proposed scalable query methods designed for high-dimensional features (Zhang et al., 2015a; 2015b). Specifically, we developed kernelized and supervised hashing methods for efficient retrieval in high dimensional feature spaces, and validated this preliminary work on the cell-level analysis of thousands of breast tissue images, for the image-guided diagnosis of intraductal breast lesions. Hashing has been widely used to compress high-dimensional features into binary codes with tens of bits (Wang et al., 2016). Therefore, such short binary features allow mapping easily into a hash table for real-time search. To improve the accuracy of traditional hashing methods, we incorporated a kernelized scheme to handle imaging data that are linearly inseparable, a common phenomenon of medical images. We also leveraged supervised information to design discriminative hash functions that are suitable for medical data retrieval. This supervised information incorporates domain knowledge into feature similarities and has the

[1] Deep neural networks have been widely investigated in this field for feature learning and/or CAD (Greenspan et al., 2016). In this paper, we do not provide a comprehensive review due to different focuses.

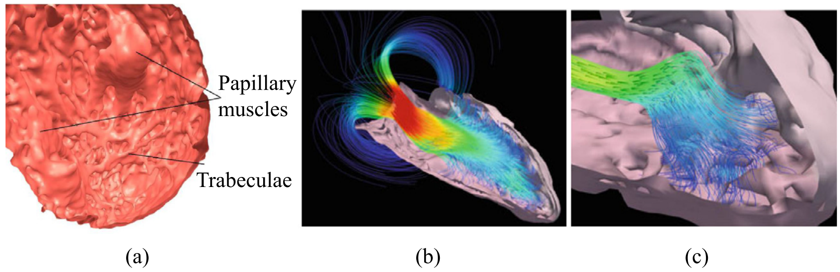

Fig. 1. (a): Our segmentation and modeling result of a healthy LV surface with accurate papillary muscles and trabeculae from CT data (Gao et al., 2013). (b) and (c): Using CT-based LV and valve reconstruction,visualization of cardiac blood flow using our method for Navier Stokes simulations with large boundary deformations. (Kulp et al., 2011).

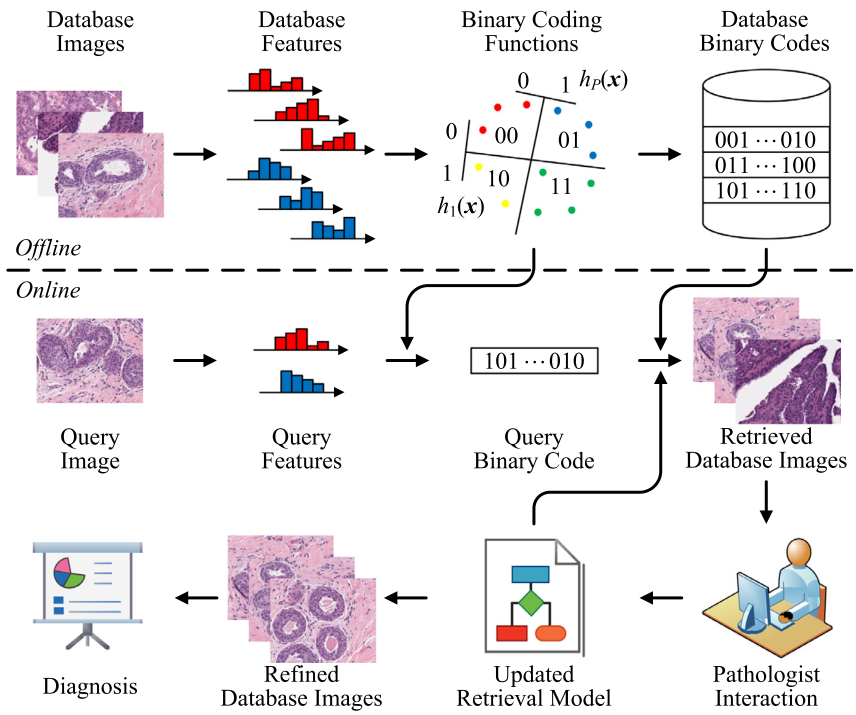

Fig. 2. Workflow of CBIR-based high-throughput histopathological image analysis. Multiple high-dimensional features are compressed into binary codes, which are used for efficient retrieval. Retrieved images can be used by pathologists to make clinical decisions.

potential to address the semantic gap. Using these compressed features, our method performs information retrieval in real time using millions of features without sacrificing the accuracy. The retrieval precision using our framework is around 88%, and the running time is less than ten milliseconds for thousands of these breast tissue images extracted from hundreds of patients, with millions of cells to analyze. In addition, we have also investigated effective approaches to integrate heterogeneous features, further boosting the performance of our system.

4. Intelligent interaction and visual analytics

To achieve the ultimate goal of assisting efficient decision making and reasoning using medical image databases, it is necessary to incorporate users in the loop to incorporate the domain knowledge of experts. While the automated methods are designed to process millions of images, human users can only reliably work with much fewer images at a time. A new challenge that arises towards further improving CBIR systems is to bridge the gap between the large-scale automated algorithms for image analytics and the knowledge that domain experts can provide even though it is at much smaller scale. To address this challenge we introduce a visual analytics system that uses a set of feature-based query, visualization, comparison, and learning methods for revealing the relevant image features and their relationships. This system supports the analysis of the retrieved relevant image sets, the extracted image features, and the feature similarities among the retrieved image sets.

We introduce two types of interactive reasoning functions that could be beneficial to display and explore analytics in large-scale medical image databases. The first set of interaction is to reason with the visualized results through feature-based selections. As doctors are often interested in specific features for different diagnosis purposes, we can allow users to select any feature combination and the visualization will be adjusted interactively to represent the relevance of feature aspects to the query. Users are also able to select and deselect retrieved instances to study details and adjust the search criteria. The interaction can be handled on individual instances as well as groups with similarity-based grouping techniques. The second set is to compare different sets of query results. This problem can be viewed as a higher level of the visual analytics for one query result from different feature aspects, as shown in Fig. 3. A flexible method could be investigated to relayout the images from different sets of query results with the consideration of generating bubble sets in various shapes for group information.

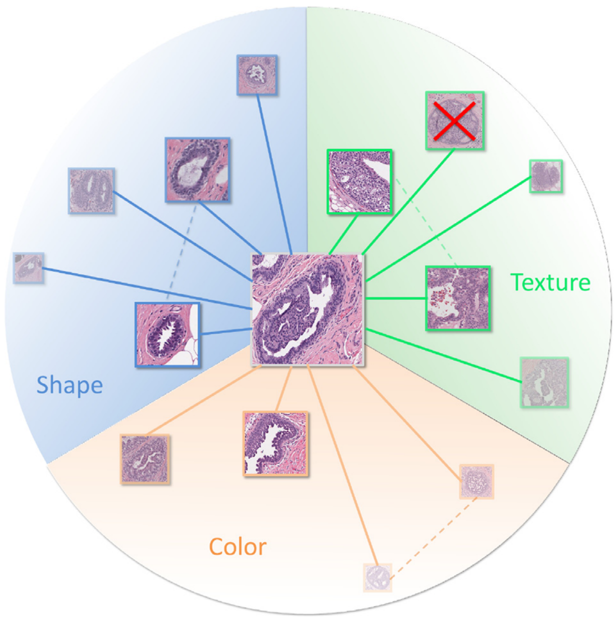

Fig. 3. Instead of plainly displaying retrieved images, we use feature-based group visualization to show them in terms of features and similarities.

Furthermore, it is also desired to provide efficient interaction methods for users to enhance the query algorithms and obtain finer-tuned results. Specifically, the image querying and visualization methods provide an initial structure to the returned images, relative to the query, based on feature similarity. However, in the general case, the visual similarity of a returned match to the query may not correspond to the semantic intent of the search. For example, in images from a nuclear scan of the liver, a returned match with similar vasculature may not be relevant if the clinician's intent was to find images with similar tumor activity. Designing a model that accounts for all of the possible variations within a large-scale database is infeasible. Incorporating the interactive aspects of data search and visualization with our underlying computational models can allow users to update the perceptual organization of the image set displayed in real-time.

5. Discussion and conclusions

In this position paper, we discussed the recent advances for scalable and interactive knowledge discovery and analytics in medical image databases. This research involves the following components, which can be investigated and evaluated in parallel and sequentially integrated: (1) robust segmentation and modeling of medical images in order to identify and model objects (e.g., tissue, organs, cells) of interest; (2) efficient learning-based reasoning in large-scale medical image databases using scalable and accurate medical image retrieval in large databases to provide real-time querying for the most relevant and consistent instances (e.g., similar morphological profiles) for clinical diagnosis and decision support; and (3) intelligent interaction and visual analytics that integrates expert feedback/knowledge and automated algorithms for efficient decision-making, comprehensive understanding of the query results and support of semantic interaction functions. These three components are coordinated for the purpose of scalable and interactive mining to provide a semantic interface between users and data.

These research activities provide the methodology that will result in improvements in applications that require large-scale medical image analytics. They include, healthcare practice, drug discovery and medical education. The efficient retrieval and visualization of relevant cases from medical databases will provide usable tools to assist and improve a clinician's diagnoses and support efficient medical image data management, such as picture archiving and communication systems (PACS). They will also pave the way for the discovery of new feature relationships and knowledge in different types of disease which has the potential to improve disease understanding and diagnosis. Finally, these methods will allow the exploration and analytics in structured and multimodal databases which may result in novel medical discoveries.

The fundamental multiscale and multimodal framework for medical image analytics, information processing, organization and retrieval we have introduced, has the potential to result in advancements in areas beyond medical image analytics such as computer vision, machine learning, and multimodal content analysis. There are several potential directions for future research. Although we mainly focus on the image-based analysis, it is also promising to use related patient and disease information, such as clinical reports. Therefore, this is similar to the image Q&A problem, i.e., given a query image, a relevant report is generated or retrieved from the database. In fact, researchers in this area have already started to connect structured image concepts to language, but it is still in a preliminary stage, and worth investigating. To achieve this, it is necessary to bring together recent advances in data mining, machine learning, nature language processing, and more, to generate relations between images and texts. Another direction is to consider and integrate into the retrieval framework additional data from the patients or cohorts, such as genomics, biomarkers and symptoms, to enhance disease diagnosis and/or prediction accuracy. For example, we can fuse phenotype and genotype data to retrieve relevant cases of disease by extracting and compressing features from both medical image and genomics data. These features are usually domain specific and often heterogeneous, while the machine learning methods are mostly information processing, organization and retrieval generic.

References

Chan, T.F., Vese, L.A., 2001. Active contours without edges. Image process., IEEE trans. 10 (2), 266–277.
Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J., 1995. Active shape models-their training and application. Comput. vision image understanding 61 (1), 38–59.
Fang, R., Pouyanfar, S., Yang, Y., Chen, S.-C., Iyengar, S.S., 2016. Computational health informatics in the big data age: a survey. ACM Comput. Surv.
Gao, M., Chen, C., Zhang, S., Qian, Z., Metaxas, D., Axel, L., 2013. Segmenting the papillary muscles and the trabeculae from high resolution cardiac ct through restoration of topological handles. In: Information Processing in Medical Imaging. Springer, pp. 184–195.
Greenspan, H., van Ginneken, B., Summers, R.M., 2016. Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique. IEEE Trans. Med. Imaging 35 (5), 1153–1159.
Heimann, T., Meinzer, H.-P., 2009. Statistical shape models for 3d medical image segmentation: a review. Med. image anal. 13 (4), 543–563.
Huang, X., Metaxas, D.N., 2008. Metamorphs: deformable shape and appearance models. Pattern Anal. Mach. Intell., IEEE Trans. 30 (8), 1444–1459.
Kass, M., Witkin, A., Terzopulos, D., 1988. Snakes: active contour models. IJCV 321–331.
Kulp, S., Gao, M., Zhang, S., Qian, Z., Voros, S., Metaxas, D., Axel, L., 2011. Using high resolution cardiac CT data to model and visualize patient-specific interactions between trabeculae and blood flow. In: MICCAI, pp. 468–475.
Metaxas, D.N., 1997. Physics-based deformable models: applications to computer vision, graphics and medical imaging, 389. Springer Science & Business Media.
Osher, S., Fedkiw, R., 2006. Level set methods and dynamic implicit surfaces, 153. Springer Science & Business Media.
Wang, J., Liu, W., Kumar, S., Chang, S.-F., 2016. Learning to hash for indexing big data: A survey. Proc. IEEE 104 (1), 34–57.
Xing, F., Yang, L., 2016. Robust nucleus/cell detection and segmentation in digital pathology and microscopy images: a comprehensive review. IEEE rev. biomed. eng.
Zhang, S., Zhan, Y., Metaxas, D.N., 2012. Deformable segmentation via sparse representation and dictionary learning. Med. Image Anal. 16 (7), 1385–1396.
Zhang, X., Liu, W., Dundar, M., Badve, S., Zhang, S., 2015. Towards large-scale histopathological image analysis: Hashing-based image retrieval. IEEE Trans. Med. Imaging 34 (2), 496–506.
Zhang, X., Xing, F., Su, H., Yang, L., Zhang, S., 2015. High-throughput histopathological image analysis via robust cell segmentation and hashing. Med. image anal. 26 (1), 306–315.