1
Medical Decision Support Using Increasingly Large Multimodal Data Sets

Henning Müller, Devrim Ünay

Abstract

Medical decision support has traditionally been using model-based (or handcrafted) approaches and small data sets for evaluation. As medical data are sometimes hard to obtain and annotations by clinicians required for learning-based approaches are very expensive they are often acquired in small amounts. Applications for such decision support include detection, classification and retrieval in several clinical scenarios. Each of these applications provides or modifies information that is essential in decision making.

This chapter analyses recent trends and techniques that make use of increasingly large data sets and thus more data driven approaches to medical decision support that have in some areas replaced the more traditional rule-based (or model-based) approaches. Besides techniques and available data sets, the text describes scientific challenges that have helped to make data and annotations available that really are the basis for advances in the field. Challenge infrastructures and approaches are explained as they are often essential to access data, and application scenarios give examples of existing applications and objectives. An outlook into ideas on how to overcome current limitations and how to tackle the upcoming challenges of image-based decision support for digital medicine ends the chapter.

1.1
Introduction

Digital medicine has become a reality over the past few years and hospitals in basically all Western countries have been using electronic medical records for several years now, making all the clinical data available in digital form [1]. This has created a major revolution in medicine that is under way, where experience is increasingly complemented by evidence based on data that are often complex to integrate manually [2, 3]. Other authors highlight the cost savings that are possible when using big data in medicine [4]. In addition to reducing costs, big data initiatives in health care can save lives and improve patient outcomes [5]. Particularly personalized or precision medicine really requires much data on individuals for optimizing treatment outcomes based on past results [6]. It is specifically stated in [7] that the prospect of applying precision medicine is dramatically improved by the recent development of large–scale biomedical databases (such as the human genome data), powerful methods for characterizing patients (such as proteomics, metabolomics, genomics and cellular assays), and computational tools for analyzing large sets of data. In order to fulfill this great potential big data analytics has to overcome some challenges like variations in the data, privacy constraints, security, and accessibility [8]. An overview of big data for health is given in [9].

Much has been written on big data analysis in the medical field and particularly medical imaging is estimated to occupy up to 30% of world storage [10]. Genomic data stored has in recent years also increased exponentially and in cancer care many hospitals now have a full genome sequencing for all patients, leading to massive amounts of data that become available for data analysis. The amount of data produced in modern imaging protocols (that can contain tens of thousands of images per patient or extremely large images such as in histopathology) and of genomic data mean that a fully manual analysis has sometimes simply become impossible. Hence, these huge heterogeneous data need to be efficiently processed (stored, distributed, indexed, and analyzed) by big data analytics to improve medical decision support [11]. [12] states that there is much potential in delivering more targeted, wide-reaching and cost-efficient health care by exploiting current big data trends and technologies.

In general, the electronic health record has as an objective to bring the right information to the right people in the right moment and the right format [1]. With massive amounts of data to be interpreted it also means that the right tools are required in this context to make a meaningful use of the supplied data and integrated interactions between the data. Historically, medical decision support has been rule-driven and decision trees have been developed based on medical knowledge for many years, where rules are derived from the data manually. In recent years data-driven approaches have aimed at complementing this in areas such as Radiomics [14], where connections in the data are learned and then often used for prediction. For visual medical decision support three large types of approaches have been used besides simple quantitative measures: *detection, classification and retrieval*. In detection the objective is to find specific regions in volumes for example nodules in the lung [15] or specific lesions [16]. Classification has looked more into determining the specific type of an entire image or image region (region of interest). Figure 1.1 gives an example for the classification of the lung tissue into a limited number of classes and its visualization. Many similar tools have been proposed and developed.

Medical image retrieval has been another active area of research with review articles describing the main approaches [17, 18]. The idea is that physicians can take images of a current case and clinical data and use these to formulate queries to find

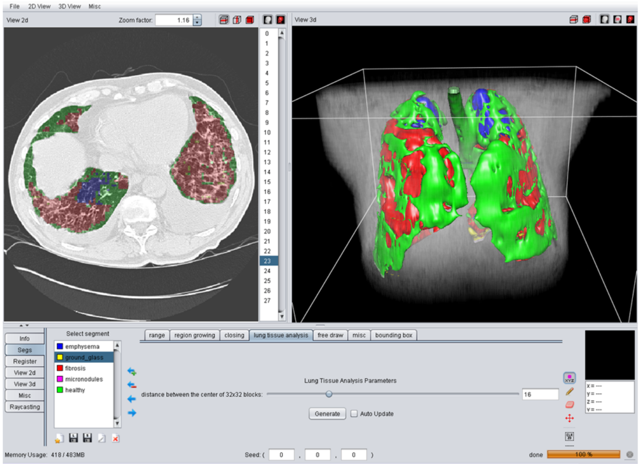

Figure 1.1 Screen shot of a tool for 3D lung tissue categorization (details in [13]).

similar cases or lesions [19, 20]. An example for a retrieval application can be seen in Figure 1.2, where queries can be formulated using text and images combined.

Whereas many early approaches relied on single sources of data, so either images, or text or structured data, it has become clear that multimodal data and fusion approaches can lead to much better results [22, 23]. An exemplary architecture of such a retrieval system can be seen in Figure 1.3. Still, most systems using visual data do not take into account any other data types. This is also shown by the relatively small number of papers found in this survey. This survey concentrates fully on multimodal approaches, thus only papers are included in the detailed analysis that include more than a single modality. mainly text and visual information. It also concentrates on data-driven approaches and thus does not include rule-based approaches.

1.2 Methodology for Reviewing the Literature of this Survey Paper

This chapter aims at reviewing the literature of mainly the past five years (2011-2016) in the field of medical visual decision support and highlights the use of multimodal data and data-driven approaches, even though not all applications really use multiple modalities at the moment. Several older papers are cited as well, mainly when we considered them important or landmark papers in the domain. We aimed to cite journal articles over conference papers when similar topics were available and we aim at having a variety of authors for the various topics, even though it is possible that papers have been missed, as the field is extremely large and consists of many communities. A concentration is on multimodal (visual + non-visual) data usage,

Figure 1.2 Screen shot of an image retrieval application that allows for text search and also image similarity search in medical data, taken from [21].

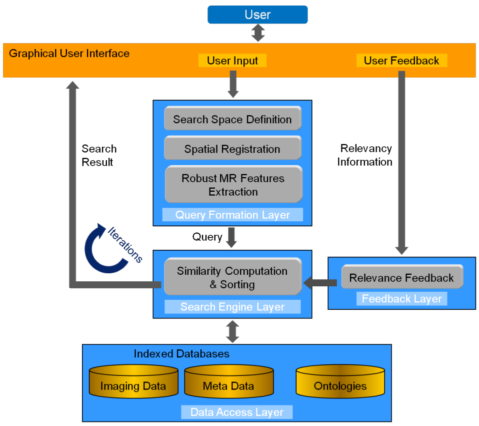

Figure 1.3 Architecture of an image retrieval system that exploits multimodal information for improved similarity search in medical data, taken from [24].

which is surprisingly small with most researchers focusing on single media, in our case only visual data without taking any other information on the patients. Many non-visual and visual-only areas are left out, as this was not considered a main objective of the text.

For the papers included in this survey we systematically searched through the Pubmed, IEEE, and Google Scholar databases and conferences such as MICCAI (Medical Image Computing and Computer Assisted Interventions), ISBI (IEEE International Symposium on Biomedical Imaging), EMBC (International Conference of the IEEE Engineering in Medicine and Biology Society), and CVPR (IEEE Conference on Computer Vision and Pattern Recognition) by using the search terms “multimodal medical decision support", “large scale medical diagnosis", “medical deep learning", “big data analytics in healthcare", “medical multimodal diagnosis", “large scale multimodal medical retrieval", and the like. The set of resulting references was then manually controlled and analyzed to keep only the articles corresponding to our criteria.

1.3 Data, Ground Truth and Scientific Challenges

This section describes the main challenges in medical data access. Medical data have specific requirements, as they can not simply be shared for research use but underly control of ethics committees and strong rules regarding the data sharing [25]. This often results in relatively small data sets and data that can rarely be shared for research use, even though data sharing is regarded as a main driver for innovation [26]. Scientific challenges have managed to make large data sets available and this is described here in more detail as well. To ease the analysis, the papers included are grouped by the size of the data sets they use as small (<1K images/patients), medium (1K–100K), or large (>100K).

1.3.1 Data Annotation and Ground Truthing

Usually, imaging data without annotations is difficult to be used for evaluating image analysis algorithms. Thus, annotation of images is an important aspect of many image-related research projects [27] to create a (sometimes subjective) ground truth. Machine learning algorithms most often require global image annotations or even better labeling of precise image regions. This is expensive and hard to obtain at a large scale, particularly in medicine where domain experts (physicians) are required, as the labeling is highly specific. Some efforts exist to make large annotated data sets available [28] and also to generate annotations at a large scale based on small manually annotated data [29]. Sometimes when enough globally annotated data are available then local regions with abnormalities can be derived as described in [30]. The type of annotations available also has a strong influence on the metrics that can be used for the evaluation and for comparing different approaches [31]. For evaluating detection detailed regions are required and for segmentation evaluation, fully annotated structures. For global classification, simple global labels are sufficient.

1.3.2 Scientific Challenges and Evaluation-as-a-Service

Historically, medical image analysis has been done on very small collections of private data, and it was thus impossible to compare any two algorithms, as baselines can be chosen fairly arbitrarily [32]. Over the years, data sharing was often promoted but only rarely put in practice on a larger scale. One of the first medical evaluation campaigns with a publicly shared image data set was the ImageCLEF medical task in 2004 [33], where several groups compared their algorithms on image retrieval. A medical task has been held almost every year since 2004 on retrieval or classification.

At the MICCAI conference the first challenge was on liver segmentation using a fairly modest data set [34] in 2007, with an on site challenge and a limited time of a few hours for computation. This has now led to many more challenges that have had an increasingly important impact also at other conferences [35, 36]. The web page on the Grand Challenges in Medical Image Analysis[1)] aims at collecting all active and past challenges in medical image analysis and these have increased strongly over the years.

Scientific challenges have improved comparability of algorithms and made medical data accessible for many researchers, which is already an important impact. Scholarly and financial impact are also high [37, 38]. But there are difficulties to distribute data when the data sets become extremely large (download impossible, sending hard disks not practical), when they are confidential (can not be shared) and when they change quickly over time (requirement to always work on the latest data), as is the case with medical routine data [39]. Approaches that bring the algorithms to the data have thus been developed to solve most of these challenges. In a larger context this is also called Evaluation-as-a-Service (EaaS) [40]. Different approaches have been implemented for this, ranging from submissions of source code, virtual machines and Docker containers.

1.3.3 Other Medical Data Resources Available

Funding organizations have realized over the past 20 years that a really large part of the research budget and time in projects is often used for acquiring and cleaning the required data but many of the data sets are only exploited in a short project lifetime and by a very small number of people. This has led to data sharing policies of funding organizations. For example, the NIH (National Institutes of Health) require funded projects to make the data available after the project end. This has led to several large archives of data that remain available for research use and are frequently

1) `https://grand-challenge.org`

employed for research, such as the TCIA[2)] (The Cancer Imaging Archive) and the TCGA[3)] (The Cancer Genome Atlas) that include many images and other clinical data or documents. The data can then be used for several years even though sometimes imaging protocols change and clinical impact on current care may be limited with very old data. This allows many research groups now to access image data sets and work on the data for research without the requirement to go through ethics approval and tedious data collection in collaboration with a clinical institution. Sharing the data acquisition can also lead to bigger and better annotated data sets if costs can be shared.

In the European Union similar ideas are underway to make sure that data created with public money can also be used by other researchers and have a maximum impact. Guidelines are made available to push for more open access publishing and data sharing, not limited to medical data[4)].

1.4
Techniques Used for Multimodal Medical Decision Support

This section summarizes the techniques employed in the papers we found through a detailed literature search, with most being published from 2011-2016.

Table 1.1 presents an overview of the papers included in this survey, while Figure 1.4 provides the breakdown of these papers into year and data set size, anatomy, application, and imaging modality. This gives a good global overview of topics that were covered in the past five years and the tendencies on multimodal data usage.

1.4.1
Visual and Non-visual Features Describing the Image Content

The literature on multimodal medical decision support consists of work exploiting various types of image features, such as intensity [43, 46, 49, 50, 56, 60, 64, 65, 69, 70], color/grey level histogram analysis [50, 52, 56, 60], texture [43, 51, 65, 69, 70], shape [24, 41, 50, 59, 60, 65, 69, 70], and volumetric information of objects of interest [42, 44, 45, 46, 47, 48, 49, 64, 65, 70]. Local binary patterns [52], and the popular scale invariant feature transform (SIFT) [55, 57] have also been used.

With the recent popularity of deep learning, several articles made use of features based on the output of deep neural networks for multimodal medical decision support [53, 54, 61, 62, 63, 66, 67, 68].

Augmenting information mined from images with that gained from non-imaging data has been a popular track to improve medical decision support. To this end, researchers benefit from demographic indicators such as age and gender [24, 49, 58, 63,

2) http://www.cancerimagingarchive.net
3) https://cancergenome.nih.gov
4) http://ec.europa.eu/research/participants/docs/h2020-funding-guide/cross-cutting-issues/open-access-data-management/open-access_en.htm

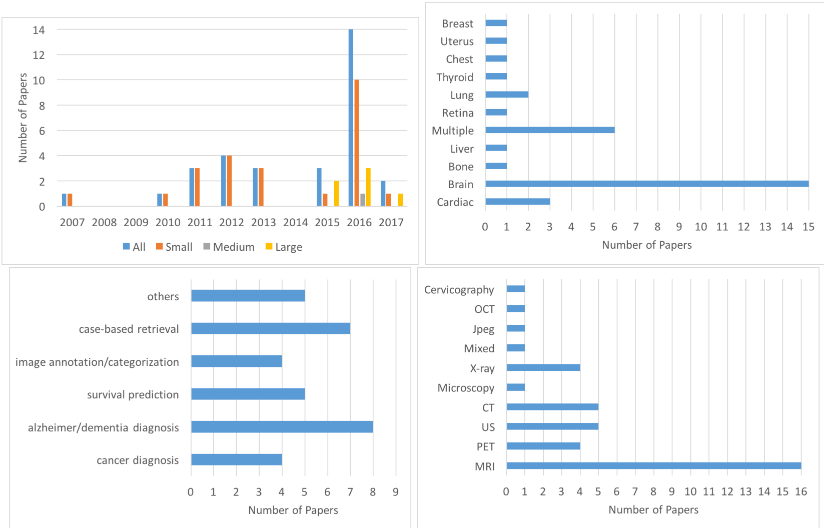

Figure 1.4 Breakdown of papers included in this survey by year and data set size, anatomy, application and imaging modality.

64, 70, 66, 68, 69, 65], medical history of subjects [24], clinical findings like biopsy results, CSF (CerebroSpinal Fluid) biomarkers, physical/blood measurements, tumor descriptors [24, 44, 45, 46, 47, 48, 49, 50, 51, 58, 59, 63, 64, 65, 68], natural language processing based analysis of clinical/radiology reports [53, 54, 56, 61, 62, 67] and biomedical journal articles [52, 55], analysis of biosignals such as EEG (Electro-EncephaloGram), ECG (Electro-CardioGram) and audio recordings [41, 42, 60], and genetic test results [43, 48, 59].

1.4.2
General Machine Learning and Deep Learning

Machine learning is an indispensable part of medical decision support, especially in diagnostic and categorization applications. In order to mine multimodal information, the literature made use of various machine learning techniques such as ensemble learners [42, 50], support vector machines [44, 45, 46, 48, 51, 55, 63], naive Bayes classifiers [48, 51], nearest neighbor classifiers [51, 56], decision trees [49], and random forests [65, 70]. Depending on the applications all these classifiers have their advantages and inconveniences. To combine the results of the different modalities, early fusion employs all features in a single space whereas late future combines classifier outcomes. More on fusion techquies between visual features and text can be found in [71].

Deep learning has had an important influence in computer vision over the past ten years [72]. The medical field has also had a strong increase in using neural networks for decision support. Several review articles have now been published such as [73] for health informatics in a more general sense and [74] really focusing on medical image analysis. Studies employing multimodal analysis of medical data covered in this survey generally employ convolutional neural networks for analysis of images and non-image data at the same time [53, 54, 62, 63, 66, 67, 68, 69]. One exception is [61] where a dual-network architecture is used with convolutional neural networks for image analysis and recurrent neural networks for extraction of text features from radiology reports.

In a guest editorial on deep learning in medical image analysis [75] the authors also highlight all the opportunities that these techniques have for the future of digital medicine. Still, in terms of the exact network architectures for combining several complementary sources of information there is still much work that is required.

1.5 Application Types of Image-Based Decision-Support

This section categorizes the papers included in this survey by their applications scenarios. We separate these into several types of applications: (1) *localization*, has as objective to find a region of interest in images with specific patterns; (2) *segmentation* aims at finding exact borders of an organ, lesion or other structure; (3) *classification* aims at attaching a finite number of class labels to images or regions of interest; (4) *prediction* aims at predicting a future event, for example treatment response or survival time; (5) *retrieval* aims at responding to information needs of users or finding similar images or regions; (6) *automatic image annotation* aims at attaching key words or semantic concepts to images or image regions. The underlying techniques (visual characteristics, machine learning, fusion) for all these application areas are similar and related but they also have differences in the exact objectives.

1.5.1 Localization

Many medical imaging tasks require the detection and localization of anatomical/pathological regions as a first step, so finding a small region that bears the important information of an image. In an exemplary multimodal work, computer-aided detection/localization of mammographic lesions from a large X-ray data set is accomplished by training a convolutional neural network with image features and demographic information [69]. The difficulty of localization is that there is much anatomic variation in medical images and anomalies are often focused on very small parts in the multidimensional data.

1.5.2 Segmentation

Exact delineation of regions-of-interests (lesions, organs) from medical images is a crucial step for building accurate decision support. To this end, researchers exploit non-image information for improved image segmentation. For example, in [53] semantic concepts mined from clinical reports are used to improve segmentation of optical coherence tomography images of the retina into intraretinal cystoid fluid, subretinal fluid and normal retinal tissue by convolutional neural networks. Several segmentation challenges of data from single modalities was provided by the VISCERAL project [28].

1.5.3 Classification

Classification of multimodal medical data into "healthy vs. disease", "normal vs. pathology", or "different disease stages" is an important task to improve medical decision support. Classification means to add each image, volume or region into one or several of a set of classes. Accordingly, most of the papers included in this survey focus on the classification task. Among these, many target diagnosis of Alzheimer's from small data sets (the ADNI data sets are publicly available, easing this application domain) by augmenting image features extracted from MRI-only [44, 47, 48, 58] or MRI+PET (Positron Emission Tomography) [42, 45, 46] with non-image features such as demographic indicators [58], clinical findings like CSF biomarkers [44, 45, 46, 47, 48, 58], biosignals [42], and genetic test results [48]. They then feed these multimodal features into machine learning algorithms such as naive Bayes [48], support vector machines [44, 45, 46, 48] or ensemble learners [42].

Cancer diagnosis is another focus of interest, where examples include glioblastoma diagnosis using microscopy images and genetic test results [43], thyroid cancer diagnosis from ultrasound images and biopsy findings [59], and cervical dysplasia diagnosis from cervicograms, demographic indicators and clinical findings [68]. All realized using small data sets.

Besides the aforementioned articles, [50] performed osteoporosis diagnosis over a small data set by combining information extracted from CT and X-ray images with physical and blood measurement results in an ensemble learner framework, while [51] realized staging of chronic liver disease over a small database by exploiting ultrasound images and laboratory/clinical findings via machine learning.

Other articles include a large-scale modality classification application where images and texts from journal articles are employed [55], and a deep learning application of image categorization into clinically/semantically relevant clusters by using CT and MRI images together with radiology reports [62].

1.5.4 Prediction

Prediction aims at predicting a future temporal event based on past data, for example predicting treatment outcome or estimating survival time. The added-value of using multimodal information in medical decision support is shown in the application of cancer survival prediction. Several small-scale studies merge information extracted from structural and functional images like diffusion/perfusion MRI with demographic indicators of subjects and clinical manifestations of tumors. Existing papers typically realize survival prediction by using machine learning or deep learning [49, 63, 64, 65, 70].

1.5.5 Retrieval

Retrieval aims at supplying information to fulfill an information need. With images this often means to find similar images or similar cases to an example. Case-based retrieval of similar patients is playing an increasingly important role for improving decision support, diagnosis, treatment planning, and physician's training. In this context, [41] presented a small-scale cardiac decision support system using ultrasound images with electroencephalography data and audio recordings. [24] proposed a small-scale retrieval based dementia diagnosis system where MRI data are used together with demographic indicators, medical history of patients and clinical findings. Recently, [56] realized case-based retrieval over a small interstitial lung disease database where CT images are employed with radiology reports. [57] implemented a small-scale case-based retrieval system by employing latent semantic topics together with MRI and PET images for Alzheimer's and CT images for lung nodules. [66] proposed and evaluated a pathology retrieval system on a small data set of chest X-rays where image information is augmented by demographic indicators.

Besides these, [52] realized medical information discovery from literature by mining Jpeg images and texts of a large number of PubMed articles. [55] adapted case-based retrieval and modality classification of multiple diseases using images and texts of a large data set of journal articles.

1.5.6 Automatic Image Annotation

Automatic image annotation aims at attaching keywords or semantic labels to visual data, so they can be searched or classified in an easier way. Semantic annotation of medical images is a prerequisite for building comprehensive archives that can be used for not only evidence-based diagnosis but also biomedical research and education of physicians. As such, recent efforts focus on exploiting multimodal data and deep learning for this problem. [54] used MRI and CT images together with clinical reports, while [61] employed X-ray images with radiology reports to annotate large databases of multiple diseases. On the other hand, [67] combined information from Doppler ultrasound images and radiology reports to annotate a small cardiac image database.

1.5.7 Other Application Types

Besides the aforementioned articles, [76] presented a clinical use-case evaluation study for effectiveness of web-based decision support tools for coronary artery disease, where perfusion scintigraphy images, electrocardiography data, the patient's medical history and clinical findings of a small data set are used. [60], on the other hand, proposed a cardiac monitoring system employing ultrasound images with electrocardiography data. These application types are different from the categories described above.

1.6 Discussion on Multimodal Medical Decision Support

The larger number of more recent papers using multimodal (visual, plus structured data or free text) medical data for decision support shows the increasing activity in this research domain. However most of these use small data sets. Many techniques have been employed and in recent years an increasing number of approaches used convolutional neural networks and similar techniques. Many variations also try to combine traditional approaches with deep learning, which often leads to good results. A big challenge remains the accessibility of large annotated data sets. Some large data sets have become available, so data itself is slowly becoming less of a problem but annotating data is critical and leveraging small amounts of manual annotations is essential. For this reason, active learning techniques and silver corpus approaches are often required to make the data really usable and useful. Learning from weakly annotated data is another direction that has started in very focused domains. Having many research data sets that are small and private also creates problems with reproducibility of the research results.

Evaluation campaigns have shown a tremendous impact in many areas [38, 77], and the medical data analysis domain is no exception. Such benchmarking campaigns are now part of all major conferences. This has made data and annotations available with a clear evaluation scenario, so it really becomes possible to compare results obtained with these techniques. Still, the size of data sets is nothing compared to what is being made available in terms of web images and maybe relatively new initiatives such as Medical ImageNET[5)] can help in the future when they are fully employed.

1.7 Outlook into the Next Steps of Multimodal Medical Decision Support

Deep learning has had an enormous impact on computer vision over the past five years and in medical image analysis or multimodal data processing the current tools such as Tensorflow[6)] or Caffe[7)] are very frequently used. With larger data sets becoming available the full potential of these techniques can likely be seen in a few years. It seems important to adapt the tools to the scenario of medical data where

5) http://langlotzlab.stanford.edu/projects/medical-image-net/
6) https://www.tensorflow.org/
7) http://caffe.berkeleyvision.org/

very small parts of an image are usually relevant for the decision making and not the entire image or volume as is often the case in stock photography. There are often 2000-4000 gray levels in the image and high resolutions, which is not foreseen in many neural network architectures and has to be added to really leverage on the full image information; or at least the right level/window levels need to be chosen depending on the organ to be analysed. New networks for visual and multimodal data that include the specific aspects of medical data are expected to be developed and also a systematic approach towards choosing the optimal networks for a specific scenario and type of data.

Fully training networks on medical data is often not possible as the amounts of data available are often too small and the variety in the data is enormous. Thus, transfer learning is used and for images often the ImageNET pre-trained networks are employed. This can be expected to change. Larger data sets have now become available and it is a question of how to best obtain annotations and limit the amount of manual work. This requires to learn from larger but weakly annotated data sets. It also pushes for active learning approaches to focus annotation efforts to the most informative examples.

Another activity that will likely gain traction is the evaluation of tools in clinical practice, so the evaluation of impact on clinical care. Most current tools cited in this paper are research prototypes and few have been evaluated in clinical scenarios even on a modest scale. To show performance in real situations clinical impact needs to be shown and tested in large-scale studies. With an increasing amount of data available on individuals the need to use tools will increase and likely also the difference in quality that can be reached. Even though some people predict that radiologists may not be necessary in the future for evaluating images we rather think that tools will improve efficiency of the clinical workflow by automating simple cases, and concentrate the clinician's work on the difficult cases and on interpreting the data and thus also the effectiveness.

It was shown that it is impossible to read all the literature in even a small focused medical domain [78], so the access to recent medical knowledge will remain an important aspect of making decision evidence-based and personalized. The half-time of medical knowledge is estimated to be between 5 and 10 years and this also highlights why decision support is needed: it is important to always include the latest knowledge into precision medicine.

With the ever–increasing amount of medical multimodal data available, big data analytics is becoming more important every day. If the associated challenges - heterogeneity, fragmentation, availability, privacy and security of data - are properly addressed [79], big data analytics has the potential to improve medical decision support by exploiting large multimodal data sets in the near future and eventually transform healthcare. This is only achievable with good data management and computerized decision support.

Table 1.1 Overview of papers using multimodal data sets for medical decision support. Data sets grouped as small (<1K patients/images), medium (1K–100K), or large (>100K). MRI=magnetic resonance imaging; PET=positron emission tomography; US=ultrasound; CT=computed tomography; OCT=optical coherence tomography. Bs=biosignals; Ge=genetics; CF=clinical findings; De=demographics; MH=medical history; Re=reports. SVM=support vector machines; NB=naive bayes; LR=logistic regression; DT=decision tree; NN=nearest neighbor; RF=random forests; CNN=convolutional neural networks; RNN=recurrent neural networks; n.a.=not available.

	Modality					
Ref.	Image	Non-image	Application	Anatomy	Data set	Learning
[41]	US	Bs	Cardiac decision support	Cardiac	Small	n.a.
[42]	MRI, PET	Bs	Alzheimer diagnosis	Brain	Small	Ensemble
[43]	Microscopy	Ge	Cancer diagnosis	Brain	Small	n.a.
[44]	MRI	CF	Alzheimer diagnosis	Brain	Small	SVM
[45]	MRI, PET	CF	Alzheimer diagnosis	Brain	Small	SVM
[46]	MRI, PET	CF	Alzheimer diagnosis	Brain	Small	SVM
[47]	MRI	CF	Alzheimer diagnosis	Brain	Small	n.a.
[48]	MRI	CF, Ge	Alzheimer diagnosis	Brain	Small	SVM, NB, LR
[49]	MRI	De, CF	Survival prediction	Brain	Small	DT
[24]	MRI	De, MH, CF	Dementia diagnosis	Brain	Small	n.a.
[50]	X-ray, CT	CF	Osteoporosis diagnosis	Bone	Small	Ensemble
[51]	US	CF	Disease staging	Liver	Small	SVM, NN, NB
[52]	Jpeg	Re	Case-based retrieval	Multiple	Large	n.a.
[53]	OCT	Re	Tissue segmentation	Retina	Small	CNN
[54]	CT, MRI	Re	Image annotation	Multiple	Large	CNN
[55]	Mixed	Re	Case-based retrieval	Multiple	Large	SVM
[55]	Mixed	Re	Modality classification	Multiple	Large	SVM
[56]	CT	Re	Case-based retrieval	Lung	Small	NN
[57]	CT	Re	Case-based retrieval	Lung	Small	n.a.
[57]	MRI, PET	Re	Case-based retrieval	Brain	Small	n.a.
[58]	MRI	De, CF	Alzheimer diagnosis	Brain	Small	n.a.
[59]	US	CF, Ge	Cancer diagnosis	Thyroid	Small	n.a.
[60]	US	Bs	Disease monitoring	Cardiac	n.a.	n.a.
[61]	X-ray	Re	Image annotation	Multiple	Medium	CNN, RNN
[62]	CT, MRI	Re	Image categorization	Multiple	Large	CNN
[63]	MRI	De, CF	Survival prediction	Brain	Small	SVM, CNN
[64]	MRI	De, CF	Survival prediction	Brain	Small	n.a.
[65]	MRI	De, CF	Survival prediction	Brain	Small	RF
[66]	X-ray	De	Pathology retrieval	Chest	Small	CNN
[67]	US	Re	Image annotation	Cardiac	Small	CNN
[68]	Cervicogram	De, CF	Cancer diagnosis	Uterus	Small	CNN
[69]	X-ray	De	Cancer diagnosis	Breast	Large	CNN
[70]	MRI	De	Survival prediction	Brain	Small	RF

References

1 Haux, R. (2005) Hospital information systems — past, present, future. *International Journal of Medical Informatics*, **75**, 268–281.

2 Topol, E. (2013) *The creative destruction of medicine – how the digital revolution will create better healthcare*, Basic Books.

3 Anthony Celi, L., Mark, R.G., Stone, D.J., and Montgomery, R.A. (2013) "big data"in the intensive care unit. closing the data loop. *American Journal of Respiratory and Critical Care Medicine*, **187** (11), 1157–1160.

4 Bates, D.W., Saria, S., Ohno-Machado, L., Shah, A., and Escobar, G. (2014) Big data in health care: Using analytics to identify and manage high-risk and high-cost patients. *Health Affairs*, **33** (7), 1123–1131.

5 Groves, P., Kayyali, B., Knott, D., and Kuiken, S.V. (2013) The 'big data' revolution in healthcare: Accelerating value and innovation, *Tech. Rep.*, McKinsey & Company, Center for US Health System Reform, Business Technology Office.

6 Mirnezami, R., Nicholson, J., and Darzi, A. (2012) Preparing for precision medicine. *New England Journal of Medicine*, **366** (6), 489–491.

7 Collins, F.S. and Varmus, H. (2015) A new initiative on precision medicine. *New England Journal of Medicine*, **372** (9), 793–795.

8 Kruse, C.S., Goswamy, R., Raval, Y., and Marawi, S. (2016) Challenges and opportunities of big data in health care: A systematic review. *JMIR Medical Informatics*, **4** (4), e38.

9 Andreu-Perez, J., Poon, C.C.Y., Merrifield, R.D., Wong, S.T.C., and Yang, G.Z. (2015) Big data for health. *IEEE Journal of Biomedical and Health Informatics*, **19** (4), 1193–1208.

10 (2010), Riding the wave: How europe can gain from the rising tide of scientific data, Submission to the European Comission, available online at `http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/hlg-sdi-report.pdf`. URL `http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/hlg-sdi-report.pdf`.

11 Wang, L., Ranjan, R., Kolodziej, J., Zomaya, A., and Alem, L. (2015) Software tools and techniques for big data computing in healthcare clouds. *Future Gener. Comput. Syst.*, **43** (C), 38–39.

12 Heinrich, A. and et al. (2016) Big data technologies in healthcare: Needs, opportunities and challenges, *Tech. Rep.*, Big Data Value Association: TF7 Healthcare Subgroup.

13 Depeursinge, A., Van De Ville, D., Platon, A., Geissbuhler, A., Poletti, P.A., and Müller, H. (2012) Near–affine–invariant texture learning for lung tissue analysis using isotropic wavelet frames. *IEEE Transactions on Information Technology in BioMedicine*, **16** (4), 665–675.

14 Aerts, H.J., Velazquez, E.R., Leijenaar, R.T., Parmar, C., Grossmann, P., Carvalho, S., Bussink, J., Monshouwer, R., Haibe-Kains, B., Rietveld, D. *et al.* (2014) Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nature communications*, **5**.

15 Gurcan, M.N., Sahiner, B., Petrick, N., Chan, H.P., Kazerooni, E.A., Cascade, P.N., and Hadjiiski, L. (2002) Lung nodule detection on thoracic computed tomography images: Preliminary evaluation of computer–aided diagnosis system. *Medical Physics*, **29** (1), 2552–2558.

16 Smith, J.T., Hawkins, R.M., Guthrie, J.A., Wilson, D.J., Arnold, P.M., Boyes, S., and Robinson, P.J. (2010) Effect of slice thickness on liver lesion detection and characterisation by multidetector CT. *Journal of Medical Imaging and Radiation Oncology*, **54** (3), 188–193.

17 Müller, H., Michoux, N., Bandon, D., and Geissbuhler, A. (2004) A review of content–based image retrieval systems in medicine–clinical benefits and future directions. *International Journal of Medical Informatics*, **73** (1), 1–23.

18 Akgül, C., Rubin, D., Napel, S., Beaulieu, C., Greenspan, H., and Acar, B. (2011) Content–based image retrieval in radiology: Current status and future directions. *Journal of Digital Imaging*, **24** (2), 208–222.

19 Depeursinge, A., Vargas, A., Gaillard, F., Platon, A., Geissbuhler, A., Poletti, P.A., and Müller, H. (2012) Case–based lung image categorization and retrieval for interstitial lung diseases: clinical workflows. *International Journal of Computer Assisted Radiology and Surgery*, **7** (1), 97–110.

20 Jimenez-del-Toro, O., Hanbury, A., Langs, G., Foncubierta-Rodríguez, A., and Müller, H. (2015) Overview of the VISCERAL Retrieval Benchmark 2015, in *Multimodal Retrieval in the Medical Domain (MRMD) 2015*, *Lecture Notes in Computer Science*, vol. 9059, Springer, *Lecture Notes in Computer Science*, vol. 9059.

21 Schaer, R. and Müller, H. (2014) A modern web interface for medical image retrieval, in *Annual Congress SGMI 2014*.

22 May, M. (2014) Life science technologies: Big biological impacts from big data. *Science*, **344** (6189), 1298.

23 Jimenez-del-Toro, O., Cirujeda, P., and Müller, H. (2017) Combining radiology images and clinical meta–data for multimodal medical case–based retrieval, in *Cloud–Based Benchmarking of Medical Image Analysis*, Springer International Publishing.

24 Soydemir, M. and Unay, D. (2013) Context-aware medical image retrieval for improved dementia diagnosis, in *Intelligent Multimedia Technologies for Networking Applications: Techniques and Tools* (ed. D. Kanellopoulos), IGI Global, Hershey, PA, USA, chap. 18, pp. 434–448.

25 Elger, B., Iavindrasana, J., Lo Iacono, L., Müller, H., Roduit, N., Summers, P., and Wright, J. (2010) Strategies for health data exchange for secondary, cross–institutional clinical research. *Computer Methods and Programs in Biomedicine*, **99** (3), 230–251.

26 Vannier, M.W. and Summers, R.M. (2003) Sharing images. *Radiology*, **228**, 23–25.

27 Grünberg, K., Jimenez-del-Toro, O., Jakab, A., Langs, G., Salas, T., Winterstein, M., Weber, M.A., and Krenn, M. (2017) Annotating medical image data, in *Cloud-Based Benchmarking of Medical Image Analysis*, Springer International Publishing.

28 Jimenez-del-Toro, O., Müller, H., Krenn, M., Gruenberg, K., Taha, A.A., Winterstein, M., Eggel, I., Foncubierta-Rodríguez, A., Goksel, O., Jakab, A., Kontokotsios, G., Langs, G., Menze, B., Salas Fernandez, T., Schaer, R., Walleyo, A., Weber, M.A., Dicente Cid, Y., Gass, T., Heinrich, M., Jia, F., Kahl, F., Kechichian, R., Mai, D., Spanier, A.B., Vincent, G., Wang, C., Wyeth, D., and Hanbury, A. (2016) Cloud–based evaluation of anatomical structure segmentation and landmark detection algorithms: VISCERAL Anatomy Benchmarks. *IEEE Transactions on Medical Imaging*, **35** (11), 2459–2475.

29 Krenn, M., Dorfer, M., Jimenez-del-Toro, O., Müller, H., Menze, B., Weber, M.A., Hanbury, A., and Langs, G. (2014) Creating a Large–Scale Silver Corpus from Multiple Algorithmic Segmentations, in *MICCAI Workshop on Medical Computer Vision: Algorithms for Big Data, MCV 2015*, vol. 8848, Springer, vol. 8848, pp. 163–170.

30 Hofmanninger, J., Krenn, M., Holzer, M., Schlegl, T., Prosch, H., and Langs, G. (2016) Unsupervised identification of clinically relevant clusters in routine imaging data, in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Springer International Publishing.

31 Aziz Taha, A. and Hanbury, A. (2015) Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. *BMC Medical Imaging*, **15** (1).

32 Armstrong, T.G., Moffat, A., Webber, W., and Zobel, J. (2009) Improvements that don't add up: ad-hoc retrieval results since 1998, in *CIKM '09: Proceeding of the 18th ACM conference on Information and knowledge management*, ACM, pp. 601–610, doi:http://doi.acm.org/10.1145/1645953.1646031.

33 Kalpathy-Cramer, J., García Seco de Herrera, A., Demner-Fushman, D., Antani, S., Bedrick, S., and Müller, H. (2015) Evaluating performance of biomedical image retrieval systems: Overview of the medical image retrieval task at ImageCLEF 2004–2014. *Computerized Medical Imaging and Graphics*, **39** (0), 55 – 61.

34 Heimann, T., Van Ginneken, B., Styner, M., Arzhaeva, Y., Aurich, V., Bauer, C., Beck, A., Becker, C., Beichel, R., Bekes, G. *et al.* (2009) Comparison and evaluation of methods for liver segmentation from CT datasets. *Medical Imaging, IEEE Transactions on*, **28** (8), 1251–1265.

35 Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., Lanczi, L., Gerstner, E., Weber, M.A., Arbel, T., Avants, B.B., Ayache, N., Buendia, P., Collins, D.L., Cordier, N., Corso, J.J., Criminisi, A., Das, T., Delingette, H., Demiralp, C., Durst, C.R., Dojat, M., Doyle, S., Festa, J., Forbes, F., Geremia, E., Glocker, B., Golland, P., Guo, X., Hamamci, A., Iftekharuddin, K.M., Jena, R., John, N.M., Konukoglu, E., Lashkari, D., Mariz, J.A., Meier, R., Pereira, S., Precup, D., Price, S.J., Raviv, T.R., Reza, S.M.S., Ryan, M., Sarikaya, D., Schwartz, L., Shin, H.C., Shotton, J., Silva, C.A., Sousa, N., Subbanna, N.K., Szekely, G., Taylor, T.J., Thomas, O.M., Tustison, N.J., Unal, G., Vasseur, F., Wintermark, M., Ye, D.H., Zhao, L., Zhao, B., Zikic, D., Prastawa, M., Reyes, M., and Van Leemput, K. (2015) The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS). *Medical Imaging, IEEE Transactions on*, **34** (10), 1993–2024.

36 Jimenez-del-Toro, O., Goksel, O., Menze, B., Müller, H., Langs, G., Weber, M.A., Eggel, I., Gruenberg, K., Holzer, M., Kotsios-Kontokotsios, G., Krenn, M., Schaer, R., Taha, A.A., Winterstein, M., and Hanbury, A. (2014) VISCERAL – VISual Concept Extraction challenge in RAdioLogy: ISBI 2014 challenge organization, in *Proceedings of the VISCERAL Challenge at ISBI* (ed. O. Goksel), Beijing, China, no. 1194 in CEUR Workshop Proceedings, pp. 6–15.

37 Rowe, B.R., Wood, D.W., Link, A.N., and Simoni, D.A. (2010) Economic impact assessment of NIST text retrieval conference (TREC) program, *Technical report project number 0211875*, National Institute of Standards and Technology.

38 Tsikrika, T., García Seco de Herrera, A., and Müller, H. (2011) Assessing the scholarly impact of ImageCLEF, in *CLEF 2011*, Springer Lecture Notes in Computer Science (LNCS), pp. 95–106.

39 Hanbury, A., Müller, H., Langs, G., Weber, M.A., Menze, B.H., and Fernandez, T.S. (2012) Bringing the algorithms to the data: cloud–based benchmarking for medical image analysis, in *CLEF conference*, Springer Lecture Notes in Computer Science.

40 Hanbury, A., Müller, H., Balog, K., Brodt, T., Cormack, G.V., Eggel, I., Gollub, T., Hopfgartner, F., Kalpathy-Cramer, J., Kando, N., Krithara, A., Lin, J., Mercer, S., and Potthast, M. (2015) Evaluation–as–a–service: Overview and outlook. *ArXiv*, **1512.07454**.

41 Syeda-Mahmood, T., Wang, F., Beymer, D., Amir, A., Richmond, M., and Hashmi, S.N. (2007) Aalim: Multimodal mining for cardiac decision support. *2007 Computers in Cardiology*, pp. 209–212.

42 Polikar, R., Tilley, C., Hillis, B., and Clark, C.M. (2010) Multimodal eeg, mri and pet data fusion for alzheimer's disease diagnosis. *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, pp. 6058–6061.

43 Kong, J., Cooper, L.A.D., Wang, F., Gutman, D.A., Gao, J., Chisolm, C., Sharma, A., Pan, T., Van Meir, E.G., Kurc, T.M., Moreno, C.S., Saltz, J.H., and Brat, D.J. (2011) Integrative, multimodal analysis of glioblastoma using tcga molecular data, pathology images, and clinical outcomes. *IEEE Transactions on Biomedical Engineering*, **58** (12), 3469–3474.

44 Davatzikos, C., Bhatt, P., Shaw, L.M., Batmanghelich, K.N., and Trojanowski, J.Q. (2011) Prediction of mci to ad conversion, via mri, csf biomarkers, and pattern classification. *Neurobiology of Aging*, **32** (12), 2322.e19–2322.e27.

45 Zhang, D., Wang, Y., Zhou, L., Yuan, H., and Shen, D. (2011) Multimodal classification of alzheimer's disease and mild cognitive impairment. *NeuroImage*, **55** (3), 856–867.

46 Zhang, D. and Shen, D. (2012) Multi-modal multi-task learning for joint prediction of multiple regression and classification variables in alzheimer's disease. *NeuroImage*, **59** (2), 895–907.

47 Westman, E., Muehlboeck, J.S., and Simmons, A. (2012) Combining mri and csf measures for classification of alzheimer's disease and prediction of mild cognitive impairment conversion. *NeuroImage*, **62** (1), 229–238.

48 Mattila, J., Koikkalainen, J., Virkki, A., van Gils, M., and Lötjönen, J.f.t.A.D.N.I. (2012) Design and application of a generic clinical decision support system for multiscale data. *IEEE Transactions on Biomedical Engineering*, **59** (1), 234–240.

49 Zacharaki, E.I., Morita, N., Bhatt, P., O'Rourke, D.M., Melhem, E.R., and Davatzikos, C. (2012) Survival analysis of patients with high-grade gliomas based on data mining of imaging variables. *American Journal of Neuroradiology*, **33** (6), 1065.

50 Tay, W.L., Chui, C.K., Ong, S.H., and Ng, A.C.M. (2013) Ensemble-based regression analysis of multimodal medical data for osteopenia diagnosis. *Expert Systems with Applications*, **40** (2), 811–819.

51 Ribeiro, R.T., Marinho, R.T., and Sanches, J.M. (2013) Classification and staging of chronic liver disease from multimodal data. *IEEE Transactions on Biomedical Engineering*, **60** (5), 1336–1344.

52 Mourão, A., Martins, F., and Magalhães, J. (2015) Multimodal medical information retrieval with unsupervised rank fusion. *Computerized Medical Imaging and Graphics*, **39**, 35–45.

53 Schlegl, T., Waldstein, S.M., Vogl, W.D., Schmidt-Erfurth, U., and Langs, G. (2015) *Predicting Semantic Descriptions from Medical Images with Convolutional Neural Networks*, Springer International Publishing, Cham, pp. 437–448.

54 Shin, H.C., Lu, L., Kim, L., Seff, A., Yao, J., and Summers, R.M. (2015) Interleaved text/image deep mining on a large-scale radiology database. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1090–1099.

55 Kitanovski, I., Strezoski, G., Dimitrovski, I., Madjarov, G., and Loskovska, S. (2016) Multimodal medical image retrieval system. *Multimedia Tools and Applications*, pp. 1–24.

56 Ramos, J., Kockelkorn, T.T.J.P., Ramos, I., Ramos, R., Grutters, J., Viergever, M.A., van Ginneken, B., and Campilho, A. (2016) Content-based image retrieval by metric learning from radiology reports: Application to interstitial lung diseases. *IEEE Journal of Biomedical and Health Informatics*, **20** (1), 281–292.

57 Zhang, F., Song, Y., Cai, W., Hauptmann, A.G., Liu, S., Pujol, S., Kikinis, R., Fulham, M.J., Feng, D.D., and Chen, M. (2016) Dictionary pruning with visual word significance for medical image retrieval. *Neurocomputing*, **177**, 75–88.

58 Rhodius-Meester, H.F.M., Koikkalainen, J., Mattila, J., Teunissen, C.E., Barkhof, F., Lemstra, A.W., Scheltens, P., Lotjonen, J., and van der Flier, W.M. (2016) Integrating biomarkers for underlying alzheimer's disease in mild cognitive impairment in daily practice: Comparison of a clinical decision support system with individual biomarkers. *J Alzheimers Dis*, **50** (1), 261–270.

59 Kim, T.H., Jeong, D.J., Hahn, S.Y., Shin, J.H., Oh, Y.L., Ki, C.S., Kim, J.W., Jang, J.Y., Cho, Y.Y., Chung, J.H., and Kim, S.W. (2016) Triage of patients with aus/flus on thyroid cytopathology: effectiveness of the multimodal diagnostic techniques. *Cancer Med*, **5** (5), 769–777.

60 Jatrniko, W., Arsa, D.M.S., Wisesa, H., Jati, G., and Ma'sum, M.A. (2016) A review of big data analytics in the biomedical field. *2016 International Workshop on Big Data and Information Security (IWBIS)*, pp. 31–41.

61 Shin, H., Roberts, K., Lu, L., Demner-Fushman, D., Yao, J., and Summers, R.M. (2016) Learning to read chest x-rays: Recurrent neural cascade model for automated image annotation. *CoRR*, **abs/1603.08486**.

62 Wang, X., Lu, L., Shin, H., Kim, L., Nogues, I., Yao, J., and Summers, R.M. (2016) Unsupervised category discovery via looped deep pseudo-task optimization using a large scale radiology image database. *CoRR*, **abs/1603.07965**.

63 Nie, D., Zhang, H., Adeli, E., Liu, L., and Shen, D. (2016) *3D Deep Learning for Multi-modal Imaging-Guided Survival Time Prediction of Brain Tumor Patients*, Springer International Publishing, Cham, pp. 212–220.

64 Burth, S., Kickingereder, P., Eidel, O., Tichy, D., Bonekamp, D., Weberling, L., Wick, A., Löw, S., Hertenstein, A., Nowosielski, M., Schlemmer, H.P., Wick, W., Bendszus, M., and Radbruch, A. (2016) Clinical parameters outweigh diffusion- and perfusion-derived mri parameters in predicting survival in newly diagnosed glioblastoma. *Neuro-Oncology*, **18** (12), 1673–1679.

65 Chang, K., Zhang, B., Guo, X., Zong, M., Rahman, R., Sanchez, D., Winder, N., Reardon, D.A., Zhao, B., Wen, P.Y., and Huang, R.Y. (2016) Multimodal imaging patterns predict survival in recurrent glioblastoma patients treated with bevacizumab. *Neuro-Oncology*, **18** (12), 1680–1687.

66 Anavi, Y., Kogan, I., Gelbart, E., Geva, O., and Greenspan, H. (2016) Visualizing and enhancing a deep learning framework using patients age and gender for chest x-ray image retrieval, vol. 9785, pp. 978 510–978 510–6.

67 Moradi, M., Guo, Y., Gur, Y., Negahdar, M., and Syeda-Mahmood, T. (2016) *A Cross-Modality Neural Network Transform for Semi-automatic Medical Image Annotation*, Springer International Publishing, Cham, pp. 300–307.

68 Xu, T., Zhang, H., Huang, X., Zhang, S., and Metaxas, D.N. (2016) *Multimodal Deep Learning for Cervical Dysplasia Diagnosis*, Springer International Publishing, Cham, pp. 115–123.

69 Kooi, T., Litjens, G., van Ginneken, B., Gubern-Mérida, A., Sánchez, C.I., Mann, R., den Heeten, A., and Karssemeijer, N. (2017) Large scale deep learning for computer aided detection of mammographic lesions. *Medical Image Analysis*, **35**, 303–312.

70 Zhang, B., Chang, K., Ramkissoon, S., Tanguturi, S., Bi, W.L., Reardon, D.A., Ligon, K.L., Alexander, B.M., Wen, P.Y., and Huang, R.Y. (2017) Multimodal mri features predict isocitrate dehydrogenase genotype in high-grade gliomas. *Neuro-Oncology*, **19** (1), 109–117.

71 Depeursinge, A. and Müller, H. (2010) Fusion techniques for combining textual and visual information retrieval, in *ImageCLEF*, *The Springer International Series On Information Retrieval*, vol. 32 (eds H. Müller, P. Clough, T. Deselaers, and B. Caputo), Springer Berlin Heidelberg, pp. 95–114.

72 LeCun, Y., Bengio, Y., and Hinton, G. (2015) Deep learning. *Nature*, **521** (7553), 436–444.

73 Ravì, D., Wong, C., Deligianni, F., Berthelot, M., Andreu-Perez, J., Lo, B., and Yang, G.Z. (2017) Deep learning for health informatics. *IEEE Journal of Biomedical and Health Informatics*, **21** (1), 4–21.

74 Litjens, G.J.S., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A.W.M., van Ginneken, B., and Sánchez, C.I. (2017) A survey on deep learning in medical image analysis. *CoRR*, **abs/1702.05747**.

75 Greenspan, H., van Ginneken, B., and Summers, R.M. (2016) Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique. *IEEE Transactions on Medical Imaging*, **35** (5), 1153–1159.

76 Lin, F.Y., Dunning, A.M., Narula, J., Shaw, L.J., Gransar, H., Berman, D.S., and Min, J.K. (2013) Impact of an automated multimodality point-of-order decision support tool on rates of appropriate testing and clinical decision making for individuals with suspected coronary artery disease: A prospective multicenter study. *Journal of the American College of Cardiology*, **62** (4), 308–316.

77 Thornley, C.V., Johnson, A.C., Smeaton, A.F., and Lee, H. (2011) The scholarly impact of TRECVid (2003–2009). *Journal of the American Society for Information Science and Technology*, **62** (4), 613–627.

78 Fraser, A.G. and Dunstan, F.D. (2010) On the impossibility of being expert. *BMJ*, **341**.

79 Kayadjanian, N., Barros, H., Billig, H., Capron, M., Cohen, A., Dubochet, G., Frackowiak, R., Grüters-Kieslich, A., Holgate, S., Kuster, S., Pozzan, T., Vingron, M., Wolff-Boenisch, B., and Radwanska, M. (2014) How to transform big data into better health: Envisioning a health big data ecosystem for advancing biomedical research and improving health outcomes in europe, *Tech. Rep.*, Science Europe: Medical Sciences Committee.