Enterprise imaging and big data: A review from a medical physics perspective

Nicholas McCarthy^a, Adil Dahlan^a, Tessa S. Cook^b, Neil O' Hare^c, Marie-Louise Ryan^a, Brendan St John^a, Aonghus Lawlor^d, Kathleen M. Curran^{a,*}

^a School of Medicine, University College Dublin, Ireland

^b Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA, USA

^c School of Physiotherapy and Performance Science, University College Dublin, Ireland

^d School of Computer Science, University College Dublin, Ireland

ABSTRACT

In recent years enterprise imaging (EI) solutions have become a core component of healthcare initiatives, while a simultaneous rise in big data has opened up a number of possibilities in how we can analyze and derive insights from large amounts of medical data. Together they afford us a range of opportunities that can transform healthcare in many fields. This paper provides a review of recent developments in EI and big data in the context of medical physics. It summarizes the key aspects of EI and big data in practice, with discussion and consideration of the steps necessary to implement an EI strategy. It examines the benefits that a healthcare service can achieve through the implementation of an EI solution by looking at it through the lenses of: compliance, improving patient care, maximizing revenue, optimizing workflows, and applications of artificial intelligence that support enterprise imaging. It also addresses some of the key challenges in enterprise imaging, with discussion and examples presented for those in systems integration, governance, and data security and privacy.

1. Introduction

An enterprise imaging (EI) strategy is an organized plan to optimize the electronic health record (EHR) so that healthcare providers and patients have intuitive and real-time access to all patient clinical images and their associated documentation, regardless of source. Primo et al. [1] proposed a road-map to achieve the goal of implementing EI for an institution that firstly requires defining and accessing all images and data used for medical decision-making and understanding the specialties and their clinical workflow challenges as related to imaging. Creating a strategy to implement an EI infrastructure, with appropriate privacy, security and governance, will improve quality of care and patient safety, for example, by eliminating the number of unnecessary repeat diagnostic tests that may include exposure to radiation. Using artificial intelligence (AI) to streamline patient diagnosis and prognosis will not only enhance patient care and advance population health but will also improve the effectiveness of the healthcare service with business intelligence solutions to optimise operations and reduce costs. Current AI predictive tools and emerging future applications will enable clinicians to make real-time remote decisions at the time of data acquisition, while the increasing adoption and prevalence of electronic health records will empower both clinicians and patients by facilitating

appropriate access to relevant data.

The main obstacles to development and clinical implementation of AI models include availability of sufficiently large, curated, and representative training data that includes expert annotations. In order for big data to transform medicine, data must be annotated and models trained, analyzed and interpreted. Machine learning algorithms as currently applied in medicine are closer to expert systems or rule sets encoding clinical knowledge, which are applied to draw conclusions about specific clinical outcomes.

The power of machine learning is in handling enormous numbers of predictors and applying them in non-linear and highly-interactive ways. An enterprise imaging framework creates an opportunity to develop models linking covariates, such as image features and genomics data with clinical outcomes across a spectrum of data at different geographical sites. We can use emerging AI deep learning technology to start with patient level observations and sift through vast numbers of complex variables to search for combinations that reliably predict clinical outcomes. However, we need to be cautious when curating existing data in EHRs and processing the data to become usable as high quality, unbiased data to train deep learning models.

The ability to derive clinical insight from big data will disrupt many areas of medicine [2]. In the following sections we discuss each of these

* Corresponding author.

E-mail address: kathleen.curran@ucd.ie (K.M. Curran).

<https://doi.org/10.1016/j.ejmp.2021.04.004>

Received 7 December 2020; Received in revised form 24 March 2021; Accepted 6 April 2021

Available online 30 April 2021

1120-1797/© 2021 Published by Elsevier Ltd on behalf of Associazione Italiana di Fisica Medica.

areas in the context of enterprise imaging. We discuss how big data and machine learning will alter the work of radiologists and pathologists, how it will improve diagnostic accuracy and reduce errors, and how the increased visibility of data in EI systems will lead to better prognostic models. In section one we introduce the key aspects of enterprise imaging, consider how EI solutions are currently implemented, and discuss emerging requirements for future adoption. Also discussed are aspects of system and information governance, legislation and jurisdictional differences, data privacy and security, how information from disparate sources can be standardised and consolidated, and how workflows and business intelligence systems can be optimized. In section two we consider the potential impact of EI on primary care, in terms of the challenges of workflow optimization and the potential benefits to multi-site patient care. Finally, in section three we discuss how current diagnostic and predictive AI tools have been deployed in the clinic and how emerging techniques could be translated into clinically operational EI systems. Fig. 1.

2. Enterprise imaging

In this section we discuss the concept of enterprise imaging and consider how EI solutions are currently implemented and emerging requirements for future adoption, including system and information governance and data privacy and security. Additionally, we discuss how information from disparate sources can be standardised and consolidated and how workflows and business intelligence systems can be optimized.

The concept of enterprise imaging in the health context, and in particular the acute hospital sector, has emerging capability but is not yet clearly defined. The original concept of enterprise imaging revolved around an expansion of radiology imaging via picture archive and communication systems (PACS) solutions to cater for growing imaging demands in cardiology, vascular surgery, and to an extent obstetrics. These were relatively easily managed as the generated images, and indeed the workflows, were often very similar or indeed identical to those found within radiology. The expansion of PACS systems to incorporate these specialities was unusual but not hugely problematic. (See Fig. 2).

The term enterprise imaging has grown in recent years, but the specific definition as to what EI actually is or represents can be varied [3]. More recently the collaborative HIMSS-SIIM member work-group (respectively, the Healthcare Information and Management Systems Society, and Society for Imaging Informatics in Medicine) defined a set of strategies, initiatives and workflows that could be implemented across a healthcare enterprise to consistently and optimally capture, index, manage, store, distribute, view, exchange and analyze all clinical

imaging and multimedia content to enhance the electronic health record [1,4,3,5,6]. The diversity of imaging that is captured ranges from radiology or radiology-type imaging, which constitutes the main component to clinical photography (e.g. dermatology), maxillo-facial reconstructions and video consultations.

Enterprise imaging solutions have an additional layer of complexity when deployed in a single hospital, across a hospital network, or on a regional or national basis. While the principles of image ingestion, indexing, linking to a report, extraction and viewing are similar in all contexts, aspects such as the wider range of imaging modalities, the size and growth of the archive, and in particular, the indexing becomes more complex as the archive extends its coverage [7,8].

Over the last decade or more there has been a growth in the implementation and use of vendor neutral archives (VNAs). These have traditionally been third-party image stores, moving away from the proprietary storage solutions by many PACS vendors. This growth came out of a need to share imaging data between departments, hospitals, and caregivers in a less complex fashion requiring a less onerous interface or integration implementations and facilitating easier and less costly future expansions. In addition, the incorporation of other image data types, such as those found in ophthalmology where the DICOM standard was not as well developed further drove the development of VNAs towards a true enterprise imaging ecosystem [9].

EI governance is becoming increasingly important in healthcare enterprises. In the following sections we discuss the decision-making body, framework, and process for optimal EI governance inclusive of five areas of focus: program governance, technology governance, information governance, clinical governance, and financial governance. It outlines the relevant parallels and differences when forming or optimizing imaging governance as compared with other established broad horizontal governance groups, such as for the electronic health record. It includes the advantages and efficiencies that implementing an EI solution can bring to revenue maximization, staff productivity, and administrative workflows, as well as some challenges faced in the areas of data governance, security, and privacy.

2.1. System and information governance

Gartner Inc. defines two elements of information governance and data governance [10]. The former is the process of making decisions about information to drive important change. The aim of an information governance program being decisions and actions, making sure that opportunities are understood and priorities are established. Data governance is defined as that which supports strategic imperatives of ensuring that data is available, accurate, safe and accessible. However, it is stressed that not all data is equal or needs to be governed in the same

Fig. 1. The 10 steps to implementing an enterprise imaging strategy as defined by Primo et. al. [1].

Fig. 2. The enterprise imaging paradigm – an organized plan to optimize the electronic health record so that healthcare providers and patients have intuitive and real-time access to all patient clinical images and their associated documentation, regardless of source [1].

way resulting in a complex matrix of data types, applications, timeframes, and data usage [10]. This is further complicated by secondary use of data to support research and education. On such secondary use, the American Medical Informatics Association (AMIA) has reflected that any such data should be “as consistent, comparable, timely, accurate, accessible, complete, and reliable as possible” [11]. Anonymization of clinical imaging datasets can be challenging where such images include ‘burned in’ data often commonly found in video content, secondary capture and other medical images [5]. The use of an enterprise imaging solution as described raises the need for a comprehensive governance framework and function. Enacting the GDPR principle of “privacy by design” with the completion of comprehensive risk assessment via the data protection impact assessment becoming essential. An EI solution that spans multiple hospitals, clinicians, systems, etc. further complicates this process.

2.2. Jurisdiction and legislative differences

In the context of healthcare, individual data must be collected, stored and processed for a number of reasons: direct clinical care, clinical research, insurance, and population health. Many applications of healthcare requires personal data to be transmitted between organisations and across national borders, and thus jurisdictional and legislative differences must be evaluated when we consider the role of “Big Data” in EI.

The General Data Protection Regulation (GDPR) came into effect in the EU in 2018 [12], introducing safeguards with respect to personal data, and places healthcare data in a special category. In the US the regulations around patient data are defined by the Health Insurance Portability and Accountability Act (HIPAA) [13], created to secure protected health information (PHI) by regulating healthcare providers.

There are numerous differences between the two regulations, not just in their geographic coverage; GDPR may still apply to EU citizens visiting a third country for treatment, whereas HIPAA does not cover healthcare data outside of the US; GDPR places a premium on patient consent, whereas under HIPAA healthcare organisations can transmit patient data to third party processors under certain circumstances and provided it is stored and transmitted securely; HIPAA does not grant the right to be forgotten, while GDPR allows patients request their personal data be deleted [14].

Other countries have their own data protection regulations that offer various degrees of protection to personal healthcare data – although GDPR is generally regarded as the most stringent regulations of its type in the world – but what’s clear is that navigating all of the jurisdiction and legislative differences is a significant challenge facing healthcare providers.

2.3. Big data and enterprise imaging

“Big Data” can be defined simply as “data that is too large to process with current database systems”. This can be further delineated into “the three Vs” of big data: the volume, velocity, and variety of data – each of which presents unique challenges to address, and which depends on the nature of the data and kind of storage, management or analysis being performed. The costs associated with providing and maintaining an on-premises system that can effectively cope with this kind of data is prohibitive to most providers. Cloud computing services are increasingly used in order manage data at this scale, as this enables ubiquitous, on-demand and configurable computing resources that can be provisioned and decommissioned with minimal effort. As the volume of electronic health records (EHRs) produced in healthcare departments continues to grow, the increased adoption of cloud computing within

healthcare facilities means that providers would have cost-effective, scalable, reliable, easy to maintain access to computing resources for data storage and processing [15], and as a result the healthcare cloud computing market is expected to grow to \$64.5bn by 2025 [16].

Cloud computing services can be placed in three categories: *Software as a Service (SaaS)*, where the customer uses an application that the provider hosts/runs on its cloud infrastructure; *Platform as a Service (PaaS)*, where the customer is provided with a cloud infrastructure to deploy consumer-acquired or created applications without worrying managing or controlling the underlying cloud infrastructure; *Infrastructure as a Service (IaaS)*, where the consumer has control over the storage, operating system and deployed application the are provided but no control over the underlying cloud infrastructure [17,15]. The increasing emphasis on cloud-based infrastructure within digital healthcare environments has led to an influx of technology companies into the space, with Amazon, Apple, Google, Microsoft and IBM all providing healthcare platforms offering a variety of services [18].

2.4. Integration of EI and electronic health records

Prior to the development of EI solutions, the prevalence of electronic health records (EHR) was established. By way of example, the Health Information Technology for Economic and Clinical Health Act (HiTECH) was enacted under the American Recovery and Reinvestment Act in 2009. Under this Act, the United States Department of Health and Human Services resolved to spend just under \$26bn to promote and expand the adoption of health information technology [19]. The result of such investment was an increase from 9% in 2008 to 96% in 2015 of hospitals that had at least a basic EHR in situ [20]. The HIMSS, founded in 1961, has performed maturity assessments on the digital adoption of hospitals for many years. One such model is the Electronic Medical Records Adoption Model (EMRAM) which is an eight-level assessment based digital adoption scale. In recent years they have brought forward a digital imaging adoption model to reflect the growth in the imaging field and in particular the movement towards enterprise imaging giving evidence of the link between EHRs and EI solutions.

2.5. Enterprise image viewing

While the DICOM standard has been addressing a range of different image data types in various supplements to the base standard [21], the viewing requirements are not met by a single viewing platform. The viewer optimised for radiology may well display still photos from dermatology or indeed endoscopy. However, their ability to display recorded video consults with required audio elements or digital pathology slides at resolutions and file sizes significantly greater than their radiology counterparts is lacking. A whole slide pathology image can have a file size ranging from 0.5 GB to over 10 GB depending on the level of compression and the extent of the tissue sample [22]. In comparison a 1,200 image multi-slice CT is about 0.6 GB. Standard radiology image viewers would struggle with the larger digital pathology file size and are not optimised for such. The potential need for multiple viewers is therefore a real requirement. Areas where image viewing and analysis capability is only required by the speciality users include: cardiology ECG/EKG waveform analysis; histopathology whole-slide imaging, with features such as cell counts; obstetric foetal-growth calculations; radiology anatomic and perfusion calculations and surgical planning. Enterprise imaging solution viewers often do not include functionality in these and other speciality specific areas requiring the need for specialised viewers.

2.6. Revenue maximization and management

The pressure of reducing costs and the focus on value is a global trend in healthcare, and value-based strategies have become a concern and core component of medical practice. In a 2014 survey of 196 healthcare

leaders it was noted that the USA wasted between \$7.5–12 billion annually on medical imaging, the reasons for this being cited as: defensive medicine, lack of appropriate testing, and patient demands for imaging [23]. Thus, there is a focus on cost reduction in radiology - and indeed across healthcare - and it forms one of the Institute for Healthcare Improvements quadruple aims [24].

In order to discuss the opportunities to reduce costs by using EI, we have broadly divided the issue into the following areas: consolidation of resources, improving workflows, staff productivity, and business analytics, with some discussion and examples presented.

2.6.1. Consolidation of HR and IT resources

There is a global trend in organizational design to move towards a decentralized structure, and the use of outsourcing and consolidation of resources. Introducing EI allows for several resources to be amalgamated, such as IT and human resources and image storage. Not only that, but costly areas such as service contracts, support, licences, vendors and maintenance can be reviewed and reduced. This does not just change the financial outlay but there is also a reduction in the number of contacts and contracts which reduce workload and increase efficiency. Sirota-Cohen et al. cited a 31.5% reduction in storage costs alone for one university hospital that merged three departments into one single VNA and a 10–15% overall decline in costs due to reductions in maintenance, supports and service contracts [9].

Implementing EI and the resulting consolidation of staff and resources should be appropriately planned. Diverse areas should be considered with appropriate timelines, including - but not limited to - staff training, legal implications of consolidation, appropriate staff job re-design and allocation. On a local level, there is certainly a challenge in implementing this in a manner that is truly a cost-effective move. Changing structures and changing job design for individuals can be a challenging task, albeit is the hallmark of successful businesses [25]. Individuals and teams naturally resist change, and any company needs staff engagement and buy in effectively, which requires strong leadership [26]. The maximisation of revenue and streamlined management, however, are a significant reward for this challenge.

2.6.2. Improving administrative workflows

Operational efficiency is defined as high-quality output for the least amount of input of which cost is a key factor [8]. Lean management and workflow standardization focus on the minutiae of day-to-day operations where cost savings can be made [27]. Significant opportunities exist with the introduction of EI. It is perhaps fundamental to walk through the patient journey when considering the impact on revenue for EI. Core to this journey is the administrative process. The administration team book patients for appointments, explain patient preparation and deal with patient queries, greet patients and 'arrive' them into the department, and carry out additional patient tasks such as dealing with patient requests for imaging, booking of follow up imaging and ensuring the accuracy of patient details. An example of this would be booking an inpatient at an appropriate time so that it does not clash with another department; ensuring the patient has not had the same imaging test recently; aligning the appointment time to clinical urgency or specific patient needs. All of these tasks can be addressed and improved with EI. One example of this is looking at encounter-based workflows, whereby all imaging studies (regardless of speciality) are represented by a single tab and thus the existence of a study is evident. Modality worklists can aid in automating imaging workflow by reducing manual data entry [8]. A small but significant improvement that is evident with EI is the reduction or elimination of physical media, namely CDs. CDs are continuously used for sharing patient data both within a hospital and between hospitals [28]. In 2019, it was reported that 80% of facilities still used CDs, 84 days per year are wasted on CD related tasks and an average of €78K per year is saved by eliminating CDs [29].

2.6.3. Staff productivity

Significant amounts of healthcare staff working hours are spent managing image and data requests from patients. There is also a noted disparity that administrators and clinical staff may likely differ on their focus on quality and efficiency [30]. Turnaround times for image reporting have been reported as being faster and more efficient (in particular in the area of mobile reporting) within EI frameworks. This was discussed specifically in light of the COVID-19 pandemic, where staff may be required to self-isolate at home. The opportunity has been afforded in these instances for staff to be able to access patient data and images and to continue working [31]. Efficiency between groups of staff has been noted: for example, if a patient attends a surgical consultation, all the information is available from other departments, therefore reducing the time required to have conversations and follow ups regarding specific patient needs [32]. It is estimated that radiologists spend about one third of their time doing administration work, and one would anticipate that with the use of EI, this would reduce significantly [33].

2.6.4. Business analytics

The role of radiology imaging has increased dramatically in recent years. It was reported pre-Covid pandemic that the demand for radiology is increasing year on year, due to imaging remaining central to diagnostic work up [34]. Never has it been more important to assess and analyse radiology departments to ensure efficient use of resources, streamlined workflows and forward planning.

Business analytics include forecasting, statistical analysis and predictive modelling. Data analysis technologies allow for clinicians and management to access dynamic data sources to support process improvement and facilitates meaningful reporting. Key performance indicators (KPIs) can be tracked effectively, which feeds into a quality-focused organisation, with clinical and non-clinical governance at the core [35]. Common applications for analytics in radiology include performance and operations assessment, dose monitoring and predictive analysis [36].

The endpoint in radiology is the production of a radiology report, but there are numerous steps along the way that are imperative to understand to run an efficient service. Shrestha et al. (2014) rightly say in their discussion on EI analytics that you can only improve what you measure [37]. Examples include:

- Patient types: patients arriving on hospital trolleys, high infection control risk and patients with certain clinical conditions impact workflow and productivity. These patients likely need more time for an examination to be completed and possibly additional staff. An understanding of this also feeds into being patient-centric.
- Productivity of radiographers/technologists: how long does it take to complete an examination, number of examinations completed per staff member, what types of examinations are being completed. These metrics can provide analysis on not just examination time optimization, but human resource areas required such as staff training.
- Assessment of scanner uptime: this is a key metric, considering current equipment costs. Use of this metric can analyse gaps in use of equipment and also help to proactively assess when new equipment is required.

Dynamic access to these metrics, together with others, allows for rules-based workflows, triage, trend analysis and task optimization. Data is available to better predict and act, both clinically and non-clinically, feeding into key advantages for organisation, staff and patients [31]. Radiology quality programs have become embedded in clinical practice in many countries. There are several KPIs which are required to be reported at national level, which often include report turnaround times, audit activity and timeliness of response for urgent reporting. Business analytics give access to these metrics, to ensure a

transparent adherence to requirements, which are often burdened with manual extraction for already busy radiologists. It allows organisations to also benchmark themselves as part of national programmes to understand their performance across multiple metrics. Because of the nature of business analytics linked to EI, it can serve as a component in root cause analysis on areas of concern.

Indeed, a further consideration is the linking of EI with radiation dose analytics which provides a key area of growth and understanding in the international focus on reducing patient radiation dose. Tracking of this data is a legislative requirement in many countries and there are several current options around dose tracking systems [38]. What is missing, and what EI offers, is for other specialities to be able to view patient radiation dose. Take for example a female of childbearing age who has had several CT examinations, yielding a high and potentially harmful cumulative dose. A cardiologist can review this on ordering and request an alternative modality or test in line with optimizing patient safety. Considering the requirements to audit these types of radiation safety initiatives, this could be tracked dynamically and with integrity with the appropriate EI system in place.

Based on the above discussion and metrics and current literature, the use of business analytics in radiology operations plays a key role in understanding business needs [39]. Manual and ad hoc understanding of KPIs is simply not sufficient in the current climate. Outsourcing of radiology services and teleradiology is on the rise. Private radiology services that provide outsourced imaging solutions are often required to provide comprehensive reporting to show their impact, which is core to them growing their business and proving that they are using best initiatives in line with cost. Radiology business needs to be able to respond to the current climate, by analysing workflows and planning resources appropriately, transforming data collection to operational correction and being able to report metrics quickly and with integrity [40].

2.7. Automated reporting

Although many of the tangible advantages of EI relate to more efficient and effective image sharing and improved patient experience, these are less evident for critically important benefits to automation of administrative reporting. The value of business analytics in a large healthcare enterprise cannot be emphasized enough [5]. However, accessing the right data to generate the reports, which often drive critical decisions around investments and hiring, may be just as difficult if not more so than accessing patients' data across treatment facilities. An effective EI strategy can decrease the inefficiency of wrangling operational data stored in silos across an enterprise [8]. While this typically requires up-front investment to standardize these data, that process should be part of an effective, well-planned EI implementation. Unifying the management of imaging data both within radiology as well as in image-producing specialties outside of radiology can reap benefits in operational efficiency, decreased liability (secondary to improved patient safety), and an eventual return on investment [3,9].

2.8. Data security

Cybersecurity is a major source of concern within any digital healthcare system [8]. Increasingly digitized healthcare environments have led to the health sector becoming a major target for cyberattacks, with a recent SANS Institute report reporting that 94% of health care organizations have been the victim of a cyberattack [41].

Hospitals are especially vulnerable to these attacks, as any disruption in operations or disclosure of patient personal information can have far-reaching and damaging consequences [42]. The increased connectivity of medical devices to computer networks leaves them vulnerable to security breaches, in much the same way as other networked computing systems are vulnerable [43].

Vendor support and assistance is provided by remote access for many health systems, which on one hand simplifies the overall support and

system diagnostic measures, but also exposes a potential site of breach. Embedded web services for system or device administration are a leading vulnerability threat. These services often remain accessible when the device or system itself is not in operation, many with default passwords remaining from the time of installation, and with insufficient – or completely without – authentication requirements for access [8].

An EI initiative helps to reduce the risk of cyberattacks, as it increases visibility to hardware and software within the imaging ecosystem [44]. This increased visibility allows devices and systems to be evaluated regularly and in an automated fashion to determine whether or not they are compliant with the institution's security policies, and non-compliant systems can be flagged and decommissioned or brought into compliance.

EI also helps to increase awareness among all personnel in the imaging ecosystem. Many healthcare data breaches are simply due to human error, and one way to mitigate against this risk is through the development of a security minded culture. The heightened awareness, increased collaboration, and communication generated via enterprise governance bodies will be a strong factor in building this culture [8].

2.9. Data privacy

Leveraging real-world “big data” in healthcare domains for statistical reporting, analysis, and machine learning tasks requires addressing many practical challenges, such as how to efficiently train models while also being privacy-aware. Healthcare data is often held in digital silos, distributed across multiple sites, and in situations where creating a centralized data repository is either unfeasible, or constrained by resource, compliance, and privacy concerns [8].

Federated learning has been proposed as a general framework and machine learning paradigm that would enable training of statistical models from distributed data sources [45–47]. This allows hospitals, medical institutions and consortiums to collaborate in a privacy-aware manner and simplifies some aspects of data governance, as a ‘consensus’ model can be trained on patient data without that data ever leaving the hospital's IT network. However, even once the data can be analyzed using a machine learning toolkit, it has been demonstrated that machine learning models can ‘leak’ the data on which they were trained, thereby violating individual privacy if the training data includes personally identifiable features [48], or ‘quasi-identifiers’ – combinations of features such as age, gender, nationality, etc that can uniquely fingerprint a patient [49]. A number of privacy models have been put forward to address this issue, the most notable of which is differential privacy [50], which is also incorporated as the primary mechanism for introducing privacy in most federated learning frameworks [46]. Among others, this has been studied in the context of its application to medical health records by Choudhury et al. [51], where it was demonstrated to work with varying privacy levels that allowed user-level privacy guarantees with only minor losses of model accuracy and utility. However, as discussed by Geyer et al. [52], achieving a high degree of model utility while maintaining an appropriate level of privacy is generally only possible when operating over thousands of sites, and as such may not be appropriate for small healthcare providers.

The application of differential privacy to training computer vision models in various medical imaging modalities and fields has also been studied, and includes work on chest X-ray [53], neuroimaging [54], and whole-slide digital pathology [55]. It should be noted that in many of these instances, beyond the removal of identifying information “burned” into the image in the form of labels there is perhaps a low probability of individual privacy being violated to any significant degree – the major exception to this is the possibility of reconstruction and subsequent identification of patient faces from CT and MRI data [56].

3. Leveraging EI in the clinic

In this section we discuss enterprise imaging in the context of its

active use by clinicians, considering its potential impact at the point of care in terms of workflow optimization, and the benefit it can bring to patient care. Also discussed is the concept of ubiquitous image access, and drawing parallels between it and the introduction of EHRs in terms of the potential advantages it can bring to multi-site patient care. Finally, we discuss some of the challenges currently found in clinical workflows.

3.1. Workflow optimization

EI has the potential to optimize clinical workflow, which in turn, can improve efficiency across the healthcare enterprise [8]. Historically, different types of imaging examinations have been siloed in different departments. Over the past three decades, radiology departments have adopted the picture archiving and communications system (PACS) and more recently, the vendor-neutral archive (VNA) for image storage. However, other image-producing specialties that use DICOM-based imaging, such as cardiology, have often used separate systems. Furthermore, specialties such as dermatology, which relies on photographs, have their own ad hoc storage systems, which may not connect to the EHR or even be as secure as PACS and other mainstream systems.

The inherent inefficiencies associated with having different types of studies stored in different systems, inaccessible to one another, and potentially susceptible to cybersecurity attack, serve to make the case for EI as a solution to multiple problems. Having a comprehensive EI solution that makes different imaging examinations accessible across the healthcare enterprise, to different users via their preferred viewers, immediately increases the efficiency of care delivery by decreasing the number of systems to which clinicians need to request access. EI further decreases the problem of data being unavailable at the point of care, which can lead to unnecessary delays or degrade the quality of care delivered.

EI also allows for a more holistic view of the patient. When subspecialty care spans more than one clinical expert, such as between radiology and cardiology, cross-enterprise image access offers clinicians a more complete picture of the patient's workup to date, and effectively increases the quality of care delivered by the clinicians in the respective departments. Thus, EI facilitates communication and collaboration between physicians in different specialties across the enterprise, which translates to better patient care and likely a more positive patient experience.

By virtue of its potential improvements to the clinical workflow – and when implemented correctly – EI becomes part of the value proposition for health care systems [8]. It has the potential to increase the completeness and utility of imaging and ancillary information available to clinicians in different departments, decrease the hoops through which patients must sometimes jump to access and transfer their images between clinicians, increase data privacy and security related to imaging and imaging equipment, and provide data compromise and loss.

One of the many practical considerations in implementing EI within the clinical workflow is the decision between orders-based and encounters-based imaging [4]. Orders-based imaging is somewhat more traditional, in that the request for the specific examination, i.e., the order, results in an independent care visit solely for the purposes of obtaining the imaging evaluation. By comparison, encounters-based imaging is often performed during a related clinical evaluation, and serves as a direct complement to the physical examination of the patient during that visit. Unlike orders-based imaging, encounters-based imaging is often unplanned, but performed when the need arises in the course of evaluating the patient.

3.2. Improving patient care

The patient perspective on EI is core to the success of how it is implemented and maintained. All considerations in EI should be patient centric, with the patient experience and safety being considered

throughout [1].

The fundamental premise of EI systems providing clinicians with a coherent system amplifies patient care [1]. If a patient is going through the hospital system in various departments, this care can be compromised. The storage of images and related data associated with one patient on multiple systems indeed poses a risk to their standard of treatment and cross disciplinary care. Payne et al. (2012) discuss this issue and conclude that streamlining patient data using one electronic medical record is an important prerequisite for risk minimization in healthcare [57].

There are several key tasks in all imaging applications where the use of EI can enhance patient care. These include reviewing previous examinations to inform the interpretation of the current exam. This could be achieved using EI for not just one speciality, but across a range of specialities. For example, interpretation of cardiac MRI and cardiac CT benefits from understanding of prior echocardiography, but the latter images are often stored in a separate system that may not be accessible to the CT and MRI readers.

Another example of this is comprehensive review of pertinent clinical information. The referral request encompasses not just clinical information, but past medical history, current medications, workup and evaluation to date, and current health status. Providing a clinician with the most comprehensive view of the patient has a role in achieving the best interpretation of the patient's imaging, and by extension, the most optimal management of the patient [58].

Furthermore, EI can result in reduction of repeat examinations, which aids in better use of resource and reduction in radiation exposure to a patient. Consider a situation where a patient is on a long waiting list for a CT examination. Their referring clinician puts them on a waiting list in two different clinical sites. The patient attends the first clinical site for their CT and they are not taken off the waiting list in the second site. Anecdotally, the patient may be unaware of the implications and attend a further CT at that second site. A similar situation may occur when the patient visits more than one A&E for the same complaint, and has multiple, duplicate imaging examinations because the results are not available via EI sharing.

3.2.1. Enhancing patient and caregiver experience

One of the biggest challenges patients often face is ensuring that the various clinicians that comprise their care team are able to review imaging that may have been done elsewhere. EHRs have gradually become better equipped to exchange patient data, for example, by encoding it using the HL-7 clinical document architecture (CDA) standard [59]. However, despite the ubiquity of the DICOM standard [60], image exchange still relies on physical media, often transferred by the patients themselves, from one care facility to another.

In a rapidly evolving technological world where social media and internet connectivity drive customer decisions, having instant access to a patient's own imaging record is a progressive, but increasingly anticipated solution. There are multiple efforts underway to decrease health care systems' reliance on physical media for image exchange, and transition to cloud-based solutions through which patients can control and initiate transfer of their images to clinicians who need them [6]. Instead of requiring patients to request image exchange, retrieve an optical disc with their examinations, deliver it to the next care facility, and hope that the correct examinations are on the disc in the correct format, patients could ideally use a web portal to specify the examinations to be transferred and the receiving facility, and electronically consent for (near) instantaneous image transfer [61]. For this model to succeed, the cost of these tools should be borne by the health systems and clinicians sharing and requesting the images, rather than by patients, who may not always have a choice of which clinician is caring for them.

One of the biggest challenges to patient-centric EI is that most EHRs are not image-enabled; while they may be able to link to a third-party image viewer, they do not inherently support storage and transfer of

imaging examinations. As such, simply transferring the patient's electronic chart from one care facility to another does not ensure successful transfer of their imaging history. Yet the availability of this imaging history can be critical to the delivery of timely, accurate, and effective care. The ability for images to be received electronically and reviewed by a specialist or emergency care provider, prior to the patient arriving, can improve outcomes for emergent care but also decrease the need for repeat imaging, unnecessary radiation exposure, and potential rejection of insurance reimbursement for duplicate imaging. It would also decrease the amount of time and effort on the part of the clinician trying to locate prior images and results.

In theory, the combination of an EHR and a robust EI solution can make access to a patient's record easier for both clinicians and patients. This becomes important not only for care delivery at the site with the data, but for easier release of information for transfer of care to other locations. However, the likelihood that the EHR and the EI solution are sufficiently integrated to allow a single data extract is low. In addition, the receiving care facility will likely have a different configuration of systems that will require customized import for the data to be clinically useful.

3.3. Ubiquitous image access

As discussed in 3.1, with the exception of radiology, most image-producing specialties store the imaging they perform within local systems housed in their departments. This results in a set of silos of imaging data that are difficult to access by clinicians and practitioners in other departments. Most radiology departments have overcome this problem by providing universal access to the examinations in PACS through a lightweight viewer - that will load images but likely not offer the full functionality of PACS. However, in order to access images from other departments, clinicians must specifically be granted access to those systems. This is neither efficient nor practical for good patient care.

Just as the EHR revolutionized access to the patient's record from multiple sites within the healthcare enterprise – and in some cases, from sites outside the enterprise that use the same EHR platform [62] – EI has the potential to do the same for imaging, both DICOM-based and non-DICOM based. Any patient data that could be considered an 'image', whether a photograph of a wound, a scan of an ECG, point-of-care ultrasound, or conventional DICOM-based imaging, should be accessible from the EHR – even if not inherently stored there [32].

Integrated delivery networks (IDNs) were introduced nearly 25 years ago in response to growing financial pressures on healthcare organizations in the United States to deliver higher quality care despite decreasing reimbursements. An IDN forms a new affiliation between multiple previously independent healthcare facilities in an effort to develop a mutually beneficial business model and strategic plan [63]. The individual healthcare facilities can vary in type from freestanding hospitals to academic medical centers to extended care facilities and standalone doctors' offices. At the beginning, IDNs often have multiple image storage and distribution systems that may replicate functionality at different locations, but communicate with one another poorly. A physician in hospital A may not be able to see imaging performed at imaging center B, because they were previously unaffiliated. Similarly, a surgeon performing an operation at academic Center C may not be able to review the preoperative imaging performed at imaging Center D. However, when A, B, C, and D are unified under enterprise-wide consolidation of tools and an EI-driven workflow, clinicians at all sites will be able to access imaging of all types included in the VNA, ideally integrated into the respective EHRs at each location. In an IDN, the EHR vendors will likely differ between sites. However, if EI is properly implemented, imaging will be accessible regardless of the particular vendor or viewer being used.

As has been discussed, there are many advantages to implementation of EI, especially within an IDN. Sirota-Cohen et al. [9] observed a decrease in care delivery cost, image storage cost, and unscheduled

outages, as well as an increase in disaster recovery support, all in the setting of increased imaging volumes and need for retrieval of images from the archive. However, they rightly note that achieving these cost savings and increased performance required an organizational strategic plan, migration of both DICOM and non-DICOM data into a centralized archive, and investment in hardware, software, and maintenance. Despite this additional expenditure, the return on investment of the EI deployment was eventually still projected to be positive. So while a decrease in cost can be expected from unnecessary duplication of resources such as local image storage management systems in different departments in different geographical locations, it is important to recognize that the benefits of an EI implementation may take a few years to fully materialize.

EI can facilitate collaboration between primary care and specialty physicians across an IDN, even if they are geographically separated from one another. Information sharing – in particular, image sharing – can improve communication during consultations, thereby allowing clinicians to deliver better care more efficiently. Encounters-based imaging can be shared as easily as orders-based imaging, making valuable data accessible to clinicians regardless of their location. Early adopters of EI report improvements in clinical decision-making, seamless access to patients' imaging across the IDN, fewer incomplete or incorrect examinations imported from 'external' sites, and fewer imaging examinations repeated. The latter decreases cost, inconvenience, delays in care, and frustration for patients and increases satisfaction on the part of both patients and clinicians [9].

3.4. Big data and clinical workflow challenges

The healthcare data generated by an enterprise clinical setting is diverse and can range from laboratory reports in the form of long-form-text, to medical images from different specialties, or other information generated from medical devices [64]. One of the key challenges in the healthcare industry is how to manage, store and exchange all of these data. In order to develop a machine learning healthcare tool using Big Data, the data must first be carefully labelled and curated, an appropriate machine learning framework selected, and best practices followed.

3.4.1. Human-in-the-loop data curation

Data curation is the process of integrating, sorting and labeling data from different sources into manageable and defined datasets, and requires input from clinicians to ensure that the clinical interpretation of the data is retained, and from data scientists who understand the limitations of the machine learning algorithms. It is essential that the data is labeled accurately and captures reliably the clinical interpretation in order for the machine learning models to work consistently and efficiently. This clinical re-classification needs to be an iterative cycle whereby prospective clinical data is extracted and refined, and the clinical workflow is designed with machine learning methods [65]. Additionally, it is important to reliably record the context of the data in terms of clinical implication, for example in the case of a note in a report about the family history of a patient versus the patient's own clinical history.

3.4.2. Data analytics and validation

In order to adhere to machine learning best practices, the structure, format, and semantics of similar data should be consistent across clinical systems [66]. Once data has been labelled, several data pre-processing methods need to be performed in order to check for errors such as missing data or duplicate entries. For example, healthcare data typically includes irregular data points, and they may need to be adjusted to fit the predetermined time points. This can be achieved by systematically imputing data at a number of days before or after the nearest predetermined time point to that time point [65].

Prospective clinical trials need to be performed to assess the effect of

machine learning tools in real world clinical settings with integrated machine learning workflows and clinicians in the loop to review the clinical trial data and refine the model [65].

3.4.3. Deep learning frameworks for clinical applications

Various data management and analytic frameworks will be more or less appropriate, depending on the nature of the data and analysis involved. For example, a large image dataset would typically require far more storage space than an equivalent text dataset, and data that pairs text and images, or even graph-based data that mixes many data types may require bespoke data management solutions to enable analysis. To train and validate a model, it must be possible to sample representative training and test, although to form a robust estimation of a model's effectiveness it may be required to select cross-validation cohorts. This process can be complicated in domains where there are relatively few samples (for example, in rare diseases), requiring stratified sampling and careful handling of data. Once this pipeline is established, the iterative process of training and refining models begins, including experimentation with various model architectures, searching for optimal model hyper-parameters, and task-specific adjustments [67].

However, a machine learning model must be demonstrated to be safe and effective for use in the clinic to be effectively employed. This is primarily achieved through implementation of a Quality Management System (QMS), one of core requirements for FDA approval (US) and CE Mark (EU). The FDA issued guidance on software as a medical device in 2017, explaining risk stratification and the analytical and clinical validation required of AI tools in healthcare. The process is broadly similar to current software engineering and machine learning best practices, but precludes certain methods common in other industries such as incremental learning, in which outcome data from a trained AI system are incorporated into a closed feedback loop and used to refine the accuracy of the systems' predictions through iterative retraining of the model require a design freeze for medical device registration. Current FDA approval of medical devices is governed by 21 CFR [68], although is currently in the process of transitioning to ISO 13485:2016 [69]. This move is a step towards global harmonization of medical device regulation as ISO 13485:2016 is already the basis for certification and QMS requirements in the EU and other countries such as Canada, Brazil, Japan and Korea.

3.4.4. Challenges to clinical deployment

The current poor inter-operability of EHRs creates challenges for performing big data analytics in healthcare [65]. It requires platforms that enable the building of multiple machine learning tools and methodologies, and linking of a patient's electronic health record and associated data governance systems are required [65]. The practices of data ingestion, modeling and visualization that are standard in other industries could then be integrated with existing tools to support a healthcare enterprise solution.

Several clinical AI tools that provide a recommendation to the clinician, who ultimately make the treatment decision have now received FDA approval [65,70]. However, there are a lack of prospective clinical trials to fully evaluate these systems. Similar to the training of junior doctors, a clinical machine learning tool is best trained by incorporating real-world medical data into the disease model, then tuned by medical experts to improve its accuracy in predicting real cases [66].

Clinical deployment of AI tools also raises several ethical concerns. These concerns include liability in cases of medical error and issues of privacy, security, and control of the patient. The emergence of explainable AI (XAI) and the more recent explainable interactive machine learning (XIL) frameworks [71] will help to bridge the gap between clinician understanding of how machine learning tools produce predictions and can potentially empower patients' understanding and control of their own health care.

3.5. Automated reporting

Automated multimedia reporting is one potential benefit of EI that has not been fully realized. Multimedia-enhanced radiology reporting (MERR) has been proposed as an alternative to text-only radiology reports, and was perceived as an appealing option by specialist physicians [72]. A recent survey of practicing radiologists found that they also felt that interactive multimedia radiology reports would be beneficial to their practice [73]. At present, most available MERR solutions require using the PACS vendor's reporting environment, rather than the radiologists' typical reporting solution. As such, this requires radiologists to be willing to make significant changes to their workflow, which may be not be realistic. However, for those practices looking for a new reporting solution or ready for a change to their reporting workflow, these vendor offerings provide seamless integration of MERR into that adjusted workflow [74,75], and enable those practices to reap the benefits of easily creating multimedia reports, but also pass the benefit of improved communication along to their referring clinicians and patients.

4. Predictive applications

In this section we will discuss AI systems in the context of their use as components of an enterprise imaging strategy. We consider this under the lens of three broad categories, focusing firstly on the current generation of operational systems, secondly on the potential for such AI systems to contribute to the monitoring and analysis of population health, and then finally looking forward to promising applications and use-cases in a number of domains.

One of the most exciting trends of recent years is the application of AI to medical imaging. This has been reflected in the dramatic rise in the number of publications on this topic [76]. Although much of this work has yet to translate into clinical practice, it is clear that machine learning and AI systems will become a core component of any enterprise imaging system, and will impact many areas of medical imaging. Although they offer substantial benefits to augment, accelerate and complement clinicians' existing practice and workflow, while enabling a higher standard of care for patients, a key issue to address is how to manage their introduction into clinical use.

Most radiologists will be familiar with computer-aided diagnosis or detection (CADx) systems. Whereas these systems operated by detecting and extracting hand-crafted features that are input to a machine learning model, the rise of deep learning models which automatically learn the relevant features has in some ways simplified the model development pipeline; rather than spending significant time in creating hand-crafted features for a specific clinical scenario, common building blocks of neural networks can instead learn them during the model training process – if given a sufficiently large and varied dataset. This change, in combination with increased access to computational resources and integration of data sources facilitates application and integration of AI into clinical decision making of increased complexity.

Integrated AI systems will be capable of handling many of the routine tasks such as detection, quantification, etc, that are currently performed by radiologists and pathologists. This will allow them to shift focus from onerous pattern recognition tasks to holistic interpretation of the available data and patient care. It seems likely that this change will also require medical imaging specialists to incorporate and become proficient in the tools of data science and statistics [77].

These systems will also remove much of the perceptual subjectivity from a radiologist or pathologist's work, mitigating against inter and intra-interpreter variations and against the debilitating effect of visual fatigue on image interpretation. It will create an objective and reproducible yard-stick for image interpretation in many medical imaging domains, which will improve overall diagnostic accuracy and reduce error.

4.1. Operational systems

It is worth considering how the renewed interest in AI due to the successes of deep learning-based computer vision is different from the previous CAD approaches that did not gain as much traction in the clinical workflow. Although CAD was hailed as a promising adjunct to the radiologist, it failed to materialize actual improvements in radiologists' accuracy or efficiency, in part because of the excessive numbers of false positives that radiologists were suddenly required to review and dismiss [78]. One of the biggest differences between deep learning and CAD is the demonstrated success of the former in countless non-medical applications. While work remains to realize the same degree of success while balancing the hype of what deep learning could do for healthcare, the technology clearly has greater true promise than CAD did [79].

A number of use cases in pixel-based AI have demonstrated early successes. In this section, we discuss some of these systems and how they can be integrated into a pre-existing enterprise imaging workflow. Some of the earliest developments in radiology have come in the field of neuroimaging, where de-identified datasets can easily be shared (after skull stripping), and labeled data has existed for some time as a result of multiple image analysis challenges for segmentation and registration. However, ophthalmologic imaging preceded radiology in both impressive performance as well as Centers for Medicare and Medicaid Services (CMS) reimbursement in the United States.

For example, studies using data from the Alzheimer's Disease Neuroimaging Initiative (ADNI) have been able to develop deep learning models to distinguish Alzheimer's disease from mild cognitive impairment or other pathologies on ^{18}F -FDG PET [80]. Ding et al. designed a convolutional neural network (CNN) that was able to suggest the diagnosis of Alzheimer's an average of six years before the diagnosis was ultimately made [80]. Qiu et al. developed a deep learning model that included a CNN to detect Alzheimer's features on brain MRIs from the ADNI dataset, and provided that as an input to a downstream model that included patient features such as age, gender and mini mental status exam (MMSE) score [81].

In the US only a handful of pixel-based AI solutions have been approved for reimbursement by CMS via their new technology add-on payment (NTAP) mechanism. These include:

1. Detection of acute large vessel occlusion strokes on CT angiography [82–84]. AI for stroke imaging has been developed not only for pixel-based analysis but for triage of the radiologist's worklist, which can integrate into an enterprise imaging solution whereby multiple specialist physicians are reviewing imaging (e.g., radiology, neurology, neurointerventional) to determine if expedient treatment of the patient via thrombectomy is necessary and feasible.
2. Detection of diabetic retinopathy or precursor findings [85,86]. Multiple studies have shown that deep learning models for detection of diabetic retinopathy on fundoscopic imaging can effectively screen patients for further evaluation by ophthalmologists, even in settings outside the healthcare enterprise (e.g., at a chemist) or outside of a specialist care setting (e.g., at the general practitioner's office).

A recent analysis of approved medical devices using A.I. and machine learning in the EU and US identified 240 and 222 such devices in each jurisdiction respectively [70], although only few were classified as high-risk. The study does also not give an indication as to the level of complexity of the AI models, which could be low. Nevertheless, since 2015 there has been a near 10-fold increase in the number of devices approved year on year, with the majority being approved in radiology (65%), and cardiovascular (15%). This does not indicate that all such devices are in active clinical use, however. The EUDAMED database will help in this regard as – once operational – will provide a publicly accessible list of the certified medical devices that are available in the EU [87].

4.2. Improving population health

In the U.S., a national incentive program for meaningful use of healthcare information technology was intended to create large quantities of electronic health data that could then be used to create benchmarks for population health analytics [88,89]. Instead, it resulted in the creation of numerous silos of electronic health data within IDNs and standalone healthcare facilities, that have to be exported from one EHR and imported into another for data exchange.

While there are some health information exchanges (HIEs) and regional health information organizations (RHIOs) across the U.S., participation in these efforts at decreasing redundant utilization of healthcare resources varies significantly [90–92]. Furthermore, these attempts at data sharing do not always incorporate imaging data, and may only include imaging reports.

In the EU, the implementation of GDPR and its strict requirement for patient consent, personal data safeguards and provenance – all generally positive aspects – has put up a number of impediments to the sharing and collecting of medical data that could lead to true ‘big data’ [93].

Because of the limitations of organized data collection approaches, there are ongoing efforts to crowd-source clinical data directly from patients, and collect not only medical histories, but actual images and even genomic analyses [94,95].

However, in light of the increased need for training and validation data with which to develop robust deep learning models, health care facilities could be incentivized to share data to support these efforts. A standardized format for data sharing, such as using HL7 CDA or FHIR, could potentially decrease the amount of data wrangling required to make the data usable for large-scale analytics, and decrease the potential for data corruption and loss during and after transfer. Furthermore, it would enable construction of models using large quantities of weakly labeled data, which have been shown to result in models of similar performance to those trained with smaller quantities of expert-labeled data [96].

4.3. Future tools

In the past few decades medical imaging has produced a huge body of research on the development of computer-aided detection (CAD) systems addressing the core problems of classification and segmentation. In addition to these standard tasks, a number of other techniques from deep learning have recently found significant applications within medical imaging. These include; super-resolution, the use of neural networks to enhance the quality of an image; the use of multi-modal models to jointly learn multiple image modalities (e.g. PET-MRI, CT, MRI), or to combine textual, time-series or even audio data within a single neural network; model interpretability and explainability, removing the ‘black-box’ nature of neural networks and gaining a better understanding of why a certain prediction is made; few-shot learning, which enables neural networks to learn from mere tens of images rather than the hundreds or thousands typically required. Due to the increased prevalence of EI, an ever-increasing volume of data, and a general shift towards connectivity and accessibility of said data, it is worth considering *how* we may take advantage of the ‘big data’ collected by EI.

In the remainder of this section we will highlight a few examples of recent work in these areas, with some discussion on their applications and the potential to enhance EI.

4.3.1. Traditional machine learning vs deep learning

Traditional medical imaging classification tasks were performed through manual definition and extraction of morphometric, statistical or digital signal features, for example describing the texture of regions of interest using features based on the gray level co-occurrence matrix (GLCM) and histogram of oriented gradients (HOG). To avoid overfitting, the feature space is then reduced using one of the various feature selection or dimensionality reduction algorithms, such as

wrapper or filter-based feature selection, principle component analysis (PCA), and linear discriminant analysis (LDA). The reduced feature space is then fed to a classifier model to assign a label to the input image (or just a region of it). Classifiers typically used include logistic regression, support vector machines (SVM) and random forests (RF).

The traditional machine learning approach for classification purposes in the medical imaging field is very time-consuming and requires huge domain expertise – not only for annotation of images but in developing hand-crafted features that may only be optimal in relatively niche domains. Instead, deep learning algorithms have been taking over due to them being structured around extracting features automatically from the input image. A variety of deep learning neural networks have been developed that have found applications in medical imaging, from simple fully connected and convolutional neural networks, to more sophisticated residual, recurrent and siamese networks [97], and even generative networks such as autoencoders and generative adversarial networks (GANs) [98].

4.3.2. Segmentation

Segmentation is the process of extracting a specific region of interest from an image, and is one of the core problems in medical imaging. Image segmentation has huge applications in different fields, such as: detecting tumours for tracking their volume in response to therapy (or for example to help determine which type of tumour the patient has, since different brain tumours have different shapes [99]), segmenting blood cells (to then classify them for detecting any increase/decrease in a specific type of blood cells that could be an indicator of some disease), detecting calcification within the breast tissue and much more [100]. In recent years, deep learning has been able to achieve human-level performance in many image recognition tasks. Deep learning image segmentation tasks can be categorised into unsupervised and supervised methods [101], as discussed below. A number of metrics are commonly used to quantify the performance of segmentation models such as the Dice score, positive predictive value (PPV), true positive rate (TPR) and absolute volume difference (AVD) [102].

b) Supervised Segmentation

Supervised segmentation relies on providing a model with a paired raw biomedical image and set of annotations during the training process. Once trained, the model is capable of performing segmentation tasks on new images without annotations. Annotating images is a highly time consuming task especially in medical fields and can only be done by an expert and with limited resources. However, over the past few years supervised learning for biomedical images segmentation has managed to achieve a human-level performance that is very promising [101]. In particular, the U-Net architecture and training strategy [103] which was originally proposed to deal with the lower number of samples commonly found in biomedical domains has been successfully applied to many segmentation problems, including abnormal tissue segmentation [104], organ segmentation in CT images [105], and tumour segmentation in brain MRI [106].

a) Unsupervised Segmentation

In contrast to supervised learning, unsupervised learning does not require paired annotations in order to train a segmentation model. However, unsupervised segmentation is generally a much more challenging task, and usually does not perform as well as supervised segmentation [101]. Over the years, different research groups have developed innovative algorithms to automatically segment a specific tissue of interest without the need of annotated images. Atlason et al. [102] developed a convolutional segmentation auto-encoder (SegAE) capable of segmenting brain lesions from T1, T2 and FLAIR images through reconstruction the FLAIR image as a canonical combination of the segmented white matter, gray matter, cerebrospinal fluid and lesions. This algorithm has the unique advantage of not requiring healthy, lesion-free, training data, which is very hard to acquire from elderly subjects. On the other hand, Dalca et al. [107] used a set of un-annotated images to build a probabilistic anatomical prior that was subsequently

used for CNN-based anatomical segmentation of the brain. Another approach for segmenting brain tumours that Bhanumurthy et al. [100] followed is by first extracting features from the brain MRI (namely, Energy, Entropy, Homogeneity, Contrast and Correlation), which are used as inputs to a neuro-fuzzy classifier in order to assign the images into normal and abnormal classes. Tumours in the abnormal MRIs are then segmented using the Region Growth Method (RGM).

4.3.3. Transfer learning

Transfer learning is a machine learning technique in which a model trained on one task is subsequently adapted to a second, typically more niche, task. This approach has been used to great effect in many applications of computer vision and natural language processing, where the use of transfer learning to adapt to a new dataset instead of training ‘from scratch’ side-steps the requirements for significant computational resources.

Open-source natural image datasets have been used to pre-train complex neural networks that can be used off-the-shelf when building neural networks for specific applications in medical imaging. ImageNet is one of the most popular open-source datasets, comprising of roughly 1 M images assigned to 1,000 classes. Several neural networks have been trained on ImageNet, including the now famous VGG, ResNet, and Inception models, achieving high classification accuracy. Transfer learning has been established as a common practice for building medical imaging classifiers, achieving very promising results - some of which have been FDA approved. However, Raghu et al. [108] have shown that using transfer learning for medical images classification is actually not significantly better than training a much smaller, lightweight, network. This is due to the fact that many medical imaging domains have very few classes – typically less than a half a dozen – and the use of neural architectures capable of representing thousands of classes leaves a large amount of redundant model parameters.

4.3.4. Multi-modality

Multi-modality in a medical imaging context usually consists of the use of different imaging modalities (MRI, CT, PET, etc) in order to aid clinical decision-making. Within the framework of an EI solution, it need not be limited to different imaging modalities, and increasingly can incorporate EMR information such as paired diagnostic reports, patient attributes, or time-series data. Recent advances in deep learning have produced models that are capable of being trained “end-to-end” using paired imaging and text data, making it possible to learn statistical models from relatively unstructured text and un-annotated images.

MRI includes various imaging sequences, including: T1 (spin-lattice relaxation), T1C (T1-contrasted), T2 (spin-spin relaxation) and FLAIR (fluid attenuation inversion recovery). Different tissues respond differently to each of these modalities and thus appear differently in each of these sequences. When building CNN-based segmentation models, ideally, all sequences are used to train the model. Tseng et al. [99] developed the first of its kind cross-modality model to segment 3D biomedical images that exploits a sequence learning method, integrating information from consecutive slices and modalities.

Zhang et al. [109] developed MDNET, a multi-modal neural network that receives paired images and diagnostic reports as input. In this fashion it learns correspondences between words in the report and image features, and is capable of automatically generating reports for novel input images. By incorporating attention mechanisms [110] into the model, visualization of the model decision can be performed, for example indicating which image features led to the classification decision, or sentence in the diagnostic report. This approach provides some sort of justification that the clinicians can use to interpret the model output.

4.3.5. Few-shot learning

Few-shot learning is the task of successfully training a model from a very limited number of samples per class. This is a direct attempt to

address data scarcity, and mitigate the data hungry nature of deep learning (requiring thousands of samples), and achieve more ‘human-like’ learning of categories from just a few examples. Few-shot learning algorithms have obvious applications in medical imaging, where datasets are frequently small, and in cases of rare diseases may be tiny. As a corollary, this can also be used in situations where there is plenty of image data, but obtaining expert annotations is costly.

Ouyang et al. [111] developed a few-shot segmentation framework that uses superpixels as a weak labeling scheme to self-supervise the model training process. This has the advantage of not requiring manual annotation of image regions. In experiments on abdominal organ segmentation for CT and MRI, and cardiac segmentation for MRI, they out-perform standard few-shot segmentation methods, while not requiring annotations.

Other attempts to reduce the number of annotations required to train a model include those by Zhou et al. [112], who developed an approach combining a small amount of pixel-level annotations, a visual attention mechanism and a multi-lesion mask generator network in order to train a lesion segmentation in a semi-supervised fashion. Zheng et al. [101] describe an algorithm called representative annotation (RA), in which the network selects representative image patches to be annotated. The selection process is done through unsupervised feature extraction from non-annotated data, and has been demonstrated to perform at a similar level to the suggested annotation (SA) method, which is the current state-of-the-art.

4.3.6. 2D vs 3D CNNs

Medical imaging is increasingly performed in 3D in order to fully visualise the organs, determine their structure and measure volume. CNN-based segmentation has a huge role in this. CNNs traditionally operate in 2 spatial dimensions, but are easily adapted to 3D inputs. However, due to practical limitations in available computing power and in particular GPU memory, there is a trade-off between the overall network size (i.e. number of feature maps), convolution filter receptive field, and therefore utilization of inter-slice information. Due to this, 2D CNN models can usually have much larger receptive fields, while 3D CNNs are smaller [103,113]).

Two approaches have been developed to work around the limited field of view of 3D CNNs. The first approach relies on using only three 2D slices from orthogonal planes (i.e the xy, xz and yz planes) instead of the entire 3D image. The second approach relies on using 2D CNNs to summarize intra-slice information and use that as additional input to the 3D CNN.

Recently, Zheng et al. [114] developed an ensemble learning framework for segmenting 3D images, using one 3D CNN and three 2D CNNs (for each of the three orthogonal 2D slices, namely xy, xz and yz). These four CNNs are used jointly as base-learners for training the meta-learner. Through combining the advantages of the 2D and 3D CNNs, this ensemble learning approach has been shown to achieve superior performance over the state-of-the-art methods.

4.3.7. Model interpretability

Neural networks have been successful in learning from vast amounts of data and producing useful results. However, they are “black box” models, in that they output results without providing a clear explanation for the decision outcome. This is in contrast to traditional classifier models that may extract specific, user-defined features to perform the classification task, or be inherently interpretable as in the case of decision trees. An area causing major interest as a method for achieving higher quality in medical imaging is in the area of explainable machine learning, which is a possible direction in which new regulations could be enforced [115]. Understanding the decision-making process of a deep learning model is essential for clinical adoption [116–118].

The use of attention mechanisms and saliency masks have gained some traction in this area, as they provide a way to visualize what region of an image was attended to that led to the predicted outcome. They

were employed recently in models used to screen chest X-rays of COVID-19 patients [119], predict lung module malignancy from longitudinal CT [120], perform abnormal tissue segmentation in natural, CT and MRI images [104] and quantification of knee osteoarthritis in X-ray images [121], while the MDNet model of Zhang et al. [109] used attention mechanisms to indicate which area of the image corresponded to the text in the generated diagnostic report.

4.3.8. Super-resolution and generative models

Super resolution (SR) is a general term used to denote methods for enhancing the quality of videos or images. Traditionally this was performed using image operations such as nearest-neighbour, bilinear and quadratic interpolation, but recent developments of deep learning super-resolution models have produced impressive results. The deep learning approach requires training of a super-resolution model on a target domain (for example faces and natural images), which then learns a ‘smart’ contextual method of up-scaling an input image. This is achieved by training the network with a number of loss functions emphasizing different image qualities, including: pixel loss, a direct comparison of pixels in the ground-truth and up-sampled image; content loss, a comparison based on perceptual quality; texture loss, a comparison of texture in each image, and; total variation loss, to suppress noise in generated images.

The family of deep learning generative models typically includes variational autoencoders (VAE) [122] and generative adversarial networks (GANs) [123]. Autoencoders are bipartite networks comprised of an encoder module that compresses the input data to a lower dimensional latent space, and a decoder module that tries to recreate the original input from the compressed latent vector representing the input data. Generative adversarial networks consist of a generator model that tries to generate realistic samples of the training data, and a discriminator that tries to distinguish the real and generated samples, creating an ‘arms race’ in which the generator learns to produce ever more realistic samples. This adversarial training framework can use many different network architectures, such as CNNs, RNNs, etc, and can therefore be used to generate everything from images, to text, to molecular graphs [124].

Although much of the focus in this area has been for artistic, rather than medical applications, a number of key opportunities for generative models have been identified, such as conditional image synthesis, super-resolution, segmentation and stain normalization [125]. Conditional image synthesis, in which a GAN is trained to generate one kind of image data from another, for example converting between MRI and CT [126], and generating PET from CT [127]. In pathology this approach has also been used to normalize stain appearance [128,129]. Work has also been performed in applying super resolution to various medical domains, including work by Chaudhari et al. [130], who applied super-resolution to double the slice resolution of 5-min quantitative double-echo steady-state (qDESS) sequences of knee MRI images, achieving comparable results with conventional MRI and arthroscopy. The same authors previously reported work investigating the use of SR to help improve MRI acquisition time by up-sampling lower resolution MRI images while maintaining signal-to-noise ratio that was acceptable for acquisition of imaging biomarkers [131]. Critically, it has been demonstrated that there can be significant interchangeability between conventional and up-sampled or reconstructed MRI images [132].

Generative models can also be used to reduce the amount of contrast agent required for some diagnostic imaging modalities. Chen et al. [133] trained a CNNs on a dataset of simultaneously acquired MRI and ultra-low-dose PET, and were able to generate full-dose-like amyloid PET images, with a variation of only 89% vs 91% accuracy in intra-reader reproducibility. Gong et al. [134] previously used similar methods to reduce the gadolinium dose required for contrast-enhanced brain MRI 10-fold. Although nascent, the use of deep learning models to expedite MRI acquisition, enhance image quality, and lower the dosage requirements of radiotracers – without degrading the quality of clinical

decision making – is very promising.

5. Conclusions

A strategy for the implementation and integration of enterprise imaging is a critical next step in the digital transformation of healthcare. Better management of digital imaging and associated reports can add real clinical value through the efficient storage and access to electronic health records. Costs can be reduced overall via infrastructure and support team consolidation. Financial, regulatory, and clinical risks can be minimized through increased training and awareness of cybersecurity principles, and implementation of data security and privacy controls. Advances in AI and machine learning will further enhance EI capabilities, ranging from real-time acquisition alerts and priority case management, to optimisation of patient-specific imaging protocols and automated reporting. The use of deep learning to expedite acquisition techniques, for example with super-resolution, will enable lower resolution images to be acquired while maintaining an acceptable signal-to-noise ratio. This will result in patients spending less time in the scanner, achieve a higher throughput of patients, and minimising waiting lists and reducing healthcare costs. The introduction of AI tools into clinical workflows will help to ensure that patient care is optimized when time-critical reporting is required, and an EI framework streamlines this communication beyond radiology and not restricted to a single site. AI techniques such as multi-view modelling, with appropriate access and regulation, will be able to use data from multiple sources and geographical locations to create highly accurate predictive tools for diagnosis.

The availability and variety of high quality annotated data is a key hurdle for implementation of AI in a clinical setting, as data source and potential biases may affect generalization of AI models to wider populations. Interactive and synoptic reporting may ultimately improve data availability for AI applications in clinical decision circumstances, thereby improving the next generation of healthcare applications. Legislative changes such as the increased harmonisation of medical device regulations, and technological approaches such as federated learning will help to make global healthcare solutions possible by making it easier to operate across organisational and national boundaries.

Acknowledgements

The authors wish to thank Nicky Dunne for proof-reading and Enterprise Ireland for funding grant CF-2019-1248-I.

References

- [1] Primo H, Bishop M, Lannum L, Cram D, Nader A, Boodoo R. 10 steps to strategically build and implement your enterprise imaging system: HIMSS-SIIM collaborative white paper. *J Digit Imag* Aug 2019;32:535–43.
- [2] Obermeyer Z, Emanuel EJ. Predicting the future – big data, machine learning, and clinical medicine. *N Engl J Med* 2016;375(13):1216–9.
- [3] Roth CJ, Lannum LM, Persons KR. A foundation for enterprise imaging: HIMSS-SIIM collaborative white paper. *J Digit Imag* 2016;29(5):530–8.
- [4] Cram D, Roth CJ, Towbin AJ. Orders- versus encounters-based image capture: Implications pre- and post-procedure workflow, technical and build capabilities, resulting, analytics and revenue capture: HIMSS-SIIM collaborative white paper. *J Digit Imag* Oct 2016;29:559–66.
- [5] Roth CJ, Lannum LM, Joseph CL. Enterprise imaging governance: HIMSS-SIIM collaborative white paper. *J Digit Imag* 2016;29(5):539–46.
- [6] Vreeland A, Persons KR, Primo HR, Bishop M, Garriott KM, Doyle MK, Silver E, Brown DM, Bashall C. Considerations for exchanging and sharing medical images for improved collaboration and patient care: HIMSS-SIIM collaborative white paper. *J Digit Imag* Oct 2016;29:547–58.
- [7] Nagels J, MacDonald D, Parker D. Foreign exam management in practice: seamless access to foreign images and results in a regional environment. *J Digit Imag* 2015;28(2):188–93.
- [8] Petersilge CA. The enterprise imaging value proposition. *J Digit Imag* 2020;33(1):37–48.
- [9] Sirota-Cohen C, Rosipko B, Forsberg D, Sunshine JL. Implementation and benefits of a vendor-neutral archive and enterprise-imaging management system in an integrated delivery network. *J Digit Imag* 2019;32(2):211–20.

- [10] Use the three rings of information governance for classifying healthcare data. url: <https://www.gartner.com/en/documents/3629832>. Accessed: 2020-11-21.
- [11] Hripacsak G, Bloomrosen M, FlatelyBrennan P, Chute CG, Cimino J, Detmer DE, Edmunds M, Embi PJ, Goldstein MM, Hammond WE, Keenan GM, Labkoff S, Murphy S, Safran C, Speedie S, Strasberg H, Temple F, Wilcox AB. Health data use, stewardship, and governance: Ongoing gaps and challenges: A report from AMIA's 2012 health policy meeting. *J Am Med Inform Assoc* 2014;21(2):204–11. *Journal of the American Medical Informatics Association*.
- [12] "2018 reform of eu data protection rules."
- [13] Centers for Medicare & Medicaid Services, The Health Insurance Portability and Accountability Act of 1996 (HIPAA). Online at <http://www.cms.hhs.gov/hipaa/>, 1996.
- [14] Tovino SA. The hipaa privacy rule and the eu gdpr: illustrative comparisons. *Seton Hall L Rev* 2016;47:973.
- [15] Abdollahzadehgan A, Gohary MM, Razak A, Hussin C, Amin M. Assessing the determinants of cloud computing services for utilizing health information systems: a case study. *J Inf Syst Res Innov (JISRI)* 2013;4(1):67–74.
- [16] Market H. Healthcare cloud computing market – global forecast to 2025. url: <https://www.marketsandmarkets.com/Market-Reports/cloud-computing-healthcare-market-347.html>. Accessed: 2020-11-21.
- [17] Mell P, Grance T. The NIST definition of cloud computing recommendations of the national institute of standards and technology, tech. rep.
- [18] The world of cloud-based services: storing health data in the cloud.
- [19] Burde JD H. Virtual mentor: health law the HITECH Act – An Overview, *AMA. J Ethics* 2011;13(3):172–5.
- [20] Henry J, Pylpuchuk Y, Searcy T, Patel V. Adoption of electronic health record systems among U.S. non-federal acute care hospitals: 2008–2015. url: <https://dashboard.healthit.gov/evaluations/data-briefs/non-federal-acute-care-hospital-ehr-adoption-2008-2015.php>. Accessed: 2020-11-21.
- [21] Dicom supplement overview. url: <https://www.dicomstandard.org/supplements>. Accessed: 2020-11-21.
- [22] Zarella MD, Bowman D, Aeffner F, Farahani N, Xthona A, Absar SF, Parwani A, et al. A practical guide to whole slide imaging: a white paper from the digital pathology association. *Arch Pathol Lab Medicine* 143(10):2018:222–234.
- [23] Up to \$12b in unnecessary medical imaging is wasted annually. url: <https://hitconsultant.net/2014/09/03/12b-in-unnecessary-medical-imaging-is-wasted-annually/>. Accessed: 2020-11-21.
- [24] Bodenheimer T, Sinsky C. From triple to quadruple aim: care of the patient requires care of the provider. *Ann Family Med* 2014;12(6):573–6.
- [25] Paton RA, McCalman J. Change management: A guide to effective implementation. Sage 2008.
- [26] By RT, Burnes B, Oswick C. Change management: Leadership, values and ethics. *J Change Manag* 2012;12(1):1–5.
- [27] How to improve healthcare operational efficiency through lean principles and predictive analytics. url: <https://www.healthitoutcomes.com/doc/how-to-improve-healthcare-operational-efficiency-through-lean-principles-and-predictive-analytics-0001>. Accessed: 2020-11-21.
- [28] Swamped with CDs. url: <https://www.radiologytoday.net/archive/rt0211p12.shtml>. Accessed: 2020-11-21.
- [29] "Discover the true cost of CDs." url: <https://ambrahealth.com/hospitals-health-systems/discover-the-true-cost-of-cds/>. Accessed: 2020-11-21.
- [30] How can radiologists effectively reduce costs?. url: <https://www.gehealthcare.com/feature-article/how-can-radiologists-effectively-reduce-costs>. Accessed: 2020-11-21.
- [31] 5 reasons why health systems should implement enterprise imaging. url: <https://www.auntminnie.com/index.aspx?sec=road&sub=pac-2020&pag=dis&itemID=130849>. Accessed: 2020-11-21.
- [32] Petersilge CA. The evolution of enterprise imaging and the role of the radiologist in the new world". *Am J Roentgenol* 2017;209(4):845–8.
- [33] Brady AP. Measuring consultant radiologist workload: method and results from a national survey. *Insights Imag* 2011;2(3):247–60.
- [34] Demand for radiology still high, average salaries up 4 url: <https://www.radiologybusiness.com/topics/healthcare-economics/demand-radiology-still-high-average-salary-4>. Accessed: 2020-11-21.
- [35] Realize true transformation with an enterprise imaging strategy. url: <https://www.changehealthcare.com/insights/enterprise-imaging-strategy-white-paper>. Accessed: 2020-11-21.
- [36] Using business analytics in radiology operations. url: <https://www.itnonline.com/article/using-business-analytics-radiology-operations>. Accessed: 2020-11-21.
- [37] Shrestha RB, et al. Analytics and value-based imaging. *Appl Radiol* 2014;43(5):34.
- [38] Rehani MM. Patient radiation exposure and dose tracking: a perspective. *J Medical Imag* 2017;4(3): 031206.
- [39] Jones S, Courneane S, Sheehy N, Hederman L. A business analytics software tool for monitoring and predicting radiology throughput performance. *J Digit Imag* 2016;29(6):645–53.
- [40] Lew C. Radiology analytics: A clear path to improved performance. *Appl Radiol* 2018;47(11):22–4.
- [41] B. Filkins. Health care cyberthreat report: Widespread compromises detected, compliance nightmare on horizon; 2014.
- [42] NEMA/MITA. Cybersecurity for Medical Imaging, 2016, p. 7.
- [43] Williams PA, Woodward AJ. Cybersecurity vulnerabilities in medical devices: a complex environment and multifaceted problem. *Medical Devices (Auckland, NZ)* 2015;8:305.
- [44] Argaw ST, Bempong N-E, Eshaya-Chauvin B, Flahault A. The state of research on cyberattacks against hospitals and available best practice recommendations: a scoping review. *BMC Medical Inf Decision Making* 2019;19(1):1–11.
- [45] Brismis TS, Chen R, Mela T, Olshkevsky A, Paschalidis IC, Shi W. Federated learning of predictive models from federated Electronic Health Records. *Int J Med Inf* 112(September):2017:59–67, 2018.
- [46] Konečný J, McMahan HB, Ramage D, Richtárik P. Federated optimization: Distributed machine learning for on-device intelligence, arXiv preprint arXiv: 1610.02527; 2016.
- [47] Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, Bakas S, Galtier MN, Landman BA, Maier-Hein K, Ourselin S, Sheller M, Summers RM, Trask A, Xu D, Baust M, Cardoso MJ. The future of digital health with federated learning. *npj Digi Medicine* 2020;3(1):1–7.
- [48] Enthoven D, Al-Ars Z. An overview of federated deep learning privacy attacks and defensive strategies, arXiv; 2020.
- [49] Rajendran K, Manoj Jayabalan Muhammad Ehsan Rana. A Study on k-anonymity, l-diversity, and t-closeness Techniques focusing Medical Data. *IJCSNS Int J Comput Sci Netw Secur* 17(12);2017.
- [50] Dwork C, Roth A. The algorithmic foundations of differential privacy". *Found Trend Theor Comput Sci* 2013;9(3–4):211–487.
- [51] Choudhury O, Gkoulalas-Divanis A, Salonidis T, Sylla I, Park Y, Hsu G, Das A. Differential Privacy-enabled federated learning for sensitive health data. *NeurIPS* 2019:1–6.
- [52] Geyer RC, Klein T, Nabi M. Differentially private federated learning: a client level perspective Nips 2017:1–7.
- [53] Yuan D, Zhu X, Wei M, Ma J. Collaborative deep learning for medical image analysis with differential privacy. In: 2019 IEEE global communications conference (GLOBECOM), IEEE, 2019, p. 1–6.
- [54] Sarwate AD, Plis SM, Turner JA, Arbabshirani MR, Calhoun VD. Sharing privacy-sensitive access to neuroimaging and genetics data: A review and preliminary validation. *Front Neuroinf* 2014;8:1–12.
- [55] Lu MY, Kong D, Lipkova J, Chen RJ, Singh R, Williamson DFK, Chen TY, Mahmood F. Federated learning for computational pathology on gigapixel whole slide images 2020.
- [56] Schwarz CG, Kremers WK, Therneau TM, Sharp RR, Gunter JL, Vemuri P, Arani A, et al. Identification of anonymous MRI research participants with face-recognition software. *New England J Medicine* 381(17):2019:1684–6.
- [57] Payne T, Fellner J, Dugowson C, Liebovitz D, Fletcher G. Use of more than one electronic medical record system within a single health care organization. *Appl Clin Inf* 2012;3(4):462.
- [58] Castillo C, Steffens T, Sim L, Caffery L. The effect of clinical information on radiology reporting: A systematic review. *J Medical Radiat Sci* 2020.
- [59] Dolin RH, Alschuler L, Beebe C, Biron PV, Boyer SL, Essin D, Kimber E, Lincoln T, et al., The hl7 clinical document architecture. *J Am Medical Inf Assoc: JAMIA* 8 (6):2001;552–569. 11687563[pmid].
- [60] Bidgood WD Jr., Horii SC, Prior FW, Van Syckle DE. Understanding and using dicom, the data interchange standard for biomedical imaging. *J Am Medical Inf Asso JAMIA* 4(3):1997:199–212. 9147339[pmid].
- [61] Greco G, Patel AS, Lewis SC, Shi W, Rasul R, Torosyan M, Erickson BJ, Hiremath A, Moskowitz AJ, Tellis WM, Siegel EL, Arenson RL, Mendelson DS. Patient-directed internet-based medical image exchange: experience from an initial multicenter implementation. *Acad Radiol* 2016;23(2):237–44.
- [62] Winden TJ, Boland LL, Frey NG, Satterlee PA, Hokanson JS, Care everywhere, a point-to-point hie tool: utilization and impact on patient care in the ed. *Appl Clin Inf* 5:2014;388–401. 25024756[pmid].
- [63] Kuperman GJ, Spurr C, Flammini S, Bates D, Glaser J. A clinical information systems strategy for a large integrated delivery network. *Proceedings AMIA Symposium* 2000:438–42. 11079921[pmid].
- [64] Willemink MJ, Koszek WA, Hardell C, Wu J, Fleischmann D, Harvey H, Folio LR, Summers RM, Rubin DL, Lungren MP. Preparing medical imaging data for machine learning. *Radiology* 2020;295(1):4–15.
- [65] Ngiam KY, Khor W. Big data and machine learning algorithms for health-care delivery. *Lancet Oncol* 2019;20(5):e262–73.
- [66] Haendel MA, Chute CG, Robinson PN. Classification, ontology, and precision medicine. *N Engl J Med* 2018;379(15):1452–62.
- [67] Erickson BJ, Korfiatis P, Akkus Z, Kline TL. Machine learning for medical imaging. *Radiographics* 2017;37(2):505–15.
- [68] Cfr - code of federal regulations title 21. url: <https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfcfr/cfrsearch.cfm>. Accessed: 2020-11-21.
- [69] Fda update transition to iso 13485:2016. url: <https://www.fda.gov/media/123488/download>. Accessed: 2020-11-21.
- [70] Muehlentatter UJ, Daniore P, Vokinger KN. Approval of artificial intelligence and machine learning-based medical devices in the USA and Europe (2015–20): a comparative analysis". *Lancet Digital Health* 2021;3(3):e195–203.
- [71] Schramowski P, Stammer W, Teso S, Brugger A, Herbert F, Shao X, Luigs H-G, Mahlein A-K, Kersting K. Making deep neural networks right for the right scientific reasons by interacting with their explanations. *Nat Mach Intell* 2020;2 (8):476–86.
- [72] Sadigh G, Hertweck T, Kao C, Wood P, Hughes D, Henry TS, Duszak R. Traditional text-only versus multimedia-enhanced radiology reporting: Referring physicians' perceptions of value. *J Am Coll Radiol* 2015;12(5):519–24.
- [73] Beesley SD, Patrie JT, Gaskin CM. Radiologist adoption of interactive multimedia reporting technology. *J Am Coll Radiol* 16(4 Part A):2019:465–471.
- [74] Rosenkrantz AB, Lui YW, Prithiani CP, Zarboulas F, Mansoubi F, Friedman KP, Ostrow D, Chandarana H, Recht MP. Development and enterprise-wide clinical

- implementation of an enhanced multimedia radiology reporting system. *J Am Coll Radiol* Dec 2014;11:1178–81.
- [75] Folio LR, Machado LB, Dwyer AJ. Multimedia-enhanced radiology reports: Concept, components, and challenges. *RadioGraphics* 2018;38(2):462–82. PMID:29528822.
- [76] Pesapane F, Codari M, Sardaneli F. Artificial intelligence in medical imaging: threat or opportunity? Radiologists again at the forefront of innovation in medicine. *Eur Radiol Exp* 2(1):2018.
- [77] Jha S, Topol EJ. Adapting to artificial intelligence: Radiologists and pathologists as information specialists. *JAMA J Am Med Assoc* 2016;316(22):2353–4.
- [78] Oakden-Rayner L. The rebirth of CAD: How is modern AI different from the CAD we know? *Radiol Artif Intell* 2019;1(3). e180089.
- [79] Neri E, de Souza N, Brady A, Bayarri AA, Becker CD, Coppola F, Visser J. What the radiologist should know about artificial intelligence – an esr white paper. *Insights Imag* 2019;10(1):1–8.
- [80] Ding Y, Sohn JH, Kawczynski MG, Trivedi H, Harnish R, Jenkins NW, Lituiev D, Copeland TP, Aboian MS, Mari Aparici C, Behr SC, Flavell RR, Huang S-Y, Zalocusky KA, Nardo L, Seo Y, Hawkins RA, Hernandez Pampaloni M, Hadley D, Franc BL. A deep learning model to predict a diagnosis of alzheimer disease by using 18f-fdg pet of the brain. *Radiology* 2019;290(2):456–64. PMID:30398430.
- [81] Qiu S, Joshi PS, Miller MI, Xue C, Zhou X, Karjadi C, Chang GH, et al., Development and validation of an interpretable deep learning framework for Alzheimer's disease classification. *Brain* 143(05):2020:1920–1933.
- [82] Lee Eun-Jae KNKD-W, Yong-Hwan Kim. Deep into the brain: Artificial intelligence in stroke imaging. *J Stroke* 19(3):2017:277–285.
- [83] Zhu G, Jiang B, Chen H, Tong E, Xie Y, Faizy TD, Heit JJ, Zaharchuk G, Wintermark M. Artificial intelligence and stroke imaging: A west coast perspective. *Neuroimag Clin* 2020;30(4):472–9.
- [84] Amukotuwa SA, Straka M, Smith H, Chandra RV, Dehkharghani S, Fischbein NJ, Bammer R. Automated detection of intracranial large vessel occlusions on computed tomography angiography. *Stroke* 2019;50(10):2790–8.
- [85] Heydon P, Egan C, Bolter L, Chambers R, Anderson J, Aldington S, Stratton IM, Scanlon PH, Webster L, Mann S, et al. Prospective evaluation of an artificial intelligence-enabled algorithm for automated diabetic retinopathy screening of 30 000 patients. *Br J Ophthalmol* 2020.
- [86] Grzybowski A, Brona P, Lim G, Ruamviboonsuk P, Tan GS, Abramoff M, Ting DS. Artificial intelligence for diabetic retinopathy screening: a review. *Eye* 2020;34(3):451–60.
- [87] Eudamed european database on medical devices. url:ec.europa.eu/tools/eudamed. Accessed: 2021-05-03.
- [88] D. o. H. Centers for Medicare & Medicaid Services (CMS) and H.S. (HHS), Medicare and medicaid programs; electronic health record incentive program. final rule, *Fed Regist* 75(144):2010;44313–588.
- [89] Heisey-Grove D, Danehy L-N, Consolazio M, Lynch K, Mostashari F. A national study of challenges to electronic health record adoption and meaningful use. *Med Care* 2014;52(2):144–8.
- [90] Menachemi N, Rahurkar S, Harle CA, Vest JR. The benefits of health information exchange: an updated systematic review. *J Am Med Inf Assoc* 25(04):2018: 1259–1265.
- [91] Kruse CS, Regier V, Rheinboldt KT, Barriers over time to full implementation of health information exchange in the united states. *JMIR Med Inform* 2(Sep):2014; e26.
- [92] Wu H, Larue E. Barriers and facilitators of health information exchange (hie) adoption in the united states. In: 2015 48th Hawaii international conference on system sciences; 2015. p. 2942–9.
- [93] Bradford L, Aboy M, Liddell K. International transfers of health data between the EU and USA: a sector-specific approach for the USA to ensure an 'adequate' level of protection. *J Law Biosci* 10:2020. Isaa055.
- [94] Whitsel LP, Wilbanks J, Huffman MD, Hall JL. The role of government in precision medicine, precision public health and the intersection with healthy living. *Progr Cardiovasc Diseases* 62(1):2019;50–54. Merging precision and healthy living medicine: tailored approaches for chronic disease prevention and treatment.
- [95] Miller KE, Lin SM. Addressing a patient-controlled approach for genomic data sharing. *Genet Med* 2017;19(11):1280–1.
- [96] Willemink MJ, Koszek WA, Hardell C, Wu J, Fleischmann D, Harvey H, Folio LR, Summers RM, Rubin DL, Lungren MP. Preparing medical imaging data for machine learning. *Radiology* 2020;295(1):4–15. PMID:32068507.
- [97] Abdelaziz Ismael SA, Mohammed A, Hefny H. An enhanced deep learning approach for brain cancer MRI images classification using residual networks". *Artif Intell Med* 2020;102. 101779.
- [98] Kazemina S, Baur C, Kuijper A, van Ginneken B, Navab N, Albarqouni S, Mukhopadhyay A. GANs for medical image analysis. 2018;1–40 arXiv.
- [99] Tseng KL, Lin YL, Hsu W, Huang CY. Joint sequence learning and cross-modality convolution for 3D biomedical segmentation. *Proceedings - 30th IEEE conference on computer vision and pattern recognition, CVPR 2017*, vol. 2017-Janua, no. c, 2017, pp. 3739–3746.
- [100] Bhanumurthy MY, Anne K. An automated detection and segmentation of tumor in brain MRI using artificial intelligence". In: 2014 IEEE international conference on computational intelligence and computing research, IEEE ICCIC; 2014. 2015..
- [101] Zheng H, Yang L, Chen J, Han J, Zhang Y, Liang P, Zhao Z, Wang C, Chen DZ. Biomedical image segmentation via representative annotation. 33rd AAAI conference on artificial intelligence, AAAI 2019, 31st innovative applications of artificial intelligence conference, IAAI 2019 and the 9th AAAI symposium on educational advances in artificial intelligence, EAAI 2019, 2019, no. 1, p. 5901–5908.
- [102] Atlason HE, Love A, Sigurdsson S, Gudnason V, Ellingsen LM. Unsupervised brain lesion segmentation from MRI using a convolutional autoencoder, 2019, p. 52.
- [103] Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. In: Medical image computing and computer-assisted intervention – MICCAI 2015 (N. Navab, J. Hornegger, W.M. Wells, and A.F. Frangi, eds.), (Cham), pp. 234–241, Springer International Publishing, 2015.
- [104] Khanh TLB, Dao DP, Ho NH, Yang HJ, Baek ET, Lee G, Kim SH, Yoo SB. Enhancing U-net with spatial-channel attention gate for abnormal tissue segmentation in medical imaging". *Appl Sci (Switzerland)* 2020;10(17):1–19.
- [105] Oktay O, Schlemper J, Folgoc LL, Lee M, Heinrich M, Misawa K, Mori K, McDonagh S, Hammerla NY, Kainz B, Glocker B, Rueckert D. Attention U-Net: Learning where to look for the pancreas. *Midl* 2018.
- [106] Dong H, Yang G, Liu F, Mo Y, Guo Y. Automatic brain tumor detection and segmentation using u-net based fully convolutional networks. In: Annual conference on medical image understanding and analysis. Springer; 2017. p. 506–17.
- [107] Dalca AV, Guttat J, Sabuncu MR. Anatomical priors in convolutional networks for unsupervised biomedical segmentation. In: *Proceedings of the IEEE computer society conference on computer vision and pattern recognition*; 2018. p. 9290–9.
- [108] Raghu M, Zhang C, Kleinberg J, Bengio S. Transfusion: Understanding transfer learning for medical imaging. *Adv Neural Inf Proc Syst* 32(NeurlPS):2019.
- [109] Zhang Z, Xie Y, Xing F, McGough M, Yang L. MDNet: A semantically and visually interpretable medical image diagnosis network. *Proceedings – 30th IEEE conference on computer vision and pattern recognition, CVPR 2017*, vol. 2017-Janua, 2017, pp. 3549–3557.
- [110] Bahdanau D, Cho K, Bengio Y. Neural machine translation by jointly learning to align and translate; 2014. arXiv preprint arXiv:1409.0473.
- [111] Ouyang C, Biffi C, Chen C, Kart T, Qiu H, Rueckert D. Self-supervision with superpixels: Training few-shot medical image segmentation without annotation 2020;1:762–80.
- [112] Zhou Y, He X, Huang L, Liu L, Zhu F, Cui S, Shao L. Collaborative learning of semi-supervised segmentation and classification for medical images. In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2019-June, 2019, p. 2074–83.
- [113] Yu L, Cheng J-Z, Dou Q, Yang X, Chen H, Qin J, Heng P-A. Automatic 3d cardiovascular mr segmentation with densely-connected volumetric convnets; 2017.
- [114] Zheng H, Zhang Y, Yang L, Liang P, Zhao Z, Wang C, Chen DZ. A new ensemble learning framework for 3D biomedical image segmentation. In: 33rd AAAI conference on artificial intelligence, AAAI 2019, 31st innovative applications of artificial intelligence conference, IAAI 2019 and the 9th AAAI symposium on educational advances in artificial intelligence, EAAI 2019, no. Wolpert, 2019. p. 5909–16.
- [115] Adadi A, Berrada M. Peeking inside the black-box: A survey on explainable artificial intelligence (xai). *IEEE Access* 2018;6:52138–60.
- [116] Jia X, Ren L, Cai J. Clinical implementation of ai technologies will require interpretable ai models. *Med Phys* 2020;47(1):1–4.
- [117] Reyes M, Meier R, Pereira S, Silva CA, Dahlweid F-M, von Tengg-Kobligh H, Summers RM, Wiest R. On the interpretability of artificial intelligence in radiology challenges and opportunities. *Radiol Artif Intell* 2020;2(3). e190043.
- [118] Interpretability of a deep learning model in the application of cardiac mri segmentation with an adc challenge dataset.
- [119] Tsiknakis N, Trivizakis E, Vassalou E, Papadakis G, Spandidos D, Tsatsakis A, Sánchez-García J, López-González R, Papanikolaou N, Karantanis A, Marias K. Interpretable artificial intelligence framework for COVID-19 screening on chest X-rays. *Exp Therapeut Med* 2020;20(2):727–35.
- [120] Veasey BP, Broadhead J, Dahle M, Seow A, Amini AA. Lung nodule malignancy prediction from longitudinal ct scans with siamese convolutional attention networks 2020.
- [121] Górriz M, Antony J, McGuinness K, Giró-i Nieto X, O'Connor NE. Assessing knee OA severity with CNN attention-based end-to-end architectures, 2019. p. 1–13.
- [122] Kingma DP, Welling M. Auto-encoding variational bayes, 2013. arXiv preprint arXiv:1312.6114.
- [123] Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, et al., Generative adversarial nets. In: *Advances in neural information processing systems*, 2014. p. 2672–80.
- [124] De Cao N, Kipf T. Molgan: An implicit generative model for small molecular graphs 2018. arXiv preprint arXiv:1805.11973.
- [125] Kazemina S, Baur C, Kuijper A, van Ginneken B, Navab N, Albarqouni S, Mukhopadhyay A. Gans for medical image analysis. *Artif Intell Med*, 2020. p. 101938.
- [126] Wolterink JM, Dinkla AM, Savenije MH, Seevinck PR, van den Berg CA, Išgum I. "Deep mr to ct synthesis using unpaired data. In: *International workshop on simulation and synthesis in medical imaging*. Springer; 2017. p. 14–23.
- [127] Bi L, Kim J, Kumar A, Feng D, Fulham M. Synthesis of positron emission tomography (pet) images via multi-channel generative adversarial networks (gans), in molecular imaging, reconstruction and analysis of moving body organs, and stroke imaging and treatment. Springer; 2017. p. 43–51.
- [128] Cho H, Lim S, Choi G, Min H. Neural stain-style transfer learning using gan for histopathological images 2017. arXiv preprint arXiv:1710.08543.
- [129] Shaban MT, Baur C, Navab N, Albarqouni S. Staining: Stain style transfer for digital histological images. In: 2019 IEEE 16th international symposium on biomedical imaging (ISBI 2019), IEEE, 2019. p. 953–6.
- [130] Chaudhari AS, Grissom MJ, Fang Z, Sveinsson B, Lee JH, Gold GE, Hargreaves BA, Stevens KJ. Diagnostic accuracy of quantitative multi-contrast 5-minute knee mri

- using prospective artificial intelligence image quality enhancement. *Am J Roentgenol* 2020.
- [131] Chaudhari A, Fang Z, Hyung Lee J, Gold G, Hargreaves B. Deep learning super-resolution enables rapid simultaneous morphological and quantitative magnetic resonance imaging. *Lecture notes in computer science (including subseries lecture notes in artificial intelligence and lecture notes in bioinformatics)*, 2018, vol. 11074 LNCS, pp. 3–11.
- [132] Recht MP, Zbontar J, Sodickson DK, Knoll F, Yakubova N, Sriram A, Murrell T, Defazio A, Rabbat M, Rybak L, Kline M, Ciavarra G, Alaia EF, Samim M, Walter WR, Lin D, Lui YW, Muckley M, Huang Z, Johnson P, Stern R, Zitnick CL. Using deep learning to accelerate knee mri at 3t: Results of an interchangeability study. *Am J Roentgenol* 2020/07/09 2020..
- [133] Chen KT, Gong E, Bezerra F, Macruz DC, Xu J. Ultra – Low-dose 18 F-florbetaben amyloid pet imaging using deep learning with multi-contrast MRI inputs, no. 10, 2019.
- [134] Gong E, Pauly JM, Wintermark M, Zaharchuk G. Deep learning enables reduced gadolinium dose for contrast-enhanced brain MRI. *J Magn Reson Imag* 2018;48 (2):330–40.