NEXT GEN RADIOLOGY AI

The Journey from an Algorithm to a Clinical Solution

TABLE OF CONTENTS

EXECUTIVE SUMMARY

The AI Ecosystem: Where we are and where we're going

As a medical specialty, radiology is extremely data dependent. Consequently, radiology is the first specialty to suffer from data overload as well as the first to benefit from the proper utilization of data via novel technologies. In this context there are great expectations for the new generation of radiological AI systems that are likely to impact the day to day radiological workflow.

From its inception, the AI ecosystem has had two driving forces. On one side were the visionaries that talked about long-term trends. They provided a vision, fostering a nurturing atmosphere and encouraged investment. On the other side lay the technical experts who focused on the "nittygritty" of algorithms. In retrospect, no progress was possible without collaboration between these two communities. The time has come to close the gap between these two groups.

For the transformation of Next-generation AI to succeed, the AI ecosystem must adopt a holistic view of the entire radiological work process. The AI ecosystem should move from talking about algorithms and models to encompass both the clinical outcomes as well as economic benefits achieved from AI augmented workflow.

To succeed, we believe that the Medical Imaging AI ecosystem will mature into a three-tier system including an Algorithmic, Product and Solution layer. Each layer would address different aspects of the overall solution (and feature different key performance indicators (KPIs)). However, only the combination of all the three layers will bring true value to this field.

This whitepaper reviews the current status of AI and explores what it will require to bring the concept of a complete three-tier solution to fruition. It is designed to help radiologists, informatics experts and other healthcare professionals understand the new direction of AI and how Next Generation AI will benefit radiologists and patients alike.

Algorithmic Layer

Second generation AI systems are expected to offer both increased accuracy levels and enhanced functionality, well beyond that of the conventional CAD systems. With the advent of deep learning, similar to other fields, next generation AI systems should offer accuracy levels similar to those of humans, especially in terms of the likelihood of detection. With access to more data sources, new algorithms promise to be more robust.

Moreover, a diverse collection of functionalities would need to be introduced. The new goals would include anomaly detection, followed by its delineation and classification. In the context of longitudinal studies, a separate set of algorithms would permit anomalies to be registered then tracked to detect and evaluate changes. Finally, a more advanced set of algorithms would permit a fusion between different systems including various imaging modalities as well as clinical and genomic data.

Algorithmic efficacy would be measured by the traditional accuracy metrics such as sensitivity and specificity. ROC curves will continue to reign supreme.

Product Layer

The Product Layer will use software to integrate the algorithms into the physician's workflow. At this stage, the KPIs would include the speed and accuracy of the combined system, where the physician is assisted by the automated AI tools. Auxiliary KPIs would measure factors such as the users' training curve, engagement and educational prerequisites (to identify use-cases where, using AI, some subtasks can be delegated from physicians to the paramedical personnel).

Solution Layer

The solution layer would address specific clinical use-cases. Note that several AI products may be needed to create a single coherent solution. The goal would be to improve clinical outcomes, increase revenues and to reduce the overall costs for the patients, payers, providers or regulators or preferably to all four of them.

In principle, all three layers are independent from each other. For instance, one can use the same product software to integrate different AI algorithms creating a unified platform. It is important, however, to consider all three layers simultaneously. For instance, one can develop an algorithm that features superior sensitivity that stems from the detection of small, clinically insignificant lesions. In this case, high sensitivity wouldn't reflect true clinical value.

Given that algorithmic ROC curves are the easiest to measure, there is a danger of the "streetlight effect." Hence, it is important to find ways to emphasize the focus on the true clinical outcomes.

THE ALGORITHMIC LAYER – ENHANCED FUNCTIONALITY OF THE NEW AI GENERATION

Traditional CAD (or CADe) systems provide lesion detection only. Newer (CADx) systems provide both lesion detection and interpretation as well. Next Generation AI systems will provide an even wider gamut of medical applications. In this section we'll review, briefly, the extended functionality of Next Gen AI algorithms.

Lesion detection

Lesion detection will continue to be one of the major functions of medical AI. The typical mode of operation here will be as follows: user enters a medical image into the system, the system returns a list of suspicious areas in the image. Optionally, the system may also indicate some measure of confidence in the lesion detection.

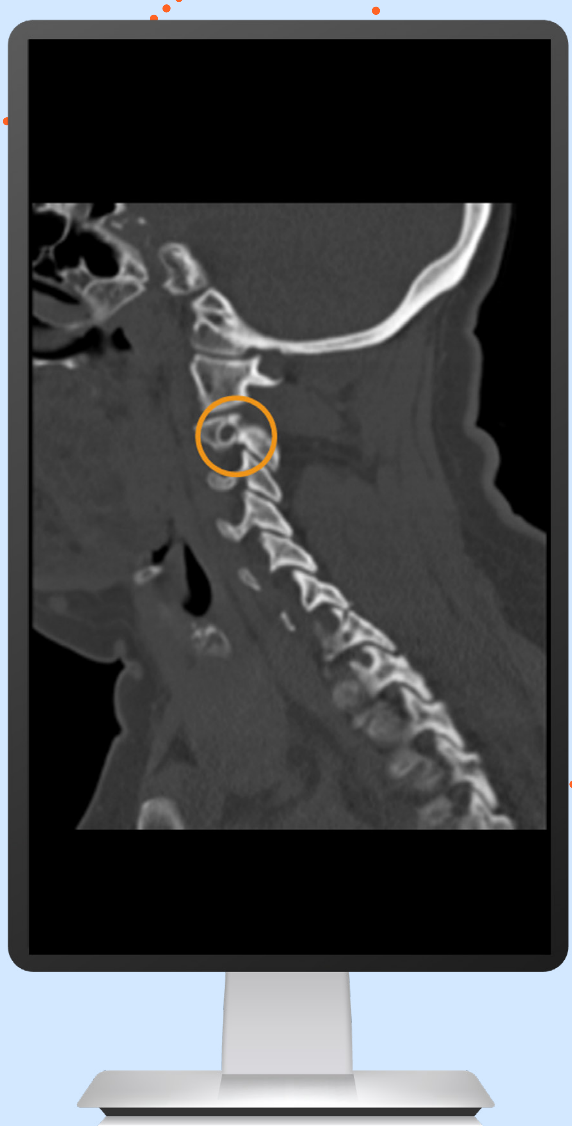

Case In Point

Example output of detection algorithm Aidoc detection of a subtle pedicle fracture in C2

Lesion Classification

In this case, in the typical mode of operation a user enters the area of interest and the system outputs lesion classification. Usually, the input to this module would be provided by automatic lesion detection, in some instances however, physicians may prefer to point manually at the area of interest to them.

Example: A physician may prefer to forgo automatic lesion detection in instances where the application offers relatively low specificity. In this case, the radiologist would be forced to review a large number of false positives, which would lead to a significant slowdown of the diagnostic process. As a result, radiologists may prefer the "point and shoot" approach by engaging the AI analytic capabilities on specific areas of interest which the radiologist believes may indicate a lesion or area of concern.

Lesion classification results can be used in a variety of applications. In some specific vertical applications (e.g. diabetic retinopathy), lesion classification would be the ultimate system output. In other words, the system receives an image of the area of interest and produces a diagnostic result. However, since there are no infallible systems, lesion classification would typically be only one step in the diagnostic process. For example, lesion classification may be used for triage allowing optimized organization of the physician's worklist or, alternatively, helping to route a given study to the most suitable radiologist.

Another approach would be to use lesion classification to reduce the workload by automatic handling of relatively simple cases. For instance, in the screening context, the automatic system can identify obviously normal cases so that expert radiologists would focus on the complex cases only. Alternatively, lesion classification can be used to identify cases (e.g. in the teleradiology setting) when additional tests are needed. As a result, the number of the recurring patient visits would be reduced.

It is important to note that each use case requires different AI settings. In some instances, very high sensitivity may be required. In this case, even 20% specificity will be useful in determining those cases that do not require additional manual analysis. In other cases, we need to determine automatically whether additional patient tests are required. As these case studies indicate, to realize AI's full potential, it can be deployed in many different formats, either as a fully automated system, detecting anomalies for the radiologist to validate, or waiting in the background to verify the radiologist's findings upon demand.

Explainability and Lesion Description

Traditional AI systems function as "black boxes" i.e. they provide useful output to the user, but they fail to provide justification for their findings. However, in many applications it is important for the user to know not only what was detected, but also why. In this context, it is advantageous for the system to provide automatic lesion delineation. Assuming that the delineation is known, the system can be enhanced to provide lesion measurements, shape and other characteristics. For instance, the AI system will draw the line following the lesion border and determine that, say, the lesion shape is oval, the border is sharp and the internal structure is heterogeneous.

An additional example of this concept would be to provide a quantitative estimate of the lesion risks. For example, the American College of Radiology (ACR) defined the standard system of BI-RADS that express apparent breast lesion malignancy. A similar approach is being extended to other domains such as LUNG-RADS for estimation of malignancy of lung cancer and LI-RADS for liver. Accordingly, an AI system can be trained to provide direct estimation of the corresponding BI-RADS parameters (or similar parameters in other clinical domains). System output would include both an overall BI-RADS estimate as well as measurements of the lesion features that were used in the BI-RADS determination. So, the radiologist will have both a malignancy estimate and an explanation of the AI thinking process.

Case In Point

Detects solid, sub-solid and ground glass nodules with high accuracy. Linear and volumetric measurement of the nodules; In sub-solid nodules the measurement is of the solid core

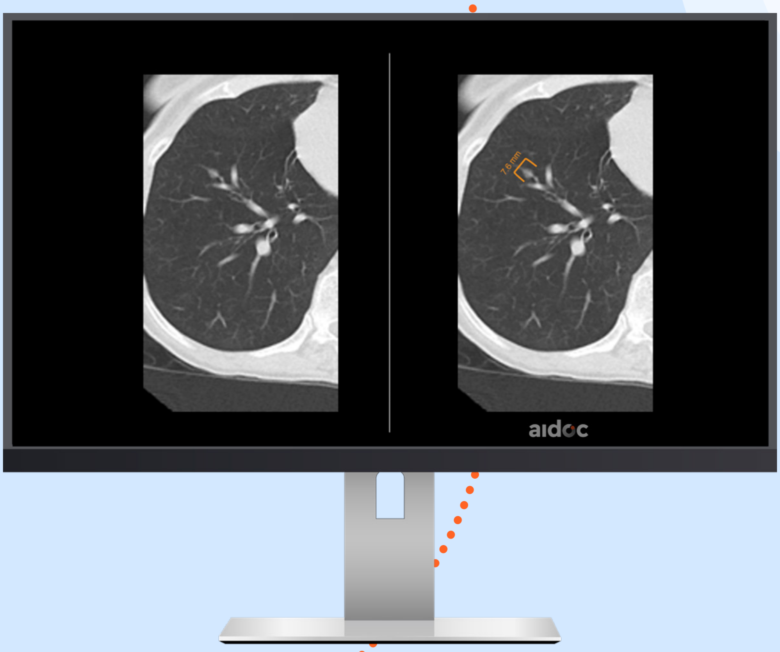

Longitudinal Studies

Given the lesion evaluation, the AI system can analyze longitudinal sequences of studies providing change detection and measurement. Change detection algorithms will naturally require registration modules which are extremely challenging as well, especially for soft organs that are likely to be distorted between different image acquisition sessions.

Multi-modal Studies

Greater challenges are posed in instances where the same lesion is imaged by different modalities; for example, CT and MRI. In this context, AI would be used not only for lesion detection and delineation, but also for registration between findings in different modalities. Note that in some cases no direct image registration is possible. For example, X-Ray, which provides body projection, can't be directly combined with ultrasound, which provides a cross section of the body. In these cases, more abstract registration is needed, based on an understanding of the objects in question. In addition, relevant information may be provided by non-imaging modalities such as clinical information, genomics or other omics. AI systems are likely to provide an ideal platform for combining disparate data relevant to the diagnosis of a given patient.

Diagnostic Support

Assuming that holistic patient information is available, it may be possible to provide diagnostic support in the form of a differential diagnosis. In this context, AI can be instrumental in providing a wide variety of suggestions regarding the differential diagnosis. However, we believe that, in the foreseeable future, the final determination would be made by the physician in charge. The question is how to enhance the utility of the AI to aid physicians in reaching a final diagnosis. Below we have outlined one mechanism, the use of AI to offer diagnostic suggestions. However other mechanisms would be very important as well.

Optimized Data Retrieval of the relevant Patient Information

In modern healthcare systems, physicians suffer from a deluge of information. For many patients, their electronic medical record (EMR) is thicker than *War and Peace*. Everybody agrees that EMR contains vital information. However, in many cases, the radiologist may only have five minutes, at most, with each patient. As a result of this time constraint, the radiologist report is often based solely on the last image. To mitigate this challenge, AI could be used to retrieve and display the most relevant pieces of information. This in turn makes accessing important information quick, feasible and convenient.

In some traditional systems, predefined templates are used. For example, in the screening scenario, the system would automatically fetch prior studies facilitating manual change detection. However, AI is likely to revolutionize this approach. For example, one can imagine an AI system that would prefetch past information (medical images, clinical signs and symptoms, lab results) that may influence the differential diagnosis. This would enable radiologists to be more effective than ever before.

Optimized Data Retrieval of the relevant Scientific Information

In addition to the EMR information deluge, physicians must cope with the deluge of scientific literature. Some physicians use Google to find sample cases of the given disease. However, such keyword-based searches are relatively ineffective. Based on AI tools, a radiologist may search for images that are diagnostically similar to the patient under consideration.

One possible scenario would be as follows: The radiologist reviews the images and suspects that the patient suffers from a rare disease "A." The AI decision support system reaches the same conclusion. However, the differential diagnosis includes also another disease "B." The radiologist is not sure. In order to assist him in the decision process, the AI would fetch several examples of disease A and disease B, each having the most similar manifestation to the case under consideration. Now, the radiologist must determine whether given patient image resembles group "A" or "B."

Index Case

45 yo woman with abdominal pain.
Suspicious growing lesion on CT. Referred for MRI evaluation.

Case In Point

AI Driven Diagnostic Support

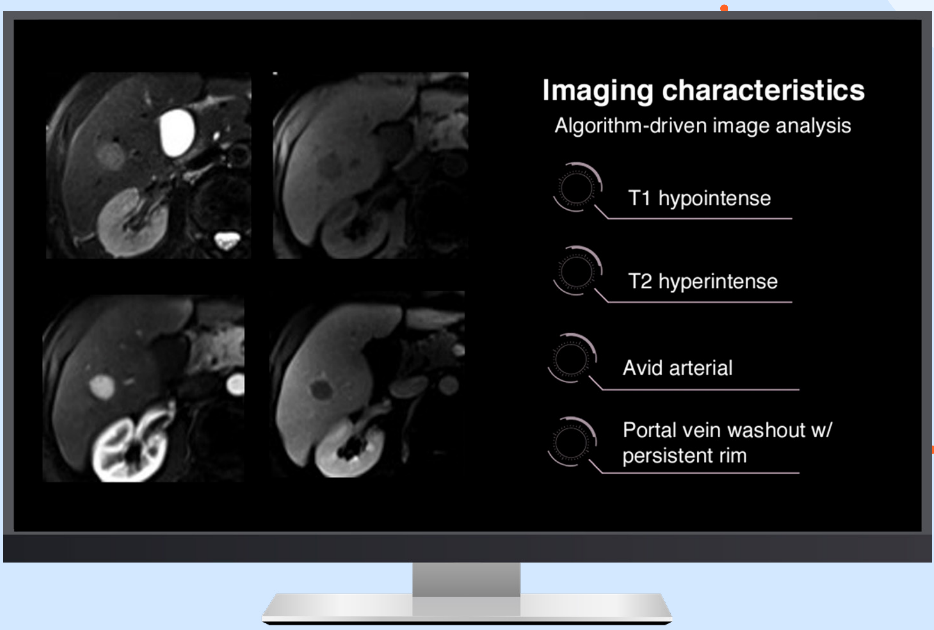

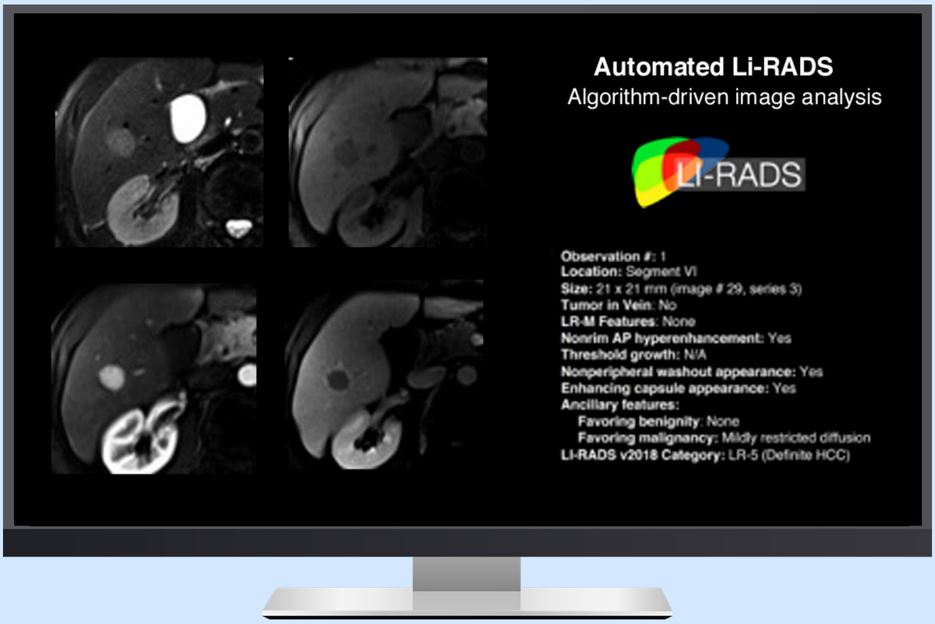

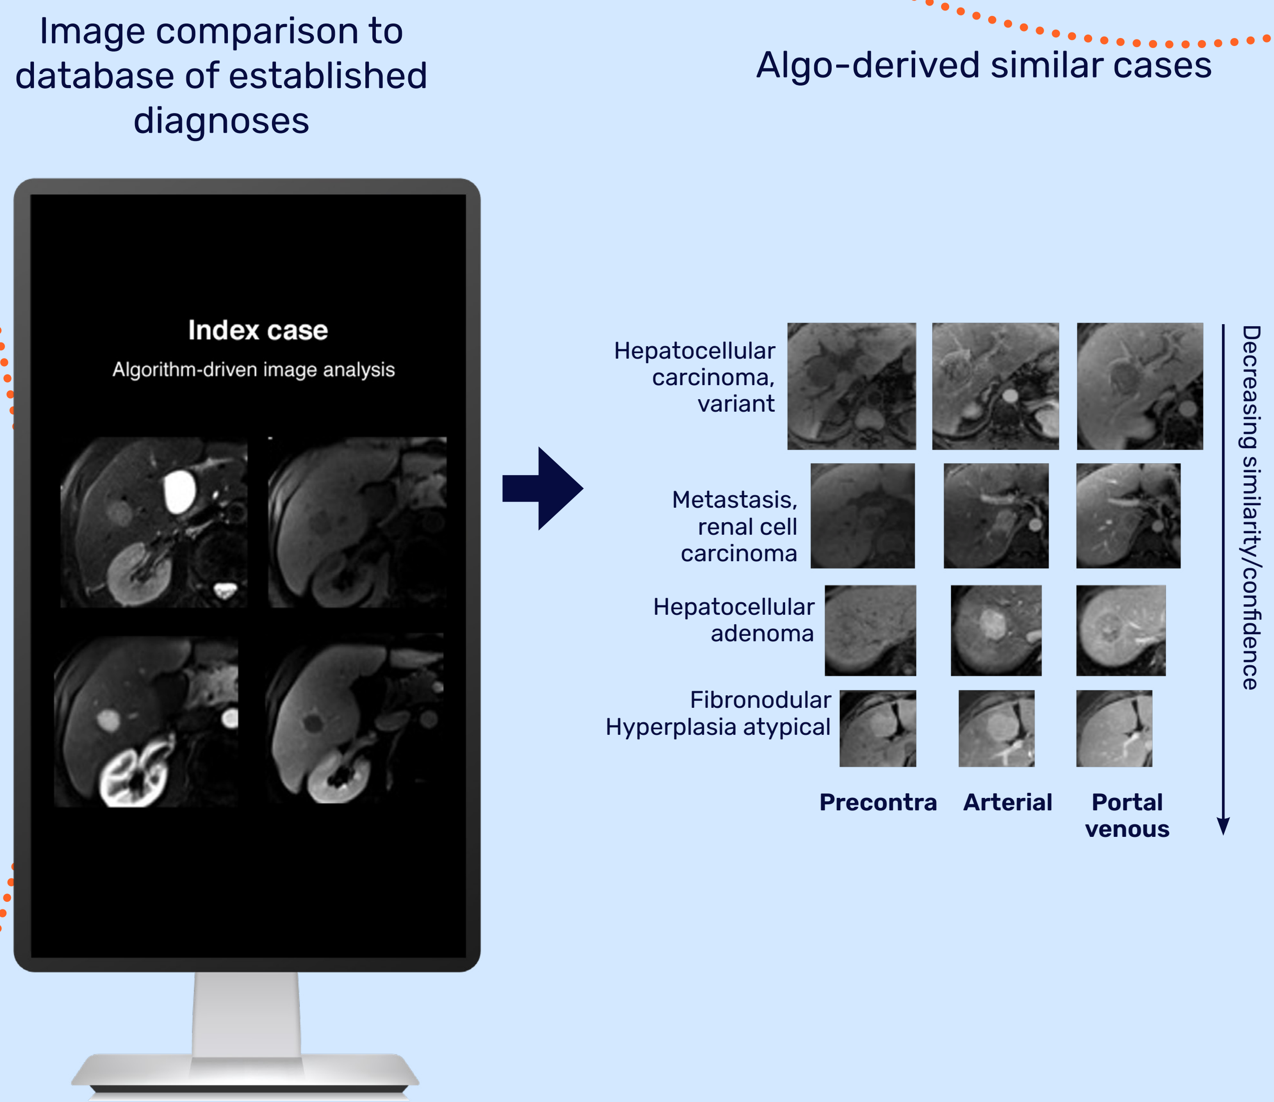

Risk assessment

Risk assessment is an essential part of each treatment plan determination. It is also an essential part of patient empowerment. However, traditional risk assessment modules are based on general information (such as patient family history) only. Such modules rarely, if ever, take into account patient specific data such as medical imaging. AI is likely to change that by enabling individualized risk assessment based on the holistic clinical picture.

Patient management recommendation

Last but not least, the physician's role is to recommend optimal patient management i.e. sequence of diagnostic and therapeutic steps. To perform this, task, physicians draw on their extensive knowledge and experience. However, AI can be instrumental in this decision-making process. The AI systems can be automated to consider the cost and efficacy of each step, taking into account patient preferences (e.g. some patients may feel claustrophobic while in an MRI machine). Then, out of a myriad of possible combinations, optimal course of action can be chosen.

Visualization

Diagnostic reasoning relies on a wide variety of information. It is of crucial importance to display it in the optimal manner. Since screen "real estate" is very important, AI is likely to play a crucial role in adapting the display taking into account both the needs of the specific patient and the personal preferences of the physician.

Another visualization approach gaining momentum is that of performing 3D printing of the affected organs in order to facilitate presurgical planning. Once again, AI is likely to be instrumental in providing reliable tissue segmentation essential in 3D organ modeling.

Interventional Radiology

Above we have focused on the potential role of AI in the diagnostic support. This reflects most of the current medical imaging AI systems. However, AI has a crucial role to play in the field of Interventional Radiology as well. This role would range from assisting in surgical navigation to facilitating automated surgery procedures.

BUILDING INTELLIGENT ALGORITHMS

"Data is the new oil"

In the conventional AI systems, it was up to the AI and Image Processing experts to determine the set of optimal classification features. For example, if a lump is discovered in mammographic breast image, an algorithm developer, in collaboration with physicians, would determine that the likelihood of malignancy can be estimated based on the lesion, border sharpness, shape and internal structure. Image processing algorithms for estimation of these features would be written. Finally, a variety of classifiers would be trained on a feature set computed for all the images in the given benchmark. After the training, for each new image, the system would be able to predict whether a given lesion is malignant or not.

In modern deep learning systems, the conventional development process has been streamlined. Now feature selection is done automatically. So, in principle, the creation of an AI system seems to be straightforward: one puts in a sufficiently large benchmark of annotated data, including both normal and abnormal cases and the system would learn, automatically, how to distinguish between these two categories. Clearly AI is only as good as the data being used for training. Hence, in some quarters, there is an idea that AI efficacy would be driven by the single factor: the amount of data available for training.

In some cases, we notice a trend where AI applications are being developed by physicians, thereby bypassing dedicated AI experts. This is a welcome trend, allowing physicians to quickly verify the feasibility of a large number of different hypothesis. However, in reality, as explained below, development of AI algorithms poses a series of significant challenges. Hence, we believe that the winning combination must rely on a close collaboration between radiologists and AI experts.

Challenges in the creation of the modern AI systems

The biggest problem of deep learning system stems from the volume of the available data. In some natural image classification problems, experts work with millions of data samples. Such data volumes seem to be impractical in the medical domain. Even with smaller benchmarks, developers face the problem of data imbalance (typically there are much more normal cases than abnormal ones). This imbalance can skew the results.

Classification results are highly dependent on the quality of annotation. In conventional image classification problems this is not a big issue (after all it is not that difficult to determine whether a given image contains cat or dog). However, in medical applications, we frequently face significant interobserver variability.

Another widely used approach is that of k-Fold Cross-Validation. A given data sample is split into k groups. The AI training process is rerun k times. Each time one group serves for testing while remaining groups serve for training. The overall result is estimated as an average of all k runs. It is a simple and useful approach for estimation of the AI system robustness, but it can't overcome inherent data weaknesses (if such exist).

Fortunately, in most applications, expert AI teams can meet all the challenges of creating an effective AI algorithm. The question remains, how do we determine whether or not an AI algorithm has reached full maturity? That "perfect annotation" is infeasible. Moreover, annotation of large data benchmarks tends to be very expensive. Typically, one has so called "weak annotation" (e.g. final diagnosis for the given patient). Indeed, such diagnostic annotation is likely to be included in the radiological report. However, radiological reports seldom include "strong annotation" providing lesion delineation (either precise or on the bounding box level). Hence, in reality, one must frequently work with benchmarks having both strong and weak annotation with various degrees of annotation certainty.

The strength of deep learning systems stems from their huge learning capacity. Such systems can feature hundreds of thousands or even millions of parameters. However, these strengths can turn into weakness if the system reaches so called overtraining (i.e. situations where the system learns the given training benchmark while losing its ability of abstraction to other benchmarks). One way to mitigate this issue is by gathering data from a wide variety of medical centers and image acquisition systems. However, doing this may be challenging given that developers must meet various regulatory constraints related to data security.

THE INTELLIGENT ALGORITHM

Each classification algorithm splits data into the following four categories: tp (true positives), fp (false positives), tn (true negatives) and fn (false negatives).

Once these values are measured one can determine the algorithm's sensitivity and specificity
Sensitivity = tp / (tp + fn)
Specificity = tn / (tn + fp)
Alternative representation is that of precision and recall
Recall = tp / (tp + fn) (the same as sensitivity)
Precision = tp /(tp + fp)

Frequently, physicians still prefer a different interpretation of the above results: they use Negative Predictive Values (NPV) and Positive Predictive Values (PPV) defined as
NPV= tn / (tn +fn)
PPV= tp / (tp + fp)
Indeed, one can view NPV as a measure of physician's "peace of mind", while determining that the patient is healthy. Similarly, PPV can be viewed as a measure of "distraction" caused by false alarms.

Naturally, by the simple change of threshold, one can trade off sensitivity at the expense of specificity. Hence, classification algorithm should be viewed as a curve of points (rather than a single number). This curve is known as the receiver operating curve (ROC). It shows mutual dependence between Sensitivity vs Specificity (Precision vs Recall, NPV vs PPV). Hence, the comparison between two AI systems boils down to the comparison between the two corresponding ROC curves. This poses a question: how do we compare two curves? Or, conversely, how do we represent a curve by a single number? Accordingly, many researchers represent system efficacy by the area under the curve (AUC). Clearly, perfect system will have AUC=1. Random system will have AUC=0.5. All the real-life systems will fit between these two numbers.

Unfortunately, the use of AUC may be very misleading. Consider for instance benchmarks obtained from the population in a cancer consulting practice. 99% of the images come from the cancer patients and only 1% from the healthy individuals. So, if we create a "classifier" that always returns cancer determination, we end up with a recall of 100% and precision of 99%. This yields an AUC of 0.99 for a completely useless system. Naturally, we can create a balanced test benchmark but then it wouldn't represent the true clinical practice. We believe that the best approach would be to choose the working point that fits actual clinical practice (say the specificity that would be acceptable).

Case In Point

An ROC curve showing Sensitivity vs. Specificity of an ICH Brain algorithm

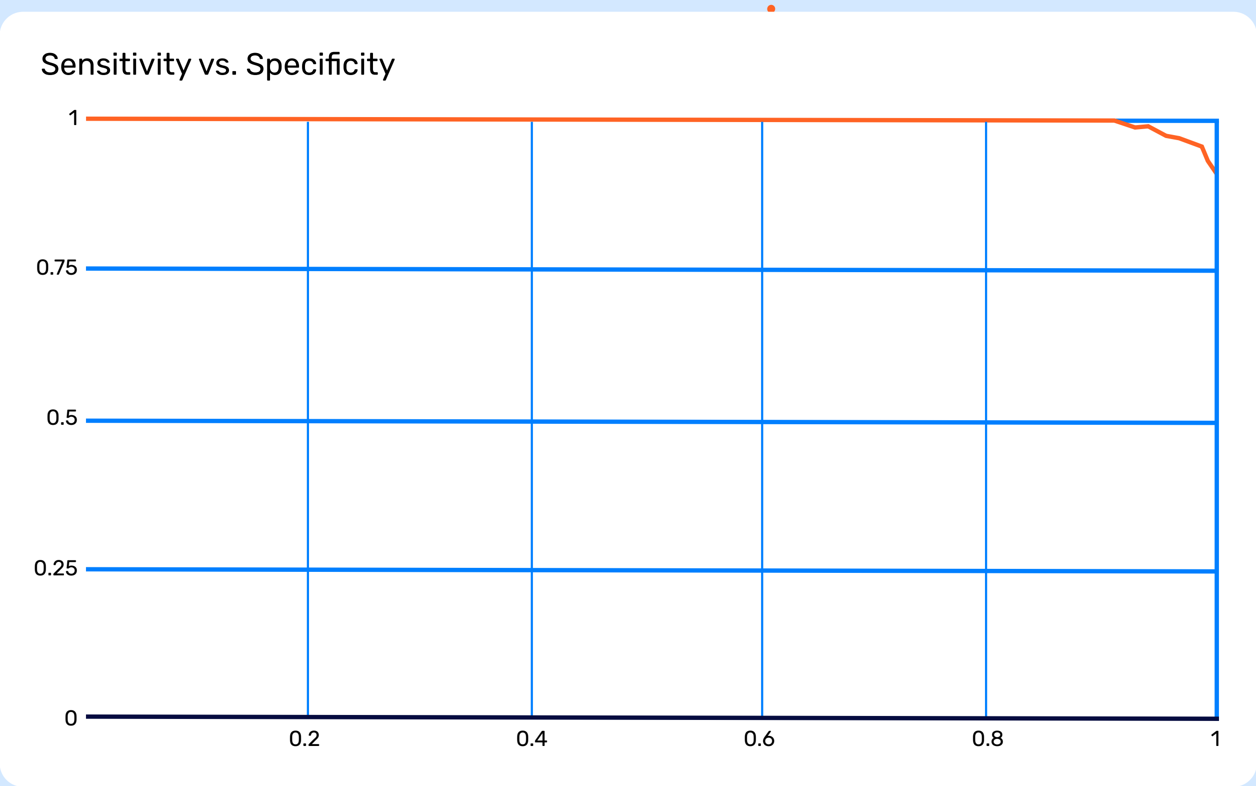

THE PRODUCT LAYER - PERFECTING THE PRODUCT

Development of an effective AI algorithm is an essential first step in the product development process. However, its usefulness would be limited without taking into account the product envelope. The first issue to consider would be that of integration with the data repositories including existing PACS and EMR systems. Even more important would be workflow integration. Indeed, if AI systems are to be widely accepted by the users, they must cause minimal disruption of the existing workflow and be intuitive enough to streamline the learning curve.

In order to achieve these goals, it is important to understand what is going to be the target user population. In some cases, the primary user would be attending physicians. In other cases, the target user population may be the residents. In other cases, the primary benefit of the AI systems would be in offloading certain tasks to paramedical personnel, thereby freeing physicians to attend more complex duties.

These distinctions are very important. For example, in some cases we have encountered department heads who noticed that introducing an AI system actually resulted in an increase in their personal workload. For example, they were required to explain to their residents some false positives detected by the AI system. Still, they found AI to be very beneficial due to the improved overall outcome and due to the steeper learning curve for the residents. Hence, the benefit of AI systems must be considered holistically in terms of its impact on physician workflow as well as clinical outcomes.

Once the Product layer is completed, its efficacy can be evaluated. However, instead of the traditional ROC curves, one must measure overall system parameters i.e. to evaluate performance of the physicians working with the AI assistant. For example, we can measure the overall system sensitivity as a function of productivity for the target user population.

Supplementary KPIs would include training time. In this context, it is advantageous to use products that require almost no training time. Indeed, some AI systems require no changes to the existing workflow. For instance, they communicate with the user via an external window.

Case In Point

AI in clinical use:
Aidoc widget

Case In Point

AI in clinical use:
Aidoc widget

Another class of AI products could be designed to solve the shortage of medical experts by delegation of certain tasks to paramedical personnel. AI may play a crucial role in automatically splitting the patient stream into five categories

(i) obviously healthy
(ii) probably healthy subject to verification by a paramedical assistant

(iii) conventional cases that needs radiologist's attention
(iv) emergency cases that require immediate attention
(v) non-typical cases (this category would include coincidental diagnosis not foreseen by the original patient referral)

Naturally, different stakeholders would focus on different KPIs. We expect that hospital administrators may focus on the overall product efficacy that may be evaluated by the cost reduction for the given quality of results.

What do we use internally for system evaluation? All of the above but our absolute favorite is that of user's engagement that measures the percentage of use by the physicians. We love it when physicians "vote with their feet" by using the product. For us, this single measure supersedes all the rest.

THE SOLUTION LAYER - THE FULL-SCALE SOLUTION

The Solution Layer would be used to address specific clinical use-cases. Note that several AI products may be needed to create a single coherent solution. In this context, the goal would be to improve clinical outcomes and to reduce overall costs for either patients, payers, providers or regulators (preferably to all four of them).

Ideally, full scale AI solution would facilitate all aspects of the radiologist's workflow. See below one example of the lifecycle of an exam, created by Nina Kottler of Radiology Partners. This represents the various aspects of the radiology workflow in which AI could have impact.

PATIENT SCHEDULING

This stage would optimize patients' handling before their arrival at the imaging facility.

IMAGE ACQUISITION

The goal would be to optimize image acquisition including such parameters as minimization of the radiation exposure.

EXTENDED TRIAGE

By this term we mean both optimization of the physician's worklist and optimization of the supplementary information to be shown to the physician both in terms of the patient's EMR and state of the art medical literature.

DIAGNOSTIC DECISION SUPPORT

This category encompasses all of the AI tools that facilitate differential diagnosis including detection of lesions, evaluation, classification and longitudinal monitoring up to the stage when the physician decides on the differential diagnosis and patient management recommendation.

TREATMENT SUPPORT

This would facilitate online interventional patient treatment

POST-TREATMENT SUPPORT

AI tools would facilitate time-consuming tasks such as reporting, coding and billing. Note that, in some cases, reporting can take up to 50% of the radiologists' work time. AI tools in this domain would provide much needed relief freeing radiologists to do their core tasks. At the same time, the quality of the reporting would be improved by enabling clear delineation of the various lesions and automatic monitoring of the KPIs. This in turn would open the gate for value-based compensation with concomitant benefits to both service providers and payers.

FOLLOW UP

That would facilitate monitoring of the subsequent treatment stages.

QUALITY ASSURANCE

That would constantly monitor all the above stages including both human practitioners and AI components. QA will drive up efficiency by providing timely feedback to all the stakeholders. It will also prove the value of the overall system for both regulators and payers.

Ideally, AI based systems will address all the above issues for the widest possible gamut of clinical domains. However, for the near future we'll have to live with a limited subset of the above. AI solutions will address more focused tasks such as – a tool for reducing a turnaround time in the ED, or tool for reducing outliers in the outpatient population. Accordingly, both users and developers will have to identify the gaps where AI would provide the biggest value and largest impact in terms of the actual clinical outcomes.

Once we have crossed into the realm of the end to end solutions, our perspective must change. It is no longer enough to talk about such esoteric parameters as sensitivity and specificity. Instead we must focus on specific clinical outcomes such as life longevity, quality of life, length of stay etc.

Consider the case of lung lesion detection. Assume that we create an effective AI system with high AUC (.99). We compare this system to the trained radiologists and notice that for the given specificity, AI performs much better than physicians (90% for the AI system compared to only 45% for the physicians). It looks too good to be true, but is it? In the clinical setting we notice that radiologists seldom miss any large clinically significant lesions (say above 5mm). Hence, improved AI sensitivity comes from small lesions (less than 3mm). So, in reality, our AI system does more harm (by over diagnosis) than good. Is it a proverbial black eye to the AI systems? Not necessarily, because if we manage to add an additional AI application providing automatic follow up after small lesions, we may have an effective solution for early lung cancer detection. The important thing is to focus on the clinical outcome (such as patients' survival rate).

Case In Point

Aidoc result of interval nodule enlargement. Aidoc algorithms detect nodule on prior and current exam, and detect change with volumetric measurement, increasing conspicuity of the findings and streamlining the workflow

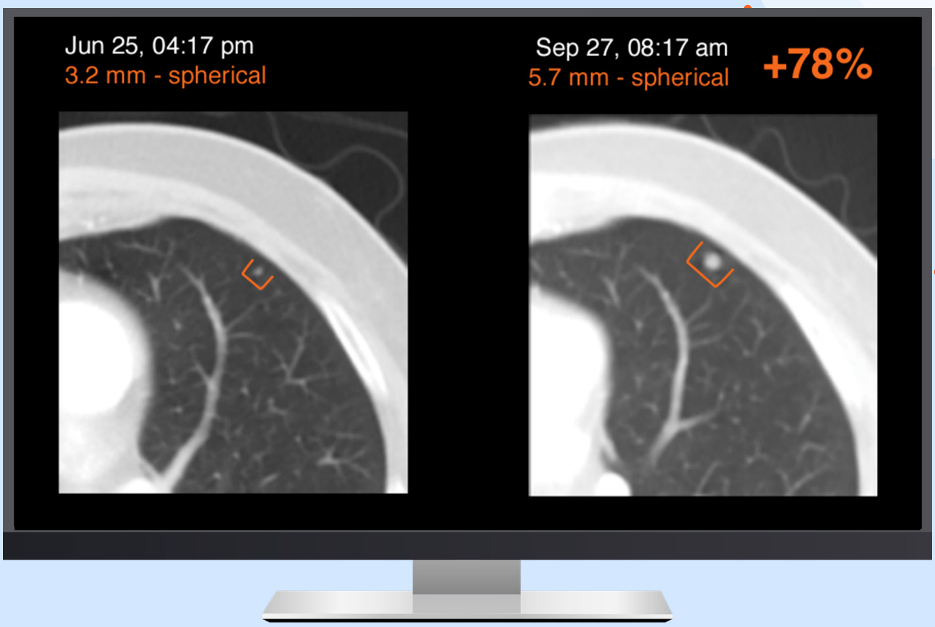

Integrated development process

It is tempting to view the above stages as a single logic flow. We start with algorithmic ideas that evolve into the products. From a combination of these products we create specific clinical solutions. Unfortunately, real life tends to be too complex for such a simplified vision. As a result, frequently, no one can foresee all the facets of the optimal solution. In such situations it is preferable to adopt a so-called agile approach. There development is done in iterations. We begin by quickly creating a first approximation of the solution. Then we experiment with it in the field. In parallel we gather additional training data, perform an evaluation of the clinical outcomes and explore ideas regarding the possible additional features. Based on these inputs, we create a second system iteration and repeat the process. In such an environment the radiologist ceases to be a mere system user that either likes or dislikes the tool at his/her disposal. Instead, the radiologist becomes a partner that guides the development process.

A Future Outlook

Next Gen AI in clinical settings holds much promise, and its application will be made possible by innovative researchers, clinicians and developers working together towards improving patient care. Aidoc continues to develop its AI product with the goal of enhancing both clinical outcomes, as well as facilitating the radiologist's workflow. However, we can't do it without your input. For any questions, input or feedback on the topic of Next Gen AI, please feel free to reach out to us through our dedicated communication channels.

ABOUT AIDOC

Aidoc is a pioneering force in clinical AI. We focus on aiding and empowering healthcare teams to optimize patient treatment, which results in improved economic value and clinical outcomes.

Since 2016, Aidoc's clinically proven AI solutions have eliminated silos, increased efficiencies and improved outcomes by delivering critical information when and where care teams need it – leading to immediate collective action.

Powered by Aidoc's exclusive aiOS™, we analyze and aggregate medical data to enable care teams to operationalize the unexpected and work seamlessly with a continued focus on the patient.

Aidoc AI is always on, running in the background to change the foreground.

Let's Connect

info@aidoc.com

www.aidoc.com

twitter.com/aidocmed

facebook.com/aidocmed

linkedin.com/company/aidoc