Enhanced Point-of-Care Ultrasound Applications by Integrating Automated Feature-Learning Systems Using Deep Learning

Hamid Shokoohi, MD, MPH, RDMS, RDCS , Maxine A. LeSaux, BS, Yusuf H. Roohani, MS, Andrew Liteplo, MD , Calvin Huang, MD, MPH, Michael Blaivas, MD, MBA

Recent applications of artificial intelligence (AI) and deep learning (DL) in health care include enhanced diagnostic imaging modalities to support clinical decisions and improve patients' outcomes. Focused on using automated DL-based systems to improve point-of-care ultrasound (POCUS), we look at DL-based automation as a key field in expanding and improving POCUS applications in various clinical settings. A promising additional value would be the ability to automate training model selections for teaching POCUS to medical trainees and novice sonologists. The diversity of POCUS applications and ultrasound equipment, each requiring specialized AI models and domain expertise, limits the use of DL as a generic solution. In this article, we highlight the most advanced potential applications of AI in POCUS tailored to high-yield models in automated image interpretations, with the premise of improving the accuracy and efficacy of POCUS scans.

Key Words—artificial intelligence; deep learning; machine learning; point-of-care ultrasound

Received August 4, 2018, from the Department of Emergency Medicine, Massachusetts General Hospital, Harvard Medical School (H.S., A.L., C.H.), Boston, Massachusetts, USA; Department of Emergency Medicine, (George Washington University School of Medicine and Health Sciences, Washington, DC, USA(M.A.L.); Platform Technology and Science, GlaxoSmithKline, Cambridge, Massachusetts, USA(Y.H.R.); Department of Emergency Medicine, University of South Carolina School of Medicine, Columbia, South Carolina, USA(M.B.); St Francis Hospital, Columbus, Georgia, USA(M.B.). Manuscript accepted for publication September 30, 2018.

Address correspondence to Hamid Shokoohi, MD, MPH, Department of Emergency Medicine, Massachusetts General Hospital, 326 Cambridge St, Suite 410, Boston, MA 02114 USA.

E-mail: hshokoohi@mgh.harvard.edu

Abbreviations
AI, artificial intelligence; CNN, convolutional neural network; CT, computed tomography; DL, deep learning; ML, machine learning; POCUS, point-of-care ultrasound; US, ultrasound

doi:10.1002/jum.14860

Deep learning (DL) is a form of machine learning (ML), which is the science of training computers to perform tasks not by being explicitly programmed but, rather, through enabling them to study patterns within data (Table 1).[1] The concept of ML dates back to the 1950s; however, it is only with the recent dramatic increase in the availability of both data and computing power that it has become possible to implement ideas such as DL.[2]

In the past several years, DL has become a burgeoning technology in the field of medical imaging.[3] This technology can boost multiple functions in an objective manner, which can assist in both image diagnostics and image enhancement. These models have well-documented use in imaging modalities such as computed tomography (CT) and radiography.[4–7]

Point-of-care ultrasound (POCUS) imaging is an optimal application for DL techniques because it encompasses a wide variety of applications and a diverse group of users with a substantial disparity of training. The binary nature of most POCUS studies removes the burden of evaluating for all possible pathologic conditions imaged by the ultrasound (US) machine and allows focus on

Table 1. List of Common Terms and Definitions in ML

Term	Definition
Algorithm	A set of rules to be followed to perform a specific task
Machine learning	A field of computer science that deals with teaching computers to perform tasks by giving them the ability to study patterns in data, without being explicitly programmed
Deep learning	A type of ML that learns on its own how best to represent data as a hierarchy of concepts, with each concept defined through its relation to simpler concepts[1]
Artificial intelligence	The theory and development of computer systems to perform tasks that are associated with human intelligence
Features	Each input piece of data or example fed into an ML model can be represented as a collection of 1 or more features
Classification	Assigning an input example to 1 or more predefined classes
Model architectures	The precise description of the computational operations that comprise an ML model and how they are connected to one another, effectively describing the flow of information through it
Neural network	A network composed of nodes or “neurons” that each perform a computational operation and through which information flows by means of weighted interconnections; to learn to perform a specific task, these weights can be tuned
Convolutional neural network	A special kind of neural network that uses a mathematical operation called convolution that is particularly suited for the type of patterns normally found in imaging data
Model training	Tuning of model parameters through repeatedly passing training data through a model and minimizing errors as measured by an objective function
Inference	Passing data through a trained model to obtain a valuable output
Model parameter	A parameter whose value is learned during training
Model hyperparameters	A parameter whose value is set before the model begins training
Learning rate	The rate at which the optimization function guiding the training process progresses at each step of training
Validation set	Data that are not directly trained on but are used for assessing the training progress
Test set	Data that are not used at all during the training process and are only used to test performance of a trained model

a limited number of diagnoses. Once trained, the model takes in an image as input and can classify pathologic findings as 0 or 1 with a measure of confidence.[8] In particular, this solution can be fully automated and is able to process images quickly for accurate and objective detection of life-threatening and time-sensitive conditions such as pneumothorax, hemothorax, cardiac standstill, pericardial effusion, tamponade, and abdominal free fluid. Accurate identification of these pathologic findings is important in management of trauma patients, informs numerous complex management decisions, and, as such, is an attractive subject for DL techniques.

The goal of this article is to highlight the use of DL in POCUS imaging and to provide an overview of the practical use of DL for purposes of POCUS applications. The diversity of POCUS applications and protocols, each requiring specialized AI models, algorithms, and software technology, limits the use of DL as a general-purpose solution for all applications. Therefore, in this article, we highlight the most advanced potential applications of AI in POCUS tailored to high-yield models in automated image enhancement and interpretations of POCUS scans. We will also discuss the novel applications of DL in POCUS practice such as its use in disaster response, prehospital care, global health, and medical education.

Deep-Learning Technology

Although artificial intelligence (AI) is the broader concept of the ability of computers or machines to perform intelligent tasks or functions, DL is an example of a narrower application of AI in which machines are fed data and are then capable of learning from that data largely by themselves (Table 1). Artificial-intelligence systems generally use ML algorithms; these consist of 2 parts. The first part, feature extraction, is the conversion of raw data into a suitable internal representation (ie, features). A learning subsystem can then use this representation to perform a second task, such as pattern classification.[9] For most of its history, ML required careful hand engineering of features and domain expertise to design feature extractors.[9] Deep learning has been able to substantially improve the accuracy of these systems because it creates an efficient mechanism to learn both optimal feature representations as well as an appropriate classifier for a given task. It does so by automatically learning a hierarchical representation of high-dimensional complex data, in which each learned concept is defined by its relationship with simpler concepts.[1] However, the trade-off for this level of automation is the requirement for much larger training data sets and greater computational power, challenges that have recently become easier to overcome.

At the implementation level, DL covers a wide class of model architectures, most of which fall under the category of a neural network. Neural networks are so named because their functionality loosely resembles a network of neurons. In these networks, each neuron receives an input, applies an activation function, and forwards it along to the neurons in the next layer of the network on the basis of certain connection weights that are learned by the system. The final output of this network is usually either a class label or a score.

There are several neural network architectures that use a DL approach. One popular class of examples is that of convolutional neural networks (CNNs).[10] They use a special kind of mathematical operation called convolution that is particularly suited for the type of patterns that are commonly present in imaging data. By using multiple layers of neurons, they are able to capture complex patterns within an image in a hierarchical fashion.[1] This allows CNNs to efficiently detect elaborate yet distinctive features in an image that are difficult to capture by using a simpler approach.

The application of CNNs in the field of image and video processing has increased considerably in the past 6 years. This has been facilitated by easily accessible, free, open-source (software code that is publicly available) software packages such as TensorFlow (TensorFlow 1.4; Google, LLC, Mountain View, CA), Cognitive Toolkit (Microsoft Corporation, Redmond, WA), PyTorch, Caffe (University of California, Berkeley, CA), and MXNet (Apache Software Foundation, Forest Hill, MD) for rapidly implementing and fine-tuning DL models.[6,11,12] Within medical imaging, a large variety of use cases have been reported from almost every aspect of medical image analysis, including detection of pleural effusion and cardiomegaly on chest radiography, mediastinal lymph nodes on CT, lung nodules on CT, and detection of tuberculosis on chest radiographs.[13–17]

Building a DL Model

Each node in a DL model (such as a CNN) represents a mathematical operation, and the connections represent the strength of the interaction between the nodes. When an image is sent through the network, it undergoes each of these operations, with the end result (such as a classification) being produced at the output.[9]

Deep-learning models need to be trained before they can be used for inference or prediction. The training of the network refers to the tuning of the weights that describe the strength of these connections. To achieve this, the training data are passed through the network in batches, and the weights are slowly optimized on the basis of the error produced at the output. Model hyperparameters that control the training process such as the parameter learning rate, initialization conditions, etc are manually set (or iterated through) to tune the training process. It is also common to cycle through different network architectures during the training process, although the search is normally limited to known well-functioning architectures (eg, VGGNet [Visual Geometry Group, University of Oxford, Oxford, England], Inception [Google, LLC], and ResNet-50 [Google, LLC]).

Usually around 20% of the training set is held out as a validation set to assess the performance of the model during training. Alternatively, a series of models could be trained while leaving out different subsets of the main training set for validation each time (cross-validation). The model that performs best would be selected for inference. Each passage of the full training set through the network is referred to as an epoch. Deep-learning models can often take more than 100 epochs to train.

Once training has been completed satisfactorily, the model can be deployed to the POCUS device or the cloud for prediction or inference. Under the inference mode, the weights are not tuned, and each piece of data (eg, image) is only passed through the model once to output the result for that data. Before deploying the model, it is required to test the performance against an independent test set. If found to be

Figure 1. Process flowchart describing the training and prediction/inference process for a DL model.

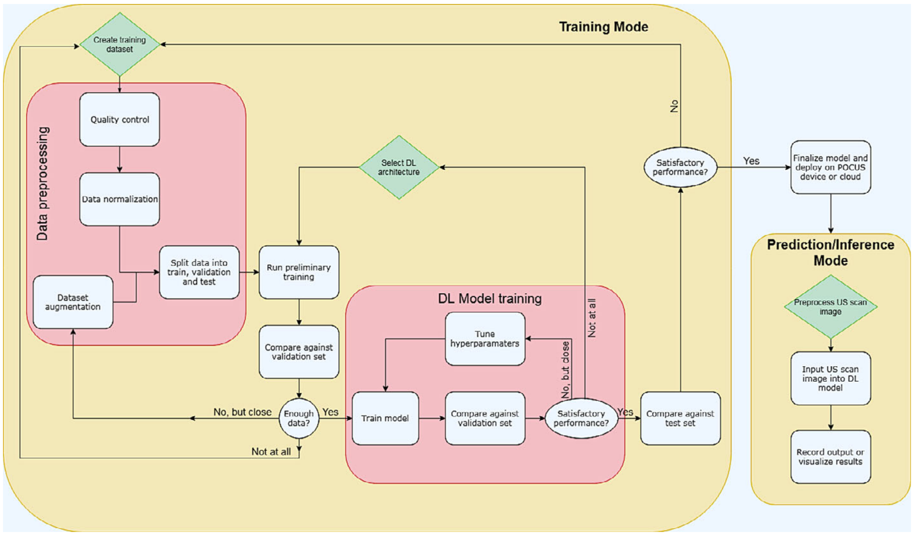

unsatisfactory, correction should not be attempted with hyperparameter tuning; instead, the whole training process should be repeated from scratch. The entire training process has been illustrated as a process flowchart in Figure 1.

Deep-Learning Models' Output for POCUS

The applicability of DL to POCUS can be quite diverse, and the output produced by these models is largely dependent on the problem that they are trained to solve. One common application is to train a DL model for detection or segmentation. In this case, a training set is put together in which the structures or regions of interest have been identified and manually delineated. The model is then trained to learn from these annotations and reproduce the same for new images. For example, Chen et al[18] used DL for segmentation of 5 different structures in US scans. Some groups have also used DL for image enhancement or to remove unnecessary aspects of an image. This is seen in chest radiographs, where a DL model can subtract the bony impression of the ribs from the image, and then analyze the remaining data: the appearance of the lungs.[19]

Alternatively, a DL model can also be trained "end to end" for classification without distinct intermediary steps. Becker et al[20] used DL to classify breast cancer on US imaging. This approach does not need finer annotation as in the previous case and could work with coarser image level labels. There are several examples of DL being used for frame labeling in US images.[20–22]

In either case, the POCUS operator would only need to provide an image, and the trained DL model would be able to immediately return the desired output, whether the outline of an organ, an enhanced US image, or the classification of the US image along with a confidence score. Thus, DL tools can be used to augment the expertise of a trained diagnostic physician by automating certain repetitive elements of the diagnosis and enabling physicians to focus on the more challenging aspects.

Deep-Learning Applications in POCUS Practices

Point-of-care US is an extremely useful diagnostic modality in emergency medicine.[23] Over time, US technology has improved to become a high-quality, rapid, and safe tool that can be used to assess diverse medical conditions. Deep-learning models have emerged as innovative technologies of choice in several POCUS applications. The premise behind DL integration is that the accuracy and efficacy of POCUS imaging can be substantially improved by automated image interpretation and by matching various algorithms for a particular clinical scenario. These models require sets of thousands of images and a partnership with a computer engineer to help design and build the model. Specifically, when building these models, it is necessary to use data from a large database of US images that include both normal and pathologic findings for the targeted conditions. The DL functionalities in US include both image enhancement and diagnostics. In the category of image enhancement, there has been image quality improvement in breast US to help a radiologist better identify breast lesions and in carotid US to improve visualization of arterial wall layers.[24,25]

The functionality of image diagnostics has been used for the evaluation of thyroid nodules, malignant breast tumors, cardiovascular conditions, prostate cancer, and myositis.[14–16,26,27] The potential applications of DL for cardiopulmonary POCUS examinations and examples are summarized in Tables 2 and 3.

Disaster Response

Adopting new strategies and using existing technology to diagnose and treat patients in mass casualty incidents will save lives and limit morbidity. Point-of-care US has become a powerful tool for health care providers to use in disaster response medicine.[28,29]

Ultrasound machines are becoming increasingly available to emergency care providers and can be critically important during a mass casualty incident, when access to other imaging modalities is limited by patient volume, time, and resources. In the setting of disaster response, POCUS has been used to diagnose bony fractures, abdominal free fluid, and pneumothorax.[29,30] Typically, these images are either interpreted on scene by a qualified sonologist or transmitted to a medical institution for interpretation.[31]

Diagnostic US techniques have been established for the detection of life-threatening cardiothoracic and abdominal injuries in individuals. Integrating DL models based on individual cases and applying them

Table 2. Applications and examples of Deep Learning in Cardiac POCUS

Detection and Prediction Automation	A. Estimation of cardiac ejection fraction with associated B-lines B. IVC caliber and collapsibility in predicting fluid responsiveness C. Assessment for cardiac tamponade in patients with pericardial effusion D. Prediction of the total volume of a pleural effusion from a 2D sample E. Detection of cardiac standstill and survival odds
Intelligence Augmentation	A. A deep 3D residual CNN for reduction of false-positive scans in assessment of regional wall motion abnormalities B. AI interpretation and guidance of TEE in the field for EMS C. DL of the spatial characteristics of wall motion abnormality and tissue doppler for risk stratification of acute coronary syndromes D. Extraction of left ventricular ejection fraction from various types of cardiac scans using DL
Automated Image Segmentation, Measurement, Labelling	A. Automated labeling and annotation of cardiac images for students' self-directed learning B. Automated determination of cardiac pacemaker capturing from cardiac ultrasound C. Automated mining of large-scale valvular lesion annotations and global lesion detection with DL
Improving Decision Support System	A. Automated prediction of fluid responsiveness using dynamic preload indices RUSH protocol-based images in patients with shock B. AI-enhanced prediction algorithm for trans-esophageal echocardiography in patients with cardiac or respiratory arrest to guide resuscitation C. Automatic determination of fine ventricular fibrillation to identify responders to cardiac defibrillation
Assessment of Image Quality	A. Instant quality assessment (QA) of an image before transmitting for a manual QA process B. Recognition of suboptimal images in real-time in predicting accuracy of the diagnosis C. Automation of QA of cardiac US using CNNs for credentialing residents and faculty
Data Mining for Research	A. Development of an image search engine which could permit searches using images directly as an input B. Creation of large databases of different patients with similar pathology such as pulmonary infarction

Table 3. Applications and examples of Deep Learning in Pulmonary POCUS

Detection and Prediction Automation	A. AI-enhanced lung ultrasound in discriminating viral and bacterial pneumonia B. Semi-supervised DL analysis of lung scans and incorporation of the BLUE algorithm C. AI-enhanced classification of lung consolidation in differentiating atelectasis from pneumonia D. Risk stratification of patients with COPD based on DL E. Detection of pneumothorax in the field by EMS providers
Intelligence Augmentation	A. Estimating the size of a pneumothorax based on the location of a lung point B. AI assessment of lung US artifacts in the diagnosis and classification of COPD C. RV assessment for the screening of pulmonary embolism D. DL-based estimation of the size and volume of pleural effusions E. DL for biomarker regression: application to BNP on lung US
Automated Image Segmentation, Labelling	A. Holistic segmentation of the lung ultrasound from cine images B. Real-time scan assistance for cardiac and lung ultrasound C. Augmenting electronic POCUS teaching files to facilitate comparison of an obtained images with normal or pathological findings
Improving Decision Support System	A. Prediction of antibiotic response from US lung images based on DL techniques B. AI-enhanced lung assessment in identifying active cases of tuberculosis C. Prediction of the best insertion site for a thoracentesis D. Comparing US images in response to therapeutic interventions such as lung pattern before and after positive pressure ventilation

in the mass casualty situation can improve triage and more effectively allocate resources.

Achieving true "real-time" image interpretation would be a valuable addition to the current practice of POCUS during disaster response. It will allow providers with limited training to use the technology in their response plan and will expand POCUS in this setting far beyond the ranks of physicians, who are often not the first responders in a disaster. This integration could improve image quality in small handheld devices, decrease the time needed for interpretation, and improve accuracy in diagnostics. These features could provide benefits in the chaotic and time-sensitive nature of a disaster response, making it easier for providers on scene to properly triage and care for their patients.

Prehospital Emergency Care

One of the most promising applications of real-time image interpretation is real-time detection of pathologic conditions, which can be used to expand and strengthen prehospital POCUS applications. For example, detection of free fluid in the chest and abdomen in trauma patients and accurate detection of cardiac standstill in patients with cardiac arrest by any responder at the scene could help improve patient care.

In some locations, prehospital POCUS is currently used during the transport of patients to emergency medicine care settings.[32,33] Ultrasound can be used in the prehospital setting to assess the severity of illness or injury during transport. This has implications on triage and resource allocation and can provide the emergency department with valuable information before patient arrival.[34] To date, US has been used in both ground ambulance and aeromedical transport of patients to a hospital.[35] Flight crews, emergency medical technicians, and paramedics are usually trained to use US through a curriculum that combines didactics and hands-on use of the technology.[36] In the aeromedical setting, in which space and noise constrain and limit physical examination and auscultation, US can be extremely useful for identifying conditions such as pneumothorax.[37] Coupled with the fact that the use of POCUS has no negative effect on flight time, this tool can be very helpful for providing information to the trauma team to better prepare for a trauma patient's arrival. Despite the benefits, the use of POCUS in the prehospital setting is not yet widespread.[38] These low use rates have been attributed mostly to equipment and training costs. In addition, there is a lack of widespread evidence for POCUS efficacy in the prehospital setting, which prevents uptake of US into practice.[38]

The use of DL models in providing automated interpretation of images in prehospital POCUS applications can help overcome challenges associated with training. If a DL model could be applied to create a diagnostic and treatment protocol using a sonologist's acquired US images, it could allow even a novice US user to make effective care management decisions using POCUS. This has obvious advantages for improving the objectivity of the analysis and reducing interexaminer variability. In practice, DL models could also improve diagnostic accuracy to help enhance decision making for appropriate transport and triage.

Medical Education

Ultrasound is increasingly being used to enhance medical student education throughout the world. Its effectiveness as a tool has been demonstrated in teaching medical student anatomy and diagnostic skills.[39–41] Medical students at varying levels of training have demonstrated proficiency with the machine and rapidly acquired skills with as little as 1 week of US training.[40] With the increase of POCUS in clinical settings, it has become essential for medical students and resident physicians to have a strong foundation in this imaging modality.

Systematic US education both requires and encourages heavy faculty participation. This requirement can be challenging, as many departments have only recently embraced POCUS and have faculty who are US learners themselves. By converging image interpretations and clinical analytics, images no longer need to undergo a quality assurance process to deliver insights. This gives instructors and administrators more time to concentrate efforts on teaching strategies, clinical integration, and decision-making processes.

Autonomous POCUS interpretation will substantially enhance medical education. The use of DL models in US for medical education has the capacity to help medical students learn more actively. A DL model that can enhance image quality and remove background "noise" could help medical students train their eyes to look for the specific features on which the model focuses. In addition, with computer-aided diagnostics, the DL model could assist medical students to better discern pathologic findings in imaging. When teaching medical students with DL models, however, it will still be extremely important to teach proper scanning techniques, including probe positioning and movements, so that the model will be able to accurately analyze the images.

The distinguishing feature of these automated learning systems is that the images obtained by trainees would be automatically labeled, and when a pathologic finding is encountered, immediate feedback could be given to the student, enhancing learning. In certain conditions, prognostic information on pathologic findings in prediction of mortality and morbidity could be determined. For example, a regression model might output a prediction or score representing the posttest probability of respiratory failure and the need for intubation and mechanical ventilation after lung imaging and cardiac scans.

Global Health

The use of POCUS from a global health perspective is advantageous as US becomes more portable and inexpensive.[42] The ability to provide US imaging in resource-limited environments can provide a unique perspective to medical providers and vastly improve their access to critical information for the medical decision-making process.[43] Ultrasound in developing countries has proved to be effective in obstetrics, trauma, cardiac and surgical emergencies, and procedural guidance,[42] coupled with the fact that medical professionals have been successful in implementing programs in developing countries to train providers with limited medical training to conduct US scans.[44–46] Perhaps with the introduction of DL models, these capabilities can be expanded even further. The training of nonmedical professionals can be bolstered with models that provide accurate interpretations of scans with minimum training needed. Additionally, automated detection of certain conditions such as pulmonary and extrapulmonary tuberculosis (periaortic lymph nodes and splenic and hepatic lesions) with POCUS may facilitate screening and evaluation efforts in tuberculosis-prevalent areas with limited access to sonologists. There are similar benefits in detecting valvular lesions in cases with suspected rheumatic heart diseases. There have been efforts in other medical imaging domains as well. For example, Zhu et al[47] looked at ways of using DL to improve imaging quality and speed for scans with a low signal-to-noise ratio. These models could also learn reconstruction without using domain expert knowledge.

Consumer Use

Although quite forward looking, it is likely that automation of US diagnostics will lead to consumer use of this technology. As the level of medical training required to perform an evaluation is driven to a minimum, consumer use is the next logical step. Much like glucometers and urine and streptococcal throat tests before them, diagnostic US devices will likely find their way to virtual or actual store shelves for consumer purchase. The price point will be a critical linchpin to enable such a transition as well as nearly full automation, including operator guidance. Both are already within sight if one only observes current technology ready to go on sale in the market and future plans announced by several US companies.

An AI-driven, highly portable, and inexpensive device that plugs into a consumer's smartphone or tablet would allow diagnosis of some conditions at home, without the need to visit a physician's office or emergency department. Additionally, and likely an intermediate step, is the use of such devices for home monitoring of high-risk patients such as those with congestive heart failure or emphysema, who are frequent users of medical facilities and would benefit from early interventions in response to data obtained at home.

It is worth mentioning that the introduction of DL models in the context of POCUS will be a disruptive innovation to current standard processes. Therefore, it is prudent to consider the effect on clinicians who will interact with this technology. From a user standpoint, the implemented DL technology must be easy to use and practically fit into the physician's normal daily work flow to maximize its use. This may include requirements for how to document the use of DL in the clinical decision-making portion of a patient chart, privacy considerations related to the centralized data set, and even reimbursement or billing rules.

Current Market Landscape

The current technological capabilities and available resources are very favorable for implementation of DL models in US practice. Critically important is that fact that US use by clinicians has expanded exponentially, making the tool widely available. As far as available revenue, a POCUS machine can be purchased for as little as $25,000, and a high-end machine may sell for up to $115,000.[48] Furthermore, the field of US has been further primed for expansion by the availability of handheld US machines. Devices combining handheld use and DL models have already been introduced into the market. For example, the Butterfly iQ (Butterfly Network, Inc, Guilford, CT), a small US device that can display US images and videos on an iPhone,[49] and the Venue US system (GE Healthcare, Chicago, IL) have integrated DL models for certain applications that allow novice users with minimal training to make interpretations for some cardiac, musculoskeletal, and obstetric applications. This ubiquity of technology has allowed their use to outweigh the training expenditure. Whereas training of a user to effectively master the skills of US can take many years of didactic and clinical application, a DL model would allow for medical professionals to quickly hone their skills while effectively using POCUS. In addition, teaching basic skills of scan acquisition does not necessarily make trainees competent in image interpretation. If a DL model is able to provide even a limited interpretation, it may increase the use and integration of US into the medical decision-making process.

From a logistic perspective, this endeavor could be implemented from preexisting infrastructure already in place. For example, there are many open-source technologies such as Microsoft's Cognitive Toolkit and Custom Vision, Google's TensorFlow, and PyTorch.[11,50–52] These technologies make DL applications with image recognition feasible in many different capacities and to users with a limited skill set in the field.

Finally, these models require large sets of images for training. Fortunately, there are multiple sources of mass data in the field. For example, QPath, one such model, allows US users to store their images in a central database, making retrospective queries and quality assurance easy.[53] These large data sources would be the perfect foundation to help train an algorithm.

How to Get Access to a POCUS DL Model

There are 3 ways to get access to DL models specific to POCUS: (1) buy a custom model developed by a vendor (expensive and not flexible); (2) use an existing model, and customize/refine it with a new training set (such as Custom Vision/Object Detection API provided by Microsoft or other vendors; less expensive but requires some work to retrain the existing model with tagged images); and (3) develop a custom model from scratch (expensive and requires specialized data scientists, data wranglers, and software developers to operationalize/integrate model into the process).

In our opinion, developing or refining a model is preferable but requires specific skills that are mastered only by few data scientists and professionals. Using the open-source ecosystem is not trivial, as it requires specific knowledge of DL and the Python programming language (Python 3.6; Python Software Foundation, Wilmington, DE), which seems to be the prevalent programming language in DL.

Discussion

Deep learning will affect all of medicine in the coming years, but imaging seems particularly ripe for revolution by DL technology. Many applications are well defined, and large amounts of imaging data already exist in radiology servers and imaging data banks around the world. Not surprisingly, body CT and magnetic resonance imaging applications were some of the first explored and already have US Food and Drug Administration–approved algorithms that aid in diagnosis. Algorithms that make mammography highly sensitive and specific instead of a diagnostic coin toss have been put into practice, and more is to come.

Although in the traditional imaging realms such as CT and magnetic resonance imaging, DL-produced imaging algorithms herald improved diagnostic accuracy and decreased costs, in POCUS, DL holds the key to widespread use at the lowest provider training level. Some POCUS enthusiasts view the current progress in US technology as a pathway toward the fabled “Tricorder” of Star Trek movies, and the most important aspect of that progress will be automation. As suggested already, increased automation to the level at which the machine can direct a completely novice user to acquire images and then interpret those images (and possibly provide a diagnosis) will democratize this imaging technology to the ultimate point: that is, where anyone could use a device, including the worried parent who scans his or her child with a newly purchased smart device from a local drug store to realize the child’s cough is from pneumonia rather than a simple cold.

This path toward increasing the use of AI algorithms in our medical tools will likely be turbulent, with resistance from multiple stakeholders who have the most to lose from changes in the status quo. For others, who can suddenly have a safe and highly accurate diagnostic imaging tool in their hands to better serve their patients, it will be a positive revolution. It is unlikely that AI will replace medical practitioners; however, it is likely that medical practitioners who use AI will replace those who do not. The ultimate benefactors are those we seem to forget first when different silos of medicine battle each other for turf and money: our patients, whether they are in midtown Manhattan or in sub-Saharan Africa. It is important for clinicians to understand how DL is likely to change their practices and how powerful tools such as smart handheld US devices are developed and ultimately find their way to practitioners’ coat pockets.

In conclusion, the use of DL models has begun to transform the capabilities of medical imaging. The potential applications in the fields of disaster response, prehospital care, global health, and medical education are promising. Deep-learning POCUS models provide the potential for automated feature-learning systems in streamlining life-threatening and time-sensitive diagnoses in a cost-effective and time-efficient manner. Furthering the implementation of such algorithms on a global basis could drastically expand POCUS applications and the use of POCUS by less experienced providers. The ultimate effect will be to improve clinical care and develop novel diagnostic work flows for screening, diagnosis, and referral.

Although DL-enhanced POCUS applications can help automate the making of diagnoses, we cannot forget that it is ultimately providers’ interpretations that should guide further workups and make clinical decisions. In fact, the emergence of these applications will only improve real-time image interpretations, enabling faster and more sophisticated clinical decisions.

References

1. Goodfellow I, Bengio Y, Courville A. *Deep Learning*. Cambridge, MA: MIT Press; 2016.
2. Samuel AL. Some studies in machine learning using the game of checkers. *IBM J Res Dev* 1959; 3:210–229.
3. Suzuki K. Overview of deep learning in medical imaging. *Radiol Phys Technol* 2017; 10:257–273.
4. Sharma K, Rupprecht C, Caroli A, et al. Automatic segmentation of kidneys using deep learning for total kidney volume quantification in autosomal dominant polycystic kidney disease. *Sci Rep* 2017; 7:2049.
5. González G, Ash SY, Vegas Sánchez-Ferrero G, et al. Disease staging and prognosis in smokers using deep learning in chest computed tomography. *Am J Respir Crit Care Med* 2018; 197:193–203.
6. Ziemlewicz TJ, Kim DH, Hinshaw JL, Lubner MG, Robbins JB, Pickhardt PJ. Computer-aided detection of colorectal polyps at CT colonography: prospective clinical performance and third-party reimbursement. *AJR Am J Roentgenol* 2017; 208:1244–1248.
7. Loog M, van Ginneken B. Segmentation of the posterior ribs in chest radiographs using iterated contextual pixel classification. *IEEE Trans Med Imaging* 2006; 25:602–611.
8. Kohli MD, Summers RM, Geis JR. Medical image data and datasets in the era of machine learning: white paper from the 2016 C-MIMI meeting dataset session. *J Digit Imaging* 2017; 30:392–399.
9. LeCun YA, Bengio Y, Hinton GE. Deep learning. *Nature* 2015; 521:436–444.
10. Shin HC, Roth HR, Gao M, et al. Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. *IEEE Trans Med Imaging* 2016; 35:1285–1298.
11. Paszke A, Chanan G, Lin Z, et al. Automatic differentiation in PyTorch. *Adv Neural Inf Process Syst* 2017; 30:1–4.
12. Jia Y, Shelhamer E, Donahue J, et al. Caffe: convolutional architecture for fast feature embedding. In: *Proceedings of the 22nd ACM International Conference on Multimedia*. New York, NY: Association for Computing Machinery; 2014:675–678.
13. Litjens G, Kooi T, Bejnordi BE, et al. A survey on deep learning in medical image analysis. *Med Image Anal* 2017; 42:60–88.
14. Henglin M, Stein G, Hushcha PV, Snoek J, Wiltschko AB, Cheng S. Machine learning approaches in cardiovascular imaging. *Circ Cardiovasc Imaging* 2017; 10:e005614.
15. Ella Hassanien A, Al-Qaheri H, Snášel V, Peters JF. Machine learning techniques for prostate ultrasound image diagnosis. In: *Advances in Machine Learning 1*. Berlin, Germany: Springer; 2010:385–403.
16. Burlina P, Billings S, Joshi N, Albayda J. Automated diagnosis of myositis from muscle ultrasound: exploring the use of machine learning and deep learning methods. *PLoS One* 2017; 12:e0184059.
17. Lakhani P, Gray D, Pett C, Nagy P, Shih G. Hello world deep learning in medical imaging. *J Digit Imaging* 2018; 31:283–289.
18. Chen H, Zheng Y, Park JH, Heng PA, Zhou S. Iterative Multi-domain regularized deep learning for anatomical structure detection and segmentation from ultrasound images. In: *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*. Berlin, Germany: Springer; 2016: 487–495.
19. Novikov AA, Lenis D, Major D, Hladuvka J, Wimmer M, Buhler K. Fully convolutional architectures for multi-class segmentation in chest radiographs. *IEEE Trans Med Imaging* 2018; 37:1–9.
20. Becker A, Mueller M, Stoffel E, Marcon M, Ghafoor S, Boss A. Classification of breast cancer in ultrasound imaging using a generic deep learning analysis software: a pilot study. *Br J Radiol* 2018; 91:20170576.
21. Chen H, Ni D, Qin J, et al. Standard plane localization in fetal ultrasound via domain transferred deep neural networks. *IEEE J Biomed Health Inform* 2015; 19:1627–1636.
22. Baumgartner CF, Kamnitsas K, Matthew J, Smith S, Kainz B, Rueckert D. Real-time standard scan plane detection and localisation in fetal ultrasound using fully convolutional neural networks. In: *Lecture Notes in Computer Science*. Berlin, Germany: Springer; 2016:203–211.
23. Whitson MR, Mayo PH. Ultrasonography in the emergency department. *Crit Care* 2016; 20:227.
24. Karimi B, Krzyżak A. A novel technique for detecting suspicious lesions in breast ultrasound images. *Concurr Comput Pract Exp* 2016; 28:2237–2260.
25. Menchón-Lara RM, Sancho-Gómez JL, Sánchez-Morales A, et al. Using machine learning techniques for the automatic detection of arterial wall layers in carotid ultrasounds. In: *Ambient Intelligence: Software and Applications*. Cham, Switzerland: Springer; 2015:193–201.
26. Lu-Cheng Z, Yun-Liang Y, Wen-Hua L, et al. A model to discriminate malignant from benign thyroid nodules using artificial neural network. *PLoS One* 2013; 8:e82211.
27. Han S, Kang HK, Jeong JY, et al. A deep learning framework for supporting the classification of breast lesions in ultrasound images. *Phys Med Biol* 2017; 62:7714–7728.
28. Wydo SM, Seamon MJ, Melanson SW, Thomas P, Bahner DP, Stawicki SP. Portable ultrasound in disaster triage: a focused review. *Eur J Trauma Emerg Surg* 2016; 42:151–159.
29. Ma OJ, Norvell JG, Subramanian S. Ultrasound applications in mass casualties and extreme environments. *Crit Care Med* 2007; 35(suppl):S275–S279.
30. Abu-Zidan FM. Ultrasound diagnosis of fractures in mass casualty incidents. *World J Orthop* 2017; 8:606–611.
31. Strode CA, Rubal BJ, Gerhardt RT, et al. Satellite and mobile wireless transmission of focused assessment with sonography in trauma. *Acad Emerg Med* 2003; 10:1411–1414.
32. El Sayed MJ, Zaghrini E. Prehospital emergency ultrasound: a review of current clinical applications, challenges, and future implications. *Emerg Med Int* 2013; 2013:531674.

33. O'Dochartaigh D, Douma M, MacKenzie M. Five-year retrospective review of physician and non-physician performed ultrasound in a Canadian critical care helicopter emergency medical service. *Prehosp Emerg Care* 2017; 21:24–31.
34. Ketelaars R, Hoogerwerf N, Scheffer GJ. Prehospital chest ultrasound by a Dutch helicopter emergency medical service. *J Emerg Med* 2013; 44:811–817.
35. Eadie L, Mulhern J, Regan L, et al. Remotely supported prehospital ultrasound: a feasibility study of real-time image transmission and expert guidance to aid diagnosis in remote and rural communities. *J Telemed Telecare* 2018; 24:616–622.
36. McCallum J, Vu E, Sweet D, Kanji HD. Assessment of paramedic ultrasound curricula: a systematic review. *Air Med J* 2015; 34: 360–368.
37. Yates JG, Baylous D. Aeromedical ultrasound: the evaluation of point-of-care ultrasound during helicopter transport. *Air Med J* 2017; 36:110–115.
38. Taylor J, McLaughlin K, McRae A, Lang E, Anton A. Use of prehospital ultrasound in North America: a survey of emergency medical services medical directors. *BMC Emerg Med* 2014; 14:6.
39. So S, Patel RM, Orebaugh SL. Ultrasound imaging in medical student education: impact on learning anatomy and physical diagnosis. *Anat Sci Educ* 2017; 10:176–189.
40. Shokoohi H, Boniface K, Kaviany P, Armstrong P, Calabrese K, Pourmand A. An experiential learning model facilitates learning of bedside ultrasound by preclinical medical students. *J Surg Educ* 2016; 73:208–214.
41. Mandavia DP, Aragona J, Chan L, Chan D, Henderson SO. Ultrasound training for emergency physicians: a prospective study. *Acad Emerg Med* 2000; 7:1008–1014.
42. Sippel S, Muruganandan K, Levine A, Shah S. Review article: use of ultrasound in the developing world. *Int J Emerg Med* 2011; 4:72.
43. Shah SP, Epino H, Bukhman G, et al. Impact of the introduction of ultrasound services in a limited resource setting: rural Rwanda 2008. *BMC Int Health Hum Rights* 2009; 9:4.
44. Shah S, Noble VE, Umulisa I, et al. Development of an ultrasound training curriculum in a limited resource international setting: successes and challenges of ultrasound training in rural Rwanda. *Int J Emerg Med* 2008; 1:193–196.
45. Henwood PC, Mackenzie DC, Rempell JS, et al. Intensive point-of-care ultrasound training with long-term follow-up in a cohort of Rwandan physicians. *Trop Med Int Health* 2016; 21: 1531–1538.
46. Henwood PC, Mackenzie DC, Liteplo AS, et al. Point-of-care ultrasound use, accuracy, and impact on clinical decision making in Rwanda hospitals. *J Ultrasound Med* 2017; 36:1189–1194.
47. Zhu B, Liu JZ, Cauley SF, Rosen BR, Rosen MS. Image reconstruction by domain-transform manifold learning. *Nature* 2018; 555:487–492.
48. Lee J. Ultrasound system prices up 20%, driven by new system rollout. Modern Healthcare website. http://www.modernhealthcare.com/article/20140701/NEWS/307019945. Accessed November 30, 2017.
49. Butterfly Network. Meet Butterfly iQ: whole body imaging, under $2k. Butterfly Network website. https://www.butterflynetwork.com/index.html. Accessed November 30, 2017.
50. Microsoft. CNTK features. Computational Network Toolkit website. https://cntk.ai/Features/Index.html. Accessed November 30, 2017.
51. Custom Vision. Custom Vision website. https://www.customvision.ai/. Accessed November 30, 2017.
52. TensorFlow. TensorFlow website. https://www.tensorflow.org/. Accessed November 30, 2017.
53. SonoSite, Inc. QPath by Telexy/SonoSite. SonoSite website. https://www.sonosite.com/products/qpath-telexy. Accessed November 30, 2017.