

Artificial Intelligence in Health in 2018: New Opportunities, Challenges, and Practical Implications

Gretchen Jackson^{1,2}, Jianying Hu³, Section Editors for the IMIA Yearbook Section on Artificial Intelligence in Health

¹ IBM Watson Health, Cambridge, Massachusetts, USA

² Vanderbilt University Medical Center, Nashville, Tennessee, USA

³ IBM Research, Yorktown Heights, New York, USA

Summary

Objective: To summarize significant research contributions to the field of artificial intelligence (AI) in health in 2018.

Methods: Ovid MEDLINE® and Web of Science® databases were searched to identify original research articles that were published in the English language during 2018 and presented advances in the science of AI applied in health. Queries employed Medical Subject Heading (MeSH®) terms and keywords representing AI methodologies and limited results to health applications. Section editors selected 15 best paper candidates that underwent peer review by internationally renowned domain experts. Final best papers were selected by the editorial board of the 2018 International Medical Informatics Association (IMIA) Yearbook.

Results: Database searches returned 1,480 unique publications. Best papers employed innovative AI techniques that incorporated domain knowledge or explored approaches to support distributed or federated learning. All top-ranked papers incorporated novel approaches to advance the science of AI in health and included rigorous evaluations of their methodologies.

Conclusions: Performance of state-of-the-art AI machine learning algorithms can be enhanced by approaches that employ a multidisciplinary biomedical informatics pipeline to incorporate domain knowledge and can overcome challenges such as sparse, missing, or inconsistent data. Innovative training heuristics and encryption techniques may support distributed learning with preservation of privacy.

Keywords

Artificial intelligence, machine learning, deep learning, feature engineering, federated learning

Yearb Med Inform 2019;52-5

<http://dx.doi.org/10.1055/s-0039-1677925>

Introduction

Enormous volumes of healthcare data combined with improved capabilities to store, process, and analyze such data have prompted rapid advances in Artificial Intelligence (AI) techniques in the medical domain. A recent review summarized the growing number of peer-reviewed publications addressing the development and use of AI algorithms in medicine and highlighted the significant need for additional advances and evidence for these innovations to be translated effectively into practice [1].

In 2018, AI researchers tackled some of the most significant challenges in practically applying AI in health, such as incomplete and inconsistent data and the need to access and analyze sensitive healthcare data sets across institutions.

Paper Selection Method

The recommended, standardized methodology for selection of best papers for the International Medical Informatics Association (IMIA) Yearbook [2] was adapted for the special section theme of “AI in Health: New Opportunities, Challenges, and Practical Implications.” The Ovid MEDLINE® and Web of Science® databases were queried in February 2019 for peer-reviewed, original research articles that addressed the special section theme and were published in the English language in 2018. Editorials,

letters, and opinion pieces were deliberately excluded from consideration as AI in health was a popular topic in health journals in 2018, and perspectives were common and often redundant.

Section editors generated Medical Subject Headings (MeSH®) and keyword search terms that represented state-of-the-art AI techniques and topics, and queries were refined iteratively for each database. Concepts incorporated included artificial intelligence, machine learning, deep learning, federated learning, reinforcement learning, neural network, convolutional network, variational autoencoder, deep segmentation, Bayesian model, latent variable, feature engineering, network embedding, gene interaction network, and automated pattern recognition. In PubMed, the exploded MeSH® term “Medical Informatics” was used as a domain filter. In Web of Science®, the search was limited by the Research Area of “Medical Informatics.” Article type filters were applied to retrieve research articles from journals or published peer-reviewed proceedings.

The MEDLINE® and Web of Science® queries retrieved 921 and 621 citations, respectively. After removal of duplicates, 1,480 citations remained. Section editors screened this initial list for relevance to the theme and scientific quality, and they rated each paper as “keep,” “pend,” or “discard.” Papers rated as “keep” by one of the section editors were independently reviewed and scored by section editors to yield the top 15 candidate best papers [3-17]. Criteria for scoring included innovation beyond established AI

techniques, work that addressed substantial challenges in the field, and rigorous scientific evaluations.

The top-ranked 15 manuscripts underwent formal peer review by at least four internationally renowned scientists from across the world. Based on ratings and comments from peer review, section editors' recommendations, and input from the IMIA Yearbook Editorial Board, three papers [3, 8, 11] were chosen as 2018 best papers in the AI for Health special section of the IMIA Yearbook.

Conclusions and Outlook

In 2018, researchers published promising advances in the science of AI to address some of the most significant challenges to its application in health. Two of the best papers enhanced performance of learning methods by incorporating domain knowledge to limit the learning space. Albers *et al.*, re-introduced the technique of data assimilation, commonly used in aeronautical applications, to the problem of predicting glucose values in type 2 diabetes. In this study, they combined machine learning with an endocrine physiological model to make accurate glucose level predictions, generated recommendations to produce desired outcomes, and learned model parameters to represent phenotypes [3]. This approach allowed accurate prediction with sparse data, which can be common in healthcare, and supported trust through explainability. The authors highlighted the need for a multidisciplinary biomedical informatics pipeline to apply data assimilation to the domain of medicine. Oktay *et al.*, presented a novel strategy to train neural networks to learn underlying anatomy from cardiac imaging data using a stacked autoencoder constrained by global anatomic knowledge [11]. In this study, anatomic rather than physiologic knowledge was leveraged to demonstrate excellent performance in image segmentation tasks using multimodal magnetic resonance imaging (MRI) and ultrasound data sets. Importantly, this approach was able to overcome the very common problems of motion artifact and a lack of internal consistency, which limit

Table 1 Best paper selection of articles for the IMIA Yearbook of Medical Informatics 2019 in the section 'Artificial Intelligence in Health'. The articles are listed in alphabetical order of the first author's surname.

Section
Artificial Intelligence in Health
<ul style="list-style-type: none"> ▪ Albers DJ, Levine ME, Stuart A, Mamykina L, Gluckman B, Hripscak G. Mechanistic Machine Learning: How Data Assimilation Leverages Physiological Knowledge Using Bayesian Inference to Forecast the Future, Infer the Present, and Phenotype. <i>J Am Med Inform Assoc</i> 2018;25(10):1392-401. ▪ Lee J, Sun J, Wang F, Wang S, Jun CH, Jiang X. Privacy-Preserving Patient Similarity Learning in a Federated Environment: Development and Analysis. <i>JMIR Med Inform</i> 2018;6(2):e20. ▪ Oktay O, Ferrante E, Kamnitsas K, Heinrich M, Bai W, Caballero J, Cook SA, de Marva A, Dawes T, O'Regan DP, Kainz B, Glocker B, Rueckert D. Anatomically Constrained Neural Networks (ACNNs): Application to Cardiac Image Enhancement and Segmentation. <i>IEEE Trans Med Imaging</i> 2018;37(2):384-95.

many approaches to image interpretation. Further, although the methodology was designed for segmentation, the algorithm showed good performance in pathological classification.

The third best paper addressed the challenging task of applying learning models in data sets that are distributed across institutions. Lee *et al.*, presented a novel privacy-preserving analytics platform for patient similarity learning in a federated setting, through a multi-hash approach for context dependent cross-institution patient representation, and incorporation of homomorphic encryption for privacy preservation [8]. This work addressed the important problem of enabling privacy-preserving learning in healthcare, where sufficient data to make inferences might be stored across a wide variety of sites.

Five of the remaining selected papers described advances in medical imaging analysis through novel application and enhancements of deep learning methods. Shi *et al.*, applied an emerging technique of stacked deep polynomial networks to a multimodal imaging-classification task [14]. These scientists used a staged approach to building machine learning algorithms, applying stacked deep polynomial networks to MRI and positron emission tomography (PET) images to fuse and learn feature representation to classify images from patients with Alzheimer's disease, normal controls, and mild cognitive impairment who subsequently did or did not progress to Alzheimer's disease. Their approach demonstrated excellent accuracy in a challenging multi-class classification

task for categories of a disease that often progresses gradually and unpredictably. While MRI has become an indispensable diagnostic tool, its slow acquisition process makes it expensive and less accessible. To address this challenge, Schlemper *et al.*, proposed a well-designed and well-executed framework for reconstructing dynamic sequences of two-dimensional cardiac MRI images from aggressively under-sampled data using a deep cascade of convolutional neural networks (CNNs) to accelerate the data acquisition process [13]. Liu *et al.*, focused on another difficult task in using MRI for diagnosis of brain diseases including Alzheimer's disease – that of automatic generation of clinically meaningful features. They proposed a novel approach of identifying discriminative regions called landmarks, followed by application of a CNN for patch-based deep feature learning [10]. Addressing challenges in other common imaging modalities, Gruezemacher *et al.*, developed a three-dimensional (3D) adaptation of deep neural nets (DNNs) for lung nodule detection from computed tomography (CT) images [6]; Hassan *et al.*, presented a novel approach combining tensor-based segmentation, Delaunay triangulation, and deep learning models to provide complete 3D presentation of macula to support automated diagnosis of various macula conditions [7].

Chang *et al.*, addressed another important problem of how to train deep learning models across distributed data sets as an alternative to data sharing [5]. They proposed three non-parallel heuristics (i.e.,

ensembling single institution models, single weight transfer, and cyclic weight transfer) for training deep learning models using distributed data and produced valuable insights through a set of very well designed simulations run on three public data sets. Cyclic weight transfer resulted in performance similar to centrally hosted data, even with low resolution images and institutional variability in class distributions.

Successful development of deep learning models requires abundant labeled data, which is particularly difficult to obtain in medical domain. Li *et al.*, proposed an approach that leverages a type of multi-task predictive neural network called generalized auxiliary-task augmented network (GATAN) to avoid overfitting [9], and they demonstrated improved performance when the learning algorithm has access to limited labeled data.

AI methods for clinical concept and knowledge extraction from text is another important topic, particularly on use of deep learning methods to reduce the need for human-engineered features and thus improve generalizability. Luo *et al.*, proposed a CNN-based method called segment convolutional networks (seg-CNNs) that used word embedding only for the extraction of entity relationships in clinical notes with promising results [17].

Causal inference is an area that is receiving increasing attention, as the field moves from initial proof-of-concept predictive models demonstrating the value of real world data analysis to methods that can provide more actionable insights. Estimating the heterogeneous effect of a treatment using observation data such as that found in electronic health records (EHRs) is one of the most common and challenging use cases of causal inference. Powers *et al.*, compared and evaluated a number of well know methods for estimating heterogeneous treatment effects through a preliminary study using the SPRINT data set [12,18]. They identified causal boosting and causal

multivariate adaptive regression splines as two of the most promising methods and emphasized the need for much more effort in this area.

Three of the selected papers covered topics in genomics, an area that has seen increasing adoption of AI methods. Uppu and Krishna proposed a hybrid approach combining DNN and random forest for the detection of multi-locus interactions among single nucleotide polymorphisms [15]. Xiao *et al.*, proposed a semi-supervised learning framework based on stacked sparse auto-encoder for cancer prediction from RNA-seq data [16]. Belciug and Gorunescu proposed a new learning method based on extreme learning machine to train a single layered feed forward network on microarray and mass spectrometry data for cancer prediction [4].

References

1. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med* 2019;25(1):44-56.
2. Lamy JB, Seroussi B, Griffon N, Kerdelhue G, Jault MC, Bouaud J. Toward a formalization of the process to select IMIA Yearbook best papers. *Methods Inf Med* 2015;54(2):135-44.
3. Albers DJ, Levine ME, Stuart A, Mamykina L, Gluckman B, Hripesak G. Mechanistic machine learning: how data assimilation leverages physiologic knowledge using Bayesian inference to forecast the future, infer the present, and phenotype. *J Am Med Inform Assoc* 2018;25(10):1392-401.
4. Belciug S, Gorunescu F. Learning a single-hidden layer feedforward neural network using a rank correlation-based strategy with application to high dimensional gene expression and proteomic spectra datasets in cancer detection. *J Biomed Inform* 2018;83:159-66.
5. Chang K, Balachandar N, Lam C, Yi D, Brown J, Beers A, et al. Distributed deep learning networks among institutions for medical imaging. *J Am Med Inform Assoc* 2018;25(8):945-54.
6. Gruetzemacher R, Gupta A, Paradise D. 3D deep learning for detecting pulmonary nodules in CT scans. *J Am Med Inform Assoc* 2018;25(10):1301-10.
7. Hassan T, Akram MU, Akhtar M, Khan SA, Yasin U. Multilayered Deep Structure Tensor Delaunay Triangulation and Morphing Based Automated Diagnosis and 3D Presentation of Human Macula. *J Med Syst* 2018;42(11):17.
8. Lee J, Sun JM, Wang F, Wang S, Jun CH, Jiang XQ. Privacy-Preserving Patient Similarity Learning in a Federated Environment: Development and Analysis. *JMIR Med Inf* 2018;6(2):4-24.
9. Li XR, Zhu DX, Levy P. Leveraging auxiliary measures: a deep multi-task neural network for predictive modeling in clinical research. *BMC Med Inform Decis Mak* 2018;18:9.
10. Liu MX, Zhang J, Nie D, Yap PT, Shen DG. Anatomical Landmark Based Deep Feature Representation for MR Images in Brain Disease Diagnosis. *IEEE J Biomed Health Inform* 2018;22(5):1476-85.
11. Oktay O, Ferrante E, Kamnitsas K, Heinrich M, Bai W, Caballero J, et al. Anatomically Constrained Neural Networks (ACNNs): Application to Cardiac Image Enhancement and Segmentation. *IEEE Trans Med Imaging* 2018;37(2):384-95.
12. Powers S, Qian JY, Jung K, Schuler A, Shah NH, Hastie T, et al. Some methods for heterogeneous treatment effect estimation in high dimensions. *Stat Med* 2018;37(11):1767-87.
13. Schlemper J, Caballero J, Hajnal JV, Price AN, Rueckert D. A Deep Cascade of Convolutional Neural Networks for Dynamic MR Image Reconstruction. *IEEE Trans Med Imaging* 2018;37(2):491-503.
14. Shi J, Zheng X, Li Y, Zhang Q, Ying SH. Multimodal Neuroimaging Feature Learning With Multimodal Stacked Deep Polynomial Networks for Diagnosis of Alzheimer's Disease. *IEEE J Biomed Health Inform* 2018;22(1):173-83.
15. Uppu S, Krishna A. A deep hybrid model to detect multi-locus interacting SNPs in the presence of noise. *Int J Med Inf* 2018;119:134-51.
16. Xiao YW, Wu J, Lin ZL, Zhao XD. A semi-supervised deep learning method based on stacked sparse auto-encoder for cancer prediction using RNA-seq data. *Comput Meth Programs Biomed* 2018;166:99-105.
17. Luo Y, Cheng Y, Uzuner O, Szolovits P, Starren J. Segment convolutional neural networks (Seg-CNNs) for classifying relations in clinical notes. *J Am Med Inform Assoc* 2018;25(1):93-8.
18. Wright JT Jr, Williamson JD, Whelton PK, Snyder JK, Sink KM, Rocco MV, et al. A Randomized Trial of Intensive versus Standard Blood-Pressure Control. *N Engl J Med* 2015;373(22):2103-16.

Correspondence to:
Gretchen Purcell Jackson
IBM Watson Health
Cambridge, Massachusetts, USA
Vanderbilt University Medical Center
Nashville, Tennessee, USA
E-mail: Gretchen.Jackson@ibm.com

Appendix: Summary of Best Papers Selected for the 2018 Edition of the IMIA Yearbook, Section AI in Health

Albers DJ, Levine ME, Stuart A, Mamykina L, Gluckman B, Hripcsak G

Mechanistic machine learning: how data assimilation leverages physiological knowledge using bayesian inference to forecast the future, infer the present, and phenotype

J Am Med Inform Assoc 2018;25(10):1392-401

Albers *et al.*, applied the computational technique of data assimilation to predict glucose values, input missing data, and learn model parameters that represent phenotypes in type 2 diabetes. Data assimilation is an established technique that has been used to solve problems in space travel. In this medical application, researchers combined machine learning with a mechanistic model of endocrine physiology that constrained the space for learning and allowed accurate prediction of future blood glucose levels, based on finger-prick glucose measurements and data about meal carbohydrates. The approach performed well even in the context of sparse data, which is a common problem in health applications. The potential for data assimilation to identify physiologically meaningful phenotypes was demonstrated through estimation of distinct model parameters for diabetic and non-diabetic individuals. Incorporation of physiological domain knowledge with machine learning algorithms supports trust and explainability, which are important

for user adoption of AI solutions. The authors beautifully articulated the need for a multidisciplinary biomedical informatics approach and pipeline to leverage the data assimilation method in healthcare.

Oktay O, Ferrante E, Kamnitsas K, Heinrich M, Bai W, Caballero J, Cook SA, de Marvao A, Dawes T, O'Regan DP, Kainz B, Glocker B, Rueckert D

Anatomically Constrained Neural Networks (ACNNs): application to cardiac image enhancement and segmentation

IEEE Trans Med Imaging 2018;37(2):384-95

Oktay *et al.*, presented a novel strategy to train neural networks to learn shape models of underlying anatomy from cardiac imaging data using a stacked autoencoder constrained by global anatomic knowledge. These anatomically constrained neural networks demonstrated excellent performance in image segmentation of multimodal two-dimensional magnetic resonance imaging (MRI) and three-dimensional ultrasound cardiac imaging data sets (i.e., the United Kingdom Digital Heart Project Dataset, the Challenge on Endocardial Three-dimensional Ultrasound Segmentation (CETUS) 2014 Challenge Dataset, and the Automated Cardiac Diagnosis Challenge (ACDC) Medical Image Computing and Computer-Assisted Intervention Society (MICCAI) 2017 Challenge Dataset). In addition to segmentation, the algorithm showed good performance for the task of pathological classification. This approach addressed the common problems of motion artifact, data corruption, and a lack of internal consistency, which limit many approaches to image interpretation, and it could be generalized to other image processing tasks for which prior knowledge can provide guidance to learning models.

Lee J, Sun J, Wang F, Wang S, Jun CH, Jiang X

Privacy-preserving patient similarity learning in a federated environment: development and analysis

JMIR Med Inform 2018;6(2):e20

Lee *et al.*, addressed the important and challenging task of enabling learning in a federated environment in healthcare, where sufficient data to make inferences might be stored across a wide variety of institutional sites. The authors presented a novel privacy-preserving analytics platform for learning patient similarity in a federated setting, using a multi-hash approach for context-dependent, cross-institution patient representation with incorporation of homomorphic encryption for privacy preservation. The authors evaluated their approach using sequential medical events extracted from the Multiparameter Intelligent Monitoring in Intensive Care (MIM-IC)-III database, a collection of data from patients admitted to the intensive care unit at Beth Israel Deaconess Hospital, to predict the incidence of five diseases common in critically ill patients: disorders of lipid metabolism, hypertensive chronic kidney disease, cardiac dysrhythmias, heart failure, and acute renal failure. They compared the performance of their federated approach with an open system, where patient data can be shared without restrictions across sites, and a closed system, where data can only be used locally. Their system with a multi-hash approach exceeded closed and approximated open system performance, and calculated similarities across sites remained the same after homomorphic encryption, demonstrating preservation of privacy.