




Artificial Intelligence for Computer-Aided Diagnosis in Medical Imaging


Jianming (Jimmy) Liang

**Arizona State University
Johnson Research Building, Mayo Clinic Arizona**

jianming.liang@asu.edu

Disclosure: Co-Founder of IMANIN and VoxelCloud, and Consultant of VisionGate





Research Goal

Develop **computer algorithms** for gleaning clinically important information from images to **support clinical decision making** and facilitate precision medicine


Computer-Aided Diagnosis (CAD) in Medical Imaging

Computer Aided Diagnosis (CAD)

CAD systems are not designed to replace physicians, but rather to enhance their capabilities through a computer-physician synergy

CAD: A Brief History (Doi, 2007)


1980s: Kunio Doi at the University of Chicago




CAD: A Brief History (Doi, 2007)

1960 – 70: Automated computer diagnosis


- Lodwick, et al. 1963;
- Myers, et al. 1964;
- Winsbarg, et al. 1967;



CAD Systems



Deep Learning



Chartrand, et al. Deep Learning: A Primer for Radiologists. Radiographics. 2017 Nov-Dec;37(7):2113-2131

Deep Learning

(Goodfellow, Bengio, Courville. 2016)

Artificial Intelligence

Machine Learning

Representation Learning

Deep Learning



6,860 views | Mar 27, 2019, 01:43pm

Turing Award And \$1 Million Given To 3 AI Pioneers



Nicole Martin Contributor @
in a big data
I write about technology, data and privacy.



Winners of Turing Award | nbc news photo

The Association for Computing Machinery (ACM) awarded Yoshua Bengio, Geoffrey Hinton and Yann LeCun with what many consider the "Nobel Prize of computing," for the innovations they've made in AI.



Courtesy: NVIDIA (<https://www.youtube.com/watch?v=YuyT2SDcYrU>).

Deep Learning


Revolutionized computer vision, self-driving cars, speed understanding, and natural language processing

The greatest potential

Deep Learning in Healthcare

A revolution is coming to computer-aided diagnosis


Deep Learning



Deep Learning: A Primer for Radiologists¹

Gabriel Chartand, PhD
Philip M. Chung, MD, MS
Eugene Vrontos, BSc, Eng Sci
Michal Denczka, PhD
Simon Turchetta, MD, MS
Christopher J. Pal, PhD
Samuel Kadlavy, PhD
An Tang, MD, MS



CRCHUM



Three Systems (Hybrid systems)

Project 1:

Pulmonary Embolism



Reading CT Angiography



Reading CT Angiography

A needle in a huge haystack

14% under-diagnosis
10% over-diagnosis


Lucassen, et al. Thromb Res. 2013 Feb;131(2):145-9.

A CAD System for PE Detection


CT Scan

PE CAD

Markers



As a first reader, concurrent, second reader



Computer-aided Pulmonary Embolism Detection Using a Novel Vessel-Aligned Multi-Planar Image Representation and Convolutional Neural Networks


Nima Tajbakhsh*, Michael B. Gotway*, and Jianming Liang*

*Department of Biomedical Informatics, Arizona State University

Department of Radiology, Mayo Clinic, Scottsdale, AZ, USA

#1 in the CAD PE competition

Our new image representation... The PE system as a 3D PE system... The PE system as a 3D PE system... The PE system as a 3D PE system...



A Finalist for a Young Scientist Award (MICCAI 2016)

Project 2

Cardiovascular Disease



Cardiovascular Disease (CVD)

#1 killer in the US

- Every minute, one American dies
- 1/2 of the deaths occur suddenly
- 1/3 of the deaths occur before the age of 65

Preventive medicine



- CVD is largely preventable
- Efficacious therapies are available

Key: identify at-risk individuals before CVD events



CIMT: Carotid Intima-Media Thickness

Quantifying subclinical atherosclerosis



Source: <http://dallasfamilymedical.com/heart-attack-stroke-risk-assessment/>

No radiation. No needles. No pain. Affordable

CIMT: Carotid Intima-Media Thickness

Each patient has 4 videos


- Two on each side

Search the right frame

- From hundreds of frames within each video


Find the right location

- Within the found frames



Measuring CIMT is tedious and time-consuming

Contribution 1: Fully Automating CIMT Video Analysis



1. Jae, et al, Automating Carotid Intima-Media Thickness Video Interpretation with Convolutional Neural Networks, CVPR 2016
2. Toghiani, et al, Automatic interpretation of CIMT videos using convolutional neural networks, Deep Learning for Medical Image Analysis, Academic Press, 2017

Clinical Evaluation (preliminary)


- 25 unseen subjects (100 CIMT videos)
- Two experts measure two times
- Automated measurement

Five measurements are not statistically distinguishable

CNNs have proven to be powerful in learning features, but to achieve a superior performance, it is critical to design meaningful image representations and post-process CNN's outputs.


Contribution 2: Boosting CNN performance via effective pre- and post-processing

Frame Selection




Contribution 3: Boosting CNN performance with an Anatomical Constraint

ROI Localization




Contribution 4: Boosting CNN performance through Active Contours

IMT measurement



Project 3

Colon Cancer





Colon Cancer

#2 leading cause of cancer death

both men and women

One of the most preventable

doctors can identify and remove the pre-cancerous growth: polyp with colonoscopy



Colonoscopy

Primary method for screening and preventing colon cancer

- contribute significantly to the decline of colon cancer

User-dependent procedure

- rely on colonopists' skills (diligence, attentiveness and thoroughness)



For every 5 polyps, ≈ 1 is missed


Reasons for missing polyps

- Unreached polyps**
 - incomplete colon examination
- Reached but unrecognizable**
 - rapid motion, too close
- Recognizable but overlooked**
 - inattentiveness




Polyps: Recognizable but overlooked

Automatically detecting polyps in colonoscopy videos




Detect Polyps:

Automatically alerting the physician to possible polyps



Polyps: Reached but unrecognizable

Automatically assessing informativeness of colonoscopy video frames



Monitor Video Quality:

Automatically monitoring the colonoscopic video quality



Polyps: Unreached

Visualizing the colon in 3D with examined regions



#1 in the polyp detection competition

A Comprehensive Computer-Aided Polyp Detection System for Colonoscopy Videos
Nima Tajbakhsh*, Suryakant R. Gurudip*, and Jianming Liang*
*Department of Biomedical Informatics, Arizona State University
*Division of Gastroenterology and Hepatology, Mayo Clinic, Arizona, USA

METHOD
Stage 1: candidate generation
Stage 2: candidate classification

OUR PATCH DESCRIPTOR
We have designed a novel patch descriptor based on 3D convolutional features, which is designed to capture the local and global context of the polyp. Our descriptor is a 3D vector of size 1024, which is used for candidate generation.

OUR VOTING SCHEME
We have designed a novel voting scheme based on 3D convolutional features, which is designed to capture the local and global context of the polyp. Our voting scheme is a 3D vector of size 1024, which is used for candidate classification.

OUR IMAGE REPRESENTATION & CONVOLUTIONAL NEURAL NETWORKS
We have designed a novel image representation based on 3D convolutional features, which is designed to capture the local and global context of the polyp. Our image representation is a 3D vector of size 1024, which is used for candidate generation and classification.

EVALUATION
For evaluation, we use 40 colonoscopy videos, which are divided into training and testing sets. The training set consists of 30 videos, and the testing set consists of 10 videos. We use the F1 score to evaluate the performance of our system. Our system achieves an F1 score of 0.95, which is the highest among all the systems in the competition.

ACKNOWLEDGMENT
This work was supported by the National Institutes of Health (NIH) under grant R01CA201010.

Other Projects

Classifying thyroid nodules based on ultrasound images

- The overall incidence of cancer in patients with thyroid nodules selected for the fine needle aspiration (FNA) is 9.2%–13.0%.
- Reducing biopsies (automatically determining malignancy)

Detecting, segmenting and classifying lung nodules


- Detecting and segmenting lung nodules at CT
- determining malignancy of lung nodules; reducing biopsies

Detecting, segmenting and classifying brain tumors

- Detecting and segmenting brain tumors at CT
- Classifying brain tumors into three types: 1p/19q, IDH, and TERT
- Reducing biopsies

Proton therapy for lung cancer

- Outlining tumors and at-risk organs
- Predicting the tumor response to proton therapy



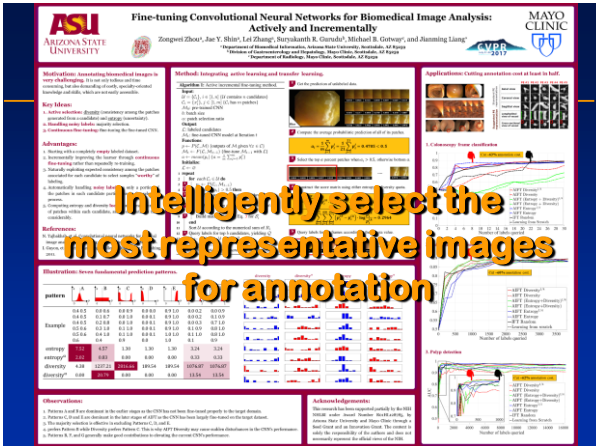


Methodologies across diseases & modalities

Deep Learning for Medical Imaging

A Significant Disadvantage

Not enough annotation

Not enough annotation



Active Fine-Tuning

Algorithm 1: Active fine-tuning

Input:
 $\mathcal{U} = \{C_i\}$, $i \in [1, n]$ $\{\mathcal{U}$ contains n AUs}
 $C_i = \{x_i^j\}$, $j \in [1, m]$ $\{C_i$ has m objects}
 \mathcal{M} : a pre-trained CNN
 b : batch size
Output:
 \mathcal{L} : the labeled AUs
 \mathcal{M}_t : the fine-tuned CNN model at Iteration t

```

1  $\mathcal{L} \leftarrow \emptyset$ 
2 repeat
3   for each  $C_i \in \mathcal{U}$  do
4      $p_i \leftarrow P(C_i, \mathcal{M}_{t-1})$  {outputs of  $\mathcal{M}_{t-1}$  given  $\forall x_i \in C_i$ }
5      $\mathcal{E}_i \leftarrow E(C_i)$  {compute entropy  $\mathcal{E}_i$  for  $C_i$  using Eq. 1}
6   end
7    $\mathcal{U}' \leftarrow S(\mathcal{U}, \mathcal{E})$  {sort  $C_i \in \mathcal{U}$  according to the value of  $\mathcal{E}_i \in \mathcal{E}$ }
8    $Q \leftarrow Q(\mathcal{U}', b)$  {associate labels for the top  $b$  AUs in the sorted  $\mathcal{U}'$ }
9    $\mathcal{L} \leftarrow \mathcal{L} \cup Q$ ;  $\mathcal{U} \leftarrow \mathcal{U} \setminus Q$ ;  $t \leftarrow t+1$ 
10   $\mathcal{M}_t \leftarrow F(\mathcal{L}, \mathcal{M})$  {fine-tune  $\mathcal{M}$  with  $\mathcal{L}$ }
11 until classification performance is satisfactory;
```

Active Fine-Tuning

Algorithm 1: Active fine-tuning

Input:
 $\mathcal{U} = \{C_i\}$, $i \in [1, n]$ $\{\mathcal{U}$ contains n AUs}
 $C_i = \{x_i^j\}$, $j \in [1, m]$ $\{C_i$ has m objects}
 \mathcal{M} : a pre-trained CNN
 b : batch size
Output:
 \mathcal{L} : the labeled AUs
 \mathcal{M}_t : the fine-tuned CNN model at Iteration t

```

1  $\mathcal{L} \leftarrow \emptyset$ 
2 repeat
3   for each  $C_i \in \mathcal{U}$  do
4      $p_i \leftarrow P(C_i, \mathcal{M}_{t-1})$  {outputs of  $\mathcal{M}_{t-1}$  given  $\forall x_i \in C_i$ }
5      $\mathcal{E}_i \leftarrow E(C_i)$  {compute entropy  $\mathcal{E}_i$  for  $C_i$  using Eq. 1}
6   end
7    $\mathcal{U}' \leftarrow S(\mathcal{U}, \mathcal{E})$  {sort  $C_i \in \mathcal{U}$  according to the value of  $\mathcal{E}_i \in \mathcal{E}$ }
8    $Q \leftarrow Q(\mathcal{U}', b)$  {associate labels for the top  $b$  AUs in the sorted  $\mathcal{U}'$ }
9    $\mathcal{L} \leftarrow \mathcal{L} \cup Q$ ;  $\mathcal{U} \leftarrow \mathcal{U} \setminus Q$ ;  $t \leftarrow t+1$ 
10   $\mathcal{M}_t \leftarrow F(\mathcal{L}, \mathcal{M})$  {fine-tune  $\mathcal{M}$  with  $\mathcal{L}$ }
11 until classification performance is satisfactory;
```

Active Fine-Tuning

Algorithm 1: Active fine-tuning

Input:
 $\mathcal{U} = \{C_i\}$, $i \in [1, n]$ $\{\mathcal{U}$ contains n AUs}
 $C_i = \{x_i^j\}$, $j \in [1, m]$ $\{C_i$ has m objects}
 \mathcal{M} : a pre-trained CNN
 b : batch size
Output:
 \mathcal{L} : the labeled AUs
 \mathcal{M}_t : the fine-tuned CNN model at Iteration t

```

1  $\mathcal{L} \leftarrow \emptyset$ 
2 repeat
3   for each  $C_i \in \mathcal{U}$  do
4      $p_i \leftarrow P(C_i, \mathcal{M}_{t-1})$  {outputs of  $\mathcal{M}_{t-1}$  given  $\forall x_i \in C_i$ }
5      $\mathcal{E}_i \leftarrow E(C_i)$  {compute entropy  $\mathcal{E}_i$  for  $C_i$  using Eq. 1}
6   end
7    $\mathcal{U}' \leftarrow S(\mathcal{U}, \mathcal{E})$  {sort  $C_i \in \mathcal{U}$  according to the value of  $\mathcal{E}_i \in \mathcal{E}$ }
8    $Q \leftarrow Q(\mathcal{U}', b)$  {associate labels for the top  $b$  AUs in the sorted  $\mathcal{U}'$ }
9    $\mathcal{L} \leftarrow \mathcal{L} \cup Q$ ;  $\mathcal{U} \leftarrow \mathcal{U} \setminus Q$ ;  $t \leftarrow t+1$ 
10   $\mathcal{M}_t \leftarrow F(\mathcal{L}, \mathcal{M})$  {fine-tune  $\mathcal{M}$  with  $\mathcal{L}$ }
11 until classification performance is satisfactory;
```

Pre-Trained CNN



Active Fine-Tuning

Algorithm 1: Active fine-tuning

Input:
 $\mathcal{U} = \{C_i\}$, $i \in [1, n]$ $\{\mathcal{U}$ contains n AUs}
 $C_i = \{x_i^j\}$, $j \in [1, m]$ $\{C_i$ has m objects}
 \mathcal{M} : a pre-trained CNN
 b : batch size
Output:
 \mathcal{L} : the labeled AUs
 \mathcal{M}_t : the fine-tuned CNN model at Iteration t

```

1  $\mathcal{L} \leftarrow \emptyset$ 
2 repeat
3   for each  $C_i \in \mathcal{U}$  do
4      $p_i \leftarrow P(C_i, \mathcal{M}_{t-1})$  {outputs of  $\mathcal{M}_{t-1}$  given  $\forall x_i \in C_i$ }
5      $\mathcal{E}_i \leftarrow E(C_i)$  {compute entropy  $\mathcal{E}_i$  for  $C_i$  using Eq. 1}
6   end
7    $\mathcal{U}' \leftarrow S(\mathcal{U}, \mathcal{E})$  {sort  $C_i \in \mathcal{U}$  according to the value of  $\mathcal{E}_i \in \mathcal{E}$ }
8    $Q \leftarrow Q(\mathcal{U}', b)$  {associate labels for the top  $b$  AUs in the sorted  $\mathcal{U}'$ }
9    $\mathcal{L} \leftarrow \mathcal{L} \cup Q$ ;  $\mathcal{U} \leftarrow \mathcal{U} \setminus Q$ ;  $t \leftarrow t+1$ 
10   $\mathcal{M}_t \leftarrow F(\mathcal{L}, \mathcal{M})$  {fine-tune  $\mathcal{M}$  with  $\mathcal{L}$ }
11 until classification performance is satisfactory;
```

Pre-Trained CNN



Select most informative and representative

Entropy & Diversity

Active Fine-Tuning

Algorithm 1: Active fine-tuning

Input:
 $\mathcal{U} = \{C_i\}$, $i \in [1, n]$ $\{\mathcal{U}$ contains n AUs}
 $C_i = \{x_i^j\}$, $j \in [1, m]$ $\{C_i$ has m objects}
 \mathcal{M} : a pre-trained CNN
 b : batch size
Output:
 \mathcal{L} : the labeled AUs
 \mathcal{M}_t : the fine-tuned CNN model at Iteration t

```

1  $\mathcal{L} \leftarrow \emptyset$ 
2 repeat
3   for each  $C_i \in \mathcal{U}$  do
4      $p_i \leftarrow P(C_i, \mathcal{M}_{t-1})$  {outputs of  $\mathcal{M}_{t-1}$  given  $\forall x_i \in C_i$ }
5      $\mathcal{E}_i \leftarrow E(C_i)$  {compute entropy  $\mathcal{E}_i$  for  $C_i$  using Eq. 1}
6   end
7    $\mathcal{U}' \leftarrow S(\mathcal{U}, \mathcal{E})$  {sort  $C_i \in \mathcal{U}$  according to the value of  $\mathcal{E}_i \in \mathcal{E}$ }
8    $Q \leftarrow Q(\mathcal{U}', b)$  {associate labels for the top  $b$  AUs in the sorted  $\mathcal{U}'$ }
9    $\mathcal{L} \leftarrow \mathcal{L} \cup Q$ ;  $\mathcal{U} \leftarrow \mathcal{U} \setminus Q$ ;  $t \leftarrow t+1$ 
10   $\mathcal{M}_t \leftarrow F(\mathcal{L}, \mathcal{M})$  {fine-tune  $\mathcal{M}$  with  $\mathcal{L}$ }
11 until classification performance is satisfactory;
```

Pre-Trained CNN



Select most informative and representative

Active Fine-Tuning

Algorithm 1: Active fine-tuning

Input:
 $\mathcal{U} = \{C_i\}$, $i \in [1, n]$ $\{\mathcal{U}$ contains n AUs}
 $C_i = \{x_i^j\}$, $j \in [1, m]$ $\{C_i$ has m objects}
 \mathcal{M} : a pre-trained CNN
 b : batch size
Output:
 \mathcal{L} : the labeled AUs
 \mathcal{M}_t : the fine-tuned CNN model at Iteration t

```


1  $\mathcal{L} \leftarrow \emptyset$ 
2 repeat
3   for each  $C_i \in \mathcal{U}$  do
4      $p_i \leftarrow P(C_i, \mathcal{M}_{t-1})$  {outputs of  $\mathcal{M}_{t-1}$  given  $\forall x_i \in C_i$ }
5      $\mathcal{E}_i \leftarrow E(C_i)$  {compute entropy  $\mathcal{E}_i$  for  $C_i$  using Eq. 1}
6   end
7    $\mathcal{U}' \leftarrow S(\mathcal{U}, \mathcal{E})$  {sort  $C_i \in \mathcal{U}$  according to the value of  $\mathcal{E}_i \in \mathcal{E}$ }
8    $Q \leftarrow Q(\mathcal{U}', b)$  {associate labels for the top  $b$  AUs in the sorted  $\mathcal{U}'$ }
9    $\mathcal{L} \leftarrow \mathcal{L} \cup Q$ ;  $\mathcal{U} \leftarrow \mathcal{U} \setminus Q$ ;  $t \leftarrow t+1$ 
10   $\mathcal{M}_t \leftarrow F(\mathcal{L}, \mathcal{M})$  {fine-tune  $\mathcal{M}$  with  $\mathcal{L}$ }
11 until classification performance is satisfactory;
```

Pre-Trained CNN



Select most informative and representative

Active Fine-Tuning



Active Fine-Tuning

Dramatically reduces the burden of annotation

AFT with ~5% = LS with whole


AFT saves > 50% relative to FT

Quality beats Quantity

AFT with ~20% = LS with whole

AFT with ~15% = LS with whole

AFT with ~20% = LS with whole



Deep Learning for Medical Imaging

A Significant Advantage

Plenty weakly labeled image data

Methodological Research across diseases and modalities

2. Deep learning with weakly labeled data


- Automatically understanding radiological reports




Methodological Research across diseases and modalities

2. Deep learning with weakly labeled data


- Automatically understanding radiological reports



- **Automatically understanding radiological reports**



- Biomedical images → radiological reports



Reinforcement learning

Flying through the colon



New Applications

Two Types of Tasks

1. **Tasks that doctors have to do but do not really want to do (low value, repetitive)**
 - Tedious, laborious, and time consuming
 - CIMT video analysis; outlining the contours of the tumors and at-risk organs
2. **Tasks that doctors really want to do but cannot do very well (high value)**
 - Lack of knowledge and skills; require subvisual features
 - Determining the malignancy of tumor (benign or malignant) before biopsy or surgery; predicting the tumor response to radiation/chemo therapy

Domains

1. **Radiotherapy**
 - Outlining tumors and organs at-risk
 - Predicting the tumor response to radiation therapy
2. **Chemotherapy**
 - Predicting the tumor response to chemotherapy
3. **Ultrasound**
 - Ultrasounds are inexpensive and widely accessible
 - User-dependent; Difficult to interpret ultrasound images
4. **Generic image segmentation**
 - Tedious, laborious, and time consuming

Clinical Projects at CRCHUM

1. Automatically outlining tumors and organs at-risk in **radiation therapy** to relieve clinicians from the tedious, repetitive, laborious, time consuming, and error prone process, dramatically reducing patient turnaround times (in collaboration with **Dr. Francois DeBlois** and **Dr. David Roberge** at **CRCHUM**)
2. Predicting liver metastases response to **chemotherapy** to triage patients for chemotherapy and surgery, improving patient life quality (in collaboration with **Dr. An Tang** and **Dr. Simon Turcotte** at **CRCHUM**)
3. Deep learning for **ultrasound imaging** to simplify ultrasound operations, clarify ultrasound images, and amplify ultrasound portability, revolutionizing ultrasound imaging in clinical practice (in collaboration with **Dr. Gilles Soulez**, **Dr. Guy Cloutier**, and **Dr. Samuel Kadoury** at **CRCHUM**)
4. Accurately segmenting abdominal aortic aneurysm for **precise diameter measurements**, realizing stent graft customization (in collaboration with **Dr. Gilles Soulez** and **Dr. Claude Kauffmann** at **CRCHUM**)

Methodological Research across diseases and modalities

5. Creating a large chest image database and building a mother model for medical imaging

- Chest may be the best area:
 - X-ray, CT scans, MRI, ...
 - Incidental findings, ...
- **Canada (Montreal) may be the best place to start this**
- **Drs. Carl Chartrand-Lefebvre and Rayyan Daghistani**



Generative Models



Courtesy: NVIDIA (<https://www.youtube.com/watch?v=Zc-ihT0DQg0>).

Generative Models




Figure 1. Multi-domain image-to-image translation results on the CelebA dataset via transferring knowledge learned from the RaFD dataset. The first and sixth columns show input images while the remaining columns are images generated by StarGAN. Note that the images are generated by a single generator network, and facial expression labels such as angry, happy, and fearful are from RaFD, not CelebA.

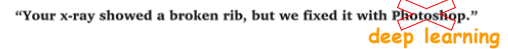
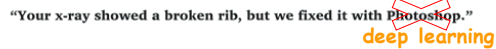
Yunjey Choi, et al. CVPR 2018.



Figure 1: Three source photos (top row) are each modified to match makeup styles in three reference photos (left column) to produce nine different outputs (3×3 lower right).

Huiwen Chang, et al. CVPR 2018







**Detecting any and every disease
(even at early stage) in images
with only patient-level annotation**



Consistent, recurrent anatomy



MICCAI 2019

Models Genesis: Generic Autoinductive Models for 3D Medical Image Analysis

Zongren Zhou¹, Varad Sathya¹, Mihai Alexandru Holian¹, Huihui Feng², Nina Tajbakhsh³, Michael B. Gotway³, and Jianming Liang¹

¹ Arizona State University, Scottsdale, AZ 85281 USA, {zongren.zhou, varad.sathya, mihai.holian, huihui.feng, nina.tajbakhsh, michael.b.gotway, jianming.liang}@asu.edu

² Mayo Clinic, Scottsdale, AZ 85259 USA, zongren.zhou@mayo.edu

³ Mayo Clinic, Scottsdale, AZ 85259 USA, zongren.zhou@mayo.edu

Abstract. Transfer learning from natural image to medical image has established as one of the most practical paradigms in deep learning for medical image analysis. However, to fit this paradigm, 2D imaging tasks in the most prominent imaging modalities (e.g., CT and MRI) have to be reformulated and solved in 2D, losing rich 3D anatomical information and inevitably compromising the performance. To overcome this limitation, we have built a set of models, called Genesis, Inductive Models, inducting Models Genesis, because they are created or re-built with no manual labeling, self-supervised learning by self-supervised, and generic (used as source models for generating applications-specific target models). Our extensive experiments demonstrate that our Models Genesis significantly outperform learning from scratch in all five target 3D applications covering both segmentation and classification. More importantly, learning a model from scratch simply in 3D may not necessarily yield performance better than transfer learning from ImageNet in 2D, but our Models Genesis consistently top any 2D approaches including fine-tuning the models pre-trained from ImageNet as well as fine-tuning the 2D versions of our Models Genesis, confirming the importance of 3D anatomical information and significance of our Models Genesis for 3D medical imaging. The performance is attributed to our unified self-supervised learning framework, built on a simple and powerful characteristic, the self-supervised or over-segmentation in natural images can serve as strong supervision signals for deep models to learn common anatomical representation automatically via self-supervision. As open sources, all pre-trained Models Genesis are available at <https://github.com/ModelGenesis/ModelGenesis>.

1 Introduction

Given the marked difference between natural images and medical images, we hypothesize that transfer learning can yield more powerful (application-specific) target models if the source models are built directly from natural images. To test this hypothesis, we have chosen chest imaging because the chest contains several

Result 1:


Models Genesis outperform 3D models trained from scratch

Task	Modality	Metric	Scratch (%)	Genesis (%)	p-value
Lung nodule false positive reduction	CT	AUC	94.25±5.07	98.20±0.51	0.0180
Lung nodule segmentation	CT	IoU	74.05±1.97	77.62±0.64	1.04e-4
PE false positive reduction	CT	AUC	79.99±8.06	88.04±1.40	0.0058
Liver segmentation	CT	IoU	74.60±4.57	79.52±4.77	0.0361
Brain tumor segmentation	MRI	IoU	90.16±0.41	90.60±0.20	0.0041

The statistical analyses are conducted between Scratch and Genesis.

Result 2:


Models Genesis consistently top any 2D approaches



Lung nodule false positive reduction in 3D CT [LUNA-2016]


Lung nodule segmentation 3D CT [LIDC-IDRI] PE false positive reduction 3D CT [PE-CAD]

Result 3: Models Genesis (2D) offer performances equivalent to supervised pre-trained models



Eight pulmonary diseases classification
2D X-ray [ChestX-ray8]

RoI/bulb/background classification
2D Ultrasound [UFL MCAEL]



A Significant Challenge


Image Data Explosion



A Manifestation of **Big Data** in Medical Imaging

A Great Opportunity

Deep Learning

A revolution is coming to
medical image interpretation



[illegible]


Automating Carotid Intima-Media Thickness Video Interpretation with Convolutional Neural Networks

Joe Wu¹, Hieu Septhidar¹, R. Sudhakar¹, Christopher A. Kessel², Jazayeri, Ian^{1*}

¹Arizona State University, Tempe, Arizona, USA; ²Massachusetts General Hospital, Boston, Massachusetts, USA

ABSTRACT: Cardiovascular disease (CVD) is the leading cause of mortality and morbidity in the United States. Carotid Intima-Media Thickness (CIMT) is a common ultrasound metric that is used to assess the risk of CVD. However, carotid CMT assessment requires expert attention and is often time-consuming. This paper presents a deep learning-based approach for automated carotid CMT measurement. The proposed method consists of three main steps: (1) Carotid CMT video segmentation, (2) Carotid CMT measurement, and (3) Carotid CMT measurement validation. The proposed method is evaluated on a dataset of 100 carotid CMT videos. The results show that the proposed method achieves a mean absolute error (MAE) of 0.15 mm, which is comparable to the results of the state-of-the-art methods. The proposed method is also evaluated on a dataset of 100 carotid CMT videos. The results show that the proposed method achieves a mean absolute error (MAE) of 0.15 mm, which is comparable to the results of the state-of-the-art methods.

INTRODUCTION

Carotid Intima-Media Thickness (CIMT) is a common ultrasound metric that is used to assess the risk of CVD. However, carotid CMT assessment requires expert attention and is often time-consuming. This paper presents a deep learning-based approach for automated carotid CMT measurement. The proposed method consists of three main steps: (1) Carotid CMT video segmentation, (2) Carotid CMT measurement, and (3) Carotid CMT measurement validation. The proposed method is evaluated on a dataset of 100 carotid CMT videos. The results show that the proposed method achieves a mean absolute error (MAE) of 0.15 mm, which is comparable to the results of the state-of-the-art methods.

Method

The proposed method consists of three main steps: (1) Carotid CMT video segmentation, (2) Carotid CMT measurement, and (3) Carotid CMT measurement validation. The proposed method is evaluated on a dataset of 100 carotid CMT videos. The results show that the proposed method achieves a mean absolute error (MAE) of 0.15 mm, which is comparable to the results of the state-of-the-art methods.

Carotid CMT Video Segmentation

The first step of the proposed method is Carotid CMT video segmentation. This step involves identifying the carotid artery in the ultrasound video and segmenting the intima-media layer. The proposed method uses a deep learning-based approach for this task. The results show that the proposed method achieves a mean absolute error (MAE) of 0.15 mm, which is comparable to the results of the state-of-the-art methods.

Carotid CMT Measurement

The second step of the proposed method is Carotid CMT measurement. This step involves measuring the thickness of the intima-media layer. The proposed method uses a deep learning-based approach for this task. The results show that the proposed method achieves a mean absolute error (MAE) of 0.15 mm, which is comparable to the results of the state-of-the-art methods.

Carotid CMT Measurement Validation

The third step of the proposed method is Carotid CMT measurement validation. This step involves validating the results of the Carotid CMT measurement. The proposed method uses a deep learning-based approach for this task. The results show that the proposed method achieves a mean absolute error (MAE) of 0.15 mm, which is comparable to the results of the state-of-the-art methods.

Conclusion

This paper presents a deep learning-based approach for automated carotid CMT measurement. The proposed method consists of three main steps: (1) Carotid CMT video segmentation, (2) Carotid CMT measurement, and (3) Carotid CMT measurement validation. The proposed method is evaluated on a dataset of 100 carotid CMT videos. The results show that the proposed method achieves a mean absolute error (MAE) of 0.15 mm, which is comparable to the results of the state-of-the-art methods.

Augmented Intelligence

Our systems are not designed to replace physicians, but rather to enhance their capabilities through a computer-physician synergy

Experts + Computer >> Experts

- Do what doctors do not want to do
- Do what doctors cannot do well

Clinical Collaborators



Cardiology

Radiology

Gastroenterology

Acknowledgements

- NIH R01 (PE CAD)
- ABRC R01-like Grant (Proton Therapy)
- Mayo Discovery Translation (CMT)
- ASU-Mayo Grant (Proton Therapy)
- Mayo Innovation Grant (Colonoscopy)
- ASU-Mayo Grant (Colonoscopy)
- Mayo CR 20 Program (CMT)
- ASU-Mayo Grant (PE CAD)
- ASU-Mayo Grant (PE)

Thank You