How to develop a quality organization of doctoral dissertations and thesis defenses?

Zongwei Zhou, PhD

The organization and content of informatics doctoral dissertations

RECEIVED 28 October 2015
REVISED 15 April 2016
ACCEPTED 19 April 2016
PUBLISHED ONLINE FIRST 6 June 2016

Edward H Shortliffe

AMIA INFORMATICS PROFESSIONALS. LEADING THE WAY.
OXFORD UNIVERSITY PRESS

ABSTRACT

This article offers suggested guidelines for graduate students who are embarking on informatics doctoral studies and anticipating the dissertation research and its documentation. Much of the guidance is pertinent for writing dissertations in other disciplines as well. The messages are largely directed at doctoral students, but some elements are also pertinent for master's students. All are relevant for faculty research advisors. The value of the dissertation is often underestimated. Too often it is seen as a hurdle to be overcome rather than an opportunity to gain insight into one's own research and to learn how to communicate effectively about it. Ideas that have been ill-formed often do not gel effectively until one tries to write about them. The main lesson is that the preparation of a carefully crafted, rigorous, logically evidence-based, and influential dissertation can be remarkably rewarding, both personally and professionally.

Keywords: academic dissertations, informatics education, research report, writing, authorship

The Value of the Dissertation

- Project 1. Fine-tuning Convolutional Neural Networks for Biomedical Image Analysis: Actively and Incrementally (2017)
- Project 2. UNet++: A Nested U-Net Architecture for Medical Image Segmentation (2018)
- Project 3. Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis (2019)
- I could have graduated in May 2020 with a decent dissertation and defense
  - Introduction-P1-P2-P3-Conclusion
  - P1-P3 consist of introduction, related work, method, experiment, result, conclusion
  - Or I could have explored Project 4
- What really happened: I started PhD in August 2016 and graduated in May 2021
  - *What was I doing in 2020-2021?*

The Value of the Dissertation

- The dissertation is archived (*forever*) once submitted, and it is your identify for quite a long time, especially if you continue to pursue academia position
  - A collection of first-authored work
  - Most innovative ideas occur in your 20s, 30s
  - Foundation of research taste, funding application, recruitment
  - Very important for faculty application: job talk, dissertation award…
- *Don't miss the lifetime opportunity* to devote a good amount of energy and time to achieve the highest possibly quality for something
  - 99% of the time we produce average-level work for many reasons
  - It indicates the upper-bound of your possibility

The last years of graduate work provide a unique opportunity: coursework is over, the research is maturing, and the student should have a singleness of purpose and be fully consumed by the research and the dissertation. In fact, as most postgraduates can attest, those few years may be the last time that trainees will have the luxury of focusing their work and their thinking almost completely on a single activity. It is accordingly the responsibility of faculty members to make sure that students are protected from other demands and distractions so that the full potential of their work, and the attendant lessons, can be achieved.

Edward H Shortliffe

ABSTRACT

This article offers suggested guidelines for graduate students who are embarking on informatics doctoral studies and anticipating the dissertation research and its documentation. Much of the guidance is pertinent for writing dissertations in other disciplines as well. The messages are largely directed at doctoral students, but some elements are also pertinent for master's students. All are relevant for faculty research advisors. The value of the dissertation is often underestimated. Too often it is seen as a hurdle to be overcome rather than an opportunity to gain insight into one's own research and to learn how to communicate effectively about it. Ideas that have been ill-formed often do not gel effectively until one tries to write about them. The main lesson is that the preparation of a carefully crafted, rigorous, logically evidence-based, and influential dissertation can be remarkably rewarding, both personally and professionally.

Keywords: academic dissertations, informatics education, research report, writing, authorship

The Contribution to the Underlying Science

The challenges may be exacerbated in biomedical informatics, where the field is young and the examples for students to emulate may be limited at their own institution. Furthermore, informatics is, by its nature, motivated by applications and by the needs of the world of biomedicine. Many informatics dissertations address a substantive problem in biomedicine, yet the dissertation cannot simply present a solution to that applied problem but must also identify how the work contributes to the underlying science of the informatics discipline.

The Contribution to the Underlying Science

- Project 1. Fine-tuning Convolutional Neural Networks for Biomedical Image Analysis: Actively and Incrementally (2017)
  - *How to select important data to annotate?*
- Project 2. UNet++: A Nested U-Net Architecture for Medical Image Segmentation (2018)
  - *How to improve AI's capability of disease segmentation?*
- Project 3. Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis (2019)
  - *How to extract visual representation from images?*
- Unfortunately, these are all solutions to specific problems.
- Graduate students often do not have a big picture—impacts to the underlying science—at the time they are working on specific projects, so am I.

The Contribution to the Underlying Science

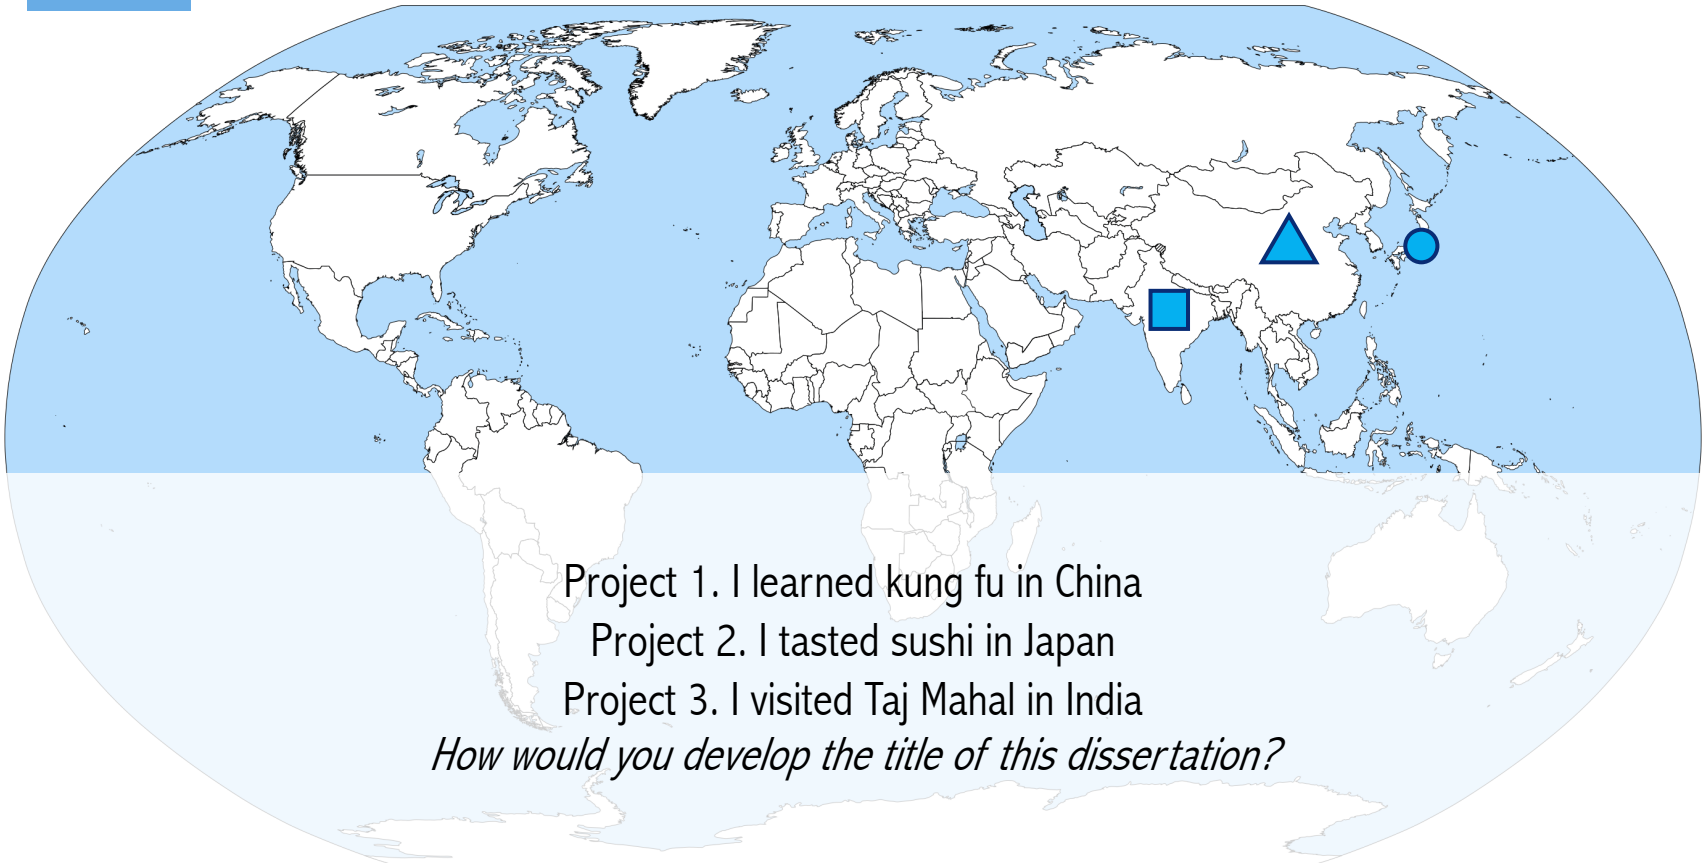

Project 1. I learned kung fu in China

Project 2. I tasted sushi in Japan

Project 3. I visited Taj Mahal in India

How would you develop the title of this dissertation?

The Contribution to the Underlying Science

- **I. Seek for common things and don't go too high level**
- Project 1. Fine-tuning Convolutional Neural Networks for Biomedical Image Analysis: Actively and Incrementally (2017)
  - *How to select important data to annotate?*
- Project 2. UNet++: A Nested U-Net Architecture for Medical Image Segmentation (2018)
  - *How to improve AI's capability of disease segmentation?*
- Project 3. Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis (2019)
  - *How to extract visual representation from images?*

The Contribution to the Underlying Science

- **I. Seek for common things and don't go too high level**
- Project 1. Fine-tuning Convolutional Neural Networks for Biomedical Image Analysis: Actively and Incrementally (2017)
  - *How to select important data to* annotate*?*
- Project 2. UNet++: A Nested U-Net Architecture for Medical Image Segmentation (2018)
  - *How to improve AI's capability of disease segmentation* given existing annotation*?*
- Project 3. Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis (2019)
  - *How to extract visual representation from* unannotated *images?*
- Annotation seems to be the key concept of my dissertation

The Contribution to the Underlying Science

- I. Seek for common things and don't go too high level
- II. Develop an intriguing story as introduction
- The introduction chapter includes the key concept definition, knowledge gap, goal, hypothesis, contribution, and so-what
- I suggest to discuss and revise the story with your advisors and peers multiple times
  - *The goal is to make your parents understand the story*
  - *An example . . .*

Imaging data account for about 90% of all healthcare data

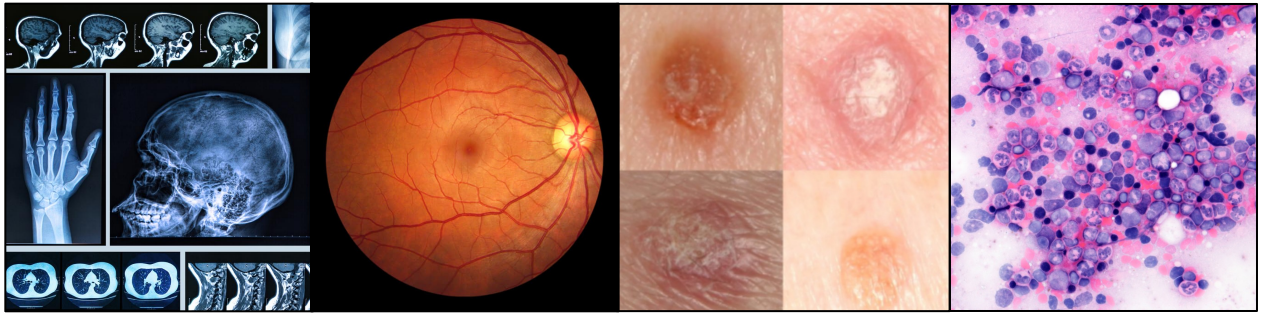

Radiology Ophthalmology Dermatology Pathology

Deep Learning has ushered in a revolution in medical imaging

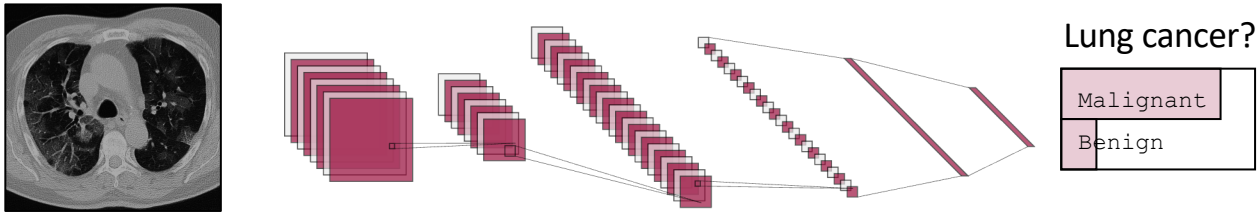

Input image Hidden layers Output

1. "The Digital Universe Driving Data Growth in Healthcare." published by EMC with research and analysis from IDC (12/13)
2. LeCun, Yann, et al. "Deep learning." Nature, 2015.

“key concept definition”

CT

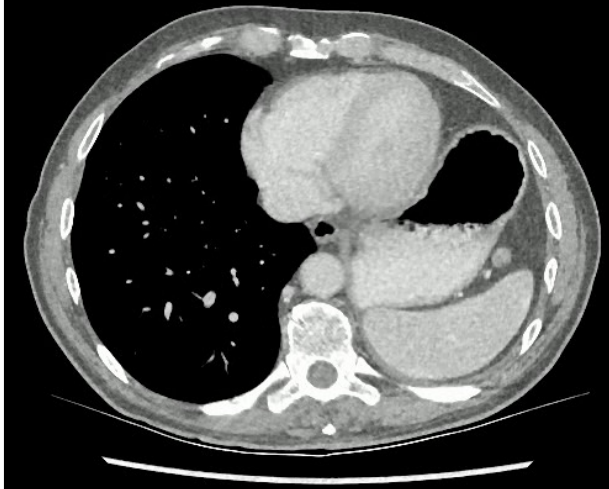

Ground Truth
annotated by human experts

Deep Learning

Liver

Liver tumor

1. Bajpai, Shivam. "Pre-Trained Models for nnUNet." Master diss., Arizona State University, 2021.
2. Zhou, Zongwei, et al. "Unet++: A nested u-net architecture for medical image segmentation." DLMIA, 2018.
3. Zhou, Zongwei, et al. "Unet++: Redesigning skip connections to exploit multiscale features in image segmentation." TMI, 2019.

To match human diagnostic precision, deep learning requires enormous annotation cost.

- **1,511,400** radiologist-annotated CT images for pancreatic cancer detection (*15 years to create*)
- **42,290** radiologist-annotated CT images for lung cancer diagnosis
- **129,450** dermatologist-annotated images for skin cancer classification

The FELIX Project at JHU

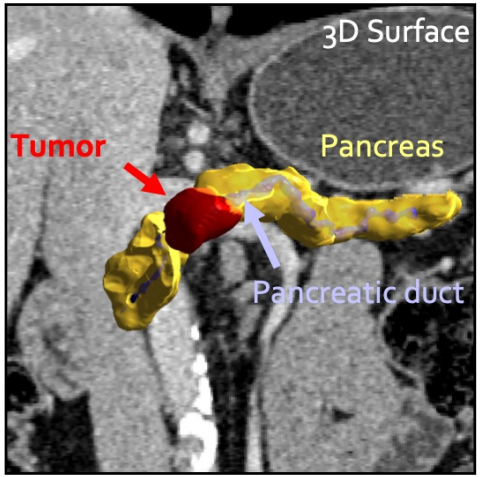

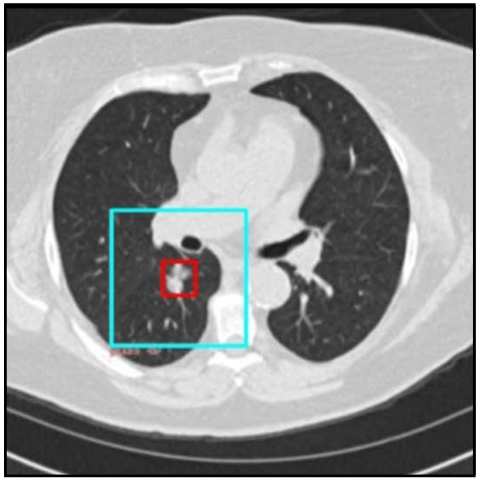

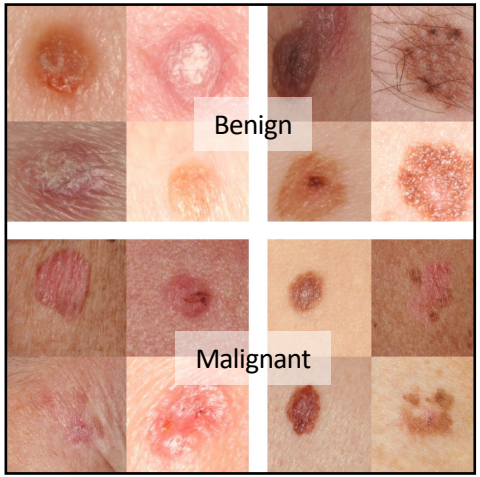

1. Xia, Yingda, et al. "The FELIX project: Deep networks to detect pancreatic neoplasms." medRxiv, 2022.
2. Ardila, Diego, et al. "End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography." Nature medicine 25.6 (2019): 954-961.
3. Esteva, Andre, et al. "Dermatologist-level classification of skin cancer with deep neural networks." Nature, 2017.

To match human diagnostic precision, deep learning requires enormous annotation cost.

- **1,511,400** radiologist-annotated CT images for pancreatic cancer detection (*15 years to create*)
- **42,290** radiologist-annotated CT images for lung cancer diagnosis
- **129,450** dermatologist-annotated images for skin cancer classification

"knowledge gap"

"How to develop annotation-efficient deep learning without such BIG annotated data?"

Significant, consider these scenarios:

- A flood of patients are waiting for imaging results during an outbreak
- Doctors do not have time to annotate every case for algorithm development
- Not many doctors have expertise for novel/rare diseases

"key concept definition"

Computer-*Aided* Diagnosis

Assisting human experts to see more patients and
to deliver more accurate diagnosis (*beyond human eye*)

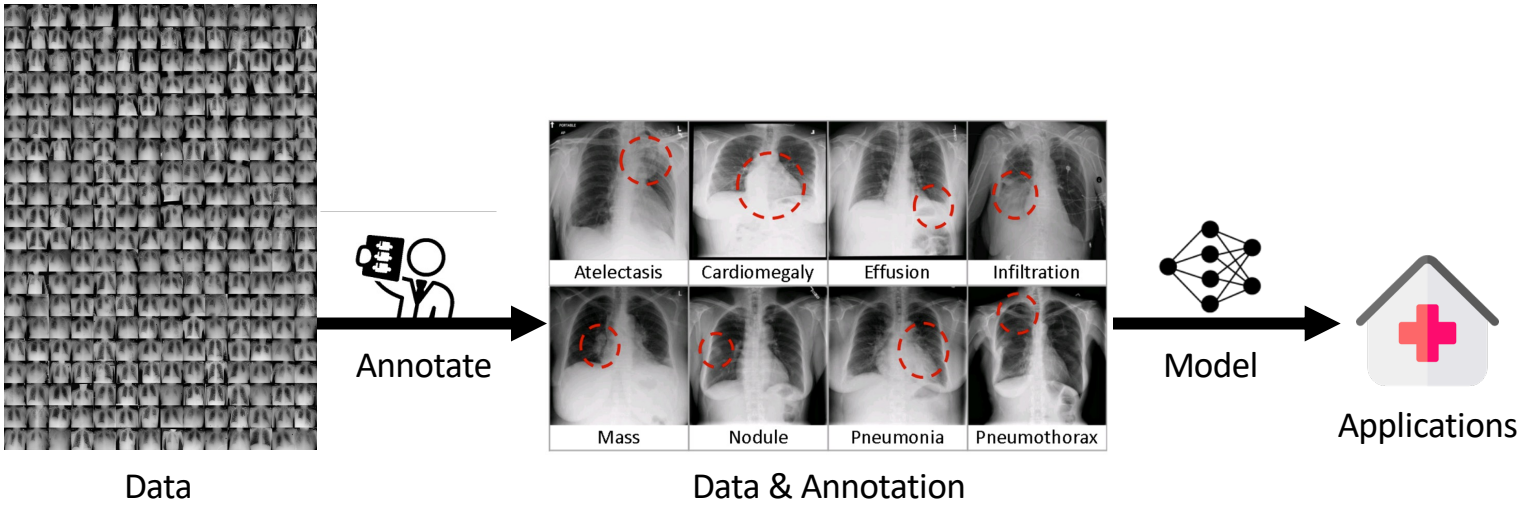

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

"goal"

This goal should cover the topics in the dissertation

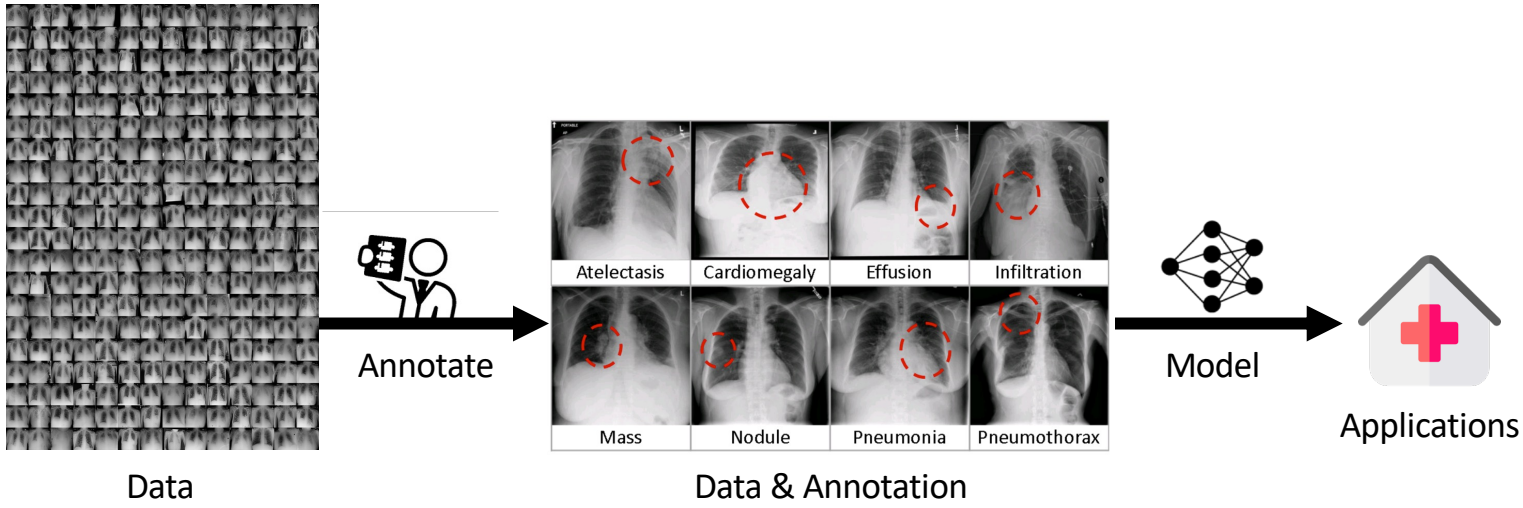

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

"Use visual strategy to ease the story"

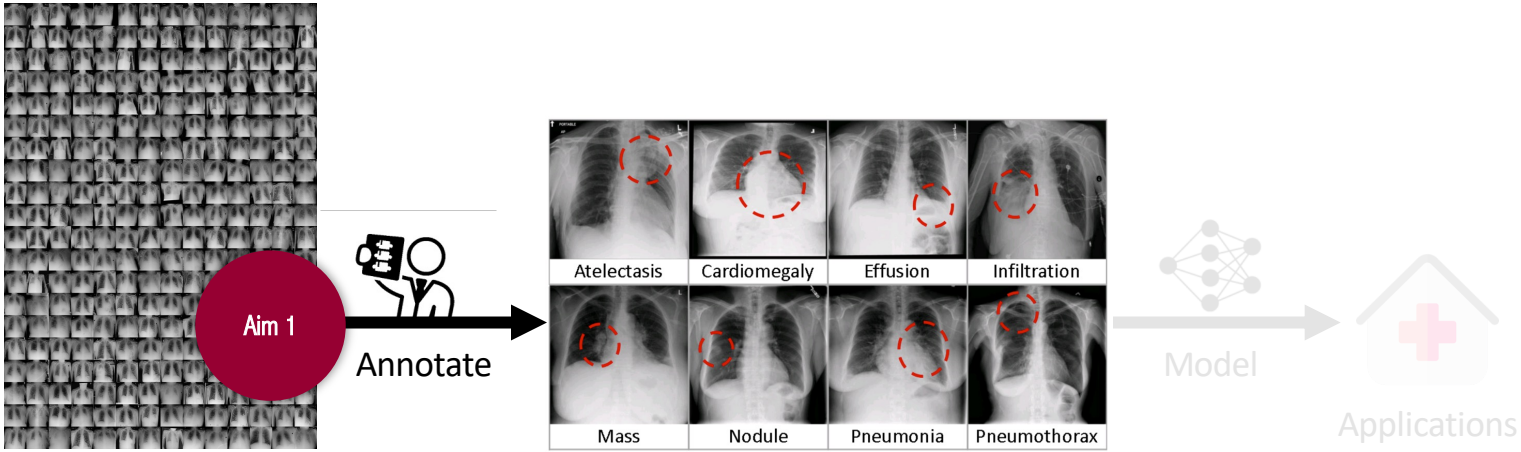

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

"Use visual strategy to ease the story"

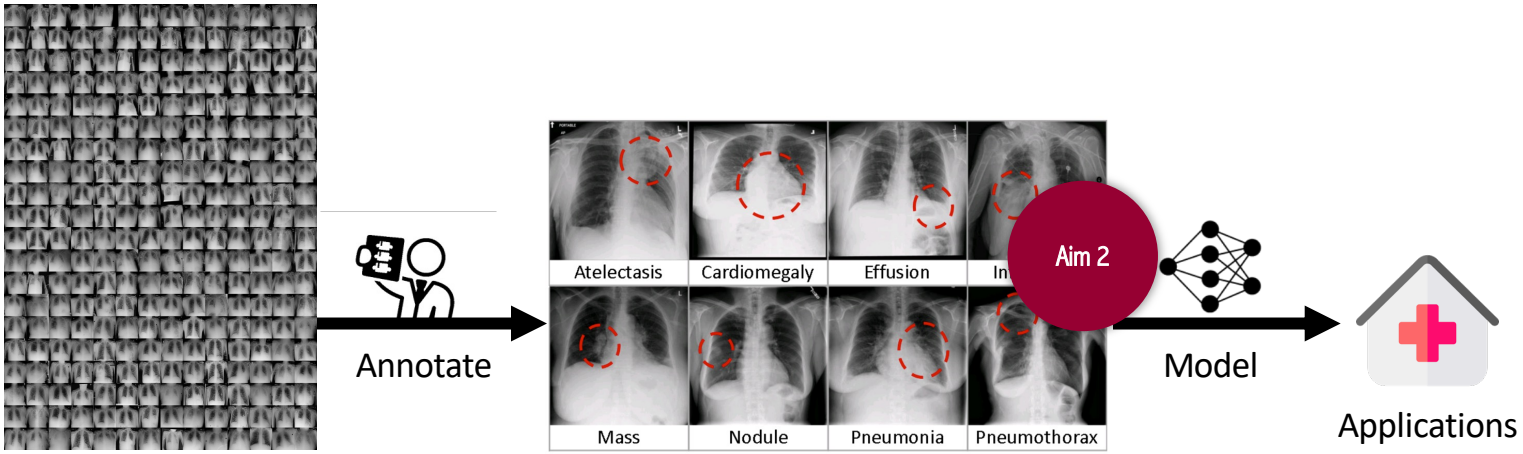

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

"Use visual strategy to ease the story"

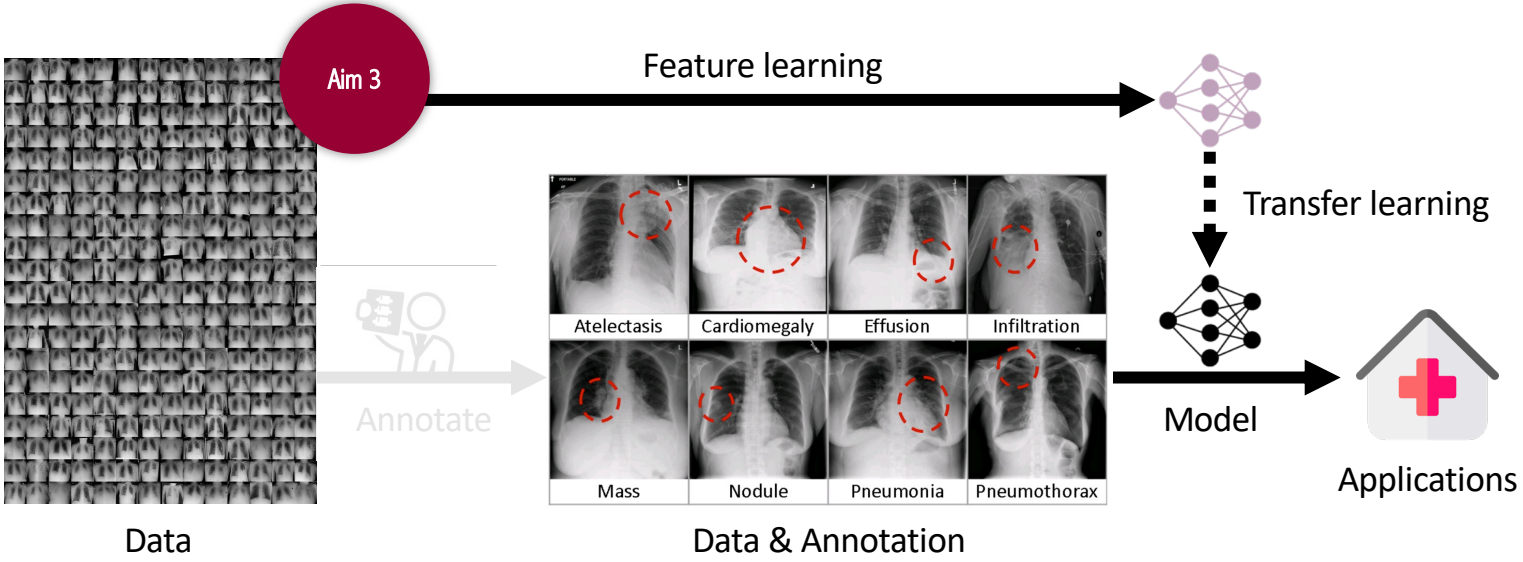

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

"Use visual strategy to ease the story"

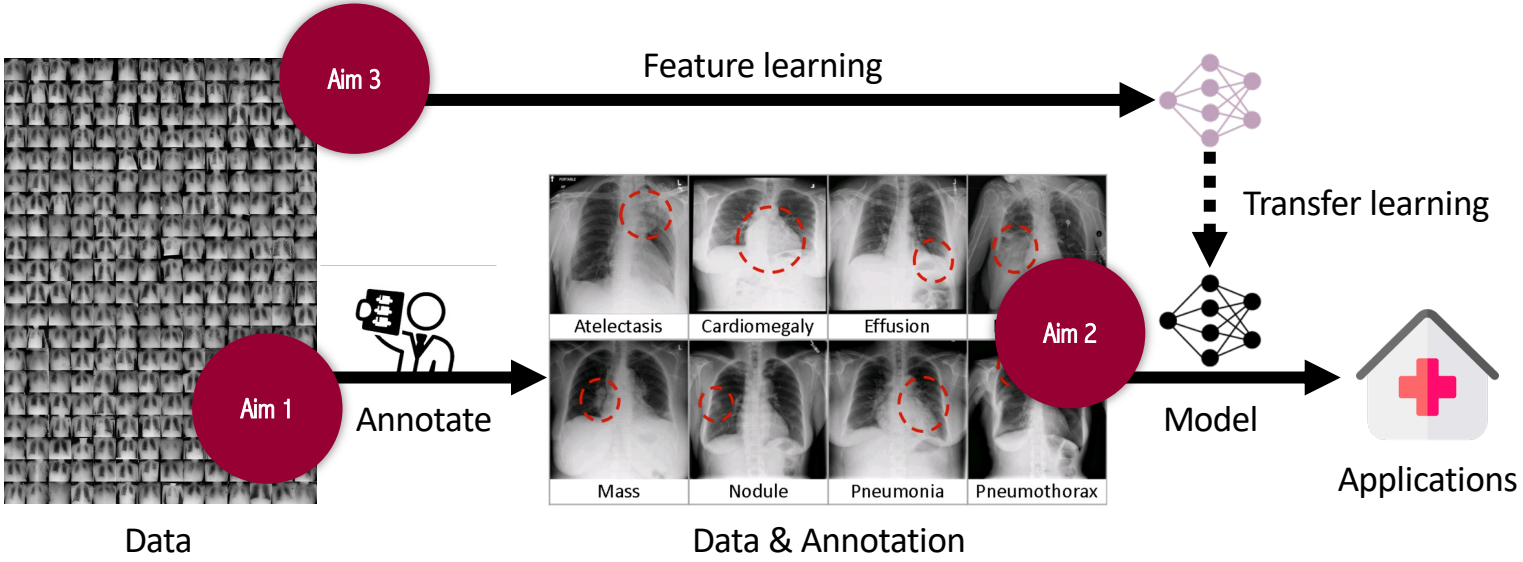

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

"Use visual strategy to ease the story"

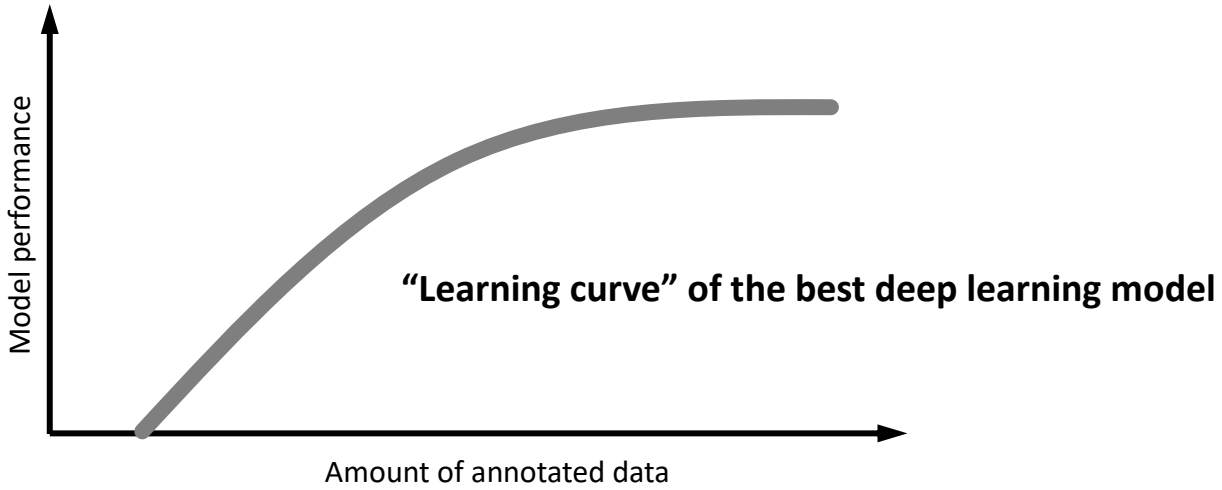

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

"Use visual strategy to ease the story"

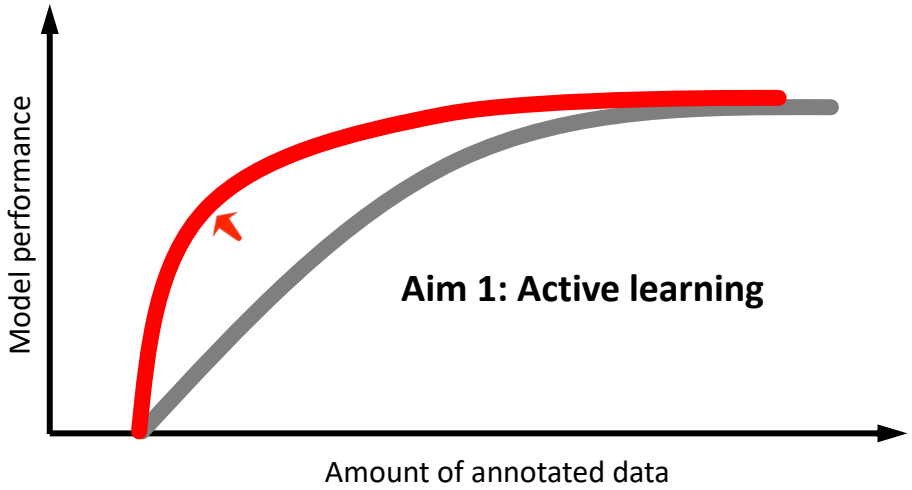

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

"Use visual strategy to ease the story"

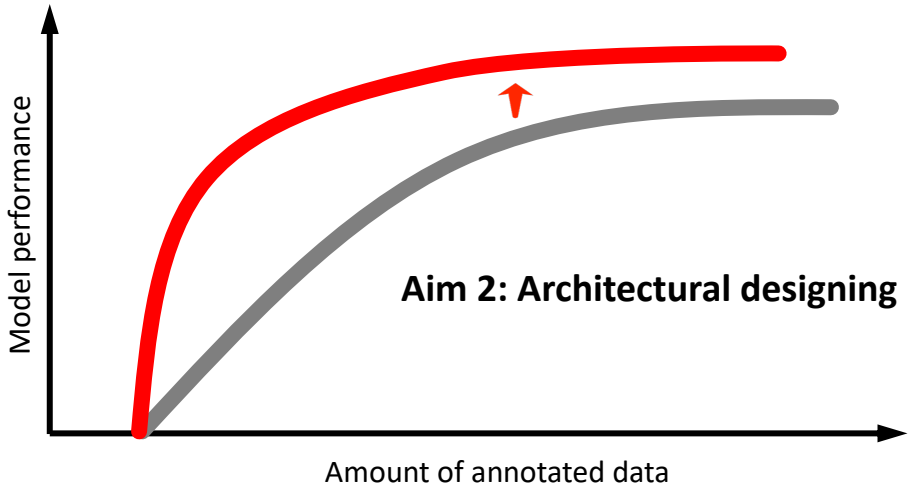

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

"Use visual strategy to ease the story"

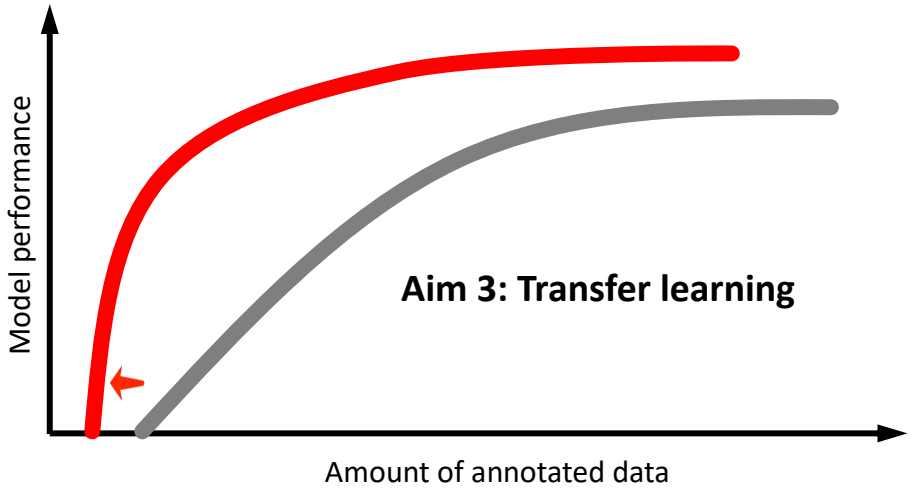

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

"Use visual strategy to ease the story"

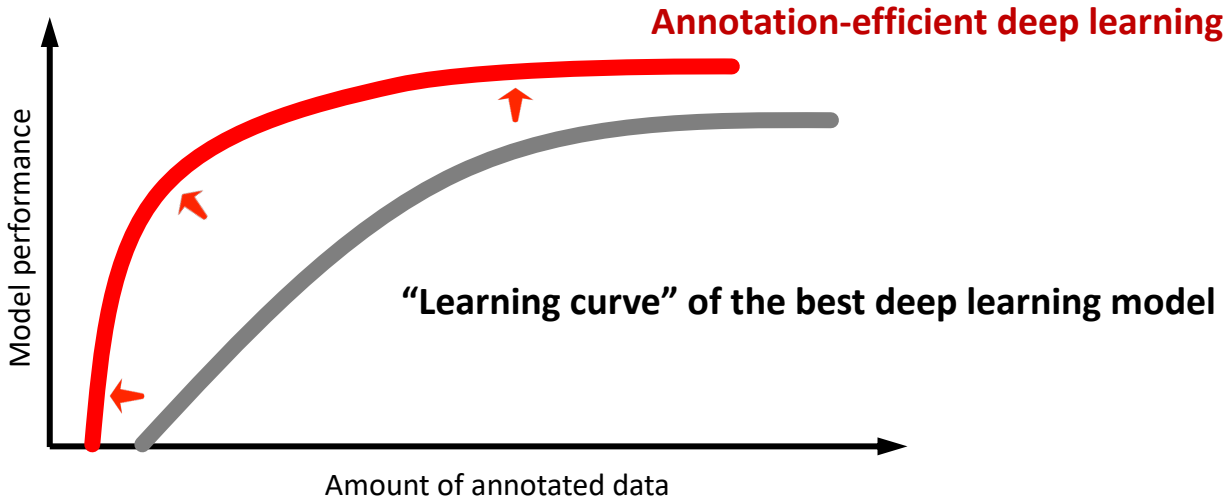

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

Hypothesis: With a small part of the dataset annotated, we can deliver deep models that approximate or even outperform those that require annotating the entire dataset.

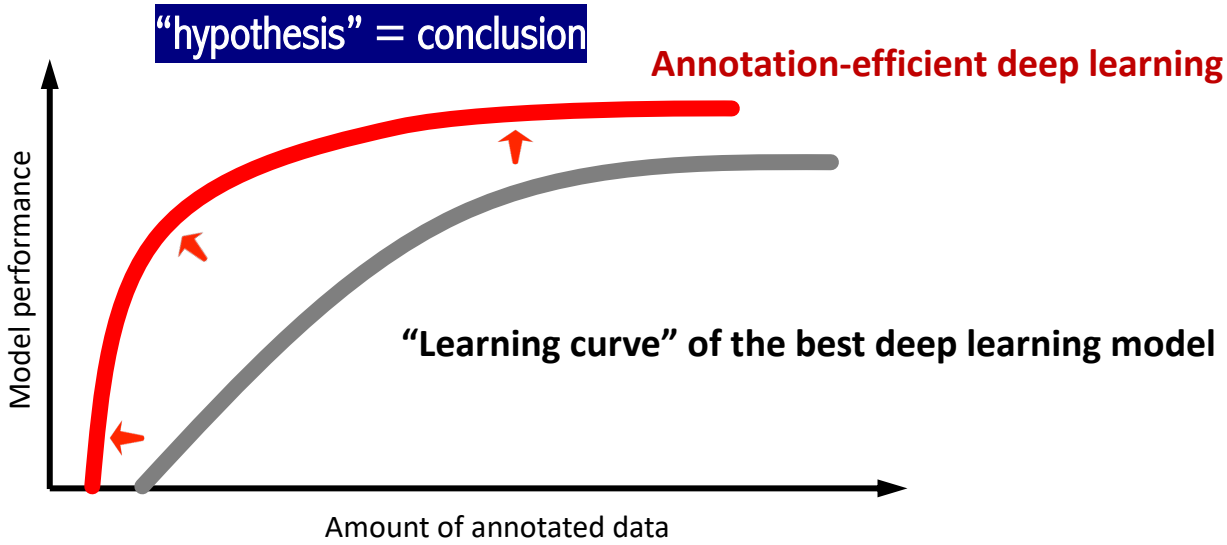

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

Hypothesis: With a small part of the dataset annotated, we can deliver deep models that approximate or even outperform those that require annotating the entire dataset.

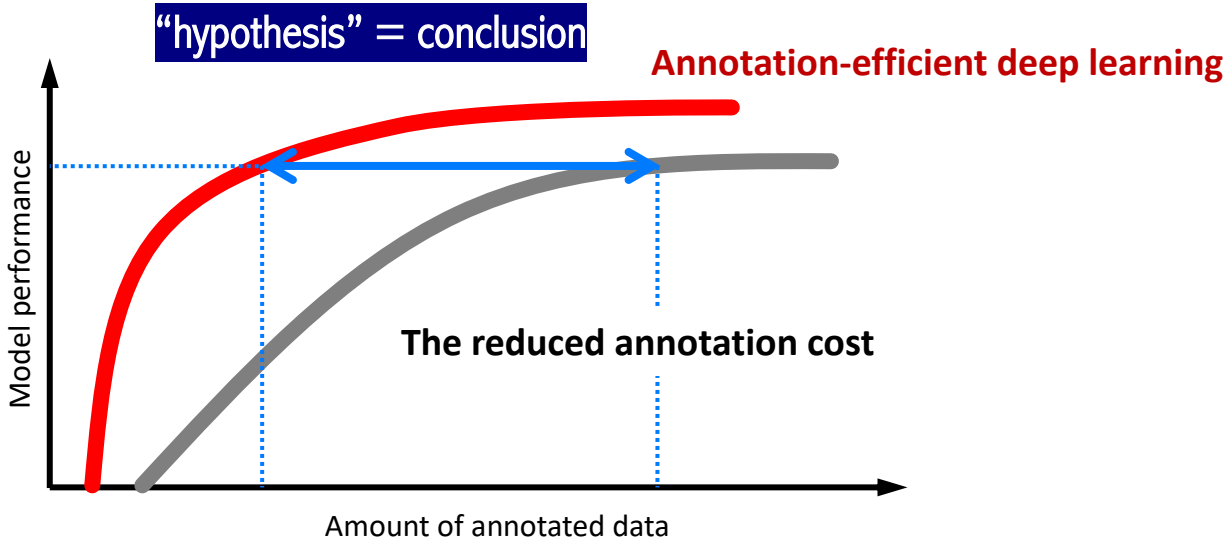

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

Hypothesis: With a small part of the dataset annotated, we can deliver deep models that approximate or even outperform those that require annotating the entire dataset.

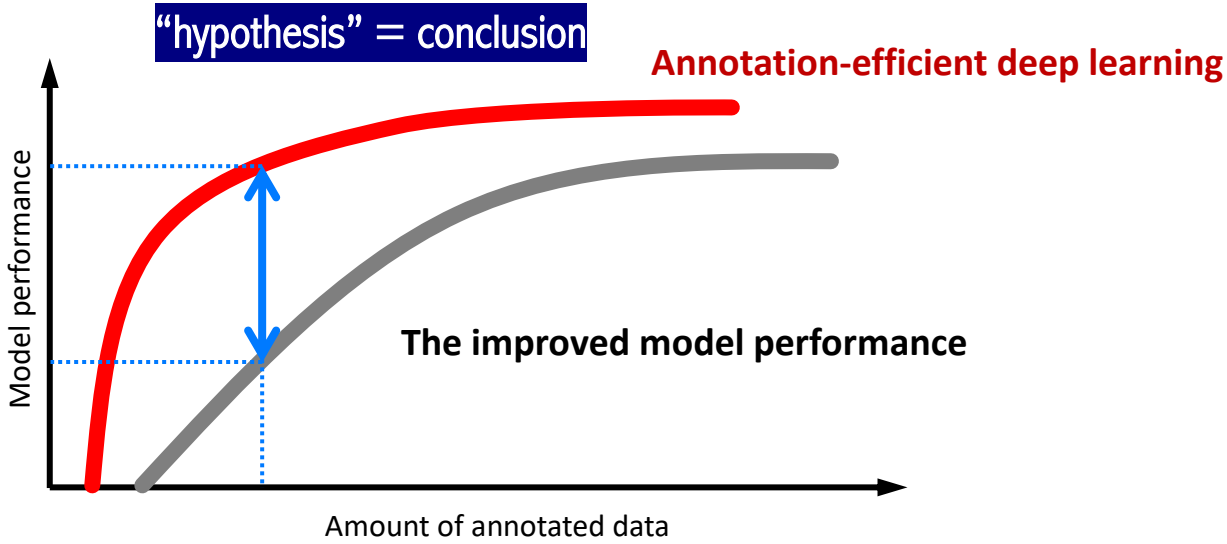

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

Hypothesis: With a small part of the dataset annotated, we can deliver deep models that approximate or even outperform those that require annotating the entire dataset.

"hypothesis" = conclusion

Chapter 1 is arguably the most important chapter in the dissertation. It will typically be written at the beginning and then revised iteratively as the document develops. I urge students to state their hypothesis within the first few pages, highlighting it in italics. That statement will guide the rest of the exposition, orient the reader to what the student has done, and lead to the assumption that success

The Contribution to the Underlying Science

- I. Seek for common things and don't go too high level
- II. Develop an intriguing story as introduction

WRITING THE DISSERTATION

The goal is to write a dissertation that others will want to read. It is important to acknowledge, however, that not all readers will understand certain technical components. Accordingly, one should provide a road map so that less technical readers can skip some of the details but still appreciate the work. I encourage students to use the first chapter to fascinate readers with the work, thereby encouraging them to keep reading so that they will learn *how* it was done!

The Contribution to the Underlying Science

- I. Seek for common things and don't go too high level
- II. Develop an intriguing story as introduction
- III. Disseminate research data for broader impact
  - What have been reported in the dissertation

> ual annotation. We have remarked our contributions in computer-aided diagnosis by supporting several aspects of medical image interpretation, including disease detection, classification, and segmentation. The experimental results on twelve distinct medical applications demonstrate that with a small part of the dataset annotated, we can deliver deep learning methods that match, or even outperform those that require annotating the entire dataset. This observation is encouraging and significant

The Contribution to the Underlying Science

- I. Seek for common things and don't go too high level
- II. Develop an intriguing story as introduction
- III. Disseminate research data for broader impact
  - What have been reported in the dissertation
  - What have been reported by other research groups—citations, applications, results

Intertwine the visual representation

Publications for Aim 2:

- Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, J. Liang, 2019. UNet++: Redesigning Skip Connections to Exploit Multi-Resolution Features in Image Segmentation. *IEEE Transactions on Medical Imaging, ranked among the most popular articles in IEEE TMI*.
- Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, J. Liang, 2018. UNet++: A Nested U-Net Architecture for Medical Image Segmentation. *Deep Learn Med Image Anal Multimodal Learn Clin Decis Support*.

U.S. Patent

- US Patent 11,164,067, Systems, Methods, and Apparatuses for Implementing a Multi-resolution Neural Network for Use with Imaging Intensive Applications Including Medical Imaging

43.9% → 58.1% (U-Net → UNet++) Covid-19 segmentation (CT) [Fan et al., IEEE TMI]	**78.6% → 82.9% (U-Net →UNet++)** Fiber tracing (corneal confocal microscopy) [Mou et al., MICCAI]	**86.5% → 89.5% (U-Net → UNet++)** Spleen segmentation (MRI) [Li et al., Computers & Graphics]

Intertwine the visual representation

Research Impacts of Aim 2: *https://github.com/MrGiovanni/UNetPlusPlus*

- We have made UNet++ open science to stimulate collaborations among the research community and to help translate the technology to clinical practice.

86.6% → 87.2% (U-Net → UNet++) SegTHOR 2019 Challenge (CT) [Zhang et al., IEEE TMI]	**90.2% → 92.0% (U-Net → UNet++)** Optic Disc & Cup Segmentation (fundus image) [Meng et al., MICCAI]	**60.3% → 71.6% (U-Net → UNet++)** Ground-glass opacity segmentation (CT) [Zheng et al., IEEE Access]
51.2% → 58.6% (U-Net → UNet++) Esophagus segmentation (CT) [Huang et al., IEEE Access]	**63.7% → 66.3% (U-Net → UNet++)** Liver tumor segmentation (CT) [Bajpai et al., Master Thesis]	**90.7% → 91.6% (U-Net → UNet++)** Heart segmentation (MRI) [Ji et al., MICCAI]

The Contribution to the Underlying Science

- I. Seek for common things and don't go too high level
- II. Develop an intriguing story as introduction
- III. Disseminate research data for broader impact
  - What have been reported in the dissertation
  - What have been reported by other groups *(citations, reproducibility, novel applications)*
  - Can the impact go beyond the field of biomedical informatics?

> Chapter *n*, then, is the final discussion that summarizes the results, the contribution of the work, and the assessment of the initial hypothesis, and lays out what lies ahead. A specific discussion of the generalizability of the method and its range of its applicability is to be encouraged. The author should bear in mind that a reader may have fully digested only Chapter 1 before reading Chapter *n*.

The Contribution to the Underlying Science

We first presented UNet++ in our DLMIA 2018 paper (Zhou *et al.*, 2018b). UNet++ has since been widely adopted by the research community, either as a strong baseline for comparison (Sun *et al.*, 2019; Fang *et al.*, 2019b,a; Meng *et al.*, 2020), or as a source of inspiration for developing newer semantic segmentation architectures (Zhang *et al.*, 2018; Chen *et al.*, 2018; Zhou *et al.*, 2018a; Wu *et al.*, 2019; Song *et al.*, 2019; Yang and Gao, 2019); it has also been utilized for multiple applications, not only for diseases/organs/tissues segmentation in biomedical images (Zyuzin and Chumarnaya, 2019; Cui *et al.*, 2019a), but also for image coloring (Di *et al.*, 2021), moon impact crater detection (Jia *et al.*, 2021), microseismic monitoring (Guo, 2021). Recently, Shenoy (2019) has independently and systematically investigated UNet++ for the task of "contact prediction model PconsC4", demonstrating significant improvement over widely-used U-Net.

The Contribution to the Underlying Science

- I. Seek for common things and don't go too high level
- II. Develop an intriguing story as introduction
- III. Disseminate research data for broader impact
  - What have been reported in the dissertation
  - What have been reported by other groups *(citations, reproducibility, novel applications)*
  - Can the impact go beyond the field of biomedical informatics?
  - Consider writing a review paper for the field

career path. And much that is written in the document, if not previously published, can be extracted and adapted for formal publication in peer-reviewed journals. For example, Chapter 2 is often a great source of content that can be edited into an influential review paper suitable for publishing.

The Contribution to the Underlying Science

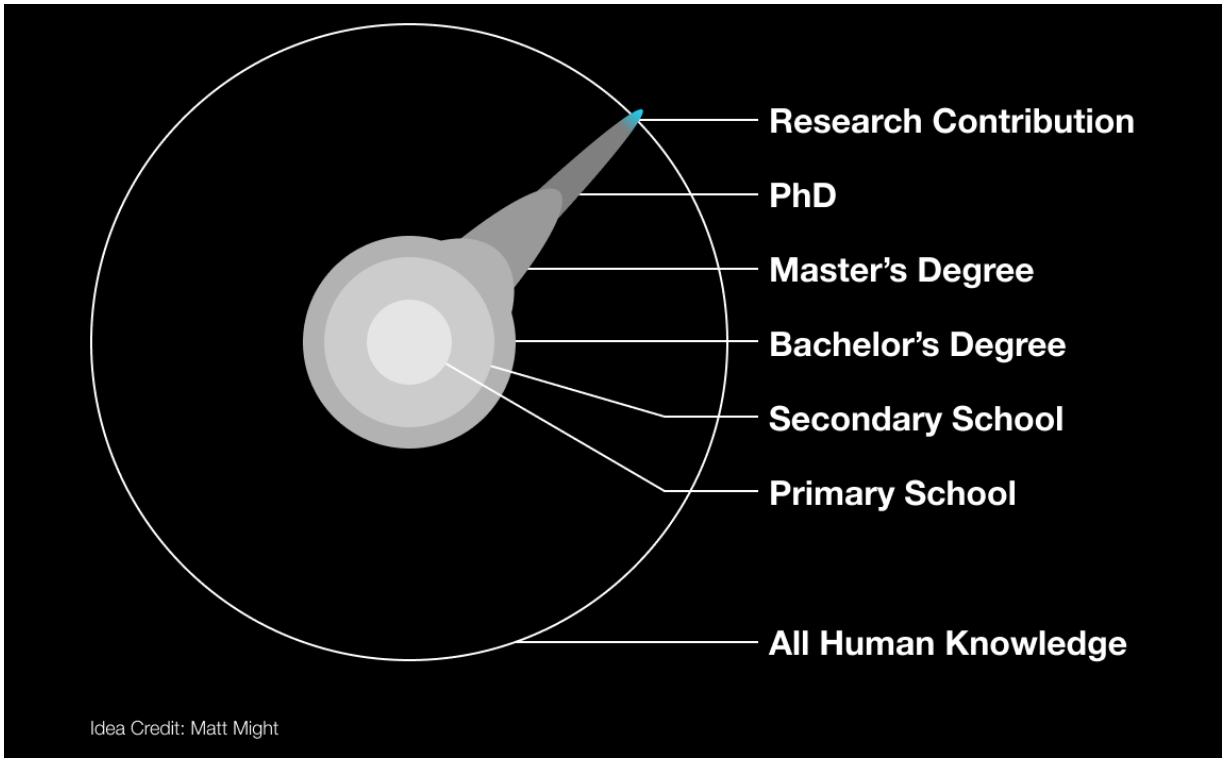

Goal: Minimize manual annotation efforts for rapid, precise computer-aided diagnosis systems

Aim 1: Acquiring necessary annotation efficiently from human experts

Aim 2: Utilizing existing annotation effectively from advanced architecture

Aim 3: Extracting generic knowledge directly from unannotated images

Interpreting Medical Images: A book chapter that overviews AI in medical image interpretation

Edward H. Shortliffe, MD, PhD

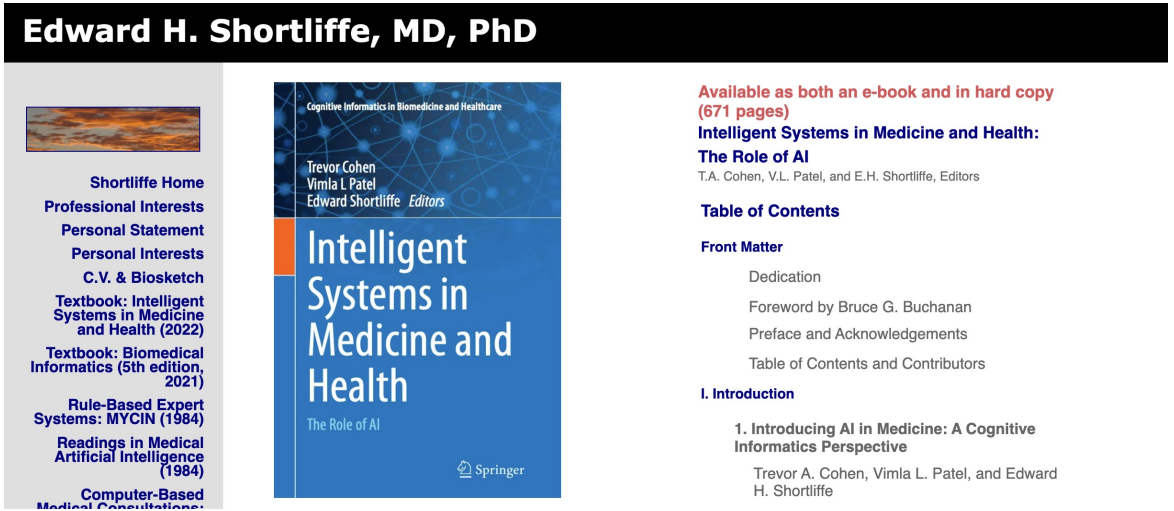

1. Zhou, Zongwei et al. “Interpreting Medical Images.” In Cognitive Informatics in Biomedicine and Healthcare. Intelligent Systems in Medicine and Health: The Role of AI. T. Cohen, V. Patel and E. Shortliffe (eds.). Springer Nature, 2022.

The Architecture of Dissertation

- I. Seek for common things and don't go too high level *(Title)*
- II. Develop an intriguing story as introduction *(Introduction)*
- III. Disseminate research data for broader impact *(Conclusion)*

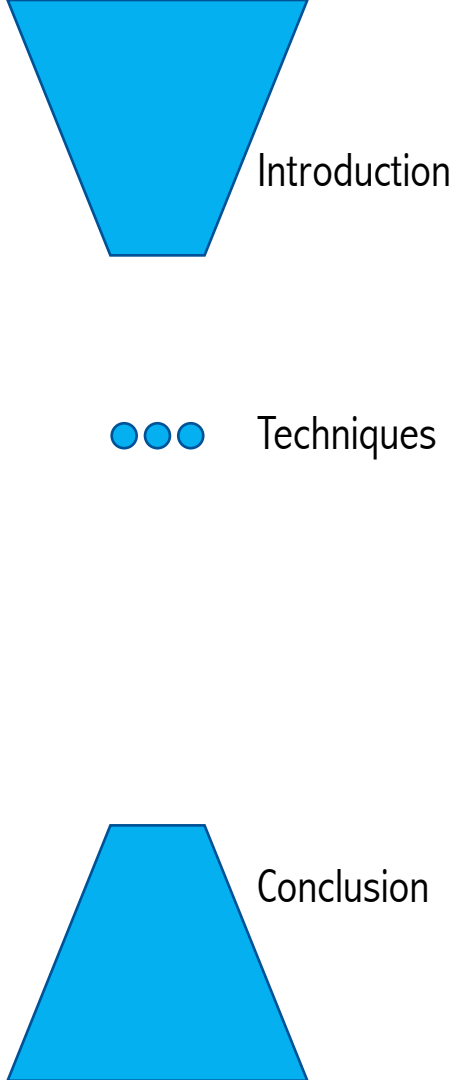

The Architecture of Dissertation

- I. Seek for common things and don't go too high level *(Title)*
- II. Develop an intriguing story as introduction *(Introduction)*
- III. Disseminate research data for broader impact *(Conclusion)*
- *How to make it better than decent?*

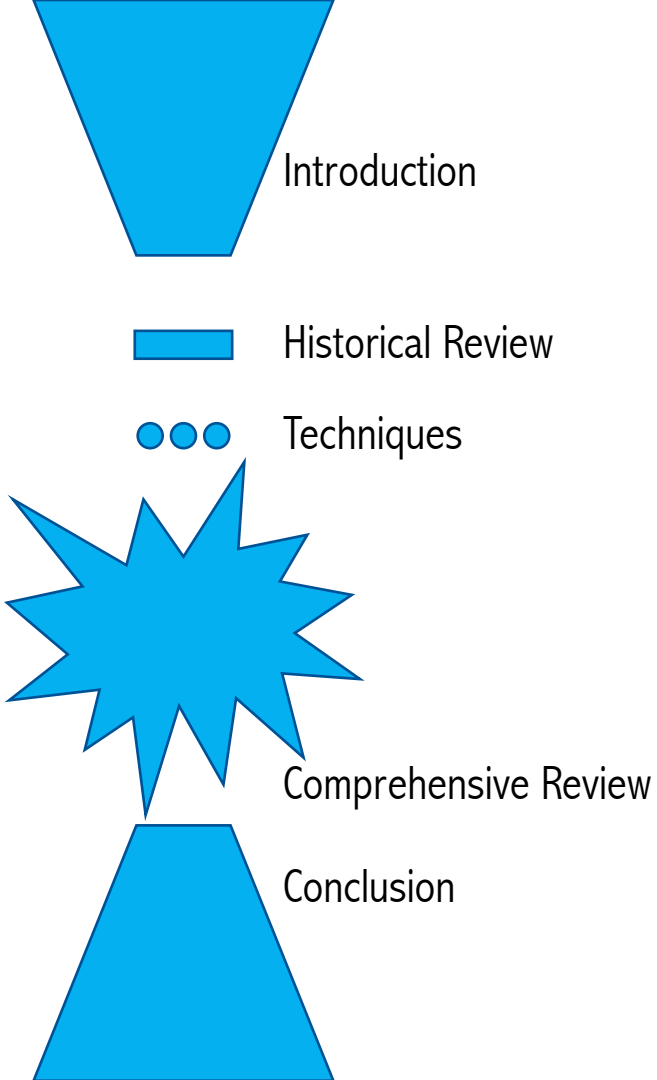

The Architecture of Dissertation

- **Historical review vs. related work**
- Both show your knowledge of the sub-field
  - Historical review is broader and long-standing

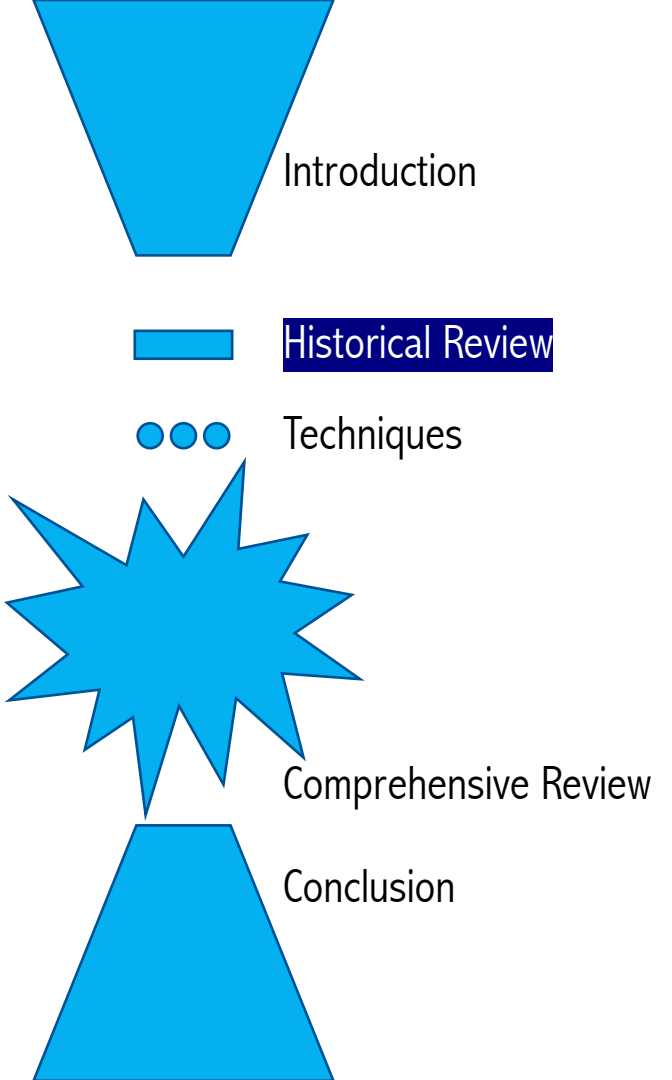

The Architecture of Dissertation

- **Historical review vs. related work**
- Both show your knowledge of the sub-field
- *Don't forget to emphasize the innovation!*

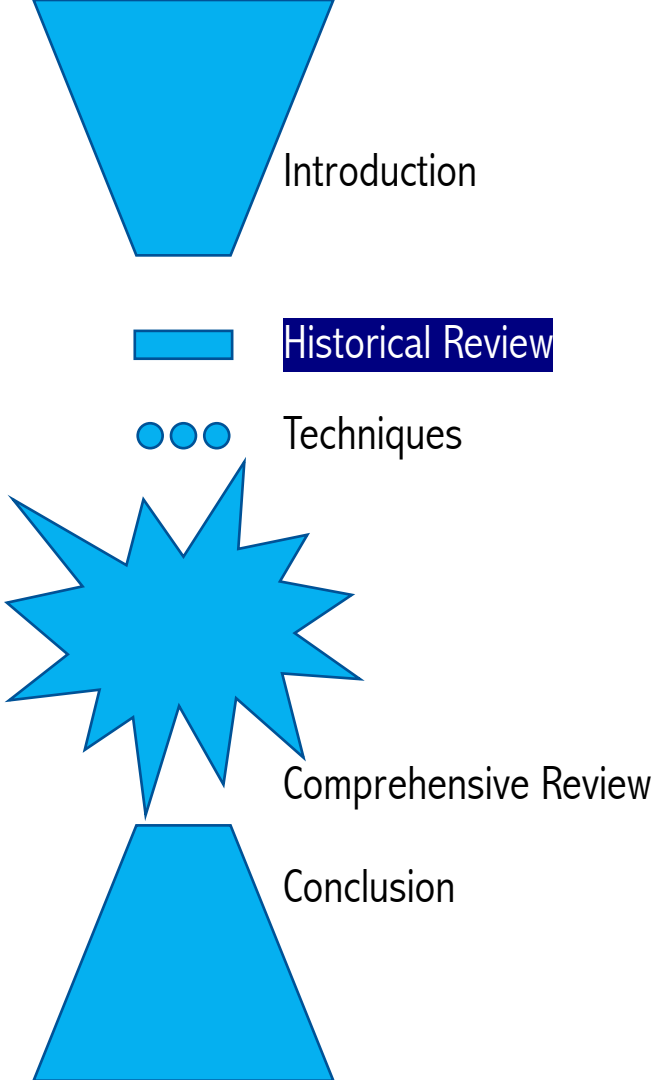

The Architecture of Dissertation

- How to build the content of three (or more) technical chapters?

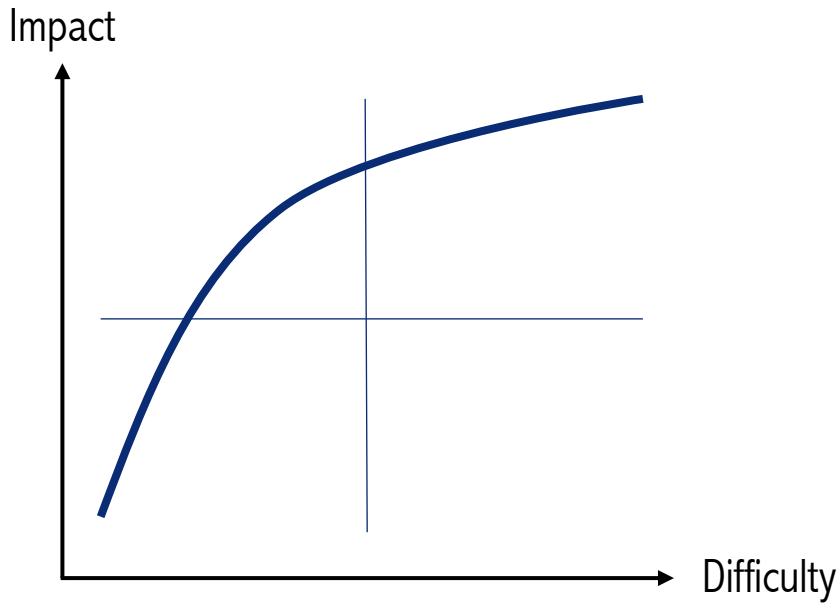

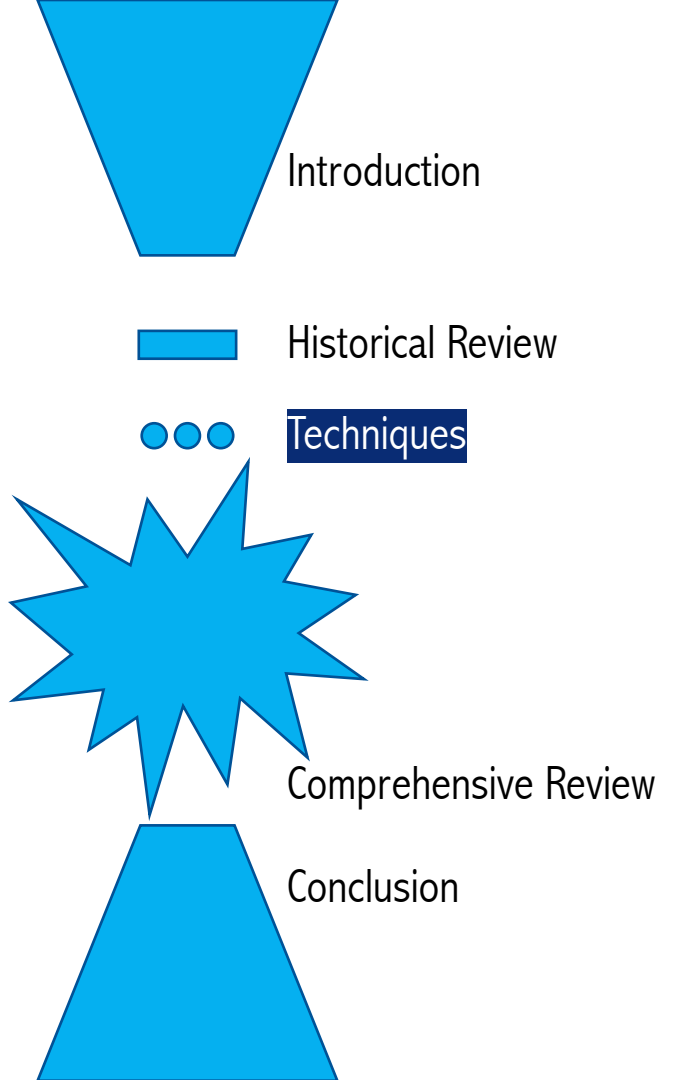

The Architecture of Dissertation

- **Make the technical chapters better**
- I used a fixed template for every technical chapter
  - Background & Motivation
  - Approach & Property
    - *No related work section; innovation is in the property section*
  - Experiment & Result
  - Discussion & Conclusion
    - *The discussion section must be insightful; usually it consists of several open questions and directs future works*
    - *Show broader impact*

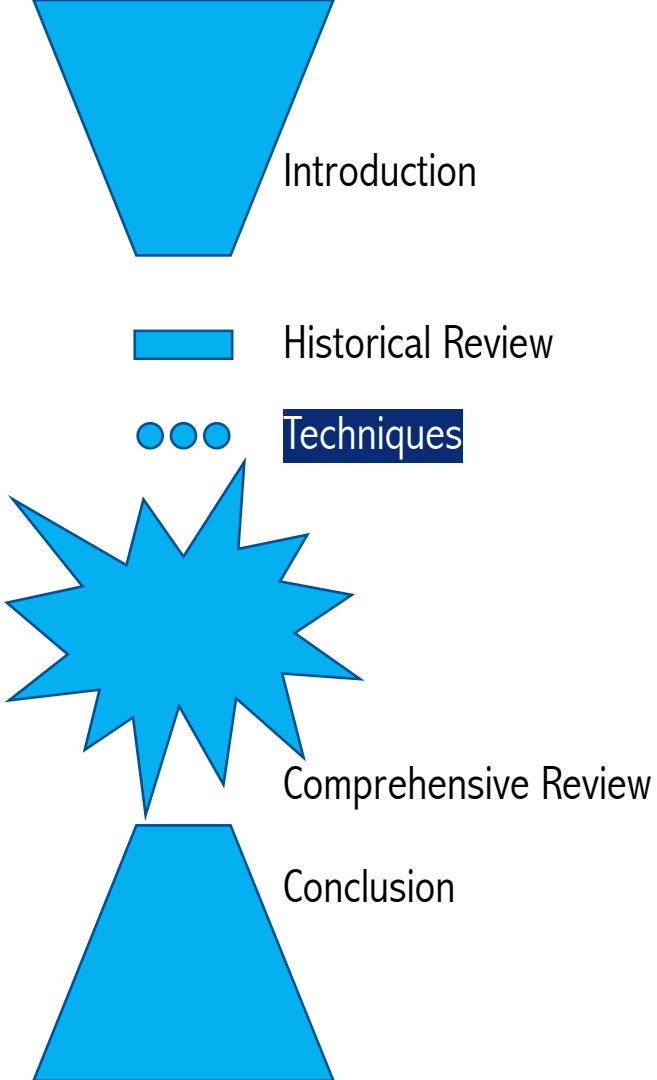

The Architecture of Dissertation

- **Bonus: Write a comprehensive review for general audience**
- Its scope is certainly broader than the focus of your dissertation
  - Dissertation title: Towards Annotation-Efficient Deep Learning for Computer-Aided Diagnosis
  - Book chapter title: Interpreting Medical Images
    - *Annotation is one of the challenges; others include reliability, usability, affordability, etc.*
    - *Computer-aided diagnosis is one of the applications; others include detection, prognosis, follow-up, surgery, reconstruction, etc.*

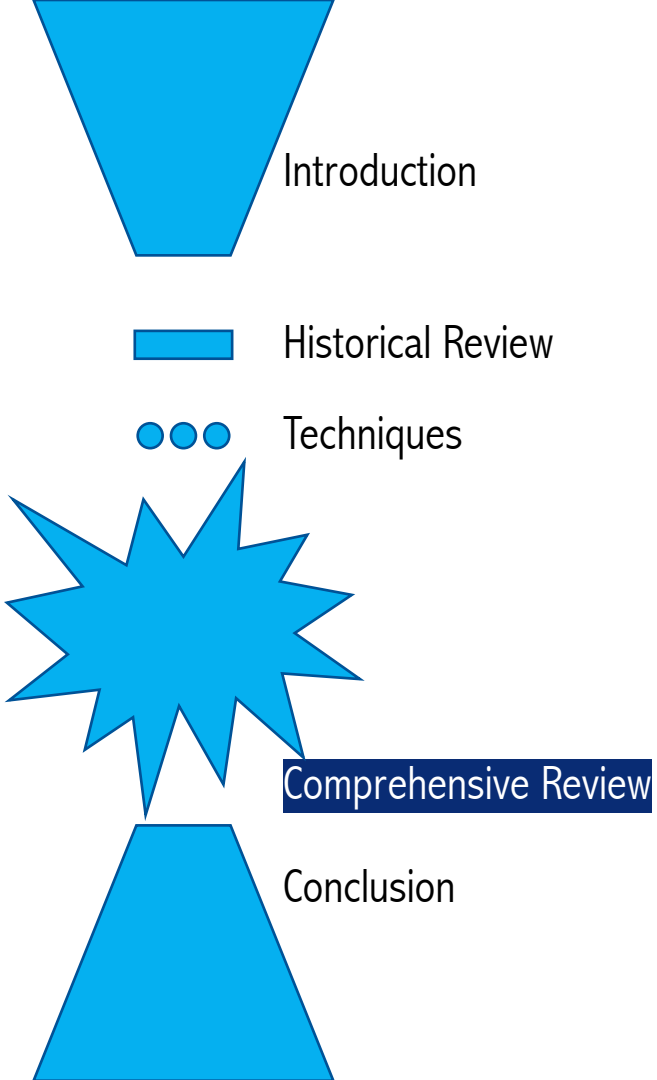

The Architecture of Dissertation

- **Bonus: Write a comprehensive review for general audience**
- Demonstrate your expertise in a relatively large field
  - Introduction to Medical Images
  - Characteristics of Medical Images
  - Historical Perspectives
  - Clinical Needs and Existing Challenges
  - Opportunities and Emerging Techniques
  - Conclusion

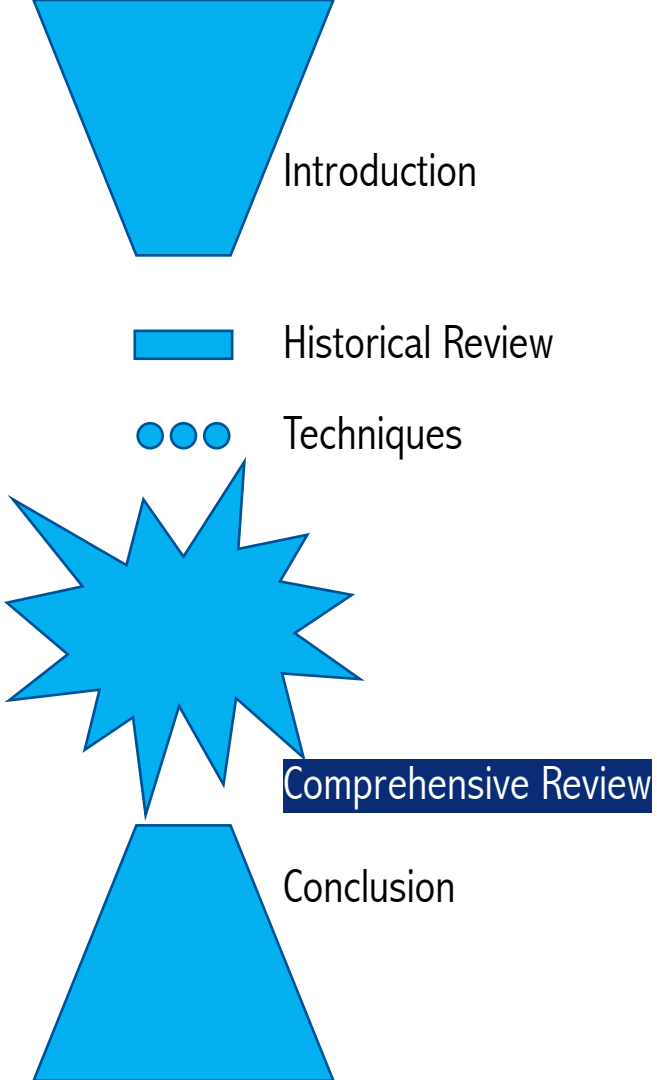

The Architecture of Dissertation

- **Bonus: Write a comprehensive review for general audience**
- Place your PhD works to the field
  - Introduction to Medical Images
  - Characteristics of Medical Images
  - Historical Perspectives
  - Clinical Needs and Existing Challenges
    - *My works are addressing annotation efficiency problem*
  - Opportunities and Emerging Techniques
    - *I developed three techniques that are significant to some degree*
  - Conclusion

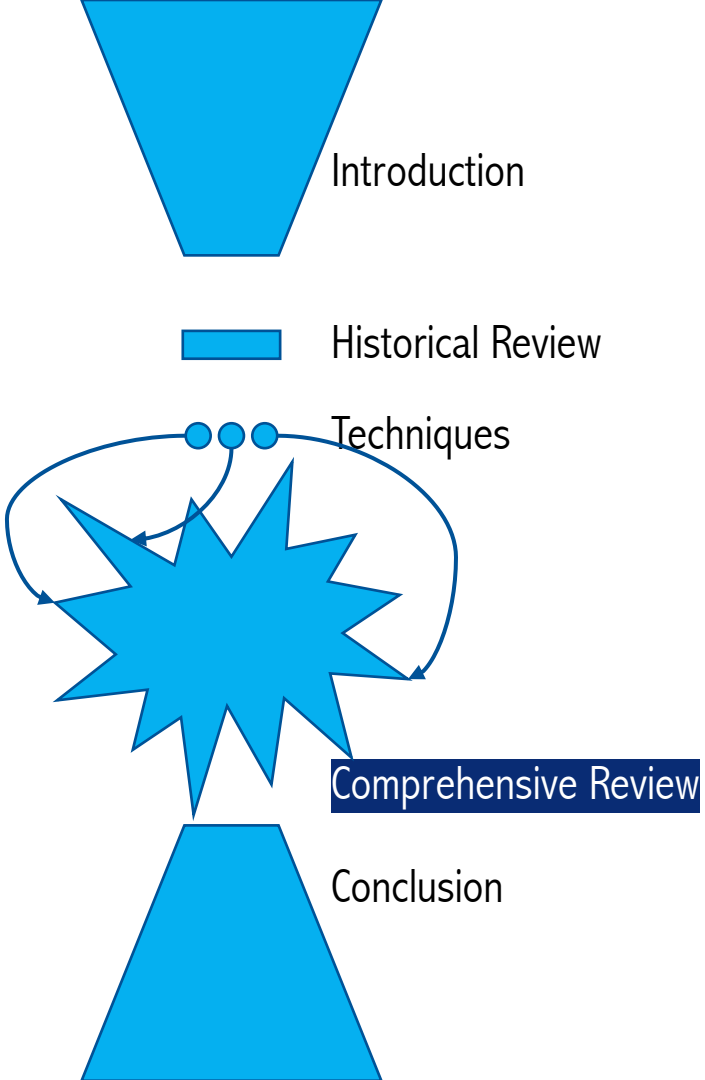

Summary

- In the final year, the value of carefully writing and organizing the dissertation is much bigger than
  - publishing other papers
  - graduating early from the university
- Check many good examples and guidelines, discuss with your advisors, friends, and parents
  - making an engaging story
  - widely disseminating research findings
  - expending your expertise in the field
- Consider AMIA Doctoral Dissertation Award
  - *AMIA Doctoral Dissertation Award is the highest honor for BMI PhD students. Each BMI program can only nominate one candidate to compete for this award in a specific year.*

Reference

Thanks & Questions?

- Shortliffe, Edward H. The organization and content of informatics doctoral dissertations. Journal of the American Medical Informatics Association 23, no. 4 (2016): 840-843.
- Cohen, Trevor A., Vimla L. Patel, and Edward H. Shortliffe, eds. Intelligent Systems in Medicine and Health: The Role of AI. Springer Nature, 2022.

- Zhou, Zongwei. Towards annotation-efficient deep learning for computer-aided diagnosis. PhD diss., Arizona State University, 2021.
- Defense: https://www.youtube.com/watch?v=Q2qRbtzLE8Y
- GitHub: https://github.com/MrGiovanni/Dissertation