



Literature Search

CSCI 16: Intro to the CS Research Process
3rd Meeting, Thursday, January 5, 2023

Kelly Shaw

Attribution



Except where otherwise noted, ERSP and all associated resources were developed by Christine Alvarado, UC San Diego, Department of Computer Science and Engineering.

These slides were originally jointly developed by Christine Alvarado and Mai ElSherief. They have been adapted by Kelly Shaw for CSCI 16 at Williams College.

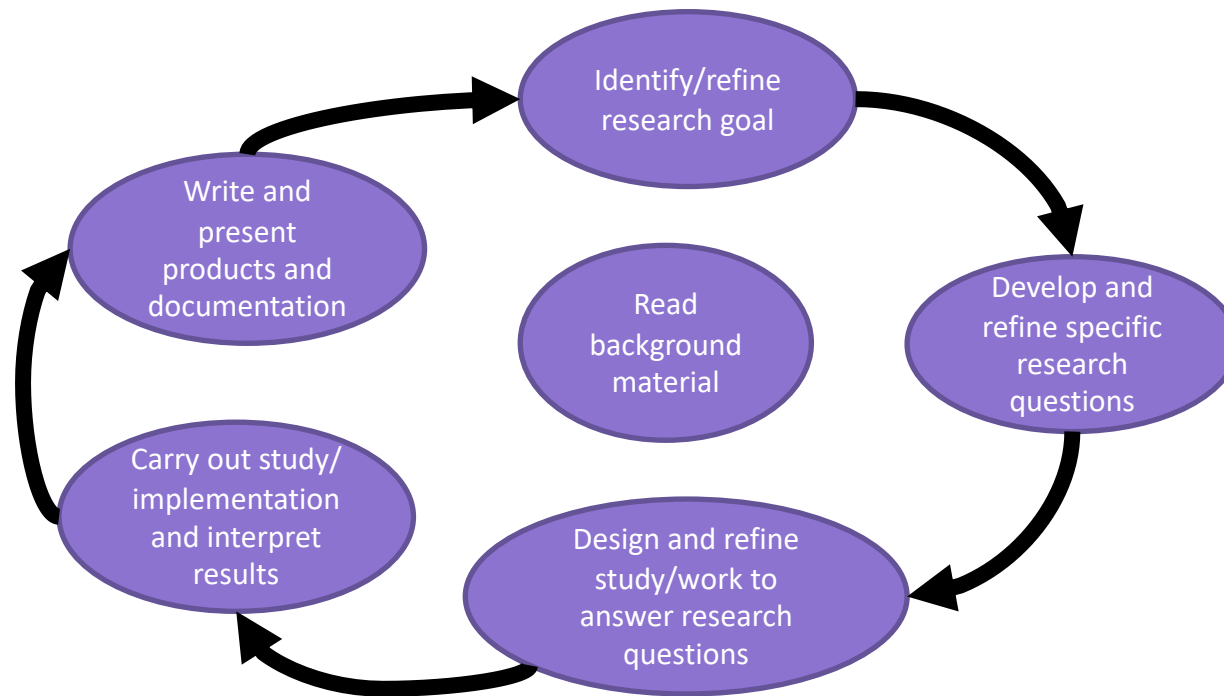


This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

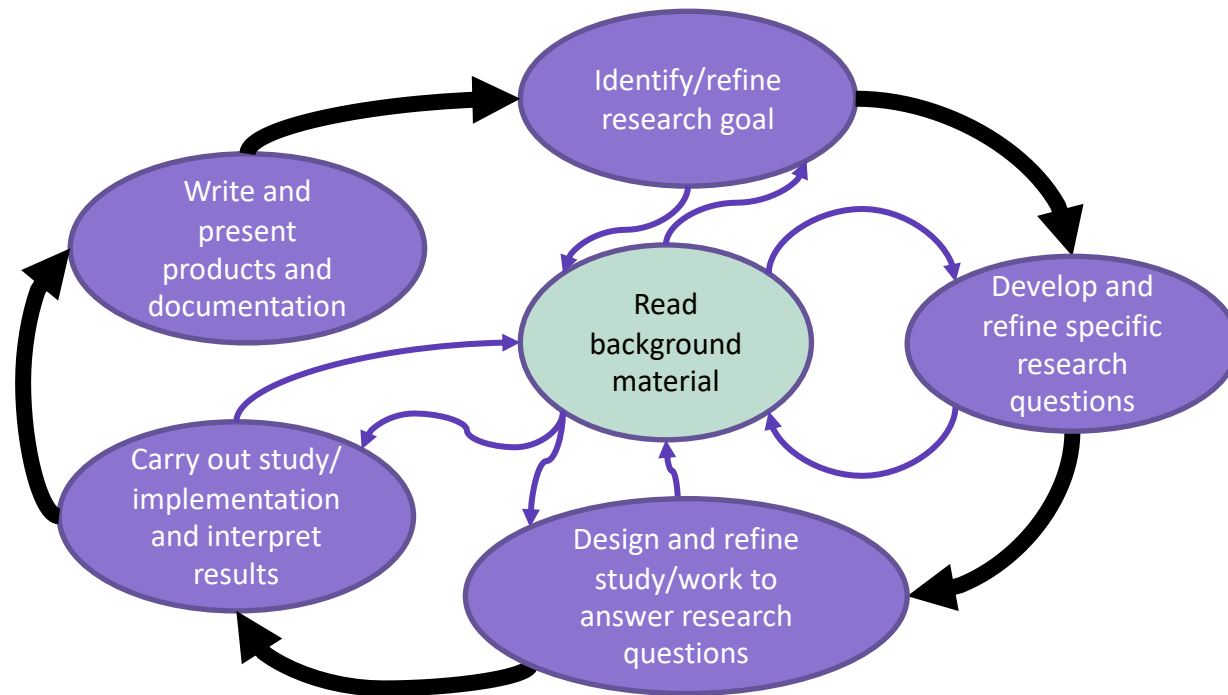
Today's Plan

- Discussion of “Fine-grained Recognition in the Wild”
- Describe important aspects of doing a literature search in computer science
- Perform a literature search in computer science

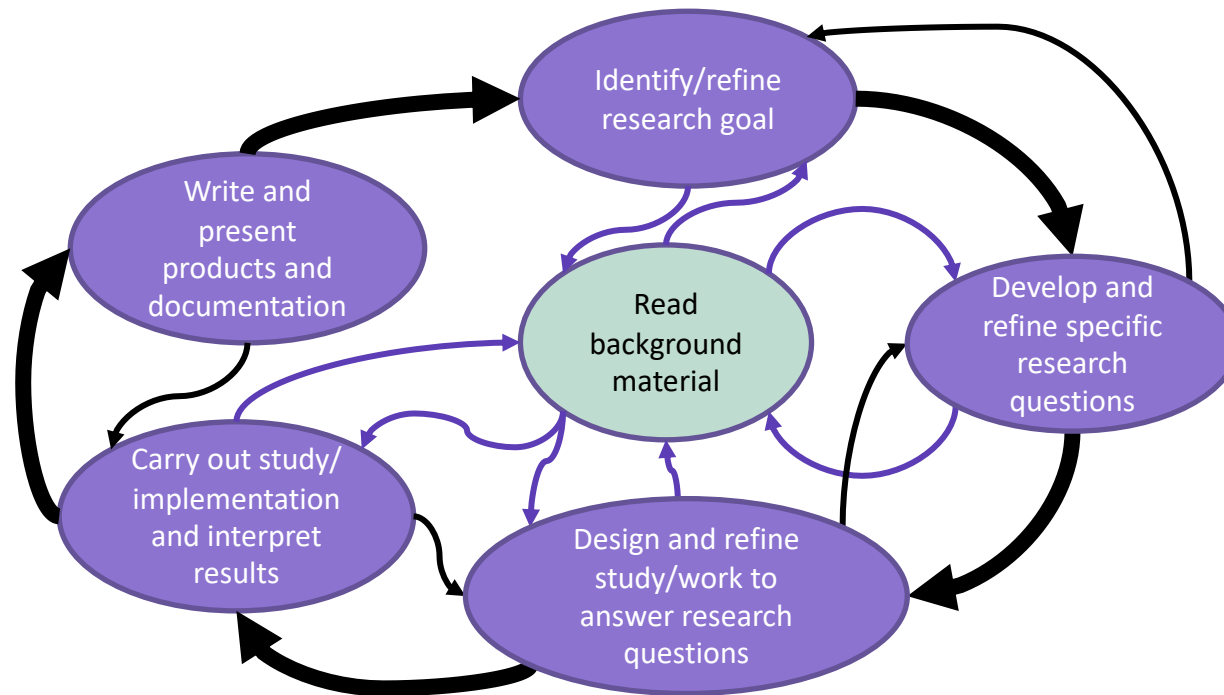
Overview of the research process



Overview of the research process



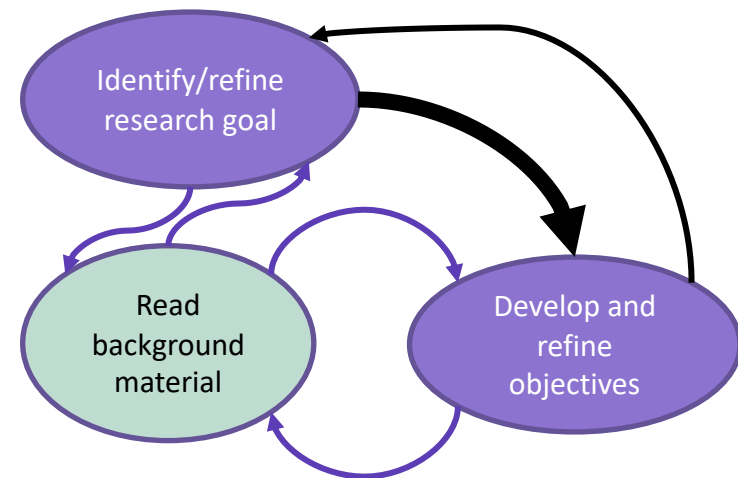
Overview of the research process



Where we are

The purpose of background reading at these two stages:

- **Develop a sense of the research field**
- **Determine which specific objectives are still open problems**
- Figure out what research problems others have addressed
- Learn about technical material you need to understand the research



A Literature Search

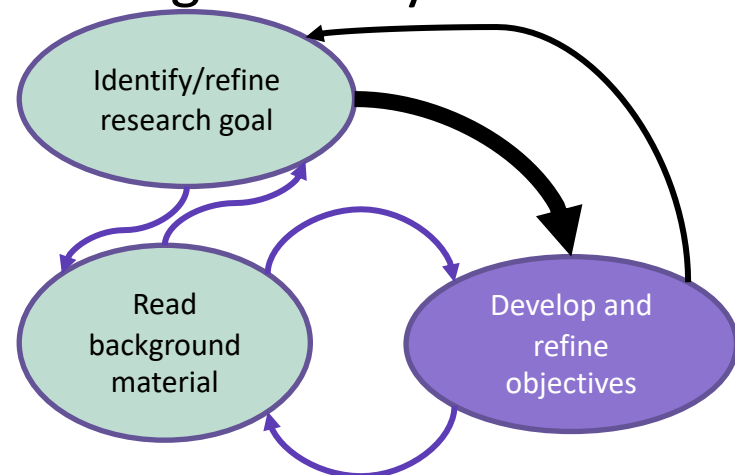
- You just read a paper and you have a brilliant idea about a follow-up study/improvement. But has it already been done? Will people care? How do you know?


Literature searching

- When working on a new research problem, you want to expand out from a paper you have read to get a broader sense of the field. This is a fundamental skill in doing research, and usually the first thing that any researcher in a new area does.
- But how to do this? Google?

Activity 1 (5 mins):

- Identify the passage in the Fine-grained Recognition in the Wild paper that states the problem the paper is solving.





In this work, we study fine-grained domain adaptation as a step towards overcoming the dataset shift between easily acquired annotated images and the real world. To our knowledge, adaptation has not been studied in the fine-grained setting where it is especially expensive to obtain image annotations. In this scenario, many of our categories may be related to one another in some known hierarchical way. For example, multiple distinct car varieties may share the same body type or the same make.



Fine-grained Recognition in the Wild: A Multi-Task Domain Adaptation Approach

Timnit Gebru Judy Hoffman Li Fei-Fei
CS Department Stanford University
`{tgebru, jhoffman, feifeili}@cs.stanford.edu`

- What other solutions previously existed to solve this problem? Where could you look to find the answer?


Fine-grained Recognition in the Wild: A Multi-Task Domain Adaptation Approach

Timnit Gebru Judy Hoffman Li Fei-Fei
CS Department Stanford University
`{tgebru, jhoffman, feifeili}@cs.stanford.edu`

- What other solutions previously existed to solve this problem? Where could you look to find the answer?

2. Related Work

- Activity 2 (~15 mins, individual first):
 - Read the related work section
 1. Identify the cited papers that relate most closely to this work
 2. Annotate how each paper relates to this work.
 - If you have time, find these papers online.



fine-grained setting. Our work builds on [49]’s method to show that attribute level softlabel transfer and domain confusion significantly boost performance in this scenario.

limited available data. Works such as [33] have used large-scale noisy data to train state-of-the-art fine-grained recognition models. However, these models are unlikely to generalize to real world photos because they are trained with images derived from field guides or product shots. Similarly,

Our method to enforce consistency between attribute and class predictions is similar in spirit to a number of works exploiting label structure [14, 12]. [14] uses Hierarchy and Exclusion (HEX) graphs to encapsulate semantic relations between pairs of labels. We use a KL divergence loss between predicted label distributions instead of hard constraints.

fine-grained setting. Our work builds on [49]'s method to show that attribute level softlabel transfer and domain confusion significantly boost performance in this scenario.

limited available data. Works such as [33] have used large-scale noisy data to train state-of-the-art fine-grained recognition models. However, these models are unlikely to generalize to real world photos because they are trained with images derived from field guides or product shots. Similarly,

Our method to enforce consistency between attribute and class predictions is similar in spirit to a number of works exploiting label structure [14, 12]. [14] uses Hierarchy and Exclusion (HEX) graphs to encapsulate semantic relations between pairs of labels. We use a KL divergence loss between predicted label distributions instead of hard constraints.

Seems to be most closely related to the current system. Current work seems to build off this previous work.

[33] focuses on fine-grained recognition. Provides a basis for the problem studied here: that current approaches fail on real-world data.

These works use a similar structure for enforcing attribute consistency.

Finding publication venues

- [12] L.-C. Chen, A. G. Schwing, A. L. Yuille, and R. Urtasun. Learning deep structured models. In *ICML*, pages 1785–1794, 2015. 2
- [14] J. Deng, N. Ding, Y. Jia, A. Frome, K. Murphy, S. Bengio, Y. Li, H. Neven, and H. Adam. Large-scale object classification using label relation graphs. In *European Conference on Computer Vision*, pages 48–64. Springer, 2014. 2
- [33] J. Krause, B. Sapp, A. Howard, H. Zhou, A. Toshev, T. Duerig, J. Philbin, and L. Fei-Fei. The unreasonable effectiveness of noisy data for fine-grained recognition. *arXiv preprint arXiv:1511.06789*, 2015. 1, 2
- [49] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko. Simultaneous deep transfer across domains and tasks. In *International Conference on Computer Vision (ICCV)*, 2015. 2, 3, 4, 5, 6, 7, 8

**What are ICML/
ECCV/ICCV/arXive??**

Where was the original paper published?



This ICCV paper is the Open Access version, provided by the Computer Vision Foundation.
Except for this watermark, it is identical to the version available on IEEE Xplore.

<https://ieeexplore.ieee.org/document/8237413>

Published in: [2017 IEEE International Conference on Computer Vision \(ICCV\)](#)

Date of Conference: 22-29 Oct. 2017

INSPEC Accession Number: 17453066

Date Added to IEEE Xplore: 25
December 2017

DOI: [10.1109/ICCV.2017.151](https://doi.org/10.1109/ICCV.2017.151)

► ISBN Information:

Publisher: IEEE

Electronic ISSN: 2380-7504

Conference Location: Venice, Italy

Where was the original paper published?

Activity 3 (10 mins, group):

- Find the top five publication venues for Computer Vision and Pattern recognition
- Gauge the quality of this paper and related work based on where they were published

https://scholar.google.com/citations?view_op=top_venues&hl=en&vq=eng_computervisionpatternrecognition

What are the primary publication venues?

Google Scholar

Top publications

Categories > Engineering & Computer Science > Computer Vision & Pattern Recognition

Publication	h5-index	h5-median
1. IEEE/CVF Conference on Computer Vision and Pattern Recognition	299	509
2. IEEE/CVF International Conference on Computer Vision		
3. European Conference on Computer Vision		
4. IEEE Transactions on Pattern Analysis and Machine Intelligence		
5. IEEE Transactions on Image Processing		
6. Pattern Recognition		
7. IEEE Computer Society Conference on Computer Vision and Pattern F		
8. International Journal of Computer Vision		

So far we searched backward looking at previously published related work

Main take aways are:


- **How to identify related work**
- **How to gauge the quality of the paper you read and that of related work based on publication outlets**

h-index

From Wikipedia, the free encyclopedia

This article is about the index of scientific research impact. For the economic measure, see [Herfindahl index](#).

The ***h*-index** is an [author-level metric](#) that attempts to measure both the [productivity](#) and [citation impact](#) of the [publications](#) of a [scientist](#) or scholar. The index is based on the set of the scientist's most cited papers and the number of citations that they have received in other publications. The index can also be applied to the productivity and impact of a [scholarly journal](#)^[1] as well as a group of scientists, such as a department or university or country.^[2] The index was suggested in 2005 by [Jorge E. Hirsch](#), a physicist at [UCSD](#), as a tool for determining [theoretical physicists'](#) relative quality^[3] and is sometimes called the *Hirsch index* or *Hirsch number*.

- 
- What other solutions have arisen since, that solve the same problems? Where could I look to figure this out?

Deep Visual-Semantic Alignments for Generating Image Descriptions

Andrej Karpathy Li Fei-Fei
Department of Computer Science, Stanford University
{karpathy, feifeili}@cs.stanford.edu

- What other solutions have arisen since, that solve the same problems? Where could I look to figure this out?

[Fine-grained recognition in the wild: A multi-task domain adaptation approach](#)


T Gebu, [J Hoffman](#), [L Fei-Fei](#) - Proceedings of the IEEE ... 2017 - [openaccess.thecvf.com](#)

While fine-grained object recognition models are unlikely to accurately need additional annotated images. However, sources annotated images for many classes towards overcoming

Activity 4 (~15 mins, individual):

- (8 minutes) Use Google Scholar to find the most relevant 3-5 papers that build off of this work. You will need to read abstracts to determine which work is most relevant.

Click here



☆ Cited by 131 Related articles All 11 versions ↗

Showing the best result for this search. [See all results](#)

Sources for finding related work

- ✓ Search backward: What was done before?
- ✓ Search forward: What has built on this work?
- ✓ Search the conference/journal: What's being done and classified as similar by the community?
- Search Databases (Google Scholar, ACM DL, IEEE Xplore, Web of Science, etc) – Must be on campus VPN
- Search the authors: What else are these same authors doing?
 - Aside: Look up your advisor and their research group to get a better understanding of what they do!

Going broader

- Searching databases (must be on campus or use VPN)
 - ACM Digital Library (CSE-specific)
 - IEEE Xplore (CSE and Electrical Engineering)
 - Web of Science (All of STEM)

Activity 5 (if time):

- Choose one of the databases above.
 - As a group, use it to look for work related to this paper.
 - What keywords did you use?
 - How easy or difficult was it to manage the results?
 - What are the advantages/disadvantages of using a broad database search?

Lit Search Summary Points (Quality, not Quantity)

- Quality can be measured by:
 - How many others cite the paper/who cites the paper
 - What conference/journal it was published in
 - Who wrote it and how many (quality) papers have they written before/what organization are they at
 - Your own assessment of the work (harder, since you're not an expert). Read the abstracts only.
- Know the important journals/conferences:
 - ACM, IEEE
 - Most fields have a “flagship” conference and/or journal. Figure out what it is.
- Search backward references *and forward*
- Tools:
 - Google scholar, databases (ACM Digital Library, IEEE Explore), Web of Science, raw Google (with care!)