

Deep Learning Project Design

Outline

1. Project Expectations
 - a. Does my project meet expectations?
 - b. FAQs
2. Picking a Project idea
 - a. Inspiration
 - b. How to read a research paper
3. Proposal, milestones, and final report
 - a. Due dates, expectation, logistics
 - b. Support

Outline

1. Project Expectations

- a. Does my project meet expectations?
- b. FAQs

2. Picking a Project idea

- a. Inspiration
- b. How to read a research paper

3. Proposal, milestones, and final report

- a. Due dates, expectation, logistics
- b. Support

Project Expectation

Open ended. Anything related to deep learning!

Project Expectation

Open ended. Anything related to deep learning!

Completed in groups of 1,2, or 3 people

- More people = higher expectations

Project Expectation

Open ended. Anything related to deep learning!

Completed in groups of 1,2, or 3 people

- More people = higher expectations

Two project options:

-Applications: Pick a new problem, and apply a deep learning solution!

- Your own data, data from another scientific field

-Model: Pick a standard problem, and find a new solution.

- Kaggle challenges etc

Project Expectation

- Title, Author(s)
- Abstract
- Related Work
- Methods
- Experiments
- Discussion
- Supplementary Material (optional)

<https://courses.cs.washington.edu/courses/cse493g1/4wi/project/>

Does my project meet expectations?

Checklist

- I am using deep learning
- I am not just download a git repository and running “train”
- I am trying to understand my results via analysis

Does my project meet expectations?

Strong projects might...

- Propose a novel variant of a technique (which takes a lot of effort)
- Adapt an existing technique to a totally new problem (which takes a lot of effort)

Weaker projects might...

- Spend several weeks collecting/cleaning data rather than testing hypotheses
- Clone an existing repo and do minimal stitching to make it work for a Kaggle competition

FAQs

Does my project need novelty?

No! Novelty is one way to fulfill the requirements but not the only way.

Do I need to get state of the art performance?

Not at all. Most research contributions do not lead to state of the art performance

How else can I show effort?

Compare and contrast different methods, show multiple design iterations leading to improved accuracy, show creative design choices to tackle a new dataset, and more

How do I show proper analysis?

Do your best to answer “why” in your discussion! What kinds of mistakes is your model making? Where is it improving? Why does the loss/accuracy curve look the way it does? etc

FAQs

Can I change my project after the proposal?

Yes! This is just a first idea for a direction

Can I change my project after the milestone?

We do not encourage this. If you feel you have to, come and speak to us about why.

Successful Past Projects

For Reports: <https://cs229.stanford.edu/proj2021spr/>

For Posters: <https://cs229.stanford.edu/proj2020spr/>

Outline

1. Project Expectations

- a. Does my project meet expectations?
- b. FAQs

2. Picking a Project idea

- a. Inspiration
- b. How to read a research paper

3. Proposal, milestones, and final report

- a. Due dates, expectation, logistics
- b. Support

Picking a project idea

One way to have novelty in your project is to take inspiration from things you care about outside of the course.

- Interested in healthcare? Robotics? Animals? Finance? Sports? Solve a problem that you are uniquely positioned to solve!

Picking a project idea

One way to have novelty in your project is to take inspiration from things you care about outside of the course.

- Interested in healthcare? Robotics? Animals? Finance? Sports? Solve a problem that you are uniquely positioned to solve!

Practical considerations:

1. Data - does data exist for your problem
2. Code - does an implementation exist?
3. Compute - do not train a 100 Billion parameter language model

Improve 1 part of the pipeline

- Create dataset to test a problem
- Design a model for this problem
- Compare to other models and analyze pros and cons

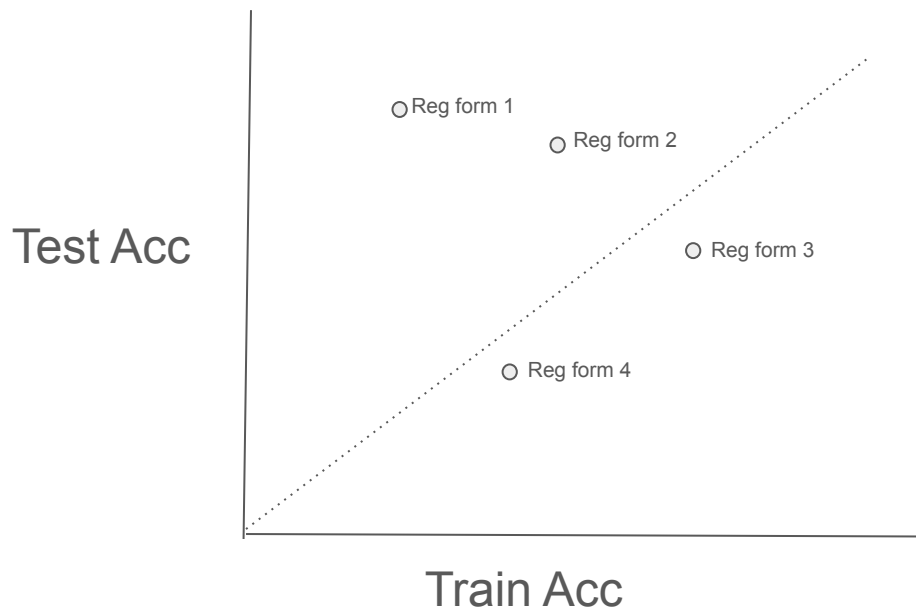
Do a deep dive into one part of this pipeline

Start with your ideal figure and work backwards

Comparing the generalization gap for various forms of regularizations

Start with your ideal figure and work backwards

Comparing the generalization gap for various forms of regularizations




Start with your ideal figure and work backwards

Does pre-training on images of non-animals help with animal classification?

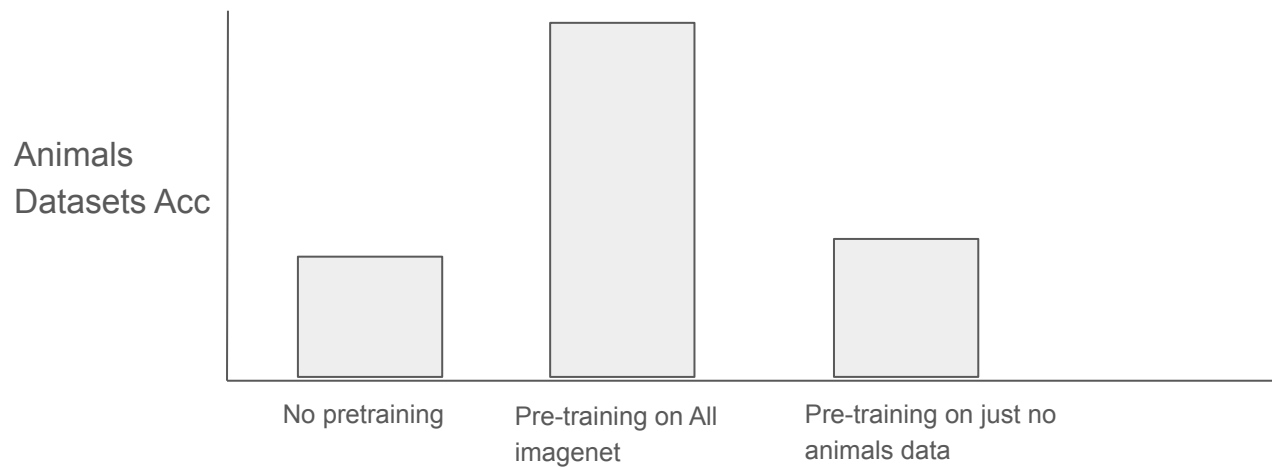
Start with your ideal figure and work backwards

Does pre-training on images of non-animals help with animal classification?



Start with your ideal figure and work backwards

Does pre-training on images of non-animals help with animal classification?



Picking a project idea

Conferences:

[CVPR](#): IEEE Conference on Computer Vision and Pattern Recognition

[ICCV](#): International Conference on Computer Vision

[ECCV](#): European Conference on Computer Vision

[NeurIPS](#): Neural Information Processing Systems

[ICLR](#): International Conference on Learning Representations

[ICML](#): International Conference on Machine Learning

(I personally like looking at best paper awards)

Picking a project idea

Labs at UW

Vision + Graphics (GRAIL) - <https://grail.cs.washington.edu/>

Vision (RAIVN) - <https://raivn.cs.washington.edu/>

NLP (H2lab) - <https://h2lab.cs.washington.edu/>

NLP (Noah Ark) - <https://noahs-ark.github.io/>

Robotics (RSE lab) - <http://rse-lab.cs.washington.edu/>

Picking a project idea

DOI: 10.1109/CVPR.2016.91 · Corpus ID: 206594738

Share This Paper

You Only Look Once: Unified, Real-Time Object Detection

Joseph Redmon, S. Divvala, +1 author, Ali Farhadi · Published 8 June 2015 · Computer Science · 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)

We present YOLO, a new approach to object detection. Prior work on object detection repurposes classifiers to perform detection. Instead, we frame object detection as a regression problem to spatially separated bounding boxes and associated class probabilities. A single neural network predicts bounding boxes and class probabilities directly from full images in one evaluation. Since the whole detection pipeline is a single network, it can be optimized end-to-end directly on detection performance... [Expand](#)

22,521 Citations

Highly Influential Citations 2,952
Background Citations 7,368
Methods Citations 8,237
Results Citations 144

[View All](#)

[\[PDF\] Semantic Reader](#) [Save to Library](#) [Create Alert](#) [Cite](#)

Figures and Tables

22,521 Citations

40 References

Related Papers

Figures and Tables from this paper



Figure 1

Real-Time Detectors	Train	mAP	FPS
YOLO: DPM1 [11]	2007	50.0	100
SOLO: DPM4 [11]	2007	26.1	30
Fast YOLO	2007-2012	52.7	165
YOLO	2007-2012	63.4	45

Less Than Real-Time	Train	mAP	FPS
Fastest DPM1 [11]	2007	30.4	15
R-CNN: Misses R [10]	2007	53.5	6
Fast R-CNN [11]	2007-2012	70.0	6.5
Fast R-CNN: VGG-16 [11]	2007-2012	73.2	7
Fast R-CNN: VGG-16 [11]	2007-2012	62.1	18
YOLO: VGG-16	2007-2012	66.4	21

Table 1




Figure 2

	mAP	Combood	Gain
Fast R-CNN	71.8	-	-
Fast R-CNN (DPM1: Data)	66.9	72.4	6
Fast R-CNN (VGG-M)	59.2	72.4	6
Fast R-CNN (CaffeNet)	57.1	72.1	3
YOLO	63.4	70.0	3.2

Table 2




Figure 3



Reading papers

Do not read a paper linearly on your first pass

- First, read the abstract (word for word) as well as the figures & captions
- Does the paper still seem relevant? If so, read the methods + results
- Finally, read the entire paper linearly (if the additional detail seems useful)

Papers are not always the most efficient way to digest an idea. Also try looking for:

- Talks, videos, or blog posts on the topics
- Github repos, containing actual code for the idea

Reading papers

You Only Look Once: Unified, Real-Time Object Detection

Joseph Redmon^{*}, Santosh Divvala^{*†}, Ross Girshick[¶], Ali Farhadi^{*†}

University of Washington^{*}, Allen Institute for AI[†], Facebook AI Research[¶]

<http://pjreddie.com/yolo/>

Outline

1. Project Expectations

- a. Does my project meet expectations?
- b. FAQs

2. Picking a Project idea

- a. Inspiration
- b. How to read a research paper

3. Proposal, milestones, and final report

- a. Due dates, expectation, logistics
- b. Support

Due Dates

Project Proposal (1 Page) – Due Feb 6th 11:59pm

Project Milestone (3-4 Pages) – Due Feb 29th 11:59pm

Project Final Report (5-6 Pages) – Due Final Exam Week

Project Poster Session – Final Exam Week

Project Proposal

- Describe the state of related work,
- Explain a problem that is unsolved given that statement,
- Introduce your ideas as an unique insight to tackle the problem or research question,
- Articulate the technical challenges you are likely to encounter,
- Plan out the experiments that justify the utility of the insight or answers the question,
- Your expected outcome

Milestone (Due Feb 29th 11:59pm)

~3-page progress report, more or less containing:

1. Literature review (3+ sources)
2. Indication that code is up and running
3. Data source explained correctly
4. What Github repo or other code you're basing your work off of
5. Ran baseline model have results
 - a. Yes, points are taken off for no model running & no preliminary results
6. Data pipeline should be in place
7. Brief discussion of your preliminary results