CS231n Project Design

Chris Waites

Agenda

1. Project expectations
    a. Does my project meet expectations?
    b. FAQs
2. Picking a project idea
    a. Sources of inspiration
    b. Reading papers efficiently
3. Proposal, milestone, and final report
    a. Due dates, expectations, logistics
    b. Support

Agenda

1. Project expectations
    a. Does my project meet expectations?
    b. FAQs
2. Picking a project idea
    a. Sources of inspiration
    b. Reading papers efficiently
3. Proposal, milestone, and final report
    a. Due dates, expectations, logistics
    b. Support

Project expectations

The course project is a (fun) way to explore concepts taught in the course on a topic of your choice!

- Fairly open-ended, anything related to vision (**link to project page**)

Completed in groups of 1, 2, or 3 people

- Project expectations are higher for groups with more people

Generally, two tracks of work:

- **Applications:** If you have a specific background or interest (e.g. biology, engineering, physics), we'd love to see you apply ConvNets to problems related to your particular domain of interest.
- **Models:** You can build a new model (algorithm) and apply it to tackle vision tasks. This track might be more challenging, and could lead to a piece of publishable work.

Project expectations

The final report has the following structure:

- Title, Author(s)
- Abstract
- Related Work
- Data Description
- Methods
- Experiments
- Conclusion
- Supplementary Material (optional)

FAQ: Does my project meet expectations?

Rule of thumb:

- ***How much effort are you putting into your project?***

Strong projects might...

- Propose a novel variant of a technique (which takes a lot of effort)
- Adapts 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

FAQ: Does my project meet expectations?

So, this **doesn't** mean:

- Your project has to be strictly novel to get a good grade (although, we encourage this!)
- You have to beat the state-of-the-art performance to get a good grade (you don't have to come up with the next best object detector to test an interesting hypothesis)

This **does** mean:

- You need to put a significant effort into your investigation, and you may have to try many different approaches

In your **analysis**, ask yourself:

- Are you *interpreting* and *understanding* your results, or merely stating them?
- Are you just plotting a loss curve, or are you evaluating the results of your approach from many different angles?

Project FAQs

Q: Can I apply convolutional networks to a purely NLP / audio / stock price problem?

- **A:** This is a computer vision course, so ***you must incorporate visual data in some form***.

Q: Can I change my project after the proposal, before the milestone?

- **A:** Yes - the proposal is to make sure you have a plausible project direction. If you need to change project directions, we understand.

Q: Can I change my project ***after*** the milestone?

- **A:** In general, we do not encourage this. At this point in the course, there will be little time to put together a sufficient project.

Agenda

1. Project expectations
    a. Does my project meet the course expectations?
    b. FAQs
2. Picking a project idea
    a. Sources of inspiration
    b. Reading papers efficiently
3. Proposal, milestone, and final report
    a. Due dates, expectations, logistics
    b. CA Support

Picking a project idea

First and foremost:

Do what is important or interesting to you, *not* what seems easiest.

- You will be far more motivated if you're invested in what you're doing
- What do you *really* care about? Healthcare? Self-driving cars? Surveillance? Sports? Ethics? You can probably find its intersection with computer vision

Practical considerations:

1. **Data:** Is there existing data for this problem? Will I need to spend weeks collecting it myself?
2. **Code & framework:** Will I have to implement this myself, or is there an existing implementation?

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

*Note: Do **not** even begin to try to read through all of these papers, or even their titles. There are far too many. Use CMD+F to find papers with relevant keywords.*

Picking a project idea

Additional resources:

- Stanford Vision Lab Publications
- Awesome Deep Vision
- Papers With Code
- Kaggle
- Previous CS229 Projects

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 as well as the 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 around for:

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

Reading papers

Example:

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/

Agenda

1. Project expectations
    a. Does my project meet the course expectations?
    b. FAQs
2. Picking a project idea
    a. Sources of inspiration
    b. Reading papers efficiently
3. Proposal, milestone, and final report
    a. Due dates, expectations, logistics
    b. CA Support

Deliverables

Due dates:

- Proposal (4/18) - Monday!
- Milestone (5/10)
- Final report (6/3)
- Final presentation (6/4) - Video submitted to gradescope, async.

Milestone (Due 5/10)

1-2 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

Support: CA areas of specialty

Day	Staff		
Mon	Kevin *Robotics, self supervised learning*	Haofeng *Detection, segmentation, multiple object tracking, video understanding, interactive annotation*	
Tue	Rachel *NLP (vision+language), Amazon Mechanical Turk, creating custom datasets, RL, robotics, 3D vision, AI for medical imaging.*	Geet *Performance (e.g. pruning, quantization), recommendation/embedding, sparse networks, transfer Learning, videos, RNNs*	Mandy *Action recognition, videos, medical imaging, optimization/learning methods*
Wed	Chris *Generative models (GANs), privacy, fairness & interpretability*	Sam *Video understanding, action recognition, speech recognition, bias, multi-task learning, medical imaging*	Nishant *Videos, unsupervised learning, self-supervised learning, 3D vision*
Thu	Guanzhi *RL, robotics, videos, self-supervised learning*	Russell *Videos, action recognition, graph neural networks, domain adaptation, image segmentation / denoising*	Sean *Graphics, image editing, virtual reality, 360 degree videos*
Fri	Yichen *3D vision, transfer learning, domain adaptation*	JQ *Image classification, image augmentation, medical imaging (classification)*	Lin *Robotics, RL, 3D vision*

Questions?