

STANFORD · CS231N · SPRING 2026

Final Project Overview

Expectations, deliverables, and a taxonomy of past projects

Cristóbal Eyzaguirre · adapted from Emily Jin & Chris Waites

CS231N

*Deep Learning for
Computer Vision*

SECTION 3

Agenda

01 Project Expectations

Scope, tracks, and what a strong project looks like

02 Picking a Project Idea

Sources of inspiration and how to read papers efficiently

03 Project Deliverables

Proposal, three milestones, final report, and poster session

04 A Taxonomy of Past Projects

Three patterns that have worked, with examples

05 Support & Logistics

CA specialties, FAQ, and due dates

Open-ended, but it has to involve vision

Projects are fairly open-ended — anything related to vision. Your project must involve visual data. Groups of 1–3; expectations scale with group size.



TRACK 1

Applications

Apply vision techniques to your particular domain or interest — medicine, physics, robotics, ecology, and beyond.

Careful data preparation · appropriate loss · training details · solid evaluation and model comparisons.



TRACK 2

Models

Build a new model or variant of an existing one and apply it to tackle vision tasks.

More challenging — but can lead to publishable work.

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 real effort to design and validate*
- **Adapt an existing technique to a totally new problem**
 - *bringing vision tools to an under-served domain*



Weaker projects might...

- **Spend several weeks just collecting and cleaning data**
 - *rather than actually testing hypotheses*
- **Clone an existing repo and stitch it together minimally**
 - *e.g. to run a Kaggle competition with no real contribution*

To be clear...

THIS DOES NOT MEAN

Your project has to be strictly novel

(Though we encourage it!)

You have to beat the state of the art

You don't need the next great object detector to test an interesting hypothesis.

THIS DOES MEAN

Put significant effort into your investigation

You may need to try many different approaches before something works.

Interpret and understand your results

Don't just plot a loss curve — evaluate your approach from multiple angles.

What to work on



Consider your interests

Do what's 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? Sports? Ethics? Climate? Art?

You can almost certainly find its intersection with computer vision.



Practical considerations

1

Data

Is there existing data, or will you burn weeks collecting it?

2

Code & framework

Do you have to implement it yourself, or is there an existing implementation you can build on?

Need a starting point? See the *Potential Project Ideas* doc linked in the course materials.

Where to look & how to read

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

ALSO WORTH BROWSING

SVL Publications · Awesome Deep Vision · TA's Research · Past Projects

How to read papers efficiently

Don't read linearly on your first pass

- 1 Abstract + figures & captions**
Read the abstract word for word, then skim the figures and their captions.
- 2 Still relevant? Read methods + results**
Focus on what the paper does and how well it works.
- 3 Entire paper, linearly**
Only if the additional detail is genuinely useful for your project.

Also try: talks, blog posts, or GitHub repos — papers aren't always the fastest way to digest an idea.

Three patterns that have worked

Looking at recent CS231N final reports, successful projects tend to fall into one of three patterns. These aren't the only options — but they're useful templates to anchor your thinking.



PATTERN 1

Domain Adaptation

Apply a strong vision or VLM to a meaningful new problem, with real task-specific data and evaluation.

TYPICAL HIGHLIGHTS

medical imaging · remote sensing · scientific imaging · sign language · geospatial



PATTERN 2

Method Improvement

Start from a recognized baseline and make a technically meaningful change — new loss, module, or training method.

TYPICAL HIGHLIGHTS

new losses · faster/more stable training · better multimodal fusion · new SSL objectives



PATTERN 3

Reproduction

Rebuild a proprietary or hard-to-reproduce capability — contribution is the implementation itself.

TYPICAL HIGHLIGHTS

no clean OSS version · making known capabilities actually run in practice

Domain Adaptation / Application of Strong Vision Models

Use a strong existing vision or vision-language model on a meaningful new problem, with real task-specific data, training, and evaluation.



MEDICAL IMAGING

Enhancing SAM for Brain Tumor Segmentation



INFRASTRUCTURE

Two-Step Deep Learning for Classifying Bridges from Street Imagery



SCIENTIFIC IMAGING

Rock Image Super-resolution: From CT to micro-CT



VISION + LANGUAGE

Bridging Vision and Language for Sign Language Understanding

Method Improvement & Reproduction



Method Improvement on a Known Baseline

Start from a recognized baseline; make a technically meaningful change — loss, module, fusion, training method, architectural tweak.

- PocketNeRF — Fast-Converging NeRF for Indoor Reconstruction from Few-Shot Mobile Images
- Total Variation Loss for Compact-Object Segmentation on Satellite Images
- Improved Mineral Detection via Spectral Attention U-Net with a Novel Hapke Layer
- Adaptive Contrastive Masked Autoencoders for Structured Representation Learning



Reproduction / Reimplementation

Rebuild an existing proprietary or hard-to-reproduce visual capability — the contribution is the implementation, analysis, and often performance optimization.

- Replicate FaceApp Effect and Enable Real-time Performance Based on GANs

Six checkpoints across the quarter

DUE	DELIVERABLE	WEIGHT	LATE DAYS	DESCRIPTION
Apr 23	Project Proposal	1%	<i>Late days</i>	One paragraph (200–400 words): problem & motivation, reading, data, model/algorithm, evaluation.
May 15	Milestone 1 · Problem + Related Work	3%	<i>Late days</i>	Define the problem precisely and situate it against prior work.
May 22	Milestone 2 · Technical Approach	3%	<i>Late days</i>	Describe the methods you intend to apply — model, data pipeline, training plan.
May 29	Milestone 3 · Preliminary Results	3%	<i>Late days</i>	Intermediate experimental results to validate your direction.
Jun 05	Final Report	20%	<i>No late days</i>	6–8 pages in the provided template. Title, abstract, intro, related work, data, methods, experiments, conclusion.
Jun 09 – 10	Poster & Code	5%	<i>No late days</i>	PDF + code due Jun 09 11:59pm PT. In-person poster session Jun 10.

Note: 4 free late days per student for the quarter, up to 2 per assignment. Late days cannot be used on the final report.

CA Areas of Specialty

Zane Durante

Computer vision, video understanding, hospital care

Aditesh Kumar

Robotics, reinforcement learning, privacy

Keshigeyan Chandrasegaran

Computer vision, machine learning

Mark Endo

Vision + language

Aniket Gupta

Video generation

Wenlong Huang

Robotics, machine learning

Chaitanya Patel

3D vision, human perception, video

Yash Shah

Cognitive, computational neuroscience

Karanpartap (Karan) Singh

Interpretability, SSL, medical applications

Bailey Trang

Image/video generation, interpretability

Heng Yu

Video generation/understanding, world modeling

Hongxing (Koven) Yu

3D vision, graphics, video generation

Cristóbal Eyzaguirre Ercilla

Video understanding

Fangrui Huang

Applications in healthcare

Jiayi (Eris) Zhang

Graphics, GenAI, HCI

Favour Nerrise

Computer vision, ambient sensors, healthcare, pose estimation

Yalcin Tur

Computer vision, embodied AI, robot learning

June Zheng

Computer vision, multi-modal foundation models, agentic systems, healthcare

Yang Zheng

3D vision, graphics, generative AI

A few we hear every year



Can I change my project after the proposal, before the milestones?

A. Yes. The proposal is to make sure you have a plausible direction. If you need to pivot, check with a CA in OH first.



Can I change my project after the milestones?

A. We do not encourage this. By that point there's little time left to put together a sufficient project.



Can I combine the Final Project with another course?

A. Yes, with the other instructor's permission. You must clearly specify the unique portion of the work counted for CS231N and prepare separate reports.



I have a question about the class — what's the best way to reach the staff?

A. Almost all questions should go on Ed. For sensitive issues, email the staff alias listed on the course site.