



Gül Varol

(Main
lecturer)



Jean Ponce

Class logistics



Cordelia
Schmid



Ivan Laptev



Vincent Lepetit

Gül Varol

IMAGINE team, École des Ponts ParisTech


gulvarols@gmail.com

<https://gulvarol.github.io/>

@RecVis, 29.09.2026



ÉCOLE NATIONALE DES
PONTS
ET CHAUSSÉES



Object recognition and computer vision 2024

Reconnaissance d'objets et vision artificielle (RecVis)

Lecturers



[Gül Varol](#)



[Jean Ponce](#)

Sep 29



[Cordelia
Schmid](#)

Dec 1



[Ivan Laptev](#)

Dec 8



[Vincent Lepetit](#)

Nov 24

Contact for other Qs
gulvarols@gmail.com

Teaching Assistants (TAs)



[Alexandros
Benetatos](#)



[Fernando Julio
Cendra](#)

Contact TAs for questions
about assignments
alexandrosbene@gmail.com
fernando.cendra@enpc.fr

Schedule

Tuesdays 15h - 18h

Location: Amphi Luton



Follow updates &
exceptions on class
webpage

#	Date	Lecturer	Topic and reading materials
Instance-level recognition			
1	Sep 29	Gül Varol Jean Ponce	Class logistics: assignments, final projects, grading; Introduction to visual recognition; Camera geometry; Image processing
2	Oct 6	Gül Varol	Instance-level recognition: local features, correspondence, image matching Assignment 1 (A1) out.
Practical	Oct 13 *hybrid*	TAs	Pytorch/Kaggle/Google Cloud tutorial. Presentations by TAs about their research topics. <i>The tutorial will take place at Imagine/ENPC in person, and online participation will be possible.</i>
3	Oct 20 *Amphi Dieulafoy*	Gül Varol	Efficient visual search Final project (FP) topics are out at the end of the lecture.
Category-level recognition			
4	Oct 27	Gül Varol	Supervised learning and deep learning; Optimization and regularization for neural networks A1 due. A2 out.
5	Nov 3	Gül Varol	Neural networks for visual recognition: CNNs and image classification A3 out.
6	Nov 10	Gül Varol	Beyond CNNs: Transformers; Beyond classification: Object detection; Segmentation; Human pose estimation A2 due.
Advanced topics			
7	Nov 17 *Amphi Dieulafoy*	Gül Varol	Generative models; Vision & language FP proposal due.
8	Nov 24	Vincent Lepetit	3D computer vision A3 due.
9	Dec 1	Cordelia Schmid	Human action recognition in videos
10	Dec 8 *Amphi Dieulafoy*	Ivan Laptev	Vision for robotics
FP	Jan 11-12	Gül Varol	FP presentations Presentations will be virtual. FP report due Jan 18.

Schedule: Add public calendar



<https://tinyurl.com/recvis26calendar>

Lecture time: Tuesdays 15:00-18:00


Lecture room: [Amphi Luton](#), 24 rue du F

*A few exceptions are denoted in the sch

Jacques, 75014 Paris ([maps](#))

The class [Google Calendar](#) is up to date

Note: Slides are provided after each lect



RecVis 2026/2027


Events shown in time zone: (GMT+02:00) Central European Time - Paris

[Add to Google Calendar](#)

Practical information: Participation



Class webpage : <https://gulvarol.github.io/teaching/recvis26/>

Google Classroom : Register with the code  to receive announcements.

Time : 15h00-18h00, Tuesdays, starting Sep 29

Location : Amphi Luton, 24 rue du Faubourg Saint-Jacques, 75014 Paris

Format : In-person lectures. Slides provided after each lecture.

For externals : You are welcome to attend the course (either for auditing or validation) provided there are enough free places in the lecture hall. If your school requires a proof of attendance, you need to get signatures from teachers after every lecture.

Practical information: Grading

- **3 programming assignments (50%)**

- A1: Instance-level recognition
- A2: Neural networks
- A3: Image classification Kaggle competition

Policy

Assignments are strictly individual

Copy-paste of the code, results, parts of the report ⇒ 0p.

- **Final project (50%)**


- More independent work, resulting in a report and a class presentation.
- We will provide Google Cloud credits for each student.

FPs can be done in groups of max 2 people

Experience with Python (numpy, pytorch) will be useful, but TAs will provide an optional crash-lecture on **Pytorch** for computer vision, as well as **Kaggle** and **Google Cloud** usage.


Assignment I: Instance-level recognition

- Part I: Sparse features for matching specific objects in images
 - Feature detector and descriptor
 - Robust match filtering techniques
 - Augmented reality
- Part II: Compact descriptors for image retrieval



Assignment II: Neural networks

- Part 1: Neural Network's theory:
 - Forward pass, Backward pass
 - Parameter update
- Part 2: Building blocks of convolutional neural networks
- Part 3: Training a CNN on CIFAR-10 dataset with PyTorch



airplane

automobile

bird

cat

deer


dog

frog

horse

ship

truck



Assignment III: Image classification competition

- Class Kaggle competition
- Example task: Bird image classification - the assignment will cover a similar task



Final Project (FP)

- Can be done individually or as a **group of max 2 people**
- The proposed project topics are from the recent top-conference publications in computer vision, see example topics from 2024 (or earlier) here: <https://gulvarol.github.io/teaching/recvis24/>
- Student-defined projects are welcome.
- Final project can be joint with another MVA course.
- We arranged \$50 Google Cloud credits for the project (with possibility to request more).
 - This will be announced through Google Classroom before projects start

- Select the topic + write project proposal
- Present the work in the class
- Write project report

Practical by TAs

Fill-in the TAs session participation form linked from the class webpage by **Tue Oct 13**.

The tutorial will be hybrid on **Tue Oct 13**, between **15h00-17h00**:

Physical: École des Ponts ParisTech (RER A Noisy Champs), Coriolis Building, IMAGINE

Online: Video link will be announced on Google Classroom <https://gulvarol.github.io/findus.html>

Who should participate?

- Students with no or limited experience with PyTorch/Kaggle/Google Cloud
- Students curious to know about life of a PhD student in computer vision
- Attendance is optional.



Agenda:

- Presentations by TAs on their PhD topics, followed by Q&A
- Introduction to PyTorch for computer vision
- Kaggle competition (used for Assignment 3)
- Using Google Cloud (used for Final Project)

Course outline


1. Instance-level recognition

- Camera geometry
- Image processing
- Image correspondence
- Large-scale image and video search



2. Category-level recognition

- Supervised learning
- Neural networks for visual recognition
- Object recognition, detection, and segmentation



3. Advanced topics

- Generative models; Vision & language
- Human action recognition in videos
- Vision for robotics
- 3D computer vision





Vibrant portrait painting of Salvador Dali with a robotic half face



Course outline


1. Instance-level recognition

- Camera geometry
- Image processing
- Image correspondence
- Large-scale image and video search



2. Category-level recognition

- Supervised learning
- Neural networks for visual recognition
- Object recognition, detection, and segmentation



3. Advanced topics

- Generative models; Vision & language
- Human action recognition in videos
- Vision for robotics
- 3D computer vision





Vibrant portrait painting of Salvador Dali with a robotic half face



Course outline


1. Instance-level recognition

- Camera geometry
- Image processing
- Image correspondence
- Large-scale image and video search



2. Category-level recognition

- Supervised learning
- Neural networks for visual recognition
- Object recognition, detection, and segmentation



3. Advanced topics

- Generative models; Vision & language
- Human action recognition in videos
- Vision for robotics
- 3D computer vision



Vibrant portrait painting of Salvador Dali with a robotic half face






3. Advanced topics

- Generative models; Vision & language
- Human action recognition in videos
- Vision for robotics
- 3D computer vision



Vibrant portrait painting of Salvador Dali with a robotic half face



Resources


- D.A. Forsyth and J. Ponce, "Computer Vision: A Modern Approach", Prentice-Hall, 2nd edition, 2011
- J. Ponce, M. Hebert, C. Schmid and A. Zisserman "Toward Category-Level Object Recognition", Lecture Notes in Computer Science 4170, Springer-Verlag, 2007
- O. Faugeras, Q.T. Luong, and T. Papadopoulos, "Geometry of Multiple Images", MIT Press, 2001.
- R. Hartley and A. Zisserman, "Multiple View Geometry in Computer Vision", Cambridge University Press, 2004.
- J. Koenderink, "Solid Shape", MIT Press, 1990
- R. Szeliski, "Computer Vision: Algorithms and Applications, 2nd ed.", 2022. [Online book.](#)
- [Computer Vision: Models, Learning, and Inference by Simon J.D. Prince \(2012\)](#)
- [Understanding Deep Learning by Simon J.D. Prince \(2023\)](#)
- [Deep Learning by I. Goodfellow, Y. Bengio and A. Courville \(2016\)](#)
- [Michael Nielsen's online book on Neural Networks and Deep Learning \(2019\)](#)
- [David Forsyth's Applied Machine Learning textbook draft \(2019\)](#)
- [Andrej Karpathy blog](#)

...

Let me know if you find a good resource (book, videos etc)
to link from the webpage.

Resources

- D.A. Forsyth and J. Ponce "Computer Vision: A Modern Approach", Prentice-Hall, 2nd edition, 2011



Jean Ponce
Today (Sep 29)

el Object Recognition", Lecture Notes in Computer Science 4170,

ges", MIT Press, 2001.

n", Cambridge University Press, 2004.

2. [Online book.](#)

[\(2012\).](#)

[\(2019\).](#)

Publication Venues

Read “good” papers

1) Conference on **C**omputer **V**ision and **P**attern **R**ecognition (**CVPR**)

2) International **C**onference on **C**omputer **V**ision (**ICCV**)

3) **E**uropean **C**onference on **C**omputer **V**ision (**ECCV**)

4) Transactions on **P**attern **A**nalysis and **M**achine **I**ntelligence (**TPAMI**)

5) International **J**ournal of **C**omputer **V**ision (**IJCV**)

*Conference
proceedings*

Journals

...

ACCV, WACV, BMVC, CVIU...


...

NeurIPS, ICLR, ICML...


EMNLP, ACL...

IROS, CoRL, ICRA...

Reading recommendation



IMAGINE computer vision team, ENPC



École des Ponts
ParisTech

imagine.enpc.fr/



Keep an eye on internships

Research

IMAGINE and WILLOW teams are active in computer vision.

<http://imagine.enpc.fr>

<http://www.di.ens.fr/willow/>

There will be master internships available. Talk to us if you are interested!

Recap

1. Register on the Google Classroom

- Assignment submissions, discussions and announcements will be done on Google Classroom.
- Assignment 1 – Instance-level recognition (due Oct 27)

2. Fill-in TAs session participation form (by Oct 6)