READING GROUP: EXPLORING LATERAL CONNECTIONS (WITH PRIZE OF 1000$)

Stefano Soatto
University of California, Los Angeles, USA

The input to Computer Vision are images. The output is either decisions (detection, recognition, categorization, etc.) or actions (navigation, manipulation, tracking etc.). In between, Computer Vision aims at computing some function of the input that is ``useful'' to the (output) task. Functions of the input, whether ``learned'', ``stored'', or ``engineered'', that are useful to the task are called Representations. In this reading group we will explore the concept of Representation in Computer Vision in relation to other disciplines.

The study of representations prompts a number of questions: are they necessary? Can we do without? What representation should we compute and maintain in memory for a computer vision algorithm? How do we measure how ``useful'' it is? Can the Representation be ``separated'' from the task? What if the representation is the task itself? Is there an optimal representation? If so, can it be computed? If not, can it be approximated? Does it depend on the task? What is the task is not known? What do algorithms designed for different tasks (e.g. categorization vs. 3D reconstruction) have to do with each other? What do they have in common? Is memory even necessary? Can a direct input-output ``reactive'' behavior work? What do animals store in memory? How?

The reading group will consist of two phases: An Analysis phase, where you will have to read up, and pay attention to preceding lectures. A collection of references is provided for convenience. This is best done before the school starts. Then there will be a Synthesis phase, where we will discuss and brainstorm live on what ideas are key, and see if connections emerge. This will be done on-site at the school.

Everyone can participate in the reading group even without prior preparation. For those willing to invest some time before the school, and to turn in a written summary of their findings, there will be a reward: The top 3 essays, as judged by the School Committee, will be called to present their findings to the rest of the group, and the audience will vote the best. The best will receive a $1,000 prize.

ANALYSIS:

Read up references in the following areas (pick any subset, including the full set). One paper is listed as a ``seed'' but feel free to roam the literatured. Some areas have two papers, you can pick either one, or both if you are curious. One is optional (Geometry). You will also be asked to pay particular attention to 4 lectures preceding the reading group:

1) Statistics:

- Sufficient Statistics, G. Lebanon (2 pages)
- (optional) Equivariant Estimation, P. D. Hoff (38 pages)

2) Information Theory:

- Data Processing Inequality: S. D. Servetto (6 pages)
- (optional) Information Theory, Cover and Thomas, Chapter 2, Section 2.8 (38 pages)
- (optional) The Information Bottleneck Method, N. Tishby et al. (11 pages)

3) Biology:

- The Biology of Memory, E. R. Kandel (9 pages)

4) Computer Architecture:

- (optional) The Impact of Memory and Architecture on Computer Performance, N. H. F. Beebe (70 pages)

5) Cognitive Science:

- What is a Knowledge Representation, R. Davis et al. (17 pages)

6) Philosophy:

- (optional) Mental Representation, SEP (19 pages)

Geometry:

- (optional) Representation theory, P. Etingof (108 pages)

*) Machine Learning:

- Learning Representations, M. A. Ranzato's lecture (Tuesday 15 July 2014, 15:00)

*) Computer Vision:

- 3D reconstruction, Y. Furukawa's lecture (Monday 14 July 2014, 11:50)
- Visual recognition, F. Perronnin's lecture (Tuesday 15 July 2014, 9:00)
- Perceptual organization, G. Medioni's lecture (Tuesday 15 July 2014, 11:20)

SYNTHESIS:

What did you learn? What are the key ideas? Where did you see connections? Could you see hints on how to answer some of the questions outlined above? Or did you come up with new questions? Be prepared to bring your thoughts and comments to the reading group.

COMPETITION:

If you can do the synthesis part ahead of time, write it down in a brief summary (2-3 pages max) and submit it to the School Organizers (see instruction below) at the start of the school. You will be evaluated based on the following three criteria:

1) New: you discovered something that was not known (to you, at least) before

2) Useful: The ``connections'' you established seem promising for the development of computer vision

3) Non-obvious: you could not have discovered these connections prior to the analysis

The 3 finalists will be called to present their findings in front of the group. The group will vote the best; the winners will receive a check for $1,000.

READING MATERIAL:

The reading material is available at the following link:

http://svg.dmi.unict.it/icvss2014/ICVSSReadingGroup2014.rar

SUBMISSION:

The report about your homework is due on the first day of the school (13/07/2014) and should be sent in PDF or Word format to icvss@dmi.unict.it. Your name, email address and university/organization should be written at the top of your report.