EECS 4422/5323 Computer Vision

Stereopsis 2

Calden Wloka

6 November, 2019

Announcements

- More site visits today
- Assignment 2 data posted Monday afternoon, updates to text as recently as Tuesday
- Reminder: Drop date is this Friday, so review your marks

White Paper (2%) + Proposal (8%) + Site Visit (5%) + Assignment 1 (15%) = 30% total grade

A Note About Grayscale

Gray is a range of colours which, in the RGB colour space, have the property that $R = G = B$.

A Note About Grayscale

Gray is a range of colours which, in the RGB colour space, have the property that $R = G = B$.

A grayscale pixel value can therefore be compactly represented by only a single value, or it can be redundantly represented by three channels. The way in which it is represented is a matter of convention and software standards, and is rather hard (at least for me) to remember and predict.

Outline

- Academic Publishing
- Correspondence Problem
- Constraints and Strategies
    - Coarse-to-fine processing
    - Epipolar geometry

Quantifying Academic Output

Now that we're in the thick of it, I thought it would be useful to refresh the possible goals and outcomes from your projects.

- Traditionally, academic output is predominantly measured in the form of published peer reviewed papers, or metrics revolving around papers

Quantifying Academic Output

Now that we're in the thick of it, I thought it would be useful to refresh the possible goals and outcomes from your projects.

- Traditionally, academic output is predominantly measured in the form of published peer reviewed papers, or metrics revolving around papers
    - Citations

Quantifying Academic Output

Now that we're in the thick of it, I thought it would be useful to refresh the possible goals and outcomes from your projects.

- Traditionally, academic output is predominantly measured in the form of published peer reviewed papers, or metrics revolving around papers
  - Citations
  - Author order/Publication venue

Quantifying Academic Output

Now that we're in the thick of it, I thought it would be useful to refresh the possible goals and outcomes from your projects.

- Traditionally, academic output is predominantly measured in the form of published peer reviewed papers, or metrics revolving around papers
  - Citations
  - Author order/Publication venue
- There are some modern alternatives to traditional publications (which can get somewhat confusing)

Quantifying Academic Output

Now that we're in the thick of it, I thought it would be useful to refresh the possible goals and outcomes from your projects.

- Traditionally, academic output is predominantly measured in the form of published peer reviewed papers, or metrics revolving around papers
  - Citations
  - Author order/Publication venue
- There are some modern alternatives to traditional publications (which can get somewhat confusing)
  - Non-peer-reviewed publication (preprints, tech reports)

Quantifying Academic Output

Now that we're in the thick of it, I thought it would be useful to refresh the possible goals and outcomes from your projects.

- Traditionally, academic output is predominantly measured in the form of published peer reviewed papers, or metrics revolving around papers
    - Citations
    - Author order/Publication venue
- There are some modern alternatives to traditional publications (which can get somewhat confusing)
    - Non-peer-reviewed publication (preprints, tech reports)
    - Code contributions (particularly GitHub activity)

Quantifying Academic Output

Now that we're in the thick of it, I thought it would be useful to refresh the possible goals and outcomes from your projects.

- Traditionally, academic output is predominantly measured in the form of published peer reviewed papers, or metrics revolving around papers
  - Citations
  - Author order/Publication venue
- There are some modern alternatives to traditional publications (which can get somewhat confusing)
  - Non-peer-reviewed publication (preprints, tech reports)
  - Code contributions (particularly GitHub activity)
- In computer science and engineering, the line between academy and industry can be somewhat blurry

Quantifying Academic Output

Now that we're in the thick of it, I thought it would be useful to refresh the possible goals and outcomes from your projects.

- Traditionally, academic output is predominantly measured in the form of published peer reviewed papers, or metrics revolving around papers
  - Citations
  - Author order/Publication venue
- There are some modern alternatives to traditional publications (which can get somewhat confusing)
  - Non-peer-reviewed publication (preprints, tech reports)
  - Code contributions (particularly GitHub activity)
- In computer science and engineering, the line between academy and industry can be somewhat blurry
  - Many industrial research positions still encourage employees to publish

Quantifying Academic Output

Now that we're in the thick of it, I thought it would be useful to refresh the possible goals and outcomes from your projects.

- Traditionally, academic output is predominantly measured in the form of published peer reviewed papers, or metrics revolving around papers
  - Citations
  - Author order/Publication venue
- There are some modern alternatives to traditional publications (which can get somewhat confusing)
  - Non-peer-reviewed publication (preprints, tech reports)
  - Code contributions (particularly GitHub activity)
- In computer science and engineering, the line between academy and industry can be somewhat blurry
  - Many industrial research positions still encourage employees to publish
  - Having publications on your CV can boost your appeal for a more research-oriented position

Project Goals

I have separated the target outcomes between Engineering and Scientific stream projects below, but they are not mutually exclusive.

- Engineering stream projects

Project Goals

I have separated the target outcomes between Engineering and Scientific stream projects below, but they are not mutually exclusive.

- Engineering stream projects
  - A software package hosted in a public repository

Project Goals

I have separated the target outcomes between Engineering and Scientific stream projects below, but they are not mutually exclusive.

- Engineering stream projects
  - A software package hosted in a public repository
  - A technical report describing the software you have produced

Project Goals

I have separated the target outcomes between Engineering and Scientific stream projects below, but they are not mutually exclusive.

- Engineering stream projects
  - A software package hosted in a public repository
  - A technical report describing the software you have produced
- Scientific stream projects

Project Goals

I have separated the target outcomes between Engineering and Scientific stream projects below, but they are not mutually exclusive.

- Engineering stream projects
  - A software package hosted in a public repository
  - A technical report describing the software you have produced
- Scientific stream projects
  - Possibly a good start on a paper

Project Goals

I have separated the target outcomes between Engineering and Scientific stream projects below, but they are not mutually exclusive.

- Engineering stream projects
  - A software package hosted in a public repository
  - A technical report describing the software you have produced
- Scientific stream projects
  - Possibly a good start on a paper
  - A codebase which could be made public

Publishing

Not all publishing venues are viewed equally.

Publishing

Not all publishing venues are viewed equally.

- Peer-reviewed
  - Conference papers
  - Journals
  - Non-peer-reviewed

Publishing

Not all publishing venues are viewed equally.

- Peer-reviewed
  - Conference papers
    - Premiere tier: CVPR, ICCV, ECCV
  - Journals
  - Non-peer-reviewed

Publishing

Not all publishing venues are viewed equally.

- Peer-reviewed
  - Conference papers
    - Premiere tier: CVPR, ICCV, ECCV
    - Mid-tier: BMVC, ACCV, WACV, and more
  - Journals
  - Non-peer-reviewed

Publishing

Not all publishing venues are viewed equally.

- Peer-reviewed
  - Conference papers
    - Premiere tier: CVPR, ICCV, ECCV
    - Mid-tier: BMVC, ACCV, WACV, and more
    - Local: *e.g.* CRV
  - Journals
  - Non-peer-reviewed

Publishing

Not all publishing venues are viewed equally.

- Peer-reviewed
  - Conference papers
    - Premiere tier: CVPR, ICCV, ECCV
    - Mid-tier: BMVC, ACCV, WACV, and more
    - Local: *e.g.* CRV
  - Journals
    - Prestige largely driven by impact factor, with PAMI and IJCV
- Non-peer-reviewed

Publishing

Not all publishing venues are viewed equally.

- Peer-reviewed
  - Conference papers
    - Premiere tier: CVPR, ICCV, ECCV
    - Mid-tier: BMVC, ACCV, WACV, and more
    - Local: *e.g.* CRV
  - Journals
    - Prestige largely driven by impact factor, with PAMI and IJCV
    - Publications are typically longer than for conference papers, review is more iterative
  - Non-peer-reviewed

Publishing

Not all publishing venues are viewed equally.

- Peer-reviewed
  - Conference papers
    - Premiere tier: CVPR, ICCV, ECCV
    - Mid-tier: BMVC, ACCV, WACV, and more
    - Local: *e.g.* CRV
  - Journals
    - Prestige largely driven by impact factor, with PAMI and IJCV
    - Publications are typically longer than for conference papers, review is more iterative
  - Non-peer-reviewed
    - Preprints (arXiv)

Publishing

Not all publishing venues are viewed equally.

- Peer-reviewed
  - Conference papers
    - Premiere tier: CVPR, ICCV, ECCV
    - Mid-tier: BMVC, ACCV, WACV, and more
    - Local: *e.g.* CRV
  - Journals
    - Prestige largely driven by impact factor, with PAMI and IJCV
    - Publications are typically longer than for conference papers, review is more iterative
  - Non-peer-reviewed
    - Preprints (arXiv)
    - Non-paper conferences

Publishing Credit

Authorship (and author order) are part of the publishing currency.

Publishing Credit

Authorship (and author order) are part of the publishing currency.

If you directly publish (or otherwise publicize) your project work, you *are not required* to include me as an author. If you feel that my supervision and suggestions have substantially contributed to your work, you are welcome to extend authorship to me, but it is entirely up to you.

Publishing Credit

Authorship (and author order) are part of the publishing currency.

If you directly publish (or otherwise publicize) your project work, you *are not required* to include me as an author. If you feel that my supervision and suggestions have substantially contributed to your work, you are welcome to extend authorship to me, but it is entirely up to you.

If you would like to continue developing your work after the end of the course to get it ready for publication, I would be happy to advise and contribute. In such a case, our relationship would at this point not be instructor-student, but would instead be one of collaborators, and I would expect authorship to be extended.

Publishing Credit

Authorship (and author order) are part of the publishing currency.

If you directly publish (or otherwise publicize) your project work, you *are not required* to include me as an author. If you feel that my supervision and suggestions have substantially contributed to your work, you are welcome to extend authorship to me, but it is entirely up to you.

If you would like to continue developing your work after the end of the course to get it ready for publication, I would be happy to advise and contribute. In such a case, our relationship would at this point not be instructor-student, but would instead be one of collaborators, and I would expect authorship to be extended.

NOTE: I must have read and agreed to the submission of anything which has my name on it as a co-author.

Correspondence in Stereo Images

Given two (or more) images, we need to find correspondence between image pixels in order to compute the three dimensional structure of the scene.

Left Image

Right Image

Two common approaches include: keypoint based, area based.

Keypoint Matching

As discussed previously in the course, keypoints can be highly discriminative and will often provide a strong set of correspondence points between images.

Keypoint Matching

As discussed previously in the course, keypoints can be highly discriminative and will often provide a strong set of correspondence points between images.

However, this correspondence is often sparse, and may therefore not provide a sufficiently dense depth map of a scene.

Area Matching

- Divide the images into overlapping patches (typically of fixed size)

Area Matching

- Divide the images into overlapping patches (typically of fixed size)
- For each patch, search for the best matching patch in the other image

Area Matching

- Divide the images into overlapping patches (typically of fixed size)
- For each patch, search for the best matching patch in the other image
- We need to pick a similarity measure to quantify degree of patch matching

Area Similarity Measure

Cross-Correlation

- Simple to compute
- Requires normalization of patch values

Sum of squared difference

- No need to normalize
- Slightly more computationally expensive than cross-correlation

Pixel-wise correspondence can be computationally expensive

As with most problems in computer vision, the number of operations required grows rapidly with image resolution and patch size or number of keypoints found.

Pixel-wise correspondence can be computationally expensive

As with most problems in computer vision, the number of operations required grows rapidly with image resolution and patch size or number of keypoints found.

One potential solution is to approximately solve the problem at a low resolution (whether reducing image resolution or skipping intervening steps), then iteratively refine the solution as one moves back toward full resolution.

Coarse-to-fine Example

Say we are attempting to match a 20×20 patch within a 200×200 image.

This requires $(180)(180)(20^2) = 12960000$ pixel-wise operations.

Coarse-to-fine Example

Say we are attempting to match a 20×20 patch within a 200×200 image.

This requires $(180)(180)(20^2) = 12960000$ pixel-wise operations.

If we instead reduce our resolution by a factor of 4, we now are searching for a 5×5 pixel patch in a 50×50 image. Finding our closest match here requires $(45)(45)(5^2) = 50625$ pixel-wise operations.

Coarse-to-fine Example

Say we are attempting to match a 20×20 patch within a 200×200 image.

This requires $(180)(180)(20^2) = 12960000$ pixel-wise operations.

If we instead reduce our resolution by a factor of 4, we now are searching for a 5×5 pixel patch in a 50×50 image. Finding our closest match here requires $(45)(45)(5^2) = 50625$ pixel-wise operations.

Of course, this low resolution match hasn't fully solved our problem, we've just found the neighbourhood where the proper match should be. But refining that match at high resolution only requires searching over sixteen different possible placements of our patch, rather than the entire image.

Epipolar Diagram

The diagram below shows two image views with the pin-hole assumption and the virtual image plane shown in blue.

When the origin of the left and right image are connected, the points where this line intersects the image plane are the *epipoles* (or *epipolar points*).

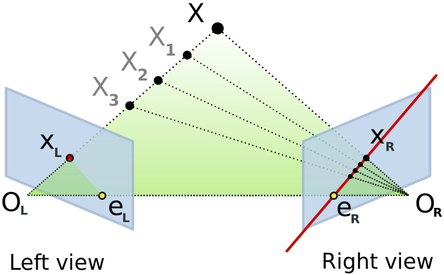

Image source: Wikipedia

Epipolar Diagram

For a point $\mathbf{X}$ which lies directly on the optical axis of the left image, it will always land on the same location within the left image regardless of distance. Depending on the distance, however, it will land at a different location in the right image along a line running through the right epipole.

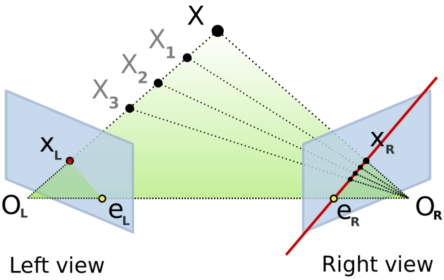

Image source: Wikipedia

Epipolar Diagram

For a point $\mathbf{X}$ which lies directly on the optical axis of the left image, it will always land on the same location within the left image regardless of distance. Depending on the distance, however, it will land at a different location in the right image along a line running through the right epipole.

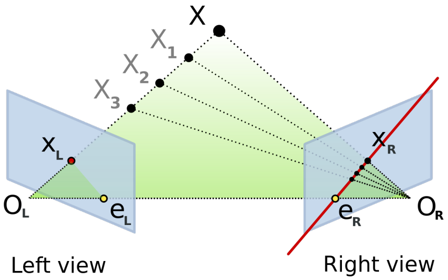

Image source: Wikipedia

This effectively reduces the search for correspondence down to a 1D search along the epipolar line.