



CS5670: Computer Vision

Noah Snavely

Lecture 15: Structure from motion



Readings

- Szeliski, Chapter 7.1 – 7.4


Structure from motion

- Given many images, how can we
 - a) figure out where they were all taken from?
 - b) build a 3D model of the scene?




This is (roughly) the **structure from motion** problem

Structure from motion



Reconstruction (side)



(top)

- Input: images with points in correspondence


$$p_{i,j} = (u_{i,j}, v_{i,j})$$

- Output
 - structure: 3D location \mathbf{x}_i for each point p_i
 - motion: camera parameters \mathbf{R}_j , \mathbf{t}_j possibly \mathbf{K}_j
- Objective function: minimize *reprojection error*


Input



Feature detection



Feature matching



Camera calibration & triangulation

- Suppose we know 3D points
 - And have matches between these points and an image
 - How can we compute the camera parameters?
- Suppose we have know camera parameters, each of which observes a point
 - How can we compute the 3D location of that point?

Structure from motion

- SfM solves both of these problems *at once*
- A kind of chicken-and-egg problem
 - (but solvable)

Photo Tourism



First step: how to get correspondence?

- Feature detection and matching


Feature detection

Detect features using SIFT [Lowe, IJCV 2004]



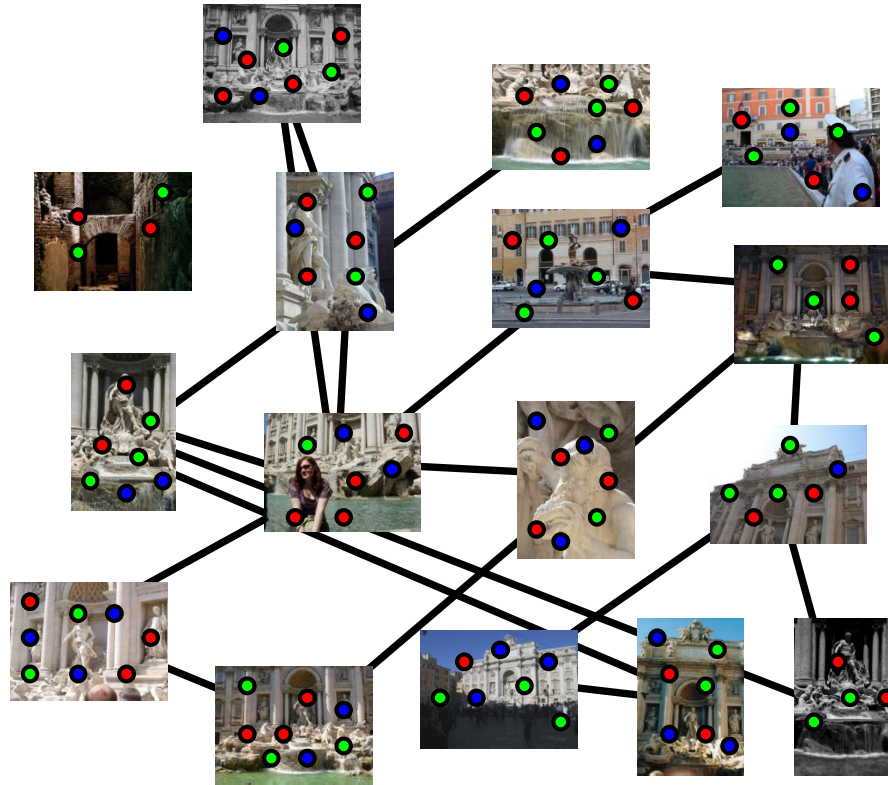
Feature detection

Detect features using SIFT [Lowe, IJCV 2004]




Feature matching

Match features between each pair of images



Feature matching

Refine matching using RANSAC to estimate fundamental matrix between each pair



Correspondence estimation

- Link up pairwise matches to form connected components of matches across several images



Image 1



Image 2




Image 3





Image 4

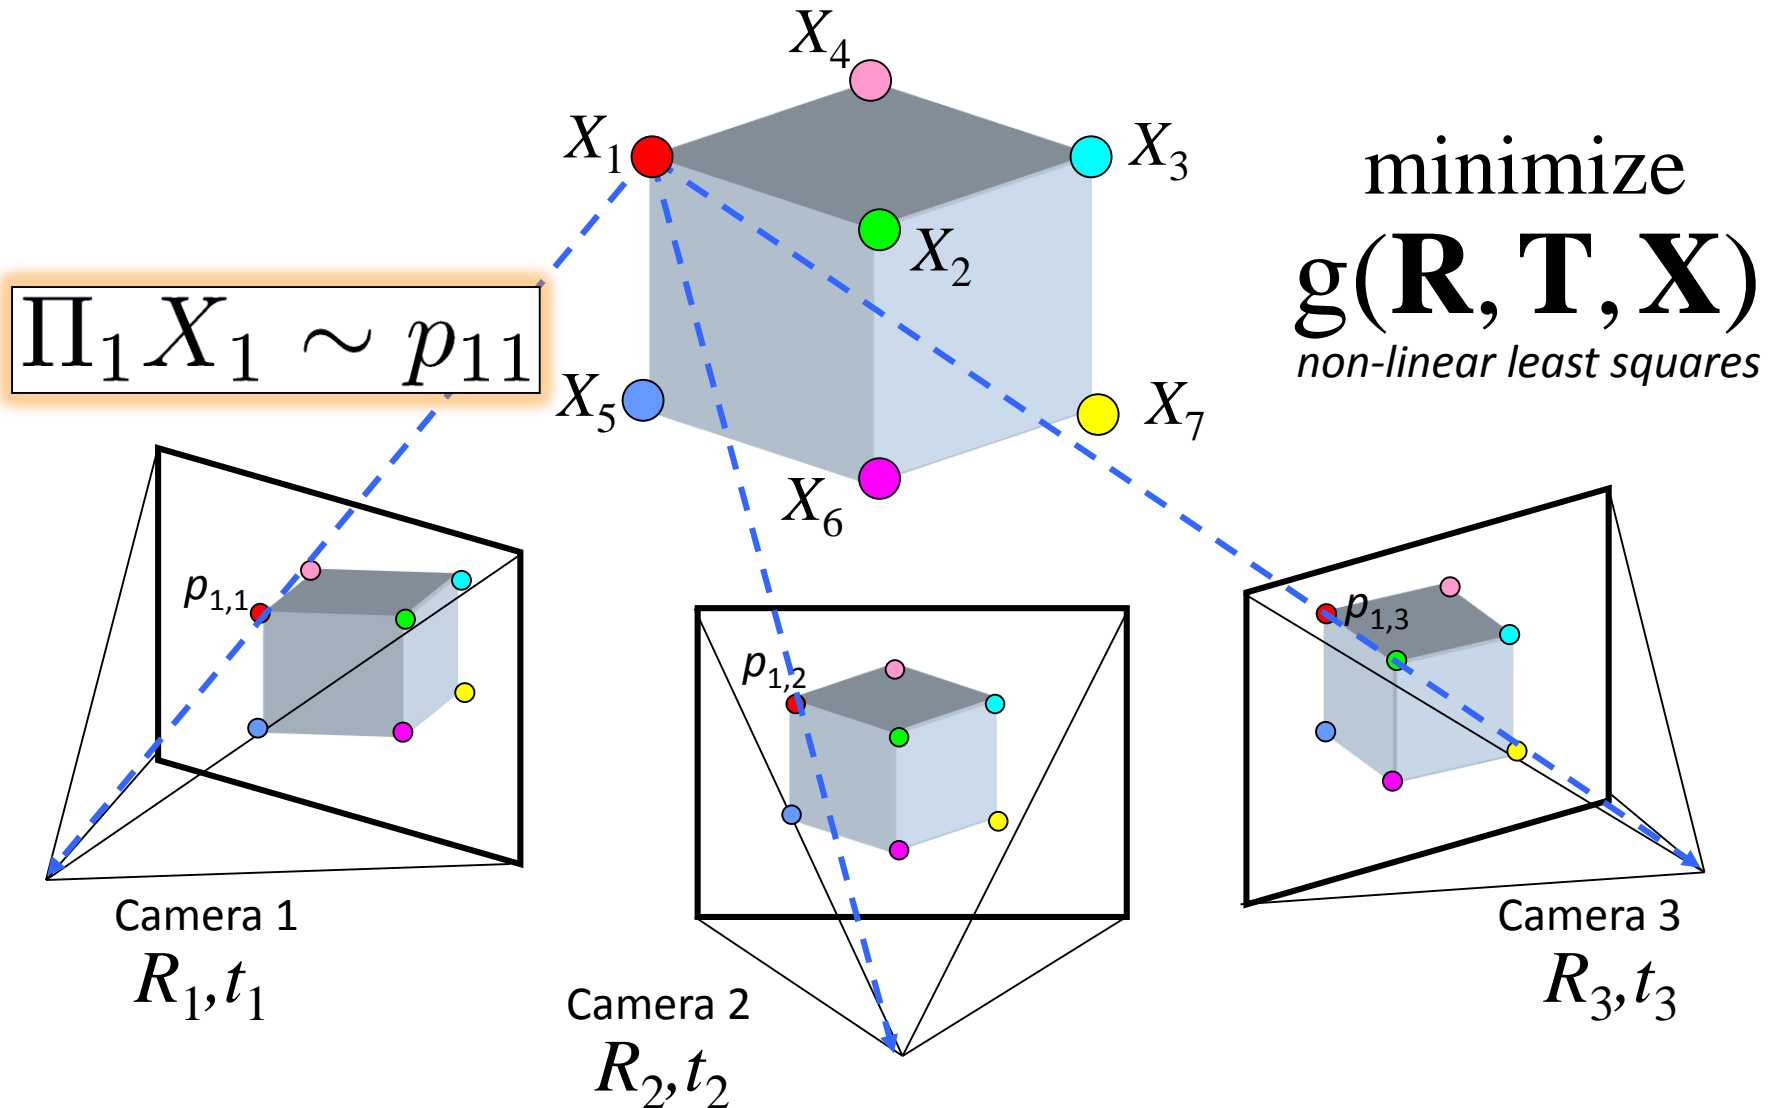


Image connectivity graph



(graph layout produced using the Graphviz toolkit: <http://www.graphviz.org/>)

Structure from motion



Problem size

- What are the variables?
- How many variables per camera?
- How many variables per point?

- Trevi Fountain collection
 - 466 input photos
 - + > 100,000 3D points
 - = very large optimization problem

Structure from motion

- Minimize sum of squared reprojection errors:


$$g(\mathbf{X}, \mathbf{R}, \mathbf{T}) = \sum_{i=1}^m \sum_{j=1}^n \underbrace{w_{ij}}_{\substack{\text{indicator variable:} \\ \text{is point } i \text{ visible in image } j?}} \cdot \left\| \underbrace{\mathbf{P}(\mathbf{x}_i, \mathbf{R}_j, \mathbf{t}_j)}_{\substack{\text{predicted} \\ \text{image location}}} - \underbrace{\begin{bmatrix} u_{i,j} \\ v_{i,j} \end{bmatrix}}_{\substack{\text{observed} \\ \text{image location}}} \right\|^2$$

- Minimizing this function is called *bundle adjustment*
 - Optimized using non-linear least squares, e.g. Levenberg-Marquardt


Is SfM always uniquely solvable?

Is SfM always uniquely solvable?

- No...



Incremental structure from motion



Incremental structure from motion




Incremental structure from motion











Incremental structure from motion

Photo Explorer



Demo





File Edit View History Bookmarks Tools Help

W http://en.wikipedia.org/wiki/Libration

Most Visited Latest Headlines Welcome to Ithaca C... Photosynth Constitution Room Lily Chaplin Cornell University Re... Photo City

Loading... Dulmage-Men... Errors-in-vari... Kim_CVPR08... CS1112 Fall 2... Libratio... Spring09 CS1... Nokia Challen... moon.jpg (JP...



- navigation
- [Main page](#)
 - [Contents](#)
 - [Featured content](#)
 - [Current events](#)
 - [Random article](#)

search

- interaction
- [About Wikipedia](#)
 - [Community portal](#)
 - [Recent changes](#)
 - [Contact Wikipedia](#)
 - [Donate to Wikipedia](#)
 - [Help](#)

- toolbox
- [What links here](#)
 - [Related changes](#)
 - [Upload file](#)
 - [Special pages](#)
 - [Printable version](#)
 - [Permanent link](#)

- [article](#) [discussion](#) [edit this page](#) [history](#)

Libration

From Wikipedia, the free encyclopedia

Not to be confused with [Liberation](#) or [Libation](#).

In **astronomy libration** (from the Latin verb *librare* "to balance, to sway", cf. *libra* "scales") refers to the various orbital conditions which make it possible to see more than 50% of the moon's surface over time, even though the front of the Moon is tidally locked to always face towards Earth. By extension, libration can also be used to describe the same phenomenon for other orbital bodies that are nominally locked to present the same face. As the orbital processes are repetitive, libration is manifested as a slow rocking back and forth (or up and down) of the face of the orbital body as viewed from the parent body, much like the rocking of a pair of scales about the point of balance.

In the specific case of the Moon's librations, this motion permits a terrestrial observer to see slightly differing halves of the Moon's surface at different times. This means that a total of 59% of the Moon's surface can be observed from Earth.

There are three types of libration:

- *Libration in longitude* is a consequence of the Moon's orbit around Earth being somewhat **eccentric**, so that the Moon's rotation sometimes leads and sometimes lags its orbital position.
- *Libration in latitude* is a consequence of the Moon's axis of rotation being slightly inclined to the **normal** to the **plane** of its **orbit** around Earth. Its origin is analogous to the way in which the **seasons** arise from Earth's revolution about the Sun.
- *Diurnal libration* is a small daily oscillation due to the Earth's rotation, which carries an observer first to one side and then to the other side of the straight line joining Earth's center to the Moon's center, allowing the observer to look first around one side of the Moon and then around the other. This is because the observer is on the surface of the Earth, not at its centre.





Simulated views of the Moon over one month, demonstrating librations in latitude and longitude.

Questions?


SfM – Failure cases

- Necker reversal



Structure from Motion – Failure cases

- Repetitive structures



SfM applications

- 3D modeling
- Surveying
- Robot navigation and mapmaking
- Visual effects (“Match moving”)
 - https://www.youtube.com/watch?v=RdYWp70P_kY

Applications – Photosynth



Applications – Hyperlapse



<https://www.youtube.com/watch?v=SOpwHaQnRSY>