ETH zürich[1]

ECCV
EUROPEAN CONFERENCE ON COMPUTER VISION

Microsoft[2]

GeoCalib

Learning Single-image Calibration with Geometric Optimization




Alexander Veicht[1]
Paul-Edouard Sarlin[1]
Philipp Lindenberger[1]
Marc Pollefeys[1,2]

github.com/cvg/GeoCalib

Camera Calibration

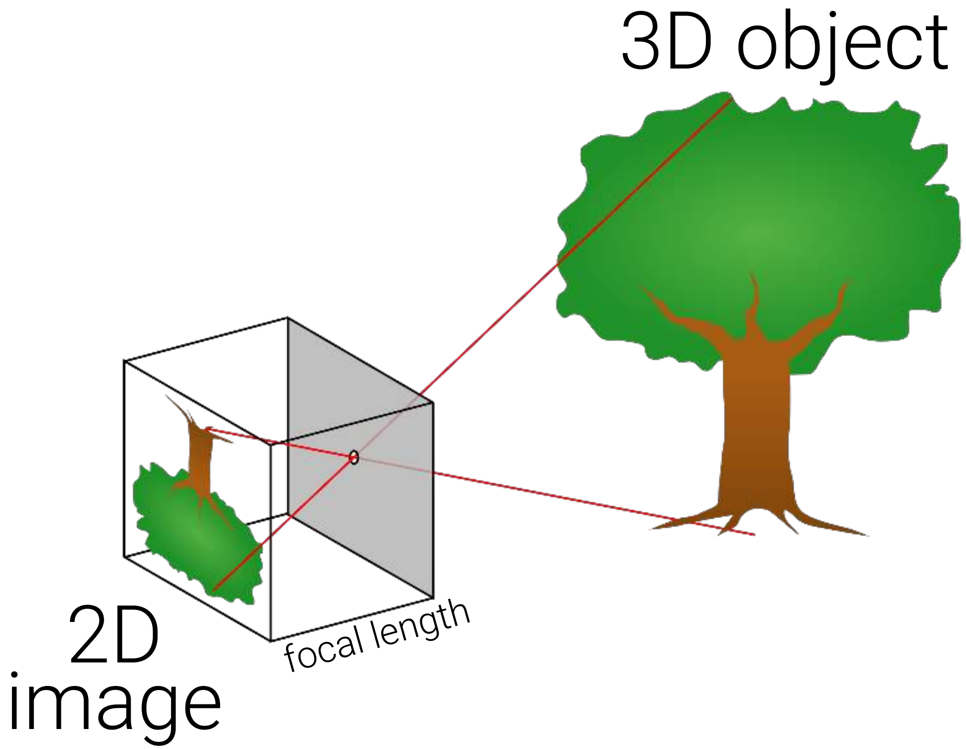

focal
length

intrinsic
parameters

lens
distortion

Controlled Conditions

OpenCV

Kalibr

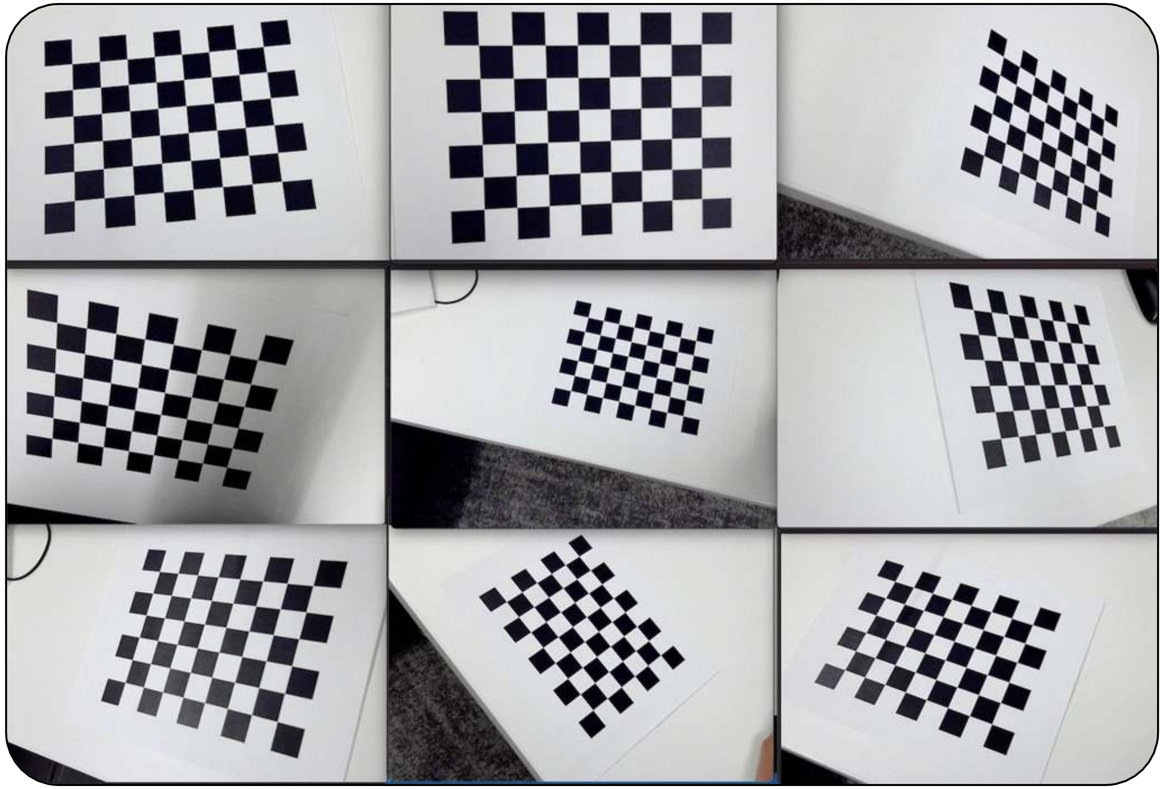

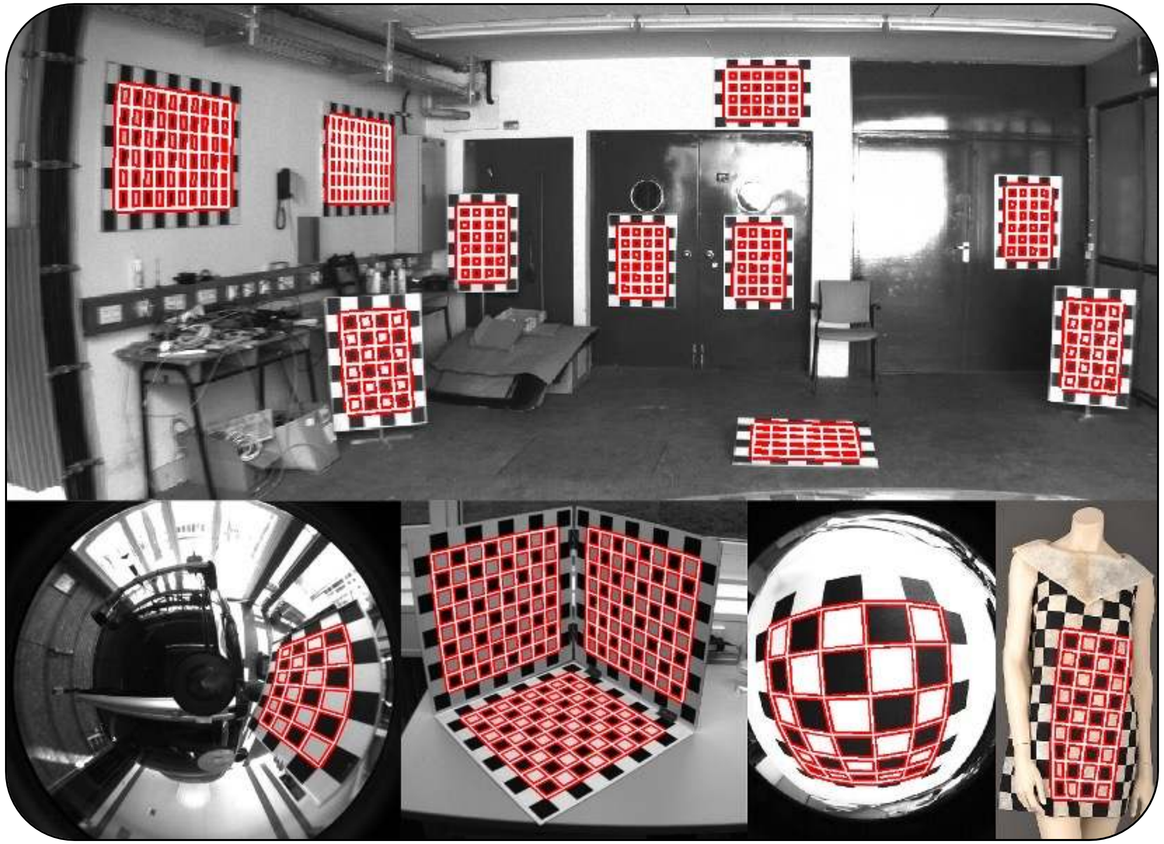

Uncontrolled Conditions

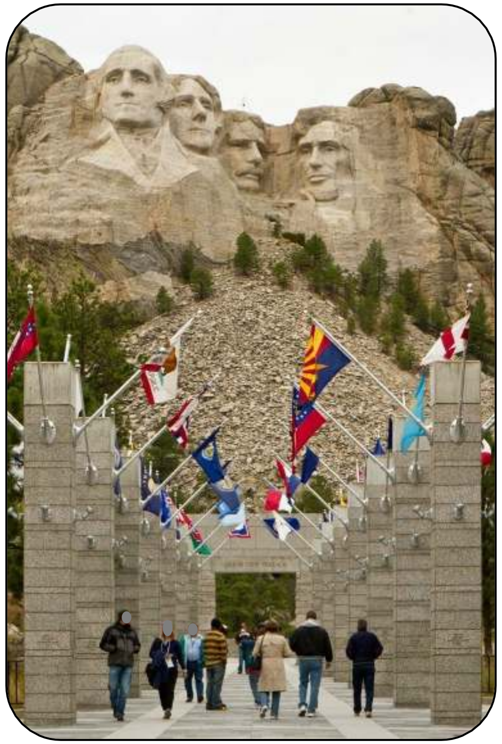




Photogrammetry
Structure-from-Motion

Uncontrolled Conditions

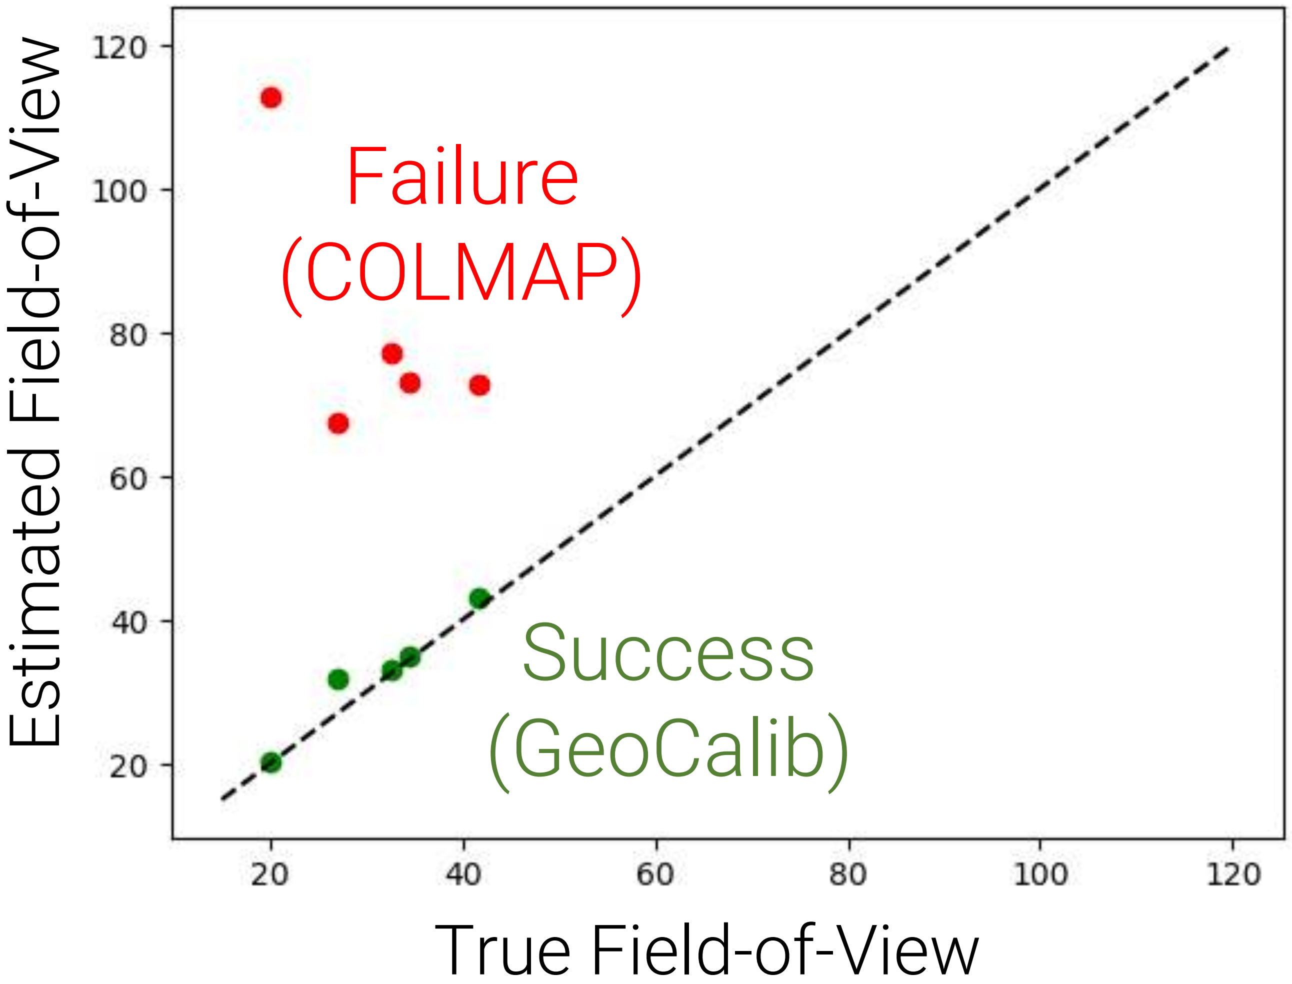

Camera Calibration

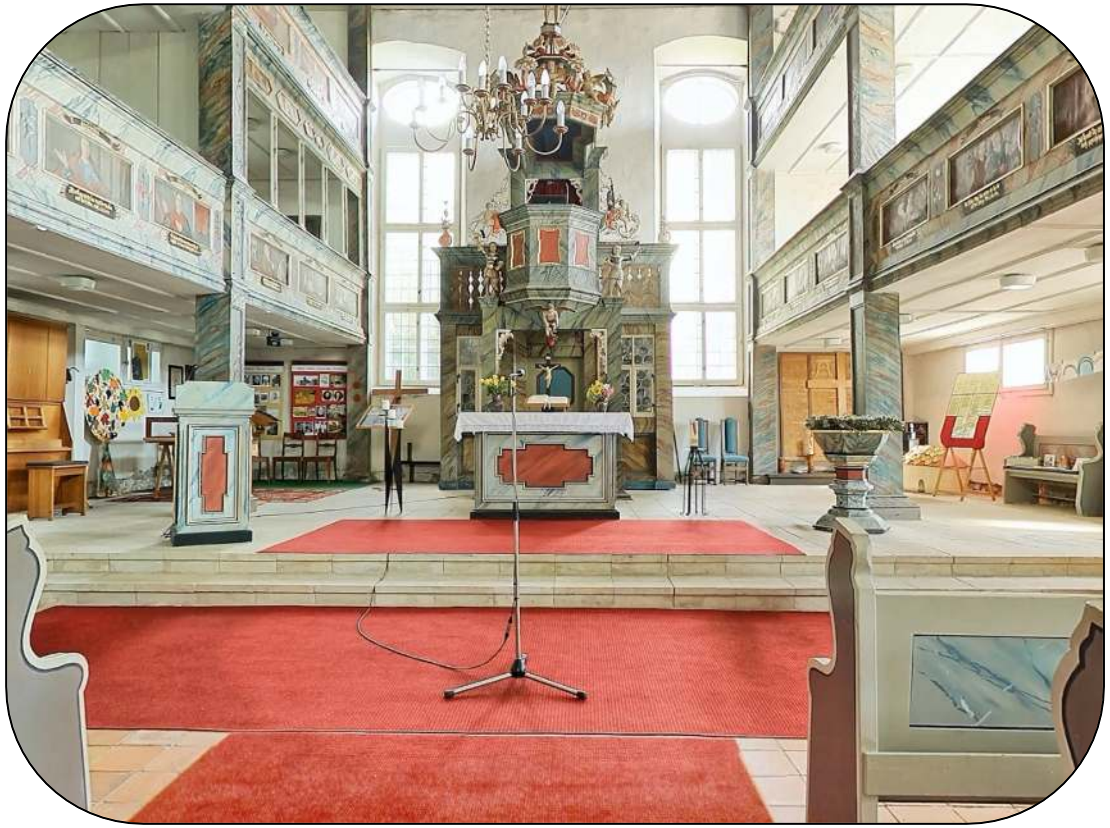

roll

pitch

extrinsic parameters = gravity direction

Single-image Camera Calibration

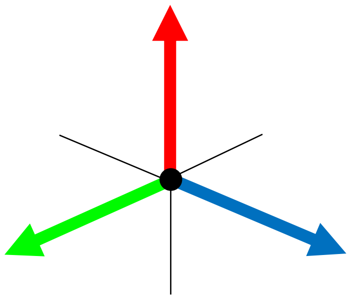

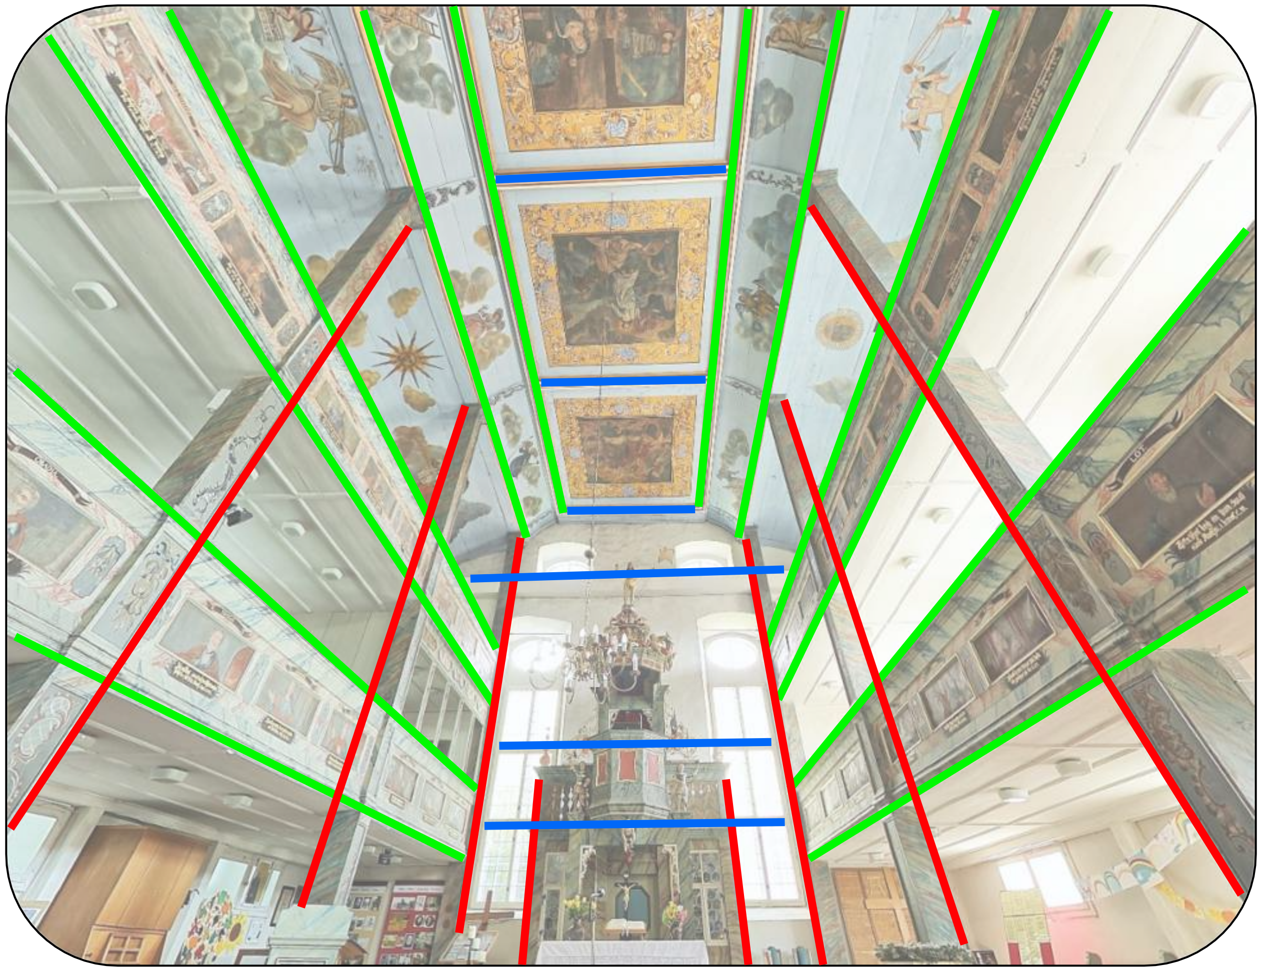

gravity & intrinsic parameters

Existing algorithms

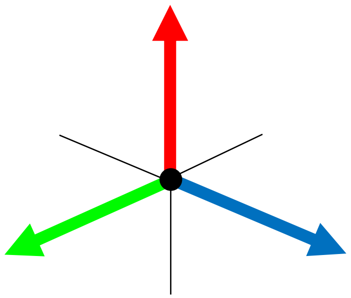

GT and estimated horizon lines

many lines → accurate

Existing algorithms

many lines → accurate

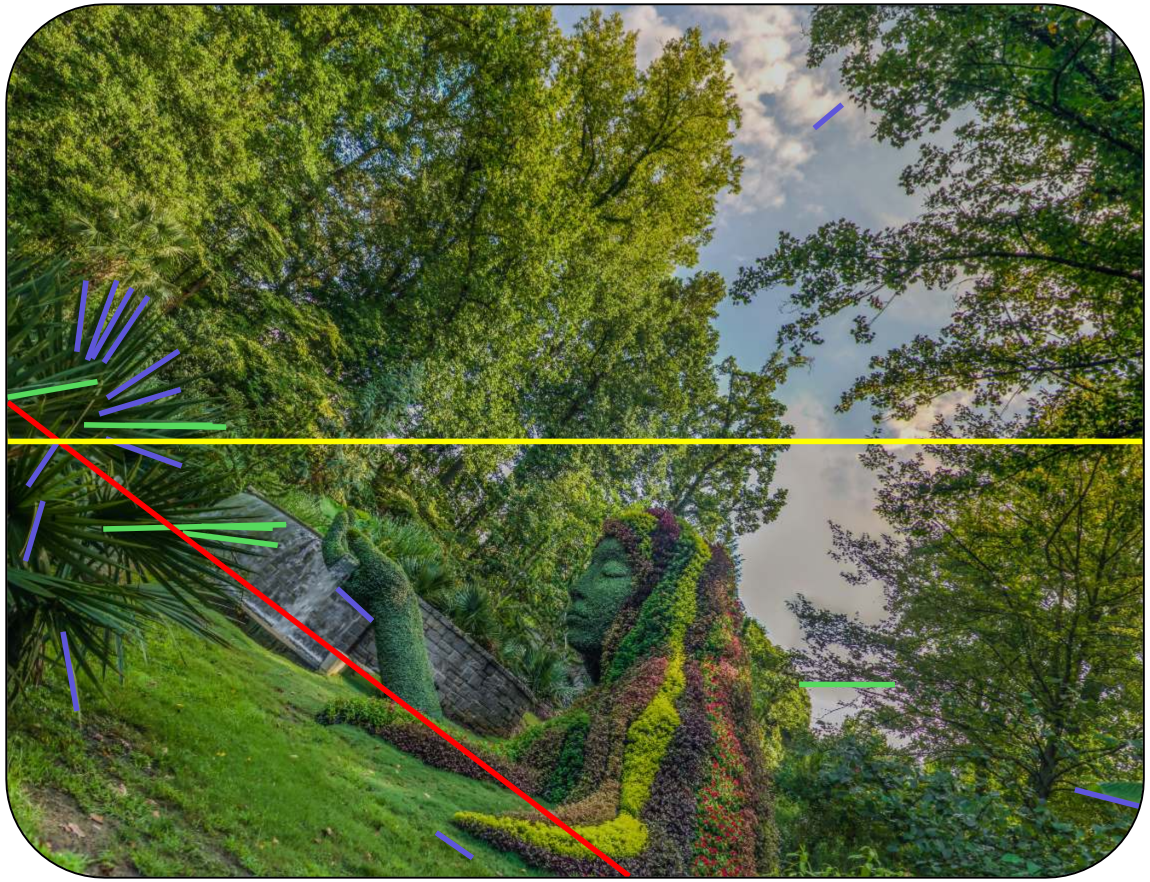

few lines → failure

Deep Learning

black-box
Deep Neural
Network

gravity
& intrinsic
parameters

Deep Learning

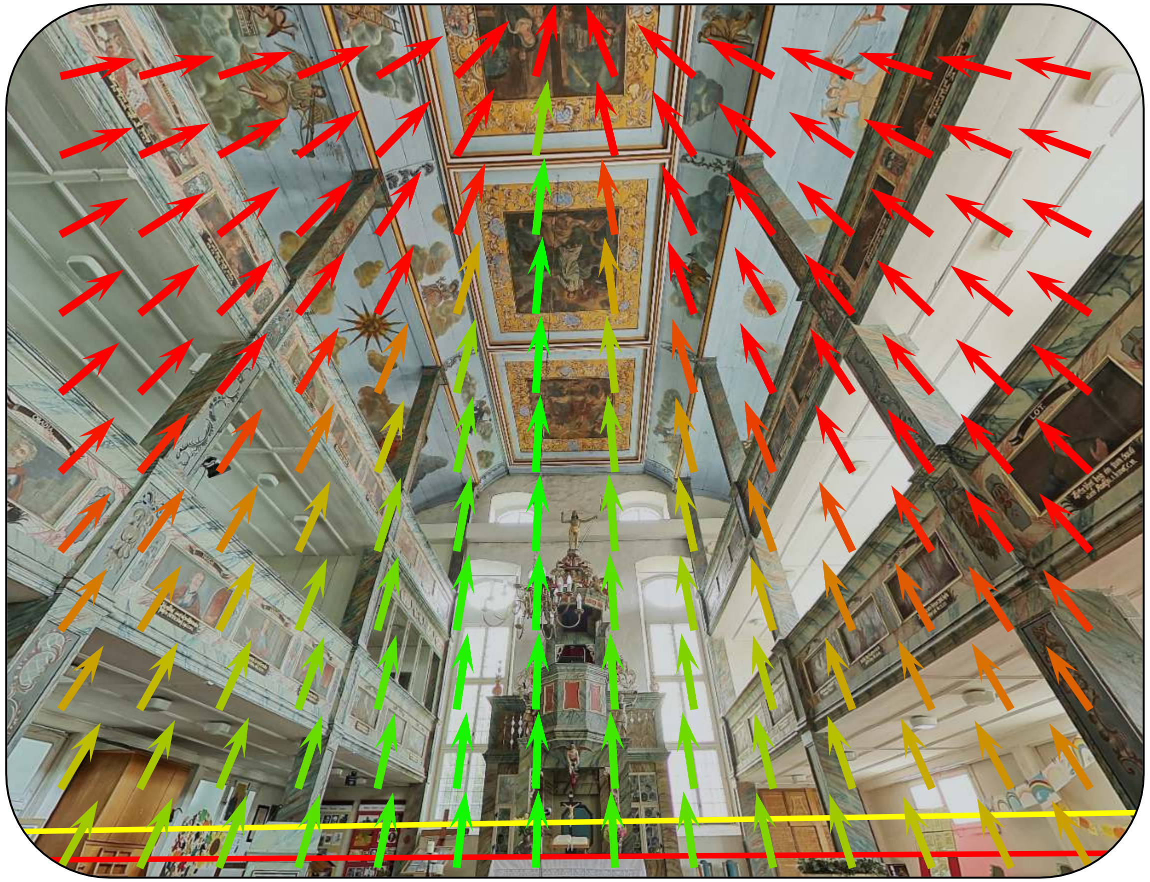

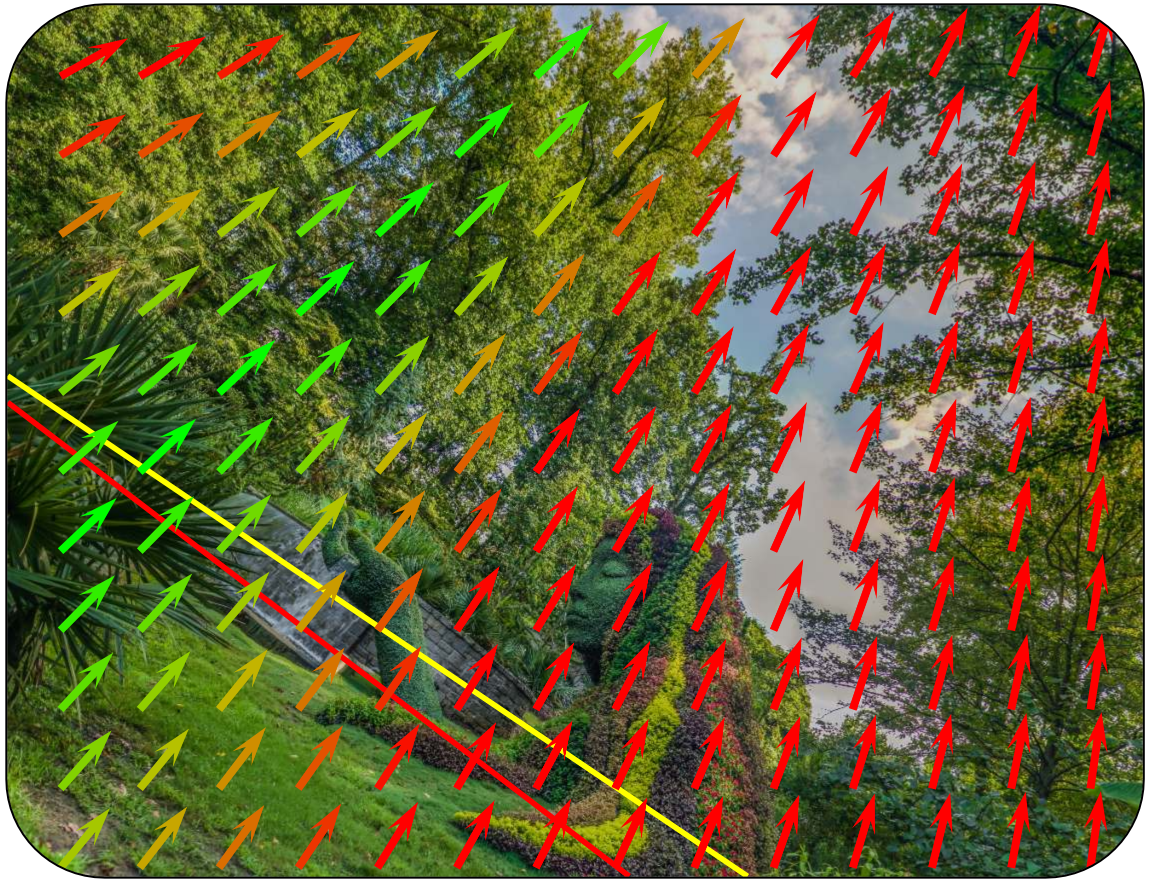

more robust but less accurate

GeoCalib = Learning + Geometry

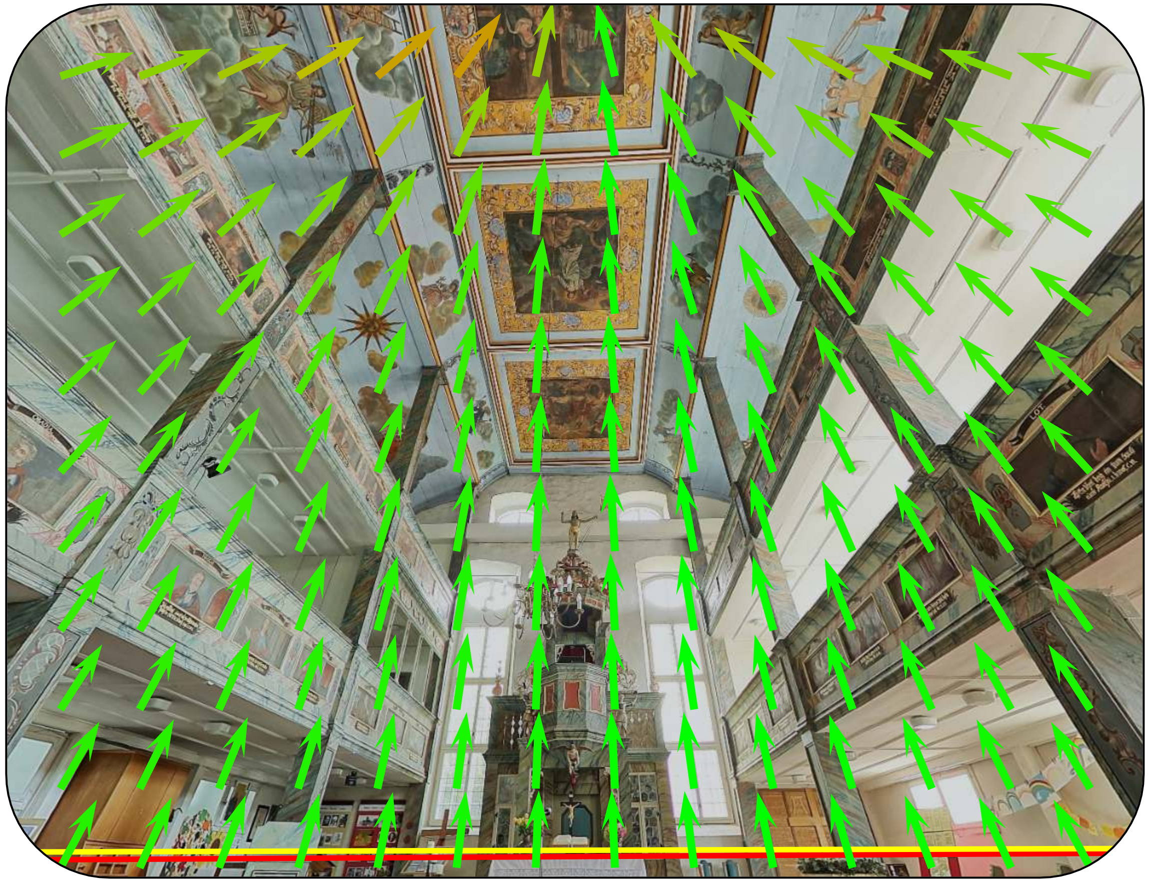

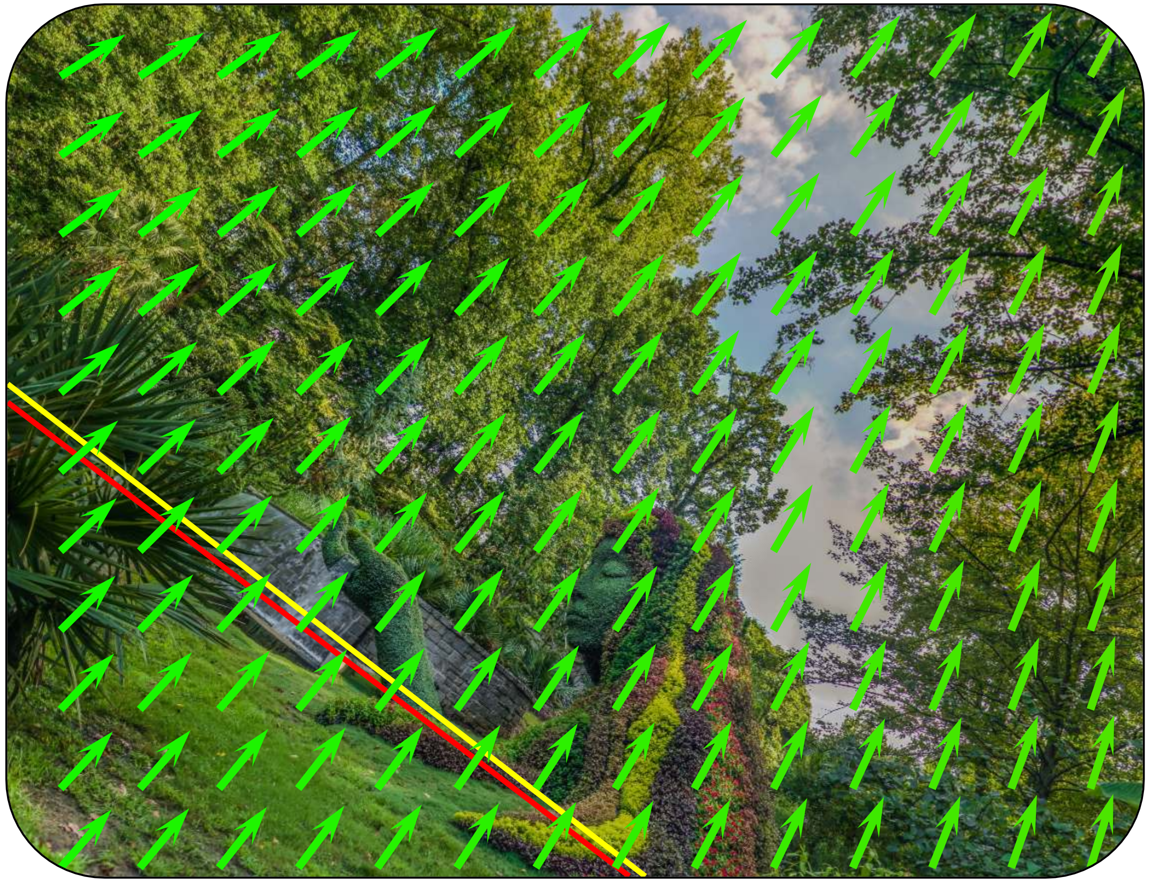

both more robust and more accurate

Perspective Field [Jin et al, CVPR 2023]

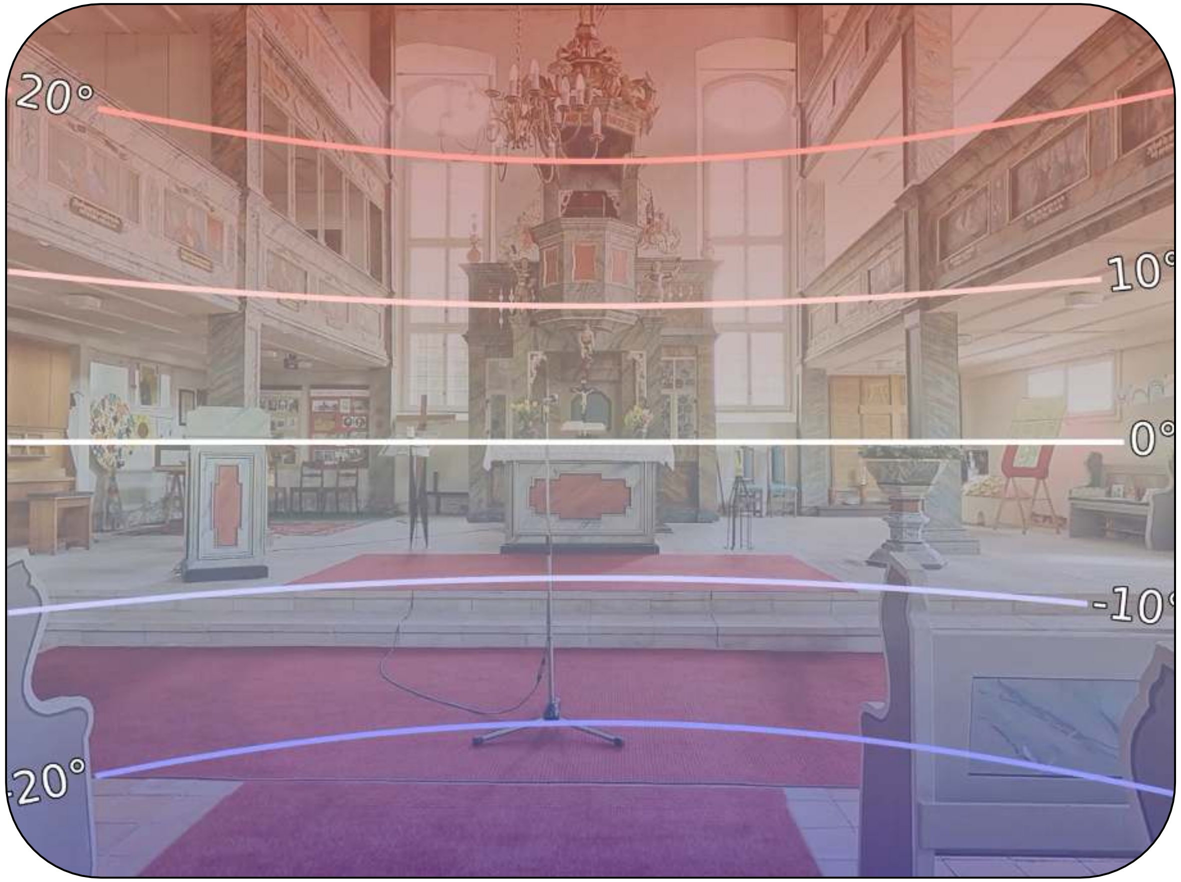

Latitude

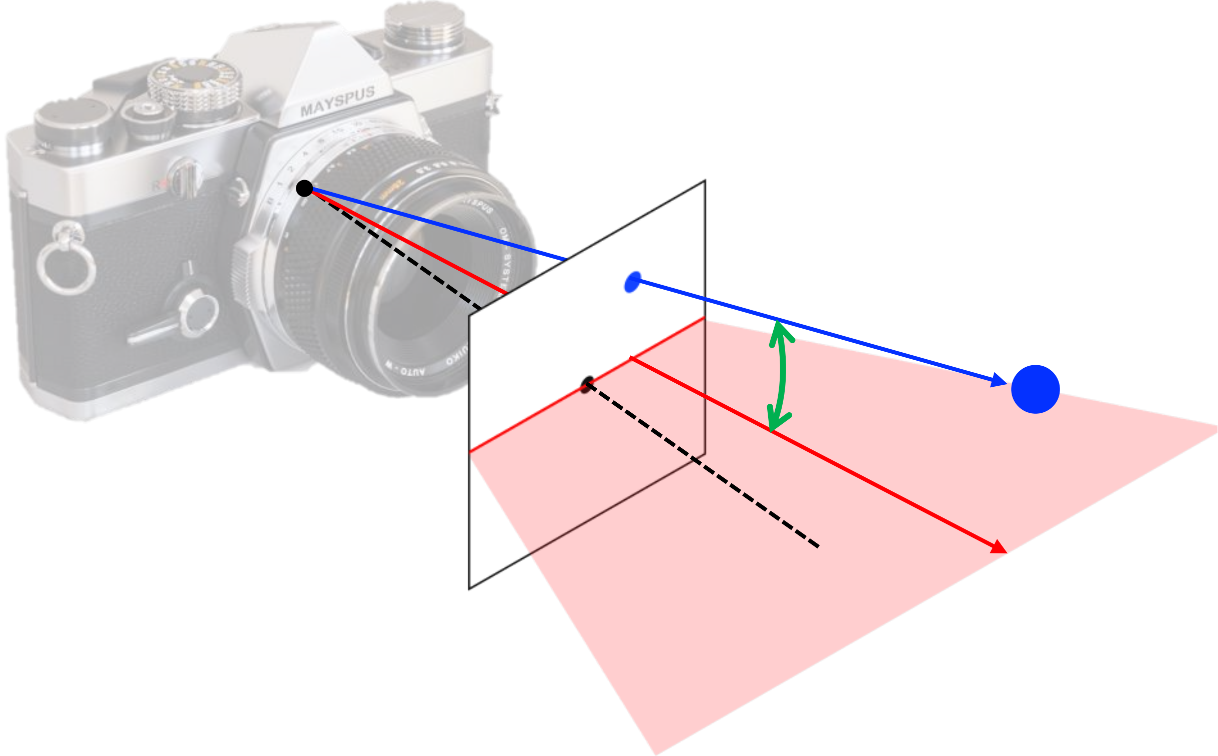

Perspective Field [Jin et al, CVPR 2023]

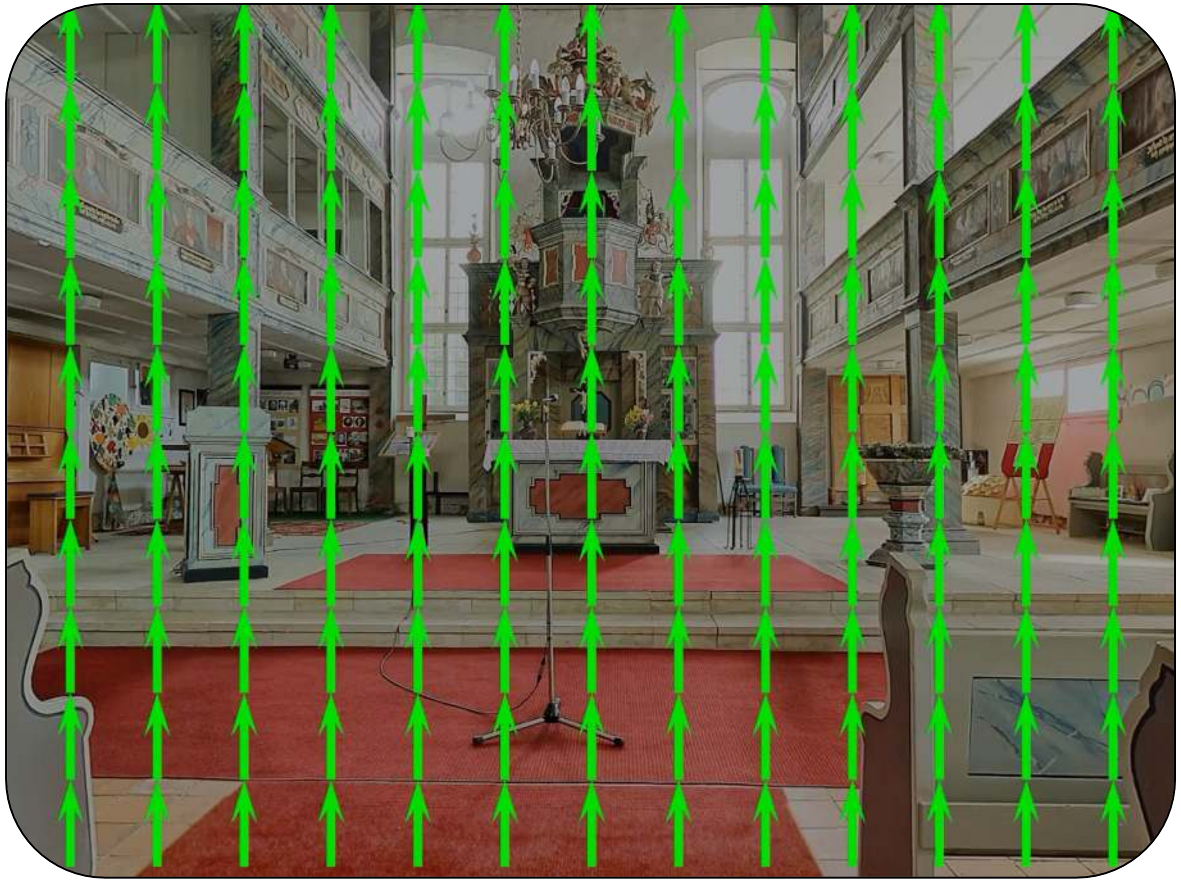

Up-vector

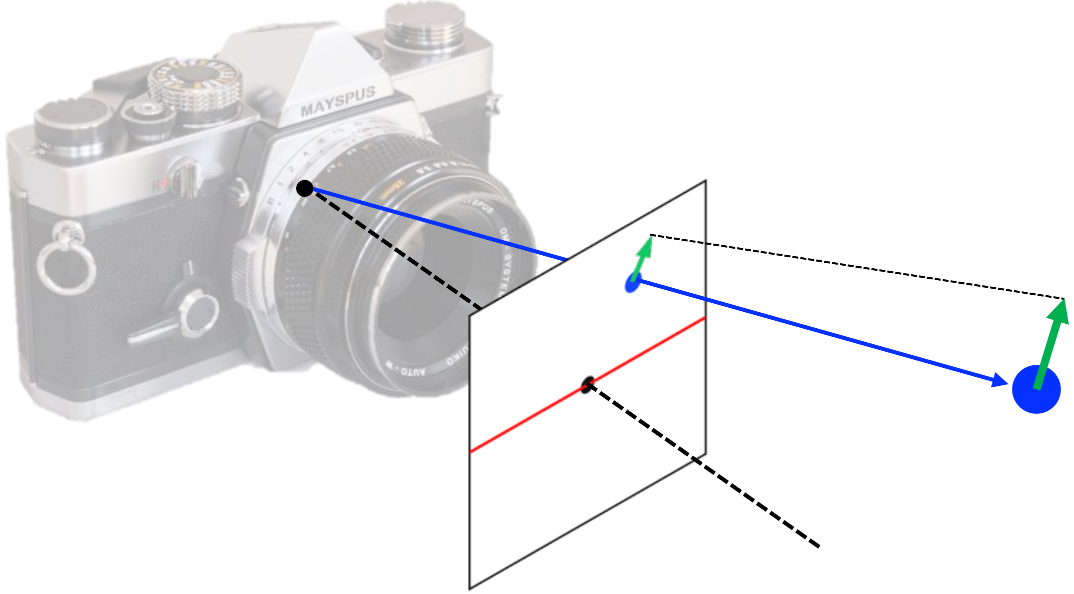

GeoCalib – Architecture

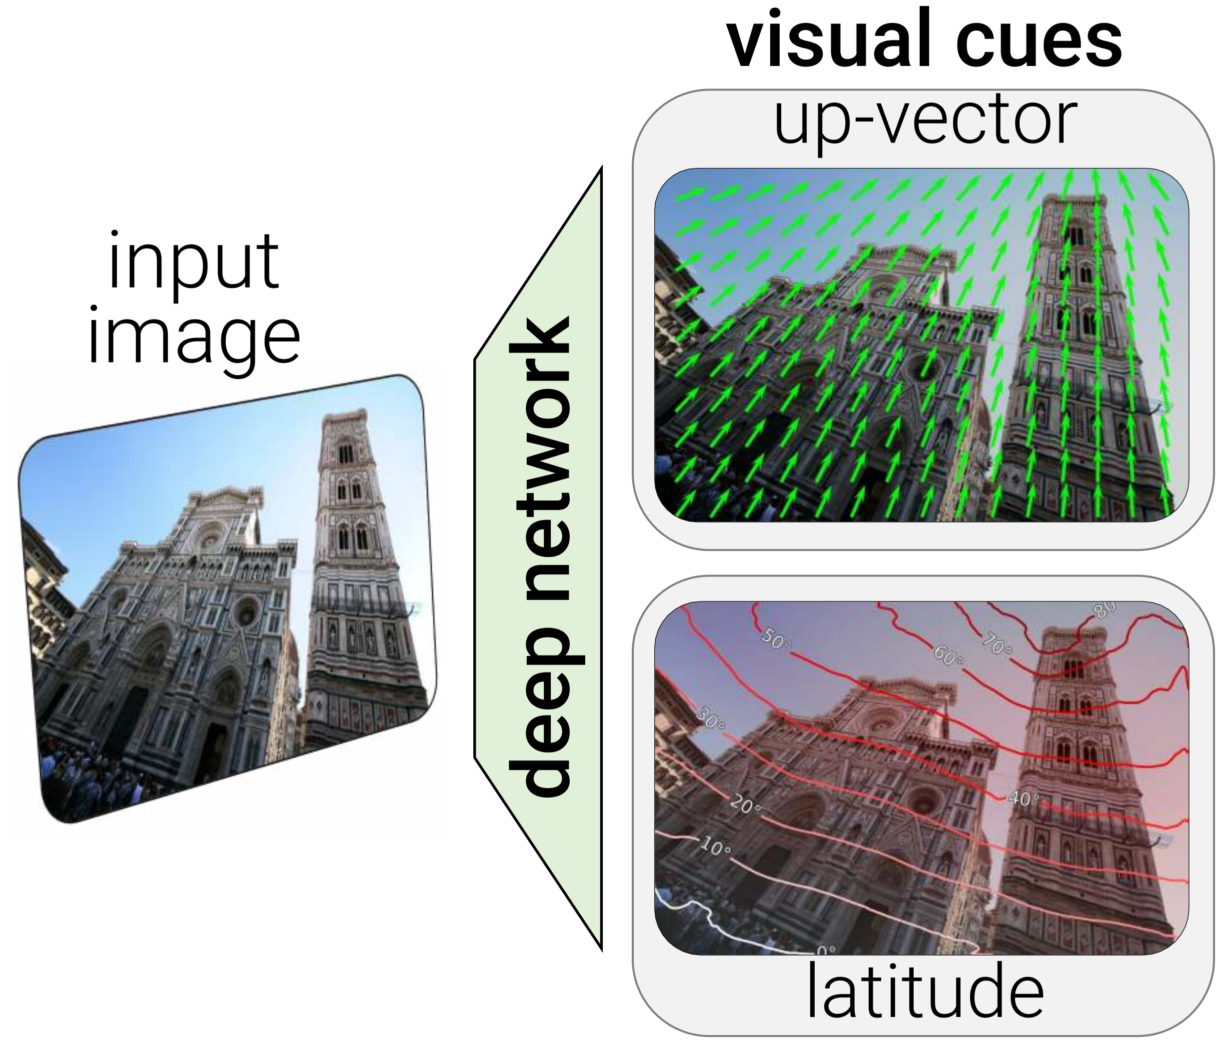

GeoCalib – Architecture

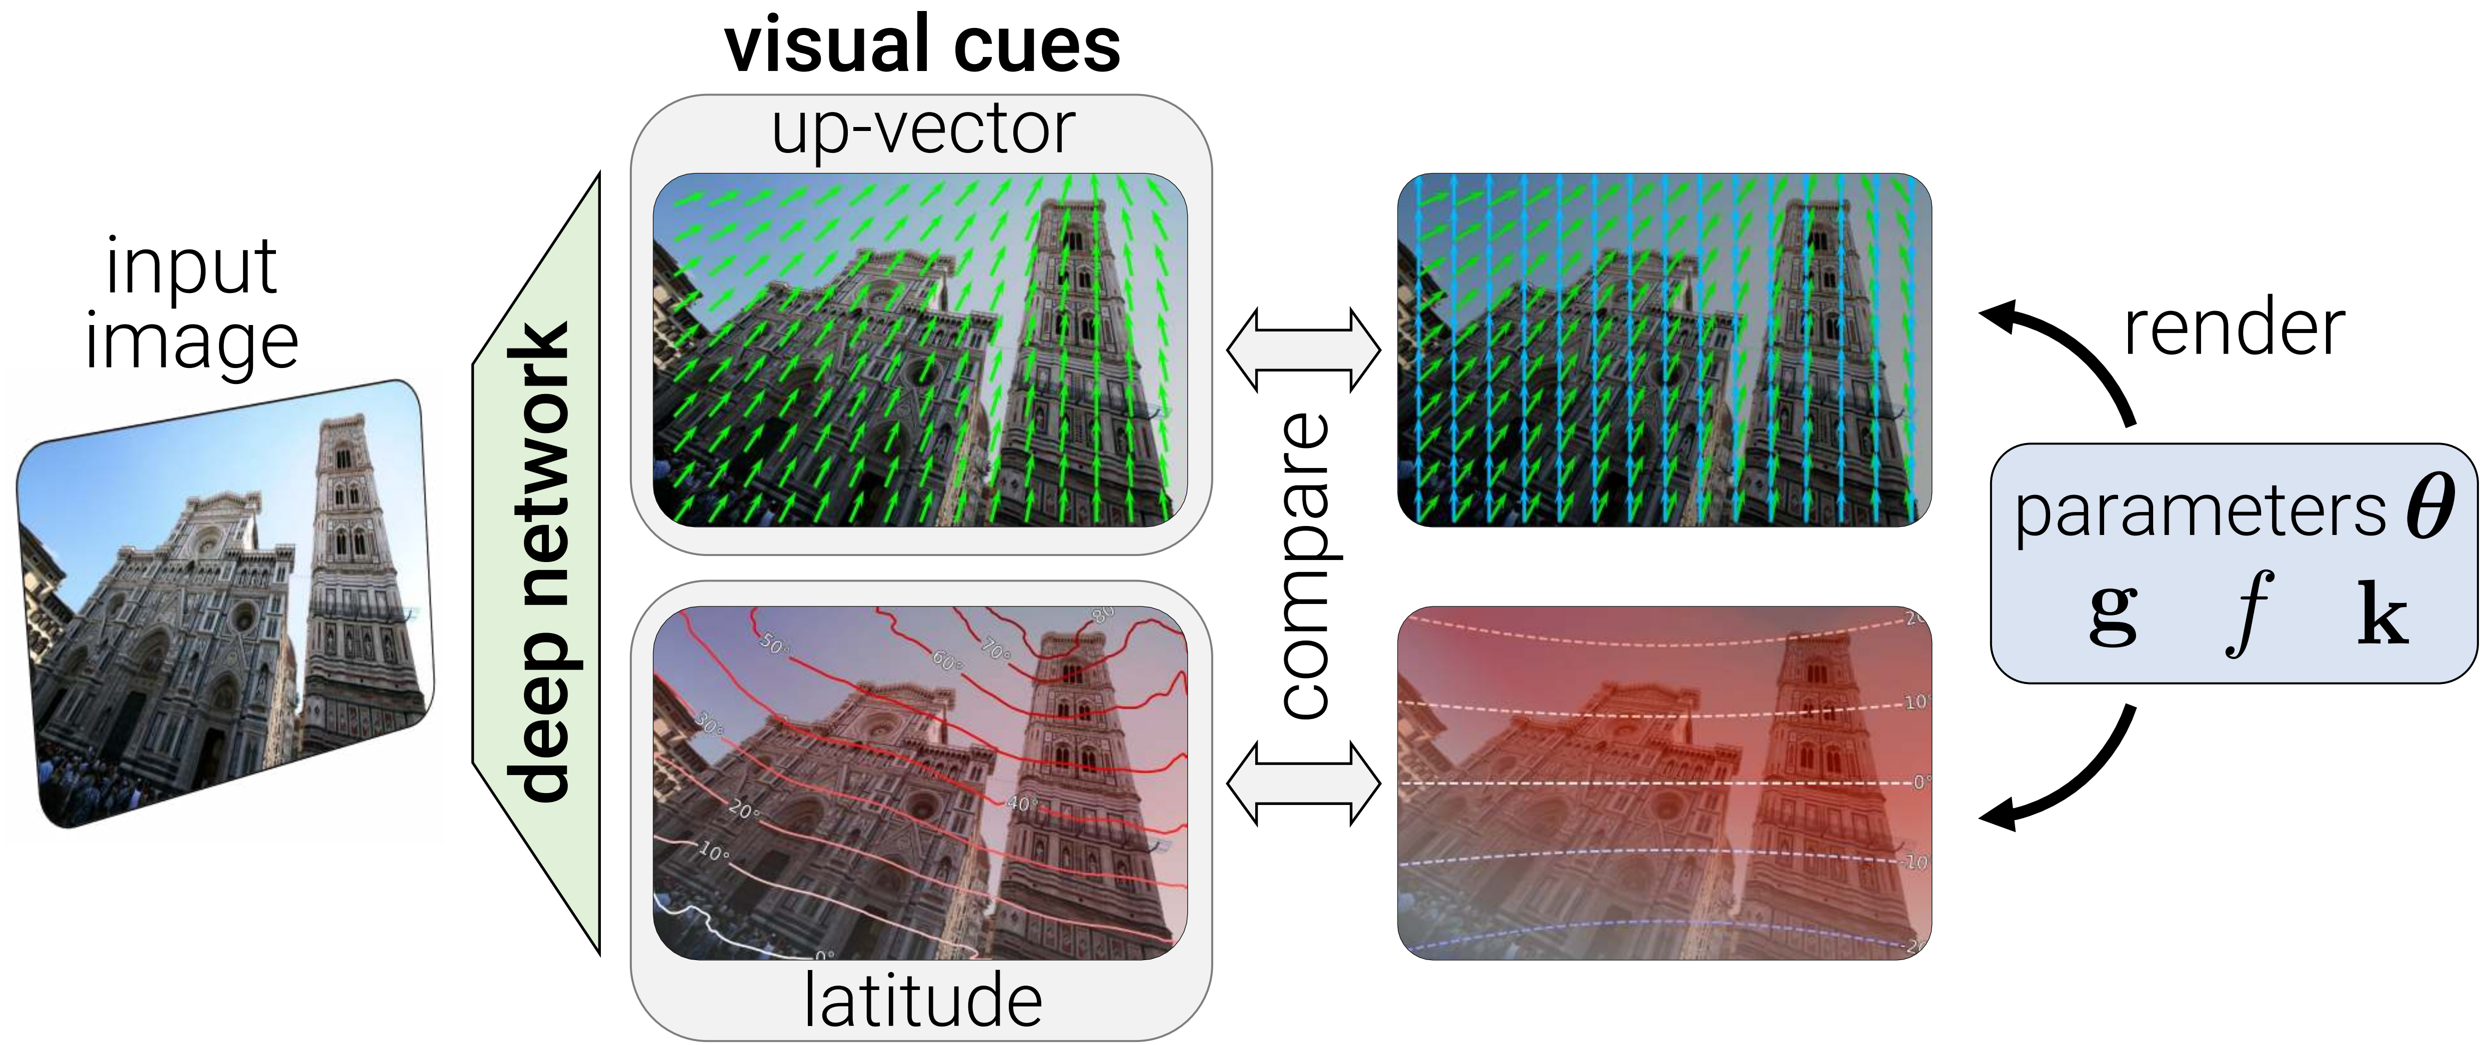

GeoCalib – Architecture

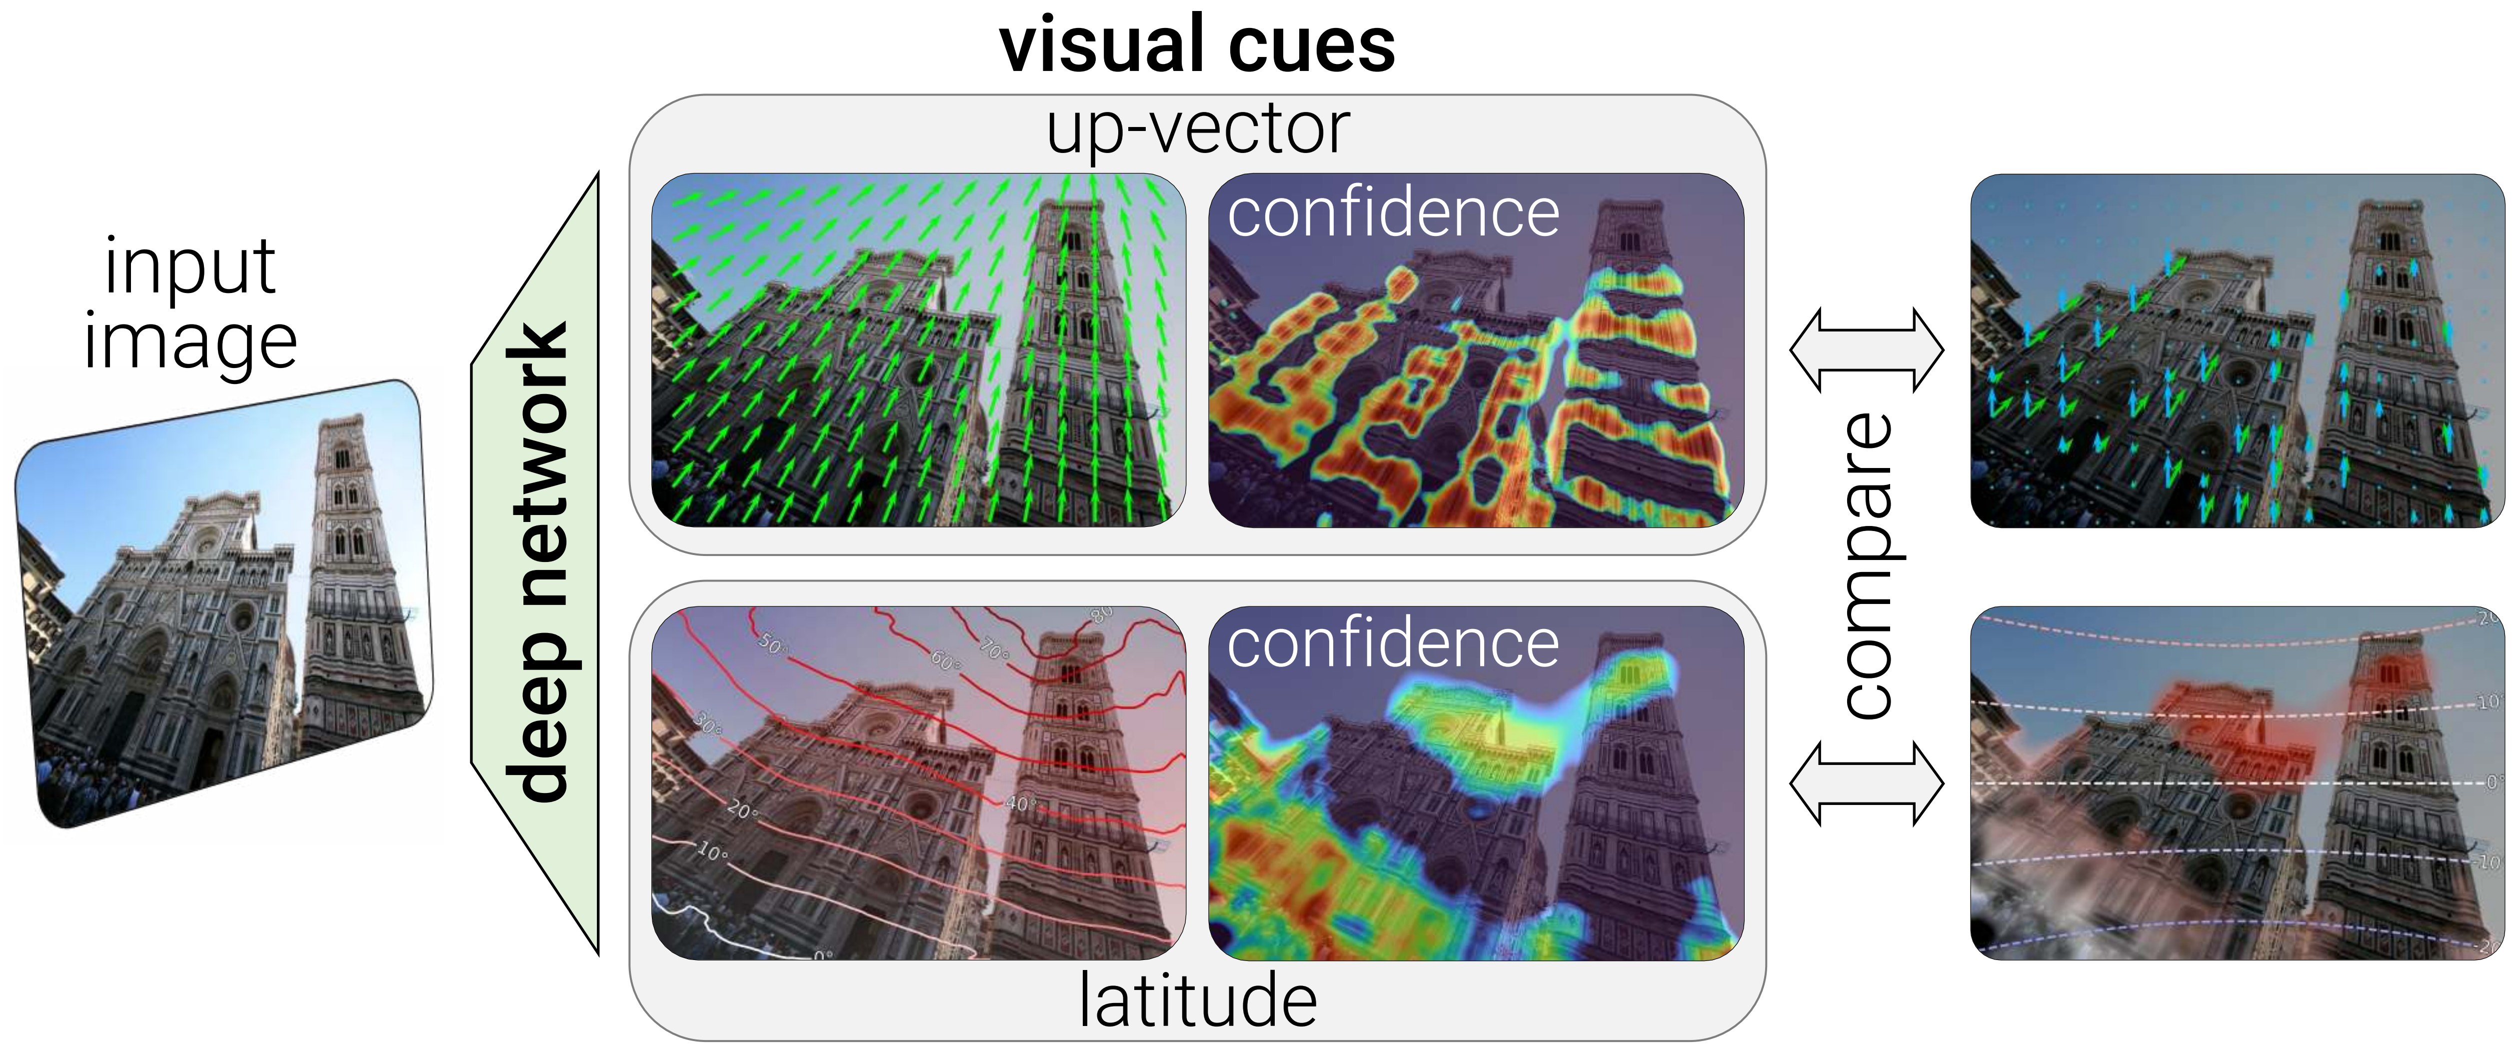

GeoCalib – Architecture

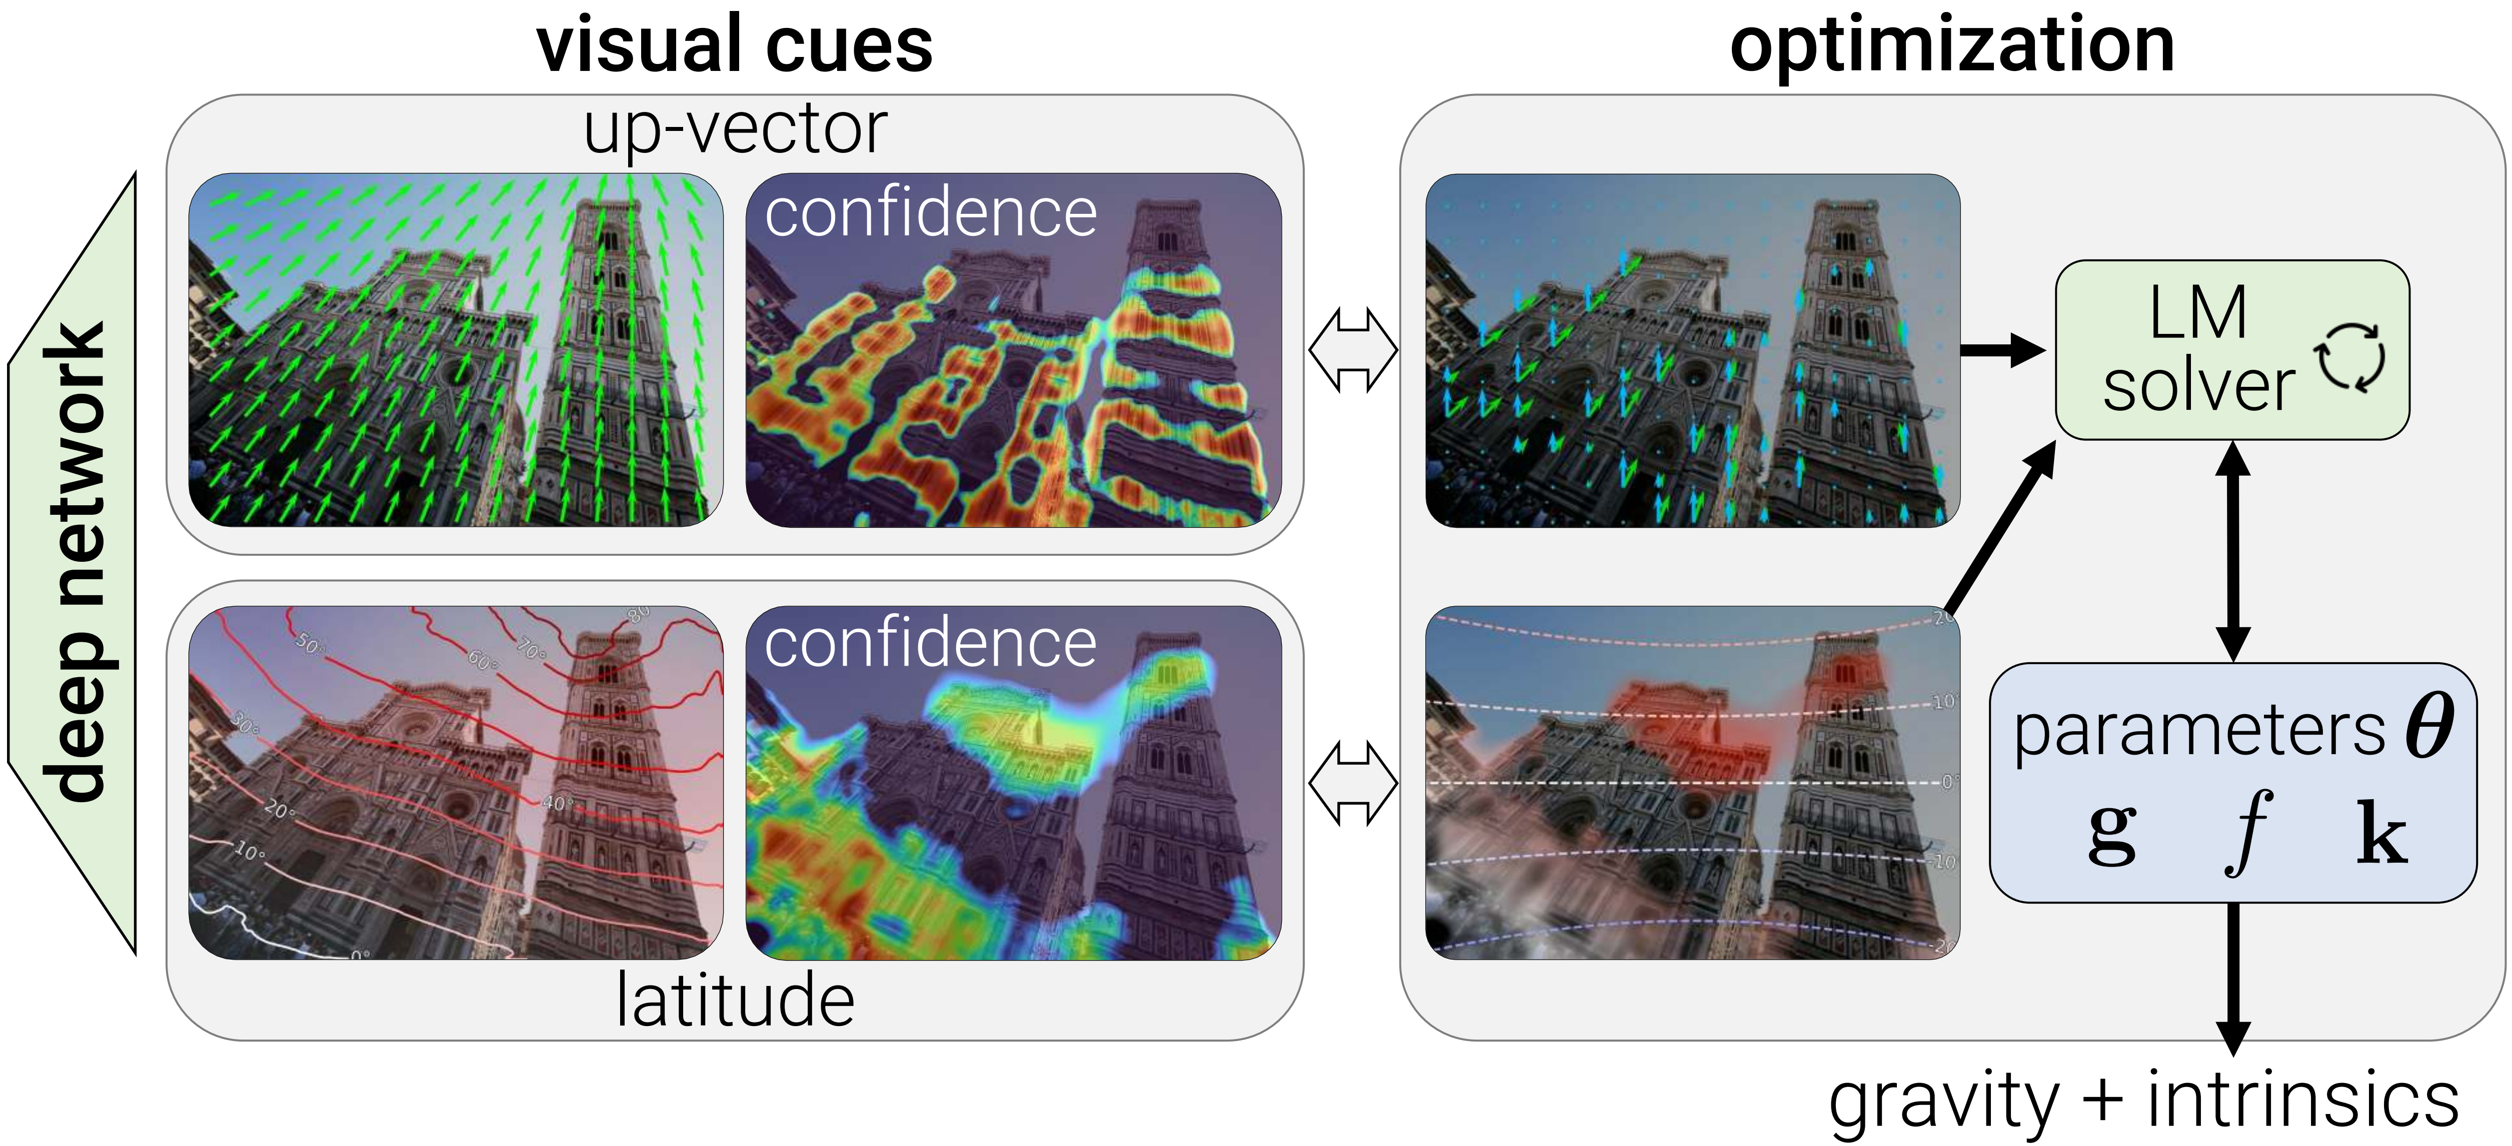

GeoCalib – Architecture

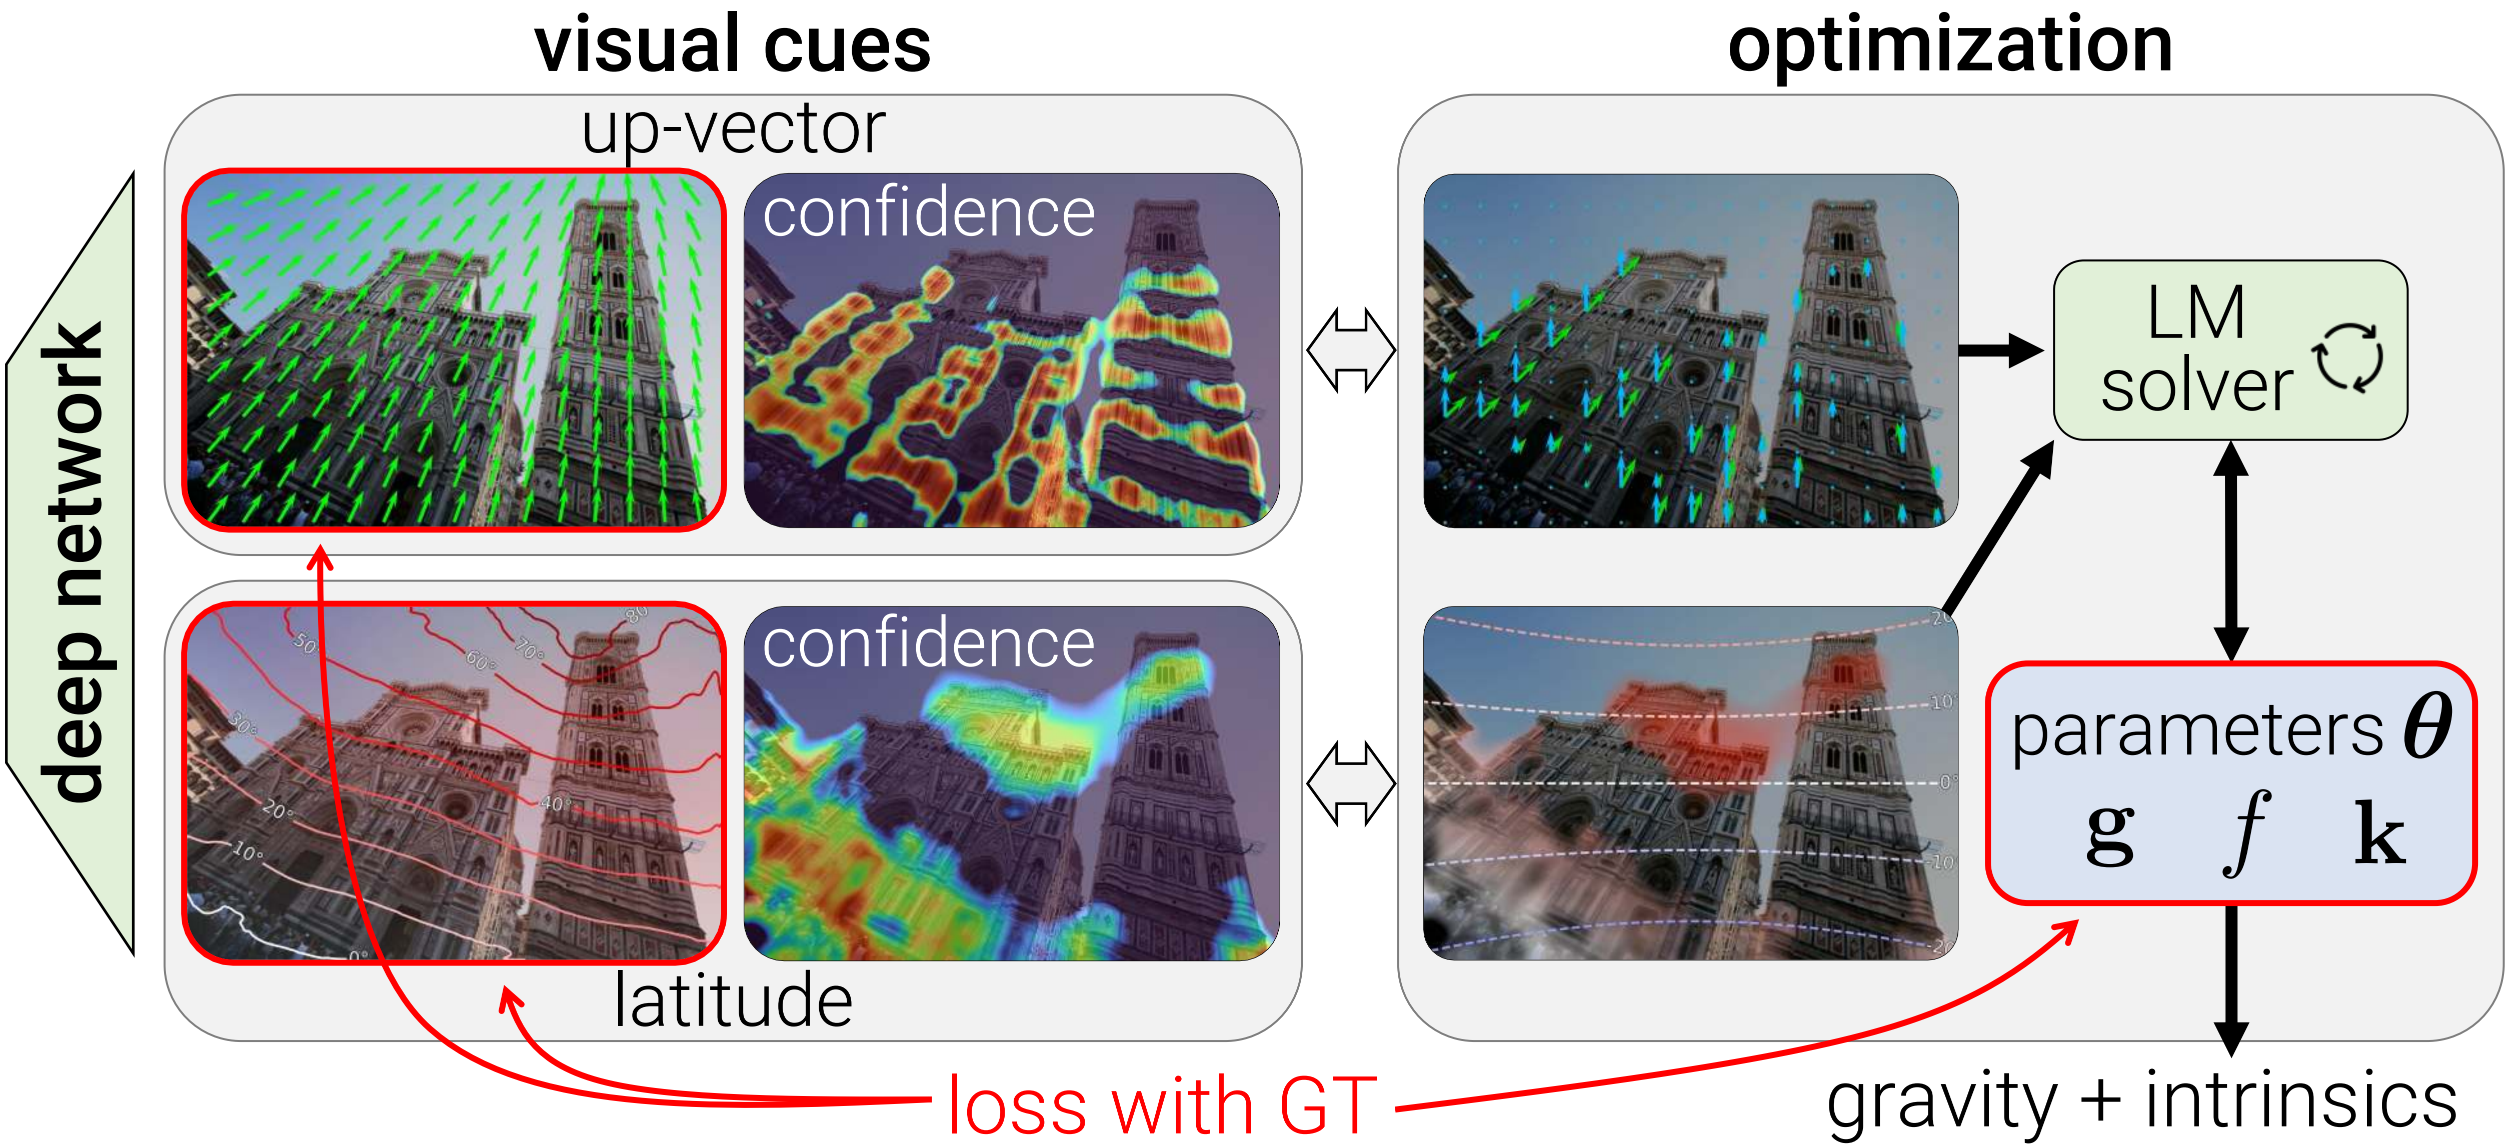

Learned confidences

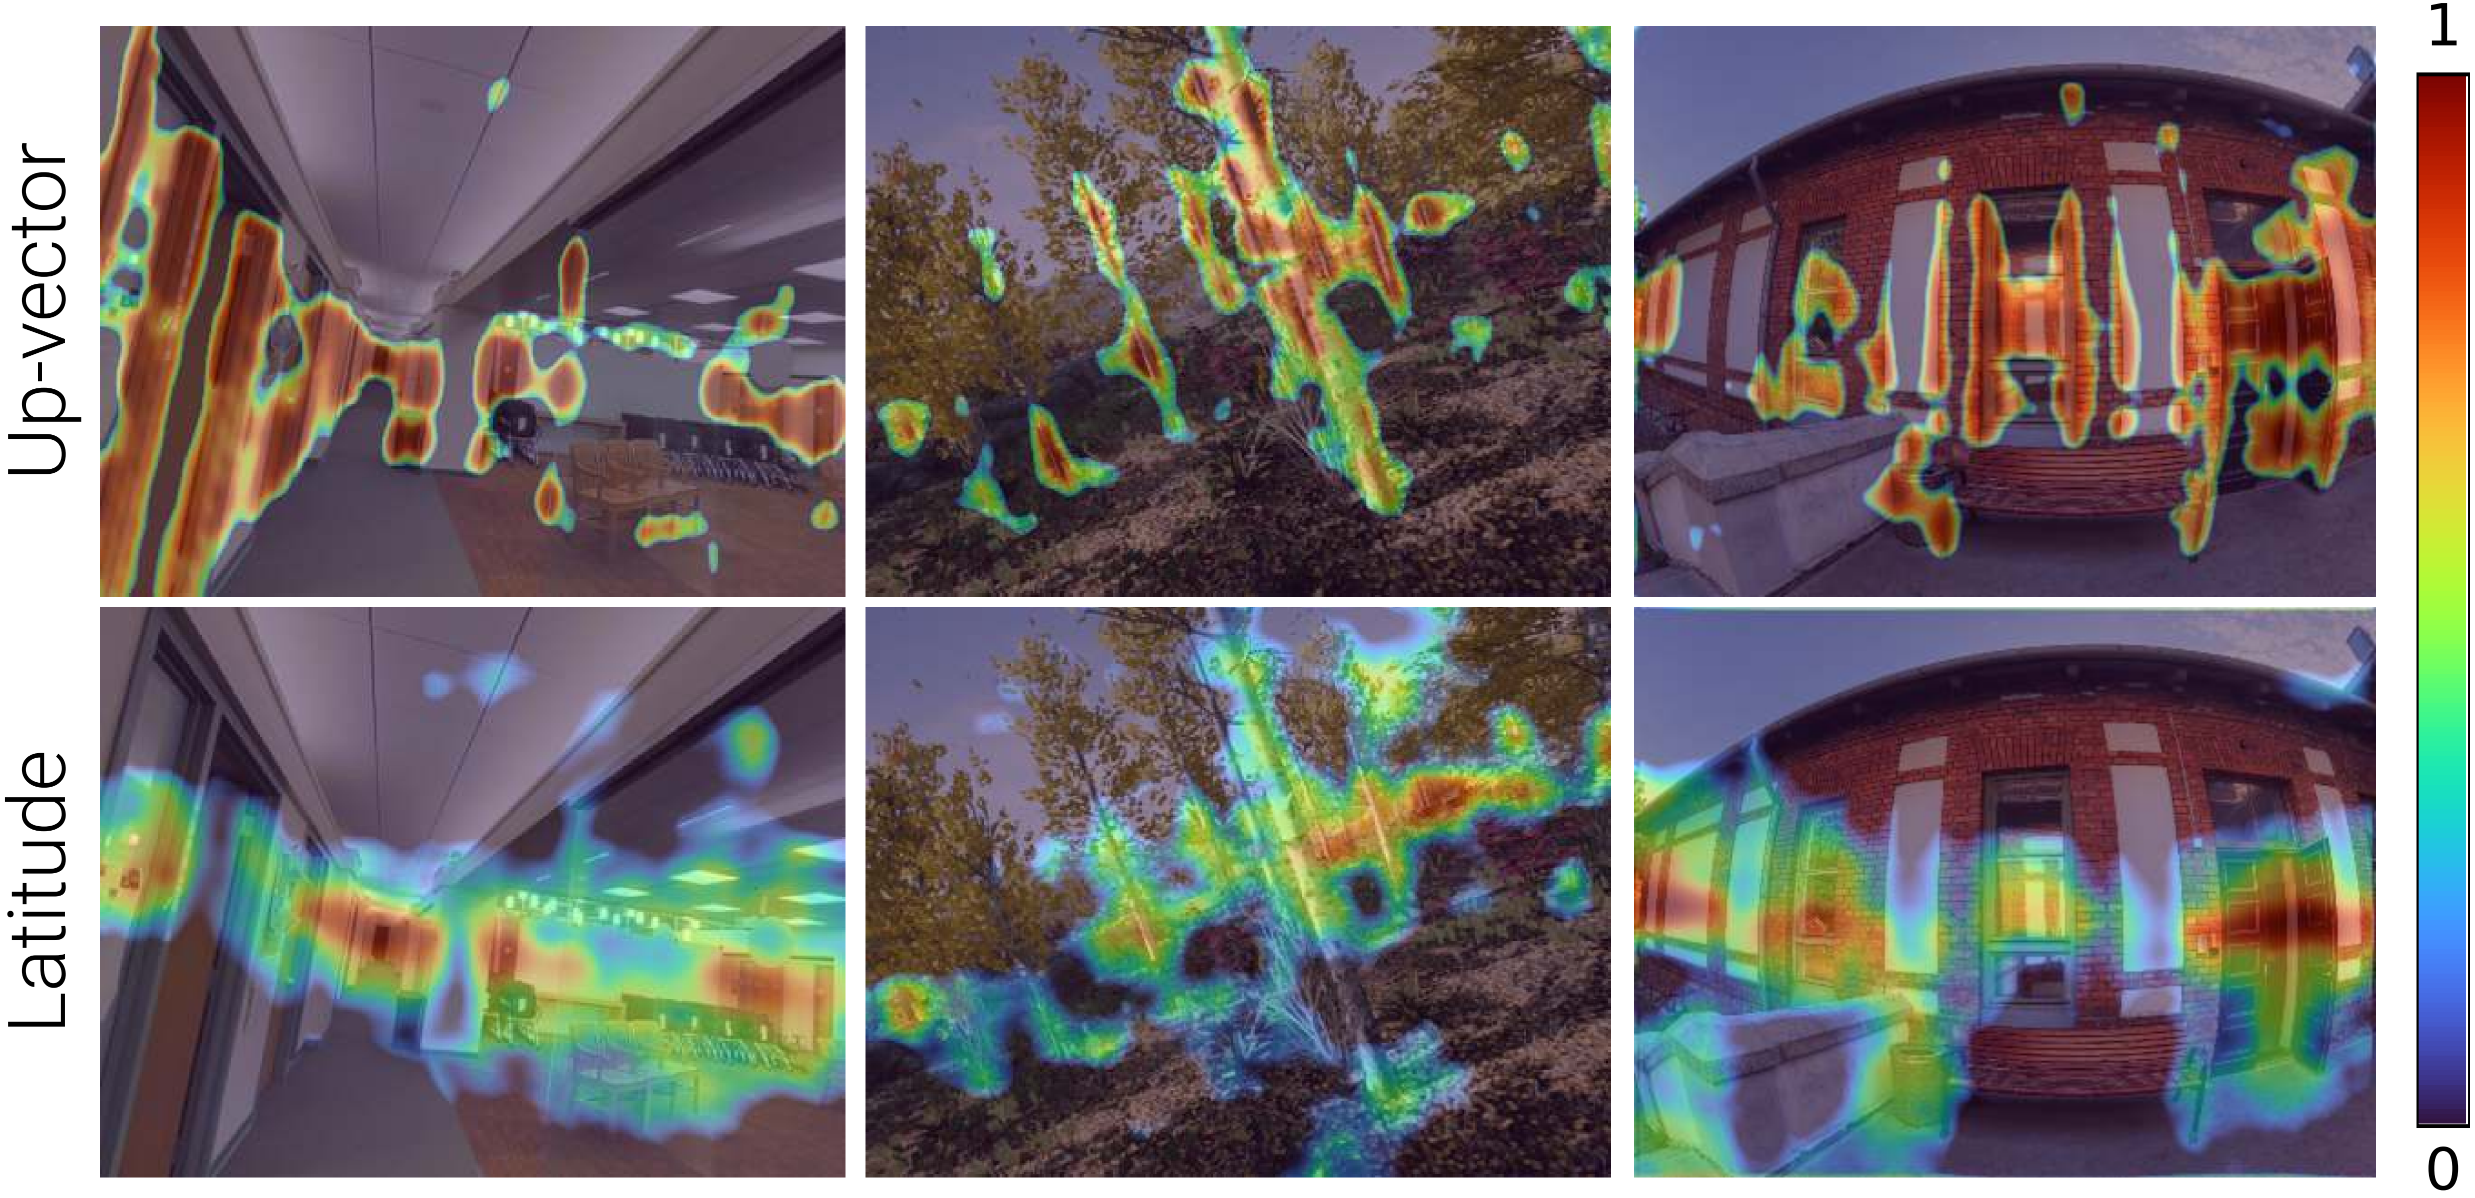

New dataset: OpenPano

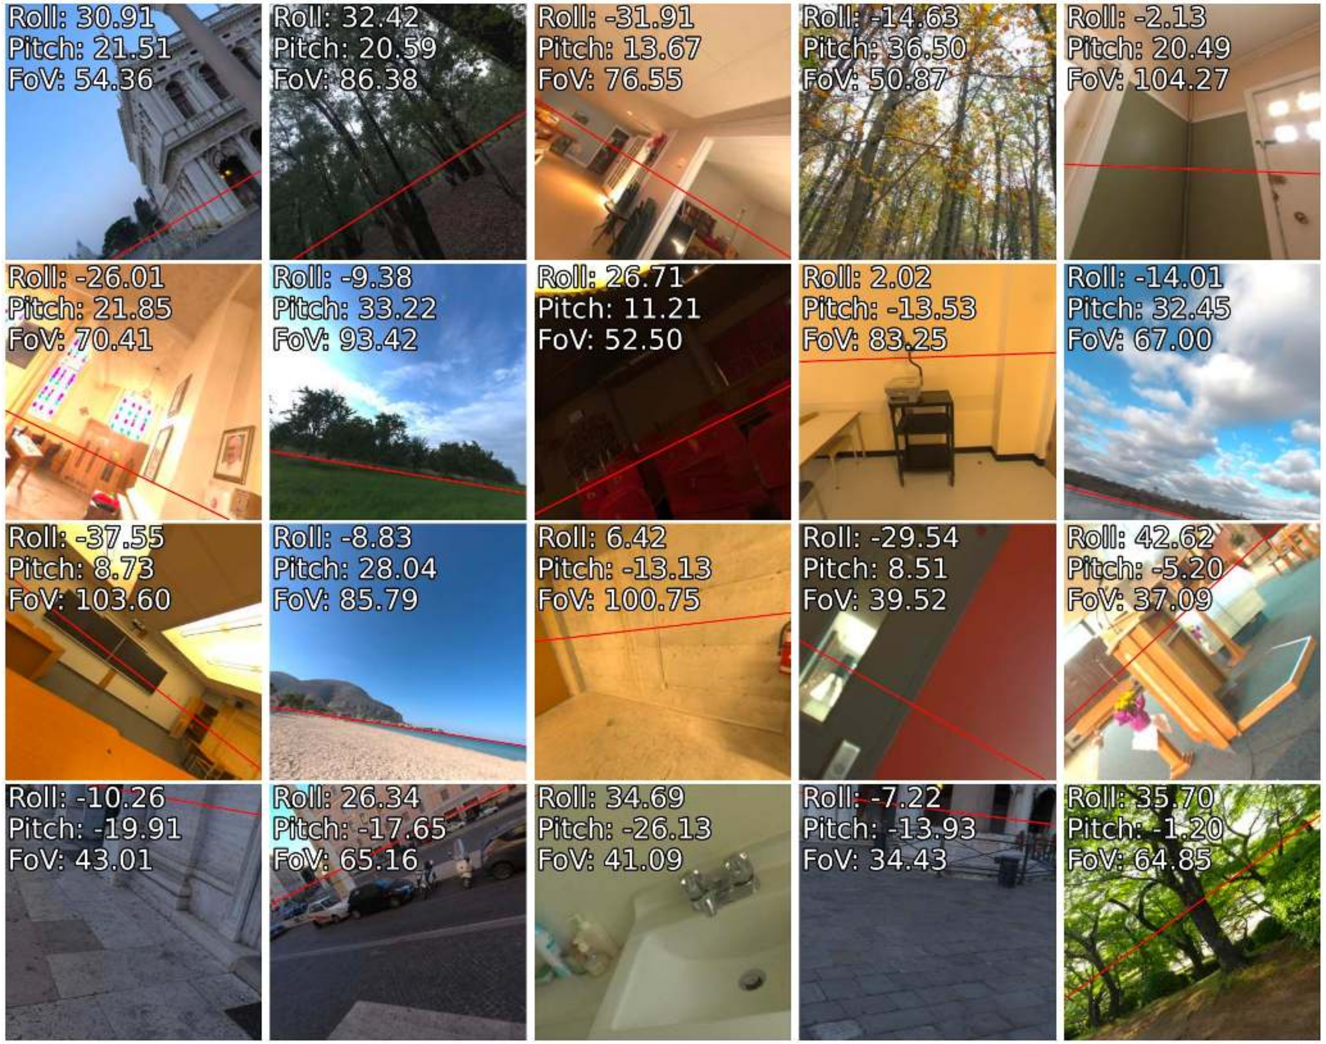

Previous research is based on proprietary datasets

- SUN360, 360Cities

OpenPano: public sources

→ GeoCalib is fully reproducible

Evaluation

LaMAR - phones

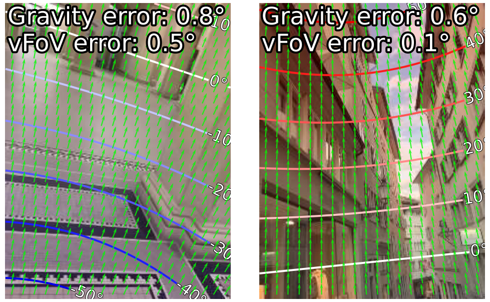

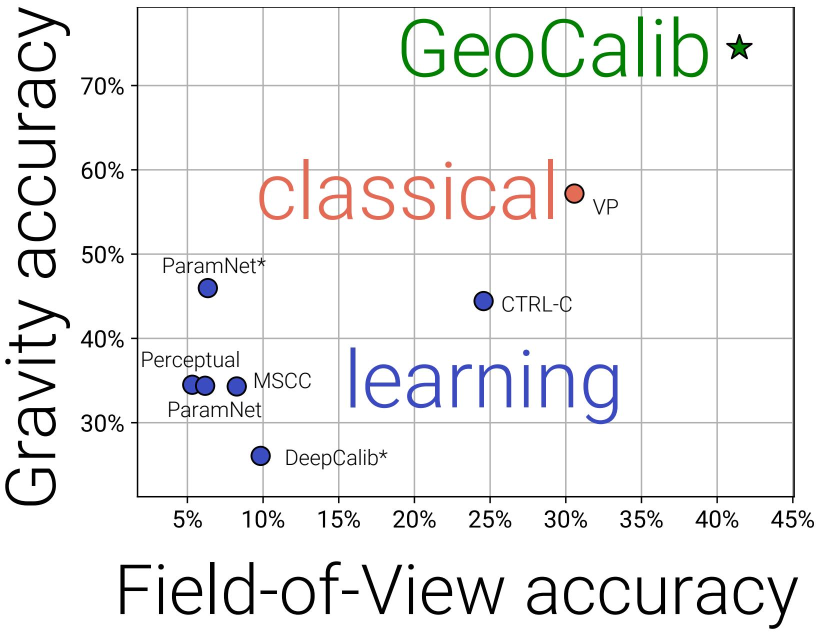

MegaDepth - outdoor

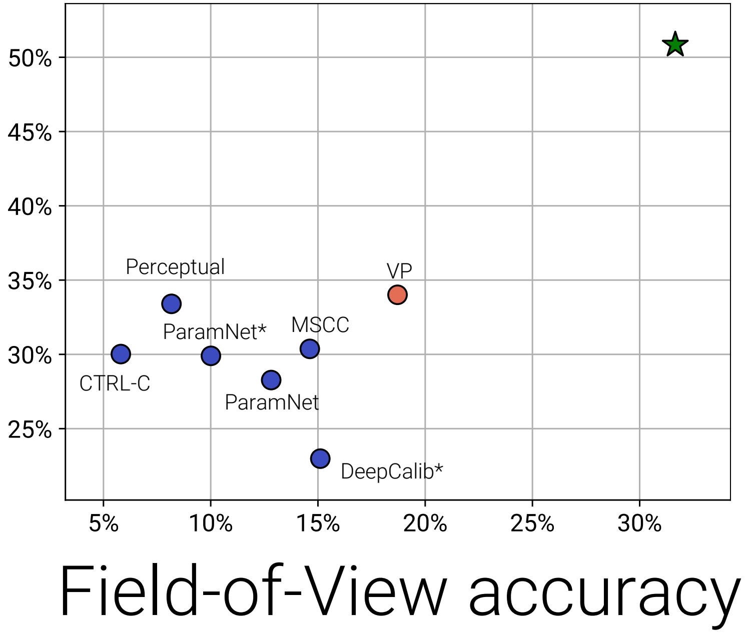

Stanford2D3D - indoor

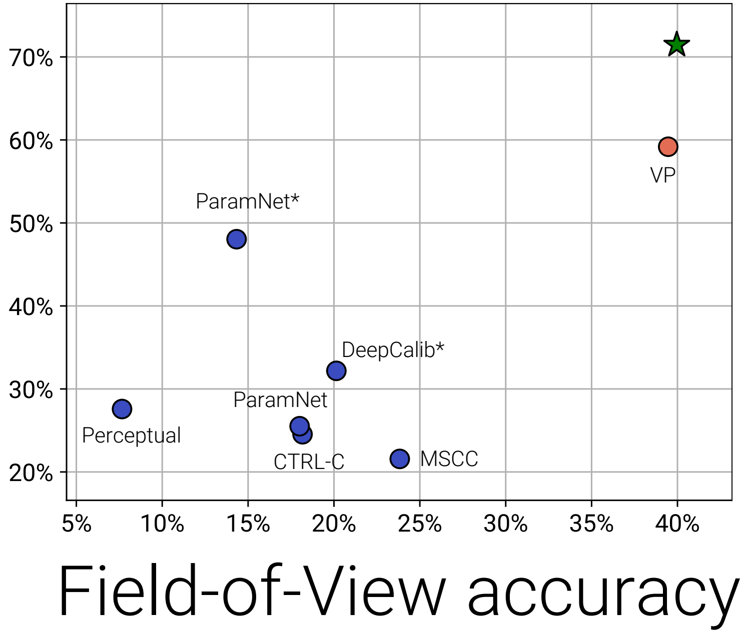

Practical benefits

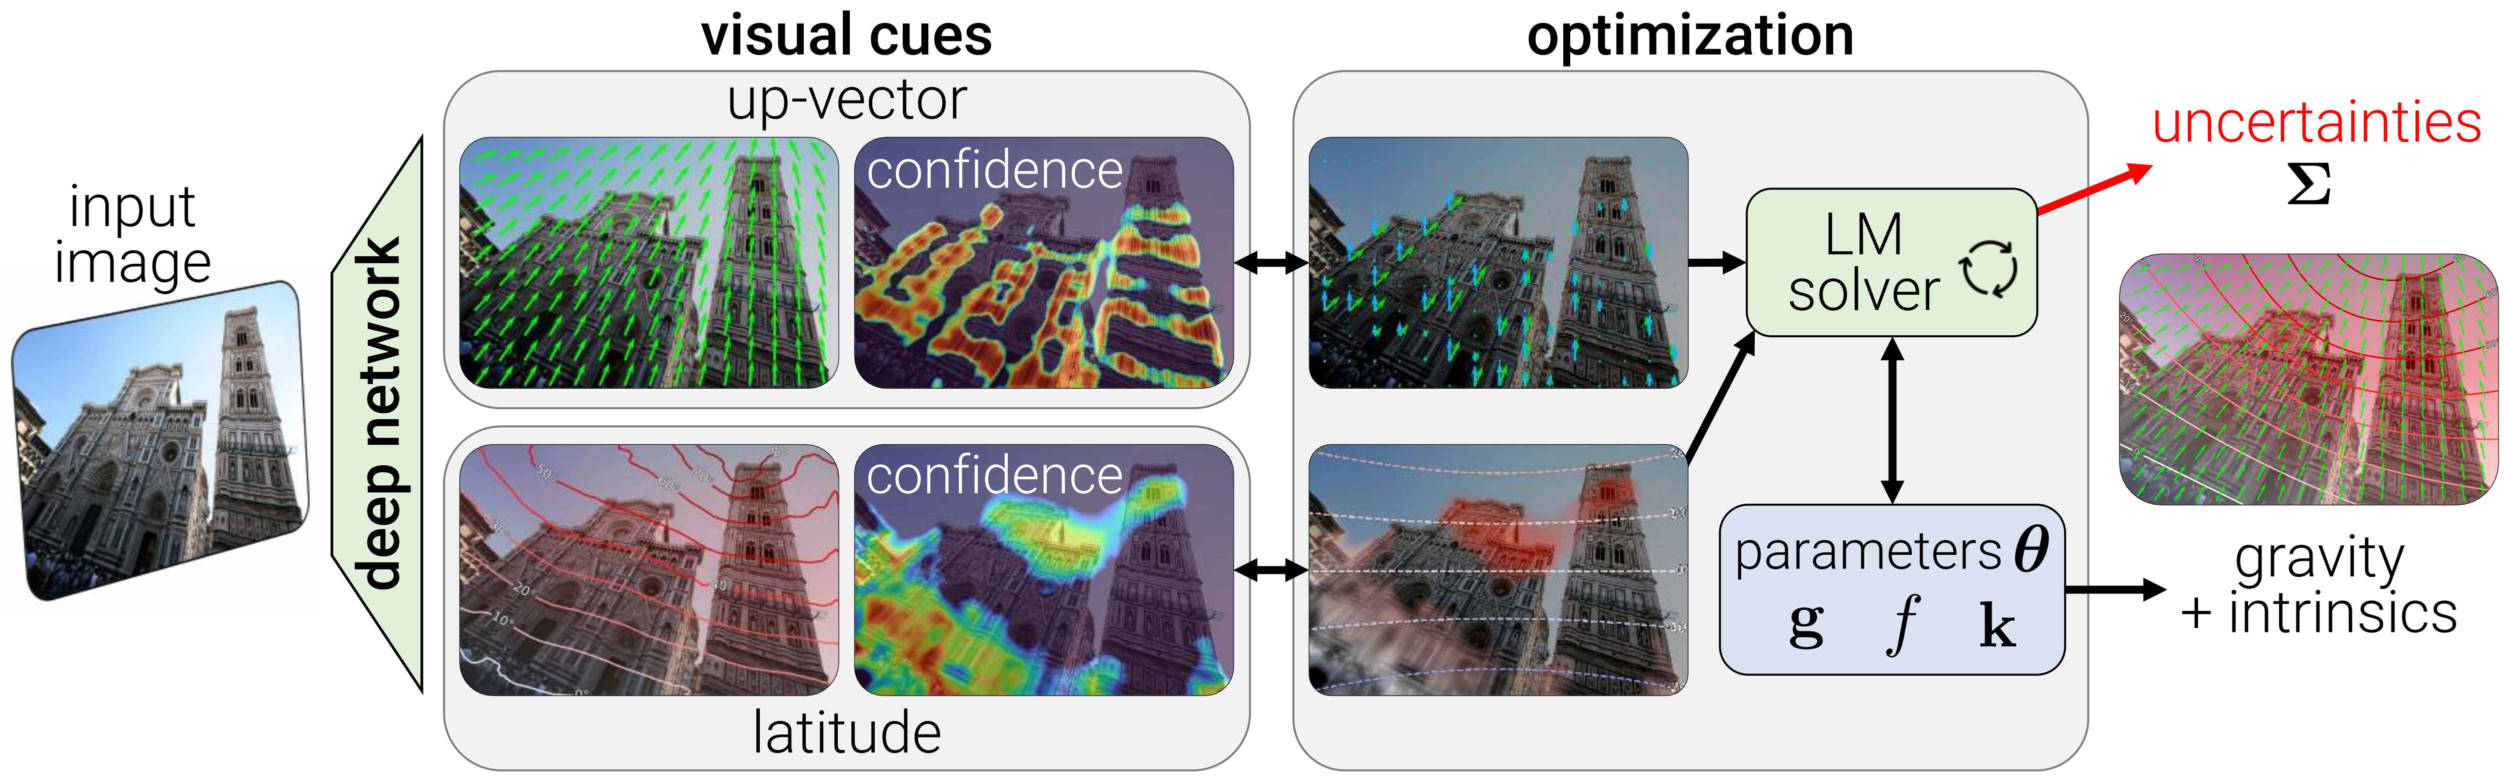

GeoCalib estimates reliable uncertainties

true gravity error *vs* predicted uncertainty

0.5°/1.2° 0.9°/1.4° 2.1°/2.2° 3.4°/3.2° 3.4°/4.7°

0.2°/0.7° 0.7°/1.0° 1.3°/1.5° 2.4°/2.6° 4.3°/4.8°

more difficult examples

Practical benefits

Any camera model
pinhole, distorted, etc.

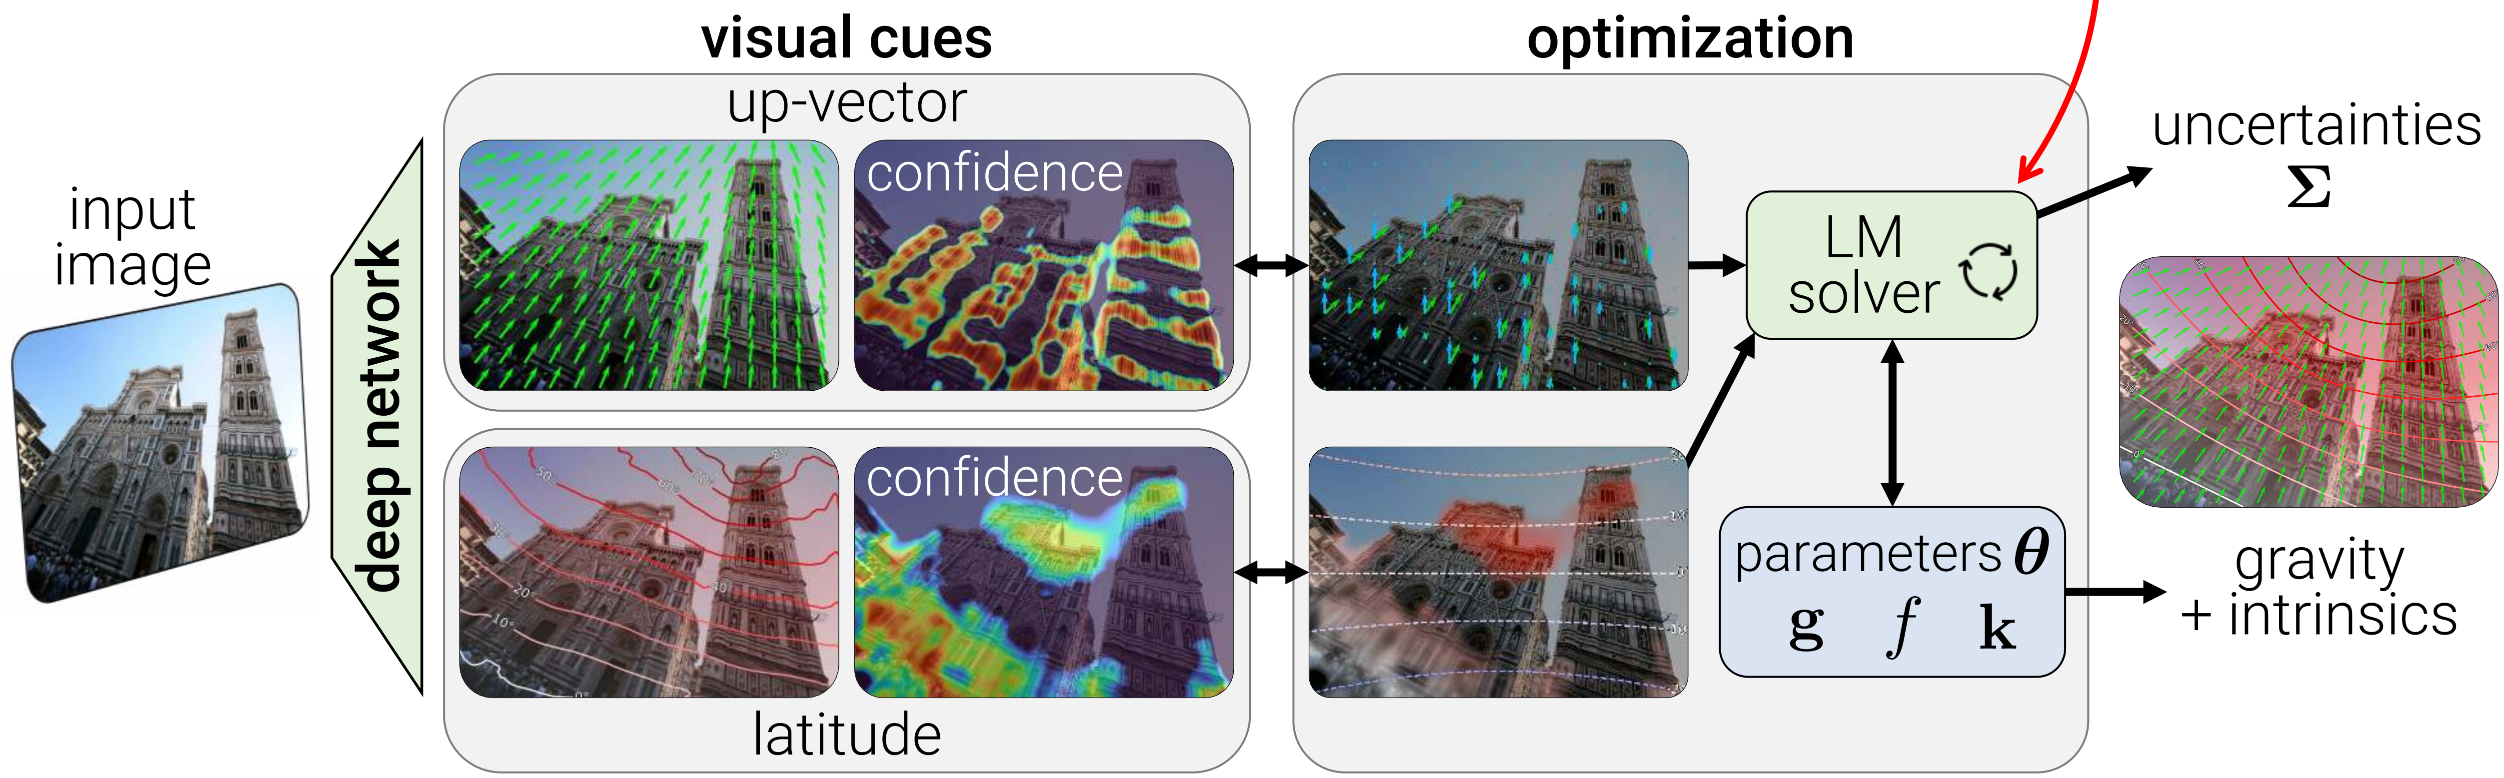

GeoCalib can handle lens distortion

input undistorted

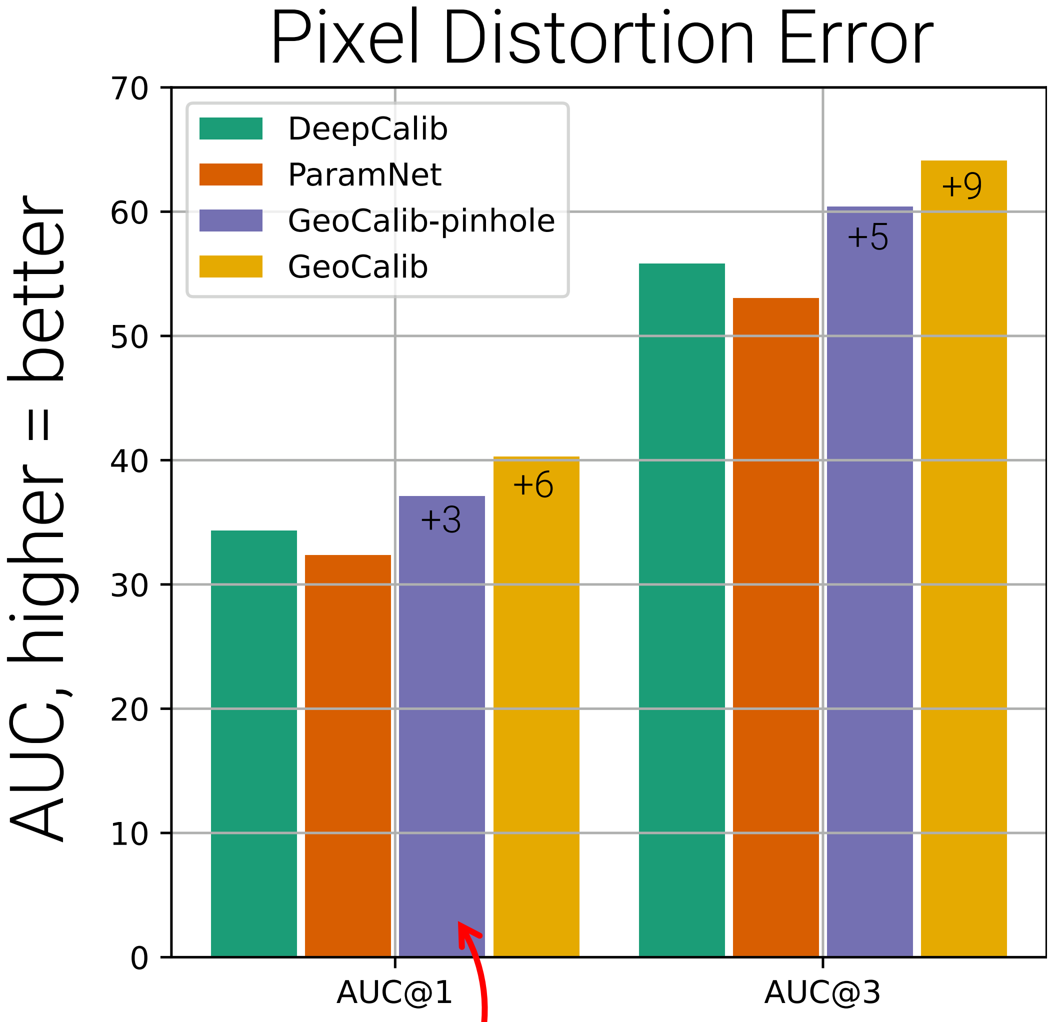

trained only on pinhole images

Practical benefits

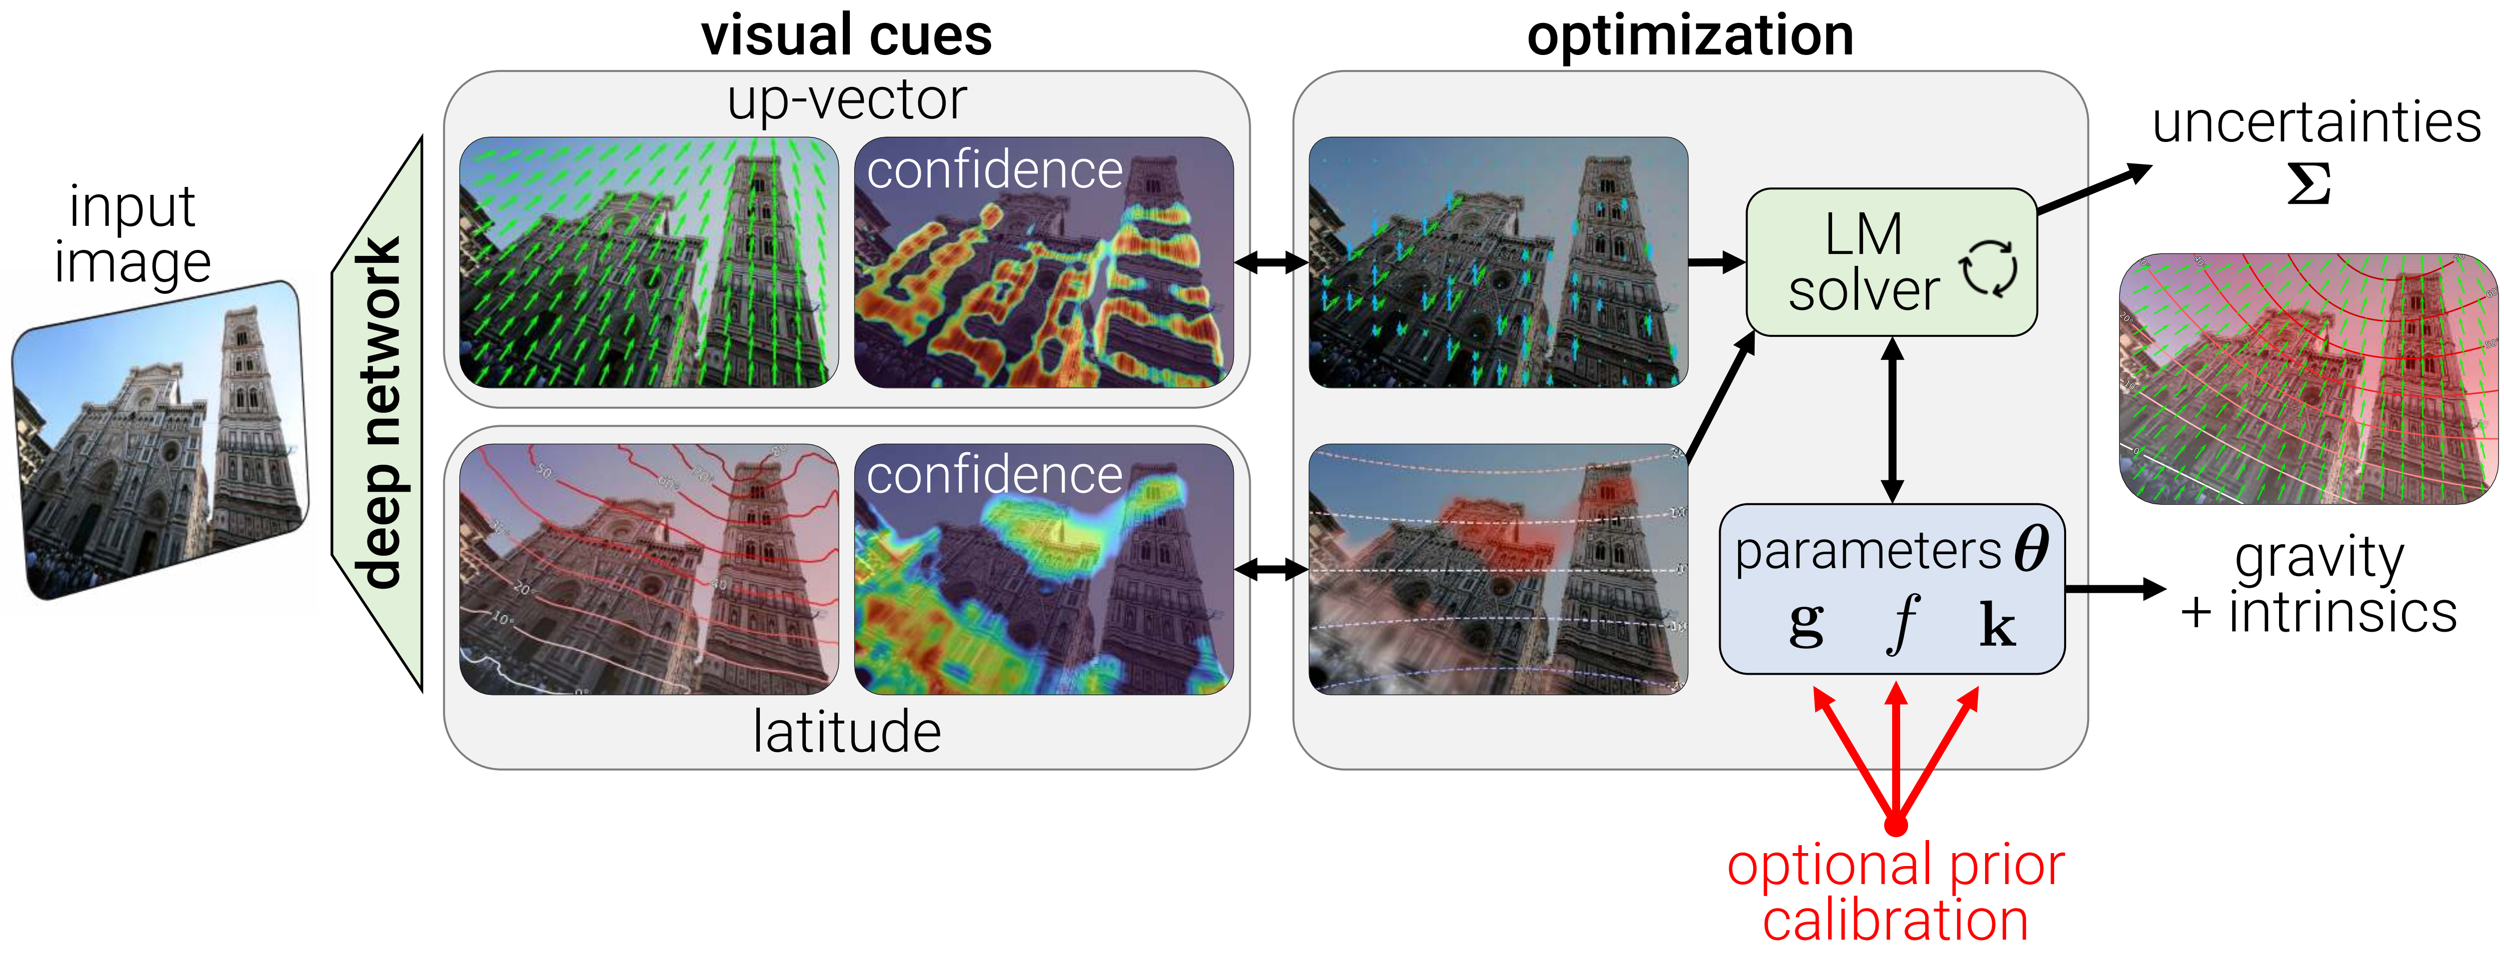

Leveraging partial priors

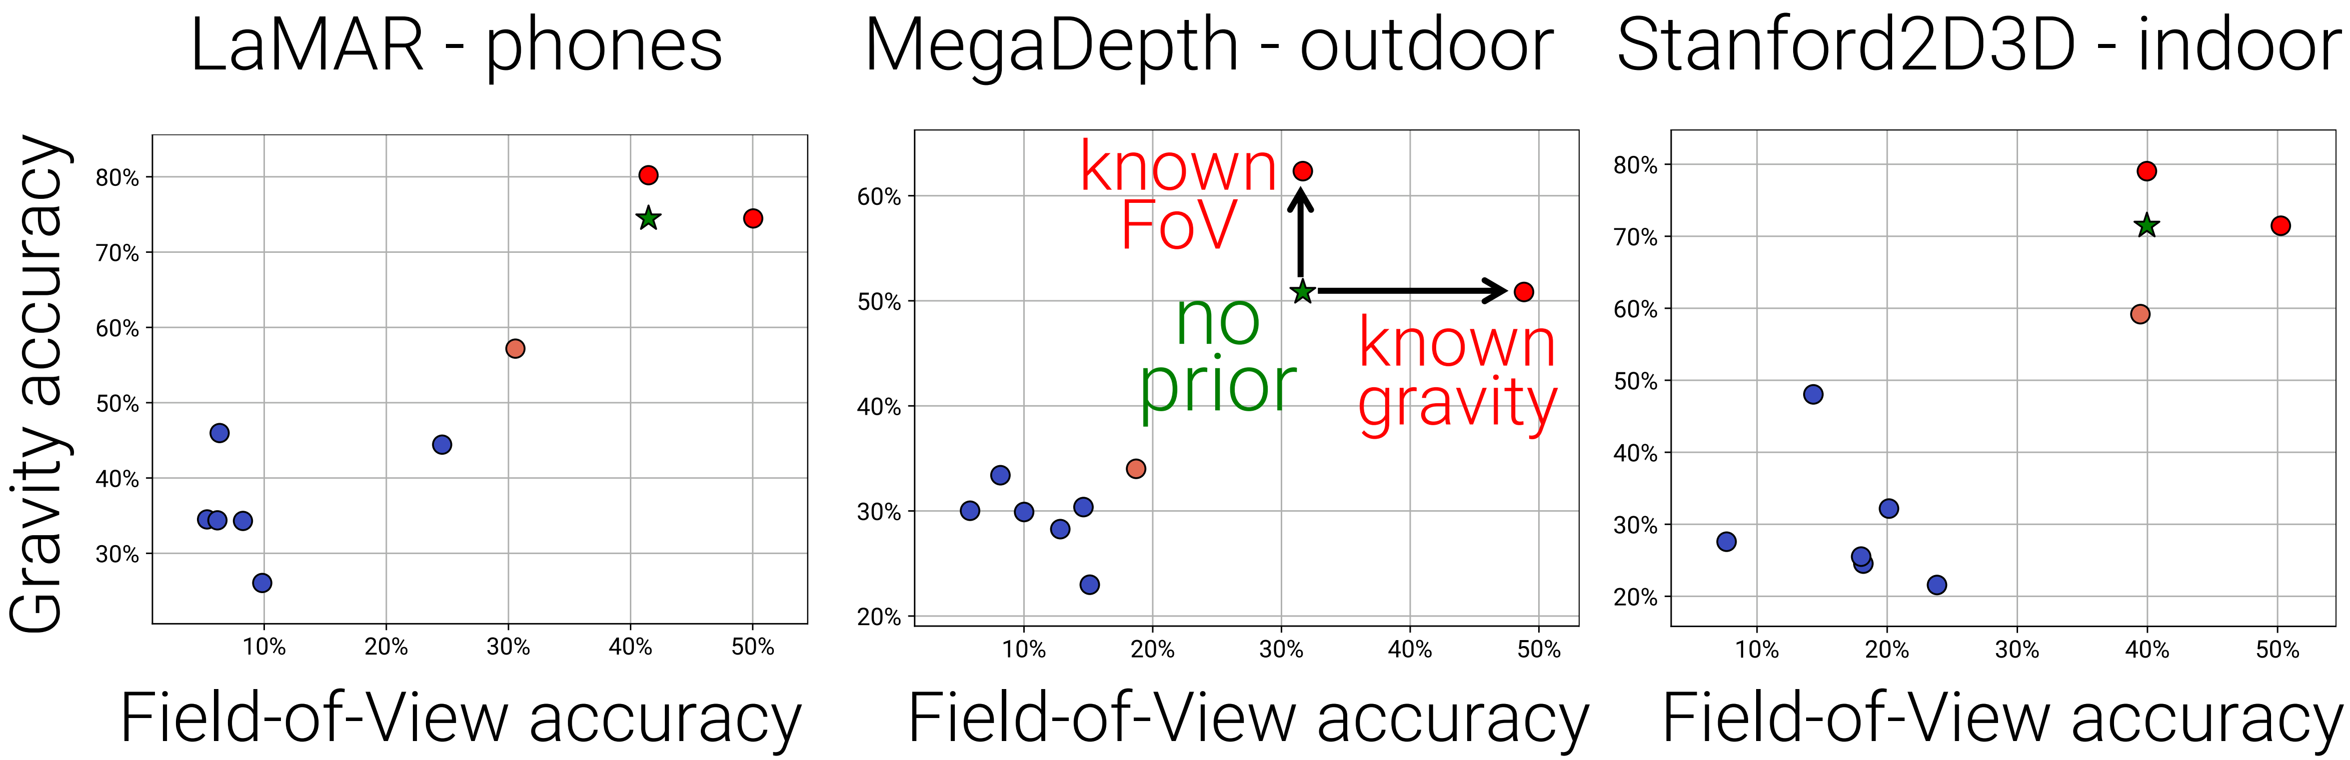

Multi-image optimization

same camera
but different gravity directions

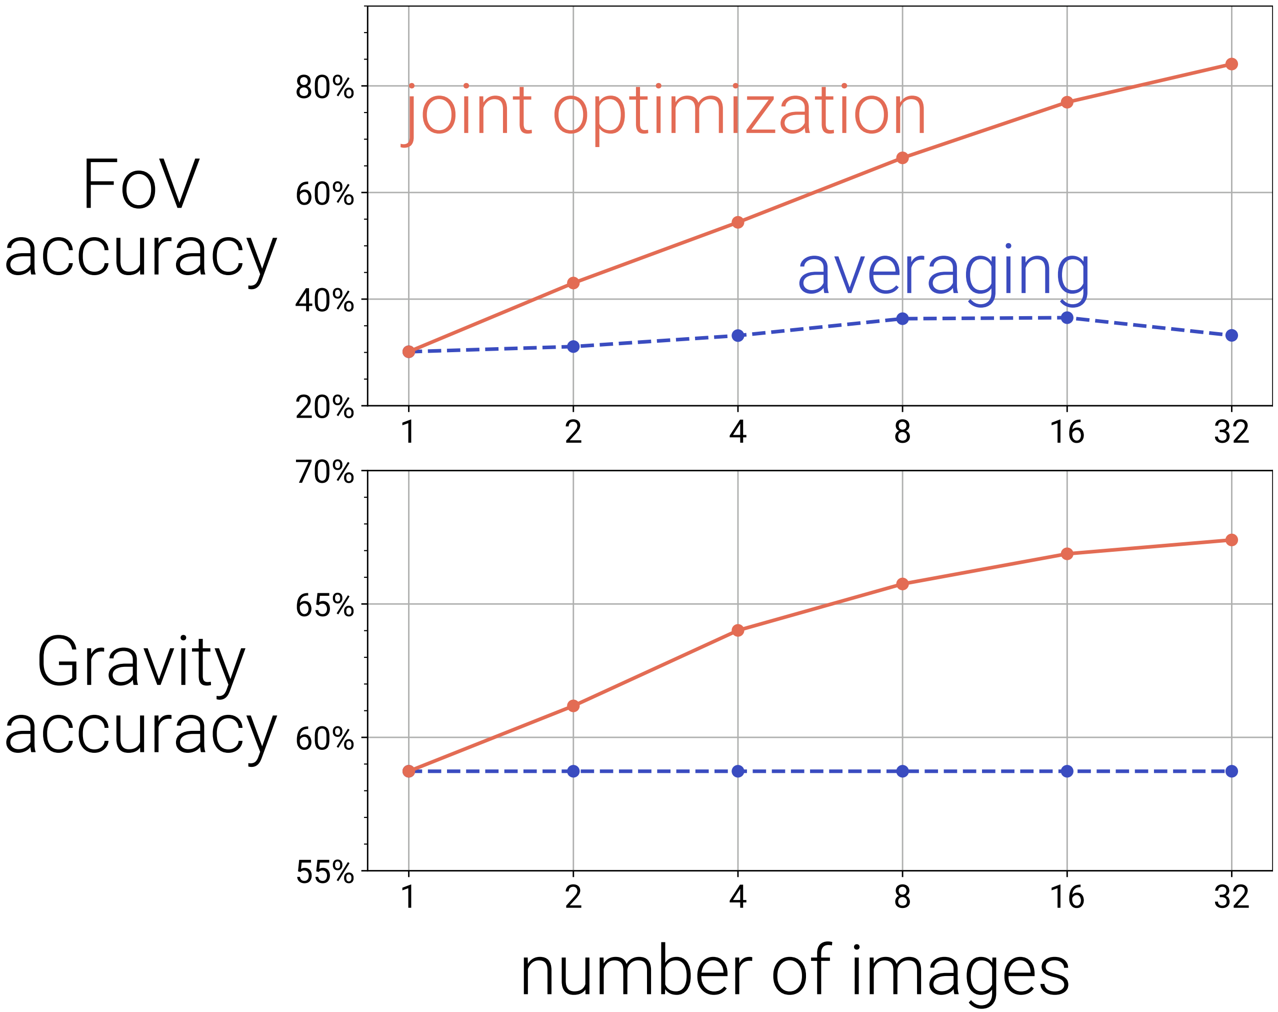

GeoCalib = Learning + Geometry

More **accurate & robust** single-view calibration

Multiple benefits of **geometric optimization**:

- Uncertainties = interpretability
- Flexible camera models
- Optional partial priors
- Multi-image optimization

github.com/cvg/GeoCalib