Single Image 3D Reconstruction of Ball Motion and Spin From Motion Blur

An Experiment in Motion from Blur

Giacomo Boracchi, *Vincenzo Caglioti*, Alessandro Giusti

Objective

From a single image, reconstruct:

- Ball 3D position
- 3D velocity
- spin axis
- angular speed

We analyze motion blur affecting the ball image exploiting known scene geometry.

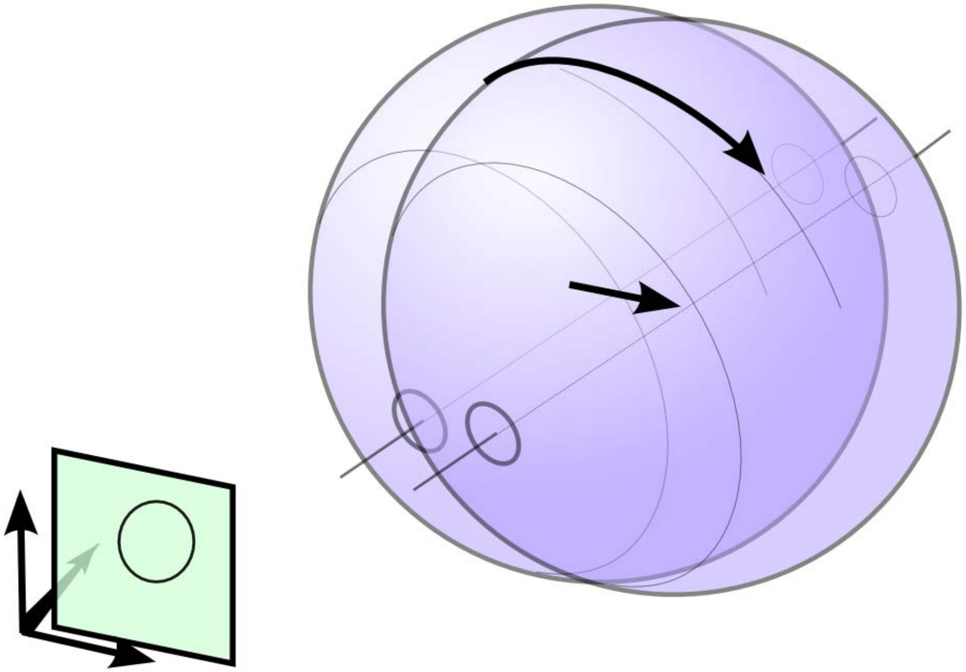

Motivations

- Multi-frame approaches often tricky due to:
  - motion blur/noise tradeoff (exp. indoors)
  - too few/too many similar features on ball surface
  - frame rate too low for spin reconstruction
  - resolution too low for 3D localization
- Our approach:
  - high resolution (one 8+ MPixel DSLR photo, triggered shutter).
  - works with any texture on ball surface
  - exploits intra-frame information

2MP
30fps

>8MP
3fps

Requirements

- Known exposure time
- Ball appears (slightly) blurred
- Ball surface color not homogeneous

Assumptions

- Exact perspective
- Spherical ball
- During exposure time:
  - Ball trajectory locally straight
  - Ball apparent translation less than 1/5 of apparent radius
  - Spin angle 1 to 10 degrees
- Usually met in practice if appropriate exposure time is used

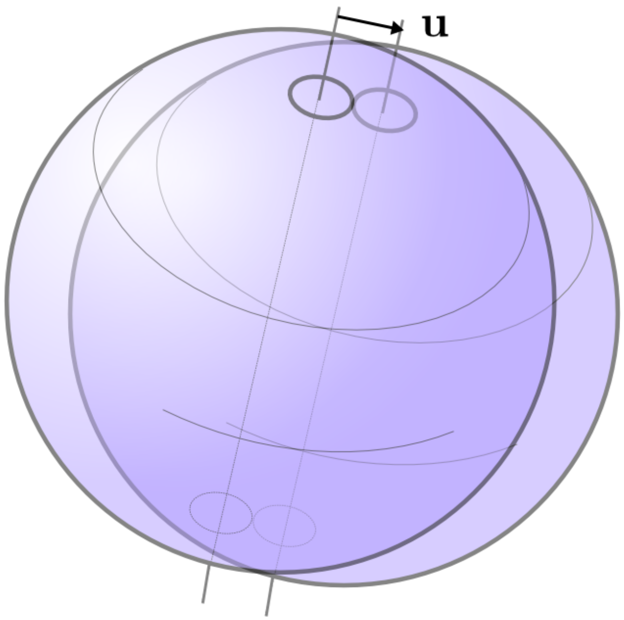

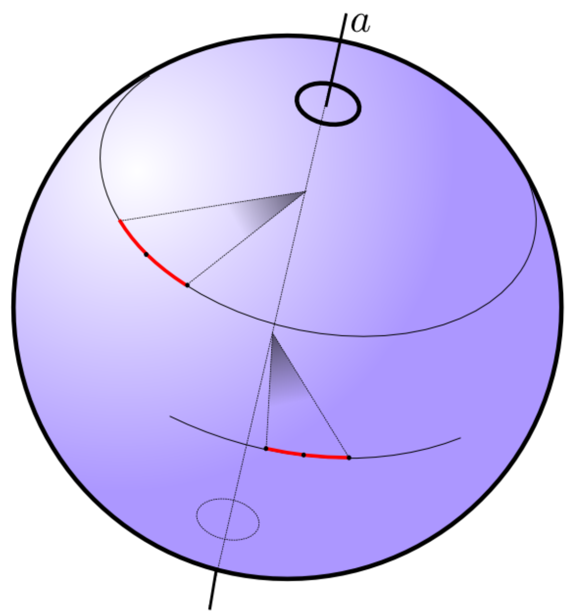

Related (1)

Shan, Qi   Xiong, Wei   Jia, Jiaya: *Rotational Motion Deblurring of a Rigid Object from a Single Image*

- Concerned with deblurring
- Considers a simpler, 2D object geometry
- Rotational-only motion
- Uses alpha map as transparency cues

Related (2)

Ball trajectory reconstruction from a single long-exposure image (CVBASE 2006)

- Exploits single high resolution image as well
- Longer exposure
- No spin recovery

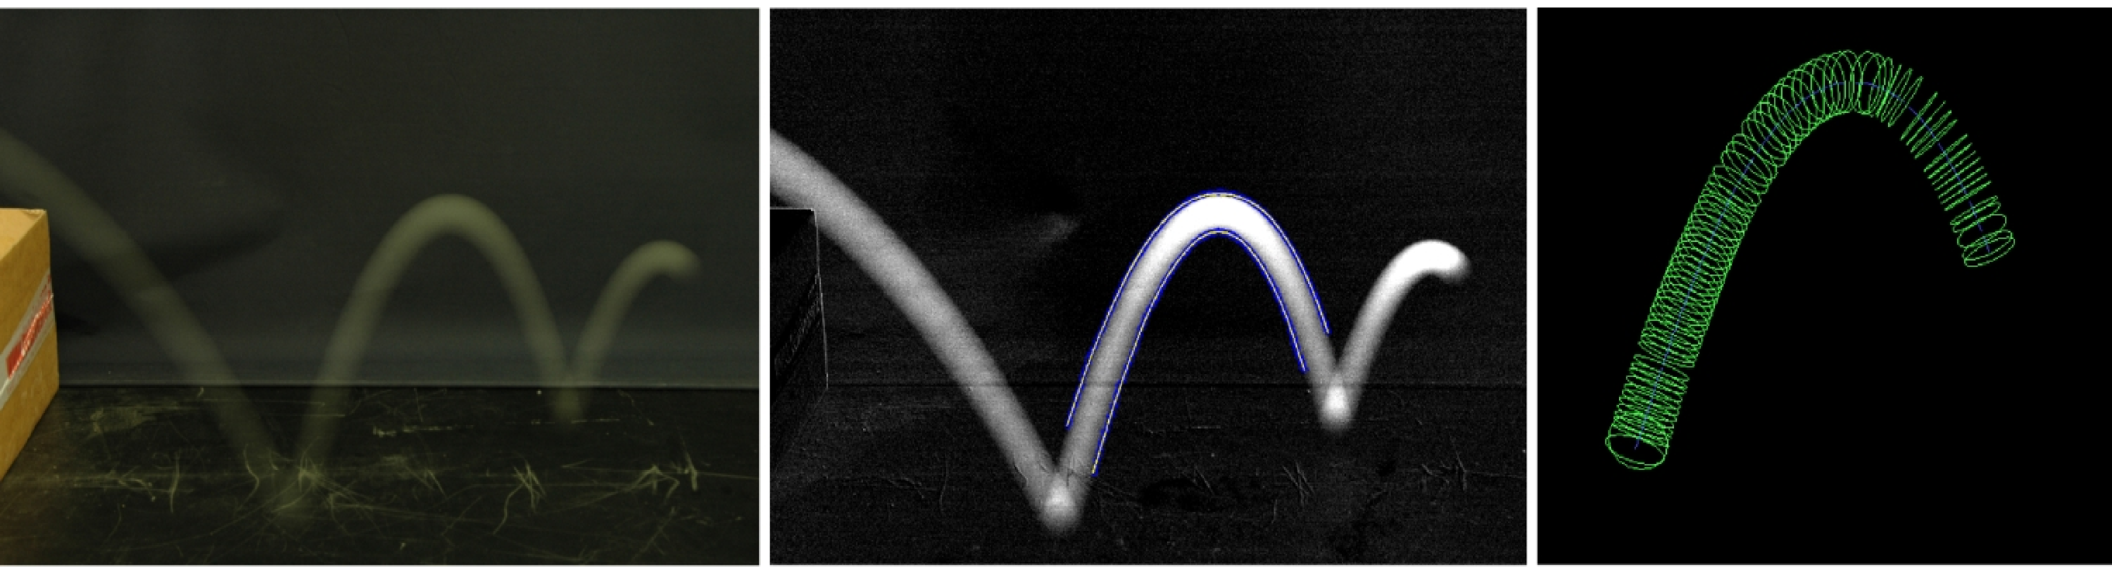

Related (3)

3D position and 3D velocity estimation (ICIAP 2007)

- Similar requirements
- Restrictive requirements on ball and background appearance (uniform color)
- No spin recovery

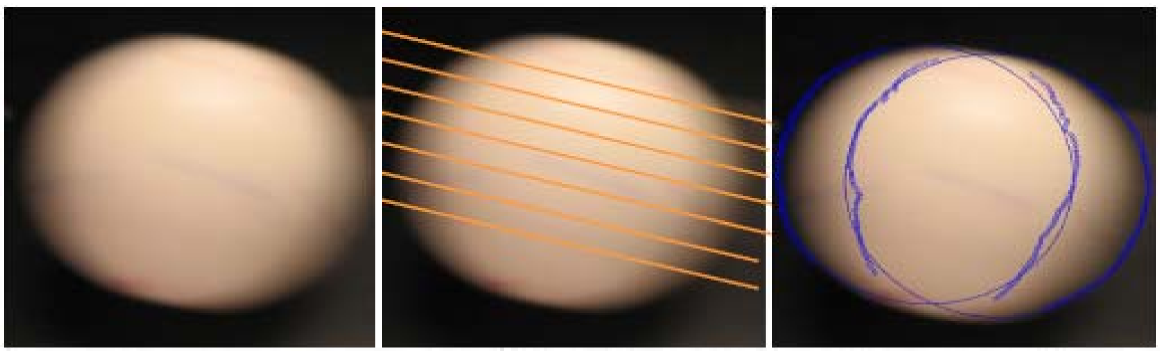

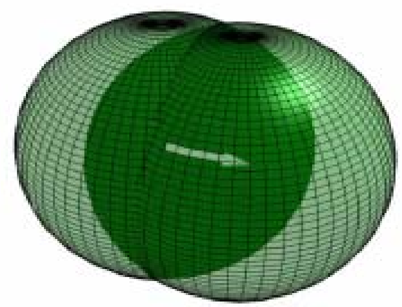

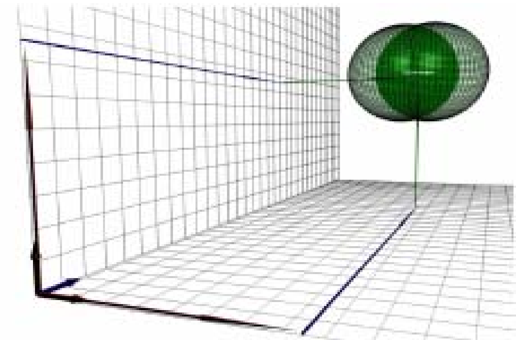

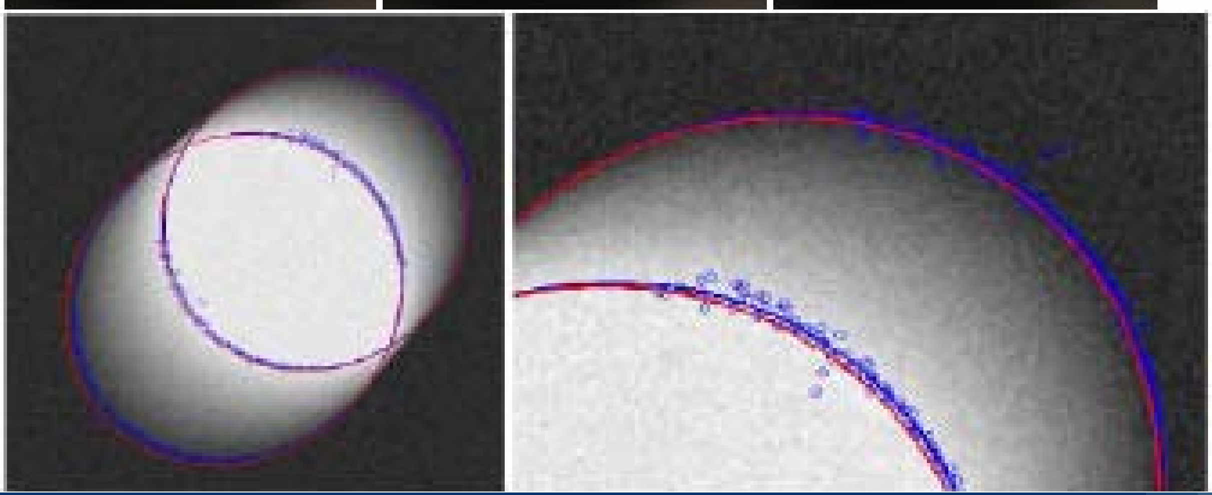

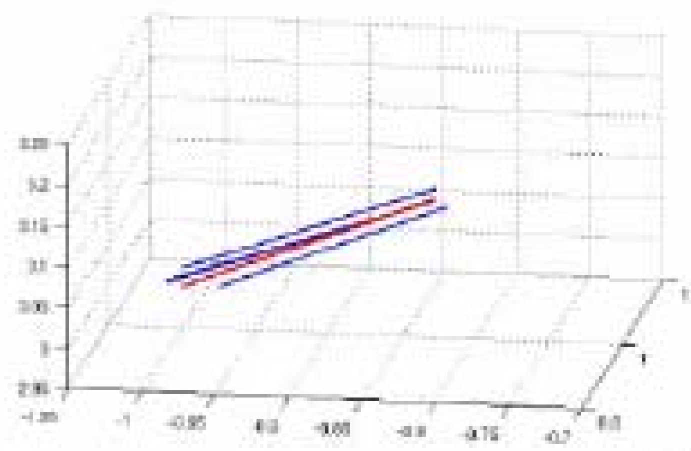

Motion blur model

The blurred image is obtained as the temporal average of many still images over the exposure time.

- Both translation and spin affect the ball appearance;
- Only translation affects the ball's apparent transparency (alpha)

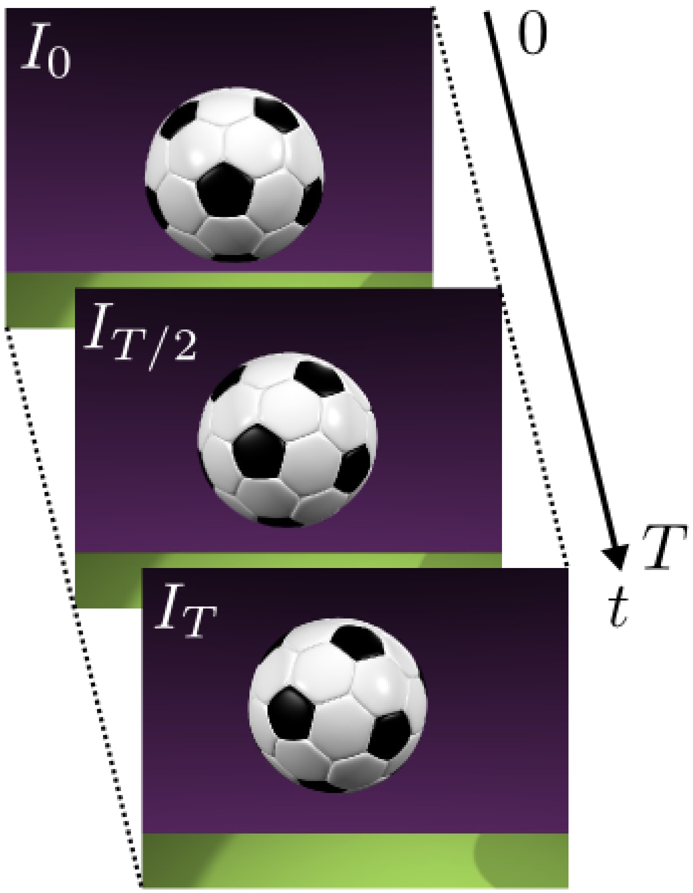

Alpha matting allows translation and spin separation

Alpha matting algorithms provide:

- An alpha map, representing the ball translation. Alpha at a pixel *p* represent the fraction of the exposure time during which the ball image overlapped *p* [Caglioti, Giusti BMVC2007]
- A foreground map, showing the ball surface without background

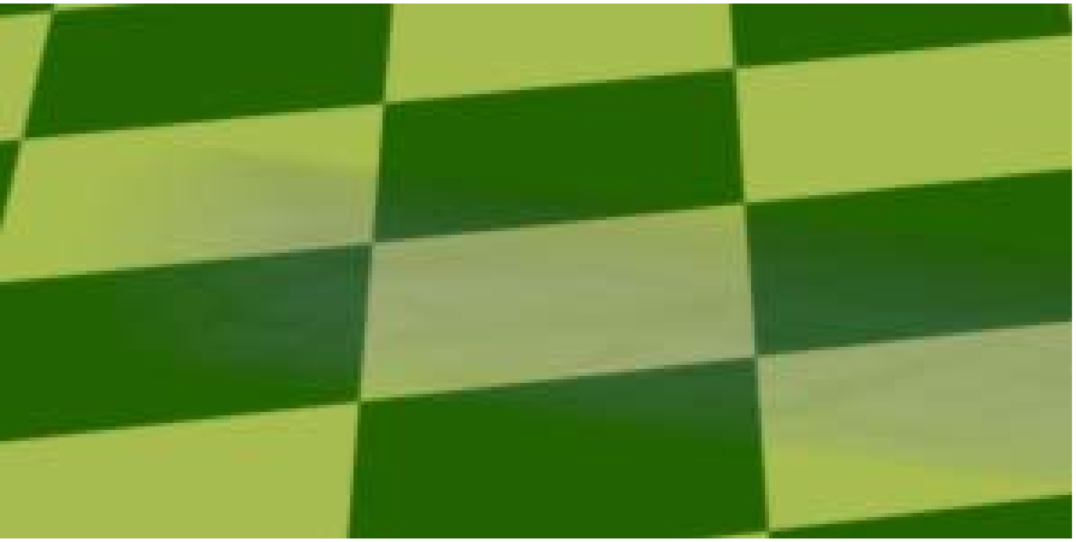

Local estimation of blur directions

- We locally estimate the direction of blur on the ball surface (from foreground map)
- Minimum derivatives algorithm

→ Reliable direction estimates

→ Unreliable extent estimates due to nonrectilinear blurring paths

- **NO INFORMATION ON ORIENTATION!**

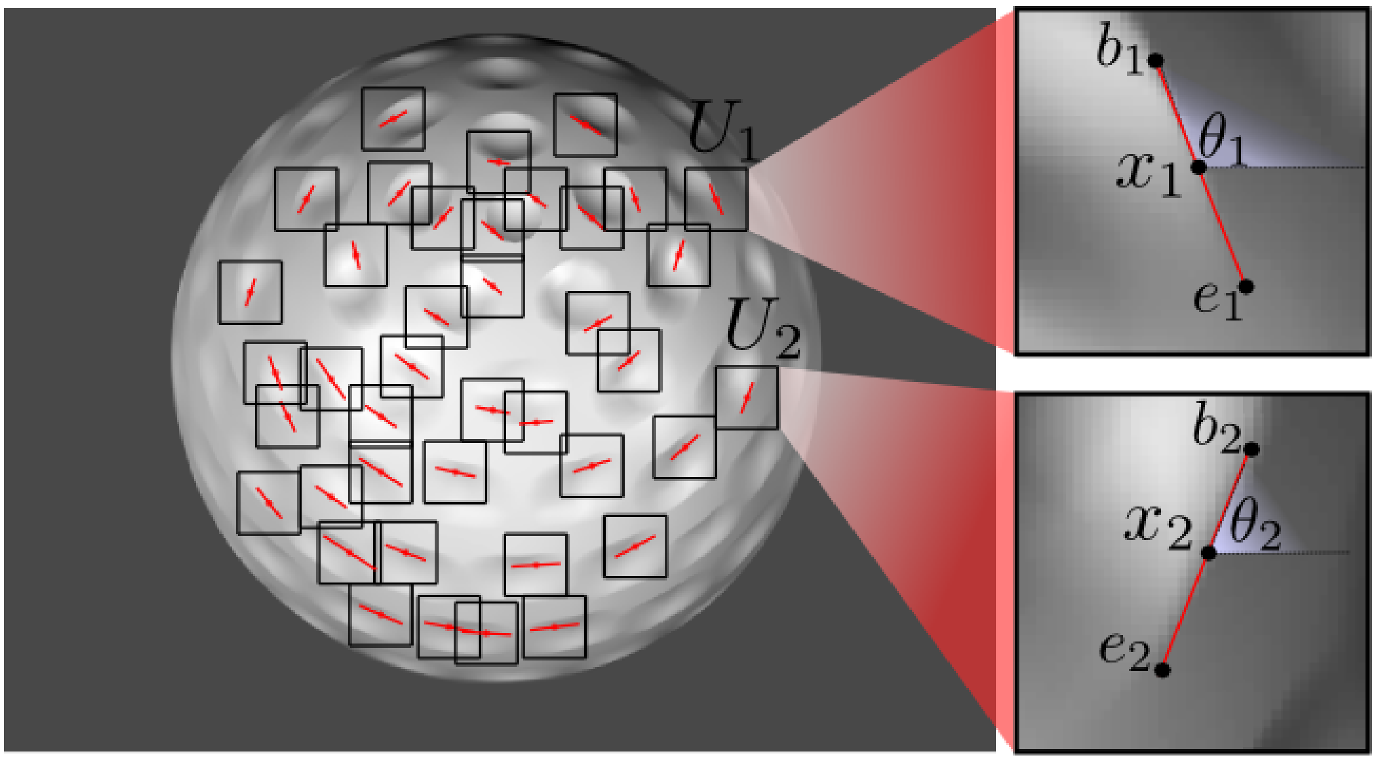

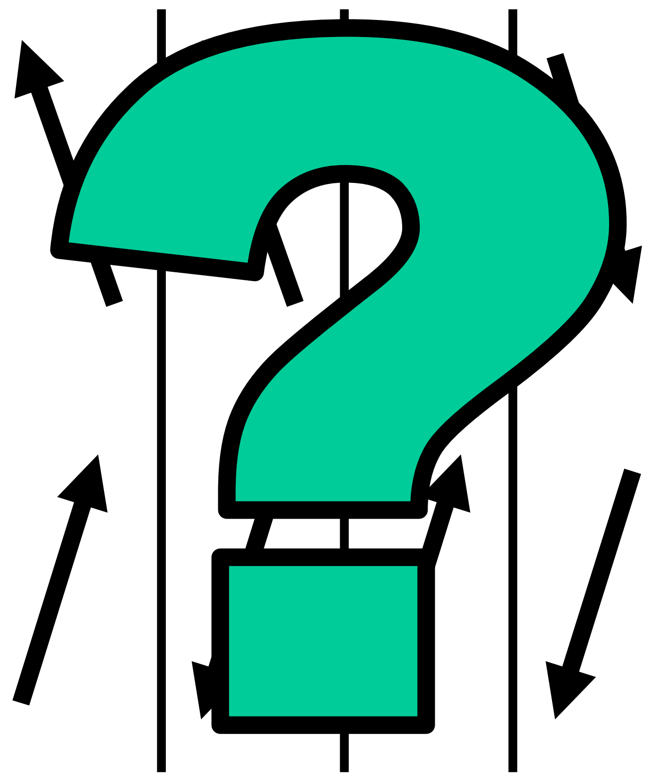

Simplified case: no apparent translation

Sharp ball contours

→ ball position from ellipse fitting and calibrated camera

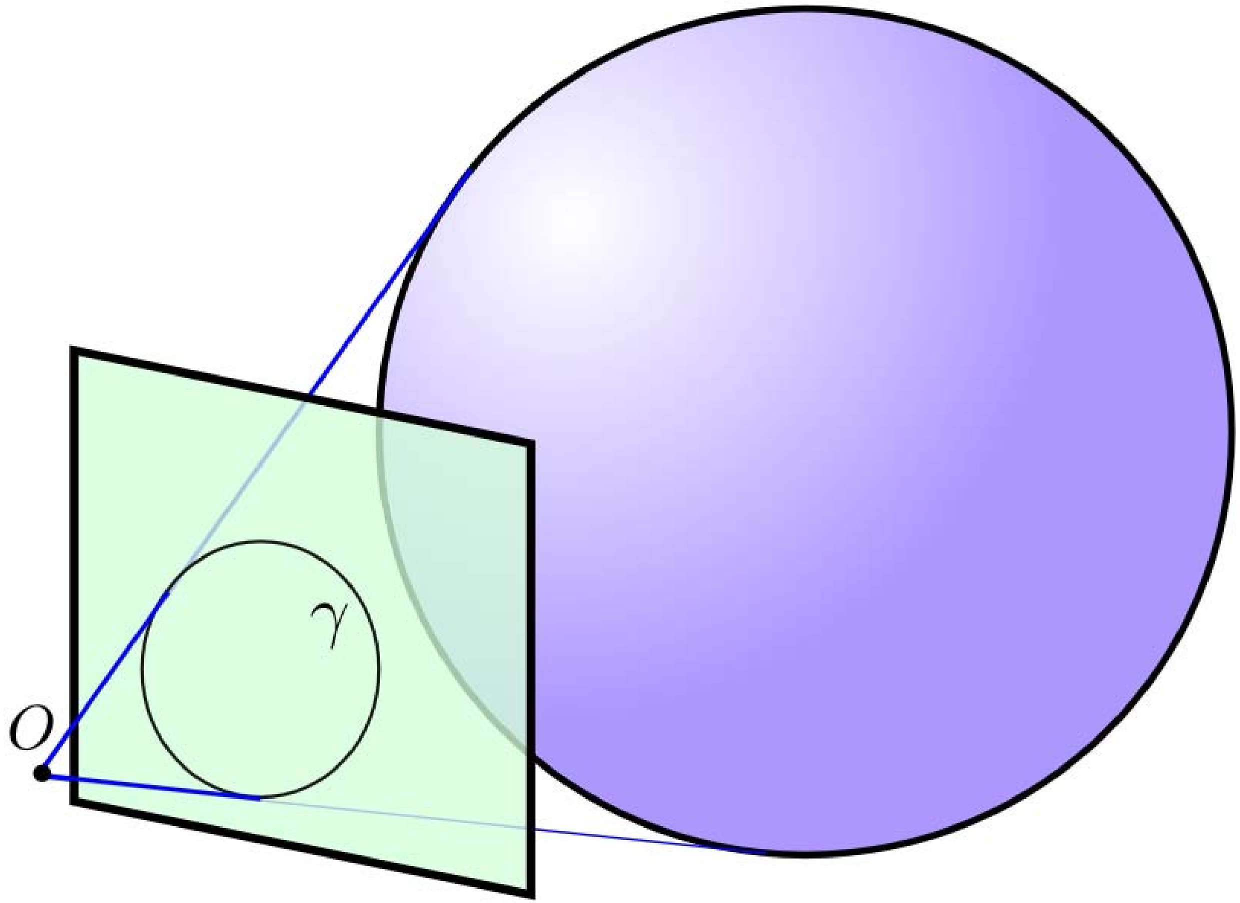

Simplified case (2) -- Spin axis recovery

- Backproject blur directions on 3D surface
- find 3D directions of ball surface displacement
- Two such directions identify spin axis (perpendicular to both)

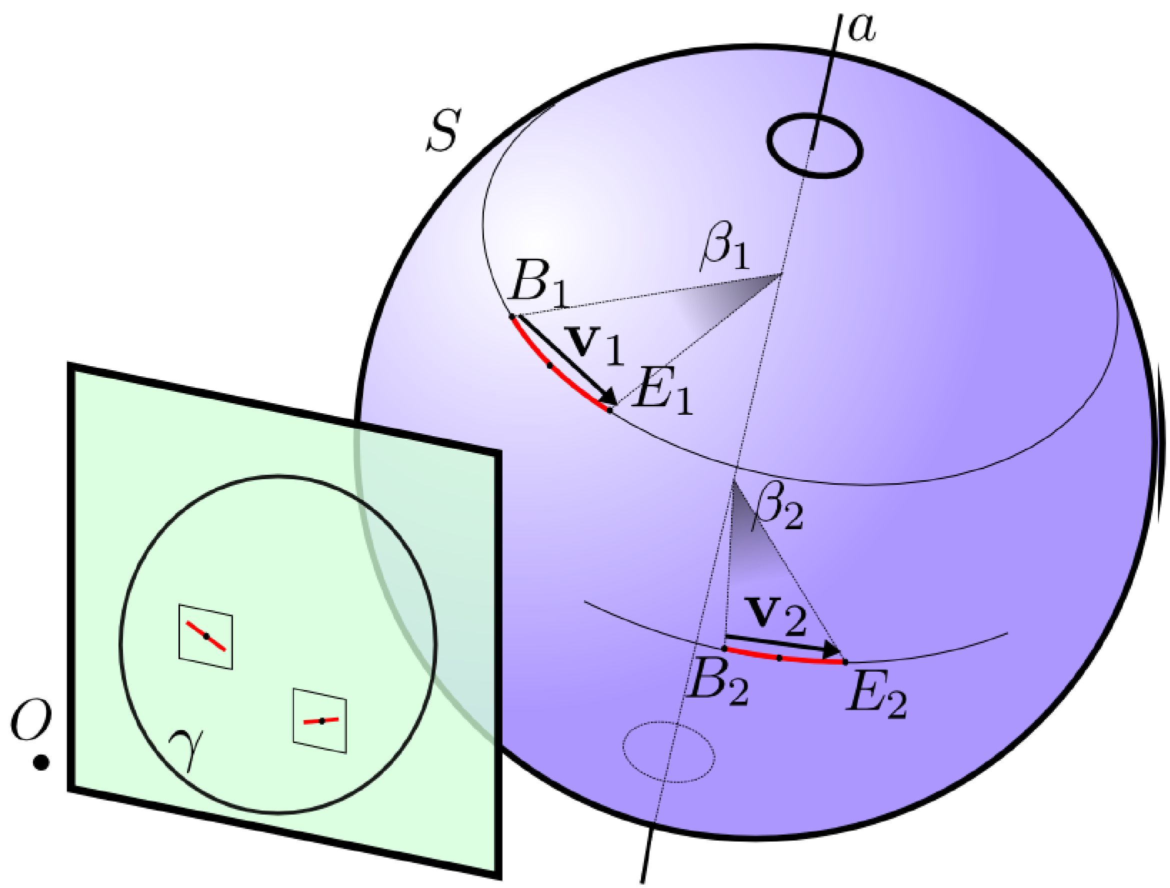

Simplified case (3) -- Robustness in spin axis recovery

- We have many local blur directions
- Corresponding 3D directions should be coplanar
- Robust fitting of a plane with iterative outlier removal. Iteratively repeat:
  - Least squares fitting of plane to 3D directions using SVD
  - Outlier identification and removal using fixed threshold

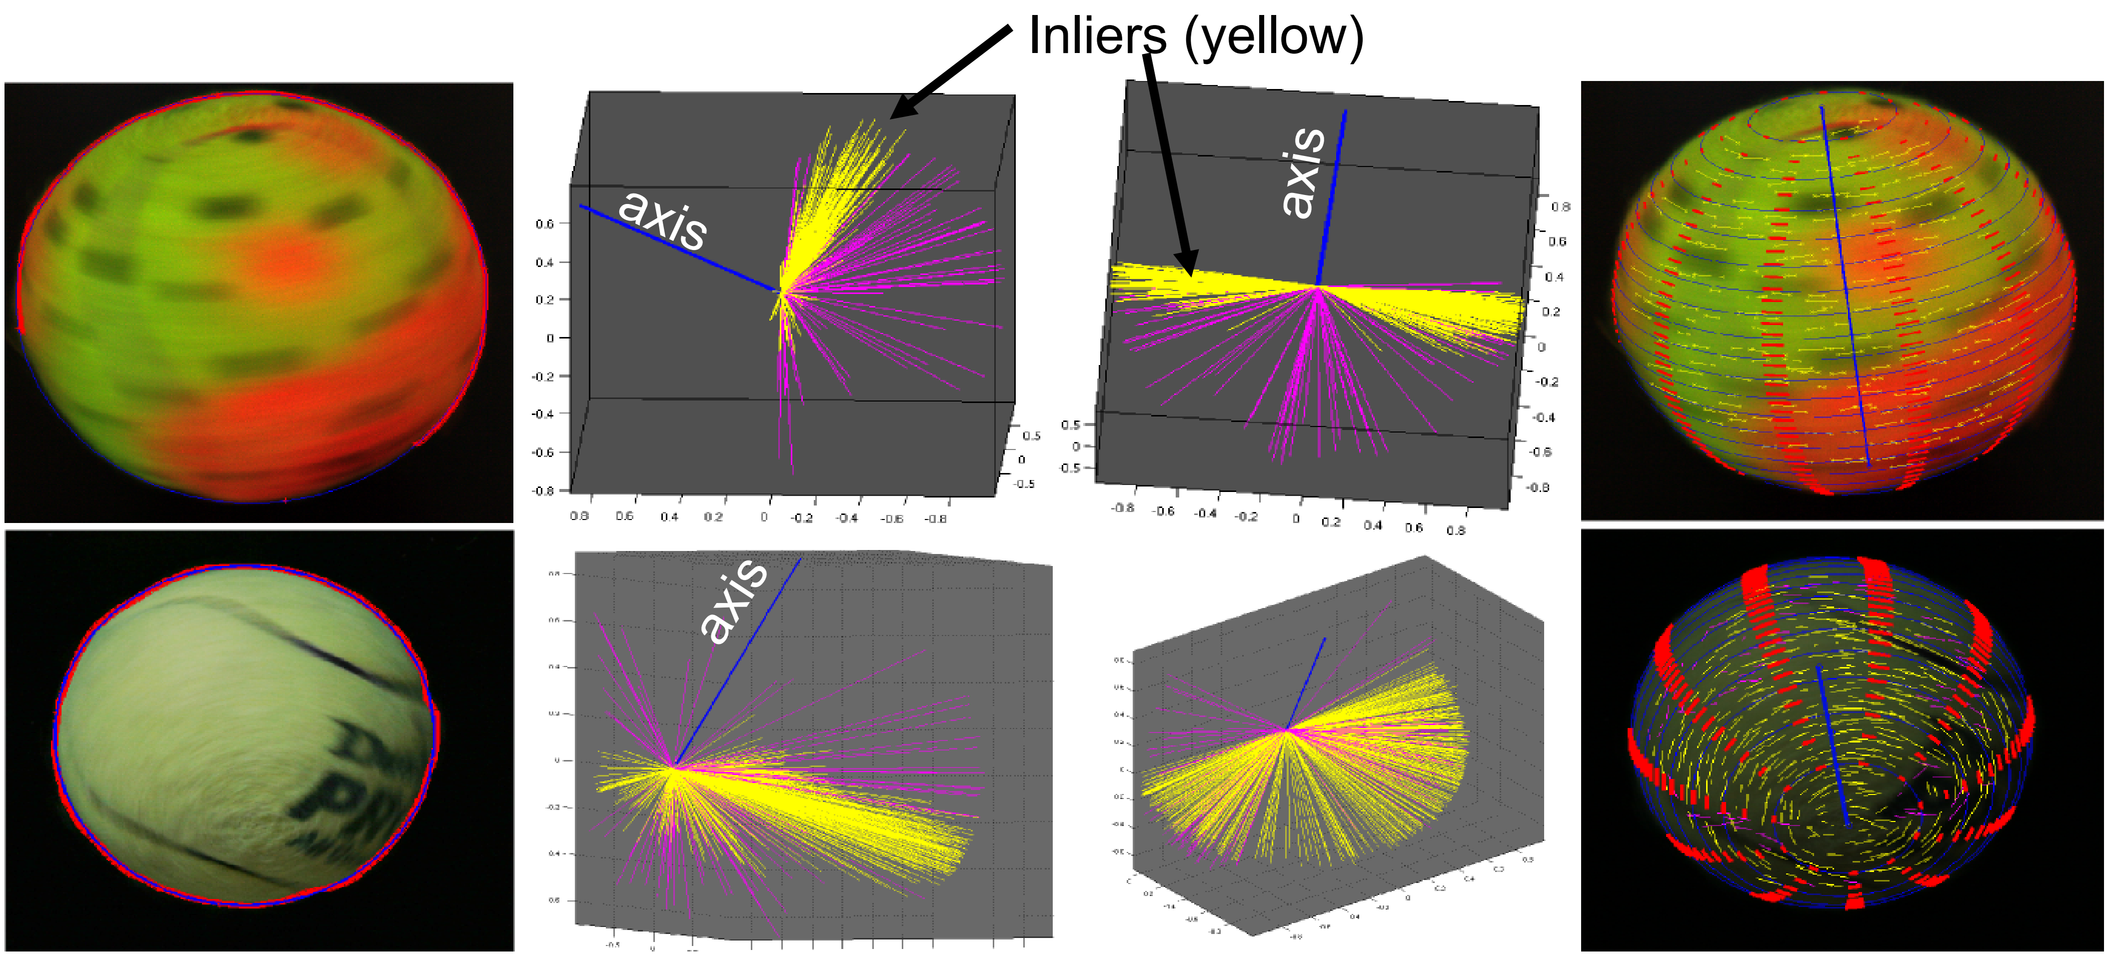

Simplified case (4) -- spin angle recovery

- Spin angle during exposure is robustly recovered as median of angles:
  - backprojection of local blur lenghts on ball surface (excluding outliers)
- spin angular speed immediately follows

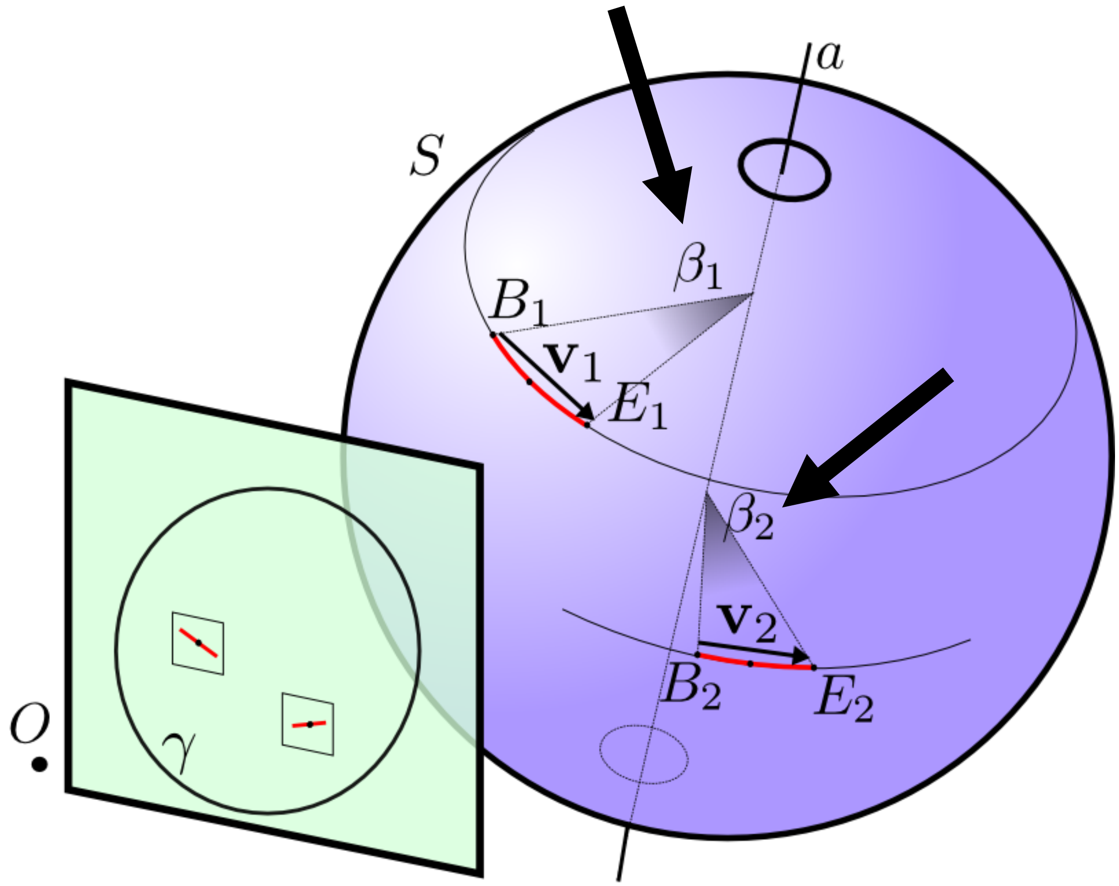

General case: translation + spin

The spin+translation case is much more difficult

General case (2): position and velocity estimation

- First step: find ball position and velocity using known technique on alpha map only [ICIAP2007]
- Find apparent motion direction (radon-like transform)
- Analyze parallel profiles
- Find gradient discontinuities along profiles → contours of the ball at beginning and end of the exposure
- Fit initial and final ellipses
- 3D localization → position and velocity

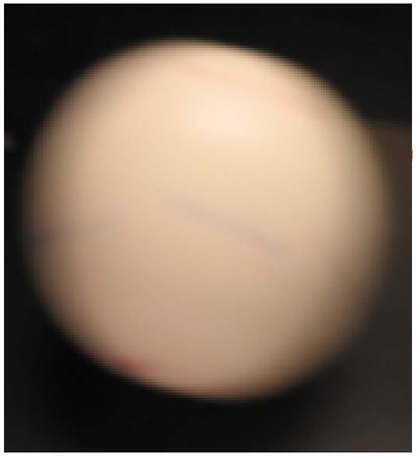

General case (3): geometry of spin recovery

Must backproject blurring path endpoints on initial/final spheres!

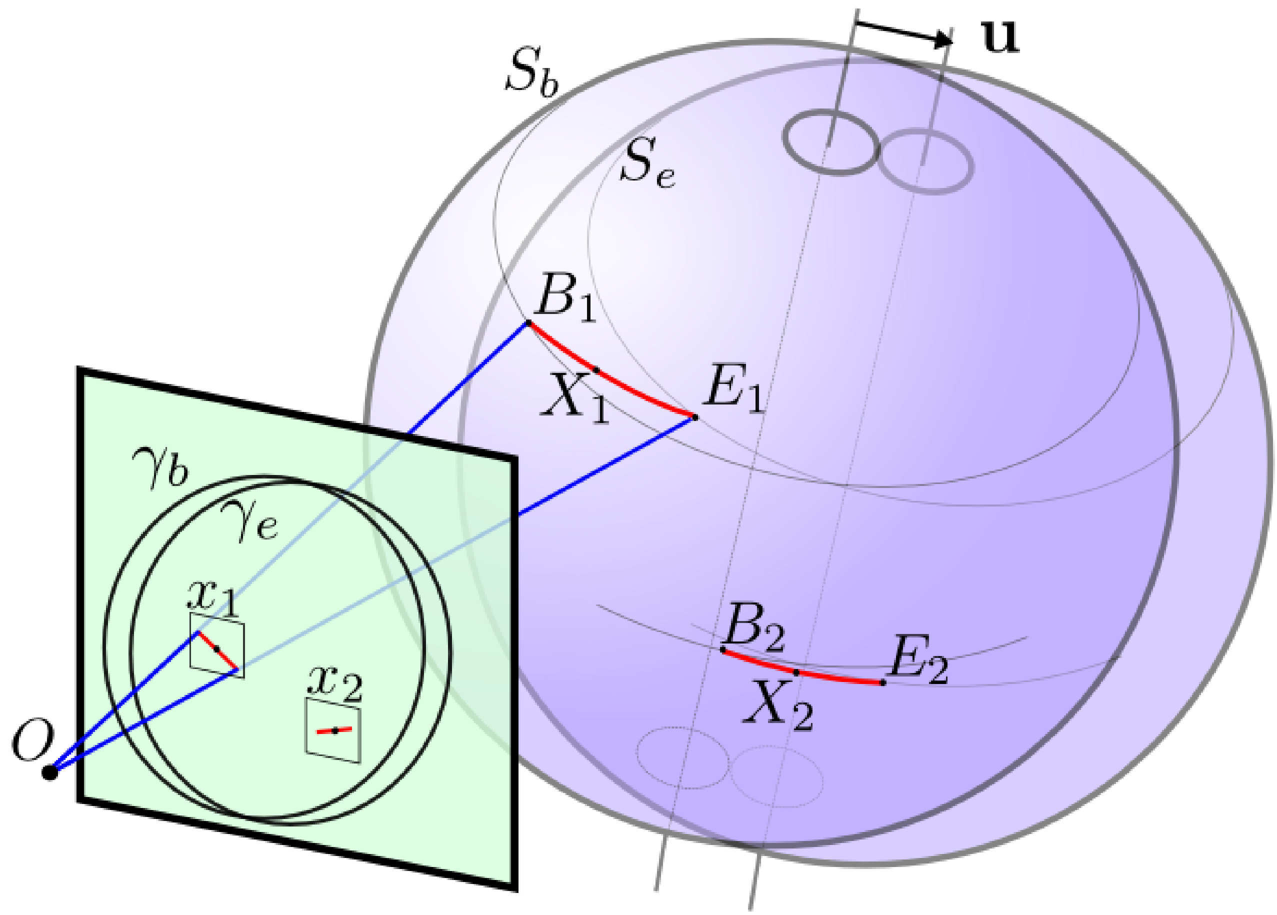

General case (4): orientation problem

- No easy way to determine which sphere each endpoint must be backprojected to.

→ Heuristics:

  - smooth vectorial field
  - consistent directions if velocity dominates spin

- - - ▶ Voting or RANSAC-like approaches possible

Experimental results

Remarkable accuracy on the spin-only case, both in simulated and real images

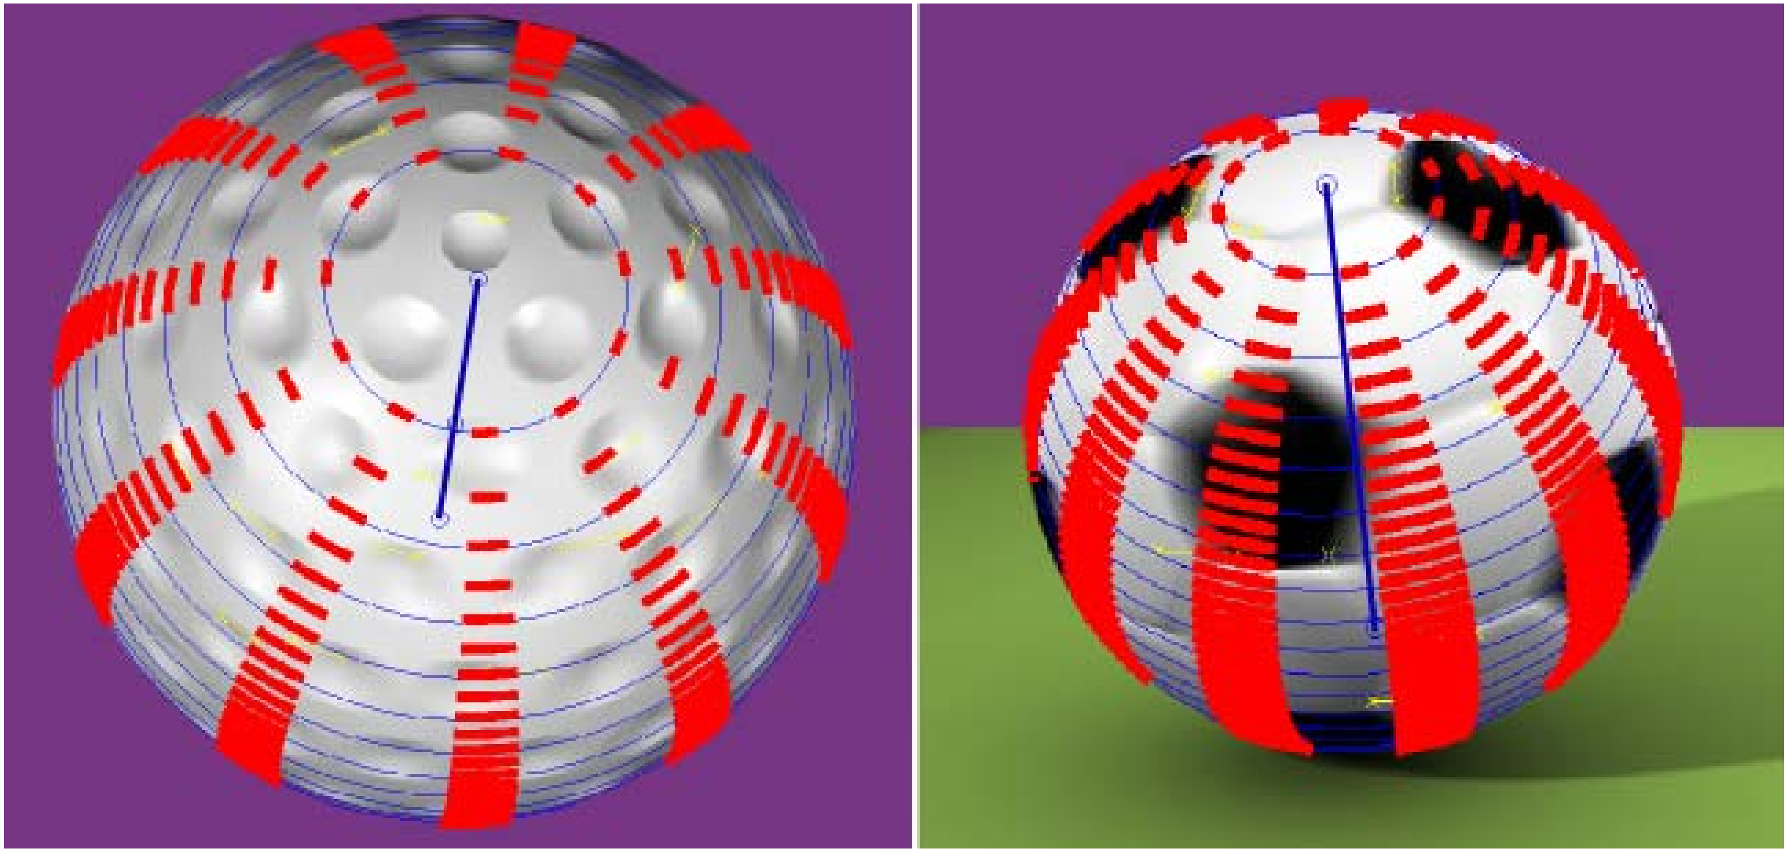

Specularities and bumps in ball surfaces are handled well although not explicitly modeled

Experimental Results (2) – Quantitative data

Mean relative error (percent) in spin angular speed estimation on synthetic images

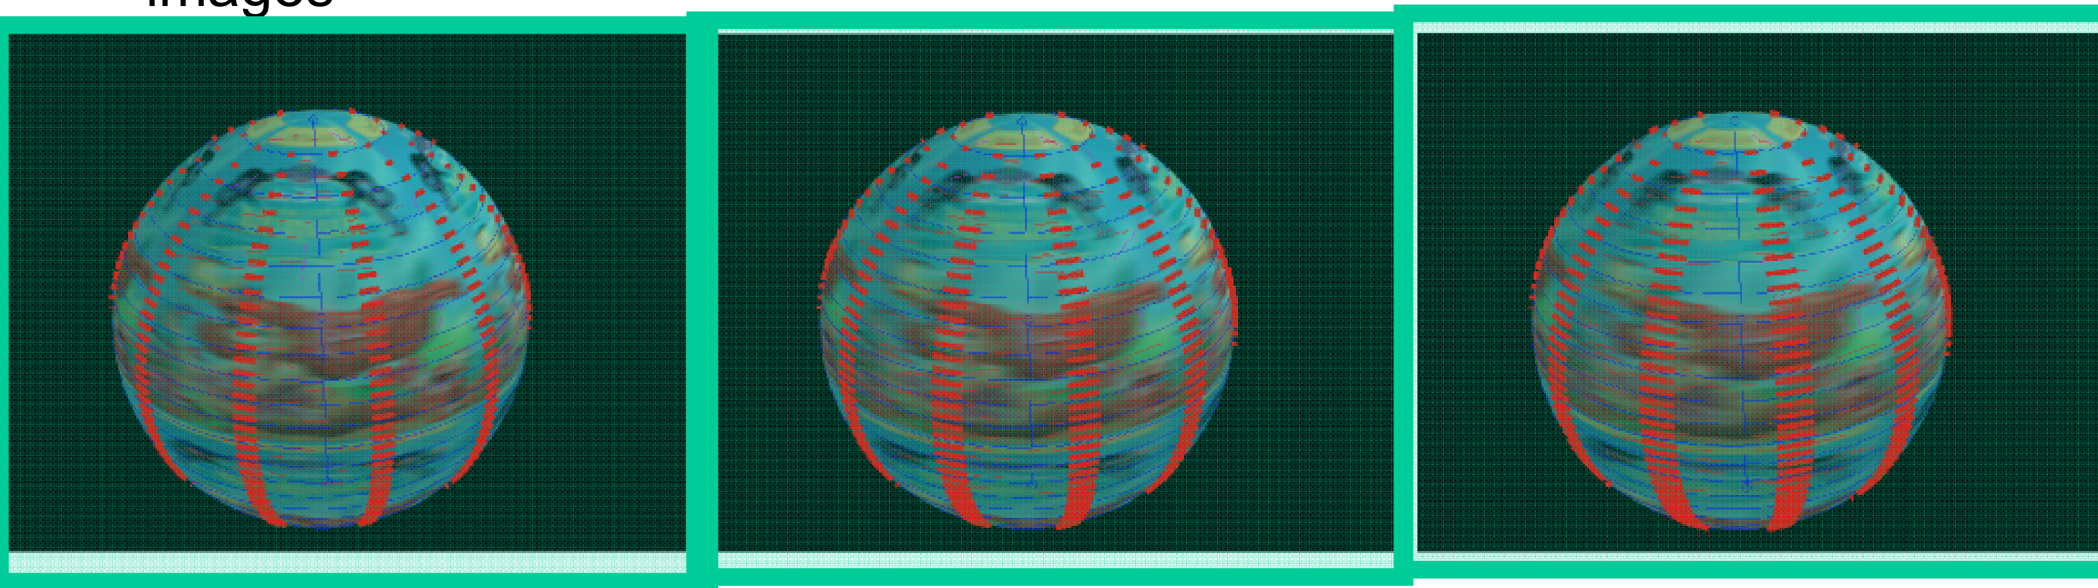

Spin angle (deg)

Noise

$\omega \cdot T \setminus \sigma$	0	1	2	3
5.00	4.31	4.6222	5.0641	3.9401
6.25	2.26	2.5562	4.7898	4.3915
7.50	2.40	3.1353	2.7236	2.0544
8.75	0.75	1.5163	2.9408	5.0431
10.00	2.15	3.3975	5.3916	11.3800

Experimental Results (3) – Real images

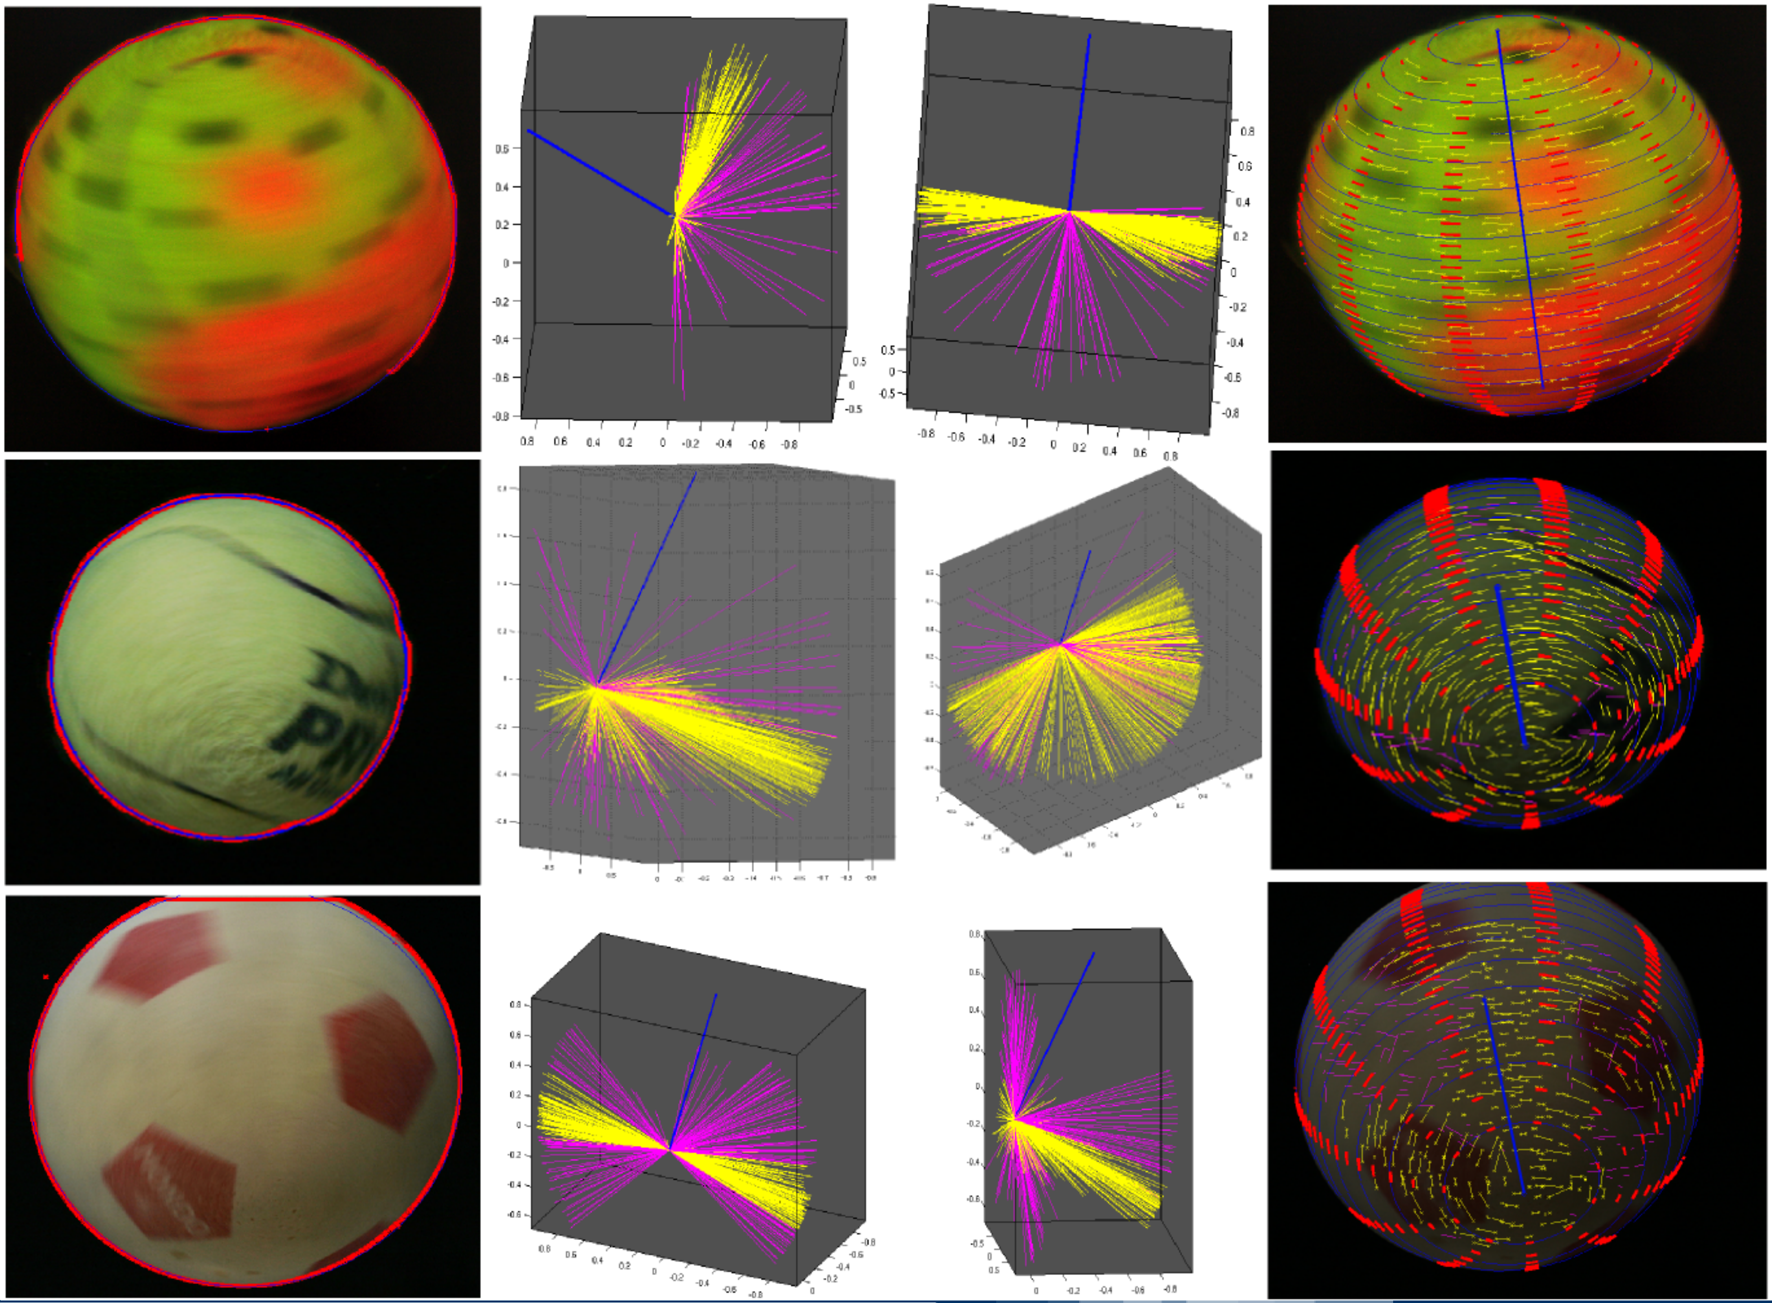

Experimental Results (4) – Real images

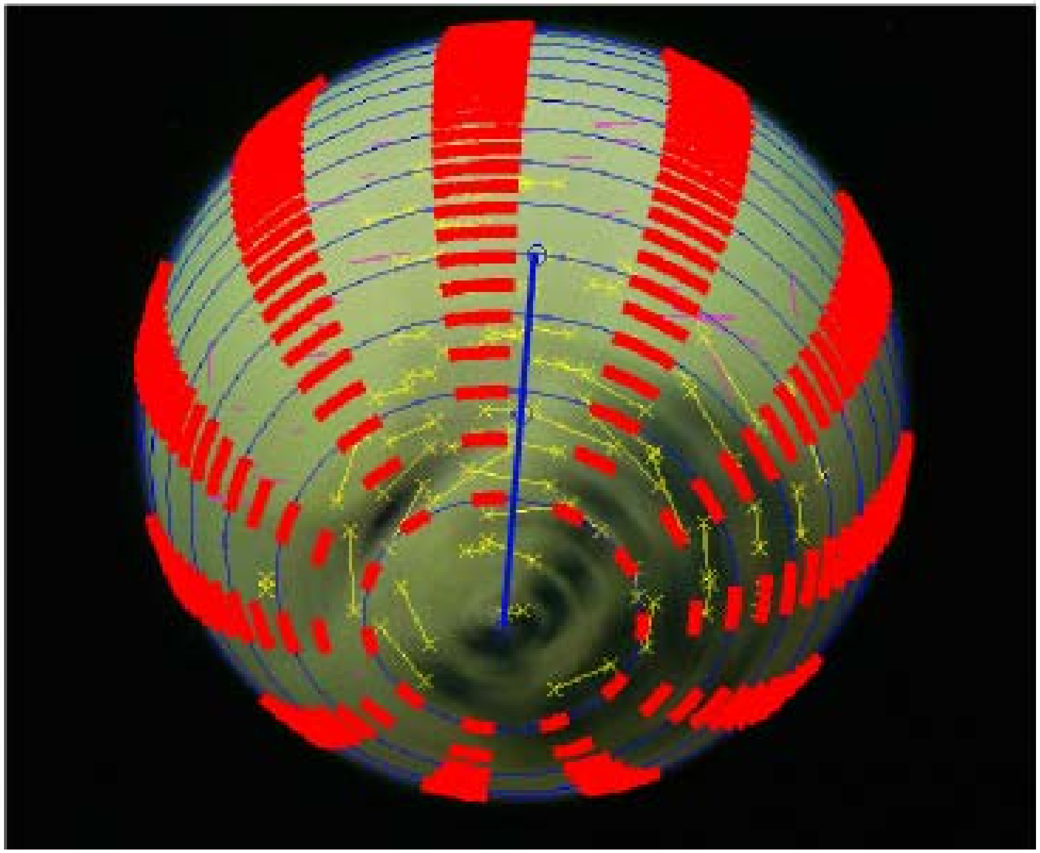

Experimental Results (5) – Real images

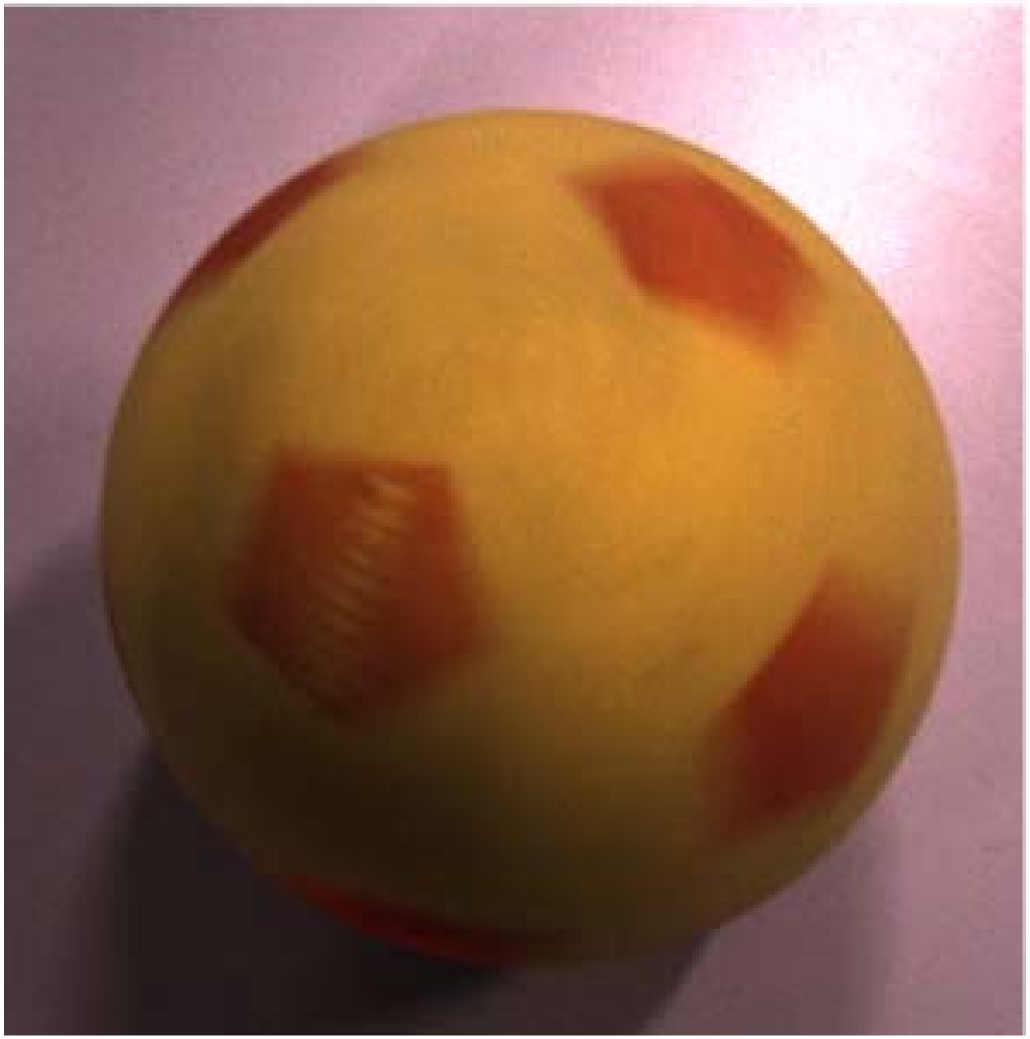

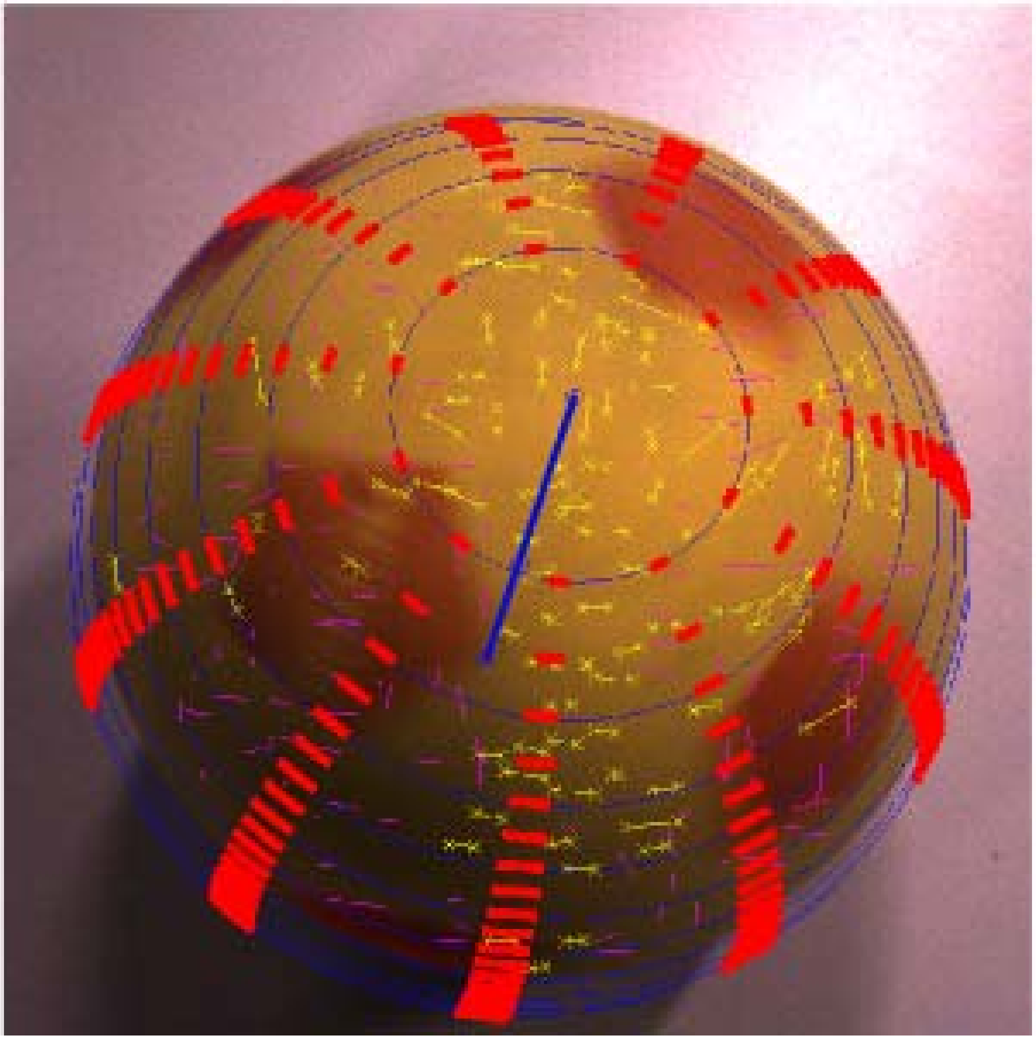

Experimental Results (6) – Real images

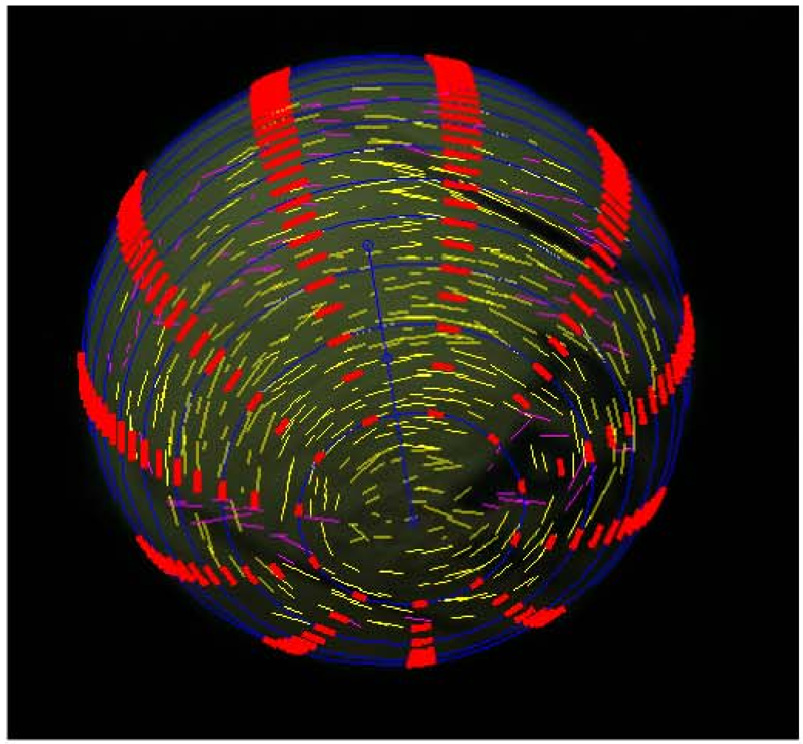

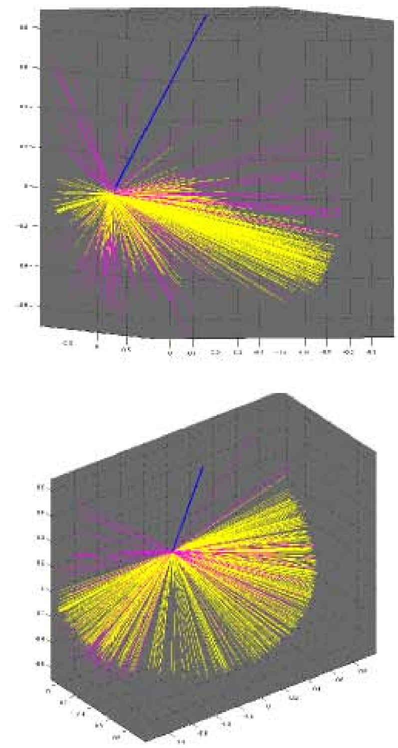

Experimental Results (7) – General case

- Correct: ball was rolling on the table:
- displacements on contact point are null
- Axis is parallel to the table

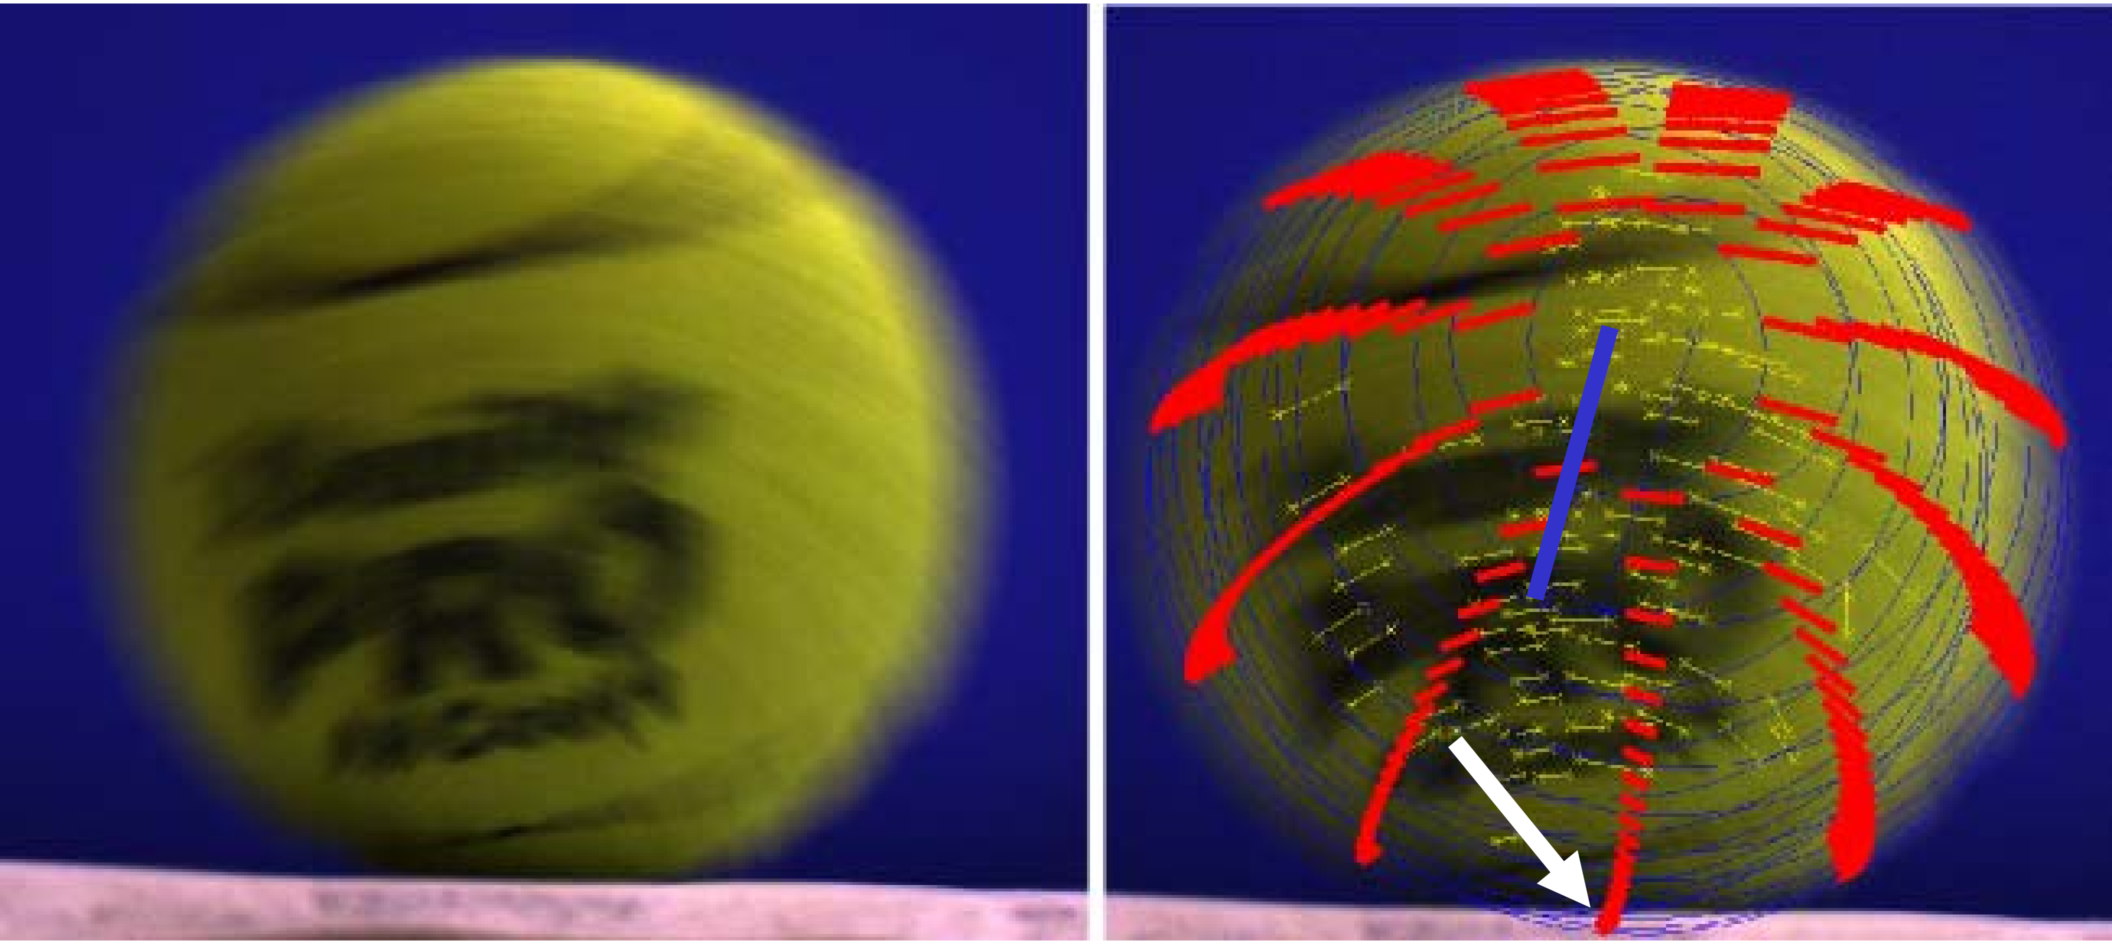

Ongoing

- The general case is not at all robust and still problematic:
  - Heuristics for blur orientations often fail
  - Current technique needs blur extents, which are not reliable due to curvilinear blurring paths
- We are developing a voting-based approach which does not require blur extents
- Alternative techniques can be used instead of minimum derivatives [Caglioti, Giusti PACV2007] with promising results, handling more difficult cases:

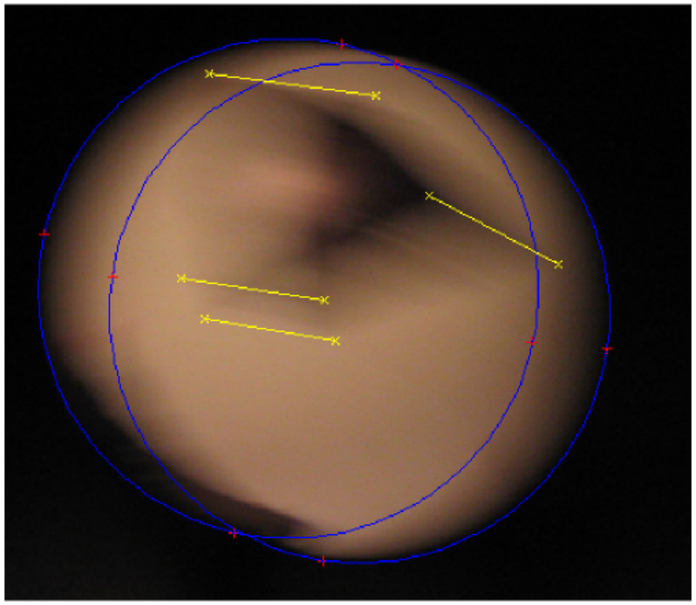

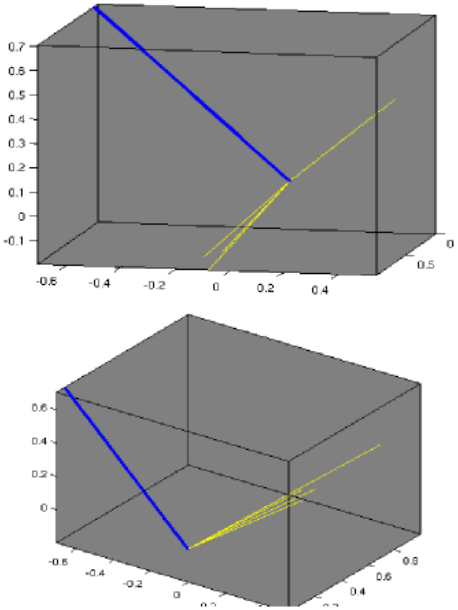

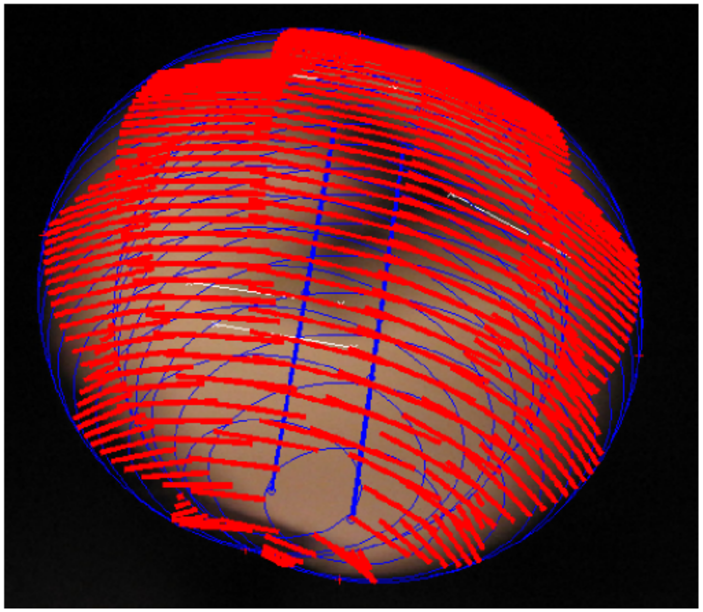

Conclusions

- A single high resolution photograph can sometimes be a viable approach for motion reconstruction
  - No need for solving the correspondence problem
- The known scene structure is an important prerequisite.
- We separate the problem into translation and spin recovery thanks to alpha matting
- The orientation of local blur estimates is unknown and poses important challanges