School of Psychology and
Clinical Language Sciences

University of Reading

LINKS BETWEEN REINFORCEMENT LEARNING AND HUMAN REPRESENTATIONS OF SPACE

Andrew Glennerster

Collaboration is good

Previous MURI talks – policy networks

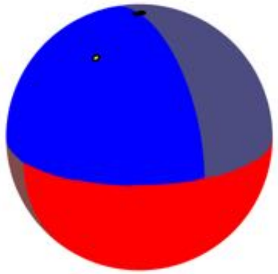
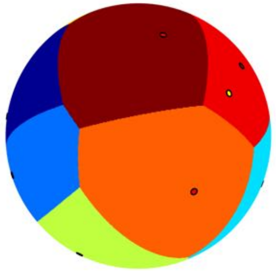
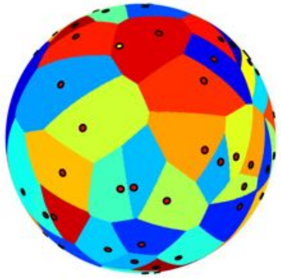
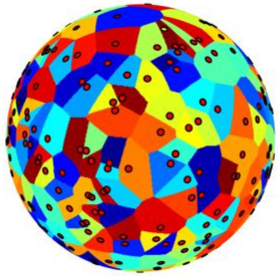
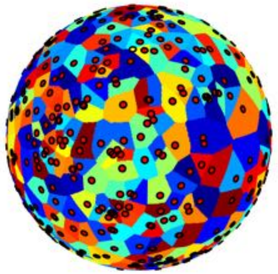

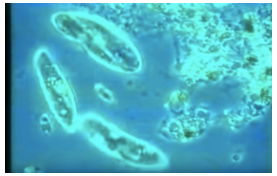

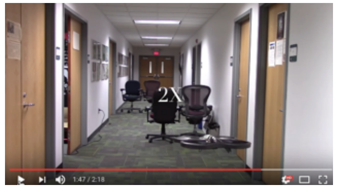


?

- bacterium
- plants
- RL
- humans
- higher intelligence

Previous MURI talks – the wrong basis set

Unsupervised Visual Representation Learning by Context Prediction

Carl Doersch[1,2] Abhinav Gupta[1] Alexei A. Efros[2]
[1] School of Computer Science, Carnegie Mellon University
[2] Dept. of Electrical Engineering and Computer Science, University of California, Berkeley

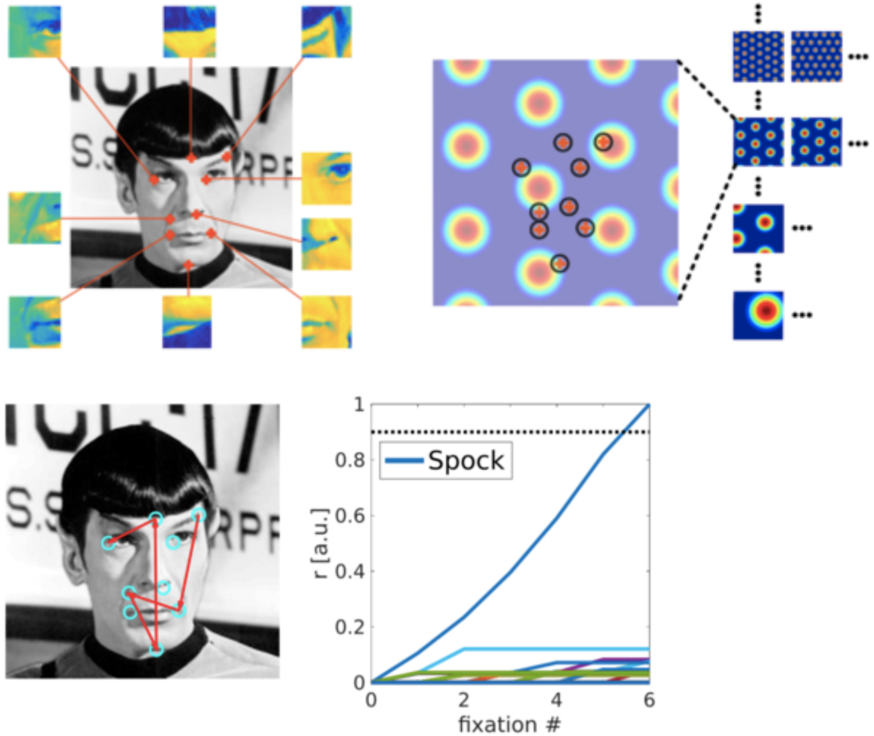

A Computational Model of Visual Recognition Memory via Grid Cells

Andrej Bicanski[1,2,*] and Neil Burgess[1,*]

Previous MURI talks – growing dimensions when required

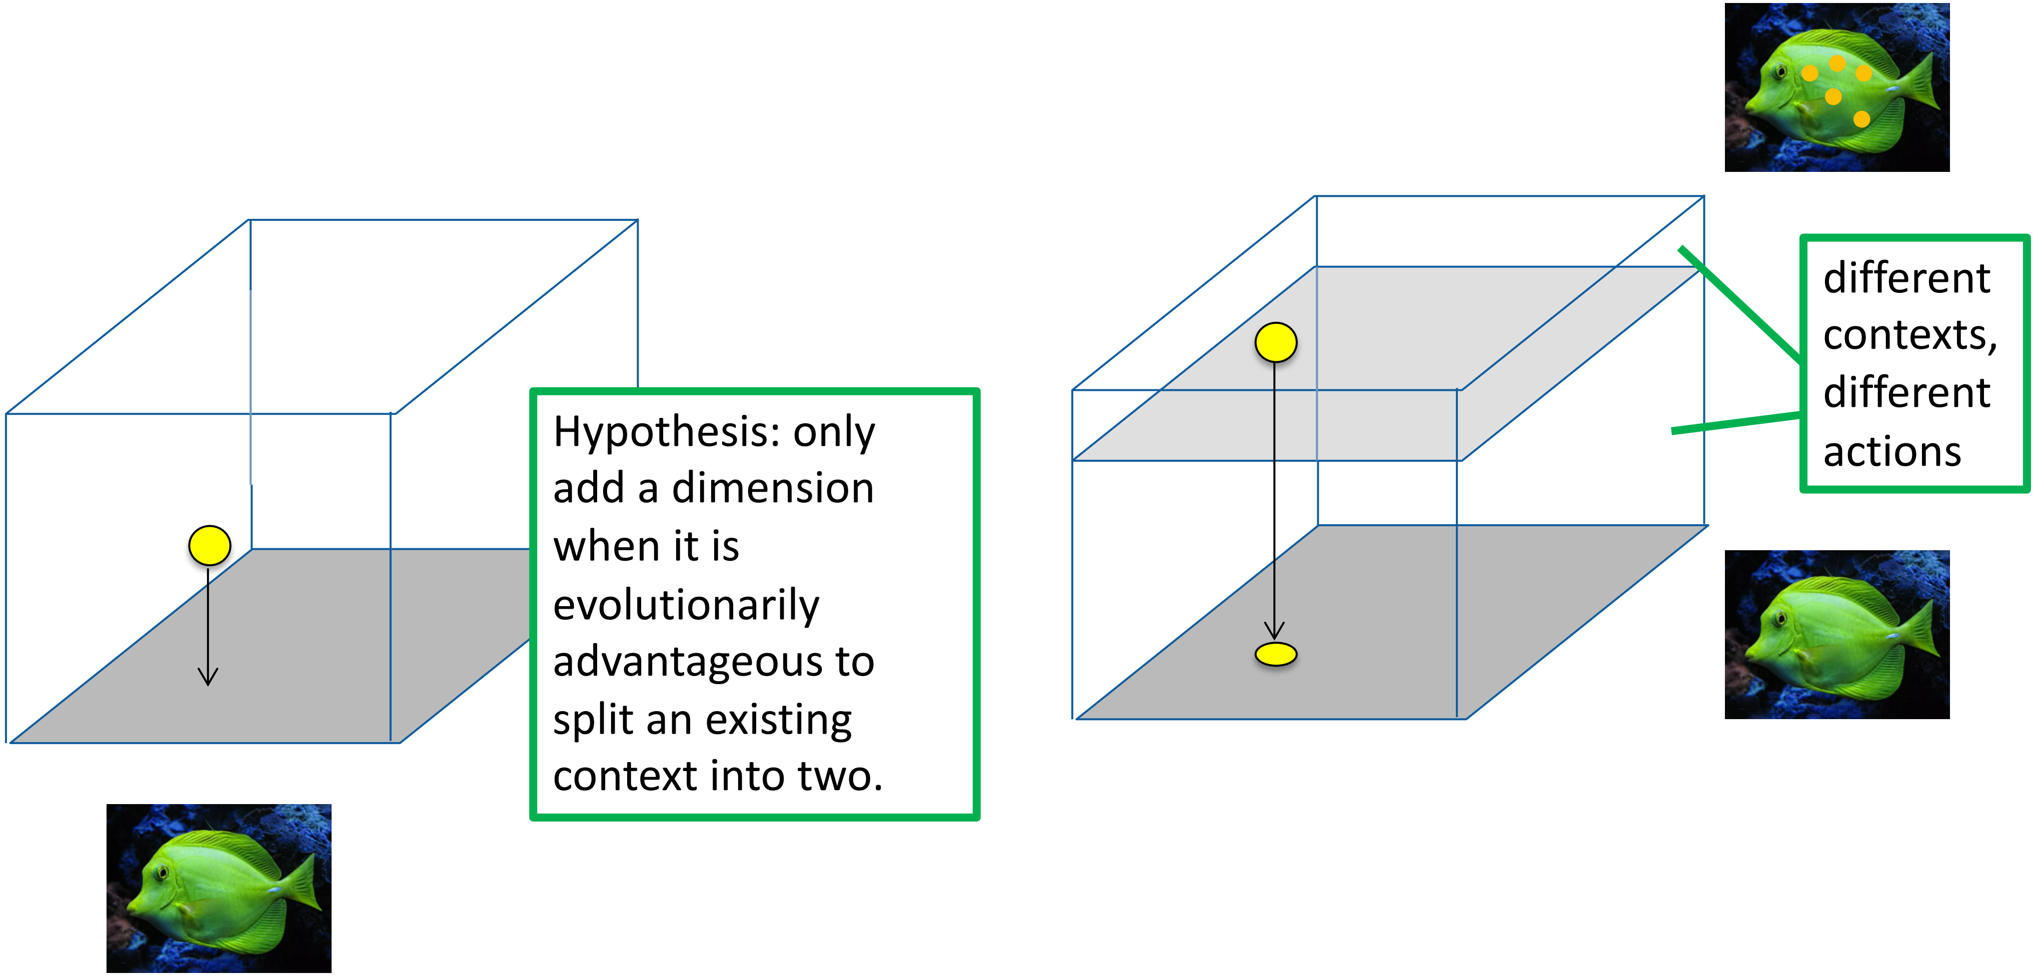

Summary

- More collaboration would be good

- Π(a|s,g) for neuroscience

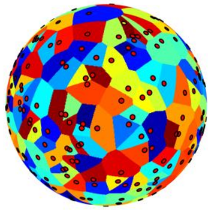

- RL/DNN use the wrong basis vectors to represent images (wrong in the DoD sense, i.e. not 'human-like').

- Learn basis vectors gradually, in tandem with gradual development of task complexity

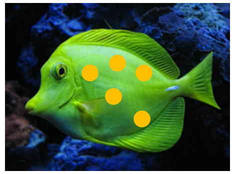

- RL for navigation could be improved by learning which features are likely to persist over longer translations of the camera

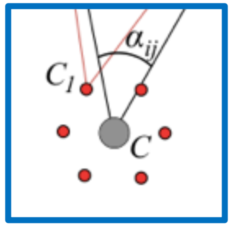

Zhu et al (2016)

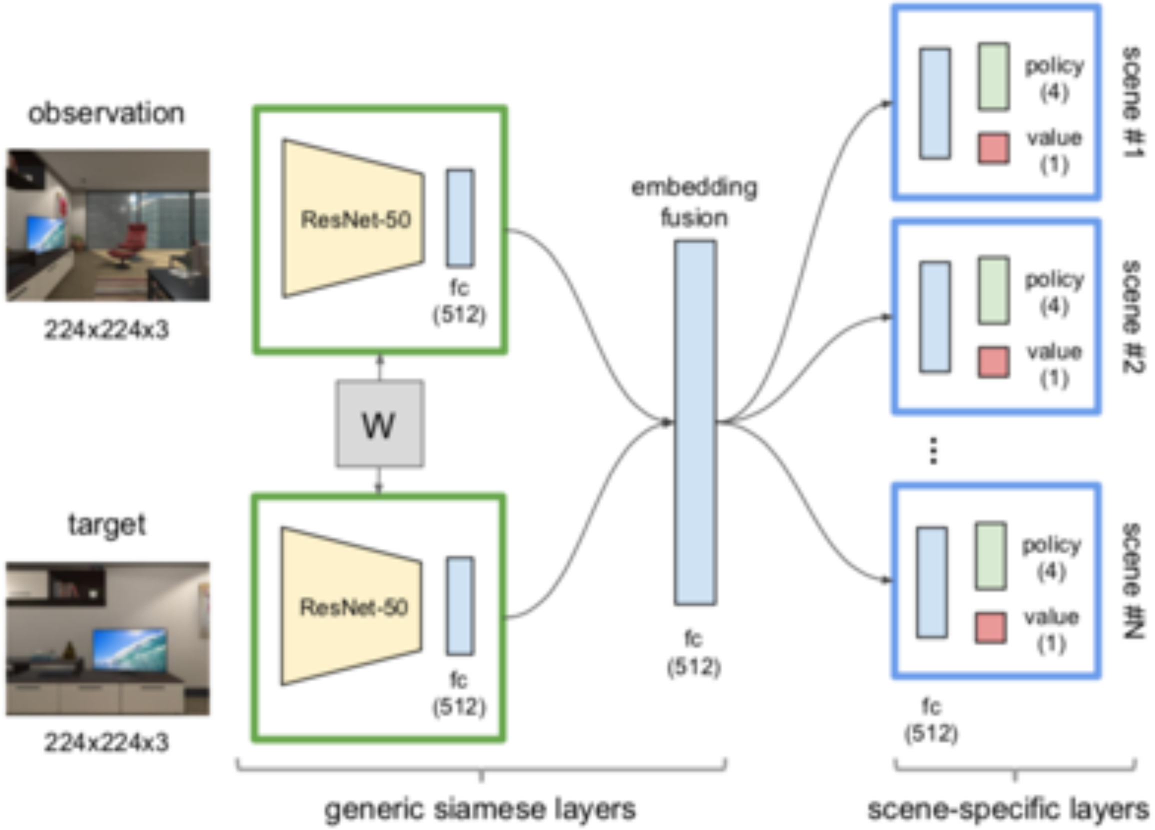

Zhu et al, 2016 https://youtu.be/SmBxMDiOrvs.

Zhu et al – tSNE (coded by target)

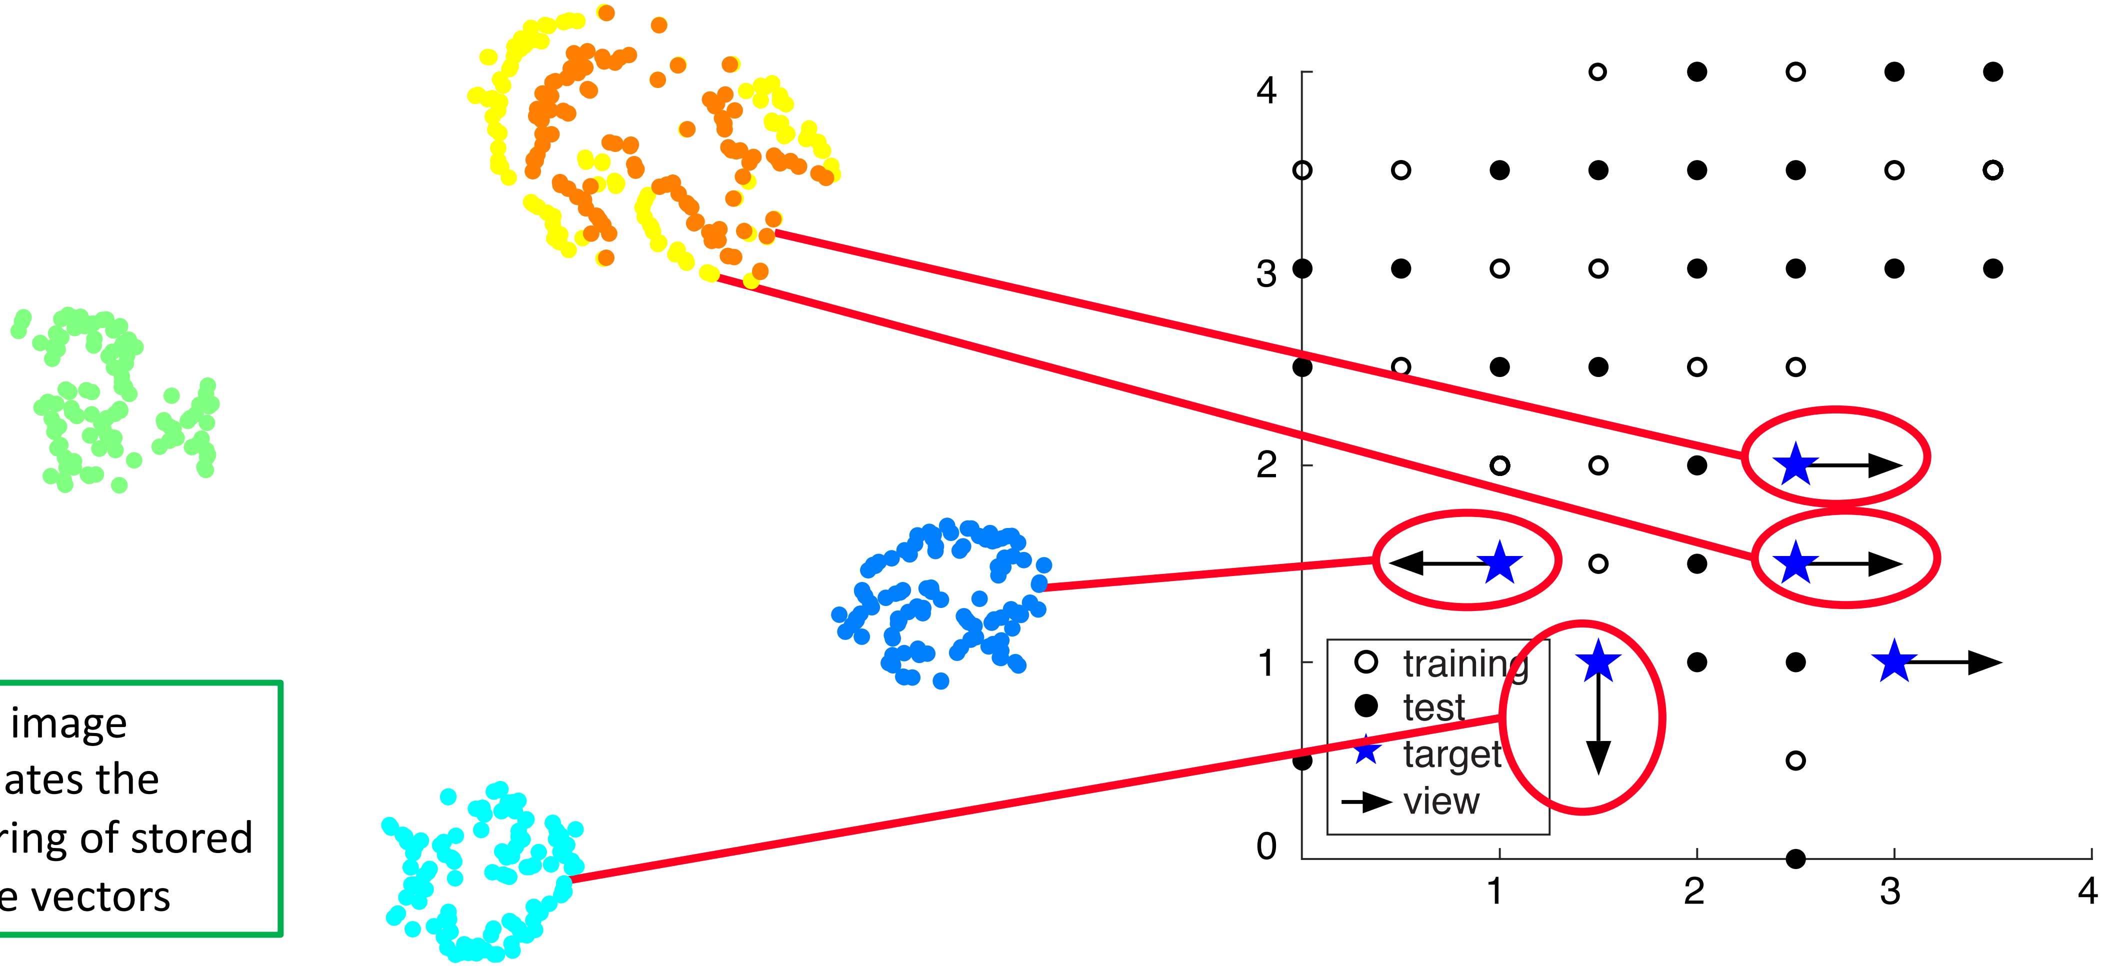

Target image dominates the clustering of stored feature vectors

Zhu et al – tSNE (coded by orient^n)

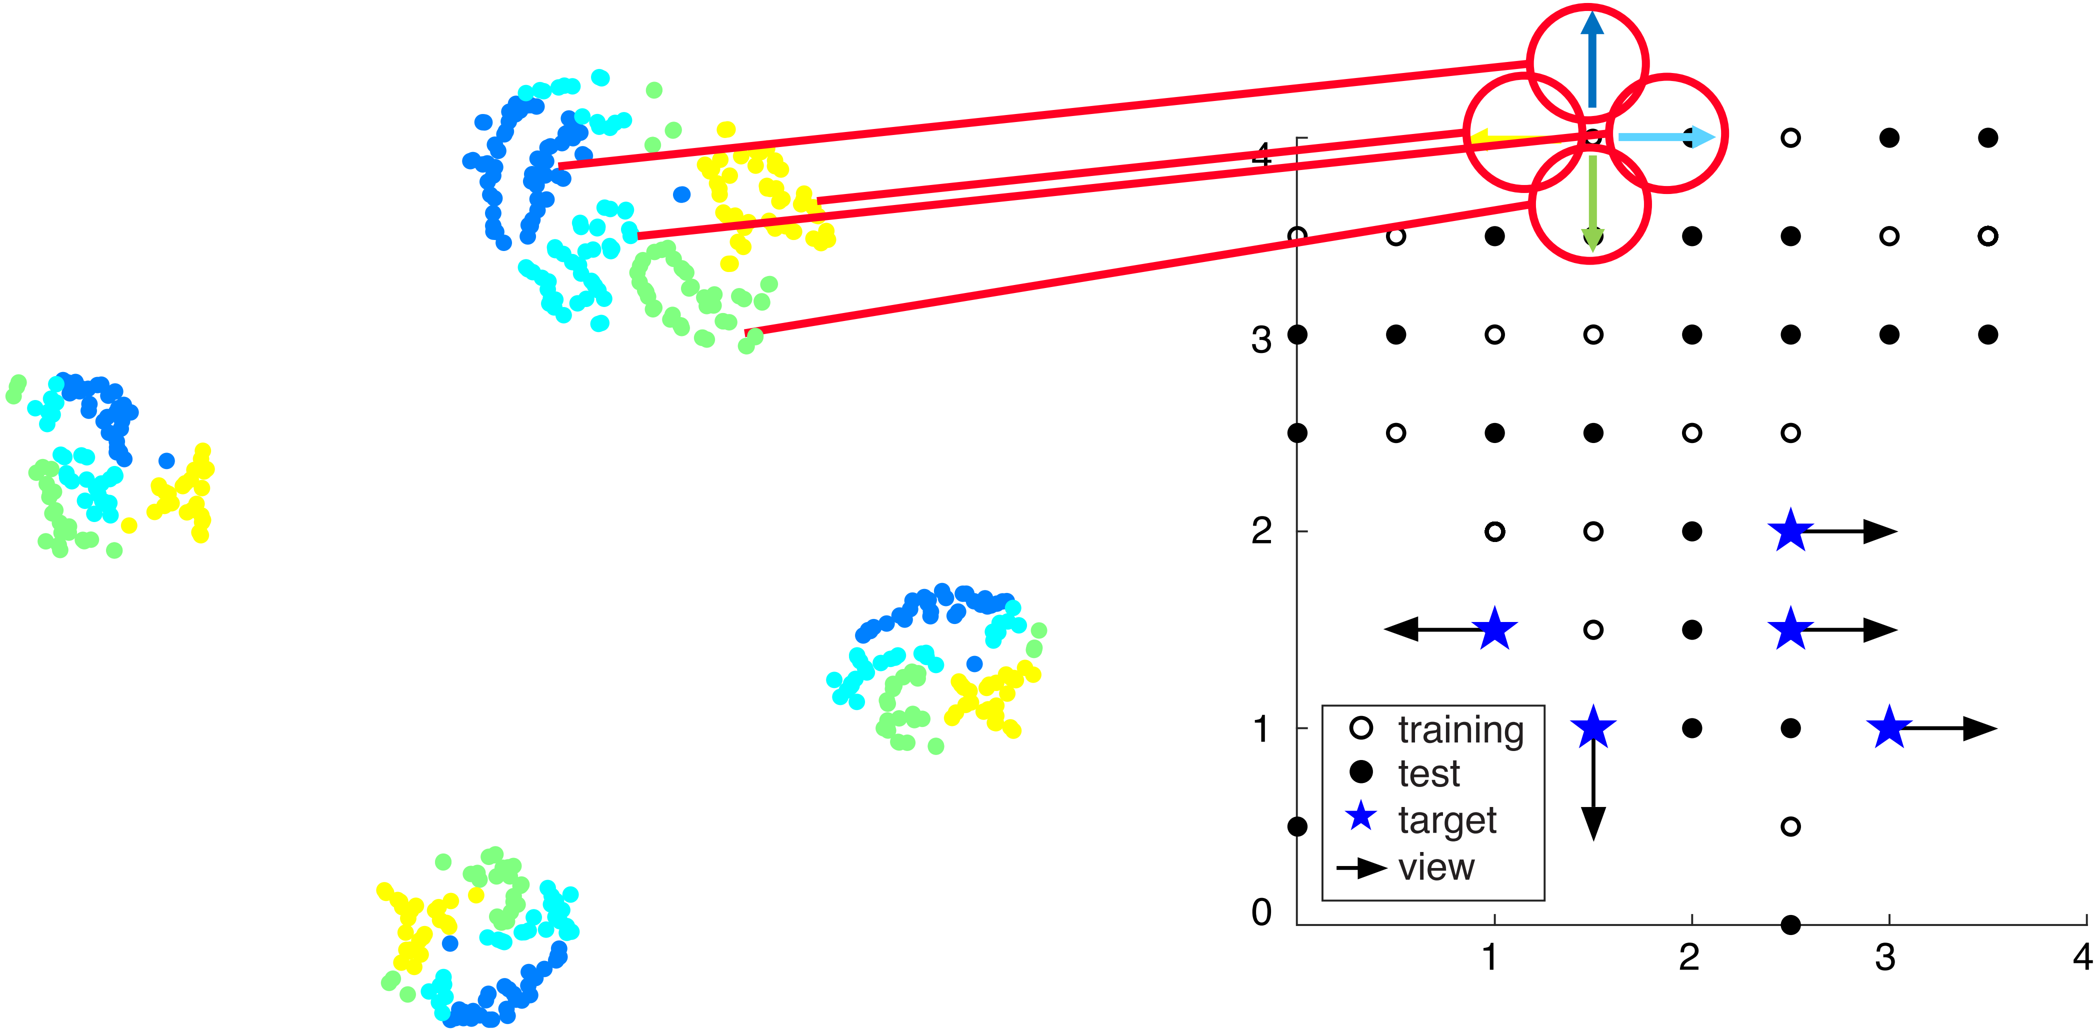

Zhu et al – tSNE (coded by distance)

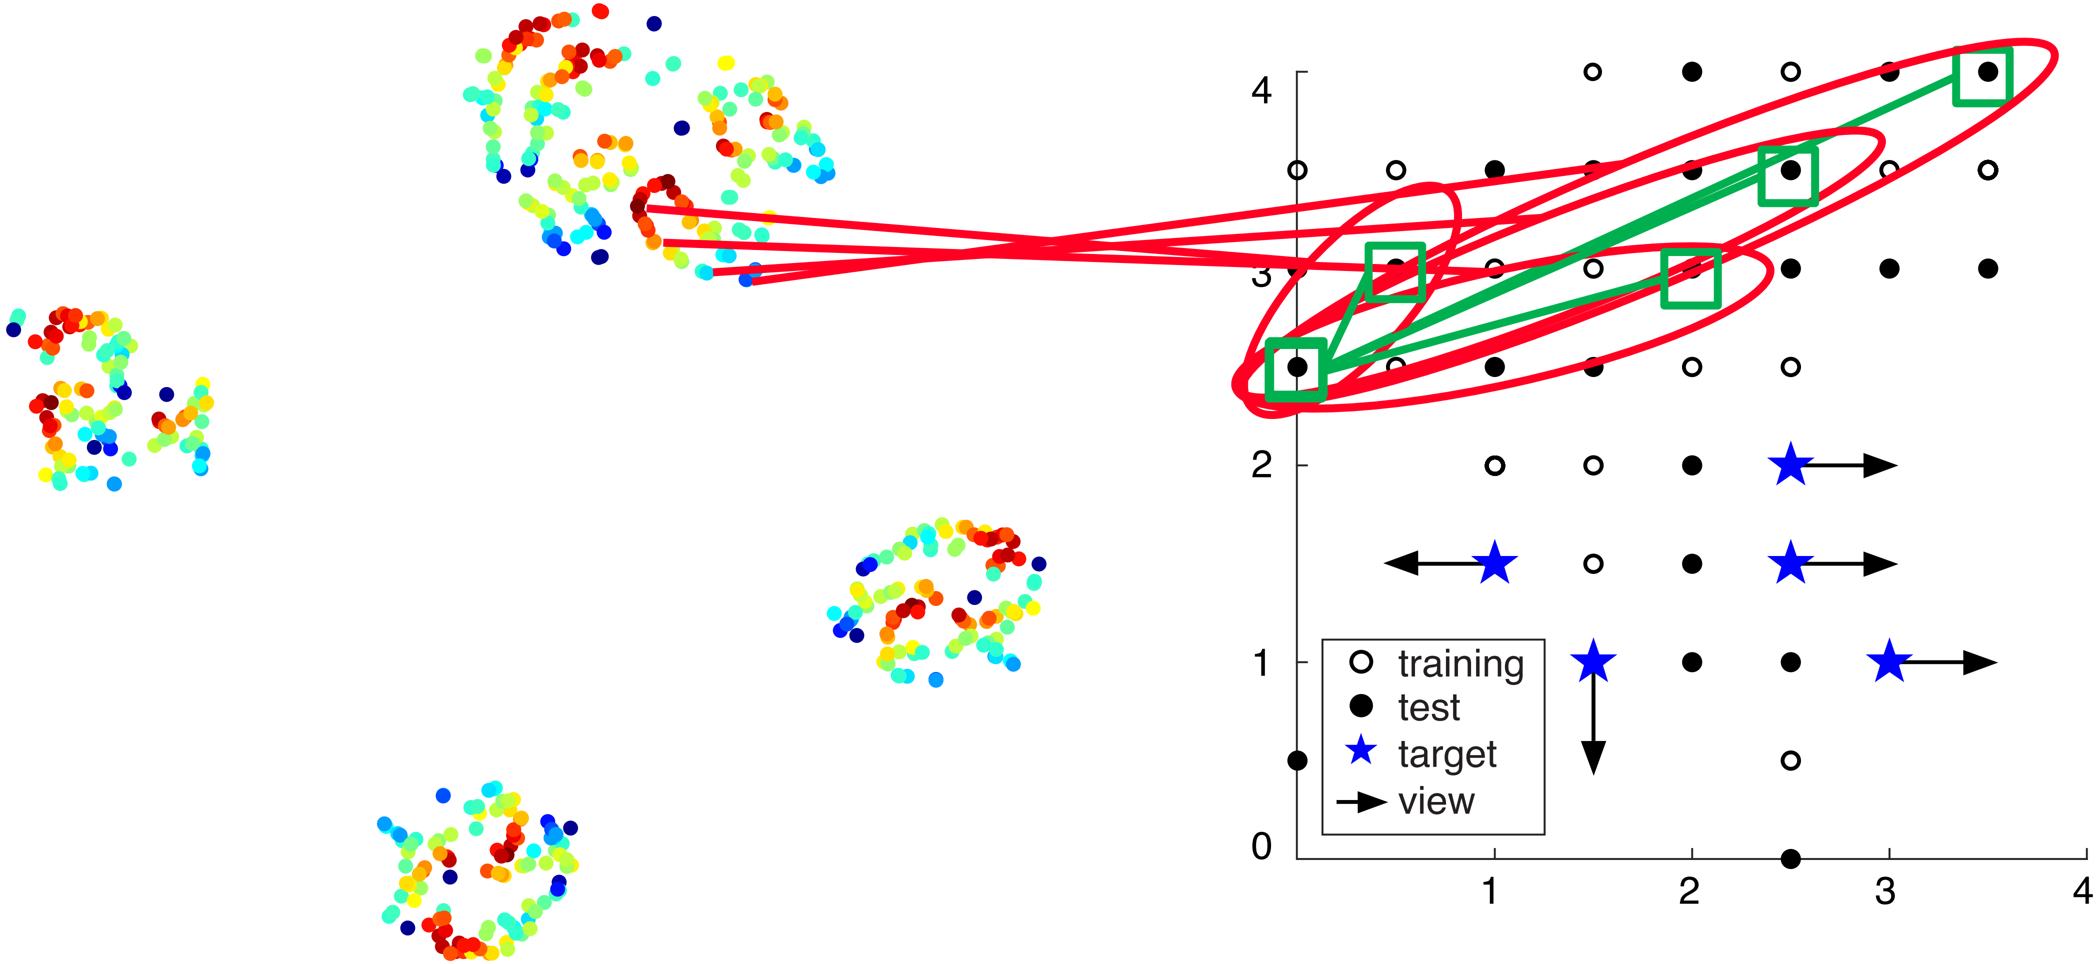

Zhu et al – geometric consistency?

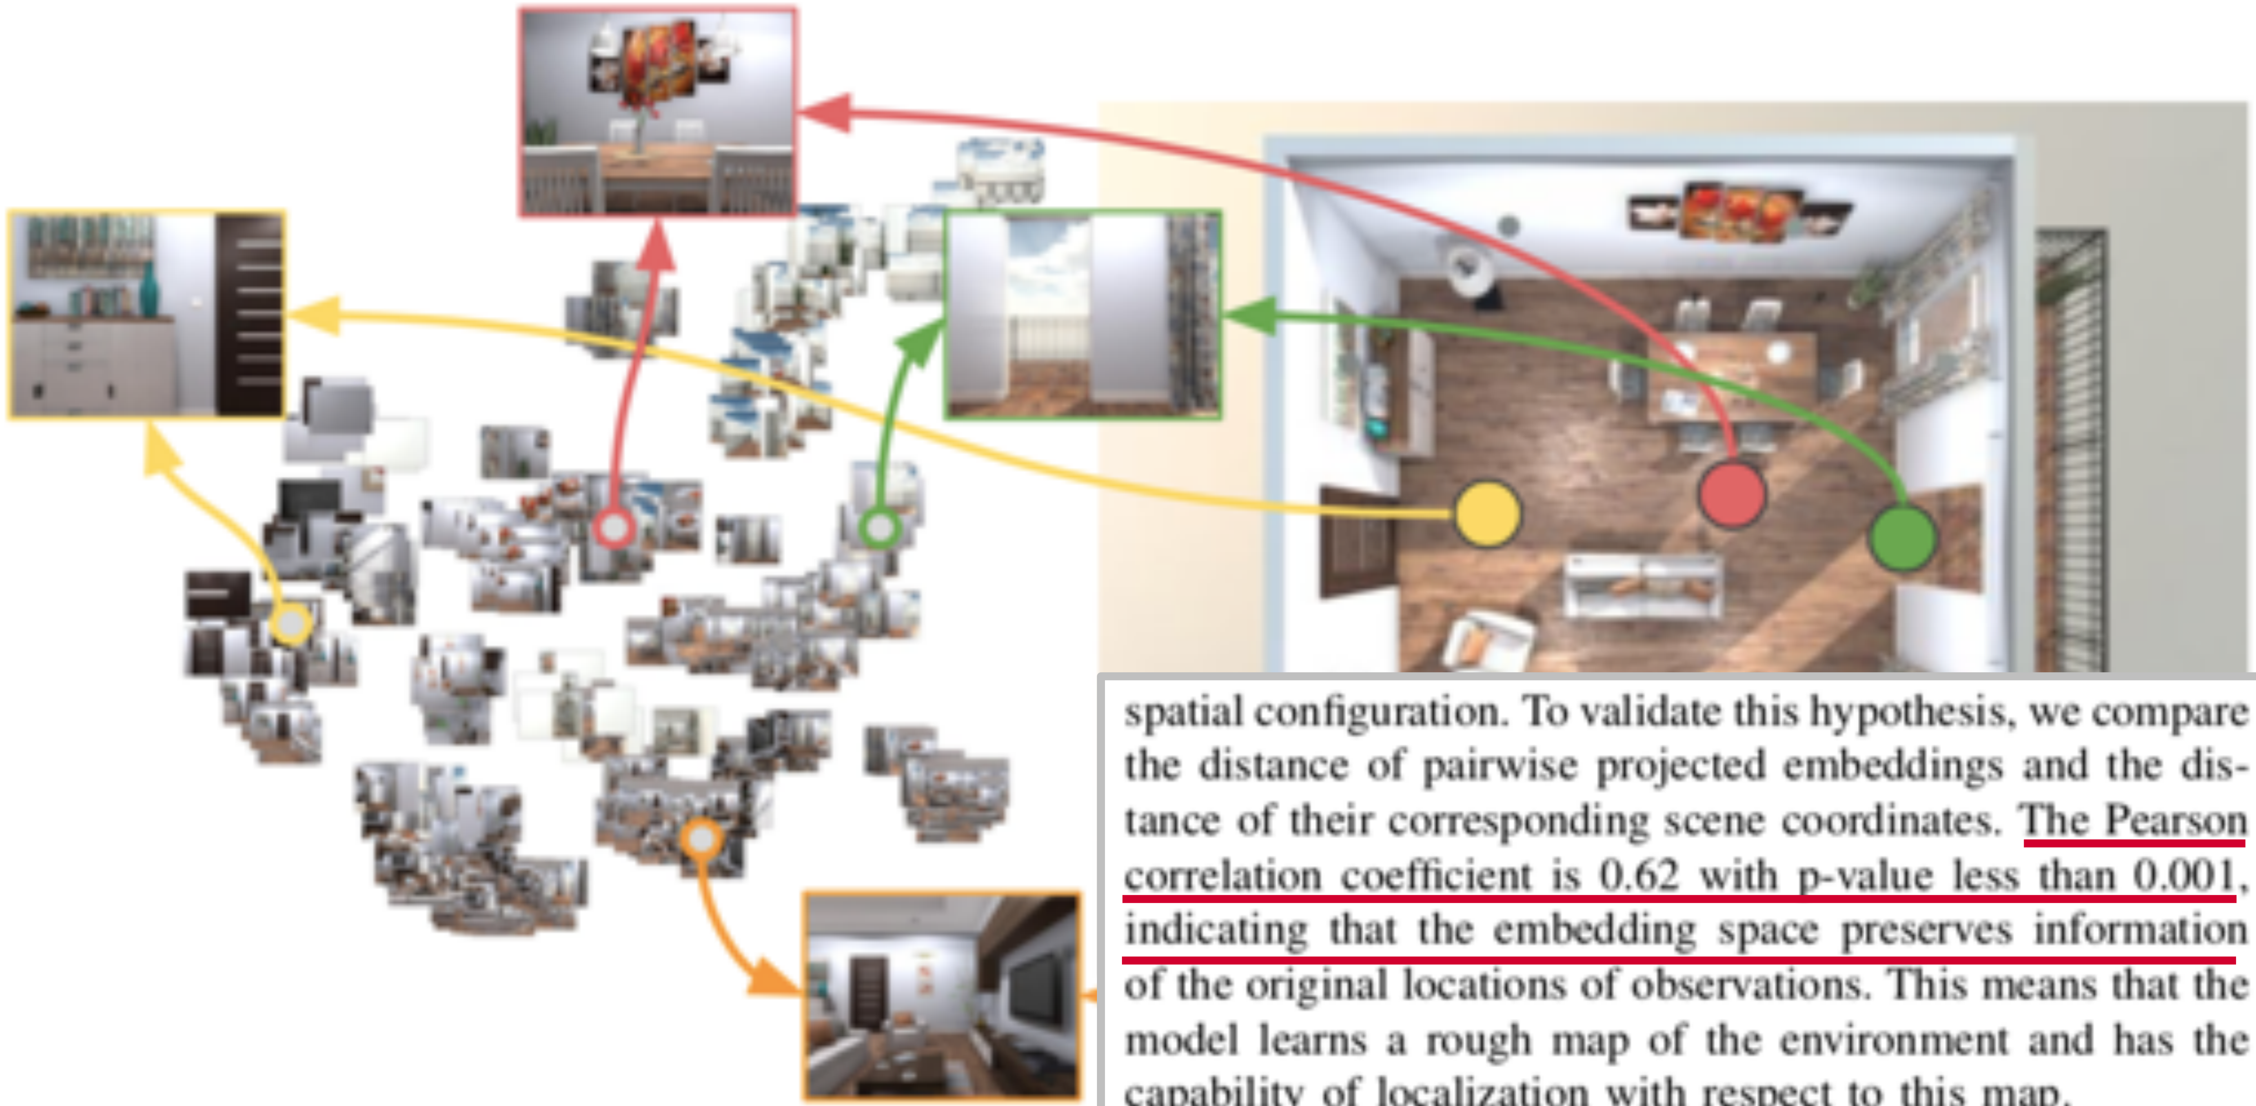

spatial configuration. To validate this hypothesis, we compare the distance of pairwise projected embeddings and the distance of their corresponding scene coordinates. The Pearson correlation coefficient is 0.62 with p-value less than 0.001, indicating that the embedding space preserves information of the original locations of observations. This means that the model learns a rough map of the environment and has the capability of localization with respect to this map.

Alex Muryy

Zhu et al – geometric consistency?

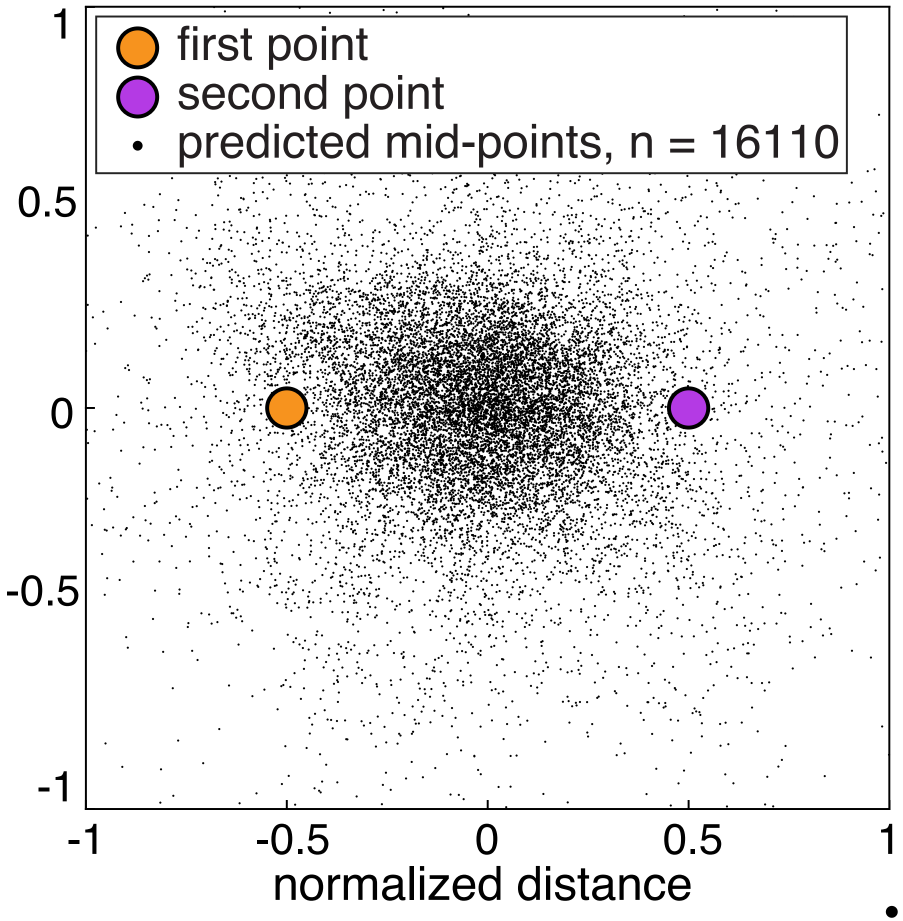

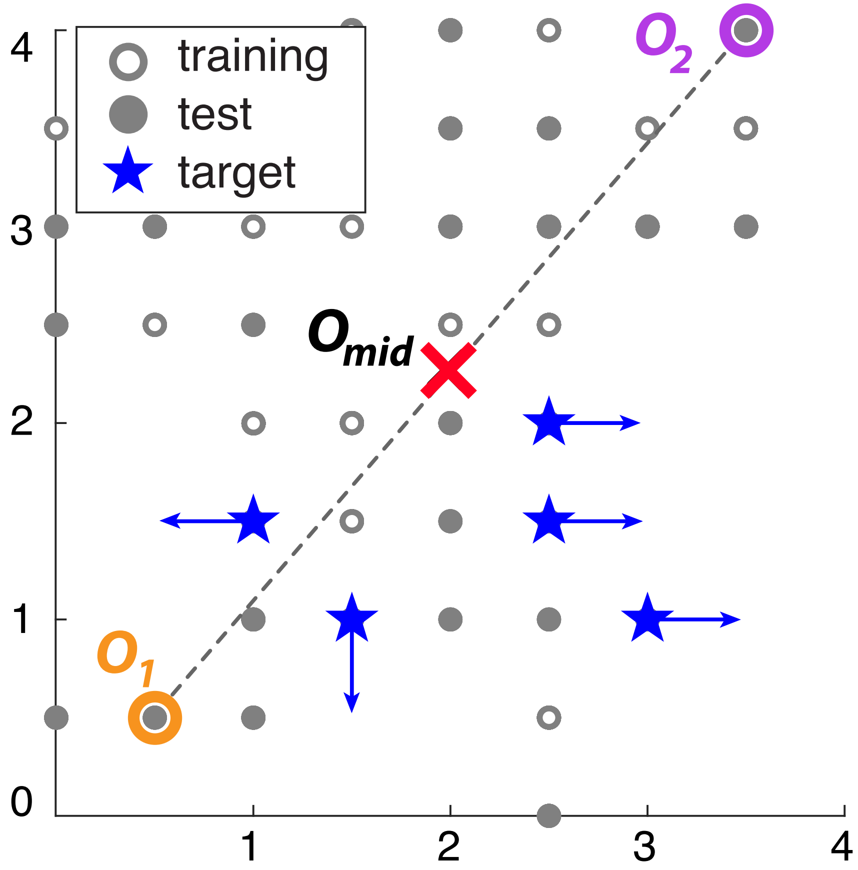

- build decoder for location
- use to calculate midpoint
- with Torr group (Siddharth Narayanaswamy, Nantas Nardelli)

Zhu et al – geometric consistency?

University of Reading

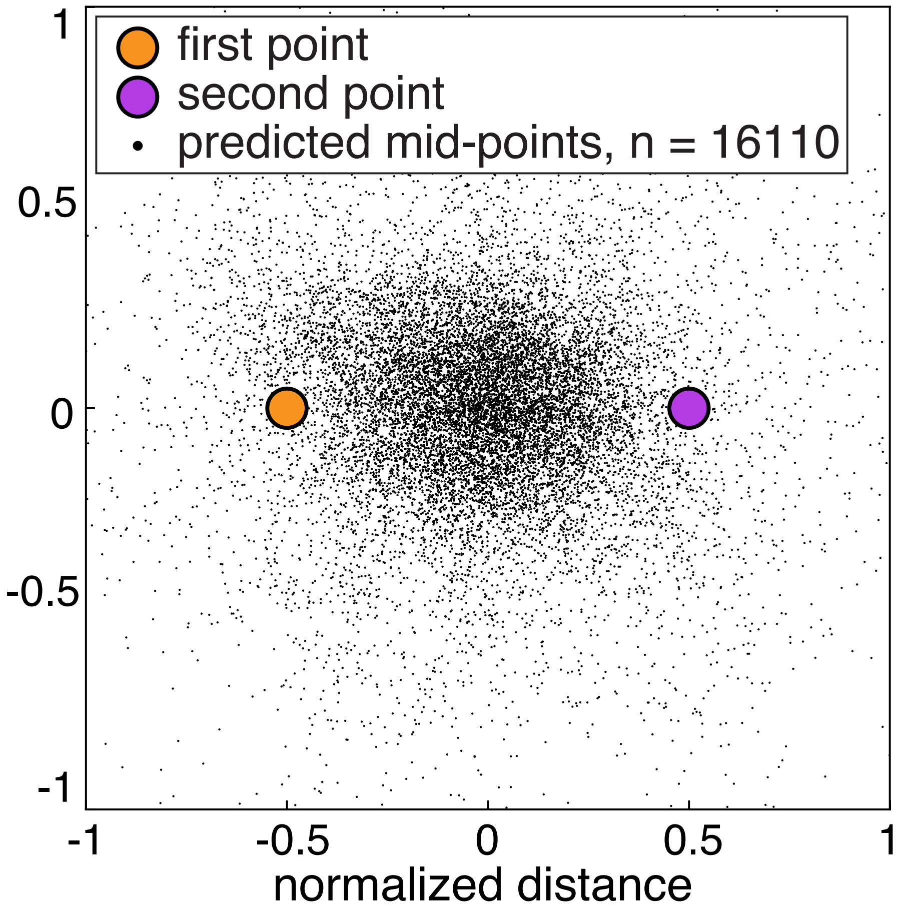

- Geometric consistency?
- Could do better….

Relative visual direction – geometric consistency by design

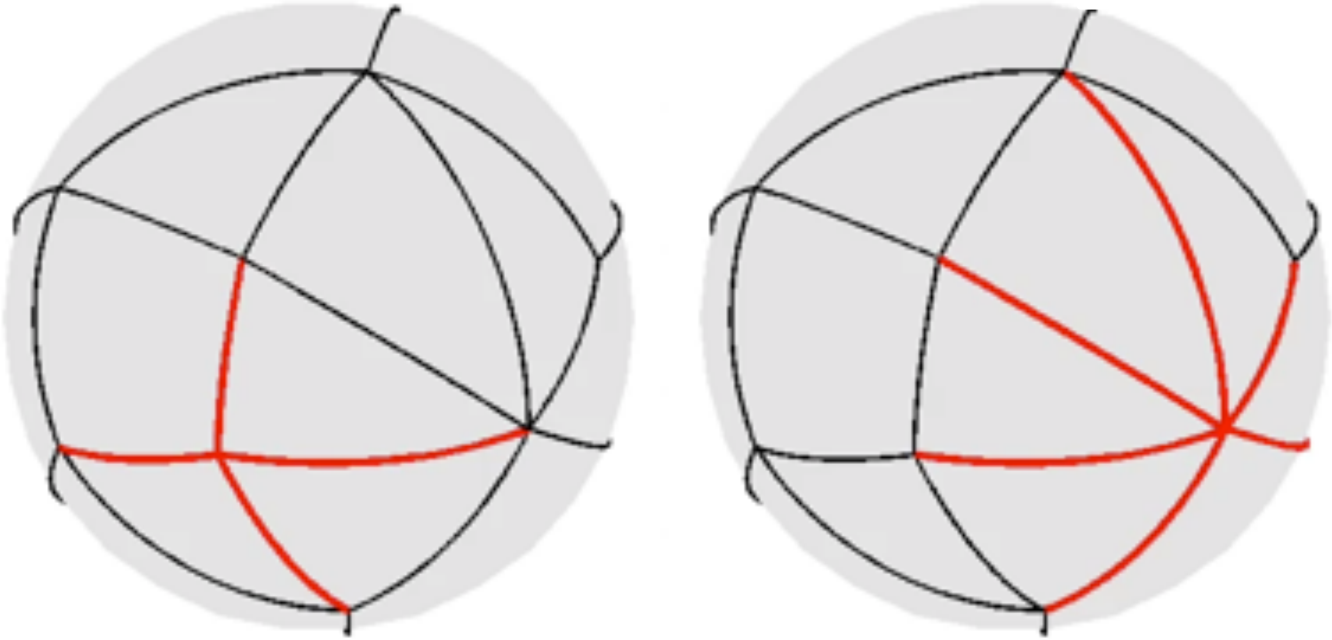

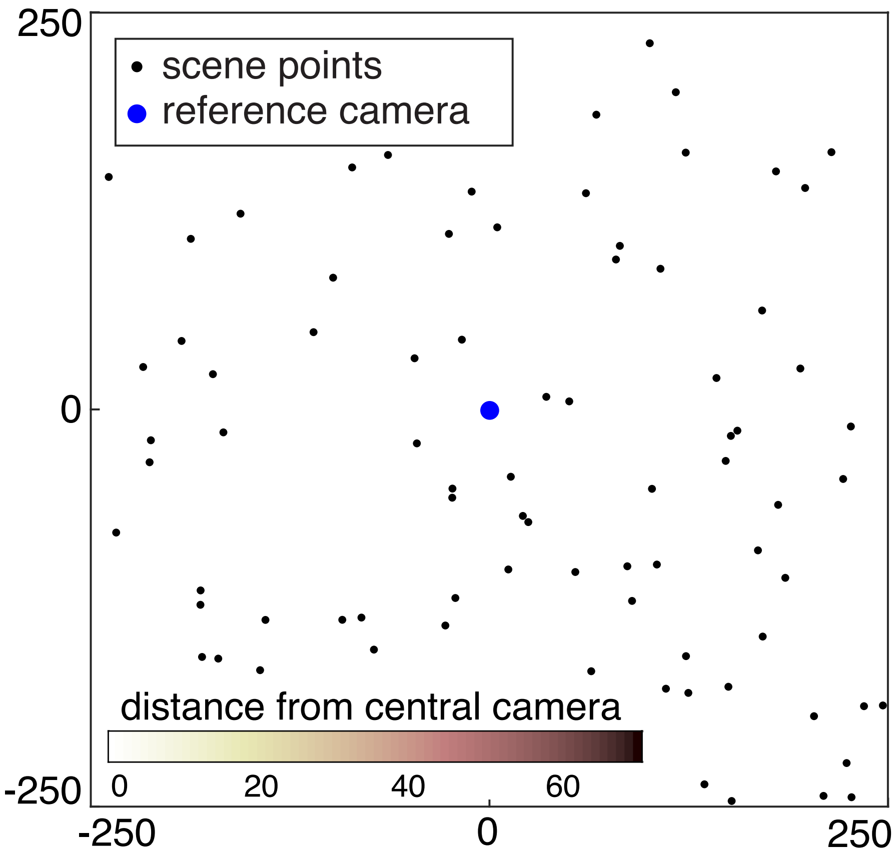

Glennerster, Hansard and Fitzgibbon (2001, 2009)

Relative visual direction – geometric consistency by design

University of Reading

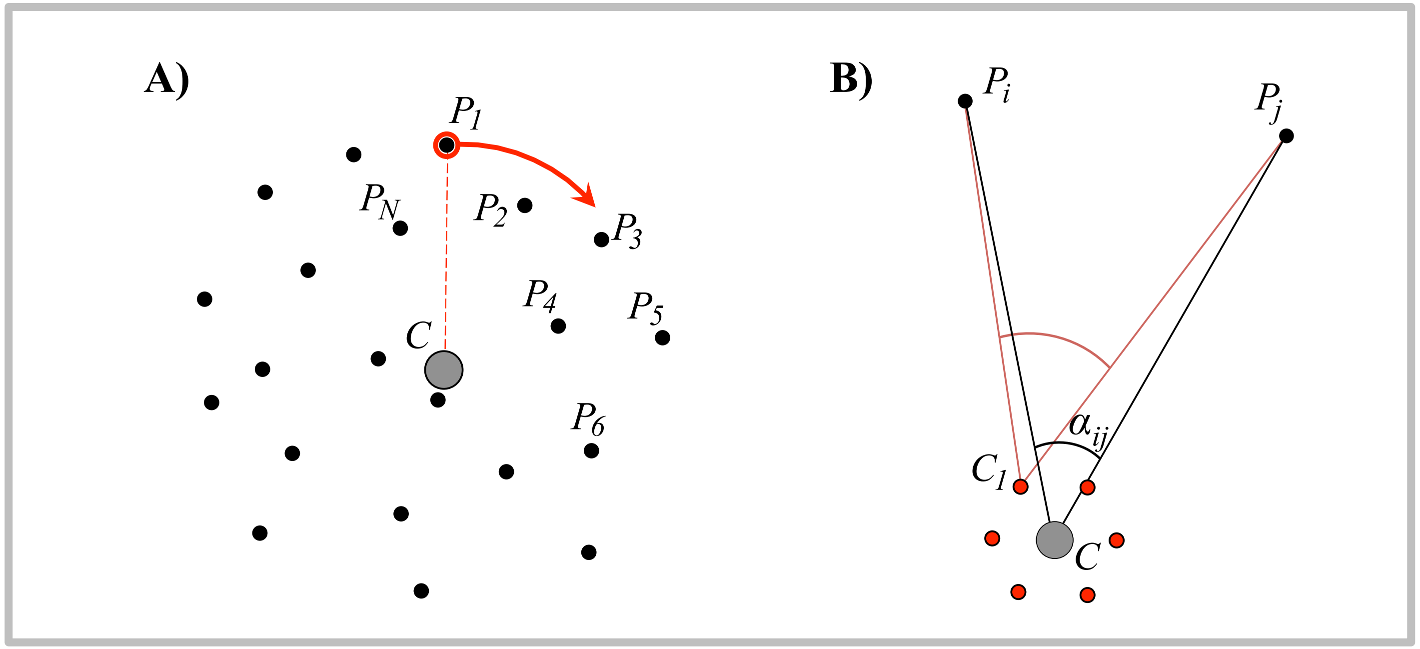

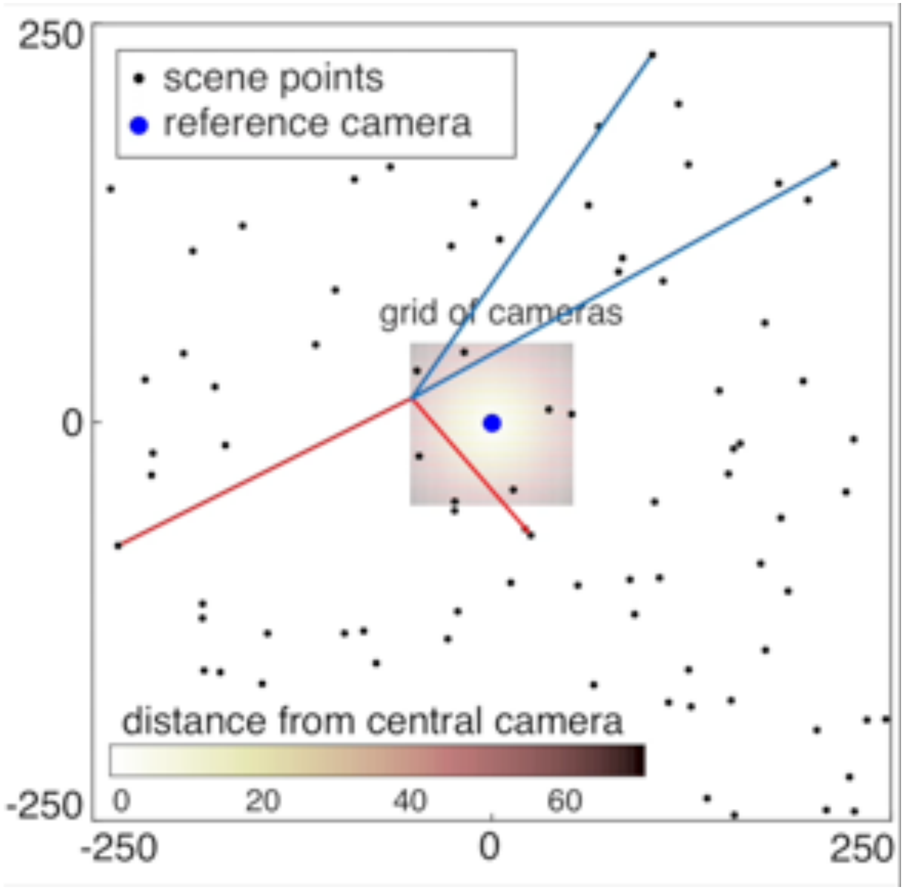

Relative visual direction – geometric consistency by design

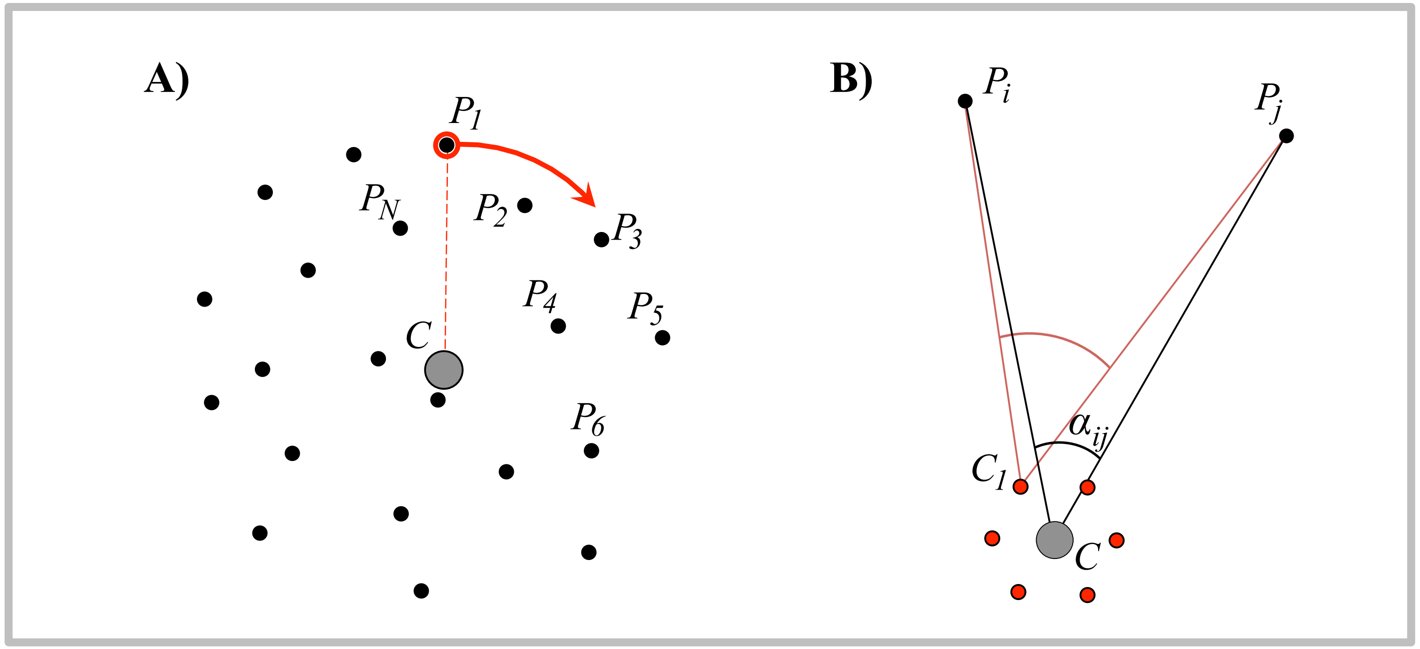

- generate a vector listing all the angles between each pair of *N* points as viewed from the camera, *C*

$$\vec{\varepsilon} = \{\alpha_{ij} : i \neq j, i = 1, \ldots, N, j = 1, \ldots, N\}.$$

$$M = N^2 - N, \; \vec{\varepsilon} \in \mathbb{R}^M$$

- calculate a parallax measure for each pair

$$\vec{\pi} = \frac{1}{n_C} \sum_{k=1}^{n_C} \frac{\vec{\varepsilon} - \vec{C}_k}{\vec{\varepsilon}} = \{\pi_1, \pi_2, \ldots, \pi_M\}$$

- find indices into $\vec{\varepsilon}$ using this parallax measure

$$\vec{\pi}' = \vec{\pi}_{\vec{\rho}} = \{\pi_k\}, \pi_k \in \vec{\pi}, \pi_k \leq T$$

$$\vec{\varepsilon}' = \vec{\varepsilon}_{\vec{\rho}}.$$

- e.g. find all pairs with low parallax and use this subset of elements to form a relatively stable feature vector (stable with translation of the camera)

- in general, use the most relevant features for the task at hand

Relative visual direction – geometric consistency by design

University of Reading

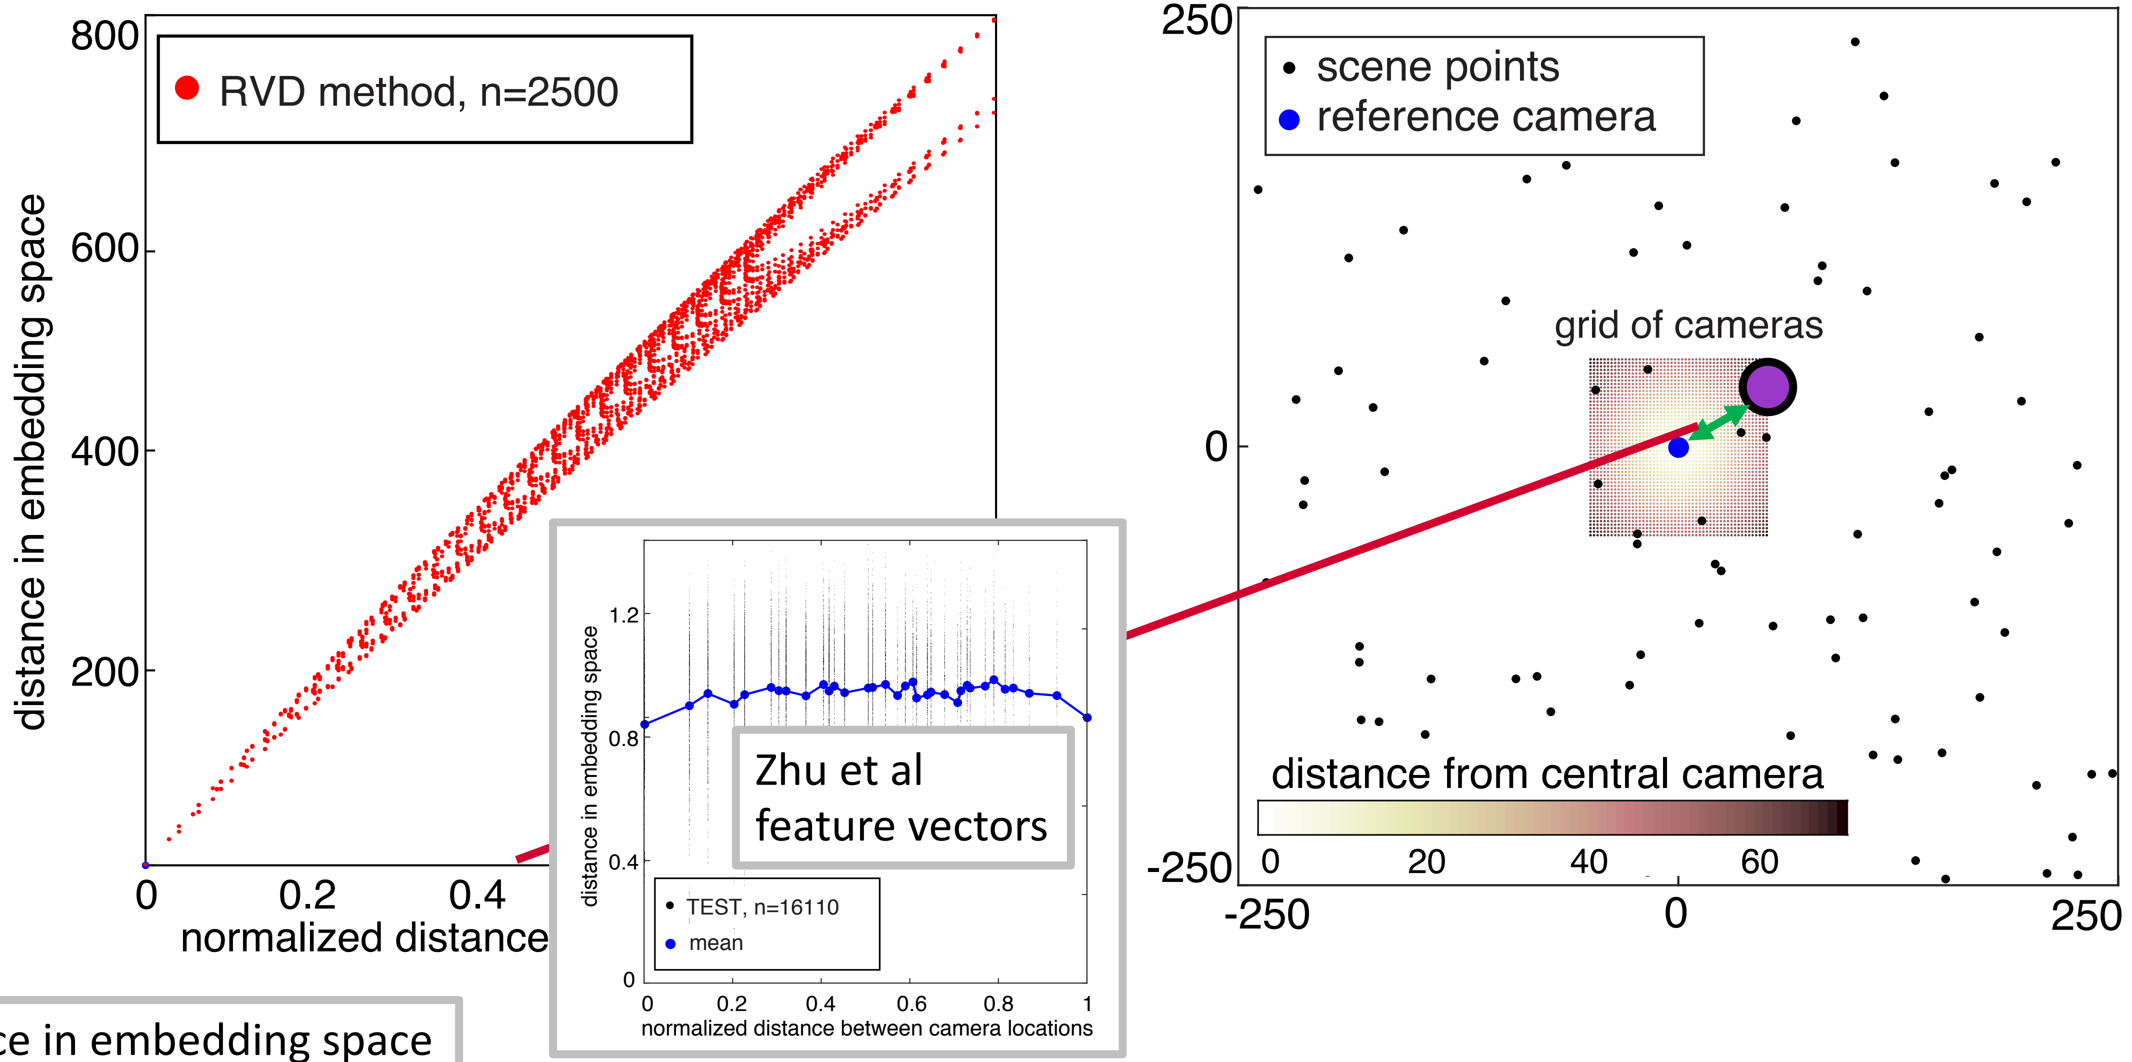

Distance in embedding space

Relative visual direction – geometric consistency by design

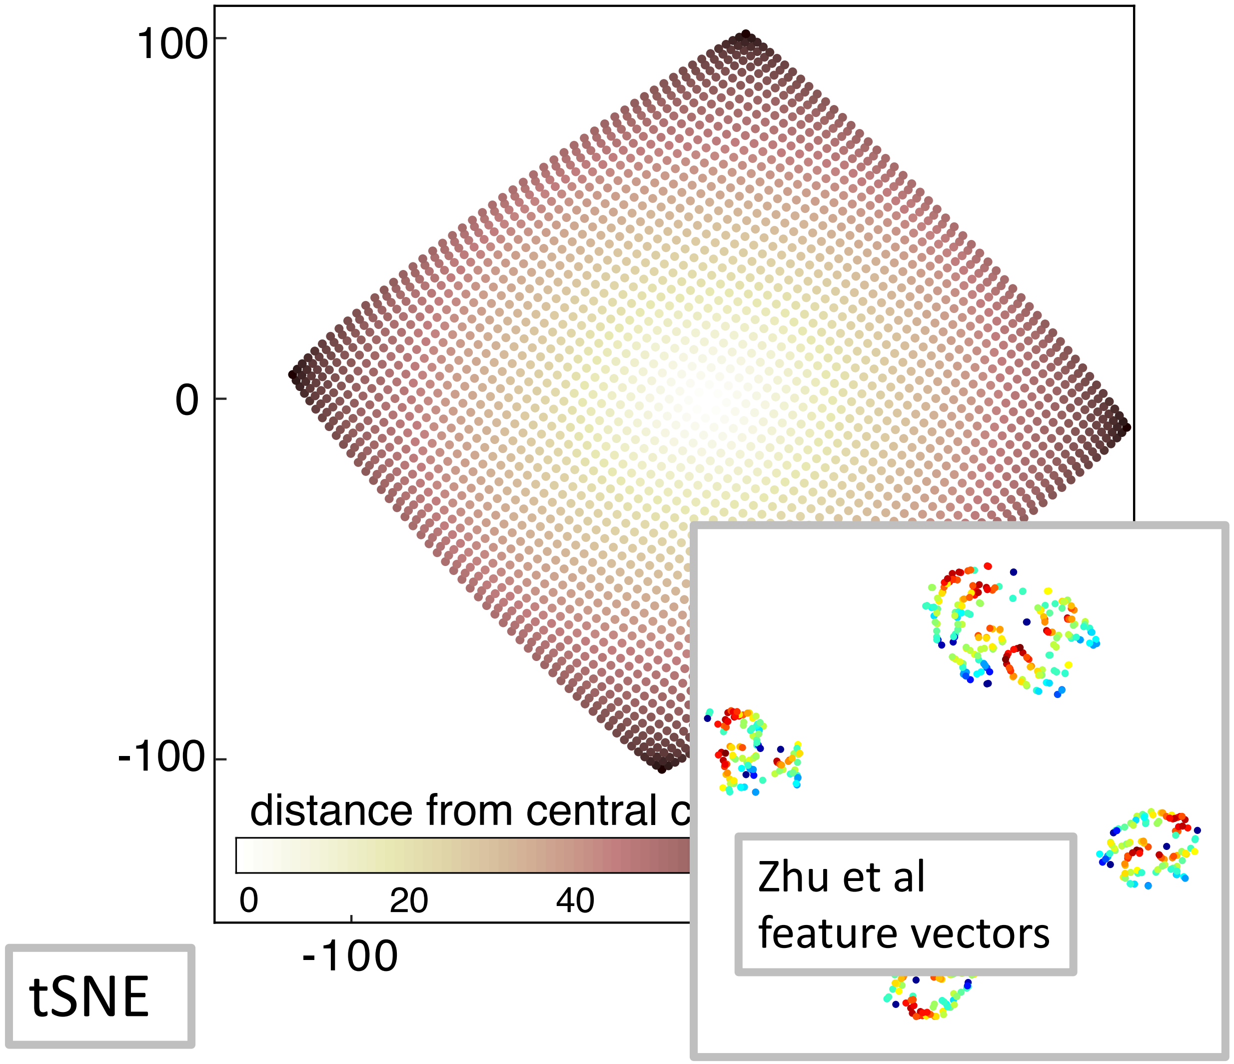

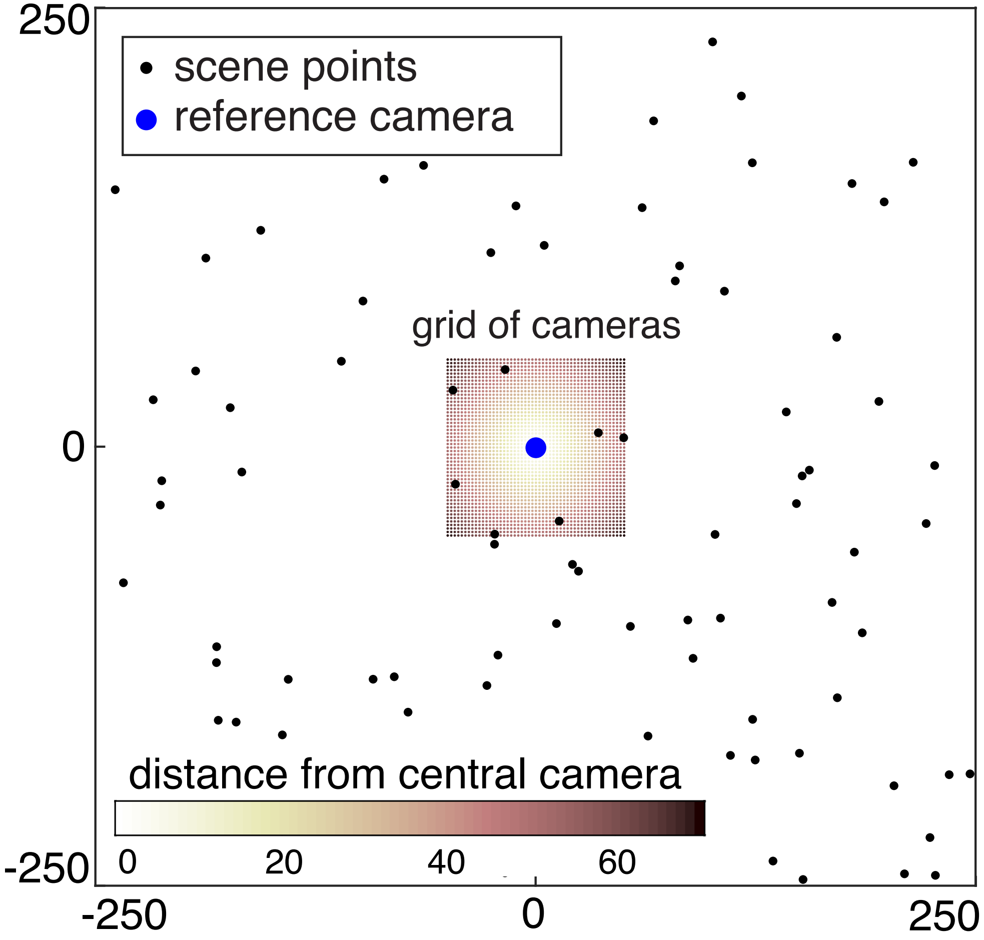

Relative visual direction – geometric consistency by design

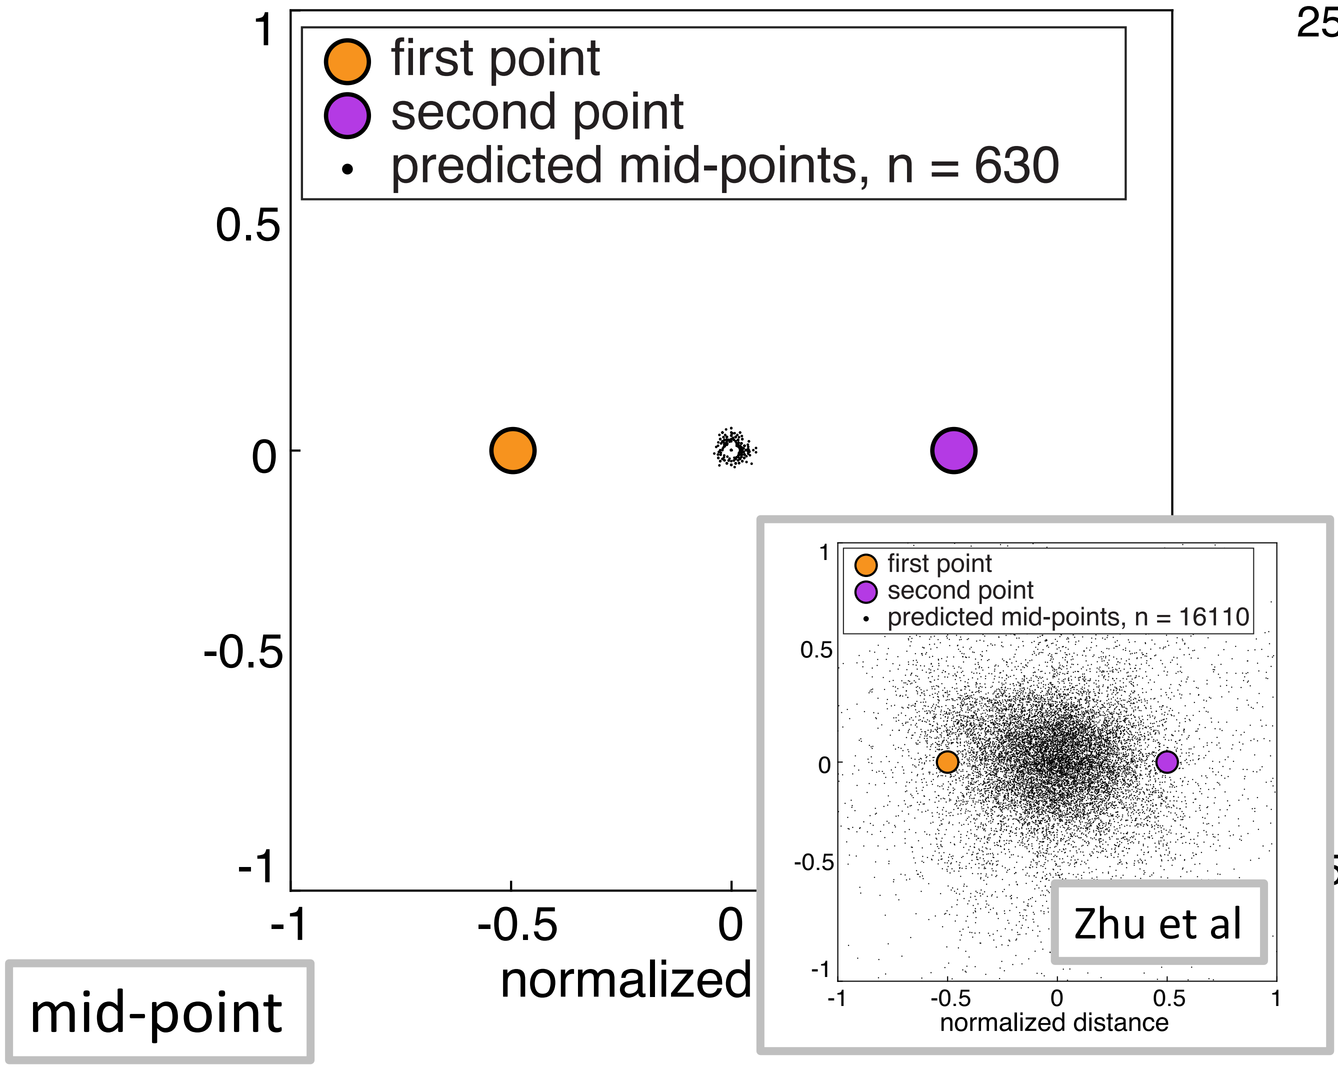

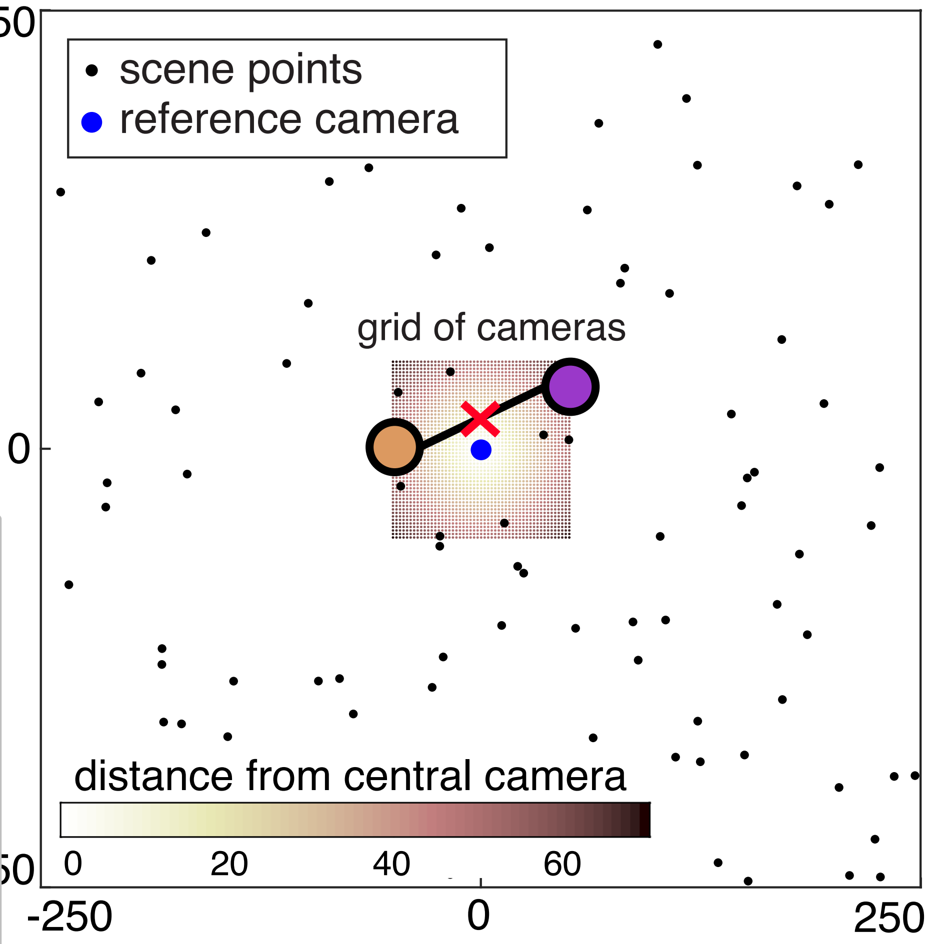

Zhu et al mid-point interpolation compared with RVD method

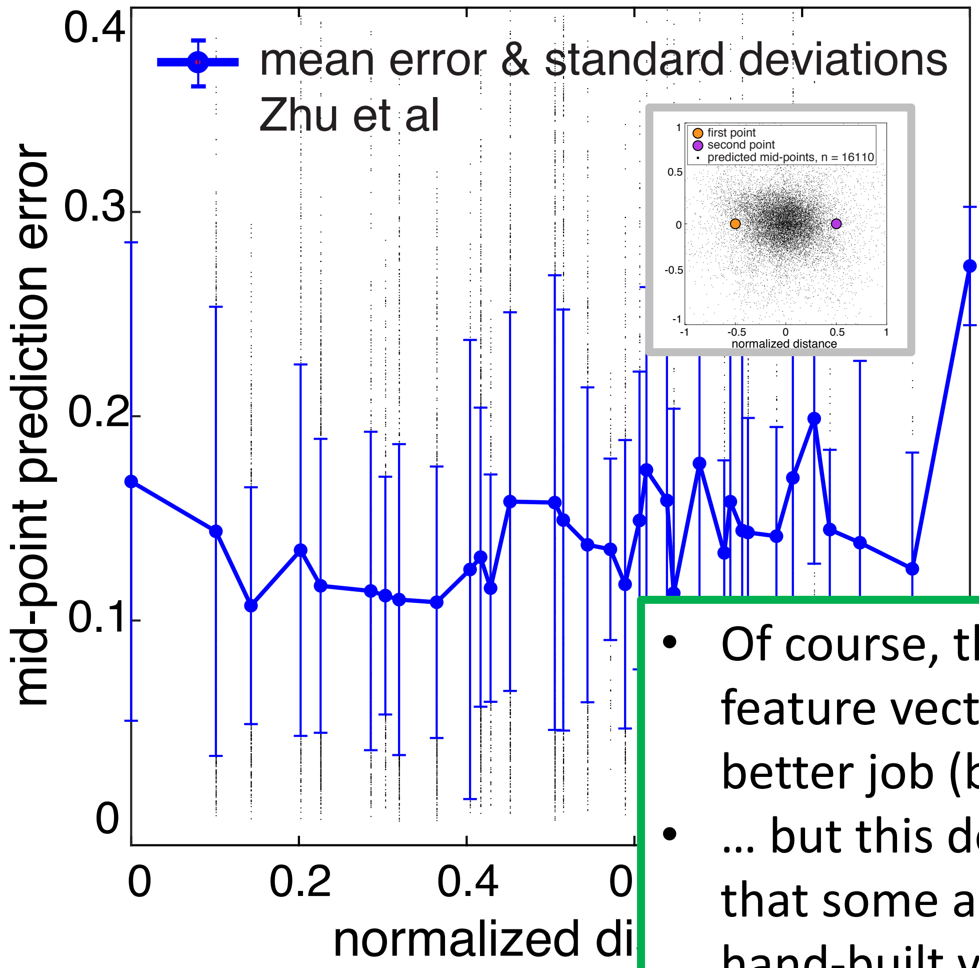

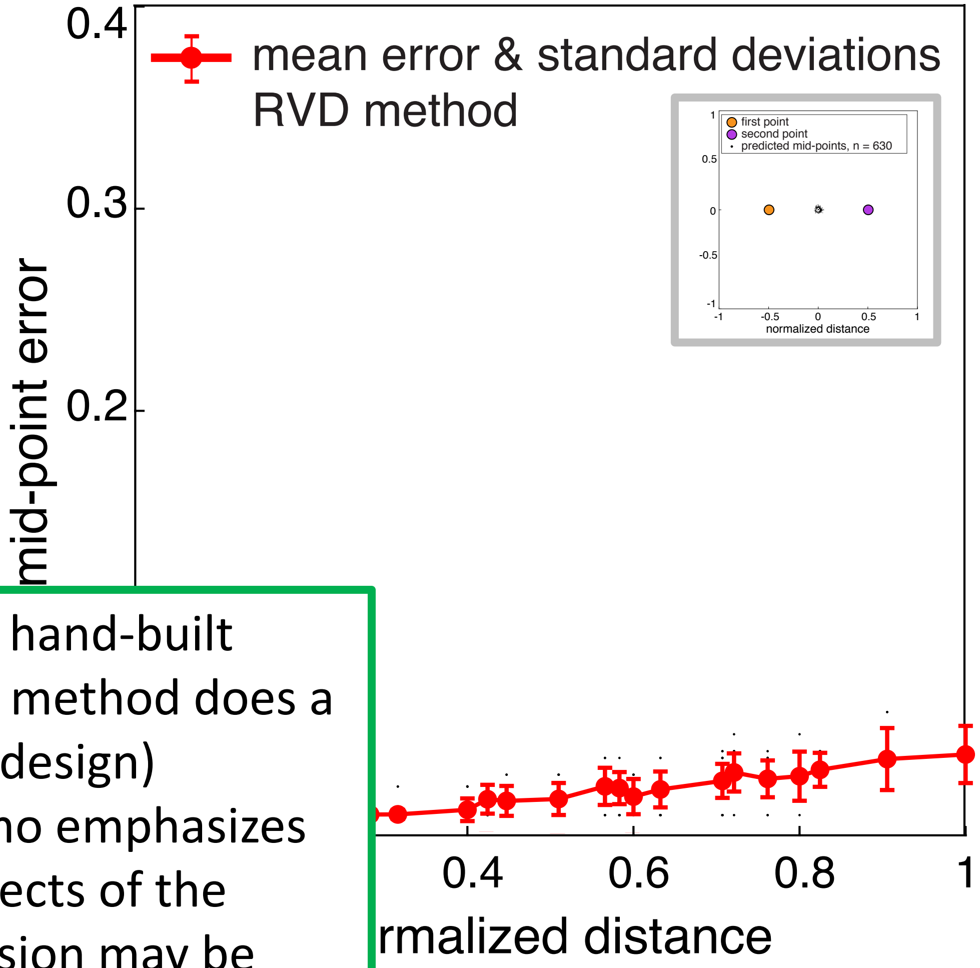

- Of course, the hand-built feature vector method does a better job (by design)
- … but this demo emphasizes that some aspects of the hand-built version may be useful to learn

Plans

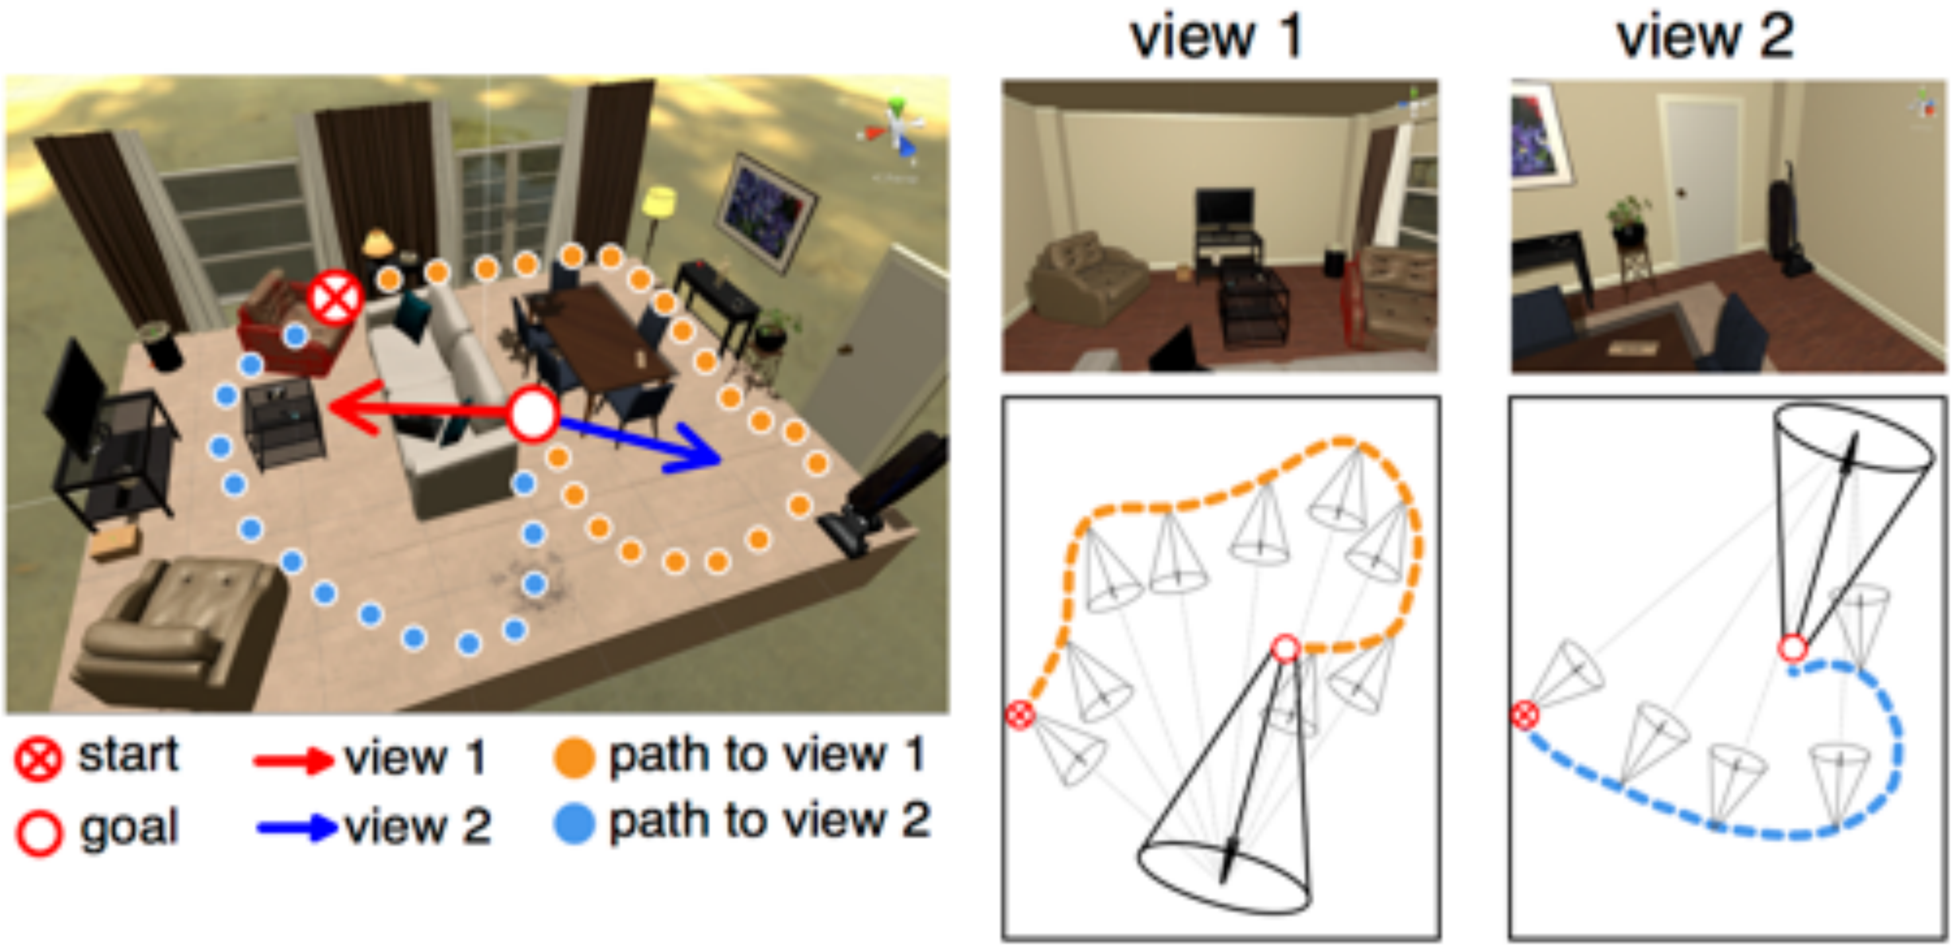

- Compare predictions of view-based and 3D reconstruction models for very different ‘learning views’

Plans

University of Reading

SCIENTIFIC REPORTS

natureresearch

No single, stable 3D representation can explain pointing biases in a spatial updating task

Jenny Vuong[1], Andrew W. Fitzgibbon[2] & Andrew Glennerster[1]

- Building on recent publication, explore predictions of neural net models of extrapolating to unseen targets, eg Eslami et al (2018).

Plans

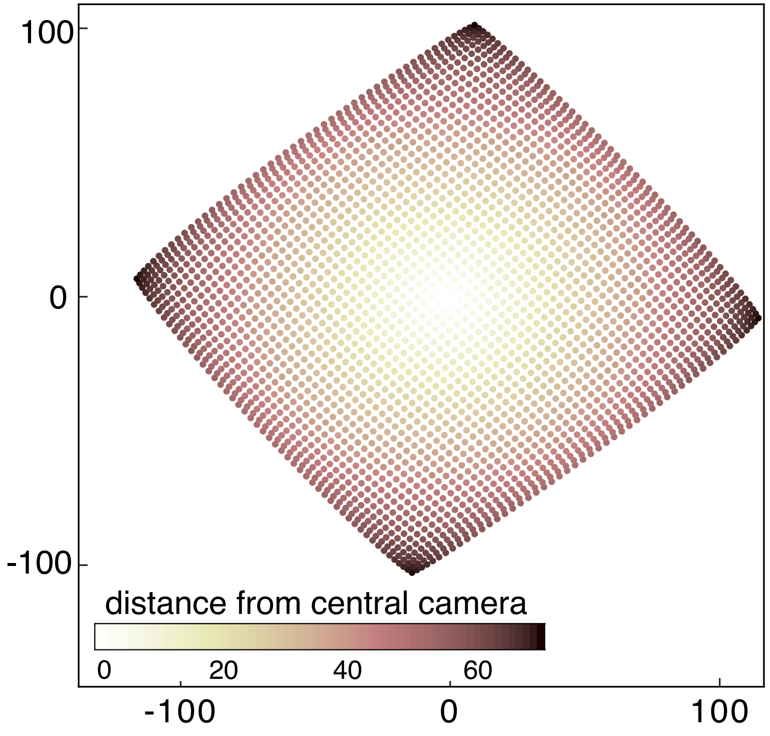

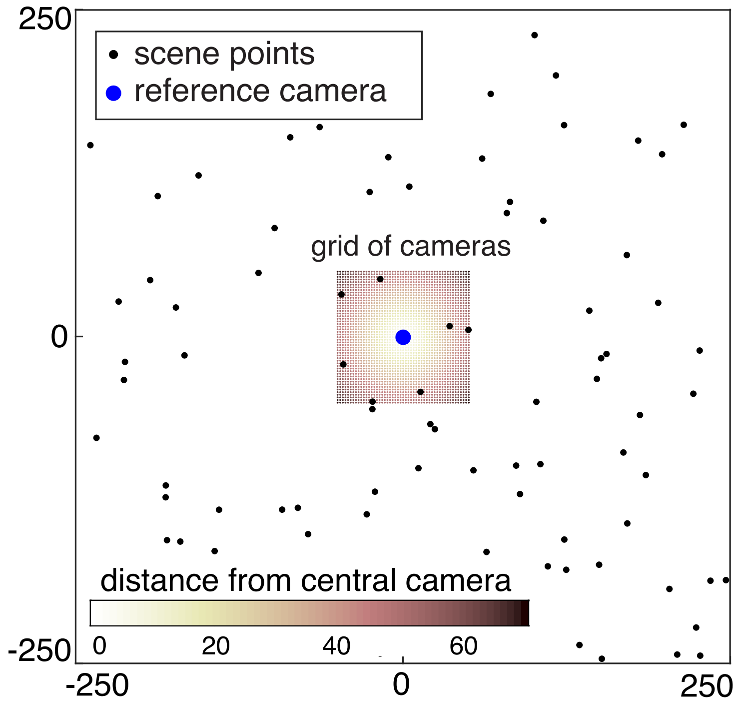

- Explore compositional representations of views as a camera translates and the implications for RL representations for navigation

Summary

- More collaboration would be good

- Π(a|s,g) for neuroscience

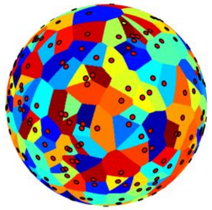

- RL/DNN use the wrong basis vectors to represent images (wrong in the DoD sense, i.e. not ‘human-like’).


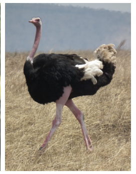

- Learn basis vectors gradually, in tandem with gradual development of task complexity

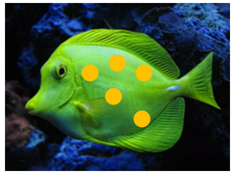

- RL for navigation could be improved by learning which features are likely to persist over longer translations of the camera

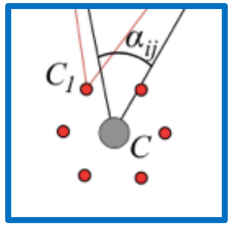

University of Reading

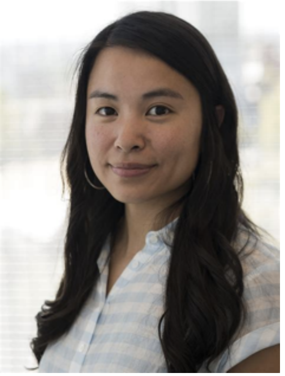

Jenny Vuong

Alex Muryy

Luise Gootjes-Dreesbach

Peter Scarfe

James Stazicker

Miles Hansard

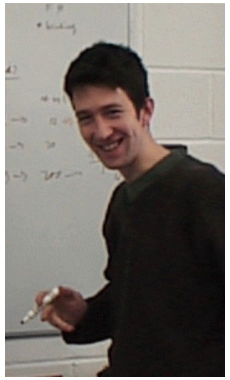

Andrew Fitzgibbon

EPSRC

[dstl]

Microsoft Research