Why you should care about sheaves

Simon Dobson

Complex and Adaptive Systems Research Group
School of Computer Science, University of St Andrews

mailto:simon.dobson@st-andrews.ac.uk
https://simondobson.org

University of
St Andrews
FOUNDED
1413

Introduction

We often find ourselves collecting different data at different places in a system

- Not necessarily the *same* data at *all* places
- Collate *local* data into a *non-local* view of the whole system

A recently-emerged approach helps us do this in a *canonical* and *well-founded* way

- And you don't need any hairy maths to understand it [1]
- May have applications to your research

[1] This claim may not be entirely accurate…

THE LOCAL-TO-GLOBAL TRANSITION

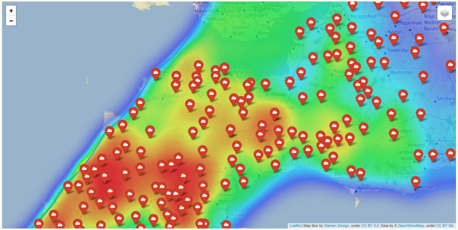

Data collected at discrete, known points [2]

- Fuse readings from related, typically nearby, points

⇒ Form a global view accounting for all the data

- There may be (local and global) structure we don't want to lose as we fuse the data

⇒ We need to control when and how fusion happens

[2] M. Robinson. Sheaves are the canonical data structure for sensor integration. *Information Fusion*, 36:208–224, 2016. doi: 10.1016/j.inffus.2016.12.002. URL `https://doi.org/10.1016/j.inffus.2016.12.002`

How to do information fusion

Suppose you find yourself tracking bears… [3]

- Several sensors
- Need to integrate the information provided by them
- The information provided will be noisy, so we need to quantify the quality of the resulting conclusions

[3] C. A. Joslyn, L. Charles, C. DePerno, N. Gould, K. Nowak, B. Praggastis, E. Purvine, M. Robinson, J. Strules, and P. Whitney. A sheaf theoretical approach to uncertainty quantification of heterogeneous geolocation information. *Sensors*, 20(12):3418, 2020. doi: 10.3390/s20123418. URL `http://dx.doi.org/10.3390/s20123418`

How to do information fusion

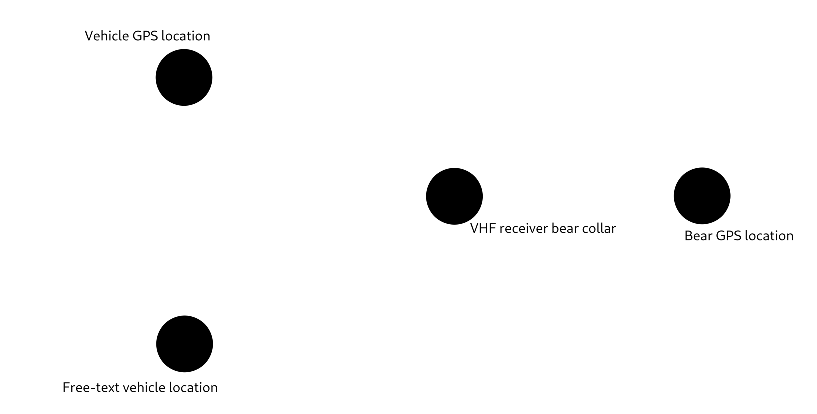

How to do information fusion

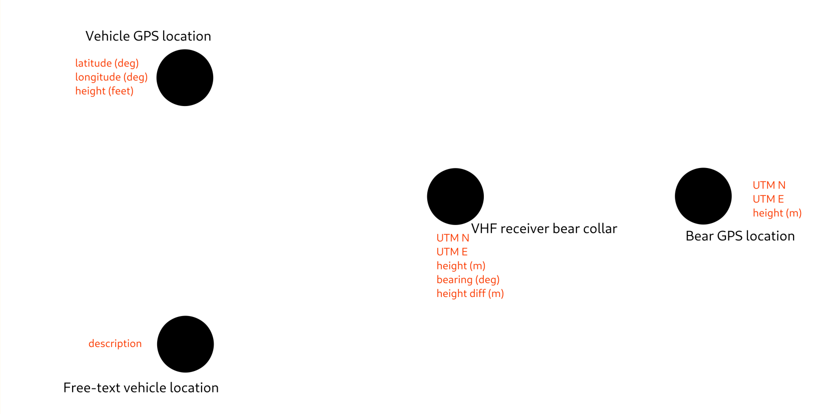

How to do information fusion

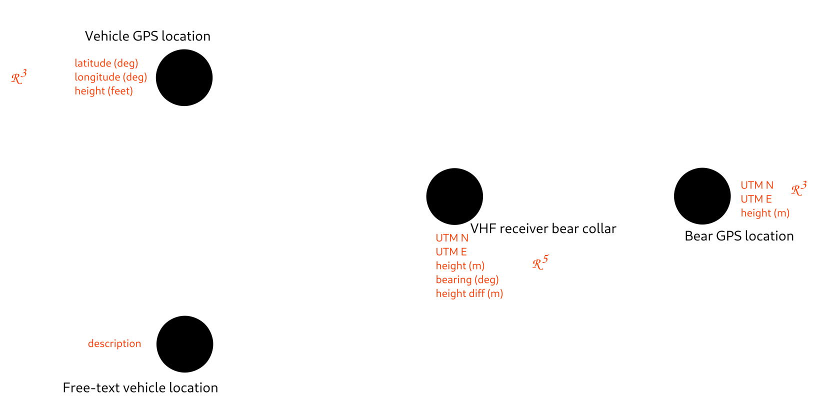

How to do information fusion

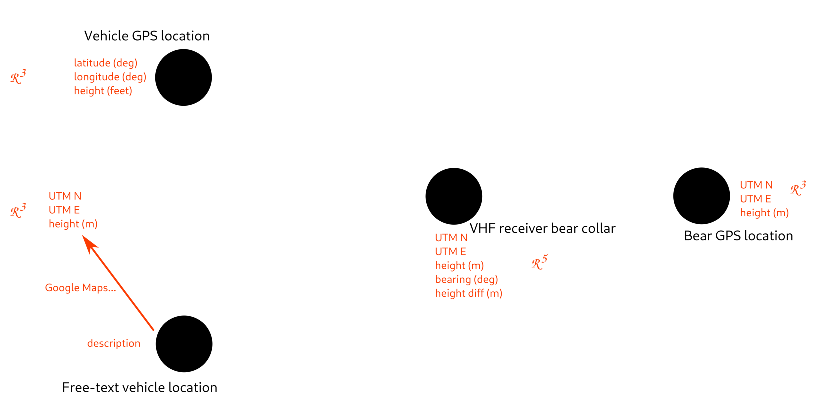

How to do information fusion

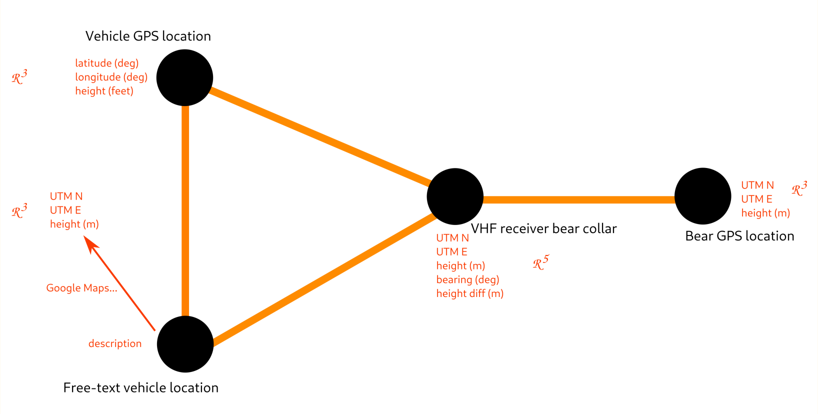

How to do information fusion

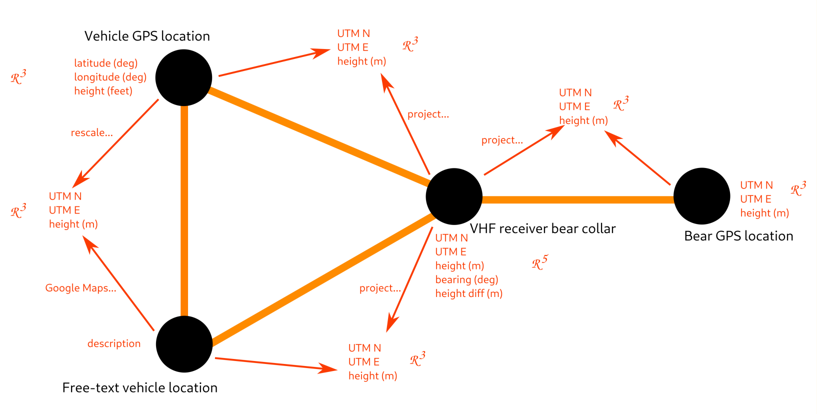

How to do information fusion

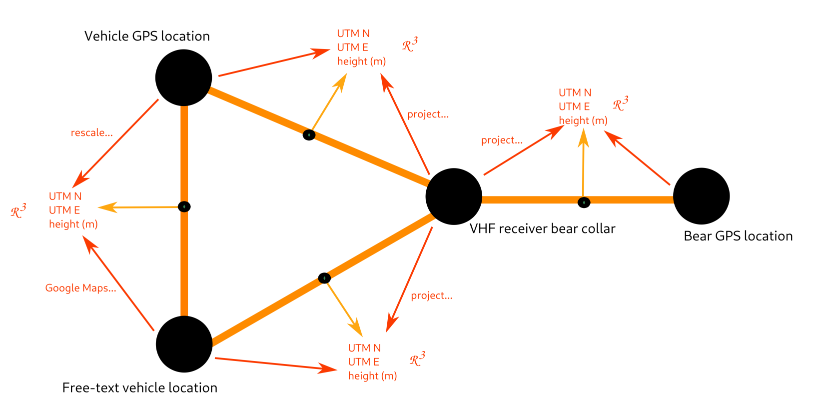

How to do information fusion

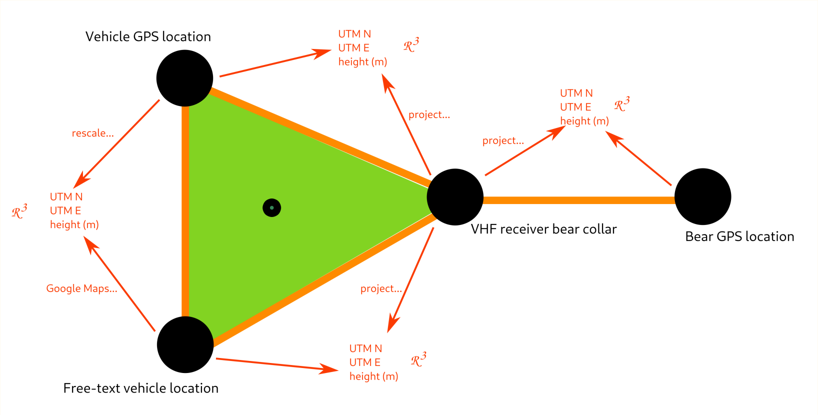

How to do information fusion

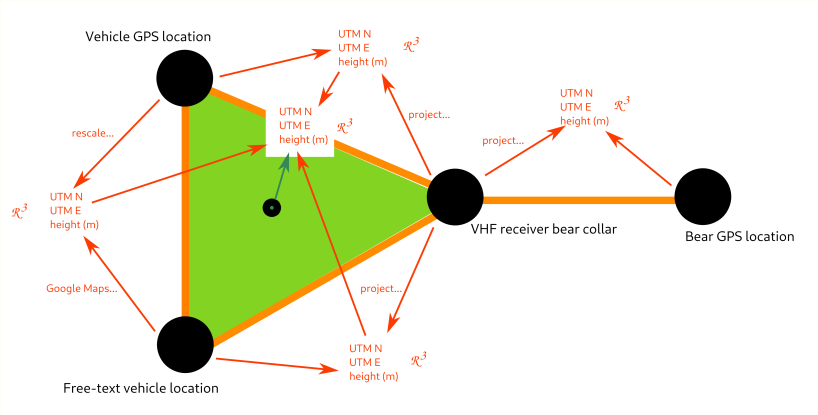

Measuring accuracy

Sensors will disagree on values observed and inject *different* values from (for example) node to edge

Measuring accuracy

Sensors will disagree on values observed and inject *different* values from (for example) node to edge

- Represent the data as a vector space

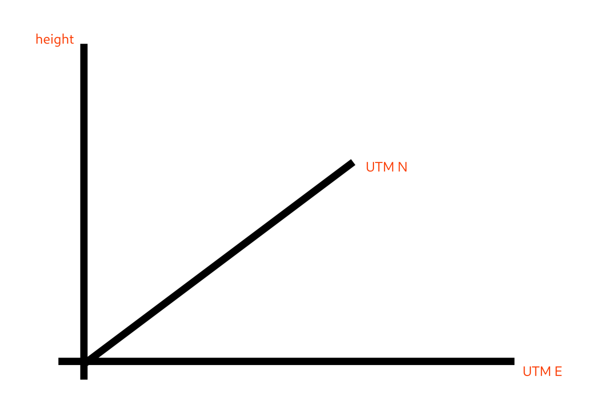

Measuring accuracy

Sensors will disagree on values observed and inject *different* values from (for example) node to edge

- Data values then form vectors, with a well-defined notion of distance between them

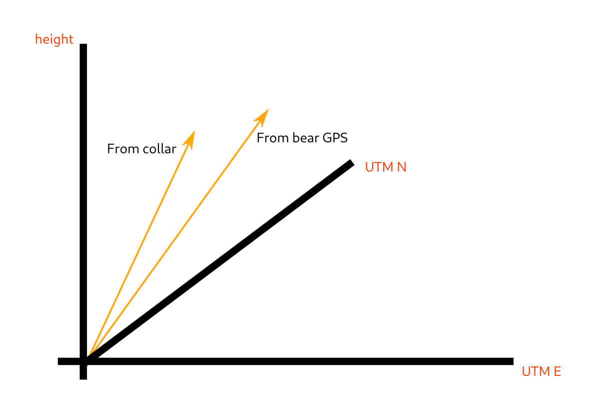

Measuring accuracy

Sensors will disagree on values observed and inject *different* values from (for example) node to edge

- A circle of minium radius that encompasses the vectors defines both the interpolated value *and* measures its accuracy

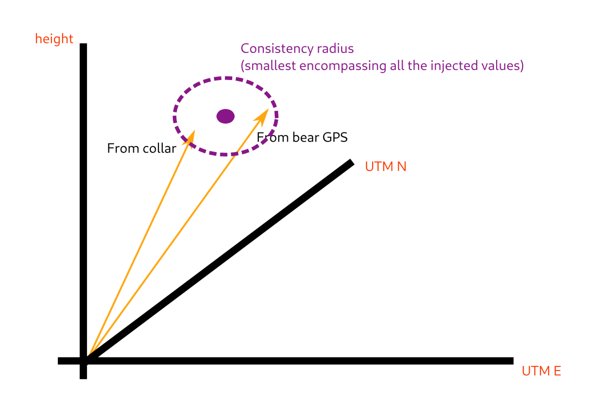

Sheaves

These structures are *sheaves* [3]

- A base space
- Each element has an attached *stalk* representing all the possible values that can be associated to the element
- Sets, posets, vector spaces, …
- *Restriction maps* controlling how data at nodes is mapped to data at edges, and edges to triangles, …
- A *section* is a choice of data at each stalk that respects all the restriction maps
- A *consistency structure* to measure agreement

[3] D. Rosiak. *Sheaf Theory through Examples*. The MIT Press, 2022. doi: 10.7551/mitpress/12581.001.0001. URL http://dx.doi.org/10.7551/mitpress/12581.001.0001

How is this different to "my favourite approach X"?

A sheaf is a very information-rich object

- Data types, mappings, consistency constraints

The sheaf model is *canonical* in a particular sense

- …and therefore subsumes many other models
- Joslyn *et alia* compare it to Bayesian Dynamic Linear Model with Kalman filtering, quantifies the differences in predictions, and explores what drives them
- DLM works on sensor self-consistency and is model-free
- Sheaf takes a model and measures divergence from it
- It's probably possible to "topologise" time to include time series in the sheaf

Making it computational

Another way to view sheaf construction [4]

- A function `restrict` to propose values for faces of a simplex
- A function `reduce` to turn the proposed values into a value for the simplex itself

```
from simplicial import SimplicialComplex, Simplex

class Sheaf[A]:

    def restrict(self, c: SimplicialComplex,
                 face: Simplex, s: Simplex) -> A:
        '''Propose a value for a face of the given simplex.'''

    def reduce(self, c: SimplicialComplex,
               vs: List[A], s: Simplex) -> A:
        '''Reduce proposed values from faces to a value for the simplex.'''
```

[4] S. Dobson. simplicial: Simplicial topology in Python. Python library, 2022. URL https://github.com/simoninireland/simplicial

Current work

Acoustic sensing to estimate bird abundance (with Alison Johnston of CREEM)

- Microphone array to hear birdsong, coupled to an "overhearing" sheaf
- Compare to extended distance sampling techniques
- Currently working on this in simulation before doing it for real

Dealing with reality (with Berné Nortier)

- Can we quantify/reduce the effects of noise?
- How do sensor placements affect sections?
- Does it scale-out to large deployments?

Conclusions

Sheaves are (arguably) the correct mathematical framework for information fusion

> *We must bear in mind, then, that there is nothing more difficult or dangerous or more doubtful of success than an attempt to introduce a new order of things.*
>
> Niccoló Machiavelli, The Prince (1515)

Fusion/inference into a single object

- Measured data, processing, measurement of accuracy
- Discrete version of the theory is easy to tool-up
- A very powerful theory to draw upon

References

S. Dobson. simplicial: Simplicial topology in Python. Python library, 2022. URL `https://github.com/simoninireland/simplicial`.

C. A. Joslyn, L. Charles, C. DePerno, N. Gould, K. Nowak, B. Praggastis, E. Purvine, M. Robinson, J. Strules, and P. Whitney. A sheaf theoretical approach to uncertainty quantification of heterogeneous geolocation information. *Sensors*, 20(12):3418, 2020. doi: 10.3390/s20123418. URL `http://dx.doi.org/10.3390/s20123418`.

M. Robinson. Sheaves are the canonical data structure for sensor integration. *Information Fusion*, 36:208–224, 2016. doi: 10.1016/j.inffus.2016.12.002. URL `https://doi.org/10.1016/j.inffus.2016.12.002`.

D. Rosiak. *Sheaf Theory through Examples*. The MIT Press, 2022. doi: 10.7551/mitpress/12581.001.0001. URL `http://dx.doi.org/10.7551/mitpress/12581.001.0001`.