

Article

A Sheaf Theoretical Approach to Uncertainty Quantification of Heterogeneous Geolocation Information

Cliff A. Joslyn ¹ *, Lauren Charles ², Chris DePerno ³, Nicholas Gould ³, Kathleen Nowak ², Brenda Praggastis ¹, Emilie Purvine ¹, Michael Robinson ⁴, Jennifer Strules ³ and Paul Whitney ²

¹ Pacific Northwest National Laboratory, Seattle, WA 98109, USA; Brenda.Praggastis@pnnl.gov (B.P.); Emilie.Purvine@pnnl.gov (E.P.)

² Pacific Northwest National Laboratory, Richland, WA 99352, USA; lauren.charles@pnnl.gov (L.C.); kathleen.nowak@cbp.dhs.gov (K.N.); pdwhitney@gmail.com (P.W.)

³ College of Natural Resources, Department of Forestry and Environmental Resources, Fisheries, Wildlife, and Conservation Biology, North Carolina State University, Raleigh, NC 27695, USA; chrisdeperno@ncsu.edu (C.D.); npgould@ncsu.edu (N.G.); urbanbearstudy@ncsu.edu (J.S.)

⁴ Department of Mathematics and Statistics, American University, Washington, DC 20016, USA; michaelr@american.edu

* Correspondence: cliff.joslyn@pnnl.gov

Received: 1 April 2020; Accepted: 8 June 2020; Published: 17 June 2020



Abstract: Integration of multiple, heterogeneous sensors is a challenging problem across a range of applications. Prominent among these are multi-target tracking, where one must combine observations from different sensor types in a meaningful and efficient way to track multiple targets. Because different sensors have differing error models, we seek a theoretically justified quantification of the agreement among ensembles of sensors, both overall for a sensor collection, and also at a fine-grained level specifying pairwise and multi-way interactions among sensors. We demonstrate that the theory of mathematical sheaves provides a unified answer to this need, supporting both quantitative and qualitative data. Furthermore, the theory provides algorithms to globalize data across the network of deployed sensors, and to diagnose issues when the data do not globalize cleanly. We demonstrate and illustrate the utility of sheaf-based tracking models based on experimental data of a wild population of black bears in Asheville, North Carolina. A measurement model involving four sensors deployed among the bears and the team of scientists charged with tracking their location is deployed. This provides a sheaf-based integration model which is small enough to fully interpret, but of sufficient complexity to demonstrate the sheaf's ability to recover a holistic picture of the locations and behaviors of both individual bears and the bear-human tracking system. A statistical approach was developed in parallel for comparison, a dynamic linear model which was estimated using a Kalman filter. This approach also recovered bear and human locations and sensor accuracies. When the observations are normalized into a common coordinate system, the structure of the dynamic linear observation model recapitulates the structure of the sheaf model, demonstrating the canonicity of the sheaf-based approach. However, when the observations are not so normalized, the sheaf model still remains valid.

Keywords: topological sheaves; information integration; consistency radius; wildlife management; stochastic linear model; Kalman filter

1. Introduction

There is a growing need for a representational framework with the ability to assimilate heterogeneous data. Partial information about a single event may be delivered in many forms.

For instance, the location of an animal in its environment may be described by its GPS coordinates, citizen-scientist camera pictures, audio streams triggered by proximity sensors, and even Twitter text. To recover a holistic picture of an event, one must combine these pieces in a meaningful and efficient way. Because different observations have differing error models, and may even be incompatible, a model capable of integrating heterogeneous data forms has become critically necessary to satisfy sensor resource constraints while providing adequate data to researchers.

Fundamentally, we seek a theoretically justified, detailed quantification of the agreement among ensembles of sensors. From this, one can identify collections of self-consistent sensors, and determine when or where certain combinations of sensors are likely to be in agreement. This quantification should support both quantitative data and qualitative data, and continue to be useful regardless of whether the sampling rate is adequate for a full reconstruction of the scene.

Working in conjunction with recent advances in signal processing [1,2], this article shows that mathematical *sheaves* support an effective processing framework with canonical analytic methods. We ground our methodological discussion with a field experiment, in which we quantified uncertainty for a network of sensors tracking black bears in Asheville, North Carolina. The experiment provided data exhibiting a variety of error models and data types, both quantitative (GPS fixes and radio direction bearings) and qualitative (text records).

With this experiment, we demonstrate that although the sheaf methodology requires careful modeling as a prerequisite, one obtains fine-grained analytics that are automatically tailored to the specific sensor deployment and also to the set of observations.

Because there is a minimum of hidden state to be estimated, relationships between sensors — which collections of sensors yield consistent observations — can be found by inspecting the *consistency filtration*, which is naturally determined by the model. These benefits are largely a consequence of the fact that sheaves are the *canonical* structure of their kind, so that any principled specification of the interaction between heterogeneous data sources will provably recapitulate some portion of sheaf theory [2]. In short, the theory provides an algorithmic way to globalize data, and to diagnose issues when the data do not globalize cleanly.

This paper is organized as follows.

- Section 2 provides context around the history of target tracking methodologies and multi-sensor fusion, including the particular role we are advocating specifically for sheaf methods.
- Then in Section 3 we describe the experimental setup for our data collection effort regarding tracking black bears in Asheville, North Carolina. This involved four sensors deployed among the bears, several “dummy” bear collars deployed at fixed locations in the field, and team of scientists charged with tracking their location. This experiment was designed specifically to identify a real tracking task where multi-sensor integration is required, but where the system complexity is not so large as to overwhelm the methodological development and demonstration, while still being sufficiently complexity to demonstrate the value of sheaf-based methods.
- Section 4 continues with a detailed mathematical development of the sheaf-based tracking model. This includes several features novel to information fusion models, including the explicit initial attention to the complexity of sensor interaction in the core sheaf models, but then the ability to equip these models with uncertainty tolerance in the form of *approximate sections*, and then to measure that both globally via a *consistency radius* and also at a more fine-grained level in terms of specific sensor dependencies in terms of a *consistency filtration*.
- Section 5 shows several aspects of our models interpreted through the measured data. First, overall results in terms of bear and dummy collars is provided; then results for a particular bear collar are examined in detail; and finally results for a particular time point in the measured data set is shown to illustrate the consistency filtration in particular.
- While our sheaf models are novel, we sought to compare them to a more traditional modeling approach, introduced in Section 6. Specifically, a statistical approach was developed to process the data once registered into a common coordinate system, a dynamic linear model which was

estimated using a Kalman filter. This approach also recovered bear and human locations and sensor accuracies. Comparisons in both form and results are obtained, which demonstrate the role that sheaves as *generic* integration models can play in conjunction with *specific* modeling approaches such as these: as noted, all integration models will provably recapitulate some portion of sheaf theory [2], even if they are not first registered into a common coordinate system.

- We conclude in Section 7 with some general observations and discussion.

2. History and Context

Situations such as wildlife tracking with a variety of sensors requires the coordinated solution of several well-studied problems, most prominently that of *target tracking* and *sensor fusion*. Although the vast majority of the solutions to these problems that are proposed in the literature are statistical in nature, the sample rates available from our sensors are generally insufficient to guarantee good performance. It is for this reason that we pursued more foundational methods grounded in the geometry and topology of *sheaves*. These methods in turn have roots in *topology* and *category theory*.

2.1. Target Tracking Methods

Tracking algorithms have a long and storied history. Several authors (most notably [3,4]) provide exhaustive taxonomies of tracker algorithms which include both probabilistic and deterministic trackers. Probabilistic methods are typically based on optimal Bayesian updates [5] or their many variations, for instance [6–8].

Closer to our approach, trackers based on optimal network flow tend to be extremely effective in video object tracking [9–11], but require high sample rates. A tracker based on network flows models targets as “flowing” along a network of detections. The network is given weights to account for either the number of targets along a given edge, or the likelihood that some target traversed that edge. (Our approach starts with a similar network of detections, but treats the weighting very differently.) Network flow trackers are optimal when target behaviors are probabilistic [12,13], but no sampling rate bounds appear to be available. There are several variations of the basic optimal network flow algorithm, such as providing local updates to the flow once solved [14], using short, connected sequences of detections as “super detections” [15], or K -shortest paths, sparsely enriched with target feature information [16].

2.2. Multi-Sensor Fusion Methods

Data fusion is the task of forming an “alliance of data originating from different sources” [17]. There are several good surveys discussing the interface between target tracking and data fusion — such as [3,4,18] — and which cover both probabilistic and deterministic methods. Various systematic experimental campaigns have also been described [19,20]. Other authors have shown that fusing detections across sensors [21–24] yields better coverage and performance.

Our idea of using a heterogenous collection of sensors, some of which are manually operated, is similar to the sensor deployment strategy discussed in [25], which describes an approach to monitoring vehicular traffic on a road network using temporary, portable sensors. Although their underlying model is linear and satisfies a conservation law, ours need not be linear. While our sheaf-based approach to uncertainty quantification is quite different from their statistical model, we compare our results with a more standard statistical model in Section 6.

Considerable effort is typically expended developing robust features, although usually the metric for selecting features is pragmatic rather than theoretical. In all cases, though, the assumption is that sample rates are sufficiently high. For situations such as our wildlife tracking problem, this basic requirement is rarely met. Furthermore, data fusion techniques that operate on quantitative data typically require that sensors be of the same type [26–28]. Some prerequisite spatial registration to a common coordinate system is generally required, especially if sensor types differ [29–31].

The lack of a common coordinate system can make all of these approaches rather brittle to changes in sensor deployment. One mitigation for this issue is to turn to a more foundational method, for example so-called “possibilistic” information theory [32–34]. Here one encodes sensor models as a set of propositions and rules of inference. Data then determine the value of logical variables, from which inferences about the scene can be drawn. Allowing the variables to be valued possibilistically — as opposed to probabilistically — supports a wider range of uncertainty models. Although these methods usually can support heterogeneous collections of sensors, they do so without the theoretical guarantees that one might expect from homogeneous collections of sensors. Although the workflow for possibilistic techniques is similar to ours, the key difference is that ours relies on *geometry*. Without the geometric structure, one is faced with combinatorial complexity that scales rapidly with the number of possible sensor outputs; this frustrates the direct application of logical techniques.

Recent probabilistic models also provide a point of comparison to our sheaf methods. In [35], a multi-sensor integration system is structured within a probabilistic information fusion system. While the resulting system is specific to probabilistic robotic sensors, it may be likely that casting this in the context of a sheaf-theoretical context could provide significant value for generalization. Alternatively, in [36] the authors consider the problem of probabilistic target identification across multiple heterogeneous sensors, each measuring potentially different characteristics of the targets being tracked. This paper takes an evidence theory approach combining data with different levels of confidence based on foundations of Dempster-Shafer theory together with a Rayleigh distribution. While these authors consider the specific question of identifying unique targets across sensors our framework provides the ability to track various aspects of targets and the agreement (or disagreement) across sensors. It could be interesting to consider a hybrid approach leveraging their confidence aggregation methods within our sheaf-based method.

2.3. Sheaf Geometry for Fusion

The typical target tracking and data fusion methods tend to defer modeling until after the observations are present. If one were to reverse this workflow, requiring careful modeling before any observations are considered, then looser sampling requirements arise. This is supported by the workflow we propose, which uses interlocking local models of consistency among the observations. These interlocking local models are canonically and conveniently formalized by *sheaves*. As we discuss in Section 4, a *sheaf* is a precise specification of which sets of local data can be fused into a consistent, more global, datum [37]. While the particular kind of sheaves we exploit in this article have been discussed sporadically in the mathematics literature [38–43], our recent work cements sheaf theory to practical application.

Sheaves are an effective organization tool for heterogeneous sensor deployments [44]. Since sheaves require modeling as a *prerequisite* before any analysis occurs, encoding sensor deployments as a sheaf can be an obstacle. Many sheaf encodings of standard models (such as dynamical systems, differential equations, and Bayesian networks) have been catalogued [45]. Furthermore, these techniques can be easily applied in several different settings, for instance in air traffic control [46,47] and in formal semantic techniques [48]. Sheaf-based techniques for fundamental tasks in signal processing have also been developed [1].

The most basic data fusion question addressed by a sheaf is whether a complete set of observations constitutes a *global section*, which is a completely consistent, unified state [49]. The interface between the geometry of a sheaf and a set of observations can be quantified by the *consistency radius* [2]. Practical algorithms for data fusion arise simply as minimization algorithms for the consistency radius, for instance [50]. In this article, as in the more theoretical treatment [51], we show that sheaves provide additional, finer-grained analysis of consistency through the *consistency filtration*.

The potential significance of sheaves has been recognized for some time in the formal modeling community [37,52,53] in part because they bridge between two competing foundations for mathematics

— logic and categories [54]. The study of categories is closely allied with formal modeling of data processing systems [55–57].

Much of the sheaf literature focuses on studying models in the absence of observations. (This is rather different from our approach, which exploits observations.) For sheaf models that have enough structure, *cohomology* is a technical tool for explaining how local observations can or cannot be fused. This has applications in several different disciplines, such as network structure [42,58,59] and quantum information [60,61]. Computation of cohomology is straightforward [1], and efficient algorithms are now available [62].

3. Tracking Experiment

Researchers worked collaboratively to adapt existing practices around bear management to an experimental setup designed specifically to capture the properties of the targeted problems in a tracking task where multi-sensor integration is required. The data gathered in this way directly support the initial sheaf models developed in Section 4 with an optimal level of complexity and heterogeneity. In particular, what was sought is the inclusion of two targets (a bear or dummy collar, and then the wildlife tracker) jointly engaged in a collection of sensors of heterogeneous types (including GPS, radiocollars, and hand-written reports).

3.1. Black Bear Study Capture and Monitoring Methodology

The black bear observational study used in this paper is located in western North Carolina and centered on the urban/suburban area in and around the city of Asheville, North Carolina. Asheville is a medium-sized city (117 km²) with approximately 83,000 people, located in Buncombe County in the southern Appalachian Mountain range ([63], Figure 1).

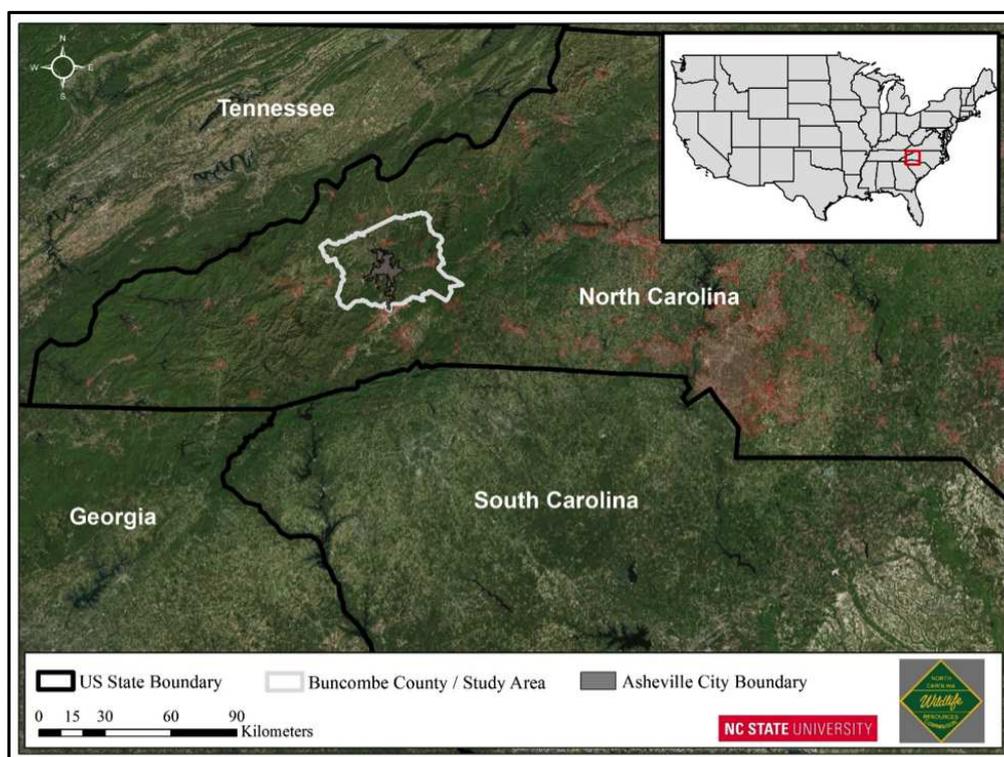


Figure 1. NCSU/NCWRC urban/suburban black bear (*Ursus americanus*) study area, Asheville, North Carolina, USA

Western North Carolina is characterized by variable mountainous topography (500–1800 m elevation), mild winters, cool summers, and high annual precipitation (130–200 cm/year), mostly in

the form of rainfall. Black bears occur throughout the Appalachian Mountains. The major forest types include mixed deciduous hardwoods with scattered pine [63], and pine-hardwood mix [64].

The Asheville City Boundary is roughly divided into four quadrants separated by two four-lane highways, Interstate 40, which runs east to west and Interstate 26, which runs north to south. Interstate 240 is a 9.1-mile (14.6 km) long Interstate Highway loop that serves as an urban connector for Asheville and runs in a semi-circle around the north of the city's downtown district. We define black bear capture sites and associated locational data as 'urban/suburban' if the locations fall within the Asheville city limit boundary and locations outside the city boundary are considered rural.

We used landowner reports of black bears on their property to target amenable landowners to establish trap lines on or near their property, while attempting to maintain a spatially balanced sample across the city of Asheville. The sampling area included suitable plots within, or adjacent to, the city limits. We set 10–14 culvert traps in selected locations that had documented bears near their properties. We checked culvert traps twice daily, once in the morning between 0800 – 1100 and again between 1830 – 2130.

We baited traps with day old pastries. Once captured, we immobilized bears with a mixture of ketamine hydrochloride (4.0 cc at 100 mg/mL), xylazine hydrochloride (1.0 cc at 100 mg/mL) and telazol (5 cc at 100 mg/mL) at a dose of 1cc per 100 lbs. We placed uniquely numbered eartags in both ears, applied a tattoo to the inside of the upper lip, removed an upper first premolar for age determination from all bears older than 12 months [65], and inserted a Passive Integrated Transponder tag (PIT tag) between the shoulder blades. Additionally, we collected blood, any ectoparasites present, and obtained hair samples. We recorded weight, sex, reproductive status (e.g., evidence of lactation, estrus, or descended testicles), morphology, date, and capture location for each bear. We fitted bears with a global positioning system (GPS) radiocollar (Vectronics, Berlin Germany) that did not exceed 2–3% of the animal's body weight. We administered a long-lasting pain-reliever and an antibiotic, and we reversed bears within approximately 60 min of immobilization with yohimbine hydrochloride (0.15 mg/kg).

Once black bears were captured and radiocollared, we used the virtual fence application on the GPS collars to obtain locational data every 15 min for individual bears that were inside the Asheville city limits (i.e., inside the virtual fence) and every hour when bears were outside the city limits. As a double-check, we monitored bears weekly with hand-held radio telemetry to ensure that the very high frequency (VHF) signal was transmitting correctly. We attempted this schedule on all bears to determine survival, proximity to roads and residential areas, home range size, habitat use, dispersal, and location of den sites. We stratified the locations throughout the day and night and programmed the GPS collars to send an electronic mortality message via iridium satellite to North Carolina State University (NCSU) if a bear remained immobile for more than 12 h. We immediately investigated any collar that emitted a signal to determine if it was mortality or a "slipped" collar.

3.2. Tracking Exercise

We used data from 12 GPS collars (6 stationary "dummy" collars and 6 "active" collars deployed on wild black bears) to design an urban tracking experiment to better estimate bear location and understand sources of error. The "dummy" collars were hidden and their locations were spatially balanced across the city limits of Asheville and hung on tree branches at various private residences. The active collars were deployed on three female black bears with the majority of their annual home range (i.e., locations) located inside the city limits. Bears N068, N083, and N024 were "tracked" three, two and one time, respectively, for the experiment.

We conducted two-hour tracking session which included 16 observations (i.e., stops) on either the dummy collar or the active bear. We stopped and recorded VHF detections on collars every 5–7 min along with data on the nearest landmark/road intersection (i.e., text description), Universal Transverse Mercator (UTM) position, elevation at each stop, compass bearing on the collar, handheld GPS position error for the observation point, and approximate distance to the collar. Also, we recorded a GPS track

log during the two-hour tracking session using a second handheld GPS unit located inside the vehicle; the vehicle GPS units logged continuous coordinates (approximately every 10 s) of our driving routes during the tracking session. Finally, the GPS collars simultaneously collected locations, with associated elevation and GPS “location error”, every 15 min for the duration of the two-hour tracking shift. GPS collar data include date, time, collar ID, latitude/longitude, satellite accuracy (horizontal dilution of precision (HDOP)), fix type (e.g., 2D, 3D, or val. 3D), and elevation.

4. Sheaf Modeling Methodology

We now introduce the fundamentals and details of our sheaf-based tracking model of our heterogeneous information integration problem. As noted above, our wildlife tracking problem is used as an instantiation of our sheaf model intended to aim at the right level of structural complexity to demonstrate the significance of the sheaf model. As such, we introduce the mathematics below in close conjunction with the example of the bear tracking problem itself, as detailed above in Section 3. For a more extensive sheaf theory introduction, the reader is directed towards [2,40].

A sheaf is a data structure used to store information over a topological space. The topological space describes the relations among the sensors while the sheaf operates on the raw input data, mapping all the sources into a common framework for comparison. In this work, the required topological space is determined by an *abstract simplicial complex* (ASC), a discrete mathematical object representing not only the available sensors as vertices (in this case, the four physical sensors involved, three on the human-vehicle tracking team, and one on the bear itself), but also their multi-way *interactions* as higher-order *faces*. Where the ASC forms the *base space* of the sheaf, the *stalks* that sit on its faces hold the recorded data. Finally, the sheaf model is completed by the specification of *restriction functions* which model interactions of the data among sensor combinations. In this context, we can then identify *assignments* as recorded readings; *sections* as assignments which are all consistent according to the sheaf model; and *partial assignments* and sections correspondingly over some partial collection of sensors.

We then introduce some techniques which are novel for sheaf modeling. Specifically, where global or partial sections indicate data which are *completely* consistent, we introduce *consistency structures* to represent data which are only somewhat consistent. While defined completely generally, consistency structures instantiated to the bear model in particular take the form of *n*-way standard deviations. Consistency structures in turn allow us to introduce *approximate sections* which can measure the degree of consistency among sensors. The *consistency radius* (more fully developed by Robinson [2] in later work) provides a native global measure of the uncertainty among the sensors present in any reading. Beyond that, the *consistency filtration* provides a detailed breakdown of the contributions of particular sensors and sensor combinations to that overall uncertainty.

4.1. Simplicial Sheaf Models

First we define the concept of an abstract simplicial complex, the type of topological space used to model our sensor network.

Definition 1. An *abstract simplicial complex* (ASC) over a finite base set U is a collection Δ of subsets of U , for which $\delta \in \Delta$ implies that every subset of δ is also in Δ . We call each $\delta \in \Delta$ with $d + 1$ elements a *d-face* of Δ , referring to the number d as its *dimension*. Zero dimensional faces (singleton subsets of U) are called *vertices*, and one dimensional faces are called *edges*. For $\gamma, \delta \in \Delta$, we say that γ is a *face* of δ (written $\gamma \rightsquigarrow \delta$) whenever γ is a proper subset of δ .

Remark 1. An abstract simplicial complex Δ over a base set U with $|U| = n$ can be represented in \mathbb{R}^n by mapping u_i to e_i and $\delta \in \Delta$ to the convex hull of its points.

An abstract simplicial complex can be used to represent the connections within a sensor network as follows. Take the base set U to be the collection of sensors for the network and take Δ to include

every collection of sensors that measure the same quantity. For example, as described in Section 3 above, there are four sensors used in the tracking experiment:

- The *GPS* reading on the Bear Collar, denoted G ;
- The *Radio VHF Device* receiver, denoted R ;
- The *Text* report, denoted T ; and
- The *Vehicle GPS*, denoted V .

As shown in Table 1, the bear collar GPS and Radio VHF Device give the location of the bear, and the text report, vehicle GPS, and Radio VHF Device give the location of the researcher. We can then let $U = \{V, R, T, G\}$ be our base sensor set. Then our ASC Δ , representing the tracking sensor network, contains the face denoted $H = \{V, R, T\}$ as the set of sensors reading off in human position; the face denoted $B = \{R, G\}$ as those reading off on bear position; and all their subsets, so nine total faces:

$$\Delta = \{ \{V, R, T\}, \{V, T\}, \{R, T\}, \{V, R\}, \{R, G\}, \{V\}, \{R\}, \{T\}, \{G\} \}$$

Table 1. Sensor matrix for the bear tracking.

	$H = \text{Human Position}$	$B = \text{Bear Position}$
$V = \text{Vehicle GPS (lat,long,ft)}$	✓	
$T = \text{Text}$	✓	
$R = \text{Receiver (UTM N,UTM E,m,deg,m)}$	✓	✓
$G = \text{GPS on Bear (UTM N, UTM E,m)}$		✓

Denoting the remaining pairwise sensor interaction faces as $X = \{V, T\}$, $Y = \{V, R\}$, and $Z = \{R, T\}$, then the ASC Δ can be shown graphically on the left side of Figure 2. Here the highest dimensional faces (the human and bear positions H and B respectively) are shown, with all the sub-faces labeled. The sensors are the singleton faces (rows), and are shown in black; and the higher dimensional faces (columns) in red.

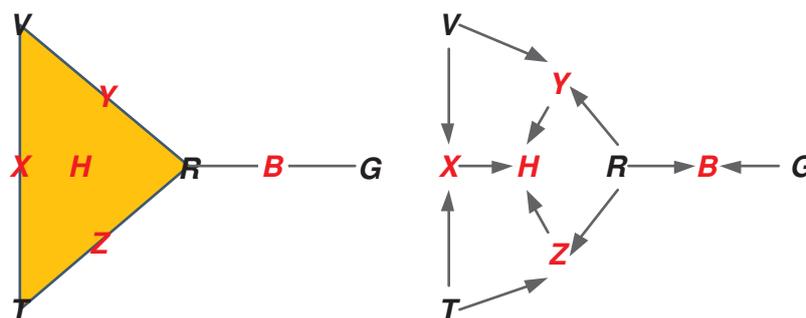


Figure 2. (Left) Simplicial complex of tracking sensors. (Right) Attachment diagram.

Note the presence of a solid triangle to indicate the three-way interaction H ; the presence of the four two-way interactions as edges; and finally the presence of each sensor individually.

As a general matter, the ASC implies a topological space representing all of these multi-way interactions. Each d -dimensional face is a d -dimensional hyper-tetrahedron, which are then “glued” together or “attached” according to the configuration of the sensor interactions. The ASC shown in this example is rather simple, consisting of the “topmost” interactions as a single 2-face (the triangle of the “human” measurements) and a 1-face (the RG edge for the bear), and then the additional seven

sub-faces. This simple structure has a single connected component, and no “open loops”. However, depending on the number of observables informed by a particular sensor, and their configuration, this structure can become arbitrarily complicated, with high order faces and complex connections including open loops or voids also of high dimension. The sheaf theoretical approach can represent all these interactions automatically. While such a “homological analysis” of the base space of our sheaf will not be the subject of this paper, it is of great interest in computational topology generally [2,66–68].

The right side of Figure 2 shows the **attachment diagram** corresponding to the ASC. This is a directed acyclic graph, where nodes are faces of the ASC, connected by a directed edge pointing up from a face to its attached face (co-face) of higher dimension.

Next, given a simplicial complex, a sheaf is an assignment of data to each face that is compatible with the inclusion of faces.

Definition 2. A *sheaf* \mathcal{S} of sets on an abstract simplicial complex Δ consists of the assignment of

1. a set $\mathcal{S}(\delta)$ to each face δ of Δ (called the **stalk** at δ), and
2. a function $\mathcal{S}(\gamma \rightsquigarrow \delta) : \mathcal{S}(\gamma) \rightarrow \mathcal{S}(\delta)$ (called the **restriction map** from γ to δ) to each inclusion of faces $\gamma \rightsquigarrow \delta$, which obeys

$$\mathcal{S}(\delta \rightsquigarrow \lambda) \circ \mathcal{S}(\gamma \rightsquigarrow \delta) = \mathcal{S}(\gamma \rightsquigarrow \lambda) \text{ whenever } \gamma \rightsquigarrow \delta \rightsquigarrow \lambda.$$

In a similar way, a **sheaf of vector spaces** assigns a vector space to each face and a linear map to each attachment. In addition, a **sheaf of groups** assigns a group to each face and a group homomorphism to each attachment.

Intuitively, the stalk over a face is the space where the data associated with that face lives; and the restriction functions establish the grounds by which interacting data can be said to be consistent or not.

Returning to the sensors of our tracking network, a bear’s GPS collar provides its position and elevation given in units of (UTM N, UTM E, m). Thus, the stalk over the $G = \text{Bear Collar GPS}$ vertex is \mathbb{R}^3 . Then when the researcher goes out to locate a bear he supplies two sets of coordinates and a text description of his location. The first coordinates, GPS position and elevation given in units of (lat, long, ft), come from the GPS in the vehicle the researcher is driving. Then when the researcher stops to make a measurement, he or she supplies a text description of his location. Additionally, the researcher records his or her position and elevation in (UTM N, UTM E, m), and the bear’s position (in relation to his own) off a handheld VHF tracking receiver that connects to the bear’s radiocollar, yielding a data point in \mathbb{R}^5 . The bear’s relative position is measured in polar coordinates (r, θ) . The data sources and stalks are summarized in Table 2, and the left side of Figure 3 shows our ASC now adorned with the stalks on the sensor faces.

Table 2. Data feeds for tracking sheaf model

Vertex	Data Format	Description	Stalk
$G = \text{Bear Collar}$	(E, N, m)	Position and elevation of bear from collar	\mathbb{R}^3
$V = \text{Vehicle GPS}$	(lat, long, ft)	Position and elevation of human from vehicle	\mathbb{R}^3
$T = \text{Text}$	string	Text description of human's location	set of strings
$R = \text{Radio VHF Device}$	(E, N, m, m, deg)	Position and elevation of human and position of bear relative to human	\mathbb{R}^5

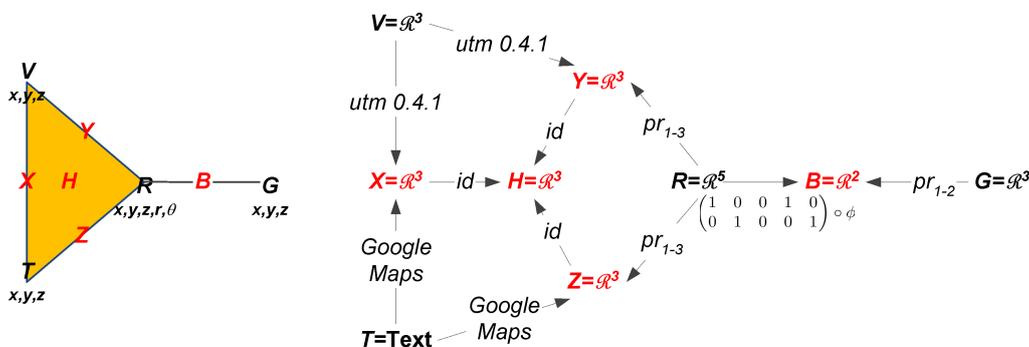


Figure 3. (Left) ASC adorned with stalks on the sensor vertices. (Right) Sheaf model.

Now that we have specified the stalks for the vertices, we must decide the stalks for the higher order faces along with the restriction maps. Please note that the triangle face of the tracking ASC is formed by sensors each measuring the researcher’s location. However, the data types do not agree. Thus, to compare these measurements the information should first be transformed into common units along the edges and then passed through to the triangular face. Since multiple sensors read out in UTM for position and meters for elevation, we choose these as the common coordinate system. Then the stalk for the higher order faces of the triangle is \mathbb{R}^3 . To convert the text descriptions we use Google Maps API [69] and to convert the (lat, long, ft) readings, coming from the vehicle, we use the Python open source package utm 0.4.1 [70]. The data from the handheld tracker is already in the chosen coordinates so the restriction map to the edges of the triangle is simply a projection.

Likewise, the edge coming off the triangle connects the two sensors that report on the bear’s location. To compare the sensors we need to map them into common coordinates through the restriction maps. The sensor on the bear’s collar provides a position and elevation for the bear, but the radio VHF device only reports the bear’s position relative to the researcher. Therefore, to compare the two readings along the adjoining edge, we must forgo a degree of accuracy and use \mathbb{R}^2 as the stalk over the edge. For the bear collar sensor this means that the restriction map is a restriction which drops the elevation reading. For the VHF data, we first convert the polar coordinates into rectangular as follows:

$$\phi : \mathbb{R}^5 \rightarrow \mathbb{R}^5, \phi((x, y, z, r, \theta)^T) = (x, y, z, r \cos(\theta), r \sin(\theta))^T. \tag{1}$$

Then to obtain the position of the bear, we add the relative bear position to the human’s position.

The right side of Figure 3 thus shows the full sheaf model on the attachment diagram of the ASC. Note the stalks on each sensor vertex, text for the text reader and numerical vectors for all the others.

Restriction functions are then labels on the edges connecting lower dimensional faces to their attached higher dimensional faces. UTM conversion and the Google Maps interface are on the edges coming from V and T respectively. Since the pairwise relations all share UTM coordinates, only id mappings are needed among them up to the three-way H face. Labels of the form pr_{x-y} are projections of the corresponding coordinates (also representable as binary matrices of the appropriate form). Finally, the restriction from R up to B is the composition of the polar conversion of the final two components with the projection on the first two to predict the bear position from the radiocollar GPS, bearing, and range.

The sheaf model informs about the agreement of sensors through time as follows. At a given time t , each vertex is assigned the reading last received from its corresponding sensor, a data point from its stalk space. These readings are then passed through the restriction maps to the higher order faces for comparison. If the two measurements are received by an edge agree, this single value is assigned to that edge and the algorithm continues. If it is possible to assign a single value $H = \langle x, y, z \rangle$ to the $\{Text, Vehicle\ GPS, Radio\ VHF\ Device\}$ face then all three measurements of the human's location agree. Similarly, obtaining a single value $B = \langle x, y \rangle$ for the $\{Radio\ VHF\ Device, Bear\ Collar\}$ edge signifies agreement among the two sensors tracking the bear. If a full assignment can be made, we call it a **global section**. Non-agreement is definitely possible, giving rise to the notion of an **assignment**.

Definition 3. Let \mathcal{S} be a sheaf on an abstract simplicial complex Δ . An **assignment** $\alpha: \Delta \rightarrow \prod_{\delta \in \Delta} \mathcal{S}(\delta)$ provides a value $\alpha(\delta) \in \mathcal{S}(\delta)$ to each face $\delta \in \Delta$. A **partial assignment**, β , provides a value for a subset $\Delta' \subset \Delta$ of faces, $\beta: \Delta' \rightarrow \prod_{\delta \in \Delta'} \mathcal{S}(\delta)$. An assignment s is called a **global section** if for each inclusion $\delta \rightsquigarrow \lambda$ of faces, $\mathcal{S}(\delta \rightsquigarrow \lambda)(s(\delta)) = s(\lambda)$.

As an example, a possible global section for our tracking sheaf is:

$$\begin{aligned}
 s(\text{Text}) &= \text{'Intersection of Victoria Rd and Meadow Rd'} \\
 s(\text{Vehicle GPS}) &= \begin{pmatrix} 35.6^\circ \text{ lat} \\ -82.6^\circ \text{ long} \\ 2019 \text{ ft} \end{pmatrix}, \\
 s(\text{Bear Collar}) &= \begin{pmatrix} 358391 \text{ E} \\ 3936750 \text{ N} \\ 581 \text{ m} \end{pmatrix}, \\
 s(\text{Radio VHF Device}) &= \begin{pmatrix} 358943 \text{ E} \\ 3936899 \text{ N} \\ 615.4 \text{ m} \\ 572 \text{ m} \\ 195^\circ \end{pmatrix}, \\
 s(\{\text{Radio VHF Device}, \text{Bear Collar}\}) &= \begin{pmatrix} 358391 \text{ E} \\ 3936750 \text{ N} \end{pmatrix} \\
 \text{and } s(\{\text{Text}, \text{Vehicle GPS}\}) &= s(\{\text{Text}, \text{Radio VHF Device}\}) \\
 &= s(\{\text{Vehicle GPS}, \text{Radio VHF Device}\}) \\
 &= s(\{\text{Text}, \text{Vehicle GPS}, \text{Radio VHF Device}\}) = \begin{pmatrix} 358943 \text{ E} \\ 3936899 \text{ N} \\ 615.4 \text{ m} \end{pmatrix}.
 \end{aligned}$$

4.2. Consistency Structures, Pseudosections, and Approximate Sections

At a given time, a global section of our tracking model corresponds to the respective sensor readings simultaneously agreeing on the location of both the researcher and the bear. This would be

ideal, however, in practice it is very unlikely that all the sensor readings will agree to the precision shown above. For certain applications, such as our tracking model, the equality constraint of a global section may be too strict. Consistency structures [71] and the consistency radius [2] address this concern. Consistency structures loosen the constraint that sensor data on the vertices match, on passing through restrictions, by instead requiring that they merely *agree*. Agreement is measured by a boolean function on each individual face. The pairing of a sheaf with a collection of these boolean functions is called a **consistency structure**.

Definition 4. A **consistency structure** is a triple $(\Delta, \mathcal{S}, \mathbf{C})$ where Δ is an abstract simplicial complex, \mathcal{S} is a sheaf over Δ , and \mathbf{C} is the assignment to each non-vertex d -face $\lambda \in \Delta, d > 0$, of a function

$$\mathbf{C}_\lambda : \left(\binom{\mathcal{S}(\lambda)}{\dim \lambda + 1} \right) \rightarrow \{0, 1\}$$

where $\binom{Z}{k}$ denotes the set of multisets of length k over Z .

The multiset in the domain of \mathbf{C}_λ represents the multiple sheaf values to be compared for all the vertices impinging on a non-vertex face λ , while the codomain $\{0, 1\}$ indicates whether they match “well enough” or not. In particular, we have the **standard consistency structure** for a sheaf \mathcal{S} , which assigns an *equality* test to each non-vertex face $\lambda = \{v_1, v_2, \dots, v_k\}$:

$$\mathbf{C}_\lambda([z_1, z_2, \dots, z_k]) = \begin{cases} 1, & \text{if } z_1 = z_2 = \dots = z_k \\ 0, & \text{otherwise} \end{cases}$$

where $[\cdot]$ denotes a multiset, and $z_i = \mathcal{S}(\{v_i\} \rightsquigarrow \lambda)(\alpha(\{v_i\}))$.

A consistency structure broadens the equality requirement natively available in a sheaf to classes of values which are considered equivalent. However, beyond that, in our tracking model, each stalk is a metric space. Thus we can use the natural metric to test that points are merely “close enough”, rather than agreeing completely or being equivalent. Using ϵ to indicate the quantitative amount of error present or tolerated in an assignment, we define the **ϵ -approximate consistency structure** for a sheaf \mathcal{S} as follows. For each non-vertex size k face $\lambda = \{v_1, v_2, \dots, v_k\} \in \Delta, k > 1$ define

$$\mathbf{C}_\lambda([z_1, z_2, \dots, z_k]) = \begin{cases} 1 & \text{if } \hat{\sigma}([z_1, z_2, \dots, z_k]) \leq \epsilon, \\ 0 & \text{otherwise,} \end{cases}$$

where

$$\hat{\sigma}(Y) = \sqrt{\frac{1}{|Y|} \sum_{y \in Y} \|y - \mu_Y\|^2} = \sqrt{\frac{1}{|Y|} \text{Tr}(\Sigma_Y)}$$

μ_Y is the mean of the set Y , and Σ_Y is the covariance matrix of the multidimensional data Y . This measure of consistency gives a general idea of the spread of the data.

The corresponding notion of a global section for sheaves is called a pseudosection for a consistency structure.

Definition 5. An assignment $s \in \prod_{\delta \in \Delta} \mathcal{S}(\delta)$ is called a $(\Delta, \mathcal{S}, \mathbf{C})$ -**pseudosection** if for each non-vertex face

$$\lambda = \{v_1, \dots, v_k\}$$

1. $\mathbf{C}_\lambda([\mathcal{S}(\{v_i\} \rightsquigarrow \lambda)s(\{v_i\}) : i = 1, \dots, k]) = 1$, and
2. $\mathbf{C}_\lambda([\mathcal{S}(\{v_i\} \rightsquigarrow \lambda)s(\{v_i\}) : i = 1, \dots, j-1, j+1, \dots, k] \cup [s(\lambda)]) = 1$ for all $j = 1, 2, \dots, k$.

An assignment s which is a pseudosection guarantees that for any non-vertex face λ , (1) the restrictions of its vertices to the face are close, and (2) the value assigned to the face is consistent with the restricted vertices.

When \mathbf{C} is the standard consistency structure, then it comes about that

$$s(\lambda) = \mathcal{S}(v_1 \rightsquigarrow \gamma)s(\{v_1\}) = \mathcal{S}(v_2 \rightsquigarrow \gamma)s(\{v_2\}) = \dots = \mathcal{S}(v_k \rightsquigarrow \gamma)s(\{v_k\})$$

for each non-vertex face $\lambda = \{v_1, \dots, v_k\}$. Thus pseudosections of a standard consistency structure are global sections of the corresponding sheaf.

On the other hand, when \mathbf{C} is the ϵ -approximate consistency structure, then we have that

1. $\hat{\sigma}([\mathcal{S}(\{v_i\} \rightsquigarrow \lambda)s(\{v_i\}) : i = 1, \dots, k]) \leq \epsilon$, and
2. $\hat{\sigma}([\mathcal{S}(\{v_i\} \rightsquigarrow \lambda)s(\{v_i\}) : i = 1, \dots, j - 1, j + 1, \dots, k] \cup [s(\lambda)]) \leq \epsilon$ for all $j = 1, 2, \dots, k$.

for each non-vertex face $\lambda = \{v_1, v_2, \dots, v_k\}$.

Pseudosections of ϵ -approximate consistency structures are precisely the generalization needed for our tracking model. The ϵ -approximate consistency structure over our tracking sheaf consists of the following assignment of functions:

$$\begin{aligned} \mathbf{C}_X, \mathbf{C}_Y, \mathbf{C}_Z & : \left(\binom{\mathbb{R}^3}{2} \right) \rightarrow \{0, 1\}, \\ \mathbf{C}_B & : \left(\binom{\mathbb{R}^2}{2} \right) \rightarrow \{0, 1\} \\ \mathbf{C}_H & : \left(\binom{\mathbb{R}^3}{3} \right) \rightarrow \{0, 1\} \end{aligned}$$

where

$$\mathbf{C}_X([z_1, z_2]) = \mathbf{C}_Y([z_1, z_2]) = \mathbf{C}_Z([z_1, z_2]) = \mathbf{C}_B([z_1, z_2]) = \begin{cases} 1 & \text{if } \hat{\sigma}([z_1, z_2]) \leq \epsilon \\ 0 & \text{otherwise} \end{cases}$$

and

$$\mathbf{C}_H([z_1, z_2, z_3]) = \begin{cases} 1 & \text{if } \hat{\sigma}([z_1, z_2, z_3]) \leq \epsilon \\ 0 & \text{otherwise.} \end{cases}$$

A pseudosection for this consistency structure is an assignment

$$\langle s(R), s(v), s(T), s(G), s(X), s(Y), s(Z), s(B), s(H)) \rangle$$

such that for each non-vertex face $\lambda \in \{X, Y, Z, B, H\}$

1. $\hat{\sigma}([\mathcal{S}(\{v\} \rightsquigarrow \lambda)s(\{v\}) : v \in \lambda]) \leq \epsilon$, and
2. $\hat{\sigma}([\mathcal{S}(\{v\} \rightsquigarrow \lambda)s(\{v\}) : v \in \lambda] \cup [s(\lambda)] \setminus [s(w)]) \leq \epsilon$ for all $w \in \lambda$.

Now $s(R), s(V), s(T)$, and $s(G)$ are measurements given by the data feeds, so we are left to assign $s(X), s(Y), s(Z), s(B)$, and $s(H)$ to minimize ϵ . Physically, a pseudosection of the ϵ -approximate consistency structure for our tracking model assigns positions for both the human and bear so that the “spread” of all measurements attributed to any face is bounded by ϵ . The question then becomes how does one minimize ϵ efficiently? The theorem below states that the minimum ϵ for which a pseudosection exists is determined only by the restricted images of the vertices.

Theorem 1. *Let \mathcal{S} be a sheaf over an abstract simplicial complex Δ such that each stalk is a metric space and let $s \in \prod_{\delta \in \Delta} \mathcal{S}(\delta)$ be an assignment. The minimum ϵ for which s is a pseudosection of the ϵ -approximate consistency structure $(\Delta, \mathcal{S}, \mathbf{C})$ is*

$$\epsilon^* = \max_{\lambda \in \Delta \setminus \{v : v \in V\}} \hat{\sigma}([\mathcal{S}(\{w\} \rightsquigarrow \lambda)s(\{w\}) : w \in \lambda]).$$

We call this minimum error value ϵ^* the **consistency radius**.

The consistency radius is a quantity that is relatively easy to interpret. If it is small, it means that there is minimal disagreement among observations. If it is large, then it indicates that at least some sensors disagree.

The proof of this theorem is an application of the following Lemma.

Lemma 1. Let $Z = \{z_1, \dots, z_k\}$ be a set of real numbers with mean μ_Z . For $z \in Z$, define $Y_z = Z \setminus z \cup \mu_Z$. Then $\forall z \in Z, \hat{\sigma}(Y_z) \leq \hat{\sigma}(Z)$.

Proof. Fix $x \in Z$. Please note that $|Y_z| = k$, and that $\hat{\sigma}(Y_z) \leq \hat{\sigma}(Z)$ if and only if $\hat{\sigma}(Y_z)^2 \leq \hat{\sigma}(Z)^2$. Next,

$$\hat{\sigma}(Y_z)^2 = \frac{1}{k} \sum_{i=1}^k (y_i - \mu_Y)^2 \leq \frac{1}{k} \sum_{i=1}^k (y_i - \mu_Z)^2$$

since the expression is minimized around the mean. Then

$$\frac{1}{k} \sum_{i=1}^k (y_i - \mu_Z)^2 \leq \hat{\sigma}(Z)^2$$

as $0 \leq (z - \mu_Z)^2$, completing the proof. \square

4.3. Maximal Consistent Subcomplexes

The association of data with a signal network, as in the case of our tracking model, lends itself naturally to the question of consistency. Data is received from various signal sources, some of which should agree but may not. Thus, one would like to identify maximal consistent portions of the network. In [72], Praggastis gave an algorithm for finding a unique set of maximally consistent subcomplexes.

Let $(\Delta, \mathcal{S}, \mathbf{C})$ be a consistency structure. For each $\tau \in \Delta$, we define $\text{star}(\tau) = \{\rho \in \Delta : \tau \subseteq \rho\}$ to be the set of faces containing τ . We can think of the star of a face as its sphere of influence. Similarly, for $\Delta' \subset \Delta$, we define $\text{star}(\Delta') = \bigcup_{\tau \in \Delta'} \text{star}(\tau)$. Next, for a subset of vertices $W \subseteq U$, we define $\Delta_W = \{\tau \in \Delta : \tau \subseteq W\}$ to be the subcomplex of Δ induced by W . Please note that Δ_W inherits the consistency structure $(\Delta_W, \mathcal{S}(\Delta_W), \mathbf{C}|_{\Delta_W})$ from $(\Delta, \mathcal{S}, \mathbf{C})$. Now let s be a sheaf partial assignment on just the vertices U . We say that s is consistent on Δ_W if $C_\gamma([\mathcal{S}(v \rightsquigarrow \gamma)s(v) : v \in \gamma]) = 1$ for each non-vertex face γ in Δ_W . The theorem below states that we can associate a unique set of maximally consistent subcomplexes to any vertex assignment.

Theorem 2. Let $(\Delta, \mathcal{S}, \mathbf{C})$ be a consistency structure and let α be a sheaf partial assignment on U . There exists a unique collection of subsets $\{W_i\}$ of U which induce subcomplexes $\{\Delta_{W_i}\}$ of Δ with the following properties:

1. The assignment α is consistent on each Δ_{W_i} , and any subcomplex on which α is consistent has some Δ_{W_i} as a supercomplex.
2. $\bigcup \text{star}(\Delta_{W_i})$ is a cover of Δ .

The proof of this result is constructive. The basic operation is described in the following lemma.

Lemma 2. Let $(\Delta, \mathcal{S}, \mathbf{C})$ be a consistency structure and let α sheaf partial assignment on U . If $C_\gamma([\mathcal{S}(v \rightsquigarrow \gamma)\alpha(v) : v \in \gamma]) = 0$ for some non-vertex face $\gamma \in \Delta$, then there exist subsets $\{W_i\}$ of U which induce subcomplexes $\{\Delta_{W_i}\}$ of Δ such that:

1. $\gamma \notin \Delta_{W_i}$ for all i .
2. Every subcomplex on which α is consistent has some Δ_{W_i} as a supercomplex.
3. $\bigcup \text{star}(\Delta_{W_i})$ is a cover of Δ .

Proof. For each vertex $v_i \in \gamma$, let $W_i = U \setminus v_i$ with induced subcomplex Δ_{W_i} . Clearly γ does not belong to any of the Δ_{W_i} , and if Δ' is a subcomplex of Δ on which α is consistent, then the vertices of Δ' is a subset of some W_i . Finally, since every vertex in U belongs to at least one W_i , we have that $\Delta = \bigcup \text{star}(\Delta_{W_i})$. \square

Proof of Theorem 2. If α is consistent on Δ , we are done, so suppose that $C_\gamma([\mathcal{S}(v \rightsquigarrow \gamma)\alpha(v) : v \in \gamma]) = 0$ for some non-vertex face $\gamma \in \Delta$. By Lemma 2, there exists a set of subcomplexes $\{\Delta_{W_i}\}$ such that (1) $\gamma \not\subseteq \Delta_{W_i}$ for all i , (2) every subcomplex on which α is consistent has some Δ_{W_i} as a supercomplex, and (3) $\bigcup \text{star}(\Delta_{W_i})$ is a cover of Δ . Next, we repeat the process for each complex Δ_{W_i} individually. If α is inconsistent on Δ_{W_i} , we apply Lemma 2 and replace W_i with $\{W_{i_k}\}$ and Δ_{W_i} with $\{\Delta_{W_{i_k}}\}$.

Since Δ is finite, the process must terminate with a list of consistent subcomplexes. It is possible to have nested sequences of subcomplexes in the final list. Since we are only interested in maximal consistent subcomplexes, we drop any complex which has a supercomplex in the list. Let $\{W_i\}$ and $\{\Delta_{W_i}\}$ be the final list of vertex sets and subcomplexes, respectively, obtained in this way. By construction, α is consistent on each Δ_{W_i} . Further, by Lemma 2, any subcomplex on which α is consistent belongs to some Δ_{W_i} and the set $\{\text{star}(\Delta_{W_i})\}$ is a cover of Δ . Finally, suppose that $\{X_j\}$ and $\{\Delta_{X_j}\}$ are another pair of vertex sets and subcomplexes which satisfy the properties of the theorem. Fix some Δ_{X_k} . Since α is consistent on Δ_{X_k} , there exists some Δ_{W_ℓ} such that $\Delta_{X_k} \subseteq \Delta_{W_\ell}$. However, by reversing the argument, there is also some Δ_{X_m} such that $\Delta_{W_\ell} \subseteq \Delta_{X_m}$. Finally, since Δ_{X_k} is maximal in $\{\Delta_{X_j}\}$, we must have that $k = m$ and $\Delta_{X_k} = \Delta_{W_\ell}$. \square

In the context of our tracking model, given a collection of sensor readings and an ϵ , Theorem 2 provides a maximal vertex cover so that each associated subcomplex is ϵ -approximate consistent.

4.4. Measures on Consistent Subcomplexes

In this section, we define a measure on the vertex cover associated with a set of maximal consistent subcomplexes by identifying the set of covers as a graded poset. We use the rank function of the poset as a measure. For a comprehensive background on posets, the interested reader is referred to [73].

Let $\mathcal{P} = \langle P, \leq \rangle$ be a poset. For a subset $X \subseteq P$, let $\downarrow X = \{s \leq x : x \in X\}$ denote the ideal of X in P . The set of all ideals, $\mathcal{I}(P)$, ordered by inclusion, forms a poset denoted $J(P)$. Additionally, there is a one-to-one correspondence between $J(P)$ and the set of all antichains $\mathcal{A}(P)$. Specifically, the functions

$$\begin{aligned} \max : J(P) &\rightarrow \mathcal{A}(P), \max(I) = \{x \in I : x \not\prec y \text{ for all } y \in I\}, \text{ and} \\ \downarrow : \mathcal{A}(P) &\rightarrow J(P), \downarrow(\mathcal{A}) = \downarrow \mathcal{A} \end{aligned}$$

are inverses of each other. Next, recall that a poset is called graded if it can be stratified.

Definition 6. A poset $\mathcal{P} = \langle P, \leq \rangle$ is **graded** if there exists a rank function $r : P \rightarrow \mathbb{N} \cup \{0\}$ such that $r(s) = 0$ if s is a minimal element of P , and $r(q) = r(p) + 1$ if $p \prec q$ in P . If $r(s) = i$, we say that s has rank i . The maximum rank, $\max_{p \in P} \{r(p)\}$, is called the **rank** of \mathcal{P} .

As a consequence of the Fundamental Theorem for Finite Distributive Lattices, it is well known that $J(P)$ is a graded poset.

Proposition 1. [73, Prop. 3.4.5] If $\mathcal{P} = \langle P, \leq \rangle$ is an n -element poset, then $J(P)$ is graded of rank n . Moreover, the rank $r(I)$ of $I \in J(P)$ is the cardinality of I .

Now consider 2^n , the Boolean lattice on n elements. Explicitly, 2^n consists of all subsets of $\{1, 2, \dots, n\}$ where for all $a, b \subseteq \{1, 2, \dots, n\}$ we define $a \leq b$ precisely when $a \subseteq b$, and the lattice

operations are defined as $a \wedge b = a \cap b$, and $a \vee b = a \cup b$. The vertex cover $\{W_i\}$ associated with a maximal set of consistent subcomplexes $\{\Delta_{W_i}\}$, provided by Theorem 2, can be viewed as an antichain of 2^n where $n = |U|$. Moreover, by design, each of these antichains is a set cover of $\{1, 2, \dots, n\}$. We will call such an antichain **full**. Specifically, $\mathcal{A} = \{a_1, a_2, \dots, a_k\}$ is said to be full if $\bigcup_{i=1}^k a_i = \{1, 2, \dots, n\}$. The set of all full ideals $\overline{\mathcal{I}(2^n)}$ forms an induced subposet $\overline{J(2^n)}$ of the graded poset $J(2^n)$. Next, we show that $\overline{J(2^n)}$ is also graded.

Proposition 2. $\overline{J(2^n)}$ is a graded poset of rank $2^n - (n + 1)$.

Proof. By Proposition 1, $J(2^n)$ is a graded poset of rank 2^n . Moreover, the rank function is given by $r(I) = |I|$. Now $\overline{J(2^n)}$ is an induced subposet of $J(2^n)$ with the single minimum element $\mathbf{0} := \{\emptyset, \{1\}, \{2\}, \dots, \{n\}\}$. Define the function

$$\bar{r} : \overline{J(2^n)} \rightarrow \mathbb{N} \cup \{0\}, \bar{r}(I) = r(I) - r(\mathbf{0}) = |I| - (n + 1).$$

Please note that if $I_1 \prec I_2$ in $\overline{J(2^n)}$, then $I_1 \prec I_2$ in $J(2^n)$. The fact that \bar{r} is a rank function then follows from the fact that r is a rank function. \square

Since $\overline{J(2^n)}$ is a graded poset, we can use the rank function as a measure on the vertex cover associated a set of maximal consistent subcomplexes.

Definition 7. Let $(\Delta, \mathcal{S}, \mathcal{C})$ be a consistency structure with $|V| = n$ and let $\mathcal{A} = \{W_i\}$ be a vertex cover obtained from Theorem 2. Then we define

$$\bar{r}(\mathcal{A}) = |\downarrow \mathcal{A}| - (n + 1).$$

Please note that $0 \leq \bar{r}(\mathcal{A}) \leq 2^n - (n + 1)$ and a larger value indicates more consistent subcomplexes. Indeed, if $\downarrow \mathcal{A}_1 \subseteq \downarrow \mathcal{A}_2$, then $\bar{r}(\mathcal{A}_1) \leq \bar{r}(\mathcal{A}_2)$.

4.5. Consistency Filtrations

Now, consider a simplicial complex Δ , a sheaf \mathcal{S} on Δ , and partial assignment α to the vertices of Δ . Given any value $\epsilon \geq 0$ we can consider the ϵ -approximate consistency structure $(\Delta, \mathcal{S}, \mathcal{C}_\epsilon)$ and use Theorem 2 to obtain the set of maximal consistent subcomplexes for the assignment α . Varying ϵ we obtain a filtration, which we call the **consistency filtration**, of vertex covers corresponding to landmark ϵ values $\epsilon_0 = 0 < \epsilon_1 < \dots < \epsilon_{\ell-1} < \epsilon_\ell = \epsilon^*$ where we recall that ϵ^* is the consistency radius, i.e., the smallest value of ϵ for which α is a $(\Delta, \mathcal{S}, \mathcal{C}_\epsilon)$ -pseudosection. The corresponding refinement of vertex covers is $C_0 \leq C_1 \leq \dots \leq C_{\ell-1} \leq C_\ell = U$ where each is a set of subcomplexes whose union is Δ . We can also compute the sequence of cover measures $p_0 < p_1 < \dots < p_{\ell-1} < p_\ell = 1$ for each C_i . (A somewhat different perspective on the consistency filtration is discussed in [51], in which the consistency filtration is shown to be both functorial and robust to perturbations, and is so in both the sheaf and the assignment.)

The consistency filtration is a useful tool for examining consistency among a collection of sensors. If the distance between two particular consecutive landmark values ϵ_i and ϵ_{i+1} is considerably larger than the rest, this indicates that there is considerable difficulty in resolving a disagreement between at least two groups of sensors. Which sensors are at fault for this disagreement can be easily determined by comparing the covers C_i and C_{i+1} .

5. Results

As detailed in Section 3, data were gathered for twelve locating sessions, six with dummy collars hidden around the city and six with live bears. The six live sessions were distributed over three female

bears with session frequencies 3, 2, and 1. Each locating session was two hours long and included 16 measurements.

We first present overall measurement results across both bear and dummy collars. We then drill down to a particular bear, N024, in the context of its overall consistency radius. We conclude with a detailed examination of a single measured assignment to this bear, particularly at minute 5.41, in order to understand the details of the consistency filtration.

5.1. Overall Measurements

Overall consistency radii ϵ^* are shown below: Figure 4 shows for the full system; Figure 5 for the human; and Figure 6 for the bears.

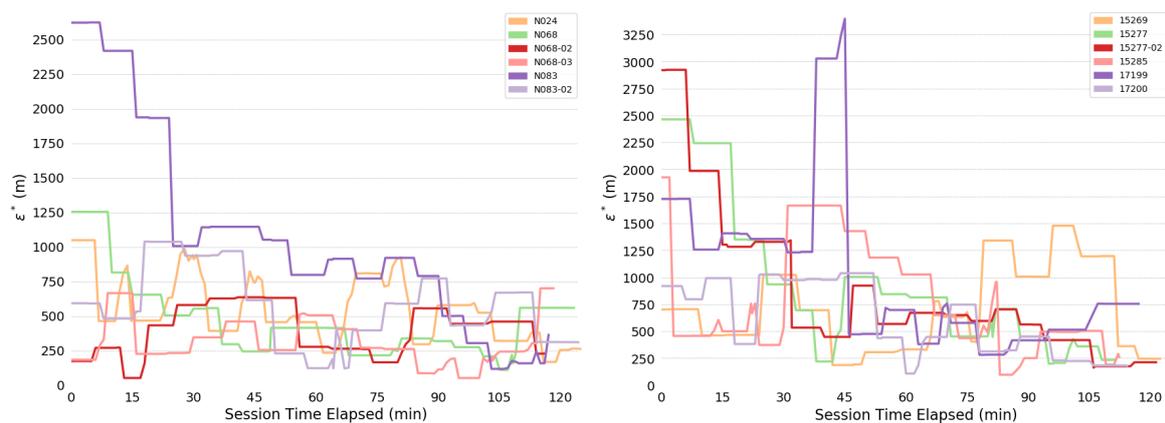


Figure 4. Consistency radii ϵ^* over time, full system. (Left) Bears. (Right) Dummy collars.

The most obvious trend in Figure 4 is a decrease in consistency radius over time. This implies that the overall error in the measurements of the bear-human system decrease over time. Because the consistency radius for the measurements of the humans alone (Figure 5 does not show this trend while the measurements of bear alone (Figure 6) does, this implies that there is a disagreement between the various measurements of the bear, but not the humans. This indicates that there is an overall improvement in agreement between measurements as the experiment progresses. In fact, as subsequent sections will show, the VHF measurements degrade with increasing range. As the experiment progressed, the humans approached the location of the bear, and so the error in the VHF measurements decreased. As the humans approach the bear, this means that the VHF measurements are more consistent with the other measurements. This tends to reduce the consistency radius since small consistency radius means that the data and model are in agreement.

On the other hand, the reader is cautioned that consistency radius does not and cannot tell you *what* the measurements are, nor does it tell you which measurements are more trustworthy. However, the consistency radius does not presuppose any distributions, nor does it require any parameter estimation in order to indicate agreement. As the next few sections indicate, the consistency filtration does a good job of blaming the sensor at fault for the time variation seen in Figures 4 and 5, again with minimal assumptions. Strictly speaking, trying to use the sensors to estimate some additional quantity is an aggregation step that brings with it additional assumptions. Using the consistency radius alone avoids making these assumptions.

The reader may ask why considering the time variation of consistency radius is more useful than some aggregate statistic, such as its maximum value. However, taking the max (or any other extremal statistic) is poor statistical practice, and worthy of being avoided on those grounds! Timeseries are almost always better at expressing system behaviors, and in our case it happens to reveal a clear trend.

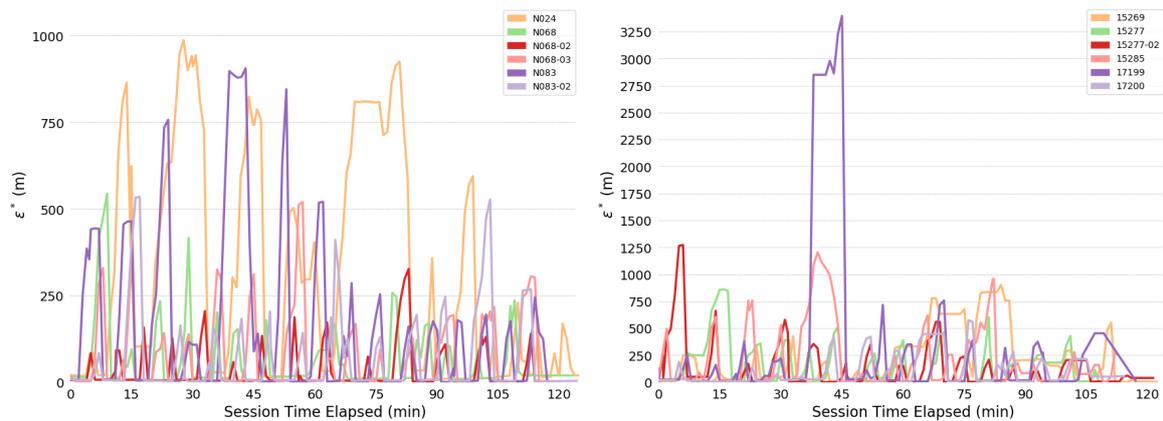


Figure 5. Consistency radii ϵ^* over time, human. (Left) Bears. (Right) Dummy collars.

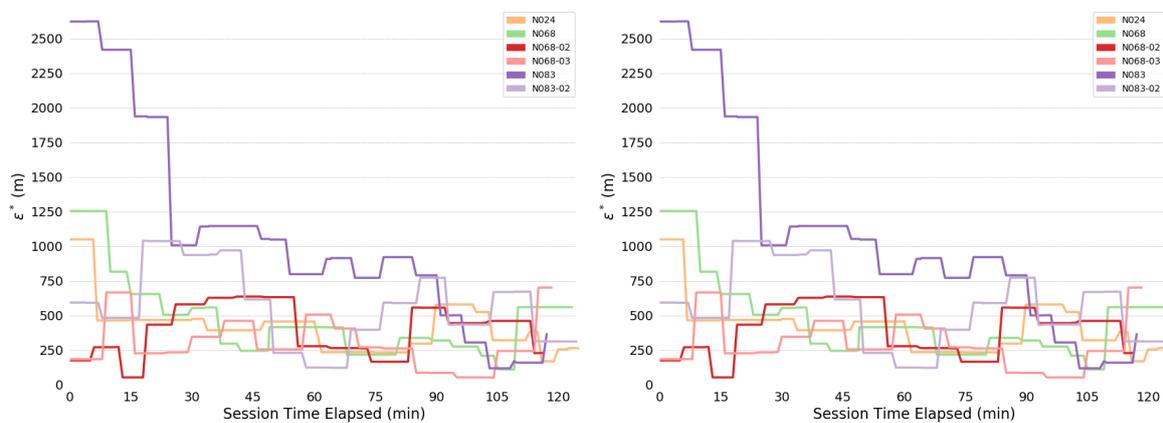


Figure 6. Consistency radii ϵ^* over time, bears. (Left) Bears. (Right) Dummy collars.

5.2. Example: Bear N024

We now consider in detail the readings for Bear N024, recorded on 6 January 2016. First we show (some of) the raw data readings. Figure 7 shows a sample of the first twelve data points for the car GPS V , showing latitude, longitude, and altitude. Figure 8 shows the readings from the bear GPS B , including UTM and elevation. Finally, Figure 9 shows a combination of the text reading T and the radiocollar R , including the five variables UTM, elevation, and bearing and range to the bear. Ranges in increments of meters according to Table 3.

Date/time	Latitude	Longitude	Elevation (ft)
1/6/2016 15:30	35.571187	-82.514005	2268
1/6/2016 15:31	35.571194	-82.514003	2267.9
1/6/2016 15:32	35.571197	-82.513999	2263.4
1/6/2016 15:32	35.571191	-82.514007	2263.7
1/6/2016 15:33	35.57118	-82.513978	2262.2
1/6/2016 15:34	35.5712	-82.513993	2259.5
1/6/2016 15:34	35.571072	-82.513998	2261.2
1/6/2016 15:34	35.570937	-82.513967	2261.9
1/6/2016 15:35	35.570563	-82.513949	2266.8
1/6/2016 15:35	35.570106	-82.514089	2256.4
1/6/2016 15:35	35.570086	-82.514123	2255.5

Figure 7. Vehical GPS readings V for Bear N024.

Observation	Time	Elevation_m	UTM_E	UTM_N
1	3:30:13 PM	577	362888	3938373
2	3:45:09 PM	616	362882	3938377
3	4:00:16 PM	559	362866	3938384
4	4:15:18 PM	580	362877	3938390
5	4:30:13 PM	603	362877	3938374
6	4:45:14 PM	585	362879	3938379
7	5:00:52 PM	594	362877	3938385
8	5:15:43 PM	603	362874	3938384
9	5:30:15 PM	576	362879	3938391
10	5:45:13 PM	622	362924	3938333

Figure 8. Bear position B for Bear N024.

Detection No.	Time	Road Intersection / Landmark	E	N	Elevation (ft)	Bearing	Distance to Bear/Collar (code: 1 - 5)
1	1530	40 Ridge Rd	362810	3937442	2192	328	2
2	1537	Wood Ave/Parker Rd	362666	3937469	2110	40	2
4	1604	29 Violet Hill Cir	362963	3938656	2166	164	4
5	1618	317 White Pine Dr	361675	3938024	2114	63	4
6	1632	387 Swannanoa River Rd	362437	3938289	2019	70	5
7	1654	470 Swannanoa River Rd	363011	3938184	2022	332	6
8	1700	Bleachery Blvd/River Ford Pkwy	362957	3937614	2088	357	4
9	1710	2 Fernlawn Rd	362811	3938791	2190	194	3
10	1714	61 Crockett Ave	362814	3938574	2145	160	5
11	1723	61 Beechwood Rd	362701	3938584	2086	120	5
12	1726	Swannanoa River Rd/Governors View Rd	362849	3938351	2021	77	6
13	1730	26 Governors View Rd	362979	3938298	2025	324	6
14	1733	40 Governors View Rd	363115	3938343	2053	286	6

Figure 9. Readings for text location T and radio receiver R for Bear N024.

Table 3. Ranges in meters for each distance code.

Distance Code	Distance (m)
2	1500
3	1000
4	750
5	500
6	375

Since the sensors read out on different time scales, interleaving of the data records is required, as illustrated in Figure 10 for the first 18 integrated readings, encompassing 909 s. Figure 11 shows a similar record, now compacted and sampled at the minute interval, showing repeated entries for sensors across all time stamps. A map is shown in Figure 12, with the red dots showing the position of the bear's GPS collar, and the blue dots and lines showing the measured positions of the human, including V , T , and R . Finally, the time sequence of the consistency radius is shown in Figure 13, where data points are adorned with the addresses visited over time, on a general basis, along with a linear trendline.

Seconds	V.E	V.N	T.E	T.N	R.E	R.N	Theta.E	Theta.N	B.E	B.N
0	362806.2	3937443	362832	3937453	362810	3937442	1272.071	-794.88	NA	NA
13	NA	NA	NA	NA	NA	NA	NA	NA	362888	3938373
60	362806.4	3937444	NA	NA	NA	NA	NA	NA	NA	NA
120	362806.8	3937444	NA	NA	NA	NA	NA	NA	NA	NA
180	362808.7	3937442	NA	NA	NA	NA	NA	NA	NA	NA
240	362807.3	3937445	NA	NA	NA	NA	NA	NA	NA	NA
300	362810.3	3937374	NA	NA	NA	NA	NA	NA	NA	NA
360	362668.2	3937471	NA	NA	NA	NA	NA	NA	NA	NA
420	362668.2	3937470	362644.3	3937491	362666	3937469	1149.067	964.1813	NA	NA
480	362668.3	3937469	NA	NA	NA	NA	NA	NA	NA	NA
540	362671.8	3937471	NA	NA	NA	NA	NA	NA	NA	NA
600	362662.6	3937467	NA	NA	NA	NA	NA	NA	NA	NA
660	362387.3	3938033	NA	NA	NA	NA	NA	NA	NA	NA
720	362251.7	3938206	NA	NA	NA	NA	NA	NA	NA	NA
780	362013.6	3938668	NA	NA	NA	NA	NA	NA	NA	NA
840	361938.1	3938935	NA	NA	NA	NA	NA	NA	NA	NA
900	361930.9	3938958	NA	NA	NA	NA	NA	NA	NA	NA
909	NA	NA	NA	NA	NA	NA	NA	NA	362882	3938377

Figure 10. Time-integrated sensors streams. Variables V , T , R , and B as in the model, with Θ being the offset to the bear on the VHF receiver (see Equation (1)), $x.N$ being northing, and $x.E$ being easting in UTM.

Minutes	Text Position T	Vehicle GPS V (UTM, elevation (m))	VHF R (UTM, elevation (ft), bearing (degrees), range code (integer))	Collar B (UTM, elevation(m))
0	40 Ridge Rd	[3.62806240e+05 3.93744310e+06 6.91294806e+02]	[3.628100e+05 3.937442e+06 2.192000e+03 3.280000e+02 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
1	40 Ridge Rd	[3.62806433e+05 3.93744387e+06 6.91264326e+02]	[3.628100e+05 3.937442e+06 2.192000e+03 3.280000e+02 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
2	40 Ridge Rd	[3.62806801e+05 3.93744420e+06 6.89892709e+02]	[3.628100e+05 3.937442e+06 2.192000e+03 3.280000e+02 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
3	40 Ridge Rd	[3.62808675e+05 3.93744229e+06 6.89526945e+02]	[3.628100e+05 3.937442e+06 2.192000e+03 3.280000e+02 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
4	40 Ridge Rd	[3.62807350e+05 3.9374453e+06 6.88703975e+02]	[3.628100e+05 3.937442e+06 2.192000e+03 3.280000e+02 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
5	40 Ridge Rd	[3.62668335e+05 3.93747030e+06 6.63344306e+02]	[3.628100e+05 3.937442e+06 2.192000e+03 3.280000e+02 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
6	40 Ridge Rd	[3.62665880e+05 3.93746978e+06 6.63923433e+02]	[3.628100e+05 3.937442e+06 2.192000e+03 3.280000e+02 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
7	Wood Ave/Parker Rd	[3.62668155e+05 3.93747041e+06 6.62643258e+02]	[3.626660e+05 3.937469e+06 2.110000e+03 4.000000e+01 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
8	Wood Ave/Parker Rd	[3.62668316e+05 3.93746908e+06 6.62734699e+02]	[3.626660e+05 3.937469e+06 2.110000e+03 4.000000e+01 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
9	Wood Ave/Parker Rd	[3.62671791e+05 3.93747102e+06 6.62612777e+02]	[3.626660e+05 3.937469e+06 2.110000e+03 4.000000e+01 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
10	Wood Ave/Parker Rd	[3.62538154e+05 3.93784585e+06 6.40331626e+02]	[3.626660e+05 3.937469e+06 2.110000e+03 4.000000e+01 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
11	Wood Ave/Parker Rd	[3.62387309e+05 3.93803268e+06 6.38533285e+02]	[3.626660e+05 3.937469e+06 2.110000e+03 4.000000e+01 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
12	Wood Ave/Parker Rd	[3.61940908e+05 3.93858887e+06 6.48835650e+02]	[3.626660e+05 3.937469e+06 2.110000e+03 4.000000e+01 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]
13	Wood Ave/Parker Rd	[3.61962649e+05 3.93892781e+06 6.55663253e+02]	[3.626660e+05 3.937469e+06 2.110000e+03 4.000000e+01 2.000000e+00]	[3.628880e+05 3.938373e+06 5.770000e+02]

Figure 11. Combined data record for collar N024 interpolated to the minute.

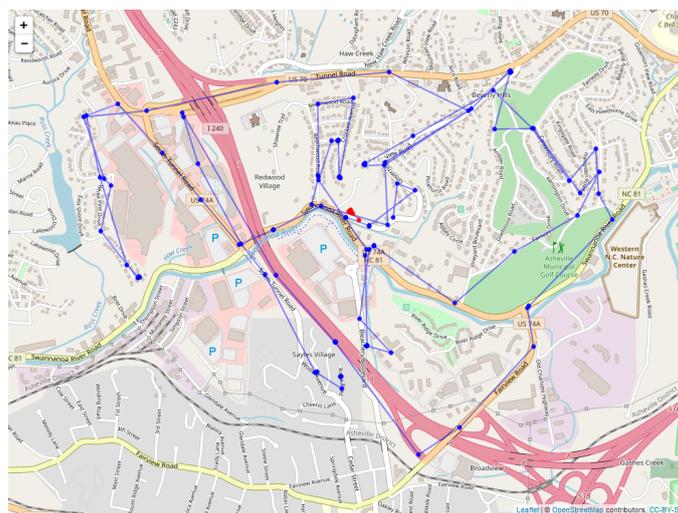


Figure 12. Map of sensor readings for bear N024.

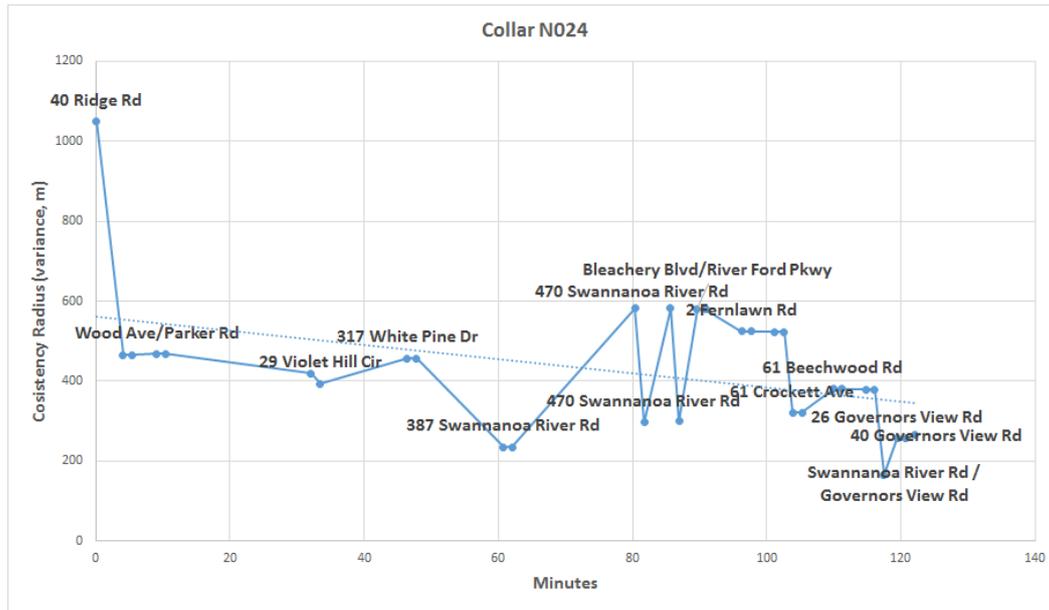


Figure 13. Consistency radius ϵ^* over time for bear N024.

5.3. Example: Minute 5.41

We now explicate the concept of a consistency filtration over a single set of measurements for N024 bear, specifically the point from minute 5.41, when the sensor readings shown on the left side of Figure 14 were recorded, as shown in the sheaf model. After this raw data is processed through the sheaf model functions shown in Figure 3, the resulting processed sensor readings are then shown on the right side of Figure 14, along with the resulting spread variance (meters) measures shown on non-vertex faces in blue.

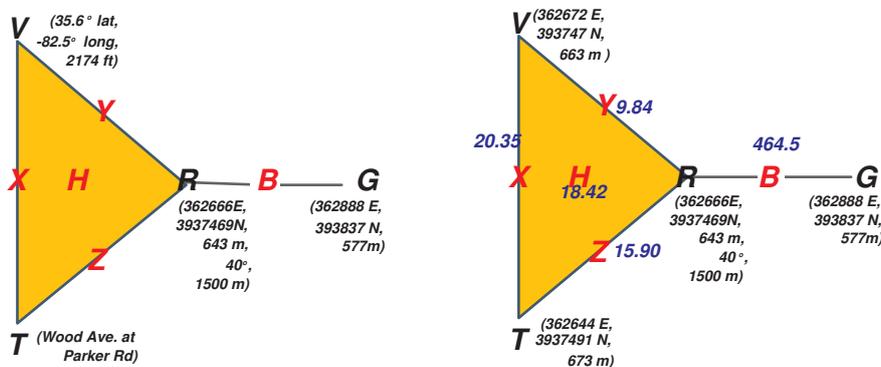


Figure 14. (Left) A raw assignment on collar N024. (Right) Processed data and spread measures (variance in m, in blue).

The consistency radius for this measurement is $\epsilon^* = 464.5$; this is the maximum variance recorded among the sensors, in particular between the receiver R and the GPS on the bear G . Our consistency parameter will range over $\epsilon \in [0, \epsilon^*] = [0, 464.5]$ over a series of steps, or landmark values, $1 \leq i \leq \ell$ of the filtration. At each step the following occur:

- There is a landmark non-zero consistency value $0 < \epsilon_i \leq \epsilon^*$ which does not exceed the consistency radius;
- There is a prior set of consistent faces Γ_{i-1} ;
- A new consistent face $\gamma \in \Delta$ is added so that $\Gamma_i = \Gamma_{i-1} \cup \{\gamma\}$;
- There is a corresponding vertex cover Λ_i , which is a coarsening of the prior Λ_{i-1} ;
- And which has a cover measure $\bar{r}(\Lambda)$.

These steps for our example are summarized in Table 4, and explained in detail below.

$\epsilon_0 = 0$: If we insist that no error be tolerated, that is that all data be consistent, then any nontrivial set in the vertex cover, produced by Theorem 2, cannot contain nontrivial faces of Δ . As such, the set of consistent faces are just the singletons $\Gamma_0 = \{\{V\}, \{R\}, \{T\}, \{G\}\}$, and the vertex cover is $\Lambda_0 = \{\{T, G\}, \{V, G\}, \{R\}\}$ with $\bar{r}(\Lambda_0) = 2/11$.

$\epsilon_1 = 9.48$: If we relax our error threshold to the next landmark value, while still well below our consistency radius, the readings on V and R are considered consistent within this tolerance, so that the face $Y = \{V, R\}$ is added, yielding the new set of consistent faces

$$\Gamma_1 = \Gamma_0 \cup \{Y\} = \{\{V\}, \{R\}, \{T\}, \{G\}, \{V, R\}\}.$$

The new vertex cover is $\Lambda_1 = \{\{V, R\}, \{V, G\}, \{G, T\}\}$, with cover measure $\bar{r}(\Lambda_1) = 3/11$.

$\epsilon_2 = 15.9$: Continuing on, next T and R come into consistency, adding the face $Z = \{R, T\}$, yielding

$$\Gamma_2 = \{\{V\}, \{R\}, \{T\}, \{G\}, \{V, R\}, \{T, R\}\},$$

$$\Lambda_2 = \{\{V, R\}, \{V, G\}, \{G, T\}, \{R, T\}\}, \quad \bar{r}(\Lambda_2) = 4/11.$$

$\epsilon_3 = 18.42$: The next landmark introduces the three-way interaction $H = \{V, T, R\}$ (for notational simplicity just note that $\Gamma_3 = \Gamma_2 \cup \{H\}$). However, the vertex cover is unchanged, yielding $\Lambda_3 = \Lambda_2$ and $\bar{r}(\Lambda_3) = \bar{r}(\Lambda_2) = 4/11$.

$\epsilon_4 = 20.35$: Next T and V are reconciled, adding $X = \{V, T\}$ to Γ_4 . Now $\Lambda_4 = \{\{V, T, R\}, \{V, T, G\}\}$, with $\bar{r}(\Lambda_4) = 7/11$.

$\epsilon_5 = \epsilon^* = 464.5$: Finally we arrive at our consistency radius with the bear collar G being reconciled to R adding the face $\{B, G\}$ to Γ_5 . Our vertex cover is naturally now the coarsest, i.e., just the set of vertices $\Lambda_5 = \{\{V, T, R, G\}\}$ as a whole, with $\bar{r}(\Lambda_5) = 1$.

These results are summarize in Table 4, while Figure 15 shows a plot of the cover coarseness as a function of consistency, and adorned by the *combination* of sensors engaged at each landmark. It is evident that the most discrepancy, by far, among sensors comes from the Radio VHF Device and Bear Collar when reporting on the bear’s location.

Table 4. Consistency filtration of our example assignment.

ϵ	New Consistent Face	Vertex Cover	Measure
0.00		$\{\{T, G\}, \{V, G\}, \{R\}\}$	2/11
9.48	$Y = \{V, R\}$	$\{\{V, R\}, \{V, G\}, \{G, T\}\}$	3/11
15.90	$Z = \{R, T\}$	$\{\{V, R\}, \{V, G\}, \{G, T\}, \{R, T\}\}$	4/11
18.42	$H = \{V, T, R\}$	$\{\{V, R\}, \{V, G\}, \{G, T\}, \{R, T\}\}$	4/11
20.35	$X = \{V, T\}$	$\{\{V, T, R\}, \{V, T, G\}\}$	7/11
464.50	$B = \{R, G\}$	$\{\{V, T, R, G\}\}$	1

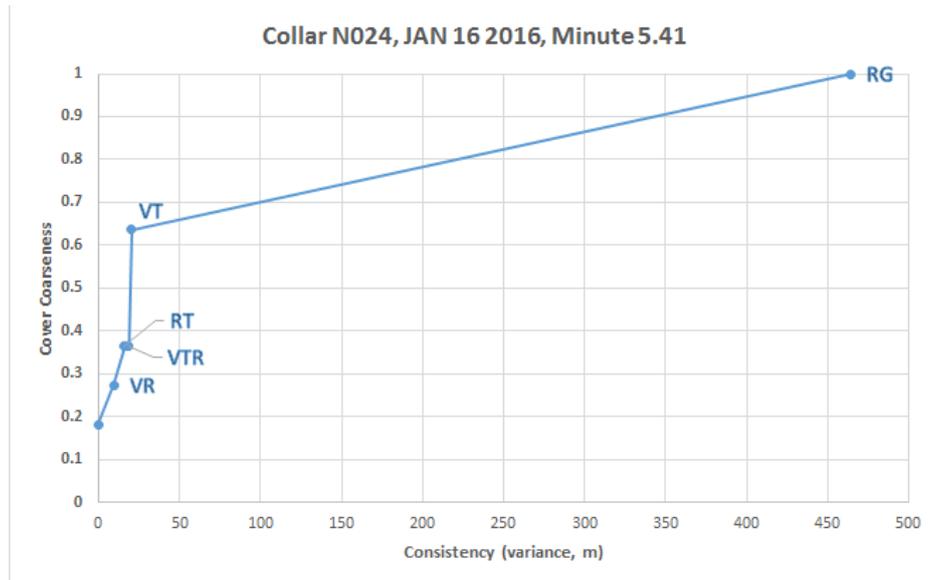


Figure 15. (Left) A raw assignment on collar N024. (Consistency vs. cover coarseness for example assignment.

Figure 15 shows the filtration for a single time stamp, $t = 5.41$, with each of the five sensor combinations possible in the simplicial complex appearing at a certain point in the filtration, with, in particular, the *RG* interaction for the bear sensors highlighted. Figure 16 shows these across all time steps for bear N024. We can now see that *RG* isn't always the least consistent sensor readings, with, at times, *VR* and *VT* also leaping to the top position over several time intervals.

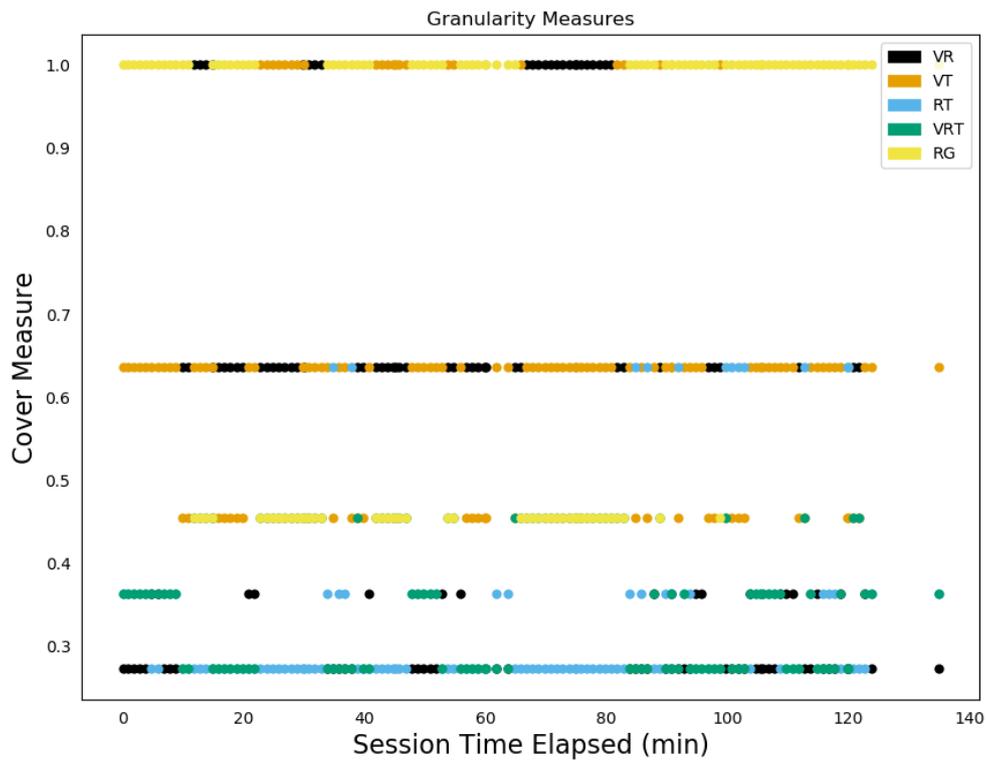


Figure 16. Consistency filtration values for entire N024 trajectory.

6. Comparison Statistical Model

Sheaf modeling generally, and especially as introduced in this paper, is a novel approach to multi-sensor fusion and target tracking. As such, it behooves us to consider alternate, or standard, approaches which analysts may bring to bear on the question of both predicting locations and quantifying uncertainty around them. For these purposes, we have also performed in parallel a statistical model of the bear tracking data. Specifically, a dynamic linear model (DLM) is introduced, whose parameters are then estimated using a Kalman filter. A comparison of these results with the sheaf model is instructive as to the role that sheaf modeling can play in the analytical landscape as *generic* integration models.

Although we perform a comparison between these two approaches, we suggest to the reader that these two approaches are complementary. Consistency radius and filtration do not and cannot tell what the measurements are, nor does it report which measurements are more trustworthy. So while we show that comparison with a Kalman filter's covariance estimates is apt, comparing consistency radius with the Kalman filter's position estimates is not appropriate. The Kalman filter provides an additional aggregation stage to the process, but one that also brings additional assumptions along with it. If those assumptions are not valid, the process may produce faulty results.

A DLM is a natural statistical approach for representing the bear/human tracking problem. A DLM separately models the movements of the bear, movements of the human, and the multiple observations on each. These model components are combined based on a conditional independence structure that is standard for hidden Markov models [74]. Once stochastic models of the movements and of the observations are obtained, then multiple standard computational approaches are available to use the model and the data to:

1. Estimate the locations, with corresponding uncertainties, over time of the bear collar and the human (see Figures 18 and 19 below).
2. Estimate accuracy parameters of the involved sensors based on each of the single runs of the experiment (see Table 5 below).
3. Combine information across the multiple experimental runs to estimate the accuracy of the sensors (see Equation (5) below).

The Kalman Filter (KF) [75,76] is a standard computational approach for estimation with a DLM from which the above three items can be computed. Monte Carlo approaches for estimating parameters in DLMs are also available [77]. In addition, this type of model is also called State Space modeling [75], and Data Augmentation - depending on the technical community. Similar tracking approaches are given in numerous references [3,78,79].

6.1. High-level state and observation equations - with examples

The two sets of equations below are a schematic for the DLM that supports bear and human tracking. First - the state equations that represent the bear and human locations are:

$$\begin{aligned}\text{bear}_{t+\Delta t} &= \text{bear}_t + \omega_{b,\Delta t,t+\Delta t} \\ \text{human}_{t+\Delta t} &= \text{human}_t + \omega_{h,\Delta t,t+\Delta t}\end{aligned}$$

The state equations represent both the locations and the movements of the bear and the human over time. In this particular set of state equations the locations of each of the bear and human are modeled as random walks that are sampled at non-equispaced times. The fact that Δt here is not constant is an important sampling feature: the measurements are not all taken at equal time intervals. The parameters to be set, and/or estimated are embedded in the two ω terms. These terms are stochastically independent Gaussian random variables. The covariance matrices for the ω 's are parameters to be set or estimated. Intuitively, a larger variance corresponds with potentially more rapid movement. Finally, these state equations are motivated by those used for general smoothing in [80].

The schematic observation equations are:

$$\begin{aligned}\text{truck gps}_t &= \text{human}_t + \epsilon_{\text{truck gps},t} \\ \text{vhf gps}_t &= \text{human}_t + \epsilon_{\text{vhf gps},t} \\ \text{streetsign}_t &= \text{human}_t + \epsilon_{\text{streetsign},t} \\ \text{bear gps}_t &= \text{bear}_t + \epsilon_{\text{bear gps},t} \\ \text{vhf}_t &= \text{human}_t - \text{bear}_t + \epsilon_{\text{vhf},t}\end{aligned}$$

These equations connect directly with the state equations, and represent a full set of observations at time t . If, as is typical for this data (see Figure 10), there is not a complete set of information at a particular time, then the observation equations are reduced to reflect only those measurements that have been made. The KF calculations for estimating bear/human locations along with sensor accuracies can proceed with that irregularly sampled data. The computations for collar N024 are in two dimensions (see Figures 18 and 19 below), but for conciseness the equations shown for the KF are in one dimension.

A general form for a DLM that shows the update from time $t - \Delta t$ to time t is:

$$\theta_t = \mathbf{G}_{\Delta t} \theta_{t-\Delta t} + \omega_t \quad (2)$$

$$\mathbf{y}_t = \mathbf{H}_t \theta_t + \epsilon_t \quad t = 1, 2, 3, \dots \quad (3)$$

where $\mathbf{G}_{\Delta t}$ is an identity matrix. The base state model is a random walk – and so the “best guess” is centered on the previous location. The covariance of ω_t is structured as a diagonal matrix with diagonals proportional to the elapsed time Δt . The parameters that are needed to be chosen to calculate a DLM are the variances of the ‘impetus’ ω_t (which depends on Δt) and the variances of the observation errors ϵ_t . The model assumes that the means and correlations of these are zero. The one-dimensional bear equations map to this abstract representation as:

$$\theta_t = \begin{pmatrix} \text{bear}_t \\ \text{human}_t \end{pmatrix}$$

$$\mathbf{y}_t = \begin{pmatrix} \text{truck gps}_t \\ \text{vhf gps}_t \\ \text{streetsign}_t \\ \text{bear gps}_t \\ \text{vhf}_t \end{pmatrix}$$

with corresponding expansions for ω_t and ϵ_t .

The data are structured so that it rarely if ever occurs that all 5 measurements are taken at the same time. For a particular time at which an observation is available the observation equations are subsetted down to the observation or observations that are available. This can be seen in the ‘time interweave’ plot for one of the experiment runs. This is represented in the general form of the DLM below (Equations (2) and (3)) by varying the matrix \mathbf{H}_t with corresponding variations in ϵ_t . For complete observations at time t the matrix \mathbf{H}_t and ϵ_t would respectively be:

$$\mathbf{H}_t = \begin{pmatrix} 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 1 & 0 \\ -1 & 1 \end{pmatrix}, \quad \epsilon_t = \begin{pmatrix} \epsilon_{\text{truck gps},t} \\ \epsilon_{\text{vhf gps},t} \\ \epsilon_{\text{streetsign},t} \\ \epsilon_{\text{bear gps},t} \\ \epsilon_{\text{vhf},t} \end{pmatrix} \quad (4)$$

The variance parameters are:

$$\sigma = \begin{pmatrix} \sigma_{\text{truck gps}} \\ \sigma_{\text{vhf gps}} \\ \sigma_{\text{streetsign}} \\ \sigma_{\text{bear gps}} \\ \sigma_{\text{vhf}} \end{pmatrix}$$

Correspondingly, if only the VHF data are available then \mathbf{H}_t and ϵ_t would respectively be:

$$\begin{pmatrix} -1 & 1 \end{pmatrix} \begin{pmatrix} \epsilon_{\text{vhf},t} \end{pmatrix},$$

and if only the bear collar GPS and truck GPS data are available then \mathbf{H}_t and ϵ_t would respectively be:

$$\begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} \epsilon_{\text{bear gps},t} \\ \epsilon_{\text{vhf},t} \end{pmatrix}.$$

6.2. Estimation of parameters

The parameters are estimated via the Kalman filter. That calculation is representable as a Bayesian update as follows. At time $t - 1$ the knowledge of the state parameter, informed by all the data up to and including that time, is represented as:

$$\hat{\theta}_{t-1} | \mathbf{y}_{t-1} \sim N(\hat{\theta}_{t-1}, \Sigma_{t-1})$$

This equation is read as ‘the distribution of $\hat{\theta}_{t-1}$ given the data \mathbf{y}_{t-1} is Gaussian with a mean of $\hat{\theta}_{t-1}$ and a covariance matrix Σ_{t-1} ’. The Kalman filter is typically written in two steps: a forecast followed by an update. With $\mathbf{e}_t = \mathbf{y}_t - \mathbf{H}_t \mathbf{G}_{\Delta t} \hat{\theta}_{t-1}$, \mathbf{V}_t the covariance of ϵ_t , and \mathbf{W}_t the covariance of ω_t , then the forecast step is:

$$\hat{\theta}_t | \mathbf{y}_{t-1} \sim N(\mathbf{G}_{\Delta t} \hat{\theta}_{t-1}, \mathbf{R}_t = \mathbf{G}_{\Delta t} \Sigma_{t-1} \mathbf{G}'_{\Delta t} + \mathbf{W}_t)$$

and the update step is:

$$\hat{\theta}_t | \mathbf{y}_t \sim N(\mathbf{G}_{\Delta t} \hat{\theta}_{t-1} + \mathbf{R}_t \mathbf{H}'_t (\mathbf{V}_t + \mathbf{H}_t \mathbf{R}_t \mathbf{H}'_t)^{-1} \mathbf{e}_t, \mathbf{R}_t - \mathbf{R}_t \mathbf{H}'_t (\mathbf{V}_t + \mathbf{H}_t \mathbf{R}_t \mathbf{H}'_t)^{-1} \mathbf{H}_t \mathbf{R}_t)$$

The above steps are executed sequentially by passing through the data items in time order.

The parameters in the DLM may be chosen as values which maximize the likelihood function. The set of variance parameters includes those for the five observation forms and the state equations — we denote this collection by σ . The values σ can be estimated from the data using as the values that maximize the *likelihood function*. The likelihood function is

$$L(\mathbf{y}_1, \dots, \mathbf{y}_n | \sigma) = \prod_{i=1}^n p(\mathbf{y}_i | \sigma, \mathbf{y}_{i-1}, \dots, \mathbf{y}_1)$$

The terms in the product are available as a byproduct of a KF calculation. Table 5 shows the estimated parameters for collar N024. These were estimated from the collar data, using the KF computation and the maximum likelihood approach. As expected, the estimated parameter for the VHF antenna measurement is larger than the other observation parameters. Also, the human (in the vehicle) is seen to move more quickly than the bear. The table shows the parameters as standard deviations so that the units are interpretable.

The KF structure supports pooling the models and measurements across the 12 experiments. This has the potential to provide more informed estimates of the sensor errors (as represented by the variances) compared with the single experiment based estimates shown in Table 5. The information

from these series can be pooled to arrive at estimates of σ under the assumption that the variances that correspond with the five observations are the same across experiments. In that case the pooled likelihood function is:

$$\prod_{j=1}^{12} L_j(\mathbf{y}_j \parallel \sigma) \quad (5)$$

In addition, the estimates of σ can be estimated as those which maximize the pooled likelihood. Another case is to take advantage of the series corresponding with the cases where the bear collar was stationary by setting that variance component to zero in the state update. The remaining variances can be estimated for those series with that assumption — essentially reducing the dimensionality of the optimization problem.

6.3. Example outputs

Figure 17 is a map showing output for the DLM for the human position plotted against the measured position, with the KF estimated human positions shown in red, and the direct measurements of human position shown in blue. Figure 18 then shows example output for the location estimates of the bear, and Figure 19 for the human. These plots also show estimated 95% confidence intervals for the locations. Human position refers to the estimated human position from the sensors at specific times. Both the sheaf and the KF output location estimates for the bear and the human. The x-axis is time(in seconds) from the first observation. The estimates for bear locations are shown as green circles for both the 'east' and 'north' coordinates. The location units are UTM on the y-axis. The observed locations for the bear collar are from the GPS measurements. Two standard deviation intervals that were estimated are shown as gray lines. The implementation of the KF that we used was set to output a location estimate plus uncertainty for each time that an observation in the overall system was made. The jumps in the estimated values are seen to primarily shift when a bear collar GPS observation is made. Jumps seen elsewhere correspond with large shifts in the related data - which in this case includes the VHF measurement and its link with the human's position.

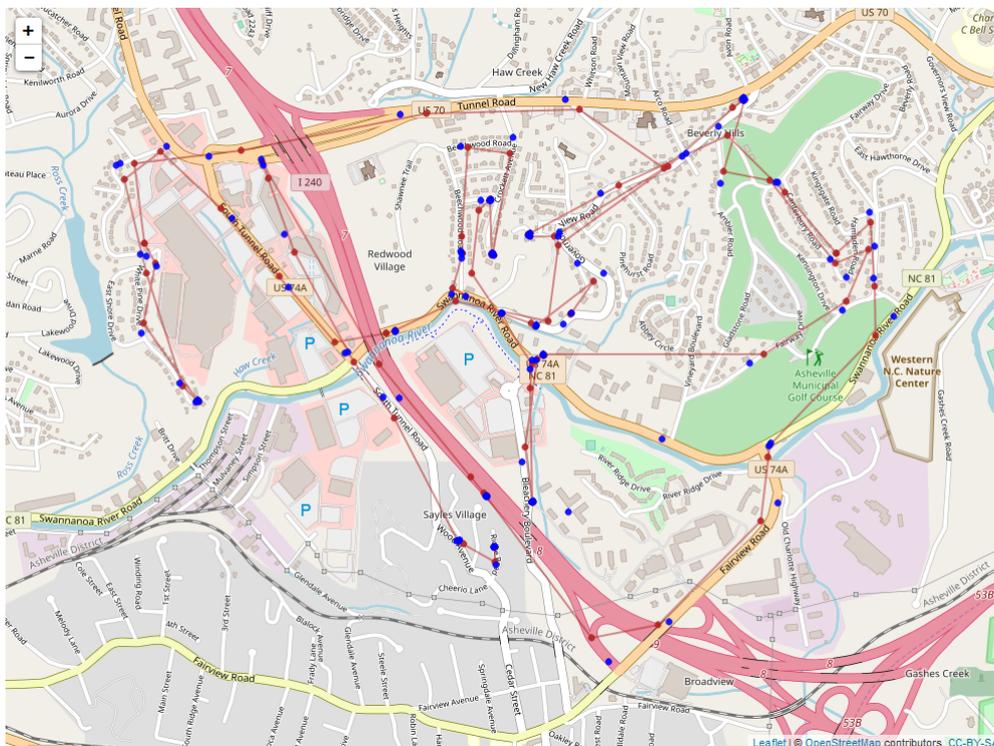


Figure 17. Output for the DLM for the human position (red) plotted against the measured position (blue).

Table 5. Estimated KF parameters for collar N024. The first two rows are the standard deviation estimates for the ω parameters. The remaining rows contain standard deviation estimates for the ϵ parameters.

Parameter	Standard Deviation	Value (m)
Bear state update		0.008
Human state update		26.524
Bear GPS Obs.		16.091
VHF GPS Obs.		0.032
Vehicle GPS Obs.		91.434
Street sign Obs.		27.597
VHF Obs.		663.998

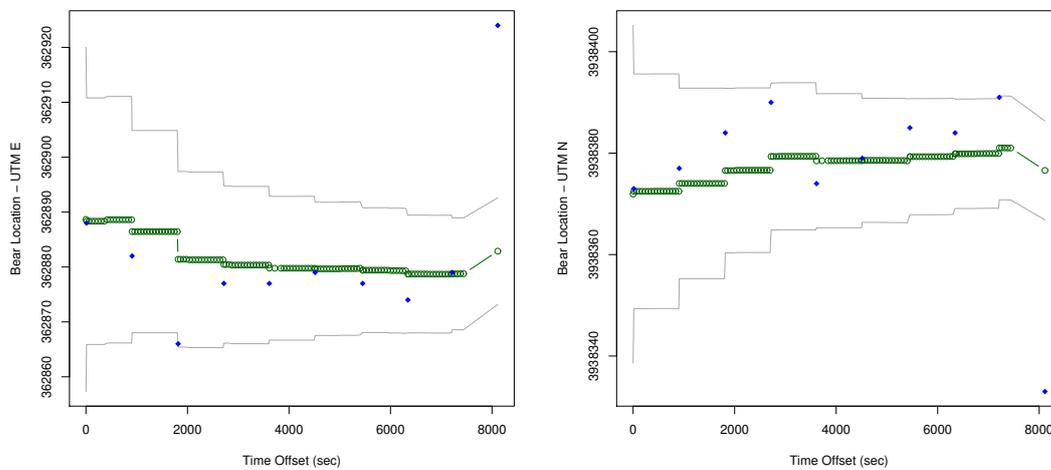


Figure 18. KF Estimates of locations of the bear for N024. (Estimates of location are shown in green, the GPS data is shown in blue, and the confidence intervals are shown in gray.)

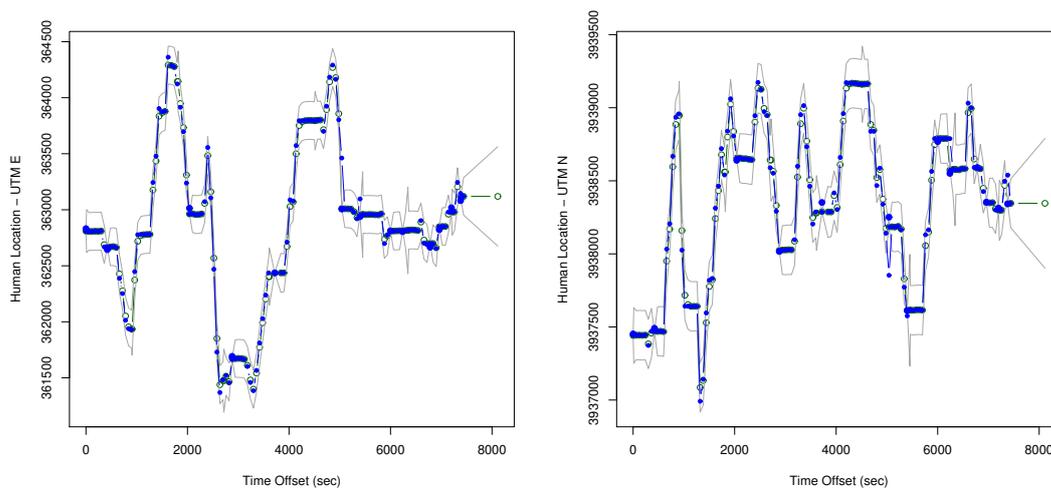


Figure 19. KF Estimates of locations of the human for N024. (Estimates of location are shown in green, the GPS data is shown in blue, and the confidence intervals are shown in gray.)

The estimates for human location are also shown in green, gps data for the human (truck gps) in blue, and confidence intervals in gray. Again - there are more data than shown in the plot, so the jumps

in the estimated locations that do not directly correspond with data points on the plots are driven by other data in the data set.

6.4. Comparisons

The *processed* sheaf model and the dynamic linear model of our example capture the same phenomenon, and even do so in relatively similar ways. Because of the canonicity of the sheaf modeling approach, any other sensor integration model must recapitulate aspects of the sheaf model. After loading processed data into the sheaf, the state space that governs the dynamic linear model is precisely the space of global sections of the sheaf model. If we do not register the text reports or the angular measurements from the VHF receiver into a common coordinate system, then while the sheaf model is still valid, the data are no longer related linearly. Without this preprocessing, the dynamic linear model is not valid, and we conclude that the sheaf carries more information. Therefore, all of the following comparisons are made considering only observations registered into a common coordinate system.

The observation equations of the dynamic linear model (represented by the matrix \mathbf{H}_t) can therefore be “read off” the restriction maps of the sheaf. The two models differ more in choices made by the modelers: (1) the sheaf model treats uncertainty parametrically, while the dynamic linear model does so stochastically, (2) the sheaf model processes time slices independently, without a dynamical model. We address these two differences in some detail after first addressing their similar treatment of the state space and observational errors.

While the space of assignments for the sheaf \mathcal{S} defined in our example is quite large, the space of global sections is comparatively small. It is parameterized by the four real numbers, specifically the east and north positions of the human and the bear. This is precisely the space of state variables θ_t at each timestep. On the other hand, the y_t matrix of observations consists of values within an assignment of the vertices. The observation matrix \mathbf{H}_t is an aggregation of all of the corresponding restriction maps, deriving the values of a global section at each vertex. The observational errors ϵ_t are then the differences between the observations we actually made (an assignment), and those predicted from the state variables (a global section). Thus the consistency radius is merely the largest of the components of ϵ_t .

Figure 20 is a map showing sheaf bear location estimates in purple and the DLM and bear location estimates in red. Figure 21 shows the human sheaf location estimates in purple and the measured locations for humans in blue. The most significant difference between the proposed sheaf-based model and dynamic linear model of sensor-to-sensor relationships is their treatment of uncertainty. The dynamic linear model assumes that the data are subject to a stochastic model, which determine independent Gaussian random variables for the positions of the bear and human. The sheaf model posits a deterministic model for these independent variables, but permits (and quantifies) violations of this model. The two models are in structural agreement on how to relate groups of non-independent variables. Both models assert exact agreement of the three observations of the human position, and that the VHF observations ought to be the vector difference between bear and human position. Specifically, the observation matrix \mathbf{H}_t (Equation (4)) and the sensor matrix determining the base space for the sheaf (Table 1) are identical up to sign (the difference in signs is due to differences in the choice of basis only). In particular, the vector difference for relating the VHF observations to bear and human positions is encoded in the restriction map $\mathcal{S}(R \rightsquigarrow B)$ shown in the right frame of Figure 3 and also in the last row of the \mathbf{H}_t matrix.

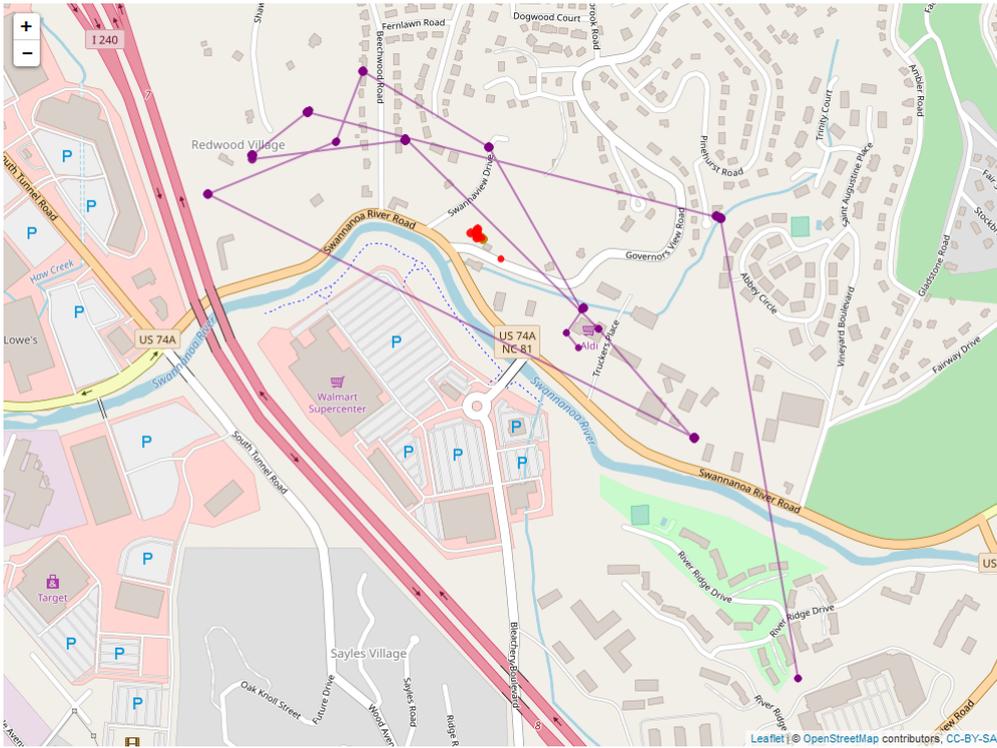


Figure 20. Comparison of the DLM (red) and sheaf model (purple) location estimates for the bear.

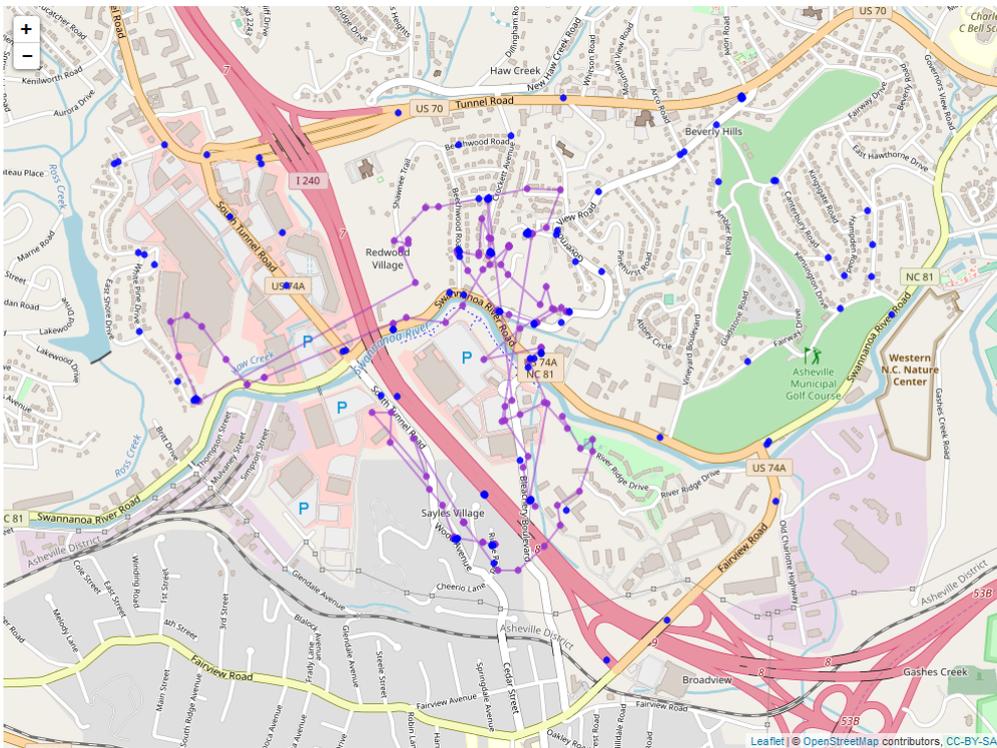


Figure 21. Comparison of the sheaf (purple) and measured data (blue) for the human location.

While the two models treat relationships between variables identically, we have chosen to treat the individual sensors differently with respect to time. The sheaf model is not intended to estimate sensor *self* consistency directly. Sensor self consistency can only be estimated in the sheaf by way of comparison against other sensors. In contrast, the dynamic linear model works through time

(Equations (2) and (3)) to estimate self consistency from the sensor's own time series. Indeed, the dynamic linear model separates observational errors ϵ_t from those arising from impetus ω_t , which can be thought of as deviations from the baseline, and can estimate relative sensor precisions. It should be noted that this distinction could be made using a sheaf by explicitly topologizing time, though this was not the focus of our study.

Considering the specific case of Collar N024, the difference between the VHF-estimated and GPS-estimated bear positions determines the consistency radius. This is shown by the presence of RG being the largest value in Figure 16 for most time values. Because of this, baseline comparisons of the observational error the VHF sensor computed by the dynamic linear model in Table 5 broadly agrees with the typical consistency radius shown in Figure 13.

Focusing on the specific time discussed in Section 5.3, namely $t = 5.41 \text{ min} = 325 \text{ s}$ as shown in Figure 15, it is clear from the sheaf analysis that the majority of the error is due to the difference between the VHF observation and the bear collar GPS. These inferences are not easily drawn from the dynamic linear model output in Figures 18 and 19, but the difference between the quality of these sensors is shown in Table 5.

Delving more finely into the timeseries, the consistency radius for the bear-human system (Figure 13) shows a steady decrease as the human approaches the bear, leading to improved VHF readings. This is also visible in Figure 18, in which the error variance steadily decreases for the bear location estimates.

The three sensors on the human show more subtle variations, as is clear from the orange curve on the left side of Figures 5 and 19. In both, there is a marked increase in human error near $t = 75 \text{ min} = 4500 \text{ s}$. Consulting the consistency filtration in Figure 15, this is due primarily to the consistency radius being determined by the edge VR in the base space of the sheaf—a difference between the GPS receiver on the vehicle and the GPS receiver attached to the VHF receiver. Shorter occurrences of the same phenomenon occur near times $t = 15 \text{ min} = 900 \text{ s}$ and $t = 32 \text{ min} = 1920 \text{ s}$.

The dynamic linear model posits that changes in the state are governed by an impetus vector ω_t , while our sheaf model is essentially not dynamic. While it may seem that these two models are at variance with one another, they are easily rectified with one another. Considering only the global sections of \mathcal{S} , which is the state space of the dynamic linear model, we build a sheaf over a line graph, namely in which I represents the identity matrix. Global sections of this sheaf consist of timeseries of the internal state θ_t with no impetus, while an assignment to the vertices corresponds to a timeseries of internal state *with* impetus. Thus again, the consistency radius measures the maximum impetus present.

7. Conclusions

Sheaves are a mathematical data structure that can assimilate heterogeneous data types, making them ideal for modeling a variety of sensor feeds. Sheaves over abstract simplicial complexes provide a global picture of the sensor network highlighting multi-way interactions between the sensors. As a consequence of the foundational nature of sheaves, we know that *any* multi-sensor fusion method can, in principle, be encoded as a sheaf. Additionally, the theory provides algorithms for assessing measurement consistencies. In addition, in this article, we demonstrated how one might use sheaves for tracking wildlife with a collection of sensors collecting types of information. We saw that sheaves can be used to produce a holistic temporal picture of the observations, and can describe the levels of agreement between different groups of sensors.

Since our approach is quite general, it could be applied to many other use cases. Perhaps the most obvious application is that of locating emergency beacons from downed aircraft [2] or other hidden radio transmitters [81]. More broadly, the idea of consistency radius can be valuable in combining disparate biochemical networks [82], analyzing the convergence of graphical models and numerical differential equation solvers [42,45], and estimating network flows [58,59,83].

The generality and expressivity of a sheaf-based approach, therefore, is to be valued, albeit in the context of consideration of its costs, potential limitations, and comparison with other methods. For example, when we compared our sheaf model with standard tracking algorithms, the sheaf model was able to provide error measurements without any upfront parameter estimation. For instance, sheaves and dynamic linear modeling encode dependence between observations starting from the same information — a relation between sensors and state variables. While broadly in quantitative agreement on our example data, these two methods format this information differently. These differences result in certain inferences being easier to make in one or the other framework. For instance, because dynamic linear models require parameter estimation, sensor self-consistency is easier to compute as a consequence. Conversely, since the sheaf model does not require parameter estimation, this simplifies the process of making fine-grained inferences about which groups of sensors provide consistent observations.

So while sheaf models obviate the need for up-front uncertainty quantification, on the other hand they carry the burden of up-front model construction and complexity. In particular, the specific sensor architecture, including detailed knowledge about pairwise sensor interactions, need to be encoded in the abstract simplicial complex. Additionally, the ability to model all of the resulting complex interactions itself carries an additional computational burden: while the initial model setup requires specification of *pairwise* sensor interactions, the resulting sheaf model calculates *all* multi-way interactions through the abstract simplicial complex.

Finally, it should be noted that the sheaf-based approach can be extended to handle targets that move at higher speeds. For instance, [50] demonstrates that the same recipe for constructing a sheaf as presented here (with different sensors: passive RF sensors and optical cameras) can be used to track moving vehicles under tree cover. Their resulting sheaf-based tracker smoothly handles sensor hand-off and maintains track custody even when the target is occluded.

Author Contributions: Conceptualization, C.A.J., L.C., E.P. and M.R.; Data curation, K.N. and P.W.; Formal analysis, C.A.J., K.N., B.P., E.P., M.R. and P.W.; Funding acquisition, C.A.J. and C.D.; Investigation, N.G. and J.S.; Methodology, C.A.J., C.D., N.G., K.N., B.P., E.P., M.R. and P.W.; Project administration, C.A.J., C.D. and N.G.; Resources, L.C. and C.D.; Software, B.P. and P.W.; Supervision, C.A.J., L.C. and C.D.; Validation, C.A.J., K.N. and M. R.; Visualization, C.A.J., K.N. and P.W.; Writing—original draft, C.A.J., K.N., B.P., E.P., M.R. and P.W.; Writing—review & editing, C.A.J., L.C., C.D., N.G., K.N., B.P., E.P., M.R., J.S. and P.W. All authors have read and agreed to the published version of the manuscript.

Funding: This work was partially funded under the High Performance Data Analytics (HPDA) program at the Department of Energy's Pacific Northwest National Laboratory. Pacific Northwest National Laboratory is operated by Battelle Memorial Institute under Contract DE-ACO6-76RL01830. This work was partially funded by the Pittman Robertson Federal Aid to Wildlife Restoration Grant and was a joint research project between the North Carolina Wildlife Resources Commission (NCWRC) and the Fisheries, Wildlife, and Conservation Biology (FWCB) Program at North Carolina State University (NCSU). We thank the homeowners who granted us permission and access to their properties. We thank numerous other staff from the NCWRC and the FWCB program at NCSU for their ongoing assistance and support.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

References

1. Robinson, M. *Topological Signal Processing*; Springer-Verlag: Berlin/Heidelberg, Germany, 2014.
2. Robinson, M. Sheaves are the canonical datastructure for information integration. *Inf. Fusion* **2017**, *36*, 208–224.
3. Luo, W.; Xing, J.; Zhang, X.; Zhao, X.; Kim, T.K. Multiple Object Tracking: A Literature Review. *arXiv* **2014**, arXiv:1409.7618.
4. Pulford, G. Taxonomy of multiple target tracking methods. *IEE Proc. Radar Sonar Navig.* **2005**, *152*, 291–304.
5. DeGroot, M.H. *Optimal Statistical Decisions*; Wiley: Hoboken, NJ, USA, 2004.
6. Deming, R.; Schindler, J.; Perlovsky, L. Multi-target/multi-sensor tracking using only range and doppler measurements. *IEEE Trans. Aerosp. Electron. Syst.* **2009**, *45*, 593–611.

7. Efe, M.; Ruan, Y.; Willett, P. Probabilistic multi-hypothesis tracker: addressing some basic issues. *IEE Proc. Radar Sonar Navig.* **2004**, *151*, 189–196.
8. Hamid Rezaatofghi, S.; Milan, A.; Zhang, Z.; Shi, Q.; Dick, A.; Reid, I. Joint Probabilistic Data Association Revisited. In Proceedings of the IEEE International Conference on Computer Vision, Santiago, Chile, 7–13 December 2015; pp. 3047–3055.
9. Berclaz, J.; Fleuret, F.; Turetken, E.; Fua, P. Multiple Object Tracking Using K-Shortest Paths Optimization. *IEEE Trans. Pattern Anal. Mach. Intell.* **2011**, *33*, 1806–1819. doi:10.1109/TPAMI.2011.21.
10. Butt, A.; Collins, R. Multi-target tracking by lagrangian relaxation to min-cost network flow. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Portland, OR, USA, 23–28 June 2013; pp. 1846–1853.
11. Pirsiaavash, H.; Ramanan, D.; Fowlkes, C.C. Globally-optimal greedy algorithms for tracking a variable number of objects. In Proceedings of the Computer Vision and Pattern Recognition (CVPR), Providence, RI, USA, 20–25 June 2011; pp. 1201–1208.
12. Perera, A.A.; Srinivas, C.; Hoogs, A.; Brooksby, G.; Hu, W. Multi-object tracking through simultaneous long occlusions and split-merge conditions. In Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, New York, NY, USA, 17–22 June 2006; pp. 666–673.
13. Zhang, L.; Li, Y.; Nevatia, R. Global data association for multi-object tracking using network flows. In Proceedings of the 2008 IEEE Conference on Computer Vision and Pattern Recognition, Anchorage, AK, USA, 23–28 June 2008; pp. 1–8.
14. Milan, A.; Gade, R.; Dick, A.; Moeslund, T.B.; Reid, I. Improving Global Multi-target Tracking with Local Updates. In *Computer Vision - ECCV 2014 Workshops: Zurich, Switzerland, September 6–7 and 12, 2014, Proceedings, Part III*; Agapito, L., Bronstein, M.M., Rother, C., Eds.; Springer International Publishing: Cham, Switzerland, 2015; pp. 174–190. doi:10.1007/978-3-319-16199-0_13.
15. Wang, B.; Wang, G.; Chan, K.; Wang, L. Tracklet association with online target-specific metric learning. In Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 1234–1241.
16. Ben Shitrit, H.; Berclaz, J.; Fleuret, F.; Fua, P. Multi-commodity network flow for tracking multiple people. *IEEE Trans. Pattern Anal. Mach. Intell.* **2014**, *36*, 1614–1627.
17. Wald, L. Some terms of reference in data fusion. *IEEE Trans. Geosci. Remote Sens.* **1999**, *37*, 1190–1193.
18. Yang, H.; Shao, L.; Zheng, F.; Wang, L.; Song, Z. Recent advances and trends in visual tracking: A review. *Neurocomputing* **2011**, *74*, 3823–3831.
19. Leal-Taixé, L.; Milan, A.; Reid, I.; Roth, S.; Schindler, K. Motchallenge 2015: Towards a benchmark for multi-target tracking. *arXiv* **2015**, arXiv:1504.01942.
20. Solera, F.; Calderara, S.; Cucchiara, R. Towards the evaluation of reproducible robustness in tracking-by-detection. In Proceedings of the 2015 12th IEEE International Conference on Advanced Video and Signal Based Surveillance, Karlsruhe, Germany, 25–28 August 2015; pp. 1–6. doi:10.1109/AVSS.2015.7301755.
21. Bailer, C.; Pagani, A.; Stricker, D. A Superior Tracking Approach: Building a Strong Tracker through Fusion. In *Computer Vision – ECCV 2014: 13th European Conference, Zurich, Switzerland, 6–12 September 2014, Proceedings, Part VII*; Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T., Eds.; Springer International Publishing: Cham, Switzerland, 2014; pp. 170–185. doi:10.1007/978-3-319-10584-0_12.
22. Hall, D.L.; McMullen, S.A. *Mathematical Techniques in Multisensor Data Fusion*; Artech House: Norwood, MA, USA, 2004.
23. Newman, A.J.; Mitzel, G.E. Upstream data fusion: History, technical overview, and applications to critical challenges. *Johns Hopkins APL Tech. Digest* **2013**, *31*, 215–233.
24. Smith, D.; Singh, S. Approaches to multisensor data fusion in target tracking: A survey. *IEEE Trans. Knowl. Data Eng.* **2006**, *18*, 1696–1710.
25. Zhu, S.; Guo, Y.; Chen, J.; Li, D.; Cheng, L. Integrating Optimal Heterogeneous Sensor Deployment and Operation Strategies for Dynamic Origin-Destination Demand Estimation. *Sensors* **2017**, *17*, 1767.
26. Alparone, L.; Aiazzi, B.; Baronti, S.; Garzelli, A.; Nencini, F.; Selva, M. Multispectral and panchromatic data fusion assessment without reference. *Photogramm. Eng. Remote Sens.* **2008**, *74*, 193–200.
27. Varshney, P.K. Multisensor data fusion. *Electron. Commun. Eng. J.* **1997**, *9*, 245–253.
28. Zhang, J. Multi-source remote sensing data fusion: status and trends. *Int. J. Image and Data Fusion* **2010**, *1*, 5–24,

29. Dawn, S.; Saxena, V.; Sharma, B. Remote Sensing Image Registration Techniques: A Survey. In *Proceedings of the Image and Signal Processing: 4th International Conference, ICISP 2010, Trois-Rivières, QC, Canada, 30 June–2 July 2010*; Elmoataz, A., Lezoray, O., Nouboud, F., Mammas, D., Meunier, J., Eds.; Springer: Berlin/Heidelberg, Germany, 2010; pp. 103–112. doi:10.1007/978-3-642-13681-8_13.
30. Guo, Z.; Sun, G.; Ranson, K.; Ni, W.; Qin, W. The Potential of Combined LIDAR and SAR Data in Retrieving Forest Parameters using Model Analysis. In *Proceedings of the IGARSS 2008—2008 IEEE International Geoscience and Remote Sensing Symposium, Boston, MA, USA, 7–11 July 2008*.
31. Koetz, B.; Sun, G.; Morsdorf, F.; Ranson, K.; Kneubühler, M.; Itten, K.; Allgöwer, B. Fusion of imaging spectrometer and LIDAR data over combined radiative transfer models for forest canopy characterization. *Remote Sens. Environ.* **2007**, *106*, 449–459.
32. Benferhat, S.; Sossai, C. Reasoning with multiple-source information in a possibilistic logic framework. *Inf. Fusion* **2006**, *7*, 80–96.
33. Benferhat, S.; Titouna, F. Fusion and normalization of quantitative possibilistic networks. *Appl. Intell.* **2009**, *31*, 135–160.
34. Crowley, J.L. Principles and techniques for sensor data fusion. In *Multisensor Fusion for Computer Vision*; Springer-Verlag: Berlin/Heidelberg, Germany, 1993.
35. Yue, Y.; Yang, C.; Wang, Y.; Senarathne, P.G.C.N.; Zhang, J.; Wen, M.; Wang, D. A Multilevel Fusion System for Multirobot 3-D Mapping Using Heterogeneous Sensors. *IEEE Syst. J.* **2020**, *14*, 1341–1352.
36. Hurley, J.D.; Johnson, C.; Dunham, J.; Simmons, J. Multilevel probabilistic target identification methodology utilizing multiple heterogeneous sensors providing various levels of target characteristics. In *Signal Processing, Sensor/Information Fusion, and Target Recognition XXVII*; International Society for Optics and Photonics: Orlando, FL, USA, 2018; Volume 10646, p. 106461M.
37. Malcolm, G. Sheaves, objects, and distributed systems. *Electron. Notes Theor. Comput. Sci.* **2009**, *225*, 3–19.
38. Bałdowski, K. Whitney numbers of geometric lattices. *Adv. Math.* **1975**, *16*, 125–138.
39. Bałdowski, K. Galois connections and the Leray spectral sequence. *Adv. Math.* **1977**, *25*, 191–215.
40. Curry, J. Sheaves, cosheaves, and applications. *arXiv* **2013**, arxiv:1303.3255.
41. Lilius, J. *Sheaf Semantics for Petri Nets*; Technical Report; Helsinki University of Technology, Digital Systems Laboratory: Espoo, Finland, 1993.
42. Robinson, M. Imaging geometric graphs using internal measurements. *J. Diff. Eqns.* **2016**, *260*, 872–896. doi:10.1016/j.jde.2015.09.014.
43. Shepard, A. A Cellular Description of the Derived Category of a Stratified Space. Ph.D. Thesis, Brown University, Providence, RI, USA, 1985.
44. Joslyn, C.; Hogan, E.; Robinson, M. Towards a topological framework for integrating semantic information sources. In *Proceedings of the Semantic Technologies for Intelligence, Defense, and Security (STIDS), Fairfax, VA, USA, 18–21 November 2014*.
45. Robinson, M. Sheaf and duality methods for analyzing multi-model systems. In *Novel Methods in Harmonic Analysis*; Pesenson, I., Gia, Q.L., Mayeli, A., Mhaskar, H., Zhou, D.X., Eds.; Birkhäuser: Basel, Switzerland, 2017.
46. Mansourbeigi, S. Sheaf Theory Approach to Distributed Applications: Analysing Heterogeneous Data in Air Traffic Monitoring. *Int. J. Data Sci. Anal.* **2017**, *3*, 34–39.
47. Mansourbeigi, S.M. Sheaf Theory as a Mathematical Foundation for Distributed Applications Involving Heterogeneous Data Sets. In *Proceedings of the 32nd International Conference on Advanced Information Networking and Applications Workshops (WAINA), Krakow, Poland, 16–18 May 2018*.
48. Zadrozny, W.; Garbayo, L. A Sheaf Model of Contradictions and Disagreements. Preliminary Report and Discussion. *arXiv* **2018**, arXiv:1801.09036.
49. Purvine, E.; Aksoy, S.; Joslyn, C.; Nowak, K.; Praggastis, B.; Robinson, M. A Topological Approach to Representational Data Models. In *Lecture Notes in Computer Science: Human Interface and the Management of Information*; Springer: Las Vegas, NV, USA, 2018.
50. Robinson, M.; Henrich, J.; Capraro, C.; Zulch, P. Dynamic Sensor Fusion Using Local Topology. In *Proceedings of the IEEE Aerospace Conference, Big Sky, MT, USA, 3–10 March 2018*.
51. Robinson, M. *Assignments to Sheaves of Pseudometric Spaces*; Technical Report; American University: Washington, DC, USA, 2018.
52. Goguen, J.A. Sheaf semantics for concurrent interacting objects. *Math. Struct. Comput. Sci.* **1992**, *2*, 159–191.

53. Nelaturi, S.; de Kleer, J.; Shapiro, V. Combinatorial models for heterogeneous system composition and analysis. In Proceedings of the 11th System of Systems Engineering Conference (SoSE), Kongsberg, Norway, 12–16 June 2016.
54. Goldblatt, R. *Topoi, the Categorical Analysis of Logic*; Dover: Mineola, NY, USA, 2006.
55. Kokar, M.; Tomasik, J.; Weyman, J. Formalizing classes of information fusion systems. *Inf. Fusion* **2004**, *5*, 189–202.
56. Kokar, M.M.; Baclawski, K.; Gao, H. Category Theory-Based Synthesis of a Higher-Level Fusion Algorithm: An Example. In Proceedings of the 9th International Conference on Information Fusion Information Fusion, Florence, Italy, 10–13 July 2006.
57. Spivak, D.I. *Category Theory for the Sciences*; MIT Press: Cambridge, MA, USA, 2014.
58. Ghrist, R.; Hiraoka, Y. Network Codings and Sheaf Cohomology. In Proceedings of the NOLTA, Kobe, Japan, 4–7 September 2011.
59. Nguemo, M.A.; Tcheka, C. Sheaf cohomology on network codings: maxflow-mincut theorem. *Appl. Gen. Topol.* **2017**, *18*, 219–230.
60. Abramsky, S.; Barbosa, R.S.; Kishida, K.; Lal, R.; Mansfield, S. Contextuality, cohomology and paradox. *arXiv* **2015**, arXiv:1502.03097.
61. Abramsky, S.; Brandenburger, A. The sheaf-theoretic structure of non-locality and contextuality. *New J. Phys.* **2011**, *13*, 113036.
62. Curry, J.; Ghrist, R.; Nanda, V. Discrete Morse theory for computing cellular sheaf cohomology. *Found. Comput. Math.* **2016**, *16*, 875–897.
63. Kirk, R.; Bolstad, P.; Manson, S. Spatio-temporal trend analysis of long-term development patterns (1900–2030) in a Southern Appalachian County. *Landsc. Urban Plann.* **2012**, *104*, 47–58.
64. Mitchell, M.; Zimmerman, J.; Powell, R. Test of a habitat suitability index for black bears in the southern appalachians. *Wildl. Soc. Bull.* **2002**, *30*, 794–808.
65. Willey, C. Aging black bears from first premolar tooth sections. *J. Wildl. Manag.* **1974**, *38*, 97–100.
66. Joslyn, C.A.; Aksoy, S.; Callahan, T.J.; Hunter, L.; Jefferson, B.; Praggastis, B.; Purvine, E.A.; Tripodi, I.J. Hypernetwork Science: From Multidimensional Networks to Computational Topology. *arXiv* **2020**, arXiv:2003.11782.
67. Edelsbrunner, H.; Harer, J.L. *Computational Topology: An Introduction*; AMS: Providence, RI, USA, 2000.
68. Ghrist, R. Barcodes: The Persistent Topology of Data. *Bull. Am. Math. Soc.* **2007**, *45:1*, 61–75.
69. Google. googlemaps 2.4.5. Available online: <https://github.com/googlemaps/google-maps-services-python> (accessed on 15 June 2020).
70. Bieniek, T. Available online: <https://github.com/Turbo87/utm> (accessed on 15 June 2020).
71. Robinson, M. *Pseudosections of Sheaves with Consistency Structures*; Technical Report; American University: Washington, DC, USA, 2015.
72. Praggastis, B. Maximal Sections of Sheaves of Data over an Abstract Simplicial Complex. Available online: <https://arxiv.org/abs/1612.00397> (accessed on 15 June 2020).
73. Stanley, R. *Enumerative Combinatorics*, 2 ed.; Cambridge University Press: New York, NY, USA, 2011, Volume 1.
74. MacDonald, I.L.; Zucchini, W. *Hidden Markov and Other Models for Discrete-Valued Time Series*; CRC Press: Washington, DC, USA, 1997, Volume 110.
75. Harvey, A.; Koopman, S.J.; Shephard, N. *State Space and Unobserved Component Models: Theory and Applications*; Cambridge University Press: Cambridge, UK, 2004.
76. Meinhold, R.J.; Singpurwalla, N.D. Understanding the Kalman filter. *Am. Stat.* **1983**, *37*, 123–127.
77. Arulampalam, M.S.; Maskell, S.; Gordon, N.; Clapp, T. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Trans. Signal Process.* **2002**, *50*, 174–188. doi:10.1109/78.978374.
78. Pearson, J.B.; Stear, E.B. Kalman Filter Applications in Airborne Radar Tracking. *IEEE Trans. Aerosp. Electron. Syst.* **1974**, *AES-10*, 319–329. doi:10.1109/TAES.1974.307824.
79. Siouris, G.M.; Chen, G.; Wang, J. Tracking an incoming ballistic missile using an extended interval Kalman filter. *IEEE Trans. Aerosp. Electron. Syst.* **1997**, *33*, 232–240.
80. Wahba, G. *Spline Models for Observational Data*; SIAM: Philadelphia, PA, USA, 1990.
81. Robinson, M. Hunting for Foxes with Sheaves. *Not. Am. Math. Soc.* **2019**, *66*, 661–676.

82. Robinson, M. Finding cross-species orthologs with local topology. In Proceedings of the ACM-Biocomputing and Bioinformatics Workshop on Topological Data Analysis in Biomedicine, Seattle, WA, USA, 2–5 October 2016.
83. Ghrist, R.; Krishnan, S. A topological max-flow-min-cut theorem. In Proceedings of the 2013 IEEE Global Conference on Signal and Information Processing, Austin, TX, USA, 3–5 December 2013.



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).