Sheaves are the canonical data structure for sensor integration

Michael Robinson
Mathematics and Statistics
American University
Washington, DC, USA
michaelr@american.edu

1. Introduction

There is increasing concern within the data processing community about “swimming in sensors and drowning in data,” [1] because unifying data across many disparate sources is difficult. This refrain is repeated throughout many scientific disciplines, because there are few treatments of unified models for complex phenomena and it is difficult to infer these models from heterogeneous data.

A sensor integration framework should be (1) sufficiently general to accurately represent all sensors of interest, and also (2) be able to summarize information in a faithful way that emphasizes important, actionable information. Few approaches adequately address these two discordant requirements. Models of specific phenomena fail when multiple sensor types are combined into a complex network, because they cannot assemble a global picture consistently. Bayesian or network theory tolerate more sensor types, but suffer a severe lack of sophisticated analytic tools.

The mathematics of *sheaves* partially addresses our two requirements and provides several powerful inferential tools that are not available to other representational frameworks. This article presents (1) a sensor-agnostic measure of data consistency, (2) a sensor-agnostic optimization method to fuse data, and (3) sensor-agnostic detection of possible “systemic blind spots.” In this article, we show that sheaves provide both theoretical and practical tools to allow representations of locally valid datasets in which the datatypes and contexts vary. Sheaves therefore provide a common, canonical language for heterogeneous datasets. We show that sheaf-based fusion methods can combine disparate sensor modalities to dramatically improve target localization accuracy over traditional methods in a series of examples, culminating in Example 24. Sheaves provide a sensor-agnostic basis for identifying when information may be inadvertently lost through processing, which we demonstrate computationally in Examples 30 – 32.

Other methods typically aggregate data either exclusively globally (on the level of whole semantic ontologies) or exclusively locally (through various maximum likelihood stages). This limits the kind of inferences that can be made by these approaches. For instance, the data association problem in tracking frustrates local approaches (such as those based on optimal state estimation) and remains essentially unsolved in the general case. The analysis of sheaves avoids both of these extremes by specifying information where it occurs – locally at each data source – and then uses global relationships to constrain how these data are interpreted.

The foundational and canonical nature of sheaves means that existing approaches that address aspects of the sensor integration problem space already illuminate some portion of sheaf theory without exploiting its full potential. In contrast to the generality that is naturally present in sheaves, existing approaches to combining heterogeneous quantitative and qualitative data tend to rely on specific domain knowledge on small collections of data sources. Even in the most limited setting of pairs quantitative data sources, considerable effort has been expended in developing models of their joint behavior. Our approach leverages these pairwise models into models of multiple-source interaction. Additionally, where joint models are not yet available, sheaf theory provides a context for understanding the properties that such a model must have.

1.1. Contribution

Sometimes data fusion is described methodologically, as it is in the Joint Directors of Laboratories (JDL) model [2, 3], which defines operationally-relevant “levels” of data fusion. A more permissive – and more theoretically useful – definition was given by Wald [4], “data fusion is a formal framework in which are expressed the means and tools for the alliance of data originating from different sources.” This article shows that this is *precisely* what is provided by sheaf theory, by carefully outlining the requisite background in sheaves alongside detailed examples of collections of sensors and their data. In Chapter 3 of Wald’s book [5], the description of the implications of this definition is strikingly similar to what is presented in this article. But we go further, showing that we obtain not just a full-fidelity representation of the data and its internal relationships, but also a jumping-off point for analysis of both the integrated sensor system and the data it produces.

We make the distinction between “sensor integration” in which a collection of *models* of sensors are combined into a single unified model (the *integrated sensor system*), and “sensor data fusion” in which *observations* from those sensors are combined. This article presents the sheaf-based workflow outlined by Figure 1, which divides the process of working with sensors and their data into several distinct stages:

1. *Sensor integration* unifies models of the individual sensors and their inter-relations into a single system model, which we will show has the mathematical structure of a *sheaf*. We emphasize that no sensor *data* are included in the sensor system model represented by the sheaf.

Figure 1: The sheaf-based multi-sensor workflow.

2. *Consistency radius computation* uses the sheaf to quantify the level of self-consistency of a set of data supplied by the sensors.
3. *Data fusion* takes the sheaf and data from the sensors to obtain a new dataset that is consistent across the sensor system through a constrained optimization. The consistency radius of the original data places a lower bound on the amount of distortion incurred by the fusion process.
4. *Cohomology* of the sheaf detects possible problems that could arise within the data fusion process. Although cohomology does not make use of any sensor data, it quantifies the possible impact that certain aspects of the data may have on the fusion process.

Through a mixture of theory and detailed examples, we will show how this workflow presents solutions to four distinct problem domains:

1. Formalizing the description of an integrated sensor system as a well-defined mathematical entity (defined by Axioms 1–6 in Section 3),
2. Quantifying the internal consistency of the data from individual sensors (the *consistency radius* in Definition 20 in Section 4),
3. Deriving globally consistent data from the data provided by the entire collection of sensors (solving Problem 19, *sheaf-based data fusion* in Section 4), and
4. Measuring the impact of the relationships between sensors on what data will be deemed consistent (using *cohomology* in Section 5).

1.2. Historical context

There are essentially two threads of research in the literature on data fusion: (1) physical sensors that return structured, numerical data (“hard” fusion), and (2) unstructured, semantic data (“soft” fusion) [6]. Hard fusion is generally performed on spatially-referenced data (for example, images), while soft fusion is generally referenced to a common ontology. Especially in hard fusion, the spatial references are taken to be absolute. Most *sensor* fusion is performed on a pixel-by-pixel basis amongst sensors of the same physical modality (for instance [7, 8, 9]). It generally requires image registration [10] as a precondition,

especially if the sensors have different modalities (for instance [11, 12]). When image extents overlap but are not coincident, *mosaics* are the resulting fused product. These are typically based on pixel- or patch-matching methods (for instance [13]). Because these methods look for areas of close agreement, they are inherently sensitive to differences in physical modality. It can be difficult to extend these ideas to heterogeneous collections of sensors.

Like hard fusion, soft fusion requires registration amongst different data sources. However, since there is no physically-apparent “coordinate system,” soft fusion must proceed without one. There are a number of approaches to align disparate ontologies into a common ontology [14, 15, 16, 17], against which analysis can proceed. There, most of the analysis derives from a combination of tools from natural language processing (for instance [18, 19, 20]) and artificial intelligence (like the methods discussed in [21, 22, 23]). That these approaches derive from theoretical logic, type theory, and relational algebras is indicative of deeper mathematical foundations. These three topics have roots in the same category theoretic machinery that drives the sheaf theory discussed in this article.

Weaving the two threads of hard and soft fusion is difficult at best, and is usually approached statistically, as discussed in [24, 25, 26]. Unfortunately, this “clouds” the issue. If a stochastic data source is combined with another data source (deterministic or not), stochastic analysis asserts that the result will be stochastic. This viewpoint is quite effective in multi-target tracking [27, 28] or in event detection from social media feeds [29], where there are sufficient data to estimate probability distributions. But if two *deterministic* data sources are combined, one numeric and one textual, why should the result be *stochastic*?

Regardless of this conundrum, information theoretic or possibilistic approaches seem to be natural and popular candidates for performing sensor integration, for instance [30, 31, 32]. They are actually subsumed by category theory [33, 34] and arise naturally when needed. These models tend to rely on the homogeneity of sensors in order to obtain strong theoretical results. Sheaf theory extends the reach of these methods by explaining that the most robust aspects of networks tend to be topological in nature. For example, one of the strengths of Bayesian methods is that strong convergence guarantees are available. However, when applied to data sources arranged in a feedback loop, Bayesian updates can converge to the wrong distribution or not at all! [35] The fact that this is possible is witnessed by the presence of a topological feature – a loop – in the relationships between sources.

Sheaf theory provides canonical computational tools that sit on the general framework and has been occasionally [36, 37, 38] used in applications. Unfortunately, it is often difficult to get the general machine into a systematic computational tool. Combinatorial advances [39, 40, 41, 42] have enabled the approach discussed in this article through the use of finite topologies and related structures.

Sheaves, and their cohomological summaries, are best thought of as being part of *category theory*. There has been a continuing undercurrent that sheaves aspire to be more than that – they provide the bridge between categories and

logic [43]. The pioneering work by Goguen [44] took this connection seriously, and inspired this paper – our axioms sets are quite similar – though the focus of this paper is somewhat different. The approach taken in this paper is that time may be implicit or explicit in the description of a data source; for Goguen, time is always explicit. More recently category theory has become a valuable tool for data scientists, for instance [45], and information fusion [46, 47].

2. Motivating examples

Although the methodology presented in this article is sufficiently general to handle many configurations of sensors and data types, we will focus on two simple scenarios for expository purposes. To emphasize that our approach is applicable for many sensor modalities, the two scenarios involve rather different signal models.

1. A notional search and rescue scenario in which multiple types of sensors are used to locate a downed airliner.
2. A computer vision task involving the tabulation of coins on a table.

We will carry these two scenarios throughout the article. We will start with the most basic stage of constructing their integrated sensor models in Section 3, we will then analyze typical sensor data in Section 4, and finally offer a data-independent analysis of the sensor models in Section 5.

2.1. Search and rescue scenario

The search and rescue scenario involves three distinct kinds of sensors (this is simplified from the guidelines produced by the Australian government [48] for the search for Malaysia Airlines Flight MH370) distributed¹ according to Figure 2.

1. The original flight plan, Air Traffic Control (ATC) radars, and the Automatic Dependent Surveillance-Broadcast (ADS-B) aircraft beacon system, all of which produce time-stamped locations of an aircraft,
2. Radio Direction Finding (RDF) sensors provided by amateur radio operators, which produce accurately time-stamped but highly uncertain bearings to an emergency beacon deployed by the aircraft near the crash site, and
3. Wide-area satellite images, which are time-stamped georeferenced intensity-valued images.

The data consists of 612 ATC and ADS-B observations. Of these, 505 observations contained latitude, longitude, altitude and speed, while the other 107 consisted of latitude and longitude only. The update rate varied considerably

¹The coastline data was obtained from <https://serc.carleton.edu/download/files/79176/coast.mat>

Figure 2: The sensors used in the search and rescue scenario for this article.

over the data, but typical observations were about 10 km apart. The RDF bearings and satellite images each consisted of exactly one observation each. Since the data were simulated, various amounts of noise and systematic biases were applied, as detailed in Example 24.

The goal of the scenario is to locate the airliner’s crash site. Due to the fact that the most accurate wide-area sensors – the ATC radars and ADS-B beacons – do not necessarily report an aircraft’s location over the water, this problem is limited by uncertainty about what happened after the airliner leaves controlled airspace. The time of the crash is (of course!) not part of the flight plan, and was not visible to the ATC radars nor to the satellite.

The time of crash is central to the problem because it connects the time of crash, the last known position, and last known velocity to an estimate of where the crash location was. When corroborated by the intersection of the RDF bearings, the crash location estimate can be used to task the satellite. Thus, the problem’s solution is to combine the bearings and time-stamps from the RDF sensors with the other sensors, a task performed by the “field office.”

2.2. Computer vision scenario

In the usual formulation of multi-camera mosaic formation (for instance [7, 8]), several cameras with similar perspectives are focused on a scene. Assuming that the perspectives and lighting conditions are similar enough, large mosaics can be formed. The problem is more complicated if the lighting or perspectives differ substantially, but it can still be addressed. In either case, the methodology aims to align *pixels* or *regions of pixels*. If some of the images are produced by sensors with different modalities – say a LIDAR and a RADAR – pixel alignment is doomed to fail. Then, the problem is fully heterogeneous, and is usually handled by trying to interpolate into a common coordinate system [11, 12]. If

Figure 3: Several mosaics (bottom row) of three photographs (top row) under different conditions. The mosaic on the left is perfect, in that the subimages on the overlap have identical pixel values. The mosaic in the middle is imperfect, in that the pixel values of the two subimages differ somewhat on the overlap. The mosaic on the right combines textual and numeric data derived from an object detector run on image 2 with the original image 3.

there is insufficient sample support or perspectives are too different, then it is necessary to match object-level summaries rather than pixel regions [9].

This example studies the task of counting coins on a table. We will introduce heterogeneity in the imaging modality through different kinds of processing:

1. All cameras yield images, with identical lighting and perspective: perfect mosaics can be formed (as shown in the lower left mosaic in Figure 3),
2. All cameras yield images, though lighting and perspective can differ: imperfect mosaics can be formed (as shown in the lower middle mosaic in Figure 3), and
3. A mixture of camera images and object-level detections derived from the images: a fully heterogeneous data fusion problem (as shown in the lower right mosaic in Figure 3).

These three examples showcase the transition from homogeneous data in problem (1) to heterogeneous data in problem (3), with homogeneous data being a clear special case. These examples highlight how fusion of heterogeneous data is to be performed, and makes the resulting structure-preserving summaries easy to interpret.

3. Sensor integration

The first stage of our workflow constructs a unified model of sensors and their interactions. To emphasize the generality of our approach from the outset – our methodology is not limited to any particular kind of sensor – we posit

a set of six Axioms to specify what this integrated system must be. A *sheaf* is precisely any mathematical object satisfying all six Axioms. We motivate the Axioms by way of examples from our two scenarios, a process that also demonstrates how the Axioms apply for other systems as well.

Axiom 1: The entities observed by the sensors lie in a set X .

The first Axiom is not controversial. “Entities” might be targets, sensor data, or even textual documents.

The second Axiom considers properties of the attributes. In order to quantify differences between attributes, we will use a *pseudometric*.

Definition 1. If A is a set, function $d : A \times A \rightarrow \mathbb{R}$ is called a *pseudometric* when it satisfies the following:

1. $d(x, x) = 0$ for all $x \in A$,
2. $d(x, y) = d(y, x) \geq 0$ for all $x, y \in A$, and
3. $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in A$.

We call the pair (A, d) a *pseudometric space*.

If additionally, $d(x, y) > 0$ whenever $x \neq y$, we call d a *metric* and (A, d) a *metric space*.

Axiom 2: All possible attributes of the entities lie in a collection² of pseudometric spaces $\mathcal{A} = \{(A_1, d_1), (A_2, d_2), \dots\}$.

We lose no generality by requiring each set of attributes to have a pseudometric, because we may choose the *discrete metric* which asserts that all pairs of distinct attributes are distance 1 apart.

Example 2. Throughout this article, we will use the following entities in the search and rescue scenario:

- *Aircraft last known position*, represented as its longitude x , its latitude y , and its altitude z ,
- *Aircraft last known velocity*, which is represented as a horizontal velocity vector (v_x, v_y) (we will assume the rate of climb or descent is not known),
- *Precise time of the crash*, represented as the time elapsed since the last known position t ,
- *Bearing to RDF sensor 1*, represented as a compass bearing θ_1 from the crash location to the sensor,
- *Bearing to RDF sensor 2*, represented as a compass bearing θ_2 from the crash location to the sensor, and

²The examples in this article all make use of a finite collection $\mathcal{A}$, but there is no particular reason why this collection should be finite or even countable.

- *The satellite image*, represented by s .

Therefore, $X = \{x, y, z, v_x, v_y, t, \theta_1, \theta_2, s\}$ is the set of entities. The attributes are given by the following sets, each of which comes with a metric:

- *Three dimensional positions*: $\mathbb{R}^3$ specifying longitude, latitude, and altitude in kilometers under the great circle distance metric,
- *Two dimensional positions*: $\mathbb{R}^2$ specifying longitude and latitude in kilometers under the great circle distance metric,
- *Two dimensional velocities*: $\mathbb{R}^2$ specifying longitude and latitude velocities in kilometers/second under the Euclidean metric,
- *Time*: $\mathbb{R}$ in units compatible with the positions and velocities,
- *Bearings*: $S^1 = \{(u, v) : u^2 + v^2 = 1\} = [0, 360)$ points on the unit circle, represented either by their coordinates or as angles in degrees, and
- *Images*: $C([0, 1] \times [0, 1], \mathbb{R})$, continuous real-valued functions on the unit square. (Satellite images are typically lowpass filtered and oversampled, continuous functions are a reasonable model. For other imaging modalities, piecewise continuity may be a better model.)

Some sensors will also return products of these sets so the set of attributes is all such products $\mathcal{A} = \{\mathbb{R}^3, \mathbb{R}^2, \mathbb{R}, S^1, C([0, 1] \times [0, 1], \mathbb{R}), \mathbb{R}^3 \times \mathbb{R}^2, \mathbb{R}^3 \times \dots\}$.

Thus far, no attributes have yet been ascribed to entities, so there is nothing to distinguish entities from one another. Yet since different sets of entities are measured by sensors, it makes sense to study these sets.

Example 3. For the search and rescue scenario, several sets of entities will be associated to the sensors as indicated in the list below:

- *Flight plan*: $\{x, y, z\}$,
- *ATC radar*: $\{x, y, z, v_x, v_y\}$,
- *RDF sensor 1*: $\{\theta_1, t\}$,
- *RDF sensor 2*: $\{\theta_2, t\}$, and
- *Satellite*: $\{\theta_1, \theta_2, s\}$.

The reader is cautioned that each of these are *unordered* sets of entities. Only when attributes are assigned to each sensor by Axiom 4, will ordering among the entities be apparent (see Example 9).

The field office (where the data fusion is performed) has access to all of the entities. The sets of entities are permitted to overlap, since the same entity may be observed by multiple sensors. We usually expect that those pairs of sensors that are providing the same entities will provide similar attributes for

these entities. For instance the flight plan and the ATC should provide similar position information unless there has been a deviation from the plan. Any such deviation is probably a cause for concern and is something that we want to detect!

One way to formalize the kinds of collections of entities is to construct a *topology* on X , which specifies which sets of entities have attributes worthy of comparison. Intuitively a topology expresses that if two entities are unrelated, they can be described in isolation of each other. If they are closely related, they will be “hard to disentangle” and will usually be described together. The way to formalize this is by enriching a set model of the universe to describe entity relationships (against which attributes of different entities can then be related) into a topological one. We recall the definition of a topology:

Definition 4. A *topology* on a set X is a collection $\mathcal{T}$ of subsets of X satisfying the following properties:

1. $X \in \mathcal{T}$.
2. $\emptyset \in \mathcal{T}$.
3. The union of any collection of elements of $\mathcal{T}$ is also in $\mathcal{T}$.
4. The intersection of any finite collection of elements of $\mathcal{T}$ is also in $\mathcal{T}$.

Usually, the elements of $\mathcal{T}$ are called *open sets*, and the pair $(X, \mathcal{T})$ is called a *topological space*. In the context of sensors, a $U \in \mathcal{T}$ is called a *sensor domain*.

The first property of a topology states that it is possible to ascribe attributes to all entities (it may be impractical, though). The second property states that it is possible to not describe any entities. The third and fourth properties are more subtle. They state that two entities can be related through their relationships with other entities. The finiteness constraint in the fourth property avoids examining an infinite regress through increasingly entity-specific relationships.³

Axiom 3: The collection of all sets of entities whose attributes are related forms a topology $\mathcal{T}$ on X .

We should emphasize that the sets listed in Example 3 do *not* form a topology. There are many intersections, for instance $\{t\} = \{\theta_1, t\} \cap \{\theta_2, t\}$, that do not correspond to a set of entities provided by a sensor. One must construct the topology by specifying the open sets as those sets of entities provided by the sensors, and then extending the collection until it is a topology.

Definition 5. If X is a set and $\mathcal{U}$ is a collection of subsets of X , the smallest topology $\mathcal{T}$ which contains $\mathcal{U}$ is given by

$$\mathcal{T} = \bigcap_{\mathcal{T}' \text{ is a topology with } \mathcal{U} \subseteq \mathcal{T}'} \mathcal{T}',$$

³In some topologies, the intersection of any collection of elements remains in the topology, but requiring this would destroy our ability to describe certain phenomena.

Figure 4: Open sets that generate the topology for the search and rescue scenario

and is called *the topology generated by $\mathcal{U}$* . In this case, we sometimes say that $\mathcal{U}$ is a *subbase* for $\mathcal{T}$.

Example 6. Figure 4 shows the relationships between the entity sets from Example 3. Using these sets as the subbase for the topology results in many other sets in the topology. The following open sets are formed as intersections:

1. $\{t\}$,
2. $\{\theta_1\}$, and
3. $\{\theta_2\}$.

There are many more open sets formed as unions of these sets, for instance

1. $\{\theta_1, \theta_2\}$,
2. $\{\theta_1, \theta_2, s, t\}$,
3. $\{x, y, z, \theta_1, \theta_2\}$,
4. $\{x, y, z, \theta_1, \theta_2, s\}$,
5. $\{x, y, z, \theta_1, \theta_2, t\}$,
6. $\{x, y, z, \theta_1, \theta_2, s, t\}$,

among others. Rather than being dismayed at all the possibilities, observe that many of these open sets represent redundant collections of information about entities. But which collections should be retained? For the moment, we will retain all of them. We will postpone a more complete answer to that question until Section 5 when we can provide a theoretical guarantee (Theorem 29) that information is not lost when only a smaller number of entity sets are used.

Following Example 6, sensor data will be supplied at some of the open sets in $\mathcal{T}$. These data are then merged on unions of these open sets whenever the data are consistent along intersections. The precise relationships among open sets in $\mathcal{T}$ are crucial and generally should have physical meaning.

Figure 5: Open sets in two different topologies for image fusion: (a) part of the usual Euclidean topology and (b) a coarser topology generated by two camera fields of view.

Remark 7. Axiom 3 is a true limitation,⁴ which posits that entities are organized into local neighborhoods. In all situations of which the author is aware, the appropriate topology $\mathcal{T}$ is either the geometric realization of a cell complex or (equivalently) a finite topology built on a preorder as described in [49]. The three axioms encompass unlabeled graphs as a special case, for instance. Finite topologies are particularly compelling from a practical standpoint even though relatively little has been written about them. The interested reader should consult [50] for a (somewhat dated) overview of the interesting problems.

Example 8. In the case of image mosaics, the open sets in $\mathcal{T}$ should have something to do with the regions observed by individual cameras. Let the set of entities X consist of the locations of points in the scene. There are two obvious possibilities for a topology:

1. The topology $\mathcal{T}_1$ of the scene itself (some of which is shown in Figure 5(a)), generally a manifold or a subspace of Euclidean space based on the pixel values or
2. The topology $\mathcal{T}_2$ which is generated by a much smaller collection of sets $\mathcal{U}$, in which each $U \in \mathcal{U}$ consists of a set of points observed by a single camera (the entirety of this topology is shown in Figure 5(b)).

Although $\mathcal{T}_1$ is rather natural – and is what is usually used when forming mosaics – it has the disadvantage of being quite large. Figure 5(a) only shows a small part of one such example! It also requires fusion to be performed at

⁴Instead of a topology, one might consider generalizing further to a Grothendieck *site* [43], but this level of generality seems excessive for practical sensor integration.

the pixel level, which is generally inappropriate when there are other heterogeneous sensors that do not return pixel values. This also can be computationally intense, explains why pixel-level fusion scales badly. In contrast, $\mathcal{T}_2$ is much smaller (usually finite) and therefore more amenable both to computations and to fusion of heterogeneous data types.

Axiom 4: Each sensor assigns one of the pseudometric spaces of attributes to an open subset of entities in the topological space $(X, \mathcal{T})$.

- This assignment is denoted by $\mathcal{S} : \mathcal{T} \rightarrow \mathcal{A}$, so that the attribute pseudometric space assigned to an open set $U \in \mathcal{T}$ is written $\mathcal{S}(U) \in \mathcal{A}$.
- The space of attributes $\mathcal{S}(U)$ is called the space of *observations* over the *sensor domain* U .
- The pseudometric on $\mathcal{S}(U)$ will be written d_U .
- An element of the Cartesian product $\prod_{U \in \mathcal{T}} \mathcal{S}(U)$ is called an *assignment*⁵ of $\mathcal{S}$.
- We call an element $s \in \mathcal{S}(X)$ a *global section*.

The distinction between a space of attributes and a space of observations is that the latter is directly associated to a particular sensor. In this article, “attributes” may not necessarily be associated to any entity, but “observations” are attributes that are assigned to a definite open set of entities.

There are some easy recipes for choosing a space of observations $\mathcal{S}(U)$ based on the kind of sensor corresponding to the open set U :

- For a camera sensor, the space of observations should be the set of possible images, perhaps denoted as vectors in $\mathbb{R}^{3mn}$ (m rows, n columns, and 3 colors for each pixel).
- For text documents, the space of observations could be the content of a document as a list of words or a word frequency vector.
- If the open set corresponds to the intersection of two sensors of the same type, then usually the observations ought to also be of that type. For instance, the observations in the overlap between two camera images are smaller images.

However, if the open set corresponds to the intersection or union of two different sensors (and therefore not a physical sensor), the set of observations needs to be more carefully chosen and may not correspond to the raw data supplied by any sensor. This situation is made well-defined by Axioms 5 – 6.

⁵When all of the sets of observations are vector spaces, *assignments* are usually called *serrations*.

Example 9. Satisfying Axioms 1-4 in the search and rescue scenario requires pairing the open sets in Example 6 with the attribute sets from Example 2. As noted in Example 6, there are *many* open sets; for brevity's sake, we will not specify them all. We will denote the assignment of observations $\mathcal{S}$.

Axiom 4 merely asks that $\mathcal{S}$ assign spaces of observations, so for instance the x , y , and z coordinates could be assigned by $\mathcal{S}(\{x, y, z\}) = \mathbb{R}^3$. Since the open set $U_1 = \{x, y, z\}$ implies no ordering on the entities, it is not yet clear which component of $\mathbb{R}^3$ corresponds to which entity. Indeed it may be that the correspondence between attributes and entities might be rather complicated. Given that consideration, those open sets that are associated to sensors according to Example 3 will be modeled as having spaces of observations as follows:

- *Flight plan*: $\mathcal{S}(U_1) = \mathcal{S}(\{x, y, z\}) = \mathbb{R}^3$ the last known position according to the flight plan.
- *ATC*: $\mathcal{S}(U_2) = \mathcal{S}(\{x, y, z, v_x, v_y\}) = \mathbb{R}^3 \times \mathbb{R}^2$ the last known position and velocity reported by the ATC.
- *RDF sensor 1*: $\mathcal{S}(U_3) = \mathcal{S}(\{\theta_1, t\}) = S^1 \times \mathbb{R}$ bearing and time of beacon's signal as reported by amateur operator 1.
- *RDF sensor 2*: $\mathcal{S}(U_4) = \mathcal{S}(\{\theta_2, t\}) = S^1 \times \mathbb{R}$ bearing and time of beacon's signal as reported by amateur operator 2.
- *Satellite*: $\mathcal{S}(U_5) = \mathcal{S}(\{s, \theta_1, \theta_2\}) = \mathbb{R}^2$ object detection from satellite image. We could have equally well chosen $\mathcal{S}(U_5)$ to be the space of raw gray scale images $C([0, 1] \times [0, 1], \mathbb{R})$ or $\mathbb{R}^{3mn}$, but following standard practice in wide-area search, we assume that the image has already been processed into a list of candidate object detections,⁶ each of which consists of a latitude and longitude only.
- *Field office*: $\mathcal{S}(X) = \mathcal{S}(\{x, y, z, t, v_x, v_y, \theta_1, \theta_2, s\}) = \mathbb{R}^3 \times \mathbb{R}^2 \times \mathbb{R}$ last known position, last known velocity, time of beacon, and satellite image. We note that because the satellite image has been processed into object detections, the unprocessed image is not represented in the attribute space.

For those open sets associated to intersections, the spaces of observations are easy to specify by reading off the entities in the open sets:

1. $\mathcal{S}(\{t\}) = \mathbb{R}$, the time of the crash,
2. $\mathcal{S}(\{\theta_1\}) = S^1$, the bearing of the beacon to amateur operator 1, and
3. $\mathcal{S}(\{\theta_2\}) = S^1$, the bearing of the beacon to amateur operator 2.

But what about sets like $\{\theta_1, \theta_2\}$? There are two obvious possibilities: $S^1 \times S^1$ (two bearings) versus $\mathbb{R}^2$ (the intersection of the bearing sight lines). Axiom 4 is mute about which is preferable, but it asserts that we must choose *some*

⁶If an object detection is consistent with a possible crash site, an analyst would be instructed to examine the unprocessed image.

set. Each open set formed by unions of the subbase is a place where such a decision must occur. Since Axioms 5 and 6 completely constrain these spaces of observations, we will delay their specification until after these Axioms are specified in Example 15.

Example 10. On an open set $U \cap V$ corresponding to the intersection between a camera image U and a radar image V , each observation in $\mathcal{S}(U \cap V)$ could be a list of target locations rather than the data from either sensor. If there are n targets with known locations in $\mathbb{R}^3$, then clearly this space is parameterized by $\mathbb{R}^{3n}$. If the number of targets is unknown, but is known to be less than n , then the space of observations would be given by

$$\{\perp\} \sqcup \mathbb{R}^3 \sqcup \mathbb{R}^6 \sqcup \dots \sqcup \mathbb{R}^{3n}$$

where $\sqcup$ represents disjoint union and $\perp$ represents the state of having no targets in the scene, and hence represents a null observation.

This example indicates that moving from the observations over one open set to a smaller open set ought to come from a particular transformation that depends on the sensors involved. Put another way, removing entities of a sensor domain from consideration should impact the set of observations in a systematic way.

Axiom 5: It is possible to transform sets of observations to new sets of observations by reducing the size of sensor domains.

- Whenever a pair of sensors assign attributes to open sets of entities $U \subseteq V \subseteq X$ in the topology $\mathcal{T}$, there is a (fixed) continuous function⁷

$$\mathcal{S}(U \subseteq V) : \mathcal{S}(V) \rightarrow \mathcal{S}(U)$$

from the pseudometric space of observations over the larger sensor domain to the pseudometric space of observations over the smaller one.

- If $U \subseteq V \subseteq W$ and $U \subseteq Y \subseteq W$, then

$$\mathcal{S}(U \subseteq V) \circ \mathcal{S}(V \subseteq W) = \mathcal{S}(U \subseteq Y) \circ \mathcal{S}(Y \subseteq W),$$

which means that we may refer to $\mathcal{S}(U \subseteq W)$ without ambiguity.

- $\mathcal{S}(U \subseteq U)$ is the identity function.
- The continuous function $\mathcal{S}(U \subseteq V)$ is called a *restriction*, because it restricts observations on an open set V to a subset U .

⁷Notationally, $\mathcal{S}(U)$ is a pseudometric space, but $\mathcal{S}(U \subseteq V)$ is a function between pseudometric spaces. $\mathcal{S}$ itself refers to *both* as a single mathematical object.

Figure 6: Some of restrictions for the search and rescue scenario; to reduce clutter, unions between open sets are not shown. (a) The open sets and their intersections, where the arrows indicate subset relations. (b) The spaces of sections, with restrictions shown as arrows. The restrictions A , B , C , D , and E are defined in Example 11.

- Any $\mathcal{S}$ that satisfies Axioms 1–5 is called a *presheaf of pseudometric spaces*. Analogously, if d_U is a metric for all $U \in \mathcal{T}$, we say $\mathcal{S}$ is a *presheaf of metric spaces*.

The topology $\mathcal{T}$ imposes a lower bound on how small sensor domains can be, and therefore how finely restrictions are able to discriminate between individual entities. Generally speaking, one wants the topology to be as small (coarse) as practical, so that fewer restrictions need to be defined (and later checked) during the process of fusing data.

Example 11. Continuing with the construction of $\mathcal{S}$ from Example 9 for the search and rescue scenario, restrictions can be constructed as shown in Figure 6, where id represents an identity function and pr_k represents a projection onto the k -th factor of a product. Except for the restrictions A , B , C , D , and E , the restrictions merely extract the spaces of observations for various entities that are retained when changing open sets. For instance, when relating the entities as viewed by the ATC to the flight plan, one merely leaves out the velocity measurements.

The restrictions A , B , C , D , and E need to perform some geometric computations to translate between compass bearings and locations. They are defined⁸ as follows:

$$\begin{aligned}
 A(x, y, z, v_x, v_y, t) &= \left(\tan^{-1} \frac{x + v_x t - r_{1x}}{y + v_y t - r_{1y}}, t \right), \\
 B(x, y, z, v_x, v_y, t) &= \left(\tan^{-1} \frac{x + v_x t - r_{2x}}{y + v_y t - r_{2y}}, t \right), \\
 C(s_x, s_y) &= \tan^{-1} \frac{s_x - r_{1x}}{s_y - r_{1y}}, \\
 D(s_x, s_y) &= \tan^{-1} \frac{s_x - r_{2x}}{s_y - r_{2y}}, \text{ and} \\
 E(x, y, z, v_x, v_y, t) &= (x + v_x t, y + v_y t)
 \end{aligned}$$

⁸We are using $\tan^{-1}$ for brevity – really something like the C language `atan2` function would be preferable to obtain angles anywhere on the circle.

Figure 7: Cropping an image corresponds to reducing the size of the sensor domain. The topology determines what crops can be performed as restriction maps: a coarse topology (left) permits fewer crops than a finer one (right).

where s_x, s_y are coordinates of an object detected in the satellite image, r_{1x}, r_{1y} are the coordinates of the first RDF sensor and r_{2x}, r_{2y} are the coordinates of the second RDF sensor. Note that restrictions involving unions of open sets are not shown, since we are still waiting to axiomatically constrain them using Axiom 6 in Example 15.

Example 12. Axiom 5 has quite a clear interpretation for collections of cameras – the restrictions perform cropping of the domains. The topology $\mathcal{T}$ controls how much cropping is permitted. Figure 7 gives two examples, in which the topology on the left is smaller than the one on the right. Indeed, if we consider the topology $\mathcal{T}_1$ from Example 8, then a huge number of crops must be defined, including many that are unnecessary to form a mosaic of the images. However, the topology $\mathcal{T}_2$ from that example requires the modeler to define only those crops necessary to assemble a mosaic from the images.

Without constraints on the nature of the restriction functions, it is in general impossible to combine observations from overlapping sensor domains because we lack a concrete definition for the sets of observations over unions. Axiom 6 supplies a necessary and sufficient constraint, by requiring there to be an assignment of observations over a union of sensor domains whenever there are observations that agree on intersections of sensor domains. Axiom 6 asserts that this “globalized” observation (a *section* of $\mathcal{S}$) is unique – a statement of the underlying determinism of the representation.

Axiom 6: Observations from overlapping sensor domains that are equal upon restricting to the intersection of the sensor domains uniquely determine an observation from the union of the sensor domains.

- Suppose that $\{U_\alpha\}_{\alpha \in I}$ is a collection of sensor domains and there are corresponding observations $a_\alpha \in \mathcal{S}(U_\alpha)$ indexed by $\alpha \in I$ taken from each

sensor domain. If the following equation⁹ is satisfied

$$(\mathcal{S}(U_\alpha \cap U_\beta \subseteq U_\alpha))(a_\alpha) = (\mathcal{S}(U_\alpha \cap U_\beta \subseteq U_\beta))(a_\beta),$$

for all α and β in I , then there is an $a \in \mathcal{S}(\cup_{\gamma \in I} U_\gamma)$ that restricts to each a_α via

$$a_\alpha = (\mathcal{S}(U_\alpha \subseteq \cup U_\gamma))(a).$$

- Suppose that U is a sensor domain and that $a, b \in \mathcal{S}(U)$ are assignments of observations to that sensor domain. If these two observations are equal on all sensor domains V in a collection $\{V_\alpha\}$ for $\alpha \in I$ of sensor domains with $\bigcup_{\alpha \in I} V_\alpha = U$,

$$\mathcal{S}(V \subseteq U)(a) = \mathcal{S}(V \subseteq U)(b),$$

then we can conclude¹⁰ that $a = b$.

- When $\mathcal{S}$ satisfies Axioms 1–6, we call $\mathcal{S}$ a *sheaf of pseudometric spaces*.

In the case of two sensor domains U_1, U_2 , Axiom 6 can be visualized as a pair of *commutative diagrams*

in which the arrows on the left represent subset relations, while the arrows on the right represent restriction maps. The Axiom asserts that pairs of elements on the middle row that restrict to the same element on the bottom row are themselves restrictions of an element in the top row. Therefore Axiom 6 can always be satisfied if $\mathcal{S}(U_1 \cup U_2)$ has not been constrained by the sensor models. One merely needs to define

$$\mathcal{S}(U_1 \cup U_2) = \{(x, y) \in \mathcal{S}(U_1) \times \mathcal{S}(U_2) : (\mathcal{S}(U_1 \cap U_2 \subseteq U_1))(x) = (\mathcal{S}(U_1 \cap U_2 \subseteq U_2))(y)\} \quad (1)$$

and let the restriction maps project out of the product in the diagram above.

⁹This equation is called *conjunctivity* in the sheaf literature.

¹⁰A presheaf satisfying this equation is called a *monopresheaf*.

Remark 13. Axiom 6 implies that the sets of observations cannot get larger as the open sets get smaller. For instance, the diagram

$$\begin{array}{ccc}
 & \mathcal{S}(U_1 \cup U_2) = \{a\} & \\
 f \swarrow & & \searrow \text{id} \\
 \mathcal{S}(U_1) = \{a, c\} & & \mathcal{S}(U_2) = \{a\} \\
 g \searrow & & \swarrow \text{id} \\
 & \mathcal{S}(U_1 \cap U_2) = \{a\} &
 \end{array}$$

in which $f(a) = a$, $g(a) = g(c) = a$, and id is the identity function violates Axiom 6. Specifically, Axiom 6 would imply that given that $c \in \mathcal{S}(U_1)$ and $a \in \mathcal{S}(U_2)$ both map to $a \in \mathcal{S}(U_1 \cap U_2)$, there should be an element of $\mathcal{S}(U_1 \cup U_2)$ mapping to $c \in \mathcal{S}(U_1)$ even though this simply does not occur.

Remark 14. If the observations consist of probability distributions, it is usually best to think of “equality” in Axiom 6 as being “equality almost surely” throughout the network of sensors, rather than pointwise equality.

Example 15. It is now possible to complete the definition of the sheaf model for the search and rescue scenario, picking up where we left off in Example 11. While enumerating all of the spaces of observations over possible unions in the topology would be tedious, a few are enlightening.

- Consider first the case of the union of U_1 (flight plan) and U_2 (ATC radar), for which we have already defined

$$\begin{aligned}
 \mathcal{S}(U_1) &= \mathbb{R}^3, \\
 \mathcal{S}(U_2) &= \mathbb{R}^3 \times \mathbb{R}^2, \text{ and} \\
 U_1 \cap U_2 &= U_1.
 \end{aligned}$$

Therefore, Equation (1) defines

$$\mathcal{S}(U_1 \cup U_2) = \mathcal{S}(U_2) = \mathbb{R}^3 \times \mathbb{R}^2.$$

- For the union of U_1 (flight plan) and U_3 (RDF sensor 1), we have that $U_1 \cap U_3 = \emptyset$, so Equation (1) defines

$$\mathcal{S}(U_1 \cup U_3) = \mathcal{S}(U_1) \times \mathcal{S}(U_3) = \mathbb{R}^3 \times S^1 \times \mathbb{R}.$$

- Finally, the open sets U_3, U_4 corresponding to the two RDF sensors have $U_3 \cap U_4 = \{t\}$. In Example 9, we defined $\mathcal{S}(\{t\}) = \mathbb{R}$ to represent the time of the crash measured by both beacons. Thus Equation (1) constructs

$$\mathcal{S}(U_3 \cup U_4) = S^1 \times \mathbb{R} \times S^1,$$

Figure 8: Sheaf models for three ways to combine information about coins on a table. (a) The topology for two overlapping camera fields of view. (b) Mosaicing the images directly. (c) processing the right image with object-level classifications. (d) processing the right image into the total value of the coins.

Figure 9: Three possible ways to combine information about coins on a table. (a) The original images. (b) forming the mosaic of two images, with no further interpretation. (c) fusing an image with object-level classifications, but no interpretation. (d) combining an image with the total value of the coins, which requires interpretation of objects in the scene.

which represents both bearings and the common time from the two RDF sensors. Global sections over $U_3 \cup U_4$ explicitly assert that the two sensors are in agreement about the time the beacon was sounded. If the two sensors do not agree, this cannot correspond to a global section!

With Axiom 6, it becomes possible to draw global inferences from data specified locally.

Example 16. Referring to the topology $\mathcal{T}_2$ shown Figure 8(a) (also in Figure 5(b)), we again consider the problem of counting coins on the table. Assume that there are three different types of coins: pennies, nickels, dimes, and quarters. Three different possibilities for processing chains lead to three distinct sheaves, shown in Figure 8:

1. No processing aside from image registration (Figure 8(b)),
2. Leaving one image (on the left) alone while running an object classifier on the other (Figure 8(c)) in which the restriction f computes the number of coins of each type present in the image, or
3. Leaving one image (again on the left) alone, while running a different process g that computes the total monetary value within an image (Figure 8(d)).

Notice that in the latter two cases, the observation in the intersection is *not* an image, but something rather more processed. Given that f and g actually perform the processing required of them, it is easy to verify that Axiom 6 will hold for each sheaf merely by verifying that Equation 1 is satisfied over the unions.

Starting with the original pair of images in Figure 9(a), which constitute observations over U_1 and U_2 in the sheaf in Figure 8(b), Figure 9(b) shows the corresponding global section. In the case of the sheaves shown in Figures 8(c) and 8(d) we can still use the image as an observation over U_1 . However, the image that was an observation over U_2 must be processed into object detections or a total dollar amount before it can constitute an observation for the sheaves in Figures 8(c) and 8(d). Once that is done, then Figures 9(c) and 9(d) constitute the resulting global sections that are guaranteed to exist by Axiom 6.

The previous examples show that observations over some of the sensor domains can be used to infer observations over all other sensor domains when there is self consistency among the observations. This extrapolation of observations into other sensor domains is the hallmark of a data fusion process, but depends on restrictions of observations being exactly equal on intersections of sensor domains. Of course, this is quite unrealistic! In Section 4 we argue that *data fusion* consists of the process of identifying the global section that is the “best match” to the available data.

Sometimes, as the next example shows, problems arise when the integrated sensor system does not satisfy Axiom 6.

Figure 10: The diagram of the sheaf modeling two tables in a passenger database owned by an airline.

Example 17. Real world examples of relational databases often do not satisfy Axiom 6, which is sometimes a source of problems. Consider a database with two tables like so:

1. a table whose columns are passenger names, flight numbers, a starting airport, and a destination airport, and
2. a table whose columns are passenger names, frequent flier miles, and flier upgrade status.

The two tables are shown at left and right middle in Figure 10. The restrictions (shown by arrows in the Figure) correspond to projecting out the columns to form subtables. Axiom 6 ensures that if the two tables each have a row with the same passenger’s name, then there really is at least one passenger with all of the associated attributes (start, destination, miles, and status). This is a table join operation: rows from the two tables that match correspond to rows in the joined table. Axiom 6 also ensures that there is *exactly one* such passenger, so in Figure 10 this means that “Bob” in Table 1 really is the same as “Bob” in Table 2. From personal experience – the author has a relatively common name – Axiom 6 does not hold for some airlines.

4. Data fusion and quality assessment

The previous section discussed how the relationships between sensors and their possible spaces of observations fit together into an integrated sensor system. Using the formal model of a sensor system as a presheaf $\mathcal{S}$ (satisfying Axioms 1–5), it becomes possible to talk about combining the data from sensors into a more unified perspective. The idea is that if two sensors describe overlapping sets of entities, their observations can be combined when they “agree” on the common entities. Since realistic data contains errors and uncertainties, it is rarely the case that “agreement” will be exact. This suggests formalizing

the distance between assignments, which has the structure of a pseudometric space.

Definition 18. The space of assignments $\prod_{U \in \mathcal{T}} \mathcal{S}(U)$ has a pseudometric, which is given by

$$D(a, b) = \sup_{U \in \mathcal{T}} d_U(a(U), b(U))$$

for two assignments a, b of $\mathcal{S}$.

Two assignments are close to one another whenever they consist of observations that are close in all sensor domains. In Example 9, assignments consist of one report from each sensor, each intersection, and from the field office. We should expect that the reports ought to be related – especially between a pair of sensors and their *joint* report on the same entities – so if one such report is awry, we should suspect that something may be amiss. Therefore, the central problem we must address is an optimization.

Problem 19. (Data Fusion) Given an assignment $a \in \prod_{U \in \mathcal{T}} \mathcal{S}(U)$, find the nearest global section s of $\mathcal{S}$, namely construct

$$\begin{aligned} \operatorname{argmin}_{s \in \mathcal{S}(X)} D(s, a) &= \operatorname{argmin}_{s \in \mathcal{S}(X)} \sup_{U \in \mathcal{T}} d_U(a(U), s(U)) \\ &= \operatorname{argmin}_{s \in \mathcal{S}(X)} \sup_{U \in \mathcal{T}} d_U(a(U), (\mathcal{S}(U \subseteq X))(s)). \end{aligned}$$

The constraint $s \in \mathcal{S}(X)$ on the argmin means that the optimization must proceed over all global sections of the presheaf $\mathcal{S}$. Thus, the presheaf structure will usually play a substantial role in constraining the solutions of Problem 19 through its space of global sections. As a practical matter, this means that the space of global sections must be determined *before* one can perform the optimization. The global sections are easily found by inspection in most of the examples in this article, as demonstrated in Examples 15 – 17. For more complicated situations, provided one has a sheaf of vector spaces (and not just a presheaf), the space of global sections can be computed via linear algebra (see Proposition 28 in Section 5). Without the presheaf structure, it is impossible to even pose Problem 19!

If each set of observations lies in a vector space, then Data Fusion is an optimization over vector spaces, which is a well-studied and well-understood setting. Our contribution is an appropriate pseudometric, which will depend intimately on the presheaf $\mathcal{S}$. If any set of observations is not a vector space, then Data Fusion will need to be solved by some other optimization method. Regardless, solving Data Fusion on a coarser topology is easier than on a finer one, since coarser topologies have fewer open sets.

In order to understand conditions under which this Problem may be solved, we make the following definition to quantify the self-consistency of an assignment.

Definition 20. An ϵ -approximate section for a presheaf of pseudometric spaces $\mathcal{S}$ is an assignment $s \in \prod_{U \in \mathcal{T}} \mathcal{S}(U)$ for which

$$d_V(s(V), \mathcal{S}(V \subseteq U)s(U)) \leq \epsilon$$

for all $V \subseteq U$. The minimum value of ϵ for which an assignment $s \in \prod_{U \in \mathcal{T}} \mathcal{S}(U)$ is an ϵ -approximate section is called the *consistency radius* of s .

Smaller consistency radii imply better agreement amongst sensors. Theoretically, the best one can do is exact equality, in which approximate sections closely correspond to global sections.

Proposition 21. *If $\mathcal{S}$ is a presheaf of pseudometric spaces, every global section $s \in \mathcal{S}(X)$ induces a 0-approximate section of $\mathcal{S}$. Furthermore, if $\mathcal{S}$ is a presheaf of metric spaces then this correspondence is a homeomorphism.*

Proof. To prove the first statement, it suffices to construct the other components in the product $\prod_{U \in \mathcal{T}} \mathcal{S}(U)$ by applying each restriction $\mathcal{S}(U \subseteq X)$ to s for $U \in \mathcal{T}$. Since d_U is a pseudometric, the distance between each such component will vanish by construction.

Observe that this defines an injective function i taking global sections to 0-approximate sections. To prove the second statement, we need only construct an inverse to i when $\mathcal{S}$ is a presheaf of metric spaces. To that end, suppose that s is instead a 0-approximate section of $\mathcal{S}$. Since s is an element of the product $\prod_{U \in \mathcal{T}} \mathcal{S}(U)$, we can apply the projection

$$j : \prod_{U \in \mathcal{T}} \mathcal{S}(U) \rightarrow \mathcal{S}(X)$$

to obtain a global section $j(s) = s(X)$ since $X \in \mathcal{T}$. It is clear that $j \circ i$ is the identity function. Since we have that s is a 0-approximate section, then

$$\begin{aligned} 0 &= d_V(s(V), \mathcal{S}(V \subseteq X)s(X)) \\ &= d_V(s(V), \mathcal{S}(V \subseteq X)j(s)) \\ &= d_V(s(V), i(j(s))_V) \end{aligned}$$

for all $V \in \mathcal{T}$, which implies that $s(V) = i(j(s))_V$ because d_V is a metric. Therefore $i \circ j$ is the identity function, and the proof is complete. $\square$

Remark 22. The technical distinction between the 0-approximate sections of presheaves of metric spaces and presheaves of pseudometric spaces runs quite deep. Since the space of 0-approximate sections is closed, the subspace of global sections is also closed in a presheaf of metric spaces, but not necessarily so in a presheaf of pseudometric spaces. Consider the topology $\mathcal{T} = \{\emptyset, \{1\}, \{1, 2\}\}$ on the two-point set $X = \{1, 2\}$ and the presheaf $\mathcal{P}$ defined by the diagram

$$\mathcal{P}(\{1, 2\}) = \{a, b\} \xrightarrow{\text{id}} \mathcal{P}(\{1\}) = \{a, b\}.$$

There exactly two global sections of $\mathcal{P}$, yet if the trivial pseudometric¹¹ is applied to all sensor domains, then there are four 0-approximate sections. Each of these approximate sections are limit points of the subspace of global sections, which shows that it is not closed. (As the reader may verify, $\mathcal{P}$ also satisfies Axiom 6, it is therefore a sheaf.)

The consistency radius of an assignment is an obstruction to it being a global section, as the following Proposition asserts.

Proposition 23. *If a is an ϵ -approximate section of a presheaf $\mathcal{P}$ of pseudometric spaces whose restrictions are Lipschitz continuous functions with Lipschitz constant K , then the distance (Definition 18) between a and the nearest global section to a is at least $\frac{\epsilon}{1+K}$.*

Assignments that are highly consistent (they have low consistency radius) are therefore easier to fuse via an optimization like Problem 19.

Proof. Suppose that s is any global section of $\mathcal{P}$ and that $V \subseteq U \in \mathcal{T}$ are open sets. Then the consistency radius is the supremum of all distances of the form

$$\begin{aligned} d_V(a(V), \mathcal{P}(V \subseteq U)a(U)) &\leq d_V(a(V), \mathcal{P}(V \subseteq U)s(U)) + \\ &\quad d_V(\mathcal{P}(V \subseteq U)s(U), \mathcal{P}(V \subseteq U)a(U)) \\ &\leq d_V(a(V), s(V)) + d_V(s(V), \mathcal{P}(V \subseteq U)s(U)) + \\ &\quad d_V(\mathcal{P}(V \subseteq U)s(U), \mathcal{P}(V \subseteq U)a(U)) \\ &\leq d_V(a(V), s(V)) + d_V(\mathcal{P}(V \subseteq U)s(U), \mathcal{P}(V \subseteq U)a(U)) \\ &\leq (1 + K)d_V(a(V), s(V)). \end{aligned}$$

Thus, taking supremums of both sides yields

$$\epsilon \leq (1 + K)D(a, s),$$

which completes the argument. $\square$

Example 24. Consider again the search and rescue scenario, recalling that we have constructed a sheaf model for the integrated sensor system in Examples 11 and 15. The data shown in Table 1 and Figures 11–13 were simulated to represent typical values for the sensors we deployed. To constitute the flight plan, we used an actual flight path [51] for Lufthansa flight 417. Locations and operating ranges of ATC radars were taken from Air Force document AFD-090615-023. The crash location was specified manually at 44.24545° N, 64.63672° W. The crash location was withheld from the data presented to the sheaf model, so that the output of its data fusion could be compared to the true value. The two RDF sensors were placed at

¹¹All attributes are distance 0 apart in the *trivial pseudometric*.

Table 1: Table of observations for Example 24.

Sensor	Entity	Case 1	Case 2	Case 3	Units
Flight plan	x	70.662	70.663	70.612	$^{\circ}$ W
	y	42.829	42.752	42.834	$^{\circ}$ N
	z	11178	11299	11237	m
ATC	x	70.587	70.657	70.617	$^{\circ}$ W
	y	42.741	42.773	42.834	$^{\circ}$ N
	z	11346	11346	11236	m
	v_x	-495	-495	-419	km/h W
	v_y	164	164	310	km/h N
RDF 1	θ_1	77.1	77.2	77.2	$^{\circ}$ true N
	t	0.943	0.930	0.985	h
RDF 2	θ_2	61.3	63.2	63.3	$^{\circ}$ true N
	t	0.890	0.974	1.05	h
Sat	s	(image)			
	s_x	64.599	64.630	62.742	$^{\circ}$ W
	s_y	44.243	44.287	44.550	$^{\circ}$ N
Field	x	70.649	70.668	70.626	$^{\circ}$ W
	y	42.753	42.809	42.814	$^{\circ}$ N
	z	11220	11431	11239	m
	v_x	-495	-495	-419	km/h W
	v_y	164	164	311	km/h N
	t	0.928	1.05	1.02	h
Crash est.	x	65.0013	64.2396	65.3745	$^{\circ}$ W
	y	44.1277	44.3721	45.6703	$^{\circ}$ N
Consist. rad. Error		15.7	11.6	152	km
		16.1	17.3	193	km

Table 2: Fused results for Example 24

Sensor	Entity	Case 1	Case 2	Case 3	Units
Field	x	70.9391	71.8296	75.7569	$^{\circ}$ W
	y	42.7849	42.7806	42.5831	$^{\circ}$ N
	z	10963	11730	12452	m
	v_x	-493.6	-497.2	-436.6	km/h W
	v_y	168.6	162.2	280.9	km/h N
	t	0.952	1.04	0.872	h
Crash est.	x	65.1704	65.4553	71.0939	$^{\circ}$ W
	y	44.2307	44.3069	44.7919	$^{\circ}$ N
Error		2.01	8.38	74.4	km

Figure 11: Case 1 in Example 24: (a) Errors in the presheaf model (b) Spatial layout.

Figure 12: Case 2 in Example 24: (a) Errors in the presheaf model (b) Spatial layout.

Figure 13: Case 3 in Example 24: (a) Errors in the presheaf model (b) Spatial layout.

- The home location for the Rensselaer Polytechnic Amateur radio club station W2SZ (the author was formerly a member) 42.7338328° N, 73.662574° W, in upstate New York and
- The author’s office at the American University Mathematics and Statistics department 38.9352387° N, 77.0897374° W in downtown Washington, DC.

We simulated three distinct cases, in which different errors were applied to the data. Measurements were perturbed by random errors that are representative of typical sensors.

We then solved Problem 19 by searching over the possible global sections to find the one closest (in the pseudometric from Definition 18) to each case. We used the Nelder-Mead simplex algorithm from the `scipy.optimize` [52] toolbox for this optimization. The results are shown in Table 2. The error between the predicted crash location (using dead reckoning from the last known position and velocity) improved rather dramatically by fusing the data.

Case 1: (See Figure 11) This case was constructed so that there was only small error present in all measurements. Applying dead reckoning to the Field office’s data results in a crash location estimate that is 16.1 km from the actual crash location. Examining Figure 11(b), it is clear that a crucial factor in obtaining this estimate was that the ATC measurement of the velocity vector and the RDF sensors’ estimate of the crash time were both reliable. Figure 11(a) shows that consistency radius of the assignment corresponding to this Case is about 15.7 km, which is fairly close to the actual error. A closer examination shows that nearly all the error is concentrated in two places: between the ATC (U_2) and flight plan (U_1) and between the Field office (X) and the satellite detection (U_5). The fused result exhibits markedly reduced error (a factor of $8\times$ reduced), largely since all of the errors were independent and unbiased.

Case 2: (See Figure 12) This case uses nearly the same setup as Case 1, but the RDF sensor 2 (from the author’s office) was biased southward by 2° . This error is immediately detectable in Figure 12(a) as an anomalously high error between the Field office (X) and the RDF sensor 2 (U_4). Since the only high errors occur between sensor domains involving RDF sensor 2 (U_4 and V_2), the sheaf model makes it clear that the errors are due exclusively to RDF sensor 2. It should be noted that Figure 12(a) does not include any *a priori* knowledge of the sensor bias. In this Case, the Field office did not use the RDF sensors for precise location – only timing – so its estimate by dead reckoning is only off by 17.3 km from the actual crash location. The consistency radius for this case is slightly smaller, at 11.6 km. This situation is a classic example of the power of data fusion, since the fused result shows that the error due to the biased RDF sensor is essentially eliminated by locating the nearest global section. The fused result exhibits around a factor of $2\times$ reduction in error.

Case 3: (See Figure 13) This case retains the RDF sensor 2 bias from Case 2, but also incorrectly records the last known velocity vector as being the one specified by the original flight plan. This is fairly realistic as target heading estimates near the edge of the usable range of a radar are often quite noisy. The consistency radius is easily read from Figure 13(a) as 152 km. This is definitely an indication that something is amiss, given that the error incurred by dead reckoning is visible on Figure 13(c) and is 193 km. Even so, the fused result is considerably closer to the actual crash location, exhibiting somewhat over a factor of $2\times$ lower error.

5. Analysis of integrated sensor systems

In his Appendix on sheaf theory, Hubbard states [53] “It is fairly easy to understand what a sheaf is, especially after looking at a few examples. Understanding what they are good for is rather harder; indeed, without cohomology theory, they aren’t good for much.” Although the previous sections of this article advocate otherwise, Hubbard makes an important point. Although there are many co-universal information representations (of which sheaves are one), the stumbling block for integration approaches is the lack of theoretically-motivated computational approaches. What sheaf theory provides is a way to perform *controlled algebraic summarization* – a weaker set of invariants that still permit certain kinds of inference. In order to perform these summarizations, the attributes must be specially encoded, which leads to our final Axiom.

Axiom 7: Each space of observations $\mathcal{S}(U)$ has the structure of a Banach space (complete normed vector space) and each restriction $\mathcal{S}(U \subseteq V)$ is a continuous linear transformation of this Banach space. We call $\mathcal{S}$ a *sheaf of Banach spaces*.

Proposition 25. *If $\mathcal{S}$ is a sheaf of Banach spaces on a finite topology $\mathcal{T}$, then Problem 19 always has a unique solution, embodied as a projection from the space of assignments to the space of global sections.*

Proof. One need only observe that the metric D on assignments (Definition 18) makes the space of assignments into a Banach space. Specifically, if $a \in \prod_{U \in \mathcal{T}} \mathcal{S}(U)$ is an assignment of $\mathcal{S}$, let

$$\|a\| = \max_{U \in \mathcal{T}} \|a(U)\|_{\mathcal{S}(U)},$$

where $\|\cdot\|_{\mathcal{S}(U)}$ is the norm on $\mathcal{S}(U)$. The maximum exists since $\mathcal{T}$ is finite. Then the metric D is given by

$$\begin{aligned} \|a - b\| &= \max_{U \in \mathcal{T}} \|a(U) - b(U)\|_{\mathcal{S}(U)} \\ &= \max_{U \in \mathcal{T}} d_U(a(U), b(U)) \\ &= D(a, b). \end{aligned}$$

□

Figure 14: The sheaf in the search and rescue example written using linear stochastic maps.

Figure 15: The sheaf in the search and rescue example written using linear stochastic maps and means.

There are a number of additional benefits once Axiom 7 is satisfied, for instance,

1. The space of global sections can be computed directly using linear algebra (Proposition 28), which greatly simplifies solving the Problem 19, the data fusion problem,
2. There is a canonical collection of invariants, against which all others can be compared (Definition 27), and
3. We need not construct sets of observations over the unions of sensor domains in order to access these invariants (Theorem 29).

5.1. Preparation: Stochastic observations

Although Axiom 7 usually does permit reasonable options for restrictions, it may also be highly constraining. Many effective algorithms that could be used as restrictions are not linear. Although this seems to dramatically limit the effectiveness of our subsequent analysis, there is a way out that actually increases the expressiveness of the model. Specifically, a non-linear restriction map can be “enriched” into one that is linear [54, 55]. The most concrete of these enrichments lifts the entities and attributes into a stochastic modeling framework, as the next Example shows.

Example 26. The sheaf model $\mathcal{S}$ defined for the search and rescue scenario in Example 15 does not satisfy Axiom 7, since the restrictions A , B , C , D , and E are not linear transformations. A different sheaf model $\mathcal{R}$ can instead be defined that does satisfy Axiom 7, by reinterpreting the attributes as probability

distributions. The restrictions then all become stochastic maps, all of which are linear.

The sheaf $\mathcal{R}$ is built upon the same topological space $(X, \mathcal{T})$ as $\mathcal{S}$, and each space of observations over an open set $U \in \mathcal{T}$ is replaced by the recipe

$$\mathcal{R}(U) = L^1(\mathcal{S}(U), \mathbb{R}),$$

where $L^1(M, \mathbb{R})$ is the space of real-valued integrable functions on a measure space M . We interpret $L^1(M, \mathbb{R})$ as the vector space generated by the probability distributions on M .

The appropriate sheaf diagram of $\mathcal{R}$ is shown in Figure 14, for which we now describe the restrictions. Every restriction of $\mathcal{S}$ translates into a linear, stochastic map in $\mathcal{R}$ according to the following procedure, which we state generally first and then specialize. Using an arbitrary measurable function $f : M \rightarrow N$, we can define a function $F : L^1(M, \mathbb{R}) \rightarrow L^1(N, \mathbb{R})$ by the following formula for each $u \in L^1(M, \mathbb{R})$ and $y \in N$

$$(F(u))(y) = \int_{f^{-1}(\{y\})} u(x) dx = \int_M u(x) \delta(y - f(x)) dx, \quad (2)$$

where δ is the Dirac distribution. If the volume of M is 1, so that

$$\int_M dx = 1,$$

then (2) is precisely Bayes' rule, and F can be shown to be a stochastic linear map acting on u , by an appeal to Fubini's theorem. Specifically, if u is nonnegative and has integral 1, then $F(u)$ will also have these properties.

In the case of the restrictions *a-d* in Figure 14 for $\mathcal{R}$, all of which are projectors of various sorts, Equation 2 specializes to

$$\begin{aligned} (a(u))(\theta) &= \int_{-\infty}^{\infty} u(\theta, t) dt, \\ (b(u))(t) &= \int_0^{2\pi} u(\theta, t) d\theta, \\ (c(u))(x, y, z) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} u(x, y, z, v_x, v_y) dv_x dv_y, \text{ and} \\ (d(u))(x, y, z, v_x, v_y) &= \int_{-\infty}^{\infty} u(x, y, z, v_x, v_y, t) dt, \end{aligned}$$

where we have used u to represent a typical function in the domain of each restriction.

Of course, we must also define A' , B' , C' , D' , and E' using Equation 2,

replacing f with each restriction from $\mathcal{S}$, so we obtain

$$\begin{aligned}
(A'(u))(\theta_1, t) &= \int \int \int \int_0^\infty u(r \sin \theta_1 + r_{1x} - v_x t, r \cos \theta_1 + r_{1y} - v_y t, z, v_x, v_y, t) r dr dz dv_x dv_y, \\
(B'(u))(\theta_2, t) &= \int \int \int \int_0^\infty u(r \sin \theta_2 + r_{2x} - v_x t, r \cos \theta_2 + r_{2y} - v_y t, z, v_x, v_y, t) r dr dz dv_x dv_y, \\
(C'(u))(\theta_1) &= \int_0^\infty u(r \sin \theta_1 + r_{1x}, r \cos \theta_1 + r_{1y}) r dr, \\
(D'(u))(\theta_2) &= \int_0^\infty u(r \sin \theta_2 + r_{2x}, r \cos \theta_2 + r_{2y}) r dr, \\
(E'(u))(s_x, s_y) &= \int \int \int \int u(s_x - v_x t, s_y - v_y t, z, v_x, v_y, t) dz dv_x dv_y dt,
\end{aligned}$$

where (s_x, s_y) are coordinates of an object detected in the satellite image, (r_{1x}, r_{1y}) are the coordinates of the first RDF sensor and (r_{2x}, r_{2y}) are the coordinates of the second RDF sensor.

It is not strictly necessary to retain probability distributions everywhere in the sheaf model. Portions of the sheaf in which there were already linear restrictions can be retained undisturbed by extracting the mean values of the distributions. This process is not unique, but one possibility is shown in Figure 15, where

$$\begin{aligned}
b'(u) &= \int_{-\infty}^\infty t(b(u))(t) dt \\
&= \int_{-\infty}^\infty \int_0^{2\pi} tu(\theta, t) d\theta dt, \\
(d'(u))(x, y, z, v_x, v_y) &= \left(\int \int \int \int x(d(u))(x, y, z, v_x, v_y) dx dy dz dv_x dv_y, \dots \right).
\end{aligned}$$

5.2. Computing cohomology

The sheaf $\mathcal{S}$ of vector spaces that results from Axiom 7 has a canonical family of invariants called *cohomology groups* $H^k(\mathcal{S})$ for $k = 0, 1, \dots$. In the most general setting, the cohomology groups can be constructed in various ways using linear algebraic manipulations, most notably the *canonical* or *Godement resolution* [56][57, II.2]. Instead, we will use a more practical construction, called *Čech cohomology* [53, 57].

Definition 27. Suppose $\mathcal{U} = \{U_1, \dots, U_n\}$ is a finite open cover for a topological space $(X, \mathcal{T})$, which means that $\mathcal{U} \subseteq \mathcal{T}$ and $\cup \mathcal{U} = X$. Given a sheaf $\mathcal{S}$ of Banach spaces on $\mathcal{T}$, the *k-cochain spaces of $\mathcal{S}$ subject to $\mathcal{U}$* are given by the Cartesian products of the stalks over $(k+1)$ -way intersections of sets in $\mathcal{U}$

$$C^k(\mathcal{U}; \mathcal{S}) = \prod_{i_0 < i_1 \dots i_k} \mathcal{S}(U_{i_0} \cap \dots U_{i_k}).$$

Notice that the coordinates of elements of $C^k(\mathcal{U}; \mathcal{S})$ are indexed by lists of $(k+1)$ increasing integers to eliminate duplications.

Beware that the k on C^k is merely an index, and that C^k does *not* mean the space of k -times continuously differentiable functions!

We also define the k -coboundary map, which is a function $d^k : C^k(\mathcal{U}; \mathcal{S}) \rightarrow C^{k+1}(\mathcal{U}; \mathcal{S})$. This function is defined by specifying its coordinates on an element $a \in C^k(\mathcal{U}; \mathcal{S})$.

$$d^k(a)_{i_0 < \dots < i_{k+1}} = \sum_{j=0}^{k+1} (-1)^j \mathcal{S}(U_{i_0} \cap \dots \cap U_{i_{k+1}} \subseteq U_{i_0} \cap \dots \widehat{U_{i_j}} \dots \cap U_{i_{k+1}}) a_{i_0 < \dots \widehat{i_j} \dots i_{k+1}},$$

where the hat indicates omission from a list.

It is a standard fact that $d^{k+1} \circ d^k = 0$, from which we can make the definition of the k -cohomology spaces

$$H^k(\mathcal{U}; \mathcal{S}) = \ker d^k / \text{image } d^{k+1}.$$

We sometimes write $H^k(\mathcal{S})$ for $H^k(\mathcal{T}; \mathcal{S})$ to reduce notational clutter.

Although computing $H^k(\mathcal{T}; \mathcal{S})$ is a bit difficult, simply due to the number of dimensions of $C^k(\mathcal{T}; \mathcal{S})$, the first consequence of this definition is that the zeroth cohomology classes are the global sections.

Proposition 28. [58, Thm 4.3] *If $\mathcal{S}$ is a sheaf of vector spaces on a finite topology $\mathcal{T}$, then $H^0(\mathcal{T}; \mathcal{S}) = \ker d^0$ consists precisely of those assignments which are global sections of $\mathcal{S}$.*

Since the $C^0(\mathcal{T}; \mathcal{S})$ is also the space of assignments, Proposition 25 implies that the solution to the Data Fusion Problem (Problem 19) is merely a projection $C^0(\mathcal{T}; \mathcal{S}) \rightarrow H^0(\mathcal{T}; \mathcal{S})$.

Unfortunately, due to the way that $H^0(\mathcal{T}; \mathcal{S})$ is constructed, $\mathcal{S}(X)$ is already a factor in $C^0(\mathcal{T}; \mathcal{S})$. This is remedied by the Leray theorem, which allows for a considerable reduction in complexity by verifying whether a smaller cover $\mathcal{U}$ than the whole topology $\mathcal{T}$ still recovers all of the cohomology spaces.

Theorem 29. (Leray theorem, [59] [53, Thm A7.2.6]/[57, Thm 4.13 in III.4]) *Suppose $\mathcal{S}$ is a sheaf on $\mathcal{T}$, and that $\mathcal{U} \subseteq \mathcal{T}$ is a collection of open sets. If for every intersection $U = U_0 \cap \dots \cap U_k$ of a collection of elements in $\mathcal{U}$ it follows that*

$$H^k(U \cap \mathcal{T}; \mathcal{S}) = 0 \text{ for all } k > 0,$$

then we can conclude that $H^p(\mathcal{U}; \mathcal{S}) \cong H^p(\mathcal{T}; \mathcal{S})$ for all p .

Practically speaking, the nontrivial elements of H^k for $k > 0$ represent obstructions to consistent sensor data. They indicate situations in which it is impossible for an assignment to some open sets to be extended to the rest of the topology. Reading from Definition 27, a nontrivial element of $H^k(\mathcal{U}; \mathcal{S})$ consists of a joint collection of observations on the k -way intersections of sensor domains

Figure 16: A collection of images in a nontrivial H^1 cohomology class.

in $\mathcal{U}$ that are consistent upon further restriction to $(k + 1)$ -way intersections, but do not arise from any joint observations from $(k - 1)$ -way intersections. They are therefore certain classes of partially self-consistent data that cannot be extended to global sections. From a data fusion standpoint, they may be considered solutions to Problem 19 on only part of the space, but these solutions cannot be extended to all of the space without incurring distortion.

When situations with nontrivial H^k for $k > 0$ arise in practice they usually indicate a failure of the model to correspond to reality. For instance, in all of the mosaic-forming examples discussed in this article, there has been an implicit assumption that all images were acquired at the same time. When time evolves between imaging events, nontrivial H^k elements indicate that inconsistency can easily arise, as Example 30 shows.

Example 30. Consider the collection of photographs shown in Figure 16. Consider the sheaf $\mathcal{S}$ for the three sensor domains $\mathcal{U} = \{U_i\}$ as shown, where the observations consist of various spaces of images and the restrictions consist of image cropping operations, as in Examples 8 and 16.

Consider the two assignments consisting of the photograph on the left (on U_1) and the bottom (U_2) and one of the photographs on the right (U_3). Neither of these assignments correspond to a global section, because one of the restriction maps out of U_3 will result in a disagreement on one of its intersections with the other two sets. In particular, the set of photographs shown on the intersections could not possibly have arisen by restricting *any* assignment of photographs to the elements of the cover, that is, the image through d^1 of an element of $C^0(\mathcal{U}; \mathcal{S})$. Therefore, the assignment of photographs to the intersections shown in Figure 16 is an example of a nontrivial class in $H^1(\mathcal{T}; \mathcal{S})$ since there are no

Figure 17: Trying to see behind a book with two different views.

nonempty intersections between all three sensor domains.

In light of Definition 27 and the example, we may be tempted to add more sensors to our collection to try to resolve the ambiguity. The essential fact is that without bringing time into the model of mosaic-forming, problems still can arise. Under fairly general conditions, such as those in the above example, the collection of sensors contains all of the possible information as is truly encoded by the model.

Example 31. Consider the case of trying to use a collection of sensors to piece together what is behind an obstacle. This occurs in optical systems as well as those based on radio or sound propagation. Usually, the process involves taking collections of images from several angles followed by an obstacle-removal step. If not enough looks are obtained from enough angles, then it will not be possible to infer anything from behind the obstacle. However, if the sensors' locations are not well known, it may be impossible to detect that this has occurred. This is an effect of a difference between the sensors' view of $H^1(\mathcal{U}; \mathcal{S})$ and the actual $H^1(\mathcal{T}; \mathcal{S})$.

Figure 17 shows the process for trying to see behind a book using two views. The original photographs on the left show the scene, while the photos on the right show the obstacle removed. Suppose that X consists of the following entities: L for the left camera, R for the right camera, V_1 for the upper overlap, and V_2 for the lower overlap. That is,

$$X = \{R, L, V_1, V_2\}.$$

There is a natural topology for the sensors, namely

$$\mathcal{T} = \{\{R, V_1, V_2\}, \{L, V_1, V_2\}, \{V_1, V_2\}, \{V_1\}, \{V_2\}, X, \emptyset\}$$

in which the overlaps are separated. Modeling the cameras as $\mathcal{U} = \{U_L, U_R\}$ where $U_L = \{L, V_1, V_2\}$, $U_R = \{R, V_1, V_2\}$, we obtain the following sheaf $\mathcal{M}$ associated to the topology generated by $\mathcal{U}$:

$$\begin{array}{ccccc}
 & & X & & \\
 & \swarrow & & \searrow & \\
 U_L & \longleftarrow \{V_1, V_2\} & \longrightarrow & U_R & \\
 & & & & \\
 & & \mathbb{R}^{3m} \oplus \mathbb{R}^{3n} & & \\
 & \swarrow & & \searrow & \\
 \mathbb{R}^{3m} & \longrightarrow \mathbb{R}^{3p} \oplus \mathbb{R}^{3q} & \longleftarrow & \mathbb{R}^{3n} &
 \end{array}$$

Figure 18: The sheaf for Example 32. (a) A representation of the topological space $\mathcal{T}$ in Examples 31 and 32. (b) The open sets in the topology. (c) The sheaf $\mathcal{P}$ and its restrictions.

where m is the number of pixels in the left image, n is the number of pixels in the right image, p is the number of pixels in the upper overlap region, and q is the number of pixels in the lower overlap region. Since the overlap region consists of two connected components, does this impact the mosaic process? It does not, according to Theorem 29, though this requires a little calculation. It is clear the only place we need to be concerned is the overlap. So we need to examine the intersection $U_L \cap U_R = \{V_1, V_2\}$ and compute

$$H^k(\mathcal{V}; \mathcal{M}),$$

where $\mathcal{V} = \{\{V_1, V_2\}, \{V_1\}, \{V_2\}\}$. The sheaf diagram for this (with the sets of $\mathcal{V}$ below for clarity) is

$$\mathbb{R}^{3p} \xleftarrow{\text{pr}_1} \mathbb{R}^{3p} \oplus \mathbb{R}^{3q} \xrightarrow{\text{pr}_2} \mathbb{R}^{3q}$$

$$\{V_1\} \longrightarrow \{V_1, V_2\} \longleftarrow \{V_2\}$$

which is easily computed to have $\dim H^k(\mathcal{V}; \mathcal{M}) = 0$ for $k > 0$. (In the diagram, pr_k is the projection onto the k -th factor of a product and id is the identity map.)

The above example shows that solving a data fusion problem applied to relatively unprocessed data works well.¹² It also works if data are processed a bit more, as the next (and final) example shows.

Example 32. Suppose instead that we are interested not in the images themselves but rather whether a single object is in the scene. Like the previous example the scene can be divided into several regions, each represented by an entity with the topology $\mathcal{T}$. If we attempt to quantify the probability that the object is in the left, right, or middle of the scene, the space of observations over U_L (or U_R) is $\mathbb{R}^2$, whose two components should be interpreted as the

¹²It works well up to the limit of having enough agreement between the sensors to deduce the overlaps, of course. This is not necessarily easy!

probability that an object is present on the left (or right) and in the overlap. We are led to consider a sheaf $\mathcal{P}$ on $\mathcal{T}$ that keeps track of whether there is an object on the left, right, or middle of the scene as shown in Figure 18(c). As before, if $\mathcal{U} = \{U_L, U_R\}$ where $U_L = \{L, V_1, V_2\}$, $U_R = \{R, V_1, V_2\}$, then $\dim H^0(\mathcal{U}; \mathcal{P}) = 3$, indicating the three probability values for the left, right, and middle of the scene. However, unlike the Example 31, $\dim H^1(\mathcal{U}; \mathcal{P}) = 1$, which indicates that there can be data that are locally self-consistent, but not globally self-consistent. This corresponds to the situation where there are different probabilities for the object's presence in V_1 and V_2 – both variants of the middle region – yet the global model only discerns a single probability for the middle region.

We can confirm that this is not an artifact by examining the refined cover $\mathcal{V} = \{\{V_1, V_2\}, \{V_1\}, \{V_2\}\}$ as in the previous example, and are once again led to the fact that $\dim H^k(\mathcal{V}; \mathcal{P}) = 0$ for $k > 0$.

The presence of nontrivial cohomology in H^1 arising after performing some data processing indicates that the processing has lost information at a global level, but that this information is still possible to detect with individual sensors. This strongly advocates for performing fusion as close to the sensors as possible, to avoid unexpected correlations across sensor domains due to inconsistent processing. This view is shared by others, for instance [28], though in some cases the artifacts can be anticipated. For instance, [60] shows how to use spectral sequences to compensate for when the hypotheses of Theorem 29 are known to fail. As a special case for graphs, [37] and [58, Sec 4.7] contain an algorithmic approach that is somewhat simpler.

6. Conclusions and new frontiers

This article has demonstrated that sheaves are an effective theoretical tool for constructing and exploiting integrated systems of sensors in a bottom-up, data-driven way. By constructing sheaf models of integrated sensor systems, one can use them as a basis for solving data fusion problems driven by sensor data. One can also apply cohomology as a tool to identify features of the integrated system as a whole, which could be useful for sensor deployment planning.

There are several mathematical frontiers that are opened by repurposing sheaf theory for information integration. For instance, although we have focused on global sections and obstructions to them, the combinatorics of *local* sections is essentially not studied. The extension of observations to global sections arises in many important problems. In homogeneous data situations, such as the reconstruction of signals from samples [61] and the solution of differential equations [62], sheaves are already useful. The generality afforded by sheaves indicates that problems of this sort can also be addressed when the data types are not homogeneous. Although persistent homology has received considerable recent attention in exploratory data analysis, persistent cohomology [63] has shown promise in signal processing and is a sheaf [58].

Although the examples in this article were large enough to be useful on their own, there is also considerable work to be done to address larger collections of sensors. Although solving the data fusion problem is straightforward computationally, computing cohomology can require significant computing resources. In homology calculations, the notion of collapses and reductions are essential for making computation practical. A similar methodology – based on discrete Morse theory [64] – exists for computing sheaf cohomology. In either case (reduced or not), efficient finite field linear algebra could be an enabler. Finally, practical categorification [55] is still very much an art, and has a large impact on practical performance of sheaf-based information integration systems that rely on cohomology.

Acknowledgment

The author would like to thank the anonymous referees for their insightful comments and questions. This work was partially supported under DARPA SIMPLEX N66001-15-C-4040. The views, opinions, and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

- [1] S. Magnuson, “Military ‘swimming in sensors and drowning in data’,” *National Defense*, vol. 94, January 2010.
- [2] D. L. Hall, M. McNeese, J. Llinas, and T. Mullen, “A framework for dynamic hard/soft fusion,” in *Information Fusion, 2008 11th International Conference on*, pp. 1–8, IEEE, 2008.
- [3] F. White, “Data fusion lexicon,” Tech. Rep. ADA529661, Joint Directors of Laboratories, Washington, DC, October 1991.
- [4] L. Wald, “Some terms of reference in data fusion,” *IEEE Trans. Geosciences and Remote Sensing*, vol. 37, no. 3, pp. 1190–1193, 1999.
- [5] L. Wald, *Data fusion: definitions and architectures: fusion of images of different spatial resolutions*. Presses des MINES, 2002.
- [6] B. Khaleghi, A. Khamis, F. O. Karray, and S. N. Razavi, “Multisensor data fusion: A review of the state-of-the-art,” *Information Fusion*, vol. 14, no. 1, pp. 28–44, 2013.
- [7] P. K. Varshney, “Multisensor data fusion,” *Electronics & Communication Engineering Journal*, vol. 9, no. 6, pp. 245–253, 1997.
- [8] L. Alparone, B. Aiazzi, S. Baronti, A. Garzelli, F. Nencini, and M. Selva, “Multispectral and panchromatic data fusion assessment without reference,” *Photogrammetric Engineering & Remote Sensing*, vol. 74, no. 2, pp. 193–200, 2008.

- [9] J. Zhang, “Multi-source remote sensing data fusion: status and trends,” *International Journal of Image and Data Fusion*, vol. 1, no. 1, pp. 5–24, 2010.
- [10] S. Dawn, V. Saxena, and B. Sharma, “Remote sensing image registration techniques: A survey,” in *Image and Signal Processing: 4th International Conference, ICISP 2010, Trois-Rivières, QC, Canada, June 30-July 2, 2010. Proceedings* (A. Elmoataz, O. Lezoray, F. Nouboud, D. Mammass, and J. Meunier, eds.), pp. 103–112, Berlin, Heidelberg: Springer Berlin Heidelberg, 2010.
- [11] B. Koetz, G. Sun, F. Morsdorf, K. Ranson, M. Kneubhler, K. Itten, and B. Allgwer, “Fusion of imaging spectrometer and LIDAR data over combined radiative transfer models for forest canopy characterization,” *Remote Sensing of Environment*, vol. 106, no. 4, pp. 449 – 459, 2007.
- [12] Z. Guo, G. Sun, K. Ranson, W. Ni, and W. Qin, “The potential of combined LIDAR and SAR data in retrieving forest parameters using model analysis,” in *Geoscience and Remote Sensing Symposium, 2008. IGARSS 2008. IEEE International*, vol. 5, pp. V – 542–V – 545, July 2008.
- [13] J. Chen, H. Feng, K. Pan, Z. Xu, and Q. Li, “An optimization method for registration and mosaicking of remote sensing images,” *Optik - International Journal for Light and Electron Optics*, vol. 125, no. 2, pp. 697 – 703, 2014.
- [14] J. Euzenat and P. Shvaiko, *Ontology Matching*. Hiedelberg: Springer-Verlag, 2007.
- [15] J. Euzenat, A. Ferrara, P. Hollink, A. Isaac, C. Joslyn, V. Malaisé, C. Meilicke, A. Nikolov, J. Pane, M. Sabou, F. Scharffe, P. Shvaiko, V. Spiliopoulos, H. Stuckenschmidt, O. Iváb-Zamazal, V. Svátek, C. T. dos Santos, G. Vouros, and S. Wang, “Results of the ontology alignment evaluation initiative 2009,” in *Proc. 4th Int. Wshop. on Ontology Matching (OM-2009), CEUR*, vol. 551, 2009.
- [16] C. Joslyn, P. Paulson, and A. White, “Measuring the structural preservation of semantic hierarchy alignments,” in *Proc. 4th Int. Wshop. on Ontology Matching (OM-2009), CEUR*, vol. 551, 2009.
- [17] E. G. Little and G. L. Rogova, “Designing ontologies for higher level fusion,” *Information Fusion*, vol. 10, no. 1, pp. 70–82, 2009.
- [18] J. Hammer, H. Garcia-Molina, J. Cho, R. Aranha, and A. Crespo, “Extracting semistructured information from the web,” 1997.
- [19] E. Riloff, R. Jones, *et al.*, “Learning dictionaries for information extraction by multi-level bootstrapping,” in *AAAI/IAAI*, pp. 474–479, 1999.

- [20] M. L. Gregory, L. McGrath, E. B. Bell, K. O'Hara, and K. Domico, "Domain independent knowledge base population from structured and unstructured data sources," in *FLAIRS Conference*, 2011.
- [21] N. Kushmerick, "Wrapper induction: Efficiency and expressiveness," *Artificial Intelligence*, vol. 118, no. 1, pp. 15–68, 2000.
- [22] G. Jakobson, L. Lewis, and J. Buford, "An approach to integrated cognitive fusion," in *Proceedings of the 7th ISIF International Conference on Information Fusion (FUSION2004)*, pp. 1210–1217, 2004.
- [23] S. Sarawagi, "Information extraction," *Foundations and trends in databases*, vol. 1, no. 3, pp. 261–377, 2008.
- [24] E. Waltz, J. Llinas, *et al.*, *Multisensor data fusion*, vol. 685. Artech house Boston, 1990.
- [25] M. H. DeGroot, *Optimal statistical decisions*. Wiley, 2004.
- [26] D. L. Hall and S. A. McMullen, *Mathematical techniques in multisensor data fusion*. Artech House, 2004.
- [27] D. Smith and S. Singh, "Approaches to multisensor data fusion in target tracking: A survey," *Knowledge and Data Engineering, IEEE Transactions on*, vol. 18, no. 12, pp. 1696–1710, 2006.
- [28] A. J. Newman and G. E. Mitzel, "Upstream data fusion: history, technical overview, and applications to critical challenges," *Johns Hopkins APL Technical Digest*, vol. 31, no. 3, pp. 215–233, 2013.
- [29] S. M. Alqhtani, S. Luo, and B. Regan, "Multimedia data fusion for event detection in twitter by using dempster-shafer evidence theory," *World Academy of Science, Engineering and Technology, International Journal of Computer, Electrical, Automation, Control and Information Engineering*, vol. 9, no. 12, pp. 2234–2238.
- [30] J. L. Crowley, "Principles and techniques for sensor data fusion," in *Multisensor Fusion for Computer Vision*, pp. 15–36, Springer, 1993.
- [31] S. Benferhat and C. Sossai, "Reasoning with multiple-source information in a possibilistic logic framework," *Information Fusion*, vol. 7, no. 1, pp. 80–96, 2006.
- [32] S. Benferhat and F. Titouna, "Fusion and normalization of quantitative possibilistic networks," *Applied Intelligence*, vol. 31, no. 2, pp. 135–160, 2009.
- [33] D. Balduzzi, "On the information-theoretic structure of distributed measurements," *EPTCS*, vol. 88, pp. 28–42, 2012.

- [34] S. Thorsen and M. Oxley, “A description of competing fusion systems,” *Information Fusion*, vol. 7, pp. 347–360, December 2006.
- [35] K. P. Murphy, Y. Weiss, and M. I. Jordan, “Loopy belief propagation for approximate inference: An empirical study,” in *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, pp. 467–475, Morgan Kaufmann Publishers Inc., 1999.
- [36] J. Lilius, “Sheaf semantics for Petri nets,” tech. rep., Helsinki University of Technology, Digital Systems Laboratory, 1993.
- [37] M. Robinson, “Imaging geometric graphs using internal measurements,” *J. Diff. Eqns.*, vol. 260, pp. 872–896, 2016.
- [38] C. Joslyn, E. Hogan, and M. Robinson, “Towards a topological framework for integrating semantic information sources,” in *Semantic Technologies for Intelligence, Defense, and Security (STIDS)*, 2014.
- [39] K. Baclawski, “Whitney numbers of geometric lattices,” *Adv. in Math.*, vol. 16, pp. 125–138, 1975.
- [40] K. Baclawski, “Galois connections and the Leray spectral sequence,” *Adv. Math.*, vol. 25, pp. 191–215, 1977.
- [41] A. Shepard, *A cellular description of the derived category of a stratified space*. PhD thesis, Brown University, 1985.
- [42] J. Curry, “Sheaves, cosheaves and applications, [arxiv:1303.3255](#),” 2013.
- [43] R. Goldblatt, *Topoi, the Categorical Analysis of Logic*. Dover, 2006.
- [44] J. A. Goguen, “Sheaf semantics for concurrent interacting objects,” *Mathematical Structures in Computer Science*, vol. 2, no. 02, pp. 159–191, 1992.
- [45] D. I. Spivak, *Category theory for the sciences*. MIT Press, 2014.
- [46] M. Kokar, J. Tomasik, and J. Weyman, “Formalizing classes of information fusion systems,” *Information Fusion*, vol. 5, pp. 189–202, 2004.
- [47] M. M. Kokar, K. Baclawski, and H. Gao, “Category theory-based synthesis of a higher-level fusion algorithm: An example,” in *Information Fusion, 2006 9th International Conference on*, pp. 1–8, IEEE, 2006.
- [48] Australian Transport Safety Bureau, “Considerations on defining the search area – MH370,” May 2014.
- [49] R. E. Stong, “Finite topological spaces,” *Trans. Amer. Math. Soc.*, vol. 123, pp. 325–340, June 1966.
- [50] J. May, “Finite topological spaces: Notes for REU,” 2003.

- [51] Flightaware.com, “Lufthansa (LH) #417 – 26-Jul-2016 – KIAD - FRA/EDDF, <http://flightaware.com/live/flight/dlh417/history/20160726/1925z/kiad/eddf>,” 2016.
- [52] T. S. community, “Optimization: scipy.optimize.” <http://docs.scipy.org/doc/scipy/reference/tutorial/optimize.html>. accessed Sept 15, 2016.
- [53] J. H. Hubbard, *Teichmüller Theory, volume 1*. Matrix Editions, 2006.
- [54] M. Robinson, “Asynchronous logic circuits and sheaf obstructions,” *Electronic Notes in Theoretical Computer Science*, pp. 159–177, 2012.
- [55] E. Purvine, M. Robinson, and C. Joslyn, “Categorification in the real world,” in *Joint Mathematics Meetings*, (Seattle, WA), January 2016.
- [56] R. Godement, *Topologie algebrique et théorie des faisceaux*. Paris: Herman, 1958.
- [57] G. Bredon, *Sheaf theory*. Springer, 1997.
- [58] M. Robinson, *Topological Signal Processing*. Springer, 2014.
- [59] J. Leray, “L’anneau spectral et l’anneau filtré d’homologie d’un espace localement compact et d’une application continue,” *Journal des Mathématiques Pures et Appliquées*, vol. 29, pp. 1–139, 1950.
- [60] S. Basu, “Computing the Betti numbers of arrangements via spectral sequences,” *Journal of Computer and System Sciences*, vol. 67, no. 2, pp. 244–262, 2003.
- [61] M. Robinson, “The Nyquist theorem for cellular sheaves,” in *Sampling Theory and Applications*, (Bremen, Germany), 2013.
- [62] L. Ehrenpreis, “Sheaves and differential equations,” *Proceedings of the American Mathematical Society*, vol. 7, no. 6, pp. 1131–1138, 1956.
- [63] V. de Silva, D. Morozov, and M. Vejdemo-Johansson, “Persistent cohomology and circular coordinates,” *Discrete & Computational Geometry*, vol. 45, no. 4, pp. 737–759, 2011.
- [64] J. Curry, R. Ghrist, and V. Nanda, “Discrete morse theory for computing cellular sheaf cohomology, [arxiv:1312.6454](https://arxiv.org/abs/1312.6454),” 2015.